

Developing Language Processing Components with GATE Version 5 (a User Guide)

For GATE version 5.2.1
(built May 6, 2010)

Hamish Cunningham
Diana Maynard
Kalina Bontcheva
Valentin Tablan
Niraj Aswani
Ian Roberts
Genevieve Gorrell
Adam Funk
Angus Roberts
Danica Damljanovic
Thomas Heitz
Mark Greenwood
Horacio Saggion
Johann Petrak
Yaoyong Li
Wim Peters
et al

©The University of Sheffield 2001-2010

<http://gate.ac.uk/>

HTML version: <http://gate.ac.uk/userguide>

Work on GATE has been partly supported by EPSRC grants GR/K25267 (Large-Scale Information Extraction), GR/M31699 (GATE 2), RA007940 (EMILLE), GR/N15764/01 (AKT) and GR/R85150/01 (MIAKT), AHRB grant APN16396 (ETCSL/GATE), Matrixware, the Information Retrieval Facility and several EU-funded projects (SEKT, TAO, NeOn, MediaCampaign, MUSING, KnowledgeWeb, PrestoSpace, h-TechSight, enIRaF).

Brief Contents

I	GATE Basics	3
1	Introduction	5
2	Installing and Running GATE	27
3	Using GATE Developer	37
4	CREOLE: the GATE Component Model	67
5	Language Resources: Corpora, Documents and Annotations	87
6	ANNIE: a Nearly-New Information Extraction System	109
II	GATE for Advanced Users	127
7	GATE Embedded	129
8	JAPE: Regular Expressions over Annotations	169
9	ANNIC: ANNotations-In-Context	203
10	Performance Evaluation of Language Analysers	213
11	Profiling Processing Resources	237
12	Developing GATE	245
III	CREOLE Plugins	257
13	Gazetteers	259
14	Working with Ontologies	281
15	Machine Learning	317
16	Tools for Alignment Tasks	367
17	Parsers and Taggers	381

18 Combining GATE and UIMA	409
19 More (CREOLE) Plugins	421
Appendices	457
A Change Log	459
B Version 5.1 Plugins Name Map	483
C Design Notes	485
D JAPE: Implementation	493
E Ant Tasks for GATE	503
F Named-Entity State Machine Patterns	511
G Part-of-Speech Tags used in the Hepple Tagger	519
References	520

Contents

I	GATE Basics	3
1	Introduction	5
1.1	How to Use this Text	8
1.2	Context	8
1.3	Overview	9
1.3.1	Developing and Deploying Language Processing Facilities	9
1.3.2	Built-In Components	11
1.3.3	Additional Facilities	11
1.3.4	An Example	12
1.4	Some Evaluations	14
1.5	Changes in this Version	15
1.5.1	Version 5.2.1 (May 2010)	15
1.5.2	Version 5.2 (April 2010)	16
1.6	Further Reading	17
2	Installing and Running GATE	27
2.1	Downloading GATE	27
2.2	Installing and Running GATE	27
2.2.1	The Easy Way	27
2.2.2	The Hard Way (1)	28
2.2.3	The Hard Way (2): Subversion	29
2.3	Using System Properties with GATE	29
2.4	Configuring GATE	30
2.5	Building GATE	31
2.5.1	Using GATE with Maven or JPF	32
2.6	Uninstalling GATE	33
2.7	Troubleshooting	33
2.7.1	I don't see the Java console messages under Windows	33
2.7.2	When I execute GATE, nothing happens	33
2.7.3	On Ubuntu, GATE is very slow or doesn't start	33
2.7.4	How to use GATE on a 64 bit system?	34
2.7.5	I got the error: Could not reserve enough space for object heap	34
2.7.6	From Eclipse, I got the error: java.lang.OutOfMemoryError: Java heap space	35

2.7.7	On MacOS, I got the error: java.lang.OutOfMemoryError: Java heap space	35
2.7.8	I got the error: log4j:WARN No appenders could be found for logger...	35
2.7.9	Text is incorrectly refreshed after scrolling and become unreadable . .	36
3	Using GATE Developer	37
3.1	The GATE Developer Main Window	38
3.2	Loading and Viewing Documents	40
3.3	Creating and Viewing Corpora	42
3.4	Working with Annotations	43
3.4.1	The Annotation Sets View	44
3.4.2	The Annotations List View	45
3.4.3	The Annotations Stack View	46
3.4.4	The Co-reference Editor	46
3.4.5	Creating and Editing Annotations	48
3.4.6	Schema-Driven Editing	50
3.4.7	Printing Text with Annotations	51
3.5	Using CREOLE Plugins	52
3.6	Loading and Using Processing Resources	54
3.7	Creating and Running an Application	54
3.7.1	Running an Application on a Datastore	56
3.7.2	Running PRs Conditionally on Document Features	56
3.7.3	Doing Information Extraction with ANNIE	57
3.7.4	Modifying ANNIE	58
3.8	Saving Applications and Language Resources	58
3.8.1	Saving Documents to File	58
3.8.2	Saving and Restoring LR in Datastores	59
3.8.3	Saving Application States to a File	60
3.8.4	Saving an Application with its Resources (e.g. GATE Teamware) . .	61
3.9	Keyboard Shortcuts	62
3.10	Miscellaneous	64
3.10.1	Stopping GATE from Restoring Developer Sessions/Options	64
3.10.2	Working with Unicode	64
4	CREOLE: the GATE Component Model	67
4.1	The Web and CREOLE	68
4.2	The GATE Framework	69
4.3	The Lifecycle of a CREOLE Resource	69
4.4	Processing Resources and Applications	70
4.5	Language Resources and Datastores	71
4.6	Built-in CREOLE Resources	72
4.7	CREOLE Resource Configuration	72
4.7.1	Configuration with XML	72
4.7.2	Configuring Resources using Annotations	78

4.7.3	Mixing the Configuration Styles	82
4.8	Tools: How to Add Utilities to GATE Developer	84
4.8.1	Putting your tools in a sub-menu	85
5	Language Resources: Corpora, Documents and Annotations	87
5.1	Features: Simple Attribute/Value Data	87
5.2	Corpora: Sets of Documents plus Features	88
5.3	Documents: Content plus Annotations plus Features	88
5.4	Annotations: Directed Acyclic Graphs	88
5.4.1	Annotation Schemas	88
5.4.2	Examples of Annotated Documents	90
5.4.3	Creating, Viewing and Editing Diverse Annotation Types	93
5.5	Document Formats	93
5.5.1	Detecting the Right Reader	95
5.5.2	XML	96
5.5.3	HTML	102
5.5.4	SGML	103
5.5.5	Plain text	104
5.5.6	RTF	105
5.5.7	Email	105
5.6	XML Input/Output	107
6	ANNIE: a Nearly-New Information Extraction System	109
6.1	Document Reset	110
6.2	Tokeniser	111
6.2.1	Tokeniser Rules	111
6.2.2	Token Types	112
6.2.3	English Tokeniser	113
6.3	Gazetteer	113
6.4	Sentence Splitter	115
6.5	RegEx Sentence Splitter	116
6.6	Part of Speech Tagger	117
6.7	Semantic Tagger	118
6.8	Orthographic Coreference (OrthoMatcher)	118
6.8.1	GATE Interface	118
6.8.2	Resources	119
6.8.3	Processing	119
6.9	Pronominal Coreference	119
6.9.1	Quoted Speech Submodule	120
6.9.2	Pleonastic It Submodule	120
6.9.3	Pronominal Resolution Submodule	120
6.9.4	Detailed Description of the Algorithm	121
6.10	A Walk-Through Example	125
6.10.1	Step 1 - Tokenisation	125

6.10.2	Step 2 - List Lookup	126
6.10.3	Step 3 - Grammar Rules	126
II	GATE for Advanced Users	127
7	GATE Embedded	129
7.1	Quick Start with GATE Embedded	129
7.2	Resource Management in GATE Embedded	130
7.3	Using CREOLE Plugins	133
7.4	Language Resources	133
7.4.1	GATE Documents	135
7.4.2	Feature Maps	135
7.4.3	Annotation Sets	137
7.4.4	Annotations	140
7.4.5	GATE Corpora	140
7.5	Processing Resources	142
7.6	Controllers	142
7.7	Duplicating a Resource	144
7.8	Persistent Applications	145
7.9	Ontologies	146
7.10	Creating a New Annotation Schema	147
7.11	Creating a New CREOLE Resource	148
7.12	Adding Support for a New Document Format	152
7.13	Using GATE Embedded in a Multithreaded Environment	153
7.14	Using GATE Embedded within a Spring Application	155
7.14.1	Duplication in Spring	156
7.14.2	Spring pooling	157
7.14.3	Further reading	159
7.15	Using GATE Embedded within a Tomcat Web Application	159
7.15.1	Recommended Directory Structure	159
7.15.2	Configuration Files	160
7.15.3	Initialization Code	160
7.16	Groovy for GATE	161
7.16.1	Groovy Scripting Console for GATE	162
7.16.2	Groovy scripting PR	163
7.16.3	Utility methods	165
7.17	Saving Config Data to gate.xml	166
7.18	Annotation merging through the API	166
8	JAPE: Regular Expressions over Annotations	169
8.1	The Left-Hand Side	171
8.1.1	Matching a Simple Text String	171
8.1.2	Matching Entire Annotation Types	172
8.1.3	Using Attributes and Values	173

8.1.4	Using Meta-Properties	174
8.1.5	Multiple Pattern/Action Pairs	174
8.1.6	LHS Macros	175
8.1.7	Using Context	176
8.1.8	Multi-Constraint Statements	178
8.1.9	Negation	179
8.1.10	Escaping Special Characters	180
8.2	LHS Operators in Detail	181
8.2.1	Compositional Operators	181
8.2.2	Matching Operators	182
8.3	The Right-Hand Side	185
8.3.1	A Simple Example	185
8.3.2	Copying Feature Values from the LHS to the RHS	185
8.3.3	RHS Macros	186
8.4	Use of Priority	187
8.5	Using Phases Sequentially	190
8.6	Using Java Code on the RHS	190
8.6.1	A More Complex Example	192
8.6.2	Adding a Feature to the Document	194
8.6.3	Finding the Tokens of a Matched Annotation	194
8.6.4	Using Named Blocks	196
8.6.5	Java RHS Overview	197
8.7	Optimising for Speed	199
8.8	Ontology Aware Grammar Transduction	200
8.9	Serializing JAPE Transducer	200
8.9.1	How to Serialize?	201
8.9.2	How to Use the Serialized Grammar File?	201
8.10	The JAPE Debugger	201
8.11	Notes for Montreal Transducer Users	201
9	ANNIC: ANNotations-In-Context	203
9.1	Instantiating SSD	204
9.2	Search GUI	206
9.2.1	Overview	206
9.2.2	Syntax of Queries	206
9.2.3	Top Section	208
9.2.4	Central Section	209
9.2.5	Bottom Section	209
9.3	Using SSD from GATE Embedded	210
9.3.1	How to instantiate a searchabledatastore	210
9.3.2	How to search in this datastore	210
10	Performance Evaluation of Language Analysers	213
10.1	Metrics for Evaluation in Information Extraction	214

10.1.1	Annotation Relations	214
10.1.2	Cohen's Kappa	215
10.1.3	Precision, Recall, F-Measure	218
10.1.4	Macro and Micro Averaging	219
10.2	The Annotation Diff Tool	220
10.2.1	Performing Evaluation with the Annotation Diff Tool	221
10.3	Corpus Quality Assurance	223
10.3.1	Description of the interface	223
10.3.2	Step by step usage	224
10.3.3	Details of the Corpus statistics table	225
10.3.4	Details of the Document statistics table	225
10.4	Corpus Benchmark Tool	225
10.4.1	Preparing the Corpora for Use	226
10.4.2	Defining Properties	227
10.4.3	Running the Tool	228
10.4.4	The Results	229
10.5	A Plugin Computing Inter-Annotator Agreement (IAA)	230
10.5.1	IAA for Classification	232
10.5.2	IAA For Named Entity Annotation	233
10.5.3	The BDM-Based IAA Scores	234
10.6	A Plugin Computing the BDM Scores for an Ontology	234
11	Profiling Processing Resources	237
11.1	Overview	237
11.1.1	Features	238
11.1.2	Limitations	238
11.2	Graphical User Interface	238
11.3	Command Line Interface	239
11.4	Application Programming Interface	240
11.4.1	Log4j.properties	240
11.4.2	Benchmark log format	241
11.4.3	Enabling profiling	241
11.4.4	Reporting tool	241
12	Developing GATE	245
12.1	Reporting Bugs and Requesting Features	245
12.2	Contributing Patches	245
12.3	Creating New Plugins	246
12.3.1	Where to Keep Plugins in the GATE Hierarchy	246
12.3.2	What to Call your Plugin	246
12.3.3	Writing a New PR	247
12.3.4	Writing a New VR	251
12.3.5	Adding Plugins to the Nightly Build	253
12.4	Updating this User Guide	253

12.4.1	Building the User Guide	254
12.4.2	Making Changes to the User Guide	254
III	CREOLE Plugins	257
13	Gazetteers	259
13.1	Introduction to Gazetteers	259
13.2	ANNIE Gazetteer	260
13.2.1	Creating and Modifying Gazetteer Lists	261
13.3	Gazetteer Visual Resource - GAZE	261
13.3.1	Display Modes	262
13.3.2	Linear Definition Pane	262
13.3.3	Linear Definition Toolbar	262
13.3.4	Operations on Linear Definition Nodes	263
13.3.5	Gazetteer List Pane	263
13.3.6	Mapping Definition Pane	263
13.4	OntoGazetteer	264
13.5	Gaze Ontology Gazetteer Editor	264
13.5.1	The Gaze Gazetteer List and Mapping Editor	264
13.5.2	The Gaze Ontology Editor	264
13.6	Hash Gazetteer	265
13.6.1	Prerequisites	266
13.6.2	Parameters	266
13.7	Flexible Gazetteer	267
13.8	Gazetteer List Collector	268
13.9	OntoRoot Gazetteer	269
13.9.1	How Does it Work?	269
13.9.2	Initialisation of OntoRoot Gazetteer	271
13.9.3	Simple steps to run OntoRoot Gazetteer	272
13.10	Large KB Gazetteer	275
13.10.1	Quick usage overview	275
13.10.2	Dictionary setup	276
13.10.3	Additional dictionary configuration	277
13.10.4	Processing Resource Configuration	277
13.10.5	Runtime configuration	277
13.10.6	Semantic Enrichment PR	278
13.11	The Shared Gazetteer for multithreaded processing	278
14	Working with Ontologies	281
14.1	Data Model for Ontologies	282
14.1.1	Hierarchies of Classes and Restrictions	282
14.1.2	Instances	283
14.1.3	Hierarchies of Properties	284
14.1.4	URIs	286

14.2	Ontology Event Model	286
14.2.1	What Happens when a Resource is Deleted?	288
14.3	The Ontology Plugin: Current Implementation	289
14.3.1	The OWLIMOntology Language Resource	290
14.3.2	The ConnectSesameOntology Language Resource	292
14.3.3	The CreateSesameOntology Language Resource	293
14.3.4	The OWLIM2 Backwards-Compatible Language Resource	294
14.4	The Ontology_OWLIM2 plugin: backwards-compatible implementation	294
14.4.1	The OWLIMOntologyLR Language Resource	294
14.5	GATE Ontology Editor	297
14.6	Ontology Annotation Tool	301
14.6.1	Viewing Annotated Text	302
14.6.2	Editing Existing Annotations	303
14.6.3	Adding New Annotations	304
14.6.4	Options	305
14.7	Using the ontology API	306
14.8	Using the ontology API (old version)	308
14.9	Ontology-Aware JAPE Transducer	309
14.10	Annotating Text with Ontological Information	310
14.11	Populating Ontologies	311
14.12	Ontology API and Implementation Changes	313
14.12.1	Differences between the implementation plugins	313
14.12.2	Changes in the Ontology API	314
15	Machine Learning	317
15.1	ML Generalities	318
15.1.1	Some Definitions	319
15.1.2	GATE-Specific Interpretation of the Above Definitions	319
15.2	Batch Learning PR	319
15.2.1	Batch Learning PR Configuration File Settings	320
15.2.2	Case Studies for the Three Learning Types	334
15.2.3	How to Use the Batch Learning PR in GATE Developer	342
15.2.4	Output of the Batch Learning PR	343
15.2.5	Using the Batch Learning PR from the API	350
15.3	Machine Learning PR	351
15.3.1	The DATASET Element	351
15.3.2	The ENGINE Element	353
15.3.3	The WEKA Wrapper	353
15.3.4	The MAXENT Wrapper	354
15.3.5	The SVM Light Wrapper	355
15.3.6	Example Configuration File	358
16	Tools for Alignment Tasks	367
16.1	Introduction	367

16.2	The Tools	367
16.2.1	Compound Document	368
16.2.2	Compound Document Editor	370
16.2.3	Composite Document	370
16.2.4	DeleteMembersPR	372
16.2.5	SwitchMembersPR	372
16.2.6	Saving as XML	373
16.2.7	Alignment Editor	373
16.2.8	Section-by-Section Processing	378
17	Parsers and Taggers	381
17.1	Verb Group Chunker	381
17.2	Noun Phrase Chunker	381
17.2.1	Differences from the Original	382
17.2.2	Using the Chunker	382
17.3	Tree Tagger	382
17.3.1	POS Tags	384
17.4	TaggerFramework	384
17.5	Chemistry Tagger	387
17.5.1	Using the Tagger	387
17.6	ABNER	387
17.7	Stemmer	389
17.7.1	Algorithms	389
17.8	GATE Morphological Analyzer	389
17.8.1	Rule File	390
17.9	MiniPar Parser	392
17.9.1	Platform Supported	395
17.9.2	Resources	395
17.9.3	Parameters	396
17.9.4	Prerequisites	396
17.9.5	Grammatical Relationships	396
17.10	RASP Parser	397
17.11	SUPPLE Parser	399
17.11.1	Requirements	400
17.11.2	Building SUPPLE	400
17.11.3	Running the Parser in GATE	400
17.11.4	Viewing the Parse Tree	401
17.11.5	System Properties	401
17.11.6	Configuration Files	402
17.11.7	Parser and Grammar	403
17.11.8	Mapping Named Entities	404
17.11.9	Upgrading from BuChart to SUPPLE	404
17.12	Stanford Parser	405
17.12.1	Input Requirements	405

17.12.2	Initialization Parameters	406
17.12.3	Runtime Parameters	406
17.13	OpenCalais, LingPipe and OpenNLP	407
18	Combining GATE and UIMA	409
18.1	Embedding a UIMA AE in GATE	410
18.1.1	Mapping File Format	410
18.1.2	The UIMA Component Descriptor	415
18.1.3	Using the <code>AnalysisEnginePR</code>	415
18.2	Embedding a GATE <code>CorpusController</code> in UIMA	416
18.2.1	Mapping File Format	417
18.2.2	The GATE Application Definition	417
18.2.3	Configuring the <code>GATEApplicationAnnotator</code>	418
19	More (CREOLE) Plugins	421
19.1	Language Plugins	422
19.1.1	French Plugin	422
19.1.2	German Plugin	422
19.1.3	Romanian Plugin	423
19.1.4	Arabic Plugin	423
19.1.5	Chinese Plugin	423
19.1.6	Hindi Plugin	424
19.2	Flexible Exporter	424
19.3	Annotation Set Transfer	425
19.4	Information Retrieval in GATE	427
19.4.1	Using the IR Functionality in GATE	429
19.4.2	Using the IR API	431
19.5	Websphinx Web Crawler	432
19.5.1	Using the Crawler PR	432
19.6	Google Plugin	435
19.7	Yahoo Plugin	435
19.7.1	Using the YahooPR	435
19.8	WordNet in GATE	436
19.8.1	The WordNet API	440
19.9	Kea - Automatic Keyphrase Detection	440
19.9.1	Using the ‘KEA Keyphrase Extractor’ PR	441
19.9.2	Using Kea Corpora	443
19.10	Ontotext JapeC Compiler	444
19.11	Annotation Merging Plugin	445
19.12	Chinese Word Segmentation	447
19.13	Copying Annotations between Documents	449
19.14	OpenCalais Plugin	450
19.15	LingPipe Plugin	451
19.15.1	LingPipe Tokenizer PR	452

19.15.2 LingPipe Sentence Splitter PR	452
19.15.3 LingPipe POS Tagger PR	452
19.15.4 LingPipe NER PR	453
19.15.5 LingPipe Language Identifier PR	453
19.16 OpenNLP Plugin	453
19.16.1 Parameters common to all PRs	454
19.16.2 OpenNLP PRs	454
19.16.3 Training new models	457
19.17 Inter Annotator Agreement	457
19.18 Balanced Distance Metric Computation	457
19.19 Schema Annotation Editor	457

Appendices 457

A Change Log	459
A.1 Version 5.2.1 (May 2010)	459
A.2 Version 5.2 (April 2010)	460
A.2.1 JAPE and JAPE-related	460
A.2.2 Other Changes	460
A.3 Version 5.1 (December 2009)	461
A.3.1 New Features	462
A.3.2 JAPE improvements	464
A.3.3 Other improvements and bug fixes	464
A.4 Version 5.0 (May 2009)	465
A.4.1 Major New Features	465
A.4.2 Other New Features and Improvements	468
A.4.3 Specific Bug Fixes	468
A.5 Version 4.0 (July 2007)	469
A.5.1 Major New Features	469
A.5.2 Other New Features and Improvements	470
A.5.3 Bug Fixes and Optimizations	472
A.6 Version 3.1 (April 2006)	474
A.6.1 Major New Features	474
A.6.2 Other New Features and Improvements	474
A.6.3 Bug Fixes	476
A.7 January 2005	476
A.8 December 2004	477
A.9 September 2004	477
A.10 Version 3 Beta 1 (August 2004)	478
A.11 July 2004	479
A.12 June 2004	479
A.13 April 2004	479
A.14 March 2004	480

A.15	Version 2.2 – August 2003	480
A.16	Version 2.1 – February 2003	481
A.17	June 2002	481
B	Version 5.1 Plugins Name Map	483
C	Design Notes	485
C.1	Patterns	485
C.1.1	Components	486
C.1.2	Model, view, controller	488
C.1.3	Interfaces	489
C.2	Exception Handling	489
D	JAPE: Implementation	493
D.1	Formal Description of the JAPE Grammar	494
D.2	Relation to CPSL	496
D.3	Initialisation of a JAPE Grammar	497
D.4	Execution of JAPE Grammars	499
D.5	Using a Different Java Compiler	501
E	Ant Tasks for GATE	503
E.1	Declaring the Tasks	503
E.2	The <code>packagegapp</code> task - bundling an application with its dependencies . . .	504
E.2.1	Introduction	504
E.2.2	Basic Usage	504
E.2.3	Handling Non-Plugin Resources	505
E.2.4	Streamlining your Plugins	508
E.2.5	Bundling Extra Resources	508
E.3	The <code>expandcreoles</code> Task - Merging Annotation-Driven Config into <code>creole.xml</code>	509
F	Named-Entity State Machine Patterns	511
F.1	<code>Main.jape</code>	511
F.2	<code>first.jape</code>	512
F.3	<code>firstname.jape</code>	513
F.4	<code>name.jape</code>	513
F.4.1	Person	513
F.4.2	Location	513
F.4.3	Organization	514
F.4.4	Ambiguities	514
F.4.5	Contextual information	514
F.5	<code>name_post.jape</code>	514
F.6	<code>date_pre.jape</code>	515
F.7	<code>date.jape</code>	515
F.8	<code>reldate.jape</code>	515
F.9	<code>number.jape</code>	515

F.10 address.jape	516
F.11 url.jape	516
F.12 identifier.jape	516
F.13 jobtitle.jape	516
F.14 final.jape	516
F.15 unknown.jape	517
F.16 name_context.jape	517
F.17 org_context.jape	517
F.18 loc_context.jape	518
F.19 clean.jape	518
G Part-of-Speech Tags used in the Hepple Tagger	519
References	520

Part I

GATE Basics

Chapter 1

Introduction

Software documentation is like sex: when it is good, it is very, very good; and when it is bad, it is better than nothing. (Anonymous.)

There are two ways of constructing a software design: one way is to make it so simple that there are obviously no deficiencies; the other way is to make it so complicated that there are no obvious deficiencies. (C.A.R. Hoare)

A computer language is not just a way of getting a computer to perform operations but rather that it is a novel formal medium for expressing ideas about methodology. Thus, programs must be written for people to read, and only incidentally for machines to execute. (The Structure and Interpretation of Computer Programs, H. Abelson, G. Sussman and J. Sussman, 1985.)

If you try to make something beautiful, it is often ugly. If you try to make something useful, it is often beautiful. (Oscar Wilde)¹

GATE² is an infrastructure for developing and deploying software components that process human language. It is nearly 15 years old and is in active use for all types of computational task involving human language. GATE excels at text analysis of all shapes and sizes. From large corporations to small startups, from €multi-million research consortia to undergraduate projects, our user community is the largest and most diverse of any system of this type, and is spread across all but one of the continents³.

GATE is open source free software; users can obtain free support from the user and developer community via GATE.ac.uk or on a commercial basis from our industrial partners. We are the biggest open source language processing project with a development team more than double the size of the largest comparable projects (many of which are integrated with

¹These were, at least, our ideals; of course we didn't completely live up to them...

²If you've read the overview on GATE.ac.uk you may prefer to skip to Section 1.1.

³Rumours that we're planning to send several of the development team to Antarctica on one-way tickets are false, libellous and wishful thinking.

GATE⁴). More than €5 million has been invested in GATE development⁵; our objective is to make sure that this continues to be money well spent for all GATE's users.

GATE has grown over the years to include a desktop client for developers, a workflow-based web application, a Java library, an architecture and a process. GATE is:

- *an IDE*, **GATE Developer**⁶: an integrated development environment for language processing components bundled with a very widely used Information Extraction system and a comprehensive set of other plugins
- *a web app*, **GATE Teamware**: a collaborative annotation environment for factory-style semantic annotation projects built around a workflow engine and a heavily-optimised backend service infrastructure
- *a framework*, **GATE Embedded**: an object library optimised for inclusion in diverse applications giving access to all the services used by GATE Developer and more
- *an architecture*: a high-level organisational picture of how language processing software composition
- *a process* for the creation of robust and maintainable services.

We also develop:

- a wiki/CMS, **GATE Wiki** (<http://gatewiki.sf.net/>), mainly to host our own websites and as a testbed for some of our experiments
- a cloud computing solution for hosted large-scale text processing, **GATE Cloud** (<http://gatecloud.net/>)

For more information see the family pages.

One of our original motivations was to remove the necessity for solving common engineering problems before doing useful research, or re-engineering before deploying research results into applications. Core functions of GATE take care of the lion's share of the engineering:

- modelling and persistence of specialised data structures
- measurement, evaluation, benchmarking (never believe a computing researcher who hasn't measured their results in a repeatable and open setting!)

⁴Our philosophy is reuse not reinvention, so we integrate and interoperate with other systems e.g.: LingPipe, OpenNLP, UIMA, and many more specific tools.

⁵This is the figure for direct Sheffield-based investment only and therefore an underestimate.

⁶GATE Developer and GATE Embedded are bundled, and in older distributions were referred to just as 'GATE'.

- visualisation and editing of annotations, ontologies, parse trees, etc.
- a finite state transduction language for rapid prototyping and efficient implementation of shallow analysis methods (JAPE)
- extraction of training instances for machine learning
- pluggable machine learning implementations (Weka, SVM Light, ...)

On top of the core functions GATE includes components for diverse language processing tasks, e.g. parsers, morphology, tagging, Information Retrieval tools, Information Extraction components for various languages, and many others. GATE Developer and Embedded are supplied with an Information Extraction system (ANNIE) which has been adapted and evaluated very widely (numerous industrial systems, research systems evaluated in MUC, TREC, ACE, DUC, Pascal, NTCIR, etc.). ANNIE is often used to create RDF or OWL (metadata) for unstructured content (*semantic annotation*).

GATE version 1 was written in the mid-1990s; at the turn of the new millennium we completely rewrote the system in Java; version 5 was released in June 2009. We believe that GATE is the leading system of its type, but as scientists we have to advise you not to take our word for it; that's why we've measured our software in many of the competitive evaluations over the last decade-and-a-half (MUC, TREC, ACE, DUC and more; see Section 1.4 for details). We invite you to give it a try, to get involved with the GATE community, and to contribute to human language science, engineering and development.

This book describes how to use GATE to develop language processing components, test their performance and deploy them as parts of other applications. In the rest of this chapter:

- Section 1.1 describes the best way to use this book;
- Section 1.2 briefly notes that the context of GATE is applied language processing, or *Language Engineering*;
- Section 1.3 gives an overview of developing using GATE;
- Section 1.4 lists publications describing GATE performance in evaluations;
- Section 1.5 outlines what is new in the current version of GATE;
- Section 1.6 lists other publications about GATE.

Note: if you don't see the component you need in this document, or if we mention a component that you can't see in the software, contact gate-users@lists.sourceforge.net⁷ – various components are developed by our collaborators, who we will be happy to put you in contact with. (Often the process of getting a new component is as simple as typing the URL into GATE Developer; the system will do the rest.)

⁷Follow the 'support' link from the GATE web server to subscribe to the mailing list.

1.1 How to Use this Text

The material presented in this book ranges from the conceptual (e.g. ‘what is software architecture?’) to practical instructions for programmers (e.g. how to deal with GATE exceptions) and linguists (e.g. how to write a pattern grammar). Furthermore, GATE’s highly extensible nature means that new functionality is constantly being added in the form of new plugins. Important functionality is as likely to be located in a plugin as it is to be integrated into the GATE core. This presents something of an organisational challenge. Our (no doubt imperfect) solution is to divide this book into three parts. Part I covers installation, using the GATE Developer GUI and using ANNIE, as well as providing some background and theory. We recommend the new user to begin with Part I. Part II covers the more advanced of the core GATE functionality; the GATE Embedded API and JAPE pattern language among other things. Part III provides a reference for the numerous plugins that have been created for GATE. Although ANNIE provides a good starting point, the user will soon wish to explore other resources, and so will need to consult this part of the text. We recommend that Part III be used as a reference, to be dipped into as necessary. In Part III, plugins are grouped into broad areas of functionality.

1.2 Context

GATE can be thought of as a **Software Architecture for Language Engineering** [Cunningham 00].

‘Software Architecture’ is used rather loosely here to mean computer infrastructure for software development, including development environments and frameworks, as well as the more usual use of the term to denote a macro-level organisational structure for software systems [Shaw & Garlan 96].

Language Engineering (LE) may be defined as:

... the discipline or act of engineering software systems that perform tasks involving processing human language. Both the construction process and its outputs are measurable and predictable. The literature of the field relates to both application of relevant scientific results and a body of practice. [Cunningham 99a]

The relevant scientific results in this case are the outputs of Computational Linguistics, Natural Language Processing and Artificial Intelligence in general. Unlike these other disciplines, LE, as an engineering discipline, entails *predictability*, both of the process of constructing LE-based software and of the performance of that software after its completion and deployment in applications.

Some working definitions:

1. **Computational Linguistics (CL)**: science of language that uses computation as an investigative tool.
2. **Natural Language Processing (NLP)**: science of computation whose subject matter is data structures and algorithms for computer processing of human language.
3. **Language Engineering (LE)**: building NLP systems whose cost and outputs are measurable and predictable.
4. **Software Architecture**: macro-level organisational principles for families of systems. In this context is also used as **infrastructure**.
5. **Software Architecture for Language Engineering (SALE)**: software infrastructure, architecture and development tools for applied CL, NLP and LE.

(Of course the practice of these fields is broader and more complex than these definitions.)

In the scientific endeavours of NLP and CL, GATE's role is to support experimentation. In this context GATE's significant features include support for automated measurement (see Chapter 10), providing a 'level playing field' where results can easily be repeated across different sites and environments, and reducing research overheads in various ways.

1.3 Overview

1.3.1 Developing and Deploying Language Processing Facilities

GATE as an architecture suggests that the elements of software systems that process natural language can usefully be broken down into various types of component, known as resources⁸. Components are reusable software chunks with well-defined interfaces, and are a popular architectural form, used in Sun's Java Beans and Microsoft's .Net, for example. GATE components are specialised types of Java Bean, and come in three flavours:

- LanguageResources (LRs) represent entities such as lexicons, corpora or ontologies;
- ProcessingResources (PRs) represent entities that are primarily algorithmic, such as parsers, generators or ngram modellers;
- VisualResources (VRs) represent visualisation and editing components that participate in GUIs.

⁸The terms 'resource' and 'component' are synonymous in this context. 'Resource' is used instead of just 'component' because it is a common term in the literature of the field: cf. the Language Resources and Evaluation conference series [LREC-1 98, LREC-2 00].

These definitions can be blurred in practice as necessary.

Collectively, the set of resources integrated with GATE is known as **CREOLE**: a Collection of REusable Objects for Language Engineering. All the resources are packaged as Java Archive (or ‘JAR’) files, plus some XML configuration data. The JAR and XML files are made available to GATE by putting them on a web server, or simply placing them in the local file space. Section 1.3.2 introduces GATE’s built-in resource set.

When using GATE to develop language processing functionality for an application, the developer uses GATE Developer and GATE Embedded to construct resources of the three types. This may involve programming, or the development of Language Resources such as grammars that are used by existing Processing Resources, or a mixture of both. GATE Developer is used for visualisation of the data structures produced and consumed during processing, and for debugging, performance measurement and so on. For example, figure 1.1 is a screenshot of one of the visualisation tools.

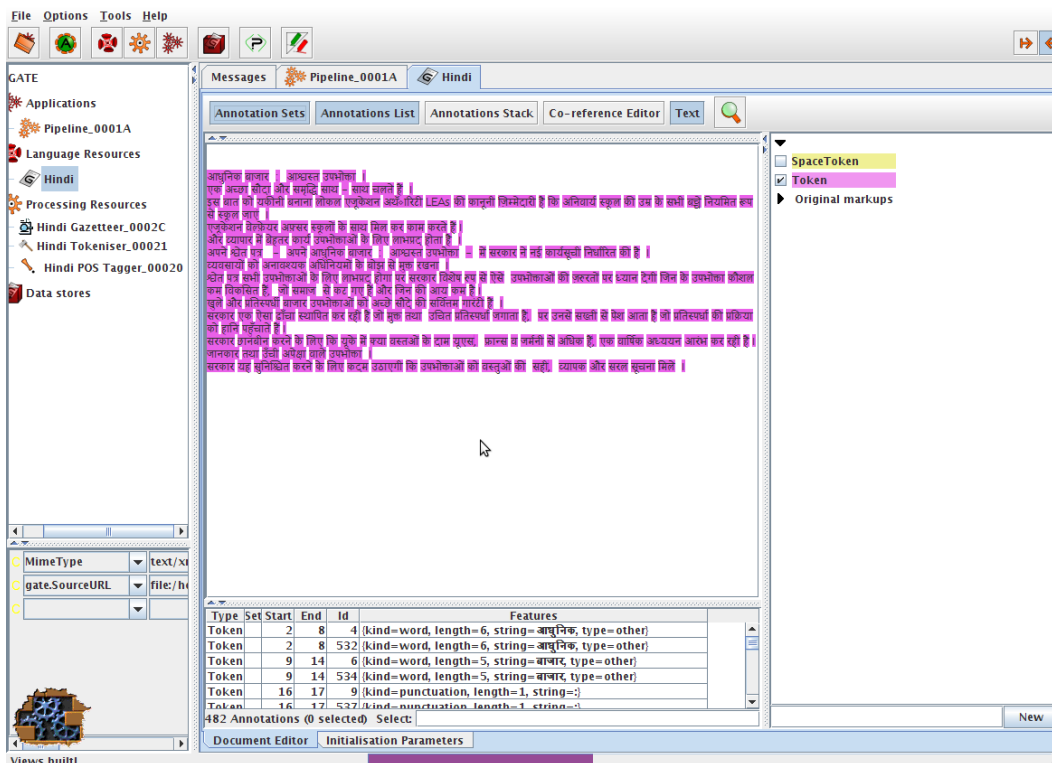


Figure 1.1: One of GATE’s visual resources

GATE Developer is analogous to systems like Mathematica for Mathematicians, or JBuilder for Java programmers: it provides a convenient graphical environment for research and development of language processing software.

When an appropriate set of resources have been developed, they can then be embedded in the target client application using GATE Embedded. GATE Embedded is supplied as a series of JAR files.⁹ To embed GATE-based language processing facilities in an application, these JAR files are all that is needed, along with JAR files and XML configuration files for the various resources that make up the new facilities.

1.3.2 Built-In Components

GATE includes resources for common LE data structures and algorithms, including documents, corpora and various annotation types, a set of language analysis components for Information Extraction and a range of data visualisation and editing components.

GATE supports documents in a variety of formats including XML, RTF, email, HTML, SGML and plain text. In all cases the format is analysed and converted into a single unified model of *annotation*. The annotation format is a modified form of the TIPSTER format [Grishman 97] which has been made largely compatible with the Atlas format [Bird & Liberman 99], and uses the now standard mechanism of ‘stand-off markup’. GATE documents, corpora and annotations are stored in databases of various sorts, visualised via the development environment, and accessed at code level via the framework. See Chapter 5 for more details of corpora etc.

A family of Processing Resources for language analysis is included in the shape of ANNIE, A Nearly-New Information Extraction system. These components use finite state techniques to implement various tasks from tokenisation to semantic tagging or verb phrase chunking. All ANNIE components communicate exclusively via GATE’s document and annotation resources. See Chapter 6 for more details. Other CREOLE resources are described in Part III.

1.3.3 Additional Facilities

Three other facilities in GATE deserve special mention:

- JAPE, a Java Annotation Patterns Engine, provides regular-expression based pattern/action rules over annotations – see Chapter 8.
- The ‘annotation diff’ tool in the development environment implements performance metrics such as precision and recall for comparing annotations. Typically a language analysis component developer will mark up some documents by hand and then use these

⁹The main JAR file (**gate.jar**) supplies the framework. Built-in resources and various 3rd-party libraries are supplied as separate JARs; for example (**guk.jar**, the GATE Unicode Kit.) contains Unicode support (e.g. additional input methods for languages not currently supported by the JDK). They are separate because the latter has to be a Java extension with a privileged security profile.

along with the diff tool to automatically measure the performance of the components. See Chapter 10.

- GUK, the GATE Unicode Kit, fills in some of the gaps in the JDK's¹⁰ support for Unicode, e.g. by adding input methods for various languages from Urdu to Chinese. See Section 3.10.2 for more details.

And by version 4 it will make a mean cup of tea.

1.3.4 An Example

This section gives a very brief example of a typical use of GATE to develop and deploy language processing capabilities in an application, and to generate quantitative results for scientific publication.

Let's imagine that a developer called Fatima is building an email client¹¹ for Cyberdyne Systems' large corporate Intranet. In this application she would like to have a language processing system that automatically spots the names of people in the corporation and transforms them into `mailto` hyperlinks.

A little investigation shows that GATE's existing components can be tailored to this purpose. Fatima starts up GATE Developer, and creates a new document containing some example emails. She then loads some processing resources that will do named-entity recognition (a tokeniser, gazetteer and semantic tagger), and creates an application to run these components on the document in sequence. Having processed the emails, she can see the results in one of several viewers for annotations.

The GATE components are a decent start, but they need to be altered to deal specially with people from Cyberdyne's personnel database. Therefore Fatima creates new 'cyber-' versions of the gazetteer and semantic tagger resources, using the 'bootstrap' tool. This tool creates a directory structure on disk that has some Java stub code, a Makefile and an XML configuration file. After several hours struggling with badly written documentation, Fatima manages to compile the stubs and create a JAR file containing the new resources. She tells GATE Developer the URL of these files¹², and the system then allows her to load them in the same way that she loaded the built-in resources earlier on.

Fatima then creates a second copy of the email document, and uses the annotation editing facilities to mark up the results that she would like to see her system producing. She saves

¹⁰JDK: Java Development Kit, Sun Microsystem's Java implementation. Unicode support is being actively improved by Sun, but at the time of writing many languages are still unsupported. In fact, Unicode itself doesn't support all languages, e.g. Sylheti; hopefully this will change in time.

¹¹Perhaps because Outlook Express trashed her mail folder again, or because she got tired of Microsoft-specific viruses and hadn't heard of Netscape or Emacs.

¹²While developing, she uses a `file:/...` URL; for deployment she can put them on a web server.

this and the version that she ran GATE on into her serial datastore. From now on she can follow this routine:

1. Run her application on the email test corpus.
2. Check the performance of the system by running the ‘annotation diff’ tool to compare her manual results with the system’s results. This gives her both percentage accuracy figures and a graphical display of the differences between the machine and human outputs.
3. Make edits to the code, pattern grammars or gazetteer lists in her resources, and recompile where necessary.
4. Tell GATE Developer to re-initialise the resources.
5. Go to 1.

To make the alterations that she requires, Fatima re-implements the ANNIE gazetteer so that it regenerates itself from the local personnel data. She then alters the pattern grammar in the semantic tagger to prioritise recognition of names from that source. This latter job involves learning the JAPE language (see Chapter 8), but as this is based on regular expressions it isn’t too difficult.

Eventually the system is running nicely, and her accuracy is 93% (there are still some problem cases, e.g. when people use nicknames, but the performance is good enough for production use). Now Fatima stops using GATE Developer and works instead on embedding the new components in her email application using GATE Embedded. This application is written in Java, so embedding is very easy¹³: the GATE JAR files are added to the project CLASSPATH, the new components are placed on a web server, and with a little code to do initialisation, loading of components and so on, the job is finished in half a day – the code to talk to GATE takes up only around 150 lines of the eventual application, most of which is just copied from the example in the `sheffield.examples.StandAloneAnnie` class.

Because Fatima is worried about Cyberdyne’s unethical policy of developing Skynet to help the large corporates of the West strengthen their strangle-hold over the World, she wants to get a job as an academic instead (so that her conscience will only have to cope with the torture of students, as opposed to humanity). She takes the accuracy measures that she has attained for her system and writes a paper for the *Journal of Nasturtium Logarithm Incitement* describing the approach used and the results obtained. Because she used GATE for development, she can cite the repeatability of her experiments and offer access to example binary versions of her software by putting them on an external web server.

And everybody lived happily ever after.

¹³Languages other than Java require an additional interface layer, such as JNI, the Java Native Interface, which is in C.

1.4 Some Evaluations

This section contains an incomplete list of publications describing systems that used GATE in competitive quantitative evaluation programmes. These programmes have had a significant impact on the language processing field and the widespread presence of GATE is some measure of the maturity of the system and of our understanding of its likely performance on diverse text processing tasks.

[Li *et al.* 07d] describes the performance of an SVM-based learning system in the NTCIR-6 Patent Retrieval Task. The system achieved the best result on two of three measures used in the task evaluation, namely the R-Precision and F-measure. The system obtained close to the best result on the remaining measure (A-Precision).

[Saggion 07] describes a cross-source coreference resolution system based on semantic clustering. It uses GATE for information extraction and the SUMMA system to create summaries and semantic representations of documents. One system configuration ranked 4th in the Web People Search 2007 evaluation.

[Saggion 06] describes a cross-lingual summarization system which uses SUMMA components and the Arabic plugin available in GATE to produce summaries in English from a mixture of English and Arabic documents.

Open-Domain Question Answering: The University of Sheffield has a long history of research into open-domain question answering. GATE has formed the basis of much of this research resulting in systems which have ranked highly during independent evaluations since 1999. The first successful question answering system developed at the University of Sheffield was evaluated as part of TREC 8 and used the LaSIE information extraction system (the forerunner of ANNIE) which was distributed with GATE [Humphreys *et al.* 99]. Further research was reported in [Scott & Gaizauskas. 00], [Greenwood *et al.* 02], [Gaizauskas *et al.* 03], [Gaizauskas *et al.* 04] and [Gaizauskas *et al.* 05]. In 2004 the system was ranked 9th out of 28 participating groups.

[Saggion 04] describes techniques for answering definition questions. The system uses definition patterns manually implemented in GATE as well as learned JAPE patterns induced from a corpus. In 2004, the system was ranked 4th in the TREC/QA evaluations.

[Saggion & Gaizauskas 04b] describes a multidocument summarization system implemented using summarization components compatible with GATE (the SUMMA system). The system was ranked 2nd in the Document Understanding Evaluation programmes.

[Maynard *et al.* 03e] and [Maynard *et al.* 03d] describe participation in the TIDES surprise language program. ANNIE was adapted to Cebuano with four person days of

effort, and achieved an F-measure of 77.5%. Unfortunately, ours was the only system participating!

[Maynard *et al.* 02b] and [Maynard *et al.* 03b] describe results obtained on systems designed for the ACE task (Automatic Content Extraction). Although a comparison to other participating systems cannot be revealed due to the stipulations of ACE, results show 82%-86% precision and recall.

[Humphreys *et al.* 98] describes the LaSIE-II system used in MUC-7.

[Gaizauskas *et al.* 95] describes the LaSIE-II system used in MUC-6.

1.5 Changes in this Version

This section logs changes in the latest version of GATE. Appendix A provides a complete change log.

1.5.1 Version 5.2.1 (May 2010)

This is a bugfix release to resolve several bugs that were reported shortly after the release of version 5.2:

- Fixed some bugs with the automatic “create instance” feature in OAT (the ontology annotation tool) when used with the new `Ontology` plugin.
- Added validation to datatype property values of the *date*, *time* and *datetime* types.
- Fixed a bug with `Gazetteer_LKB` that prevented it working when the `dictionaryPath` contained spaces.
- Added a utility class to handle common cases of encoding URIs for use in ontologies, and fixed the example code to show how to make use of this. See chapter 14 for details.
- The annotation set transfer PR now copies the feature map of each annotation it transfers, rather than re-using the same `FeatureMap` (this means that when used to copy annotations rather than move them, the copied annotation is independent from the original and modifying the features of one does not modify the other). See section 19.3 for details.
- The Log4J log files are now created by default in the `.gate` directory under the user’s home directory, rather than being created in the current directory when GATE starts, to be more friendly when GATE is installed in a shared location where the user does not have write permission.

This release also fixes some shortcomings in the Groovy support added by 5.2, in particular:

- The `corpora` variable in the console now includes persistent corpora (loaded from a datastore) as well as transient corpora.
- The subscript notation for annotation sets works with long values as well as ints, so `someAS[annotation.start()..annotation.end()]` works as expected.

1.5.2 Version 5.2 (April 2010)

JAPE and JAPE-related

Introduced a utility class `gate.Utils` containing static utility methods for frequently-used idioms such as getting the string covered by an annotation, finding the start and end offsets of annotations and sets, etc. This class is particularly useful on the right hand side of JAPE rules (section 8.6.5).

Added type parameters to the `bindings` map available on the RHS of JAPE rules, so you can now do `AnnotationSet as = bindings.get("label")` without a cast (see section 8.6.5).

Fixed a bug with JAPE's handling of features called "class" in non-ontology-aware mode. Previously JAPE would always match such features using an equality test, even if a different operator was used in the grammar, i.e. `{SomeType.class != "foo"}` was matched as `{SomeType.class == "foo"}`. The correct operator is now used. Note that this does not affect the ontology-aware behaviour: when an ontology parameter is specified, "class" features are always matched using ontology subsumption.

Custom JAPE operators and annotation accessors can now be loaded from plugins as well as from the `lib` directory (see section 8.2.2).

Other Changes

Added a mechanism to allow plugins to contribute menu items to the "Tools" menu in GATE Developer. See section 4.8 for details.

Enhanced Groovy support in GATE: the Groovy console and Groovy Script PR (in the Groovy plugin) now import many GATE classes by default, and a number of utility methods are mixed in to some of the core GATE API classes to make them more natural to use in Groovy. See section 7.16 for details.

Modified the batch learning PR (in the `Learning` plugin) to make it safe to use several instances in `APPLICATION` mode with the same configuration file and the same learned model at the same time (e.g. in a multithreaded application). The other modes (including

training and evaluation) are unchanged, and thus are still *not* safe for use in this way. Also fixed a bug that prevented APPLICATION mode from working anywhere other than as the last PR in a pipeline when running over a corpus in a datastore.

Introduced a simple way to create duplicate copies of an existing resource instance, with a way for individual resource types to override the default duplication algorithm if they know a better way to deal with duplicating themselves. See section 7.7.

Enhanced the Spring support in GATE to provide easy access to the new duplication API, and to simplify the configuration of the built-in Spring pooling mechanisms when writing multi-threaded Spring-based applications. See section 7.14.

The GAPP packager Ant task now respects the ordering of mapping hints, with earlier hints taking precedence over later ones (see section E.2.3).

Bug fix in the UIMA plugin from Roland Cornelissen - `AnalysisEnginePR` now properly shuts down the wrapped `AnalysisEngine` when the PR is deleted.

Patch from Matt Nathan to allow several instances of a gazetteer PR in an embedded application to share a single copy of their internal data structures, saving considerable memory compared to loading several complete copies of the same gazetteer lists (see section 13.11).

1.6 Further Reading

Lots of documentation lives on the GATE web server, including:

- movies of the system in operation;
- the main system documentation tree;
- JavaDoc API documentation;
- HTML of the source code;
- parts of the requirements analysis that version 3 was based on.

For more details about Sheffield University's work in human language processing see the NLP group pages or *A Definition and Short History of Language Engineering* ([Cunningham 99a]). For more details about Information Extraction see *IE, a User Guide* or the GATE IE pages.

A list of publications on GATE and projects that use it (some of which are available on-line):

2009

[**Bontcheva et al. 09**] is the ‘Human Language Technologies’ chapter of ‘Semantic Knowledge Management’ (John Davies, Marko Grobelnik and Dunja Mladeni eds.)

[**Damljanovic et al. 09**] - to appear.

[**Laclavik & Maynard 09**] reviews the current state of the art in email processing and communication research, focusing on the roles played by email in information management, and commercial and research efforts to integrate a semantic-based approach to email.

[**Li et al. 09**] investigates two techniques for making SVMs more suitable for language learning tasks. Firstly, an SVM with uneven margins (SVMUM) is proposed to deal with the problem of imbalanced training data. Secondly, SVM active learning is employed in order to alleviate the difficulty in obtaining labelled training data. The algorithms are presented and evaluated on several Information Extraction (IE) tasks.

2008

[**Agatonovic et al. 08**] presents our approach to automatic patent enrichment, tested in large-scale, parallel experiments on USPTO and EPO documents.

[**Damljanovic et al. 08**] presents Question-based Interface to Ontologies (QuestIO) - a tool for querying ontologies using unconstrained language-based queries.

[**Damljanovic & Bontcheva 08**] presents a semantic-based prototype that is made for an open-source software engineering project with the goal of exploring methods for assisting open-source developers and software users to learn and maintain the system without major effort.

[**Della Valle et al. 08**] presents ServiceFinder.

[**Li & Cunningham 08**] describes our SVM-based system and several techniques we developed successfully to adapt SVM for the specific features of the F-term patent classification task.

[**Li & Bontcheva 08**] reviews the recent developments in applying geometric and quantum mechanics methods for information retrieval and natural language processing.

[**Maynard 08**] investigates the state of the art in automatic textual annotation tools, and examines the extent to which they are ready for use in the real world.

[**Maynard et al. 08a**] discusses methods of measuring the performance of ontology-based information extraction systems, focusing particularly on the Balanced Distance Metric (BDM), a new metric we have proposed which aims to take into account the more flexible nature of ontologically-based applications.

[**Maynard et al. 08b**] investigates NLP techniques for ontology population, using a combination of rule-based approaches and machine learning.

[**Tablan et al. 08**] presents the QuestIO system – a natural language interface for accessing structured information, that is domain independent and easy to use without training.

2007

[**Funk et al. 07a**] describes an ontologically based approach to multi-source, multilingual information extraction.

[**Funk et al. 07b**] presents a controlled language for ontology editing and a software implementation, based partly on standard NLP tools, for processing that language and manipulating an ontology.

[**Maynard et al. 07a**] proposes a methodology to capture (1) the evolution of metadata induced by changes to the ontologies, and (2) the evolution of the ontology induced by changes to the underlying metadata.

[**Maynard et al. 07b**] describes the development of a system for content mining using domain ontologies, which enables the extraction of relevant information to be fed into models for analysis of financial and operational risk and other business intelligence applications such as company intelligence, by means of the XBRL standard.

[**Saggion 07**] describes experiments for the cross-document coreference task in SemEval 2007. Our cross-document coreference system uses an in-house agglomerative clustering implementation to group documents referring to the same entity.

[**Saggion et al. 07**] describes the application of ontology-based extraction and merging in the context of a practical e-business application for the EU MUSING Project where the goal is to gather international company intelligence and country/region information.

[**Li et al. 07a**] introduces a hierarchical learning approach for IE, which uses the target ontology as an essential part of the extraction process, by taking into account the relations between concepts.

[**Li et al. 07b**] proposes some new evaluation measures based on relations among classification labels, which can be seen as the label relation sensitive version of important measures such as averaged precision and F-measure, and presents the results of applying the new evaluation measures to all submitted runs for the NTCIR-6 F-term patent classification task.

[**Li et al. 07c**] describes the algorithms and linguistic features used in our participating system for the opinion analysis pilot task at NTCIR-6.

[**Li et al. 07d**] describes our SVM-based system and the techniques we used to adapt the approach for the specifics of the F-term patent classification subtask at NTCIR-6 Patent Retrieval Task.

[**Li & Shawe-Taylor 07**] studies Japanese-English cross-language patent retrieval using Kernel Canonical Correlation Analysis (KCCA), a method of correlating linear relationships between two variables in kernel defined feature spaces.

2006

[**Aswani et al. 06**] (Proceedings of the 5th International Semantic Web Conference (ISWC2006)) In this paper the problem of disambiguating author instances in ontology is addressed. We describe a web-based approach that uses various features such as publication titles, abstract, initials and co-authorship information.

[**Bontcheva et al. 06a**] ‘Semantic Annotation and Human Language Technology’, contribution to ‘Semantic Web Technology: Trends and Research’ (Davies, Studer and Warren, eds.)

[**Bontcheva et al. 06b**] ‘Semantic Information Access’, contribution to ‘Semantic Web Technology: Trends and Research’ (Davies, Studer and Warren, eds.)

[**Bontcheva & Sabou 06**] presents an ontology learning approach that 1) exploits a range of information sources associated with software projects and 2) relies on techniques that are portable across application domains.

[**Davis et al. 06**] describes work in progress concerning the application of Controlled Language Information Extraction - CLIE to a Personal Semantic Wiki - Semper- Wiki, the goal being to permit users who have no specialist knowledge in ontology tools or languages to semi-automatically annotate their respective personal Wiki pages.

[**Li & Shawe-Taylor 06**] studies a machine learning algorithm based on KCCA for cross-language information retrieval. The algorithm is applied to Japanese-English cross-language information retrieval.

[**Maynard et al. 06**] discusses existing evaluation metrics, and proposes a new method for evaluating the ontology population task, which is general enough to be used in a variety of situation, yet more precise than many current metrics.

[**Tablan et al. 06a**] describes an approach that allows users to create and edit ontologies simply by using a restricted version of the English language. The controlled language described is based on an open vocabulary and a restricted set of grammatical constructs.

[**Tablan et al. 06b**] describes the creation of linguistic analysis and corpus search tools for Sumerian, as part of the development of the ETCSL.

[**Wang et al. 06**] proposes an SVM based approach to hierarchical relation extraction, using features derived automatically from a number of GATE-based open-source language processing tools.

2005

- [**Aswani et al. 05**] (Proceedings of Fifth International Conference on Recent Advances in Natural Language Processing (RANLP2005)) It is a full-featured annotation indexing and search engine, developed as a part of the GATE. It is powered with Apache Lucene technology and indexes a variety of documents supported by the GATE.
- [**Bontcheva 05**] presents the ONTOSUM system which uses Natural Language Generation (NLG) techniques to produce textual summaries from Semantic Web ontologies.
- [**Cunningham 05**] is an overview of the field of Information Extraction for the 2nd Edition of the Encyclopaedia of Language and Linguistics.
- [**Cunningham & Bontcheva 05**] is an overview of the field of Software Architecture for Language Engineering for the 2nd Edition of the Encyclopaedia of Language and Linguistics.
- [**Dowman et al. 05a**] (Euro Interactive Television Conference Paper) A system which can use material from the Internet to augment television news broadcasts.
- [**Dowman et al. 05b**] (World Wide Web Conference Paper) The Web is used to assist the annotation and indexing of broadcast news.
- [**Dowman et al. 05c**] (Second European Semantic Web Conference Paper) A system that semantically annotates television news broadcasts using news websites as a resource to aid in the annotation process.
- [**Li et al. 05a**] (Proceedings of Sheffield Machine Learning Workshop) describe an SVM based IE system which uses the SVM with uneven margins as learning component and the GATE as NLP processing module.
- [**Li et al. 05b**] (Proceedings of Ninth Conference on Computational Natural Language Learning (CoNLL-2005)) uses the uneven margins versions of two popular learning algorithms SVM and Perceptron for IE to deal with the imbalanced classification problems derived from IE.
- [**Li et al. 05c**] (Proceedings of Fourth SIGHAN Workshop on Chinese Language processing (Sighan-05)) a system for Chinese word segmentation based on Perceptron learning, a simple, fast and effective learning algorithm.
- [**Polajnar et al. 05**] (University of Sheffield-Research Memorandum CS-05-10) User-Friendly Ontology Authoring Using a Controlled Language.
- [**Saggion & Gaizauskas 05**] describes experiments on content selection for producing biographical summaries from multiple documents.
- [**Ursu et al. 05**] (Proceedings of the 2nd European Workshop on the Integration of Knowledge, Semantic and Digital Media Technologies (EWIMT 2005)) Digital Media Preservation and Access through Semantically Enhanced Web-Annotation.

[**Wang et al. 05**] (Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence (WI 2005)) Extracting a Domain Ontology from Linguistic Resource Based on Relatedness Measurements.

2004

[**Bontcheva 04**] (LREC 2004) describes lexical and ontological resources in GATE used for Natural Language Generation.

[**Bontcheva et al. 04**] (JNLE) discusses developments in GATE in the early naughties.

[**Cunningham & Scott 04a**] (JNLE) is the introduction to the above collection.

[**Cunningham & Scott 04b**] (JNLE) is a collection of papers covering many important areas of Software Architecture for Language Engineering.

[**Dimitrov et al. 04**] (Anaphora Processing) gives a lightweight method for named entity coreference resolution.

[**Li et al. 04**] (Machine Learning Workshop 2004) describes an SVM based learning algorithm for IE using GATE.

[**Maynard et al. 04a**] (LREC 2004) presents algorithms for the automatic induction of gazetteer lists from multi-language data.

[**Maynard et al. 04b**] (ESWS 2004) discusses ontology-based IE in the hTechSight project.

[**Maynard et al. 04c**] (AIMSA 2004) presents automatic creation and monitoring of semantic metadata in a dynamic knowledge portal.

[**Saggion & Gaizauskas 04a**] describes an approach to mining definitions.

[**Saggion & Gaizauskas 04b**] describes a sentence extraction system that produces two sorts of multi-document summaries; a general-purpose summary of a cluster of related documents and an entity-based summary of documents related to a particular person.

[**Wood et al. 04**] (NLDB 2004) looks at ontology-based IE from parallel texts.

2003

[**Bontcheva et al. 03**] (NLPXML-2003) looks at GATE for the semantic web.

[**Cunningham et al. 03**] (Corpus Linguistics 2003) describes GATE as a tool for collaborative corpus annotation.

[**Kiryakov 03**] (Technical Report) discusses semantic web technology in the context of multimedia indexing and search.

- [**Manov et al. 03**] (HLT-NAACL 2003) describes experiments with geographic knowledge for IE.
- [**Maynard et al. 03a**] (EACL 2003) looks at the distinction between information and content extraction.
- [**Maynard et al. 03c**] (Recent Advances in Natural Language Processing 2003) looks at semantics and named-entity extraction.
- [**Maynard et al. 03e**] (ACL Workshop 2003) describes NE extraction without training data on a language you don't speak (!).
- [**Saggion et al. 03a**] (EACL 2003) discusses robust, generic and query-based summarisation.
- [**Saggion et al. 03b**] (Data and Knowledge Engineering) discusses multimedia indexing and search from multisource multilingual data.
- [**Saggion et al. 03c**] (EACL 2003) discusses event co-reference in the MUMIS project.
- [**Tablan et al. 03**] (HLT-NAACL 2003) presents the OLLIE on-line learning for IE system.
- [**Wood et al. 03**] (Recent Advances in Natural Language Processing 2003) discusses using parallel texts to improve IE recall.

2002

- [**Baker et al. 02**] (LREC 2002) report results from the EMILLE Indic languages corpus collection and processing project.
- [**Bontcheva et al. 02a**] (ACL 2002 Workshop) describes how GATE can be used as an environment for teaching NLP, with examples of and ideas for future student projects developed within GATE.
- [**Bontcheva et al. 02b**] (NLIS 2002) discusses how GATE can be used to create HLT modules for use in information systems.
- [**Bontcheva et al. 02c**], [**Dimitrov 02a**] and [**Dimitrov 02b**] (TALN 2002, DAARC 2002, MSc thesis) describe the shallow named entity coreference modules in GATE: the orthomatcher which resolves pronominal coreference, and the pronoun resolution module.
- [**Cunningham 02**] (Computers and the Humanities) describes the philosophy and motivation behind the system, describes GATE version 1 and how well it lived up to its design brief.
- [**Cunningham et al. 02**] (ACL 2002) describes the GATE framework and graphical development environment as a tool for robust NLP applications.

- [**Dimitrov 02a**, **Dimitrov et al. 02**] (DAARC 2002, MSc thesis) discuss lightweight coreference methods.
- [**Lal 02**] (Master Thesis) looks at text summarisation using GATE.
- [**Lal & Ruger 02**] (ACL 2002) looks at text summarisation using GATE.
- [**Maynard et al. 02a**] (ACL 2002 Summarisation Workshop) describes using GATE to build a portable IE-based summarisation system in the domain of health and safety.
- [**Maynard et al. 02c**] (AIMSA 2002) describes the adaptation of the core ANNIE modules within GATE to the ACE (Automatic Content Extraction) tasks.
- [**Maynard et al. 02d**] (Nordic Language Technology) describes various Named Entity recognition projects developed at Sheffield using GATE.
- [**Maynard et al. 02e**] (JNLE) describes robustness and predictability in LE systems, and presents GATE as an example of a system which contributes to robustness and to low overhead systems development.
- [**Pastra et al. 02**] (LREC 2002) discusses the feasibility of grammar reuse in applications using ANNIE modules.
- [**Saggion et al. 02b**] and [**Saggion et al. 02a**] (LREC 2002, SPLPT 2002) describes how ANNIE modules have been adapted to extract information for indexing multimedia material.
- [**Tablan et al. 02**] (LREC 2002) describes GATE's enhanced Unicode support.

Older than 2002

- [**Maynard et al. 01**] (RANLP 2001) discusses a project using ANNIE for named-entity recognition across wide varieties of text type and genre.
- [**Bontcheva et al. 00**] and [**Brugman et al. 99**] (COLING 2000, technical report) describe a prototype of GATE version 2 that integrated with the EUDICO multimedia markup tool from the Max Planck Institute.
- [**Cunningham 00**] (PhD thesis) defines the field of Software Architecture for Language Engineering, reviews previous work in the area, presents a requirements analysis for such systems (which was used as the basis for designing GATE versions 2 and 3), and evaluates the strengths and weaknesses of GATE version 1.
- [**Cunningham et al. 00a**], [**Cunningham et al. 98a**] and [**Peters et al. 98**] (OntoLex 2000, LREC 1998) presents GATE's model of Language Resources, their access and distribution.

- [**Cunningham et al. 00b**] (LREC 2000) taxonomises Language Engineering components and discusses the requirements analysis for GATE version 2.
- [**Cunningham et al. 00c**] and [**Cunningham et al. 99**] (COLING 2000, AISB 1999) summarise experiences with GATE version 1.
- [**Cunningham et al. 00d**] and [**Cunningham 99b**] (technical reports) document early versions of JAPE (superseded by the present document).
- [**Gambäck & Olsson 00**] (LREC 2000) discusses experiences in the Svensk project, which used GATE version 1 to develop a reusable toolbox of Swedish language processing components.
- [**Maynard et al. 00**] (technical report) surveys users of GATE up to mid-2000.
- [**McEnergy et al. 00**] (Vivek) presents the EMILLE project in the context of which GATE's Unicode support for Indic languages has been developed.
- [**Cunningham 99a**] (JNLE) reviewed and synthesised definitions of Language Engineering.
- [**Stevenson et al. 98**] and [**Cunningham et al. 98b**] (ECAI 1998, NeMLaP 1998) report work on implementing a word sense tagger in GATE version 1.
- [**Cunningham et al. 97b**] (ANLP 1997) presents motivation for GATE and GATE-like infrastructural systems for Language Engineering.
- [**Cunningham et al. 96a**] (manual) was the guide to developing CREOLE components for GATE version 1.
- [**Cunningham et al. 96b**] (TIPSTER) discusses a selection of projects in Sheffield using GATE version 1 and the TIPSTER architecture it implemented.
- [**Cunningham et al. 96c, Cunningham et al. 96d, Cunningham et al. 95**] (COLING 1996, AISB Workshop 1996, technical report) report early work on GATE version 1.
- [**Gaizauskas et al. 96a**] (manual) was the user guide for GATE version 1.
- [**Gaizauskas et al. 96b, Cunningham et al. 97a, Cunningham et al. 96e**] (ICTAI 1996, TIPSTER 1997, NeMLaP 1996) report work on GATE version 1.
- [**Humphreys et al. 96**] (manual) describes the language processing components distributed with GATE version 1.
- [**Cunningham 94, Cunningham et al. 94**] (NeMLaP 1994, technical report) argue that software engineering issues such as reuse, and framework construction, are important for language processing R&D.

Chapter 2

Installing and Running GATE

2.1 Downloading GATE

To download GATE point your web browser at <http://gate.ac.uk/download/>.

2.2 Installing and Running GATE

GATE will run anywhere that supports Java 5 or later, including Solaris, Linux, Mac OS X and Windows platforms. We don't run tests on other platforms, but have had reports of successful installs elsewhere.

2.2.1 The Easy Way

The easy way to install is to use one of the platform-specific installers (created using the excellent IzPack). Download a 'platform-specific installer' and follow the instructions it gives you. Once the installation is complete, you can start GATE Developer using `gate.exe` (Windows) or `GATE.app` (Mac) in the top-level installation directory, or `gate.sh` in the bin directory (other platforms).

Note for Mac users: on 64-bit-capable systems, `GATE.app` will run as a 64-bit application. It will use the first listed *64-bit* JVM in your Java Preferences, even if your highest priority JVM is a 32-bit one. Thus if you want to run using Java 5 rather than 6 you must ensure that "J2SE 5.0 64-bit" is listed ahead of "Java SE 6 64-bit".

2.2.2 The Hard Way (1)

Download the Java-only release package or the binary build snapshot, and follow the instructions below.

Prerequisites:

- A conforming Java 2 environment,
 - version 1.4.2 or above for GATE 3.1
 - version 5.0 for GATE 4.0 beta 1 or later.

available free from Sun Microsystems or from your UNIX supplier. (We test on various Sun JDKs on Solaris, Linux and Windows XP.)

- Binaries from the GATE distribution you downloaded: `gate.jar`, `lib/ext/guk.jar` (Unicode editing support) and a suitable script to start Ant, e.g. `ant.sh` or `ant.bat`. These are held in a directory called `bin` like this:

```
.../bin/
  gate.jar
  ant.sh
  ant.bat
```

You will also need the `lib` directory, containing various libraries that GATE depends on.

- An open mind and a sense of humour.

Using the binary distribution:

- Unpack the distribution, creating a directory containing jar files and scripts.
- To run GATE Developer: on Windows, start a Command Prompt window, change to the directory where you unpacked the GATE distribution and run `'bin/ant.bat run'`; on UNIX or mac open a terminal window and run `'bin/ant run'`.
- To embed GATE as a library (GATE Embedded), put `gate.jar` and all the libraries in the `lib` directory in your `CLASSPATH` and tell Java that `guk.jar` is an extension (`-Djava.ext.dirs=path-to-guk.jar`).

The Ant scripts that start GATE Developer (`ant.bat` or `ant`) require you to set the `JAVA_HOME` environment variable to point to the top level directory of your JAVA installation. The value of `GATE_CONFIG` is passed to the system by the scripts using either a `-i` command-line option, or the Java property `gate.config`.

2.2.3 The Hard Way (2): Subversion

The GATE code is maintained in a Subversion repository. You can use a Subversion client to check out the source code – the most up-to-date version of GATE is the trunk:

```
svn checkout https://gate.svn.sourceforge.net/svnroot/gate/gate/trunk gate
```

Once you have checked out the code you can build GATE using Ant (see Section 2.5)

You can browse the complete Subversion repository online at <http://gate.svn.sourceforge.net/>.

2.3 Using System Properties with GATE

During initialisation, GATE reads several Java system properties in order to decide where to find its configuration files.

Here is a list of the properties used, their default values and their meanings:

gate.home sets the location of the GATE install directory. This should point to the top level directory of your GATE installation. This is the only property that is required. If this is not set, the system will display an error message and then it will attempt to guess the correct value.

gate.plugins.home points to the location of the directory containing installed plugins (a.k.a. CREOLE directories). If this is not set then the default value of `{gate.home}/plugins` is used.

gate.site.config points to the location of the configuration file containing the site-wide options. If not set this will default to `{gate.home}/gate.xml`. The site configuration file must exist!

gate.user.config points to the file containing the user's options. If not specified, or if the specified file does not exist at startup time, the default value of `gate.xml` (`.gate.xml` on Unix platforms) in the user's home directory is used.

gate.user.session points to the file containing the user's saved session. If not specified, the default value of `gate.session` (`.gate.session` on Unix) in the user's home directory is used. When starting up GATE Developer, the session is reloaded from this file if it exists, and when exiting GATE Developer the session is saved to this file (unless the user has disabled 'save session on exit' in the configuration dialog). The session is not used when using GATE Embedded.

load.plugin.path is a path-like structure, i.e. a list of URLs separated by ';'. All directories listed here will be loaded as CREOLE plugins during initialisation. This has similar functionality with the `-d` command line option.

`gate.builtin.creole.dir` is a URL pointing to the location of GATE's built-in CREOLE directory. This is the location of the `creole.xml` file that defines the fundamental GATE resource types, such as documents, document format handlers, controllers and the basic visual resources that make up GATE. The default points to a location inside `gate.jar` and should not generally need to be overridden.

When using GATE Embedded, you can set the values for these properties before you call `Gate.init()`. Alternatively, you can set the values programmatically using the static methods `setGateHome()`, `setPluginsHome()`, `setSiteConfigFile()`, etc. before calling `Gate.init()`. See the Javadoc documentation for details. If you want to set these values from the command line you can use the following syntax for setting `gate.home` for example:

```
java -Dgate.home=/my/new/gate/home/directory -cp... gate.Main
```

When running GATE Developer, you can set the properties by creating a file `build.properties` in the top level GATE directory. In this file, any system properties which are prefixed with 'run.' will be passed to GATE. For example, to set an alternative user config file, put the following line in `build.properties`¹:

```
run.gate.user.config=${user.home}/alternative-gate.xml
```

This facility is not limited to the GATE-specific properties listed above, for example the following line changes the default temporary directory for GATE (note the use of forward slashes, even on Windows platforms):

```
run.java.io.tmpdir=d:/bigtmp
```

2.4 Configuring GATE

When GATE Developer is started, or when `Gate.init()` is called from GATE Embedded, GATE loads various sorts of configuration data stored as XML in files generally called something like `gate.xml` or `.gate.xml`. This data holds information such as:

- whether to save settings on exit;
- whether to save session on exit;
- what fonts GATE Developer should use;
- plugins to load at start;
- colours of the annotations;

¹In this specific case, the alternative config file must already exist when GATE starts up, so you should copy your standard `gate.xml` file to the new location.

- locations of files for the file chooser;
- and a lot of other GUI related options;

This type of data is stored at two levels (in order from general to specific):

- the site-wide level, which by default is located the `gate.xml` file in top level directory of the GATE installation (i.e. the `GATE home`. This location can be overridden by the Java system property `gate.site.config`;
- the user level, which lives in the user's HOME directory on UNIX or their profile directory on Windows (note that parts of this file are overwritten when saving user settings). The default location for this file can be overridden by the Java system property `gate.user.config`.

Where configuration data appears on several different levels, the more specific ones overwrite the more general. This means that you can set defaults for all GATE users on your system, for example, and allow individual users to override those defaults without interfering with others.

Configuration data can be set from the GATE Developer GUI via the 'Options' menu then 'Configuration'. The user can change the appearance of the GUI in the 'Appearance' tab, which includes the options of font and the 'look and feel'. The 'Advanced' tab enables the user to include annotation features when saving the document and preserving its format, to save the selected Options automatically on exit, and to save the session automatically on exit. The 'Input Methods' submenu from the 'Options' menu enables the user to change the default language for input. These options are all stored in the user's `.gate.xml` file.

When using GATE Embedded, you can also set the site config location using `Gate.setSiteConfigFile(File)` prior to calling `Gate.init()`.

2.5 Building GATE

Note that you don't need to build GATE unless you're doing development on the system itself.

Prerequisites:

- A conforming Java environment as above.

- A copy of the GATE sources and the build scripts – either the SRC distribution package from the nightly snapshots or a copy of the code obtained through Subversion (see Section 2.2.3).
- An appreciation of natural beauty.

GATE now includes a copy of the ANT build tool which can be accessed through the scripts included in the `bin` directory (use `ant.bat` for Windows 98 or ME, `ant.cmd` for Windows NT, 2000 or XP, and `ant.sh` for Unix platforms).

To build gate, cd to `gate` and:

1. Type:
`bin/ant`
2. [optional] To test the system:
`bin/ant test`
3. [optional] To make the Javadoc documentation:
`bin/ant doc`
4. You can also run GATE Developer using Ant, by typing:
`bin/ant run`
5. To see a full list of options type: `bin/ant help`

(The details of the build process are all specified by the `build.xml` file in the `gate` directory.)

You can also use a development environment like Borland JBuilder (click on the `gate.jpx` file), but note that it's still advisable to use ant to generate documentation, the jar file and so on. Also note that the run configurations have the location of a `gate.xml` site configuration file hard-coded into them, so you may need to change these for your site.

2.5.1 Using GATE with Maven or JPF

This section is based on contributions by Georg Öttl and William Oberman.

To use GATE with Maven you need a definition of the dependencies in POM format. There's an example POM here.

To use GATE with JPF (a Java plugin framework) you need a plugin definition like this one.

2.6 Uninstalling GATE

If you have used the installer, run:

```
java -jar uninstaller.jar
```

or just delete the whole of the installation directory (the one containing bin, lib, Uninstaller, etc.). The installer doesn't install anything outside this directory, but for completeness you might also want to delete the settings files GATE creates in your home directory (.gate.xml and .gate.session).

2.7 Troubleshooting

2.7.1 I don't see the Java console messages under Windows

Note that the `gate.bat` script uses `javaw.exe` to run GATE which means that you will see no console for the java process. If you have problems starting GATE and you would like to be able to see the console to check for messages then you should edit the `gate.bat` script and replace `javaw.exe` with `java.exe` in the definition of the `JAVA` environment variable.

2.7.2 When I execute GATE, nothing happens

You might get some clues if you start GATE from the command line, using:

```
bin/ant -Druntime.spawn=false run
```

which will allow you to see all error messages GATE generates.

2.7.3 On Ubuntu, GATE is very slow or doesn't start

GATE and many other Java applications are known not to work with GCJ, the open-source Java SDK or others non SUN Java SDK.

Make sure you have the official version of Java installed. Provided by Sun, the package is named 'sun-java6-jdk' in Synaptic. GATE also works with Java version 5 so 'sun-java5-jdk'.

To install it, run in a terminal:

```
sudo apt-get install sun-java6-jdk
```

Make sure that your default Java version is the one from SUN. You can do this by running:

```
sudo update-java-alternatives -l
```

This will list the installed Java VMs. You should see 'java-6-sun' as one of the options.

Then you should run :

```
sudo update-java-alternatives -s java-6-sun
```

to set the 'java-6-sun' as your default.

Finally, try GATE again.

2.7.4 How to use GATE on a 64 bit system?

32-bit vs. 64-bit is a matter of the JVM rather than the build of GATE -

For example, on Mac OS X, either use Applications/Utilities/Java Preferences and put one of the 64-bit options at the top of the list, or run GATE from the terminal using Java 1.6.0 (which is 64-bit only on Mac OS):

```
export JAVA_HOME=/System/Library/Frameworks/JavaVM.framework/Versions/1.6.0/Home  
bin/ant run
```

2.7.5 I got the error: Could not reserve enough space for object heap

GATE doesn't use the JAVA_OPTS variable. The default memory allocations are defined in the gate/build.xml file but you can override them by creating a file called build.properties in the same directory containing

```
runtime.start.memory=256m  
runtime.max.memory=1048m
```

If you don't use ant to start GATE but your own application directly with the 'java' executable then you must use something like:

```
java -Xmx512m -classpath <your classpath here> <yourClassName>
```

2.7.6 From Eclipse, I got the error: `java.lang.OutOfMemoryError: Java heap space`

Configuring `xms` and `xmx` parameters in `eclipse.ini` file just adds memory to your Eclipse process. If you start a Java application from within Eclipse, that will run in a different process.

To give more memory to your application, as opposed to just to Eclipse, you need to add those values in the 'VM Arguments' section of the run application dialog: lower pane, in the second tab of 'Run Configurations' dialog.

2.7.7 On MacOS, I got the error: `java.lang.OutOfMemoryError: Java heap space`

You can try to set the environment variable `ANT_OPTS` to allow for more memory as follows:

```
export ANT_OPTS=-Xmx1024m
```

Another cause can be when compiling with Java 6 on Mac. It builds OK using Java 5 with

```
export JAVA_HOME=/System/Library/Frameworks/JavaVM.framework/Versions/1.5.0/Home
```

and once built it runs fine with Java 6. Adding

```
fork="true"  
memoryMaximumSize="500M"
```

or similar to the `<javac>` in `build.xml` might fix it but if you're making changes that are to be committed to subversion you really ought to be building with Java 5 anyway :-)

2.7.8 I got the error: `log4j:WARN No appenders could be found for logger...`

You need to copy the 'gate/bin/log4j.properties' file to the directory from which you execute your project.

2.7.9 Text is incorrectly refreshed after scrolling and become unreadable

Change the look and feel used in GATE with menu 'Options' then 'Configuration'. Restart GATE and try again. We use mainly 'Metal' and 'Nimbus' without problem.

Change the video driver you use.

Update Java.

Chapter 3

Using GATE Developer

‘The law of evolution is that the strongest survives!’

‘Yes; and the strongest, in the existence of any social species, are those who are most social. In human terms, most ethical. . . . There is no strength to be gained from hurting one another. Only weakness.’

The Dispossessed [p.183], Ursula K. le Guin, 1974.

This chapter introduces GATE Developer, which is the GATE graphical user interface. It is analogous to systems like Mathematica for mathematicians, or Eclipse for Java programmers, providing a convenient graphical environment for research and development of language processing software. As well as being a powerful research tool in its own right, it is also very useful in conjunction with GATE Embedded (the GATE API by which GATE functionality can be included in your own applications); for example, GATE Developer can be used to create applications that can then be embedded via the API. This chapter describes how to complete common tasks using GATE Developer. It is intended to provide a good entry point to GATE functionality, and so explanations are given assuming only basic knowledge of GATE. However, probably the best way to learn how to use GATE Developer is to use this chapter in conjunction with the demonstrations and tutorials movies. There are specific links to them throughout the chapter. There is also a complete new set of video tutorials [here](#).

The basic business of GATE is annotating documents, and all the functionality we will introduce relates to that. Core concepts are;

- the **documents** to be annotated,
- **corpora** comprising sets of documents, grouping documents for the purpose of running uniform processes across them,
- **annotations** that are created on documents,

- **annotation types** such as ‘Name’ or ‘Date’,
- **annotation sets** comprising groups of annotations,
- **processing resources** that manipulate and create annotations on documents, and
- **applications**, comprising sequences of processing resources, that can be applied to a document or corpus.

What is considered to be the end result of the process varies depending on the task, but for the purposes of this chapter, output takes the form of the annotated document/corpus. Researchers might be more interested in figures demonstrating how successfully their application compares to a ‘gold standard’ annotation set; Chapter 10 in Part II will cover ways of comparing annotation sets to each other and obtaining measures such as F1. Implementers might be more interested in using the annotations programmatically; Chapter 7, also in Part II, talks about working with annotations from GATE Embedded. For the purposes of this chapter, however, we will focus only on creating the annotated documents themselves, and creating GATE applications for future use.

GATE includes a complete information extraction system that you are free to use, called ANNIE (a Nearly-New Information Extraction System). Many users find this is a good starting point for their own application, and so we will cover it in this chapter. Chapter 6 talks in a lot more detail about the inner workings of ANNIE, but we aim to get you started using ANNIE from inside of GATE Developer in this chapter.

We start the chapter with an exploration of the GATE Developer GUI, in Section 3.1. We describe how to create documents (Section 3.2) and corpora (Section 3.3). We talk about viewing and manually creating annotations (Section 3.4).

We then talk about loading the plugins that contain the processing resources you will use to construct your application, in Section 3.5. We then talk about instantiating processing resources (Section 3.6). Section 3.7 covers applications, including using ANNIE (Section 3.7.3). Saving applications and language resources (documents and corpora) is covered in Section 3.8. We conclude with a few assorted topics that might be useful to the GATE Developer user, in Section 3.10.

3.1 The GATE Developer Main Window

Figure 3.1 shows the main window of GATE Developer, as you will see it when you first run it. There are five main areas:

1. at the top, the *menus bar* and *tools bar* with menus ‘File’, ‘Options’, ‘Tools’, ‘Help’ and icons for the most frequently used actions;




Figure 3.1: Main Window of GATE Developer

2. on the left side, a tree starting from ‘GATE’ and containing ‘Applications’, ‘Language Resources’ etc. – this is the *resources tree*;
3. in the bottom left corner, a rectangle, which is the *small resource viewer*;
4. in the center, containing tabs with ‘Messages’ or the name of a resource from the resources tree, the *main resource viewer*;
5. at the bottom, the *messages bar*.

The menu and the messages bar do the usual things. Longer messages are displayed in the messages tab in the main resource viewer area.

The resource tree and resource viewer areas work together to allow the system to display diverse resources in various ways. The many resources integrated with GATE can have either a small view, a large view, or both.

At any time, the main viewer can also be used to display other information, such as messages, by clicking on the appropriate tab at the top of the main window. If an error occurs in processing, the messages tab will flash red, and an additional popup error message may also occur.

In the options dialogue from the Options menu you can choose if you want to link the selection in the resources tree and the selected main view.

3.2 Loading and Viewing Documents

Name	Type	Required	Value
collectRepositioningInfo	Boolean	<input checked="" type="checkbox"/>	false
encoding	String	<input type="checkbox"/>	
markupAware	Boolean	<input checked="" type="checkbox"/>	true
mimeType	String	<input type="checkbox"/>	
preserveOriginalContent	Boolean	<input checked="" type="checkbox"/>	false
sourceUrl	URL	<input checked="" type="checkbox"/>	<input type="text"/>
sourceUrlEndOffset	Long	<input type="checkbox"/>	
sourceUrlStartOffset	Long	<input type="checkbox"/>	

OK Cancel Help

Figure 3.2: Making a New Document

If you right-click on ‘Language Resources’ in the resources pane, select ‘New’ then ‘GATE Document’, the window ‘Parameters for the new GATE Document’ will appear as shown in figure 3.2. Here, you can specify the GATE document to be created. Required parameters are indicated with a tick. The name of the document will be created for you if you do not specify it. Enter the URL of your document or use the file browser to indicate the file you wish to use for your document source. For example, you might use ‘http://www.gate.ac.uk’, or browse to a text or XML file you have on disk. Click on ‘OK’ and a GATE document will be created from the source you specified.

See also the movie for creating documents.

The document editor is contained in the central tabbed pane in GATE Developer. Double-click on your document in the resources pane to view the document editor. The document editor consists of a top panel with buttons and icons that control the display of different views and the search box. Initially, you will see just the text of your document, as shown in figure 3.3. Click on ‘Annotation Sets’ and ‘Annotations List’ to view the annotation sets to the right and the annotations list at the bottom. You will see a view similar to figure 3.4. In place of the annotations list, you can also choose to see the annotations stack. In place of the annotation sets, you can also choose to view the co-reference editor. More information about this functionality is given in Section 3.4.

Text in a loaded document can be edited in the document viewer. The usual platform specific cut, copy and paste keyboard shortcuts should also work, depending on your operating

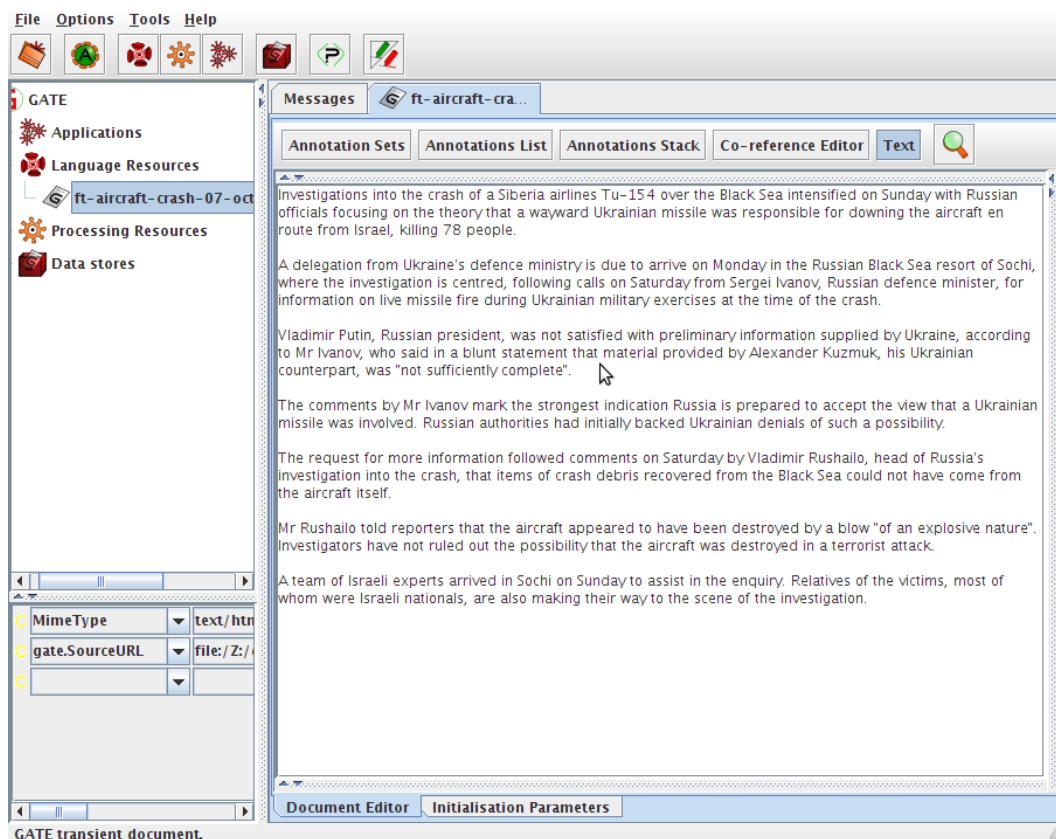


Figure 3.3: The Document Editor

system (e.g. CTRL-C, CTRL-V for Windows). The last icon, a magnifying glass, at the top of the document editor is for searching in the document. To prevent the new annotation windows popping up when a piece of text is selected, hold down the CTRL key. Alternatively, you can hide the annotation sets view by clicking on its button at the top of the document view; this will also cause the highlighted portions of the text to become un-highlighted.

You can set the document editor to be read-only in the options dialogue from the 'Options' menu. If enabled, you won't be able to edit the text but you will still be able to edit annotations.

Another options is to choose if the insertion when editing text should be before or after the caret.

See also Section 16.2.2 for the compound document editor.

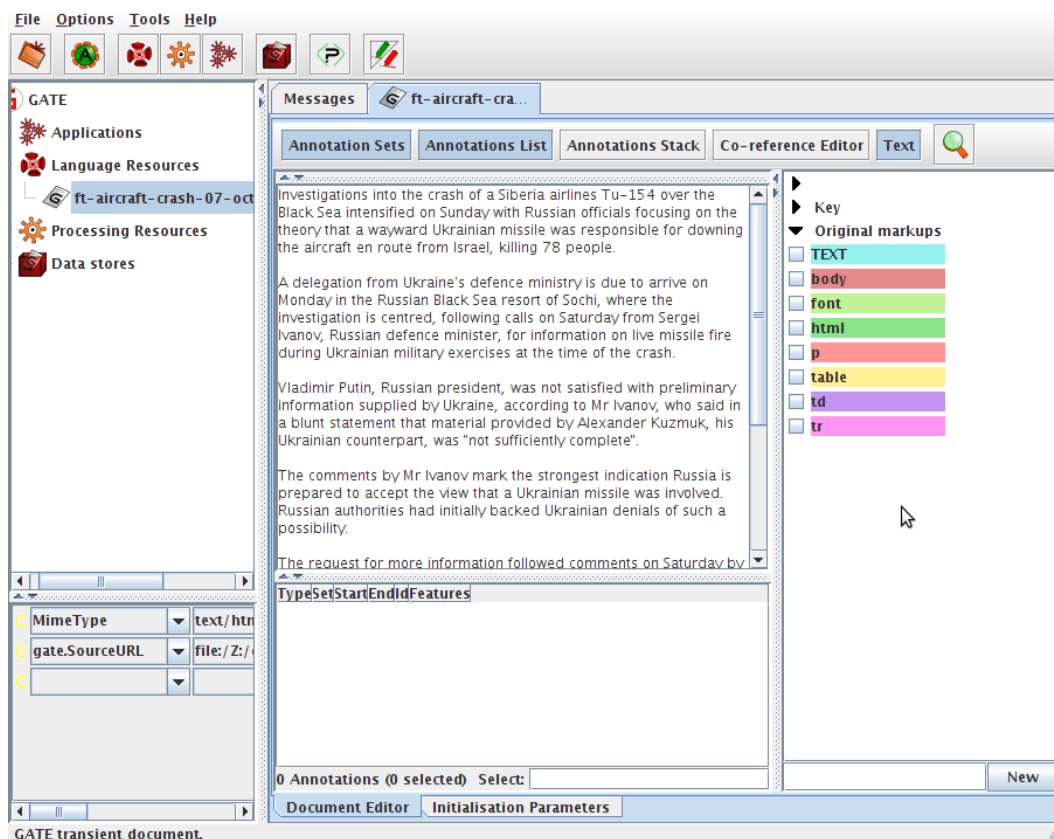


Figure 3.4: The Document Editor with Annotation Sets and Annotations List

3.3 Creating and Viewing Corpora

You can create a new corpus in a similar manner to creating a new document; simply right-click on ‘Language Resources’ in the resources pane, select ‘New’ then ‘GATE corpus’. A brief dialogue box will appear in which you can optionally give a name for your corpus (if you leave this blank, a corpus name will be created for you) and optionally add documents to the corpus from those already loaded into GATE.

There are three ways of adding documents to a corpus:

1. When creating the corpus, clicking on the icon next to the “documentsList” input field brings up a popup window with a list of the documents already loaded into GATE Developer. This enables the user to add any documents to the corpus.
2. Alternatively, the corpus can be loaded first, and documents added later by double clicking on the corpus and using the + and - icons to add or remove documents to the corpus. Note that the documents must have been loaded into GATE Developer before they can be added to the corpus.

3. Once loaded, the corpus can be populated by right clicking on the corpus and selecting ‘Populate’. With this method, documents do not have to have been previously loaded into GATE Developer, as they will be loaded during the population process. If you right-click on your corpus in the resources pane, you will see that you have the option to ‘Populate’ the corpus. If you select this option, you will see a dialogue box in which you can specify a directory in which GATE will search for documents. You can specify the extensions allowable; for example, XML or TXT. This will restrict the corpus population to only those documents with the extensions you wish to load. You can choose whether to recurse through the directories contained within the target directory or restrict the population to those documents contained in the top level directory. Click on ‘OK’ to populate your corpus. This option provides a quick way to create a GATE Corpus from a directory of documents.

Additionally, right-clicking on a loaded document in the tree and selecting the ‘New corpus with this document’ option creates a new transient corpus named `Corpus for document name` containing just this document.

See also the movie for creating and populating corpora.

Double click on your corpus in the resources pane to see the corpus editor, shown in figure 3.5. You will see a list of the documents contained within the corpus.

In the top left of the corpus editor, plus and minus buttons allow you to add documents to the corpus from those already loaded into GATE and remove documents from the corpus (note that removing a document from a corpus does not remove it from GATE).

Up and down arrows at the top of the view allow you to reorder the documents in the corpus. The rightmost button in the view opens the currently selected document in a document editor.

At the bottom, you will see that tabs entitled ‘Initialisation Parameters’ and ‘Corpus Quality Assurance’ are also available in addition to the corpus editor tab you are currently looking at. Clicking on the ‘Initialisation Parameters’ tab allows you to view the initialisation parameters for the corpus. The ‘Corpus Quality Assurance’ tab allows you to calculate agreement measures between the annotations in your corpus. Agreement measures are discussed in depth in Chapter 10. The use of corpus quality assurance is discussed in Section 10.3.

3.4 Working with Annotations

In this section, we will talk in more detail about viewing annotations, as well as creating and editing them manually. As discussed in at the start of the chapter, the main purpose of GATE is annotating documents. Whilst applications can be used to annotate the documents entirely automatically, annotation can also be done manually, e.g. by the user, or semi-automatically, by running an application over the corpus and then correcting/adding new annotations




Figure 3.5: Corpus Editor

manually. Section 3.4.5 focuses on manual annotation. In Section 3.6 we talk about running processing resources on our documents. We begin by outlining the functionality around viewing annotations, organised by the GUI area to which the functionality pertains.

3.4.1 The Annotation Sets View

To view the annotation sets, click on the ‘Annotation Sets’ button at the top of the document editor, or use the F3 key (see Section 3.9 for more keyboard shortcuts). This will bring up the annotation sets viewer, which displays the annotation sets available and their corresponding annotation types.

The annotation sets view is displayed on the left part of the document editor. It’s a tree-like view with a root for each annotation set. The first annotation set in the list is always a nameless set. This is the default annotation set. You can see in figure 3.4 that there is a

drop-down arrow with no name beside it. Other annotation sets on the document shown in figure 3.4 are ‘Key’ and ‘Original markups’. Because the document is an XML document, the original XML markup is retained in the form of an annotation set. This annotation set is expanded, and you can see that there are annotations for ‘TEXT’, ‘body’, ‘font’, ‘html’, ‘p’, ‘table’, ‘td’ and ‘tr’.

To display all the annotations of one type, tick its checkbox or use the space key. The text segments corresponding to these annotations will be highlighted in the main text window. To delete an annotation type, use the delete key. To change the color, use the enter key. There is a context menu for all these actions that you can display by right-clicking on one annotation type, a selection or an annotation set.

If you keep shift key pressed when you open the annotation sets view, GATE Developer will try to select any annotations that were selected in the previous document viewed (if any); otherwise no annotation will be selected.

Having selected an annotation type in the annotation sets view, hovering over an annotation in the main resource viewer or right-clicking on it will bring up a popup box containing a list of the annotations associated with it, from which one can select an annotation to view in the annotation editor, or if there is only one, the annotation editor for that annotation. Figure 3.6 shows the annotation editor.

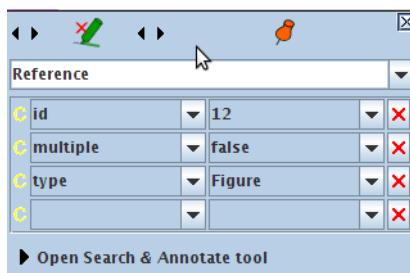


Figure 3.6: The Annotation Editor

3.4.2 The Annotations List View

To view the list of annotations and their features, click on the ‘Annotations list’ button at the top of the main window or use F4 key. The annotation list view will appear below the main text. It will only contain the annotations selected from the annotation sets view. These lists can be sorted in ascending and descending order for any column, by clicking on the corresponding column heading. Moreover you can hide a column by using the context menu by right-clicking on the column headings. Selecting rows in the table will blink the respective annotations in the document. Right-click on a row or selection in this view to delete or edit an annotation. Delete key is a shortcut to delete selected annotations.

3.4.3 The Annotations Stack View

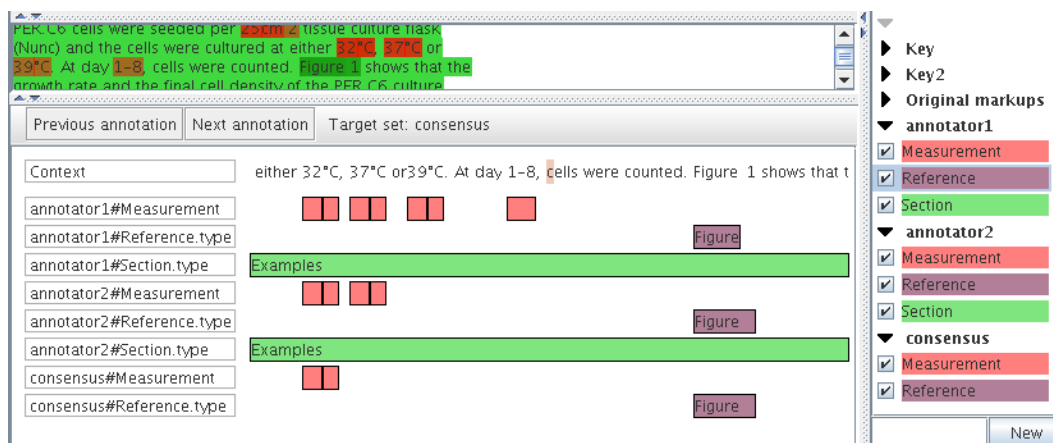


Figure 3.7: Annotations stack view centred on the document caret.

This view is similar to the ANNIC view described in section 9.2. It displays annotations at the document caret position with some context before and after. The annotations are stacked from top to bottom, which gives a clear view when they are overlapping.

As the view is centred on the document caret, you can use the conventional keypresses to move it and update the view: notably the keys left and right to skip one letter; control + left/right to skip one word; up and down to go one line up or down; and use the document scrollbar then click in the document to move further. There are also two buttons at the top of the view that centre the view on the closest previous/next annotation boundary among all displayed. This is useful when you want to skip a region without annotation or when you want to reach the beginning or end of a very long annotation.

The annotation types displayed correspond to those selected in the annotation sets view. You can display feature values for an annotation rectangle by hovering the mouse on it or select only one feature to display by double-clicking on the annotation type in the first column.

Right-clicking on an annotation in the annotations stack view gives the option to edit that annotation.

3.4.4 The Co-reference Editor

The co-reference editor allows co-reference chains (see Section 6.9) to be displayed and edited in GATE Developer. To display the co-reference editor, first open a document in GATE Developer, and then click on the **Co-reference Editor** button in the document viewer.

The combo box at the top of the co-reference editor allows you to choose which annotation set to display co-references for. If an annotation set contains no co-reference data, then the

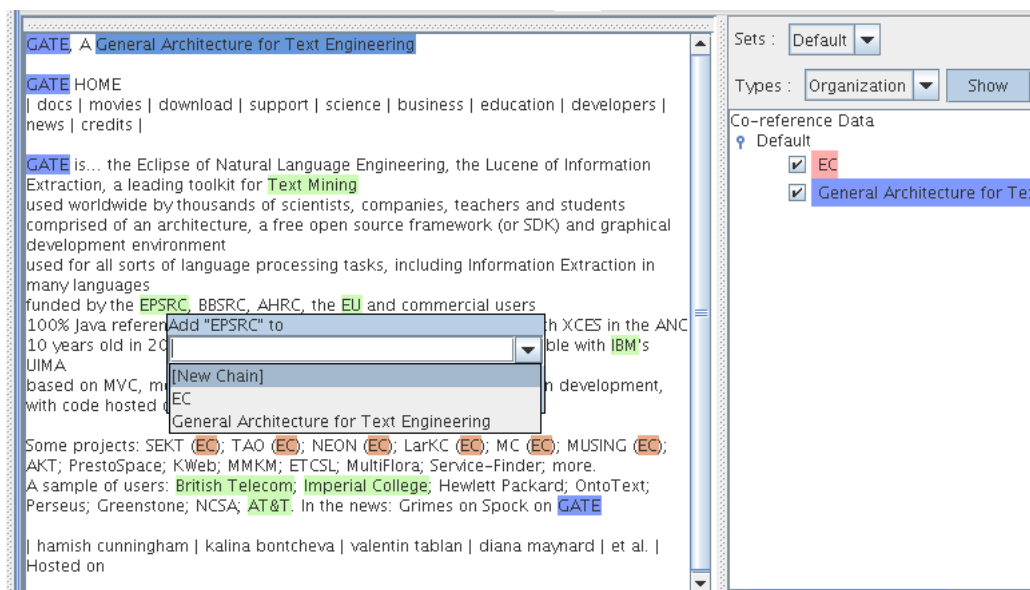


Figure 3.8: Co-reference editor inside a document editor. The popup window in the document under the word ‘EPSRC’ is used to add highlighted annotations to a co-reference chain. Here the annotation type ‘Organization’ of the annotation set ‘Default’ is highlighted and also the co-references ‘EC’ and ‘GATE’.

tree below the combo box will just show ‘Coreference Data’ and the name of the annotation set. However, when co-reference data does exist, a list of all the co-reference chains that are based on annotations in the currently selected set is displayed. The name of each co-reference chain in this list is the same as the text of whichever element in the chain is the longest. It is possible to highlight all the member annotations of any chain by selecting it in the list.

When a co-reference chain is selected, if the mouse is placed over one of its member annotations, then a pop-up box appears, giving the user the option of deleting the item from the chain. If the only item in a chain is deleted, then the chain itself will cease to exist, and it will be removed from the list of chains. If the name of the chain was derived from the item that was deleted, then the chain will be given a new name based on the next longest item in the chain.

A combo box near the top of the co-reference editor allows the user to select an annotation type from the current set. When the **Show** button is selected all the annotations of the selected type will be highlighted. Now when the mouse pointer is placed over one of those annotations, a pop-up box will appear giving the user the option of adding the annotation to a co-reference chain. The annotation can be added to an existing chain by typing the name of the chain (as shown in the list on the right) in the pop-up box. Alternatively, if the user presses the down cursor key, a list of all the existing annotations appears, together with the option **[New Chain]**. Selecting the **[New Chain]** option will cause a new chain to be created containing the selected annotation as its only element.

Each annotation can only be added to a single chain, but annotations of different types can be added to the same chain, and the same text can appear in more than one chain if it is referenced by two or more annotations.

The movie for inspecting results is also useful for learning about viewing annotations.

3.4.5 Creating and Editing Annotations

To create annotations manually, select the text you want to annotate and hover the mouse on the selection or use control+E keys. A popup will appear, allowing you to create an annotation, as shown in figure 3.9

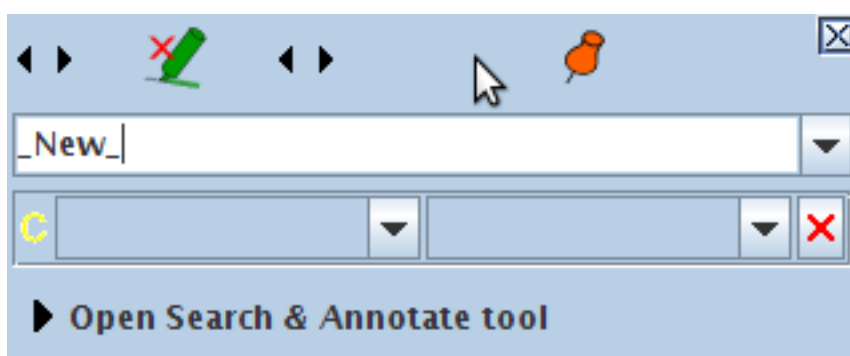


Figure 3.9: Creating a New Annotation

The type of the annotation, by default, will be the same as the last annotation you created, unless there is none, in which case it will be ‘_New_’. You can enter any annotation type name you wish in the text box, unless you are using schema-driven annotation (see Section 3.4.6). You can add or change features and their values in the table below.

To delete an annotation, click on the red X icon at the top of the popup window. To grow/shrink the span of the annotation at its start use the two arrow icons on the left or right and left keys. Use the two arrow icons next on the right to change the annotation end or alt+right and alt+left keys. Add shift and control+shift keys to make the span increment bigger. The red X icon is for removing the annotation.

The pin icon is to pin the window so that it remains where it is. If you drag and drop the window, this automatically pins it too. Pinning it means that even if you select another annotation (by hovering over it in the main resource viewer) it will still stay in the same position.

The popup menu only contains annotation types present in the Annotation Schema and those already listed in the relevant Annotation Set. To create a new Annotation Schema, see Section 3.4.6. The popup menu can be edited to add a new annotation type, however.

The new annotation created will automatically be placed in the annotation set that has been selected (highlighted) by the user. To create a new annotation set, type the name of the new set to be created in the box below the list of annotation sets, and click on ‘New’.

Figure 3.10 demonstrates adding a ‘Organization’ annotation for the string ‘EPSRC’ (highlighted in green) to the default annotation set (blank name in the annotation set view on the right) and a feature name ‘type’ with a value about to be added.

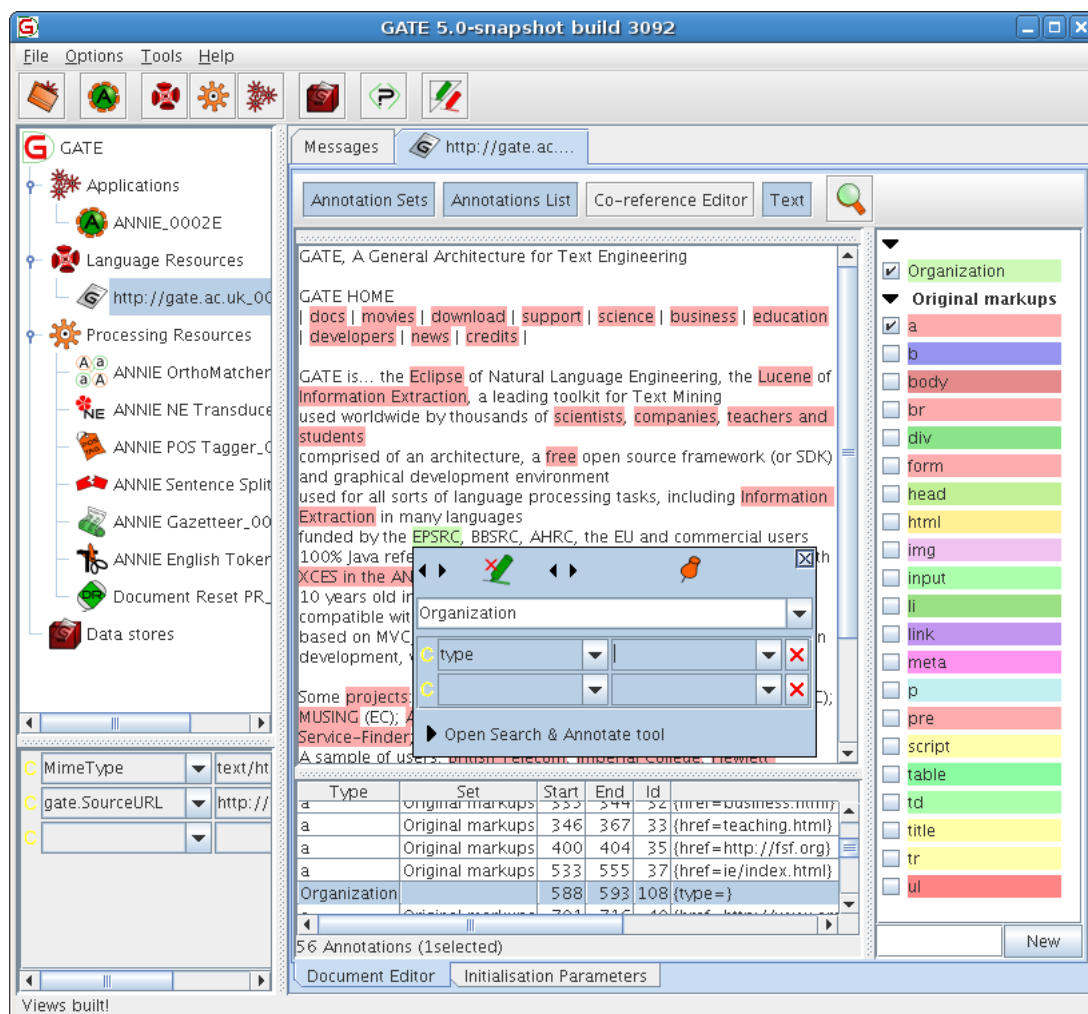


Figure 3.10: Adding an Organization annotation to the Default Annotation Set

To add a second annotation to a selected piece of text, or to add an overlapping annotation to an existing one, press the CTRL key to avoid the existing annotation popup appearing, and then select the text and create the new annotation. Again by default the last annotation type to have been used will be displayed; change this to the new annotation type. When a piece of text has more than one annotation associated with it, on mouseover all the annotations will be displayed. Selecting one of them will bring up the relevant annotation popup.

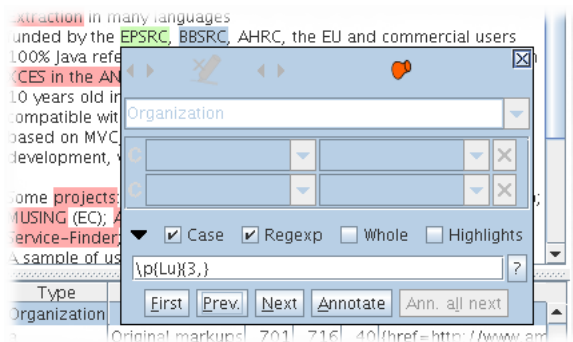


Figure 3.11: Search and Annotate Function of the Annotation Editor.

To search and annotate the document automatically, use the search and annotate function as shown in figure 3.11:

- Create and/or select an annotation to be used as a model to annotate.
- Open the panel at the bottom of the annotation editor window.
- Change the expression to search if necessary.
- Use the [First] button or Enter key to select the first expression to annotate.
- Use the [Annotate] button if the selection is correct otherwise the [Next] button. After a few cycles of [Annotate] and [Next], Use the [Ann. all next] button.

Note that after using the [First] button you can move the caret in the document and use the [Next] button to avoid continuing the search from the beginning of the document. The [?] button at the end of the search text field will help you to build powerful regular expressions to search.

3.4.6 Schema-Driven Editing

Annotation schemas allow annotation types and features to be pre-specified, so that during manual annotation, the relevant options appear on the drop-down lists in the annotation editor. You can see some example annotation schemas in Section 5.4.1. Annotation schemas provide a means to define types of annotations in GATE Developer. Basically this means that GATE Developer ‘knows about’ annotations defined in a schema.

Annotation schemas are supported by the ‘Annotation schema’ language resource in ANNIE, so to use them you must first ensure that the ‘ANNIE’ plugin is loaded (see Section 3.5). This will load a set of default schemas, as well as allowing you to load schemas of your own.

The default annotation schemas contain common named entities such as Person, Organisation, Location, etc. You can modify the existing schema or create a new one, in order to tell GATE Developer about other kinds of annotations you frequently use. You can still create annotations in GATE Developer without having specified them in an annotation schema, but you may then need to tell GATE Developer about the properties of that annotation type each time you create an annotation for it.

To load a schema of your own, right-click on ‘Language Resources’ in the resources pane. Select ‘New’ then ‘Annotation schema’. A popup box will appear in which you can browse to your annotation schema XML file.

An alternative annotation editor component is available which constrains the available annotation types and features much more tightly, based on the annotation schemas that are currently loaded. This is particularly useful when annotating large quantities of data or for use by less skilled users.

To use this, you must load the `Schema_Annotation_Editor` plugin. With this plugin loaded, the annotation editor will *only* offer the annotation types permitted by the currently loaded set of schemas, and when you select an annotation type only the features permitted by the schema are available to edit¹. Where a feature is declared as having an enumerated type the available enumeration values are presented as an array of buttons, making it easy to select the required value quickly.

3.4.7 Printing Text with Annotations

We suggest you to use your browser to print a document as GATE don’t propose a printing facility for the moment.

First save your document by right clicking on the document in the left resources tree then choose ‘Save Preserving Format’. You will get an XML file with all the annotations highlighted as XML tags plus the ‘Original markups’ annotations set.

Then add a stylesheet processing instruction at the beginning of the XML file:

```
<?xml version="1.0" encoding="UTF-8" ?>
<?xml-stylesheet type="text/css" href="gate.css"?>
```

And create a file ‘gate.css’ in the same directory:

```
BODY, body { margin: 2em } /* or any other first level tag */
P, p { display: block } /* or any other paragraph tag */
```

¹existing features outwith the schema, e.g. those created by previously-run processing resources, are not editable but not modified or removed by the editor.

```

/* ANNIE tags but you can use whatever tags you want */
/* be careful that XML tags are case sensitive */
Date          { background-color: rgb(230, 150, 150) }
FirstPerson   { background-color: rgb(150, 230, 150) }
Identifier     { background-color: rgb(150, 150, 230) }
JobTitle      { background-color: rgb(150, 230, 230) }
Location      { background-color: rgb(230, 150, 230) }
Money         { background-color: rgb(230, 230, 150) }
Organization  { background-color: rgb(230, 200, 200) }
Percent       { background-color: rgb(200, 230, 200) }
Person        { background-color: rgb(200, 200, 230) }
Title         { background-color: rgb(200, 230, 230) }
Unknown       { background-color: rgb(230, 200, 230) }
Etc           { background-color: rgb(230, 230, 200) }

```

Finally open the XML file in your browser and print it.

Note that overlapping annotations, cannot be expressed correctly with inline XML tags and thus won't be displayed correctly.

3.5 Using CREOLE Plugins

In GATE, processing resources are used to automatically create and manipulate annotations on documents. We will talk about processing resources in the next section. However, we must first introduce CREOLE plugins. In most cases, in order to use a particular processing resource (and certain language resources) you must first load the CREOLE plugin that contains it. This section talks about using CREOLE plugins. Then, in Section 3.6, we will talk about creating and using processing resources.

The definitions of CREOLE resources (e.g. processing resources such as taggers and parsers, see Chapter 4) are stored in CREOLE directories (directories containing an XML file describing the resources, the Java archive with the compiled executable code and whatever libraries are required by the resources).

Starting with version 3, CREOLE directories are called 'CREOLE plugins' or simply 'plugins'. In previous versions, the CREOLE resources distributed with GATE used to be included in the monolithic `gate.jar` archive. Version 3 includes them as separate directories under the `plugins` directory of the distribution. This allows easy access to the linguistic resources used without the requirement to unpack the `gate.jar` file.

Plugins can have one or more of the following states in relation with GATE:

known plugins are those plugins that the system knows about. These include all the plugins

in the `plugins` directory of the GATE installation (the so-called *installed* plugins) as well all the plugins that were manually loaded from the user interface.

loaded plugins are the plugins currently loaded in the system. All CREOLE resource types from the loaded plugins are available for use. All known plugins can easily be loaded and unloaded using the user interface.

auto-loadable plugins are the list of plugins that the system loads automatically during initialisation.

The default location for installed plugins can be modified using the `gate.plugins.home` system property while the list of auto-loadable plugins can be set using the `load.plugin.path` property, see Section 2.3 above.

The CREOLE plugins can be managed through the graphical user interface which can be activated by selecting ‘Manage CREOLE Plugins’ from the ‘File’ menu. This will bring up a window listing all the known plugins. For each plugin there are two check-boxes – one labelled ‘Load now’, which will load the plugin, and the other labelled ‘Load always’ which will add the plugin to the list of auto-loadable plugins. A ‘Delete’ button is also provided – which will remove the plugin from the list of known plugins. This operation does not delete the actual plugin directory. Installed plugins are found automatically when GATE is started; if an installed plugin is deleted from the list, it will re-appear next time GATE is launched.

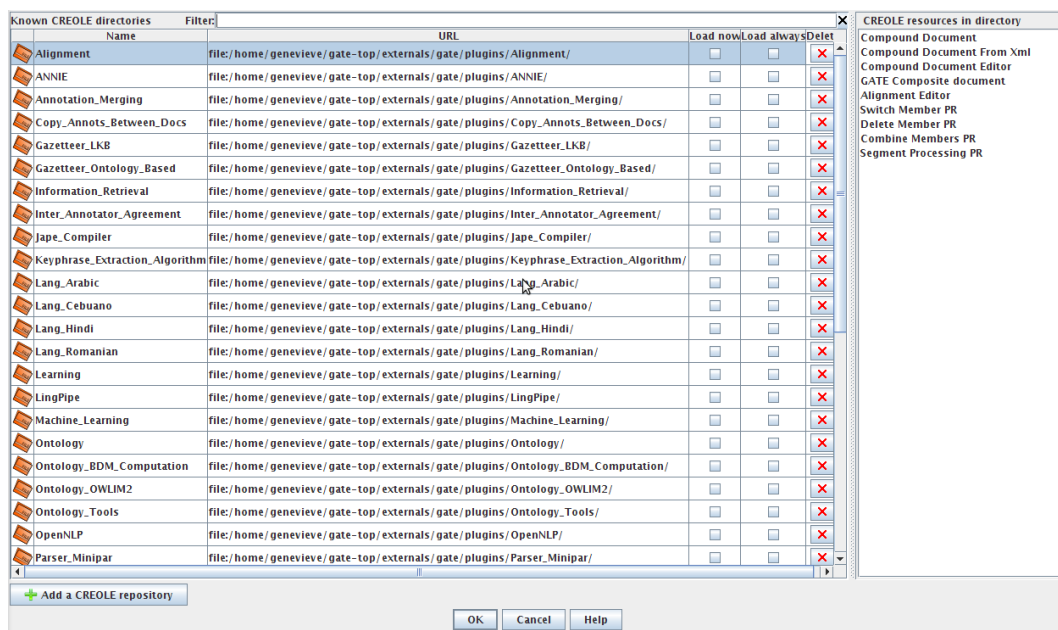


Figure 3.12: Plugin Management Console

If you select a plugin, you will see in the pane on the right the list of resources that plugin contains. For example, in figure 3.12, the ‘Alignment’ plugin is selected, and you can see that

it contains ten processing resources; ‘Compound Document’, ‘Compound Document From Xml’, ‘Compound Document Editor’, ‘GATE Composite document’ etc. If you wish to use a particular resource you will have to ascertain which plugin contains it. This list can be useful for that. Alternatively, the GATE website provides a directory of plugins and their processing resources.

Having loaded the plugins you need, the resources they define will be available for use. Typically, to the GATE Developer user, this means that they will appear on the ‘New’ menu when you right-click on ‘Processing Resources’ in the resources pane, although some special plugins have different effects; for example, the `Schema_Annotation_Editor` (see Section 3.4.6).

3.6 Loading and Using Processing Resources

This section describes how to load and run CREOLE resources not present in ANNIE. To load ANNIE, see Section 3.7.3. For technical descriptions of these resources, see the appropriate chapter in Part III (e.g. Chapter 19). First ensure that the necessary plugins have been loaded (see Section 3.5). If the resource you require does not appear in the list of Processing Resources, then you probably do not have the necessary plugin loaded. Processing resources are loaded by selecting them from the set of Processing Resources: right click on Processing Resources or select ‘New Processing Resource’ from the File menu.

For example, use the Plugin Console Manager to load the ‘Tools’ plugin. When you right click on ‘Processing Resources’ in the resources pane and select ‘New’ you have the option to create any of the processing resources that plugin provides. You may choose to create a ‘GATE Morphological Analyser’, with the default parameters. Having done this, an instance of the GATE Morphological Analyser appears under ‘Processing Resources’. This processing resource, or PR, is now available to use. Double-clicking on it in the resources pane reveals its initialisation parameters, see figure 3.13.

This processing resource is now available to be added to applications. It must be added to an application before it can be applied to documents. You may create as many of a particular processing resource as you wish, for example with different initialisation parameters. Section 3.7 talks about creating and running applications.

See also the movie for loading processing resources.

3.7 Creating and Running an Application

Once all the resources you need have been loaded, an application can be created from them, and run on your corpus. Right click on ‘Applications’ and select ‘New’ and then either ‘Corpus Pipeline’ or ‘Pipeline’. A pipeline application can only be run over a single document,

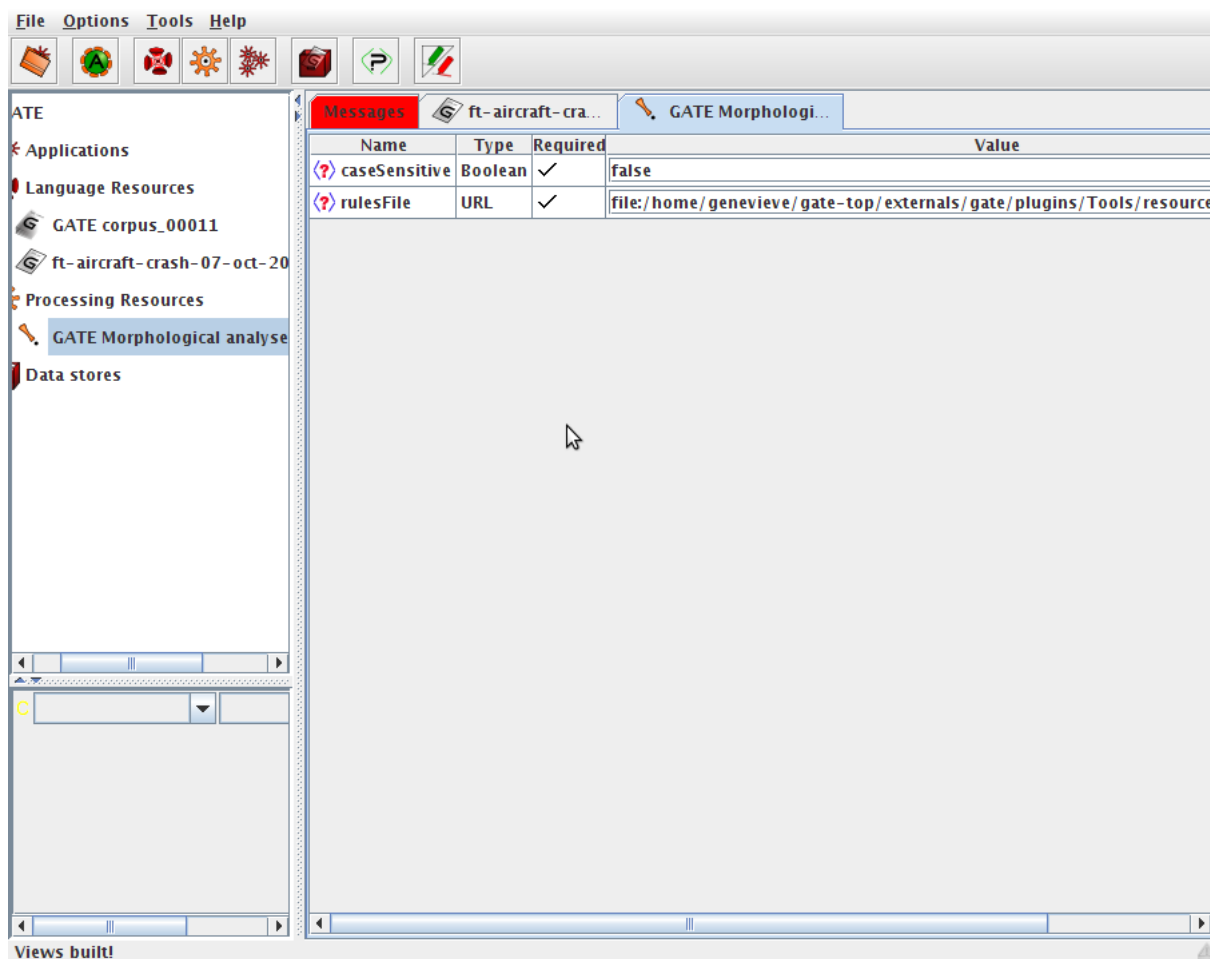


Figure 3.13: GATE Morphological Analyser Initialisation Parameters

while a corpus pipeline can be run over a whole corpus.

To build the pipeline, double click on it, and select the resources needed to run the application (you may not necessarily wish to use all those which have been loaded).

Transfer the necessary components from the set of ‘loaded components’ displayed on the left hand side of the main window to the set of ‘selected components’ on the right, by selecting each component and clicking on the left and right arrows, or by double-clicking on each component.

Ensure that the components selected are listed in the correct order for processing (starting from the top). If not, select a component and move it up or down the list using the up/down arrows at the left side of the pane.

Ensure that any parameters necessary are set for each processing resource (by clicking on the resource from the list of selected resources and checking the relevant parameters from

the pane below). For example, if you wish to use annotation sets other than the Default one, these must be defined for each processing resource.

Note that if a corpus pipeline is used, the corpus needs only to be set once, using the drop-down menu beside the ‘corpus’ box. If a pipeline is used, the document must be selected for each processing resource used.

Finally, right-click on ‘Run’ to run the application on the document or corpus.

See also the movie for loading and running processing resources.

For how to use the *conditional* versions of the pipelines see Section 3.7.2 and for saving/restoring the configuration of an application see Section 3.8.3.

3.7.1 Running an Application on a Datastore

To avoid loading all your documents at the same time you can run an application on a datastore corpus.

To do this you need to load your datastore, see section 3.8.2, and to load the corpus from the datastore by double clicking on it in the datastore viewer.

Then, in the application viewer, you need to select this corpus in the drop down list of corpora.

When you run the application on the corpus datastore, each document will be loaded, processed, saved then unloaded. So at any time there will be only one document from the datastore corpus loaded. This prevent memory shortage but is also a little bit slower than if all your documents were already loaded.

The processed documents are automatically saved back to the datastore so you may want to use a copy of the datastore to experiment.

Be very careful that if you have some documents from the datastore corpus already loaded before running the application then they will not be unloaded nor saved. To save such document you have to right click on it in the resources tree view and save it to the datastore.

3.7.2 Running PRs Conditionally on Document Features

The ‘Conditional Pipeline’ and ‘Conditional Corpus Pipeline’ application types are conditional versions of the pipelines mentioned in Section 3.7 and allow processing resources to be run or not according to the value of a feature on the document. In terms of graphical interface, the only addition brought by the conditional versions of the applications is a box situated underneath the lists of available and selected resources which allows the user to

choose whether the currently selected processing resource will run always, never or only on the documents that have a particular value for a named feature.

If the *Yes* option is selected then the corresponding resource will be run on all the documents processed by the application as in the case of non-conditional applications. If the *No* option is selected then the corresponding resource will never be run; the application will simply ignore its presence. This option can be used to temporarily and quickly disable an application component, for debugging purposes for example.

The *If value of feature* option permits running specific application components conditionally on document features. When selected, this option enables two text input fields that are used to enter the name of a feature and the value of that feature for which the corresponding processing resource will be run. When a conditional application is run over a document, for each component that has an associated condition, the value of the named feature is checked on the document and the component will only be used if the value entered by the user matches the one contained in the document features.

3.7.3 Doing Information Extraction with ANNIE

This section describes how to load and run ANNIE (see Chapter 6) from GATE Developer. ANNIE is a good place to start because it provides a complete information extraction application, that you can run on any corpus. You can then view the effects.

From the File menu, select ‘Load ANNIE System’. To run it in its default state, choose ‘with Defaults’. This will automatically load all the ANNIE resources, and create a corpus pipeline called ANNIE with the correct resources selected in the right order, and the default input and output annotation sets.

If ‘without Defaults’ is selected, the same processing resources will be loaded, but a popup window will appear for each resource, which enables the user to specify a name, location and other parameters for the resource. This is exactly the same procedure as for loading a processing resource individually, the difference being that the system automatically selects those resources contained within ANNIE. When the resources have been loaded, a corpus pipeline called ANNIE will be created as before.

The next step is to add a corpus (see Section 3.3), and select this corpus from the drop-down corpus menu in the Serial Application editor. Finally click on ‘Run’ from the Serial Application editor, or by right clicking on the application name in the resources pane and selecting ‘Run’. (Many people prefer to switch to the messages tab, then run their application by right-clicking on it in the resources pane, because then it is possible to monitor any messages that appear whilst the application is running.)

To view the results, double click on one of the document contained in the corpus processed in the left hand tree view. No annotation sets nor annotations will be shown until annotations are selected in the annotation sets; the ‘Default’ set is indicated only with an unlabelled

right-arrowhead which must be selected in order to make visible the available annotations. Open the default annotation set and select some of the annotations to see what the ANNIE application has done.

See also the movie for loading and running ANNIE.

3.7.4 Modifying ANNIE

You will find the ANNIE resources in `gate/plugins/ANNIE/resources`. Simply locate the existing resources you want to modify, make a copy with a new name, edit them, and load the new resources into GATE as new Processing Resources (see Section 3.6).

3.8 Saving Applications and Language Resources

In this section, we will describe how applications and language resources can be saved for use outside of GATE and for use with GATE at a later time. Section 3.8.1 talks about saving documents to file. Section 3.8.2 outlines how to use datastores. Section 3.8.3 talks about saving application states (resource parameter states), and Section 3.8.4 talks about exporting applications together with referenced files and resources to a ZIP file.

3.8.1 Saving Documents to File

There are three main ways to save annotated documents:

1. preserving the original markup, with optional added annotations;
2. in GATE's own XML serialisation format (including all the annotations on the document);
3. by writing your own dump algorithm as a processing resource.

This section describes how to use the first two options.

Both types of data export are available in the popup menu triggered by right-clicking on a document in the resources tree (see Section 3.1): type 1 is called 'Save Preserving Format' and type 2 is called 'Save as XML'. In addition, all documents in a corpus can be saved as individual XML files into a directory by right-clicking on the corpus in the resources tree and choosing the option 'Save as XML'.

Selecting the save as XML option leads to a file open dialogue; give the name of the file you want to create, and the whole document and all its data will be exported to that file. If you

later create a document from that file, the state will be restored. (**Note:** because GATE's annotation model is richer than that of XML, and because our XML dump implementation sometimes cuts corners², the state may not be identical after restoration. If your intention is to store the state for later use, use a DataStore instead.)

The 'Save Preserving Format' option also leads to a file dialogue; give a name and the data you require will be dumped into the file. The action can be used for documents that were created from files using the XML or HTML format. It will save all the original tags as well as the document annotations that are currently displayed in the 'Annotations List' view. This option is useful for selectively saving only some annotation types.

The annotations are saved as normal document tags, using the annotation type as the tag name. If the advanced option 'Include annotation features for "Save Preserving Format"' (see Section 2.4) is set to true, then the annotation features will also be saved as tag attributes.

Using this operation for GATE documents that were not created from an HTML or XML file results in a plain text file, with in-line tags for the saved annotations.

Note that GATE's model of annotation allows graph structures, which are difficult to represent in XML (XML is a tree-structured representation format). During the dump process, annotations that cross each other in ways that cannot be represented in legal XML will be discarded, and a warning message printed.

3.8.2 Saving and Restoring LRs in Datastores

Where corpora are large, the memory available may not be sufficient to have all documents open simultaneously. The datastore functionality provides the option to save documents to disk and open them only one at a time for processing. This means that much larger corpora can be used. A datastore can also be useful for saving documents in an efficient and lossless way.

To save a text in a datastore, a new datastore must first be created if one does not already exist. Create a datastore by right clicking on Datastore in the left hand pane, and select the option 'Create Datastore'. Select the data store type you wish to use. Create a directory to be used as the datastore (note that the datastore is a directory and not a file).

You can either save a whole corpus to the datastore (in which case the structure of the corpus will be preserved) or you can save individual documents. The recommended method is to save the whole corpus. To save a corpus, right click on the corpus name and select the 'Save to...' option (giving the name of the datastore created earlier). To save individual documents to the datastore, right clicking on each document name and follow the same procedure.

²Gorey details: features of annotations and documents in GATE may be any virtually any Java object; serialising arbitrary binary data to XML is not simple; instead we serialise them as strings, and therefore they will be re-loaded as strings.

To load a document from a datastore, do not try to load it as a language resource. Instead, open the datastore by right clicking on Datastore in the left hand pane, select ‘Open Datastore’ and choose the datastore to open. The datastore tree will appear in the main window. Double click on a corpus or document in this tree to open it. To save a corpus and document back to the same datastore, simply select the ‘Save’ option.

See also the movie for creating a datastore and the movie for loading corpus and documents from a datastore.

3.8.3 Saving Application States to a File

Resources, and applications that are made up of them, are created based on the settings of their parameters (see Section 3.6). It is possible to save the data used to create an application to a file and re-load it later. To save the application to a file, right click on it in the resources tree and select ‘Save application state’, which will give you a file creation dialogue. Choose a file name that ends in `gapp` as this file dialog and the one for loading application states age displays all files which have a name ending in `gapp`. A common convention is to use `.gapp` as a file extension.

To restore the application later, select ‘Restore application from file’ from the ‘File’ menu.

Note that the data that is saved represents how to *recreate* an application – not the resources that make up the application itself. So, for example, if your application has a resource that initialises itself from some file (e.g. a grammar, a document) then that file must still exist when you restore the application.

In case you don’t want to save the corpus configuration associated with the application then you must select ‘<none>’ in the corpus list of the application before saving the application.

The file resulting from saving the application state contains the values of the initialisation and runtime parameters for all the processing resources contained by the stored application as well as the values of the initialisation parameters for all the language resources referenced by those processing resources. Note that if you reference a document that has been created with an empty URL and empty string content parameter and subsequently been manually edited to add content, that content will not be saved. In order for document content to be preserved, load the document from an URL, specify the content as for the string content parameter or use a document from a datastore.

For the parameters of type URL (which are typically used to select external resources such as grammars or rules files) a transformation is applied so that the paths are stored relative to the location of the file used to store the state. This means that the resource files which are not part of GATE and used by an application do not need to be in the same location as when the application was initially created but rather in the same *location relative to the location of the application file*. It also means that the resource files which are part of GATE should be found correctly no matter where GATE is installed. This allows the creation and

deployment of portable applications by keeping the application file and the resource files used by the application together.

If you want to save your application along with all the resources it requires you can use the ‘Export for Teamware’ option (see Section 3.8.4).

See also the movie for saving and restoring applications.

3.8.4 Saving an Application with its Resources (e.g. GATE Teamware)

When you save an application using the ‘Save application state’ option (see Section 3.8.3), the saved file contains references to the plugins that were loaded when the application was saved, and to any resource files required by the application. To be able to reload the file, these plugins and other dependencies must exist at the same locations (relative to the saved state file). While this is fine for saving and loading applications on a single machine it means that if you want to package your application to run it elsewhere (e.g. deploy it to a GATE Teamware installation) then you need to be careful to include all the resource files and plugins at the right locations in your package. The ‘Export for Teamware’ option on the right-click menu for an application helps to automate this process.

When you export an application in this way, GATE Developer produces a ZIP file containing the saved application state (in the same format as ‘Save application state’). Any plugins and resource files that the application refers to are also included in the zip file, and the relative paths in the saved state are rewritten to point to the correct locations within the package. The resulting package is therefore self-contained and can be copied to another machine and unpacked there, or passed to your Teamware Administrator for deployment.

As well as selecting the location where you want to save the package, the ‘Export for Teamware’ option will also prompt you to select the annotation sets that your application uses for input and output. For example, if your application makes use of the unpacked XML markup in source documents and creates annotations in the default set then you would select ‘Original markups’ as an input set and the ‘<Default annotation set>’ as an output set. GATE Developer will try to make an educated guess at the correct sets but you should check and amend the lists as necessary.

There are a few important points to note about the export process:

- The complete contents of all the plugin directories that are loaded when you perform the export will be included in the resulting package. Use the plugin manager to unload any plugins your application is not using before you export it.
- If your application refers to a resource file in a directory that is not under one of the loaded plugins, the entire contents of this directory will be recursively included in the

package. If you have a number of unrelated resources in a single directory (e.g. many sets of large gazetteer lists) you may want to separate them into separate directories so that only the relevant ones are included in the package.

- The packager only knows about resources that your application refers to directly in its parameters. For example, if your application includes a multi-phase JAPE grammar the packager will only consider the main grammar file, not any of its sub-phases. If the sub-phases are not contained in the same directory as the main grammar you may find they are not included. If indirect references of this kind are all to files under the same directory as the ‘master’ file it will work OK.

If you require more flexibility than this option provides you should read Section E.2, which describes the underlying Ant task that the exporter uses.

3.9 Keyboard Shortcuts

You can use various keyboard shortcuts for common tasks in GATE Developer. These are listed in this section.

General (Section 3.1):

- **F1** Display a help page for the selected component
- **Alt+F4** Exit the application without confirmation
- **Tab** Put the focus on the next component or frame
- **Shift+Tab** Put the focus on the previous component or frame
- **F6** Put the focus on the next frame
- **Shift+F6** Put the focus on the previous frame
- **Alt+F** Show the File menu
- **Alt+O** Show the Options menu
- **Alt+T** Show the Tools menu
- **Alt+H** Show the Help menu
- **F10** Show the first menu

Resources tree (Section 3.1):

- **Enter** Show the selected resources
- **Ctrl+H** Hide the selected resource
- **Ctrl+Shift+H** Hide all the resources
- **F2** Rename the selected resource
- **Ctrl+F4** Close the selected resource

Document editor (Section 3.2):

- **Ctrl+F** Show the search dialog for the document
- **Ctrl+S** Save the document in a file
- **F3** Show/Hide the annotation sets
- **Shift+F3** Show the annotation sets with preselection
- **F4** Show/Hide the annotations list
- **F5** Show/Hide the coreference editor
- **F7** Show/Hide the text

Annotation editor (Section 3.4):

- **Right/Left** Grow/Shrink the annotation span at its start
- **Alt+Right/Alt+Left** Grow/Shrink the annotation span at its end
- **+Shift/+Ctrl+Shift** Use a span increment of 5/10 characters
- **Alt+Delete** Delete the currently edited annotation

Annac/Lucene datastore (Chapter 9):

- **Alt+Enter** Search the expression in the datastore
- **Alt+Backspace** Delete the search expression
- **Alt+Right** Display the next page of results
- **Alt+Left** Display the row manager
- **Alt+E** Export the results to a file

Annic/Lucene query text field (Chapter 9):

- **Ctrl+Enter** Insert a new line
- **Enter** Search the expression
- **Alt+Top** Select the previous result
- **Alt+Bottom** Select the next result

3.10 Miscellaneous

3.10.1 Stopping GATE from Restoring Developer Sessions/Options

GATE can remember Developer options and the state of the resource tree when it exits. The options are saved by default; the session state is not saved by default. This default behaviour can be changed from the ‘Advanced’ tab of the ‘Configuration’ choice on the ‘Options’ menu.

If a problem occurs and the saved data prevents GATE Developer from starting, you can fix this by deleting the configuration and session data files. These are stored in your home directory, and are called `gate.xml` and `gate.session` or `.gate.xml` and `.gate.session` depending on platform. On Windows your home is:

95, 98, NT: Windows Directory/profiles/username

2000, XP: Windows Drive/Documents and Settings/username

3.10.2 Working with Unicode

GATE provides various facilities for working with Unicode beyond those that come as default with Java³:

1. a Unicode editor with input methods for many languages;
2. use of the input methods in all places where text is edited in the GUI;

³Implemented by Valentin Tablan, Mark Leisher and Markus Kramer. Initial version developed by Mark Leisher.

3. a development kit for implementing input methods;
4. ability to read diverse character encodings.

1 using the editor:

In GATE Developer, select ‘Unicode editor’ from the ‘Tools’ menu. This will display an editor window, and, when a language with a custom input method is selected for input (see next section), a virtual keyboard window with the characters of the language assigned to the keys on the keyboard. You can enter data either by typing as normal, or with mouse clicks on the virtual keyboard.

2 configuring input methods:

In the editor and in GATE Developer’s main window, the ‘Options’ menu has an ‘Input methods’ choice. All supported input languages (a superset of the JDK languages) are available here. Note that you need to use a font capable of displaying the language you select. By default GATE Developer will choose a Unicode font if it can find one on the platform you’re running on. Otherwise, select a font manually from the ‘Options’ menu ‘Configuration’ choice.

3 using the development kit:

GUK, the GATE Unicode Kit, is documented at:

<http://gate.ac.uk/gate/doc/javadoc/guk/package-summary.html>.

4 reading different character encodings:

When you create a document from a URL pointing to textual data in GATE, you have to tell the system what character encoding the text is stored in. By default, GATE will set this parameter to be the empty string. This tells Java to use the default encoding for whatever platform it is running on at the time – e.g. on Western versions of Windows this will be ISO-8859-1, and Eastern ones ISO-8859-9. A popular way to store Unicode documents is in UTF-8, which is a superset of ASCII (but can still store all Unicode data); if you get an error message about document I/O during reading, try setting the encoding to UTF-8, or some other locally popular encoding. (To see a list of available encodings, try opening a document in GATE’s unicode editor – you will be prompted to select an encoding.)

Chapter 4

CREOLE: the GATE Component Model

... Noam Chomsky's answer in *Secrets, Lies and Democracy* (David Barsamian 1994; Odonian) to 'What do you think about the Internet?'

'I think that there are good things about it, but there are also aspects of it that concern and worry me. This is an intuitive response – I can't prove it – but my feeling is that, since people aren't Martians or robots, direct face-to-face contact is an extremely important part of human life. It helps develop self-understanding and the growth of a healthy personality.

'You just have a different relationship to somebody when you're looking at them than you do when you're punching away at a keyboard and some symbols come back. I suspect that extending that form of abstract and remote relationship, instead of direct, personal contact, is going to have unpleasant effects on what people are like. It will diminish their humanity, I think.'

Chomsky, quoted at <http://photo.net/wtr/dead-trees/53015.htm>.

The GATE architecture is based on components: reusable chunks of software with well-defined interfaces that may be deployed in a variety of contexts. The design of GATE is based on an analysis of previous work on infrastructure for LE, and of the typical types of software entities found in the fields of NLP and CL (see in particular chapters 4–6 of [Cunningham 00]). Our research suggested that a profitable way to support LE software development was an architecture that breaks down such programs into components of various types. Because LE practice varies very widely (it is, after all, predominantly a research field), the architecture must avoid restricting the sorts of components that developers can plug into the infrastructure. The GATE framework accomplishes this via an adapted version of the *Java Beans* component framework from Sun, as described in section 4.2.

GATE components may be implemented by a variety of programming languages and databases, but in each case they are represented to the system as a Java class. This class

may do nothing other than call the underlying program, or provide an access layer to a database; on the other hand it may implement the whole component.

GATE components are one of three types:

- LanguageResources (LRs) represent entities such as lexicons, corpora or ontologies;
- ProcessingResources (PRs) represent entities that are primarily algorithmic, such as parsers, generators or ngram modellers;
- VisualResources (VRs) represent visualisation and editing components that participate in GUIs.

The distinction between language resources and processing resources is explored more fully in section C.1.1. Collectively, the set of resources integrated with GATE is known as **CREOLE**: a Collection of REusable Objects for Language Engineering.

In the rest of this chapter:

- Section 4.3 describes the lifecycle of GATE components;
- Section 4.4 describes how Processing Resources can be grouped into applications;
- Section 4.5 describes the relationship between Language Resources and their datasources;
- Section 4.6 summarises GATE's set of built-in components;
- Section 4.7 describes how configuration data for Resource types is supplied to GATE.

4.1 The Web and CREOLE

GATE allows resource implementations and Language Resource persistent data to be distributed over the Web, and uses Java annotations and XML for configuration of resources (and GATE itself).

Resource implementations are grouped together as 'plugins', stored at a URL (when the resources are in the local file system this can be a `file:/` URL). When a plugin is loaded into GATE it looks for a configuration file called `creole.xml` relative to the plugin URL and uses the contents of this file to determine what resources this plugin declares and where to find the classes that implement the resource types (typically these classes are stored in a JAR file in the plugin directory). Configuration data for the resources may be stored directly in the `creole.xml` file, or it may be stored as Java annotations on the resource classes themselves; in either case GATE retrieves this configuration information and adds the resource definitions

to the CREOLE register. When a user requests an instantiation of a resource, GATE creates an instance of the resource class in the virtual machine.

Language resource data can be stored in binary serialised form in the local file system.

4.2 The GATE Framework

We can think of the GATE framework as a backplane into which users can plug CREOLE components. The user gives the system a list of URLs to search when it starts up, and components at those locations are loaded by the system.

The backplane performs these functions:

- component discovery, bootstrapping, loading and reloading;
- management and visualisation of native data structures for common information types;
- generalised data storage and process execution.

A set of components plus the framework is a deployment unit which can be embedded in another application.

At their most basic, all GATE resources are *Java Beans*, the Java platform's model of software components. Beans are simply Java classes that obey certain interface conventions:

- beans must have no-argument constructors.
- beans have *properties*, defined by pairs of methods named by the convention `setProp` and `getProp`.

GATE uses Java Beans conventions to construct and configure resources at runtime, and defines interfaces that different component types must implement.

4.3 The Lifecycle of a CREOLE Resource

CREOLE resources exhibit a variety of forms depending on the perspective they are viewed from. Their implementation is as a Java class plus an XML metadata file living at the same URL. When using GATE Developer, resources can be loaded and viewed via the resources tree (left pane) and the 'create resource' mechanism. When programming with GATE Embedded, they are Java objects that are obtained by making calls to GATE's `Factory` class. These various incarnations are the phases of a CREOLE resource's 'lifecycle'.

Depending on what sort of task you are using GATE for, you may use resources in any or all of these phases. For example, you may only be interested in getting a graphical view of what GATE's ANNIE Information Extraction system (see Chapter 6) does; in this case you will use GATE Developer to load the ANNIE resources, and load a document, and create an ANNIE application and run it on the document. If, on the other hand, you want to create your own resources, or modify the Java code of an existing resource (as opposed to just modifying its grammar, for example), you will need to deal with all the lifecycle phases.

The various phases may be summarised as:

Creating a new resource from scratch (bootstrapping). To create the binary image of a resource (a Java class in a JAR file), and the XML file that describes the resource to GATE, you need to create the appropriate `.java` file(s), compile them and package them as a `.jar`. GATE provides a bootstrap tool to start this process – see Section 7.11. Alternatively you can simply copy code from an existing resource.

Instantiating a resource in GATE Embedded. To create a resource in your own Java code, use GATE's `Factory` class (this takes care of parameterising the resource, restoring it from a database where appropriate, etc. etc.). Section 7.2 describes how to do this.

Loading a resource into GATE Developer. To load a resource into GATE Developer, use the various 'New ... resource' options from the `File` menu and elsewhere. See Section 3.1.

Resource configuration and implementation. GATE's bootstrap tool will create an empty resource that does nothing. In order to achieve the behaviour you require, you'll need to change the configuration of the resource (by editing the `creole.xml` file) and/or change the Java code that implements the resource. See section 4.7.

4.4 Processing Resources and Applications

PRs can be combined into *applications*. Applications model a control strategy for the execution of PRs. In GATE, applications are called 'controllers' accordingly.

Currently only sequential, or pipeline, execution is supported. There are two main types of pipeline:

Simple pipelines simply group a set of PRs together in order and execute them in turn. The implementing class is called `SerialController`.

Corpus pipelines are specific for LanguageAnalysers – PRs that are applied to documents and corpora. A corpus pipeline opens each document in the corpus in turn, sets that document as a runtime parameter on each PR, runs all the PRs on the corpus, then closes the document. The implementing class is called **SerialAnalyserController**.

Conditional versions of these controllers are also available. These allow processing resources to be run conditionally on document features. See Section 3.7.2 for how to use these.

Controllers are themselves PRs – in particular a simple pipeline is a standard PR and a corpus pipeline is a LanguageAnalyser – so one pipeline can be nested in another. This is particularly useful with conditional controllers to group together a set of PRs that can all be turned on or off as a group.

There is also a real-time version of the corpus pipeline. When creating such a controller, a `timeout` parameter needs to be set which determines the maximum amount of time (in milliseconds) allowed for the processing of a document. Documents that take longer to process, are simply ignored and the execution moves to the next document after the timeout interval has lapsed.

All controllers have special handling for processing resources that implement the interface `gate.creole.ControllerAwarePR`. This interface provides methods that are called by the controller at the start and end of the whole application's execution – for a corpus pipeline, this means before any document has been processed and after all documents in the corpus have been processed, which is useful for PRs that need to share data structures across the whole corpus, build aggregate statistics, etc. For full details, see the JavaDoc documentation for `ControllerAwarePR`.

4.5 Language Resources and Datastores

Language Resources can be stored in Datastores. Datastores are an abstract model of disk-based persistence, which can be implemented by various types of storage mechanism. Here are the types implemented:

Serial Datastores are based on Java's serialisation system, and store data directly into files and directories.

Lucene Datastores is a full-featured annotation indexing and retrieval system. It is provided as part of an extension of the Serial Datastores. See Section 9 for more details.

4.6 Built-in CREOLE Resources

GATE comes with various built-in components:

- Language Resources modelling Documents and Corpora, and various types of Annotation Schema – see Chapter 5.
- Processing Resources that are part of the ANNIE system – see Chapter 6.
- Gazetteers – see Chapter 13.
- Ontologies – see Chapter 14.
- Machine Learning resources – see Chapter 15.
- Alignment tools – see Chapter 16.
- Parsers and taggers – see Chapter 17.
- Other miscellaneous resources – see Chapter 19.

4.7 CREOLE Resource Configuration

This section describes how to supply GATE with the configuration data it needs about a resource, such as what its parameters are, how to display it if it has a visualisation, etc. Several GATE resources can be grouped into a single *plugin*, which is a directory containing an XML configuration file called `creole.xml`. Configuration data for the plugin's resources can be given in the `creole.xml` file or directly in the Java source file using Java 5 annotations.

A `creole.xml` file has a root element `<CREOLE-DIRECTORY>`, but the further contents of this element depend on the configuration style. The following three sections discuss the different styles – all-XML, all-annotations and a mixture of the two.

4.7.1 Configuration with XML

To configure your resources in the `creole.xml` file, the `<CREOLE-DIRECTORY>` element should contain one `<RESOURCE>` element for each resource type in the plugin. The `<RESOURCE>` elements may optionally be contained within a `<CREOLE>` element (to allow a single `creole.xml` file to be built up by concatenating multiple separate files). For example:


```
<CREOLE-DIRECTORY>
```

```
<CREOLE>
```

```
<RESOURCE>
```

```
<NAME>Minipar Wrapper</NAME>
```

```
<JAR>MiniparWrapper.jar</JAR>
```

```
<CLASS>minipar.Minipar</CLASS>
```

```
<COMMENT>MiniPar is a shallow parser. It determines the  
dependency relationships between the words of a sentence.</COMMENT>
```

```
<HELPURL>http://gate.ac.uk/cgi-bin/userguide/sec:parsers:minipar</HELPURL>
```

```
<PARAMETER NAME="document"
```

```
RUNTIME="true"
```

```
COMMENT="document to process">gate.Document</PARAMETER>
```

```
<PARAMETER NAME="miniparDataDir"
```

```
RUNTIME="true"
```

```
COMMENT="location of the Minipar data directory">
```

```
java.net.URL
```

```
</PARAMETER>
```

```
<PARAMETER NAME="miniparBinary"
```

```
RUNTIME="true"
```

```
COMMENT="Name of the Minipar command file">
```

```
java.net.URL
```

```
</PARAMETER>
```

```
<PARAMETER NAME="annotationInputSetName"
```

```
RUNTIME="true"
```

```
OPTIONAL="true"
```

```
COMMENT="Name of the input Source">
```

```
java.lang.String
```

```
</PARAMETER>
```

```
<PARAMETER NAME="annotationOutputSetName"
```

```
RUNTIME="true"
```

```
OPTIONAL="true"
```

```
COMMENT="Name of the output AnnotationSetName">
```

```
java.lang.String
```

```
</PARAMETER>
```

```
<PARAMETER NAME="annotationTypeName"
```

```
RUNTIME="false"
```

```
DEFAULT="DepTreeNode"
```

```
COMMENT="Annotations to store with this type">
```

```
java.lang.String
```

```
</PARAMETER>
```

```
</RESOURCE>
```

```
</CREOLE>
```

```
</CREOLE-DIRECTORY>
```

Basic Resource-Level Data

Each resource must give a name, a Java class and the JAR file that it can be loaded from. The above example is taken from the `Parser_Minipar` plugin, and defines a single resource with a number of parameters.

The full list of valid elements under `<RESOURCE>` is as follows:

NAME the name of the resource, as it will appear in the ‘New’ menu in GATE Developer. If omitted, defaults to the bare name of the resource class (without a package name).

CLASS the fully qualified name of the Java class that implements this resource.

JAR names JAR files required by this resource (paths are relative to the location of `creole.xml`). Typically this will be the JAR file containing the class named by the `<CLASS>` element, but additional `<JAR>` elements can be used to name third-party JAR files that the resource depends on.

COMMENT a descriptive comment about the resource, which will appear as the tooltip when hovering over an instance of this resource in the resources tree in GATE Developer. If omitted, no comment is used.

HELPURL a URL to a help document on the web for this resource. It is used in the help browser inside GATE Developer.

INTERFACE the interface type implemented by this resource, for example new types of document would specify `<INTERFACE>gate.Document</INTERFACE>`.

ICON the icon used to represent this resource in GATE Developer. This is a path inside the plugin’s JAR file, for example `<ICON>/some/package/icon.png</ICON>`. If the path specified does not start with a forward slash, it is assumed to name an icon from the GATE default set, which is located in `gate.jar` at `gate/resources/img`. If no icon is specified, a generic language resource or processing resource icon (as appropriate) is used.

PRIVATE if present, this resource type is hidden in the GATE Developer GUI, i.e. it is not shown in the ‘New’ menus. This is useful for resource types that are intended to be created internally by other resources, or for resources that have parameters of a type that cannot be set in the GUI. `<PRIVATE/>` resources can still be created in Java code using the `Factory`.

AUTOINSTANCE (and HIDDEN-AUTOINSTANCE) tells GATE to automatically create instances of this resource when the plugin is loaded. Any number of auto instances may be defined, GATE will create them all. Each `<AUTOINSTANCE>` element may optionally contain `<PARAM NAME="..." VALUE="..." />` elements giving parameter values to use when creating the instance. Any parameters not specified explicitly

will take their default values. Use `<HIDDEN-AUTOINSTANCE>` if you want the auto instances not to show up in GATE Developer – this is useful for things like document formats where there should only ever be a single instance in GATE and that instance should not be deleted.

TOOL if present, this resource type is considered to be a “tool”. Tools can contribute items to the Tools menu in GATE Developer.

For visual resources, a `<GUI>` element should also be provided. This takes a `TYPE` attribute, which can have the value `LARGE` or `SMALL`. `LARGE` means that the visual resource is a large viewer and should appear in the main part of the GATE Developer window on the right hand side, `SMALL` means the VR is a small viewer which appears in the space below the resources tree in the bottom left. The `<GUI>` element supports the following sub-elements:

RESOURCE_DISPLAYED the type of GATE resource this VR can display. Any resource whose type is assignable to this type will be displayed with this viewer, so for example a VR that can display all types of document would specify `gate.Document`, whereas a VR that can only display the default GATE document implementation would specify `gate.corpora.DocumentImpl`.

MAIN_VIEWER if present, GATE will consider this VR to be the ‘most important’ viewer for the given resource type, and will ensure that if several different viewers are all applicable to this resource, this viewer will be the one that is initially visible.

For annotation viewers, you should specify an `<ANNOTATION_TYPE_DISPLAYED>` element giving the annotation type that the viewer can display (e.g. `Sentence`).

Resource Parameters

Resources may also have parameters of various types. These resources, from the GATE distribution, illustrate the various types of parameters:

```
<RESOURCE>
  <NAME>GATE document</NAME>
  <CLASS>gate.corpora.DocumentImpl</CLASS>
  <INTERFACE>gate.Document</INTERFACE>
  <COMMENT>GATE transient document</COMMENT>
  <OR>
    <PARAMETER NAME="sourceUrl"
      SUFFIXES="txt;text;xml;xhtm;xhtml;html;htm;sgml;sgm;mail;email;eml;rtf"
      COMMENT="Source URL">java.net.URL</PARAMETER>
    <PARAMETER NAME="stringContent"
```

```

    COMMENT="The content of the document">java.lang.String</PARAMETER>
</OR>
<PARAMETER
  COMMENT="Should the document read the original markup"
  NAME="markupAware" DEFAULT="true">java.lang.Boolean</PARAMETER>
<PARAMETER NAME="encoding" OPTIONAL="true"
  COMMENT="Encoding" DEFAULT="">java.lang.String</PARAMETER>
<PARAMETER NAME="sourceUrlStartOffset"
  COMMENT="Start offset for documents based on ranges"
  OPTIONAL="true">java.lang.Long</PARAMETER>
<PARAMETER NAME="sourceUrlEndOffset"
  COMMENT="End offset for documents based on ranges"
  OPTIONAL="true">java.lang.Long</PARAMETER>
<PARAMETER NAME="preserveOriginalContent"
  COMMENT="Should the document preserve the original content"
  DEFAULT="false">java.lang.Boolean</PARAMETER>
<PARAMETER NAME="collectRepositioningInfo"
  COMMENT="Should the document collect repositioning information"
  DEFAULT="false">java.lang.Boolean</PARAMETER>
<ICON>lr.gif</ICON>
</RESOURCE>

<RESOURCE>
  <NAME>Document Reset PR</NAME>
  <CLASS>gate.creole.annotdelete.AnnotationDeletePR</CLASS>
  <COMMENT>Document cleaner</COMMENT>
  <PARAMETER NAME="document" RUNTIME="true">gate.Document</PARAMETER>
  <PARAMETER NAME="annotationTypes" RUNTIME="true"
    OPTIONAL="true">java.util.ArrayList</PARAMETER>
</RESOURCE>

```

Parameters may be optional, and may have default values (and may have comments to describe their purpose, which is displayed by GATE Developer during interactive parameter setting).

Some PR parameters are execution time (RUNTIME), some are initialisation time. E.g. at execution time a doc is supplied to a language analyser; at initialisation time a grammar may be supplied to a language analyser.

The <PARAMETER> tag takes the following attributes:

NAME: name of the JavaBean property that the parameter refers to, i.e. for a parameter named 'someParam' the class must have `setSomeParam` and `getSomeParam` methods.¹

¹The JavaBeans spec allows `is` instead of `get` for properties of the primitive type `boolean`, but GATE

DEFAULT: default value (see below).

RUNTIME: doesn't need setting at initialisation time, but must be set before calling `execute()`. Only meaningful for PRs

OPTIONAL: not required

COMMENT: for display purposes

ITEM_CLASS_NAME: (only applies to parameters whose type is `java.util.Collection` or a type that implements or extends this) this specifies the type of elements the collection contains, so GATE can use the right type when parameters are set. If omitted, GATE will pass in the elements as Strings.

SUFFIXES: (only applies to parameters of type `java.net.URL`) a semicolon-separated list of file suffixes that this parameter typically accepts, used as a filter in the file chooser provided by GATE Developer to select a local file as the parameter value.

It is possible for two or more parameters to be mutually exclusive (i.e. a user must specify one or the other but not both). In this case the `<PARAMETER>` elements should be grouped together under an `<OR>` element.

The type of the parameter is specified as the text of the `<PARAMETER>` element, and the type supplied must match the return type of the parameter's `get` method. Any reference type (class, interface or enum) may be used as the parameter type, including other resource types – in this case GATE Developer will offer a list of the loaded instances of that resource as options for the parameter value. Primitive types (`char`, `boolean`, ...) are not supported, instead you should use the corresponding wrapper type (`java.lang.Character`, `java.lang.Boolean`, ...). If the getter returns a parameterized type (e.g. `List<Integer>`) you should just specify the raw type (`java.util.List`) here².

The `DEFAULT` string is converted to the appropriate type for the parameter - `java.lang.String` parameters use the value directly, primitive wrapper types e.g. `java.lang.Integer` use their respective `valueOf` methods, and other built-in Java types can have defaults specified provided they have a constructor taking a `String`.

The type `java.net.URL` is treated specially: if the default string is not an absolute URL (e.g. `http://gate.ac.uk/`) then it is treated as a path relative to the location of the `creole.xml` file. Thus a `DEFAULT` of `'resources/main.jape'` in the file `file:/opt/MyPlugin/creole.xml` is treated as the absolute URL `file:/opt/MyPlugin/resources/main.jape`.

For `Collection`-valued parameters multiple values may be specified, separated by semicolons, e.g. `'foo;bar;baz'`; if the parameter's type is an interface – `Collection` or one of

does not support parameters with primitive types. Parameters of type `java.lang.Boolean` (the wrapper class) are permitted, but these have `get` accessors anyway.

²In this particular case, as the type is a collection, you would specify `java.lang.Integer` as the `ITEM_CLASS_NAME`.

its sub-interfaces (e.g. `List`) – a suitable concrete class (e.g. `ArrayList`, `HashSet`) will be chosen automatically for the default value.

For parameters of type `gate.FeatureMap` multiple `name=value` pairs can be specified, e.g. `'kind=word;orth=upperInitial'`. For enum-valued parameters the default string is taken as the name of the enum constant to use. Finally, if no `DEFAULT` attribute is specified, the default value is `null`.

4.7.2 Configuring Resources using Annotations

As an alternative to the XML configuration style, GATE provides Java 5 annotation types to embed the configuration data directly in the Java source code. `@CreoleResource` is used to mark a class as a GATE resource, and parameter information is provided through annotations on the JavaBean `set` methods. At runtime these annotations are read and mapped into the equivalent entries in `creole.xml` before parsing. The metadata annotation types are all marked `@Documented` so the CREOLE configuration data will be visible in the generated JavaDoc documentation.

For more detailed information, see the JavaDoc documentation for `gate.creole.metadata`.

To use annotation-driven configuration a `creole.xml` file is still required but it need only contain the following:

```
<CREOLE-DIRECTORY>
  <JAR SCAN="true">myPlugin.jar</JAR>
  <JAR>lib/thirdPartyLib.jar</JAR>
</CREOLE-DIRECTORY>
```

This tells GATE to load `myPlugin.jar` and scan its contents looking for resource classes annotated with `@CreoleResource`. Other JAR files required by the plugin can be specified using other `<JAR>` elements without `SCAN="true"`.

Basic Resource-Level Data

To mark a class as a CREOLE resource, simply use the `@CreoleResource` annotation (in the `gate.creole.metadata` package), for example:

```
1 import gate.creole.AbstractLanguageAnalyser;
2 import gate.creole.metadata.*;
3
4 @CreoleResource(name = "GATE Tokeniser",
5                 comment = "Splits text into tokens and spaces")
6 public class Tokeniser extends AbstractLanguageAnalyser {
7     ...
```

The `@CreoleResource` annotation provides slots for all the values that can be specified under `<RESOURCE>` in `creole.xml`, except `<CLASS>` (inferred from the name of the annotated class) and `<JAR>` (taken to be the JAR containing the class):

name (String) the name of the resource, as it will appear in the ‘New’ menu in GATE Developer. If omitted, defaults to the bare name of the resource class (without a package name). (XML equivalent `<NAME>`)

comment (String) a descriptive comment about the resource, which will appear as the tooltip when hovering over an instance of this resource in the resources tree in GATE Developer. If omitted, no comment is used. (XML equivalent `<COMMENT>`)

helpURL (String) a URL to a help document on the web for this resource. It is used in the help browser inside GATE Developer. (XML equivalent `<HELPURL>`)

isPrivate (boolean) should this resource type be hidden from the GATE Developer GUI, so it does not appear in the ‘New’ menus? If omitted, defaults to false (i.e. not hidden). (XML equivalent `<PRIVATE/>`)

icon (String) the icon to use to represent the resource in GATE Developer. If omitted, a generic language resource or processing resource icon is used. (XML equivalent `<ICON>`, see the description above for details)

interfaceName (String) the interface type implemented by this resource, for example a new type of document would specify `"gate.Document"` here. (XML equivalent `<INTERFACE>`)

autoInstances (array of `@AutoInstance` annotations) definitions for any instances of this resource that should be created automatically when the plugin is loaded. If omitted, no auto-instances are created by default. (XML equivalent, one or more `<AUTOINSTANCE>` and/or `<HIDDEN-AUTOINSTANCE>` elements, see the description above for details)

tool (boolean) is this resource type a tool?

For visual resources only, the following elements are also available:

guiType (GuiType enum) the type of GUI this resource defines. (XML equivalent `<GUI TYPE="LARGE|SMALL">`)

resourceDisplayed (String) the class name of the resource type that this VR displays, e.g. `"gate.Corpus"`. (XML equivalent `<RESOURCE_DISPLAYED>`)

mainViewer (boolean) is this VR the ‘most important’ viewer for its displayed resource type? (XML equivalent `<MAIN_VIEWER/>`, see above for details)

For annotation viewers, you should specify an `annotationTypeDisplayed` element giving the annotation type that the viewer can display (e.g. `Sentence`).

Resource Parameters

Parameters are declared by placing annotations on their JavaBean `set` methods. To mark a setter method as a parameter, use the `@CreoleParameter` annotation, for example:

```
@CreoleParameter(comment = "The location of the list of abbreviations")
public void setAbbrListUrl(URL listUrl) {
    ...
}
```

GATE will infer the parameter's name from the name of the JavaBean property in the usual way (i.e. strip off the leading `set` and convert the following character to lower case, so in this example the name is `abbrListUrl`). The parameter name is *not* taken from the name of the method parameter. The parameter's type is inferred from the type of the method parameter (`java.net.URL` in this case).

The annotation elements of `@CreoleParameter` correspond to the attributes of the `<PARAMETER>` tag in the XML configuration style:

comment (String) an optional descriptive comment about the parameter. (XML equivalent `COMMENT`)

defaultValue (String) the optional default value for this parameter. The value is specified as a string but is converted to the relevant type by GATE according to the conversions described in the previous section. Note that relative path default values for URL-valued parameters are still relative to the location of the `creole.xml` file, not the annotated class. (XML equivalent `DEFAULT`)

suffixes (String) for URL-valued parameters, a semicolon-separated list of default file suffixes that this parameter accepts. (XML equivalent `SUFFIXES`)

collectionElementType (Class) for Collection-valued parameters, the type of the elements in the collection. This can usually be inferred from the generic type information, for example `public void setIndices(List<Integer> indices)`, but must be specified if the `set` method's parameter has a raw (non-parameterized) type. (XML equivalent `ITEM_CLASS_NAME`)

Mutually-exclusive parameters (such as would be grouped in an `<OR>` in `creole.xml`) are handled by adding a `disjunction="label"` to the `@CreoleParameter` annotation – all parameters that share the same label are grouped in the same disjunction.

Optional and runtime parameters are marked using extra annotations, for example:

```
1  @Optional
2  @RunTime
3  @CreoleParameter
4  public void setAnnotationSetName(String asName) {
5      ...
}
```


Inheritance

Unlike with pure XML configuration, when using annotations a resource will inherit any configuration data that was not explicitly specified from annotations on its parent class and on any interfaces it implements. Specifically, if you do not specify a comment, interfaceName, icon, annotationTypeDisplayed or the GUI-related elements (guiType and resourceDisplayed) on your `@CreoleResource` annotation then GATE will look up the class tree for other `@CreoleResource` annotations, first on the superclass, its superclass, etc., then at any implemented interfaces, and use the first value it finds. This is useful if you are defining a family of related resources that inherit from a common base class.

The resource name and the `isPrivate` and `mainViewer` flags are *not* inherited.

Parameter definitions are inherited in a similar way. This is one of the big advantages of annotation configuration over pure XML – if one resource class extends another then with pure XML configuration all the parent class’s parameter definitions must be duplicated in the subclass’s `creole.xml` definition. With annotations, parameters are inherited from the parent class (and its parent, etc.) as well as from any interfaces implemented. For example, the `gate.LanguageAnalyser` interface provides two parameter definitions via annotated `set` methods, for the `corpus` and `document` parameters. Any `@CreoleResource` annotated class that implements `LanguageAnalyser`, directly or indirectly, will get these parameters automatically.

Of course, there are some cases where this behaviour is not desirable, for example if a subclass calculates a value for a superclass parameter rather than having the user set it directly. In this case you can hide the parameter by overriding the `set` method in the subclass and using a marker annotation:

```

1  @HiddenCreoleParameter
2  public void setSomeParam(String someParam) {
3      super.setSomeParam(someParam);
4  }
```

The overriding method will typically just call the superclass one, as its only purpose is to provide a place to put the `@HiddenCreoleParameter` annotation.

Alternatively, you may want to override some of the configuration for a parameter but inherit the rest from the superclass. Again, this is handled by trivially overriding the `set` method and re-annotating it:

```

1  // superclass
2  @CreoleParameter(comment = "Location of the grammar file",
3                  suffixes = "jape")
4  public void setGrammarUrl(URL grammarLocation) {
5      ...
6  }
7
8  @Optional
9  @RunTime
```

```

10  @CreoleParameter(comment = "Feature to set on success")
11  public void setSuccessFeature(String name) {
12      ...
13  }

1  //-----
2  // subclass
3
4  // override the default value, inherit everything else
5  @CreoleParameter(defaultValue = "resources/defaultGrammar.jape")
6  public void setGrammarUrl(URL url) {
7      super.setGrammarUrl(url);
8  }
9
10 // we want the parameter to be required in the subclass
11 @Optional(false)
12 @CreoleParameter
13 public void setSuccessFeature(String name) {
14     super.setSuccessFeature(name);
15 }

```

Note that for backwards compatibility, data is only inherited from superclass annotations if the subclass is itself annotated with `@CreoleResource`. If the subclass is not annotated then GATE assumes that *all* its configuration is contained in `creole.xml` in the usual way.

4.7.3 Mixing the Configuration Styles

It is possible and often useful to mix and match the XML and annotation-driven configuration styles. The rule is always that *anything specified in the XML takes priority over the annotations*. The following examples show what this allows.

Overriding Configuration for a Third-Party Resource

Suppose you have a plugin from some third party that uses annotation-driven configuration. You don't have the source code but you would like to override the default value for one of the parameters of one of the plugin's resources. You can do this in the `creole.xml`:

```

<CREOLE-DIRECTORY>
  <JAR SCAN="true">acmePlugin-1.0.jar</JAR>

  <!-- Add the following to override the annotations -->
  <RESOURCE>
    <CLASS>com.acme.plugin.UsefulPR</CLASS>
    <PARAMETER NAME="listUrl"
      DEFAULT="resources/myList.txt">java.net.URL</PARAMETER>
  </RESOURCE>

```

```
</RESOURCE>
</CREOLE-DIRECTORY>
```

The default value for the `listUrl` parameter in the annotated class will be replaced by your value.

External AUTOINSTANCES

For resources like document formats, where there should always and only be one instance in GATE at any time, it makes sense to put the auto-instance definitions in the `@CreoleResource` annotation. But if the automatically created instances are a convenience rather than a necessity it may be better to define them in XML so other users can disable them without re-compiling the class:

```
<CREOLE-DIRECTORY>
  <JAR SCAN="true">myPlugin.jar</JAR>

  <RESOURCE>
    <CLASS>com.acme.AutoPR</CLASS>
    <AUTOINSTANCE>
      <PARAM NAME="type" VALUE="Sentence" />
    </AUTOINSTANCE>
    <AUTOINSTANCE>
      <PARAM NAME="type" VALUE="Paragraph" />
    </AUTOINSTANCE>
  </RESOURCE>
</CREOLE-DIRECTORY>
```

Inheriting Parameters

If you would prefer to use XML configuration for your own resources, but would like to benefit from the parameter inheritance features of the annotation-driven approach, you can write a normal `creole.xml` file with all your configuration and just add a blank `@CreoleResource` annotation to your class. For example:

```
1 package com.acme;
2 import gate.*;
3 import gate.creole.metadata.CreoleResource;
4
5 @CreoleResource
6 public class MyPR implements LanguageAnalyser {
7     ...
8 }
```

```

<!-- creole.xml -->
<CREOLE-DIRECTORY>
  <CREOLE>
    <RESOURCE>
      <NAME>My Processing Resource</NAME>
      <CLASS>com.acme.MyPR</CLASS>
      <COMMENT>...</COMMENT>
      <PARAMETER NAME="annotationSetName"
        RUNTIME="true" OPTIONAL="true">java.lang.String</PARAMETER>
    <!--
      don't need to declare document and corpus parameters, they
      are inherited from LanguageAnalyser
    -->
    </RESOURCE>
  </CREOLE>
</CREOLE-DIRECTORY>

```

N.B. Without the `@CreoleResource` the parameters would not be inherited.

4.8 Tools: How to Add Utilities to GATE Developer

Visual Resources allow a developer to provide a GUI to interact with a particular resource type (PR or LR), but sometimes it is useful to provide general utilities for use in the GATE Developer GUI that are not tied to any specific resource type. Examples include the annotation diff tool and the Groovy console (provided by the `Groovy` plugin), both of which are self-contained tools that display in their own top-level window. To support this, the CREOLE model has the concept of a *tool*.

A resource type is marked as a tool by using the `<TOOL/>` element in its `creole.xml` definition, or by setting `tool = true` if using the `@CreoleResource` annotation configuration style. If a resource is declared to be a tool, and written to implement the `gate.gui.ActionsPublisher` interface, then whenever an instance of the resource is created its published actions will be added to the “Tools” menu in GATE Developer.

Since the published actions of *every* instance of the resource will be added to the tools menu, it is best not to use this mechanism on resource types that can be instantiated by the user. The “tool” marker is best used in combination with the “private” flag (to hide the resource from the list of available types in the GUI) and one or more hidden autoinstance definitions to create a limited number of instances of the resource when its defining plugin is loaded. See the `GroovySupport` resource in the `Groovy` plugin for an example of this.

4.8.1 Putting your tools in a sub-menu

If your plugin provides a number of tools (or a number of actions from the same tool) you may wish to organise your actions into one or more sub-menus, rather than placing them all on the single top-level tools menu. To do this, you need to put a special value into the actions returned by the tool's `getActions()` method:

```
1 action.putValue(GateConstants.MENU_PATH_KEY,  
2     new String[] {"Acme toolkit", "Statistics"});
```

The key must be `GateConstants.MENU_PATH_KEY` and the value must be an array of strings. Each string in the array represents the name of one level of sub-menus. Thus in the example above the action would be placed under “Tools → Acme toolkit → Statistics”. If no `MENU_PATH_KEY` value is provided the action will be placed directly on the Tools menu.

Chapter 5

Language Resources: Corpora, Documents and Annotations

Sometimes in life you've got to dance like nobody's watching.

...

I think they should introduce 'sleeping' to the Olympics. It would be an excellent field event, in which the 'athletes' (for want of a better word) all lay down in beds, just beyond where the javelins land, and the first one to fall asleep and not wake up for three hours would win gold. I, for one, would be interested in seeing what kind of personality would be suited to sleeping in a competitive environment.

...

Life is a mystery to be lived, not a problem to be solved.

Round Ireland with a Fridge, Tony Hawks, 1998 (pp. 119, 147, 179).

This chapter documents GATE's model of corpora, documents and annotations on documents. Section 5.1 describes the simple attribute/value data model that corpora, documents and annotations all share. Section 5.2, Section 5.3 and Section 5.4 describe corpora, documents and annotations on documents respectively. Section 5.5 describes GATE's support for diverse document formats, and Section 5.5.2 describes facilities for XML input/output.

5.1 Features: Simple Attribute/Value Data

GATE has a single model for information that describes documents, collections of documents (corpora), and annotations on documents, based on attribute/value pairs. Attribute names are strings; values can be any Java object. The API for accessing this feature data is Java's Map interface (part of the Collections API).

5.2 Corpora: Sets of Documents plus Features

A Corpus in GATE is a Java Set whose members are Documents. Both Corpora and Documents are types of LanguageResource (LR); all LRs have a FeatureMap (a Java Map) associated with them that stored attribute/value information about the resource. FeatureMaps are also used to associate arbitrary information with ranges of documents (e.g. pieces of text) via the annotation model (see below).

Documents have a DocumentContent which is a text at present (future versions may add support for audiovisual content) and one or more AnnotationSets which are Java Sets.

5.3 Documents: Content plus Annotations plus Features

Documents are modelled as content plus annotations (see Section 5.4) plus features (see Section 5.1). The content of a document can be any subclass of DocumentContent.

5.4 Annotations: Directed Acyclic Graphs

Annotations are organised in graphs, which are modelled as Java sets of Annotation. Annotations may be considered as the arcs in the graph; they have a start Node and an end Node, an ID, a type and a FeatureMap. Nodes have pointers into the sources document, e.g. character offsets.

5.4.1 Annotation Schemas

Annotation schemas provide a means to define types of annotations in GATE. GATE uses the XML Schema language supported by W3C for these definitions. When using GATE Developer to create/edit annotations, a component is available (`gate.gui.SchemaAnnotationEditor`) which is driven by an annotation schema file. This component will constrain the data entry process to ensure that only annotations that correspond to a particular schema are created. (Another component allows unrestricted annotations to be created.)

Schemas are resources just like other GATE components. Below we give some examples of such schemas. Section 3.4.6 describes how to create new schemas.

Date Schema

```
<?xml version="1.0"?>
<schema
xmlns="http://www.w3.org/2000/10/XMLSchema">
  <!-- XSchema deffinition for Date-->
  <element name="Date">
    <complexType>
      <attribute name="kind" use="optional">
        <simpleType>
          <restriction base="string">
            <enumeration value="date"/>
            <enumeration value="time"/>
            <enumeration value="dateTime"/>
          </restriction>
        </simpleType>
      </attribute>
    </complexType>
  </element>
</schema>
```

Person Schema

```
<?xml version="1.0"?>
<schema
xmlns="http://www.w3.org/2000/10/XMLSchema">
  <!-- XSchema definition for Person-->
  <element name="Person" />
</schema>
```

Address Schema

```
<?xml version="1.0"?> <schema
xmlns="http://www.w3.org/2000/10/XMLSchema">
  <!-- XSchema deffinition for Address-->
  <element name="Address">
    <complexType>
      <attribute name="kind" use="optional">
        <simpleType>
          <restriction base="string">
            <enumeration value="email"/>
            <enumeration value="url"/>
            <enumeration value="phone"/>
            <enumeration value="ip"/>
          </restriction>
        </simpleType>
      </attribute>
    </complexType>
  </element>
</schema>
```

```

        <enumeration value="street"/>
        <enumeration value="postcode"/>
        <enumeration value="country"/>
        <enumeration value="complete"/>
    </restriction>
</simpleType>
</attribute>
</complexType>
</element>
</schema>

```

5.4.2 Examples of Annotated Documents

This section shows some simple examples of annotated documents.

This material is adapted from [Grishman 97], the TIPSTER Architecture Design document upon which GATE version 1 was based. Version 2 has a similar model, although annotations are now graphs, and instead of multiple spans per annotation each annotation now has a single start/end node pair. The current model is largely compatible with [Bird & Liberman 99], and roughly isomorphic with "stand-off markup" as latterly adopted by the SGML/XML community.

Each example is shown in the form of a table. At the top of the table is the document being annotated; immediately below the line with the document is a ruler showing the position (byte offset) of each character (see TIPSTER Architecture Design Document).

Underneath this appear the annotations, one annotation per line. For each annotation is shown its Id, Type, Span (start/end offsets derived from the start/end nodes), and Features. Integers are used as the annotation Ids. The features are shown in the form name = value.

The first example shows a single sentence and the result of three annotation procedures: tokenization with part-of-speech assignment, name recognition, and sentence boundary recognition. Each token has a single feature, its part of speech (pos), using the tag set from the University of Pennsylvania Tree Bank; each name also has a single feature, indicating the type of name: person, company, etc.

Annotations will typically be organized to describe a hierarchical decomposition of a text. A simple illustration would be the decomposition of a sentence into tokens. A more complex case would be a full syntactic analysis, in which a sentence is decomposed into a noun phrase and a verb phrase, a verb phrase into a verb and its complement, etc. down to the level of individual tokens. Such decompositions can be represented by annotations on nested sets of spans. Both of these are illustrated in the second example, which is an elaboration of our first example to include parse information. Each non-terminal node in the parse tree is represented by an annotation of type parse.

Text				
Cyndi savored the soup.				
^0...^5...^10...^15...^20				
Annotations				
Id	Type	SpanStart	Span End	Features
1	token	0	5	pos=NP
2	token	6	13	pos=VBD
3	token	14	17	pos=DT
4	token	18	22	pos=NN
5	token	22	23	
6	name	0	5	name_type=person
7	sentence	0	23	

Table 5.1: Result of annotation on a single sentence

Text				
Cyndi savored the soup.				
^0...^5...^10...^15...^20				
Annotations				
Id	Type	SpanStart	Span End	Features
1	token	0	5	pos=NP
2	token	6	13	pos=VBD
3	token	14	17	pos=DT
4	token	18	22	pos=NN
5	token	22	23	
6	name	0	5	name_type=person
7	sentence	0	23	constituents=[1],[2],[3].[4],[5]

Table 5.2: Result of annotations including parse information

Text				
To: All Barnyard Animals				
^0...^5...^10...^15...^20.				
From: Chicken Little				
^25...^30...^35...^40..				
Date: November 10,1194				
...^50...^55...^60...^65.				
Subject: Descending Firmament				
.^70...^75...^80...^85...^90...^95				
Priority: Urgent				
.^100...^105...^110.				
The sky is falling. The sky is falling.				
....^120...^125...^130...^135...^140...^145...^150.				
Annotations				
Id	Type	SpanStart	Span End	Features
1	Addressee	4	24	
2	Source	31	45	
3	Date	53	69	ddmmyy=101194
4	Subject	78	98	
5	Priority	109	115	
6	Body	116	155	
7	Sentence	116	135	
8	Sentence	136	155	

Table 5.3: Annotation showing overall document structure

In most cases, the hierarchical structure could be recovered from the spans. However, it may be desirable to record this structure directly through a constituents feature whose value is a sequence of annotations representing the immediate constituents of the initial annotation. For the annotations of type parse, the constituents are either non-terminals (other annotations in the parse group) or tokens. For the sentence annotation, the constituents feature points to the constituent tokens. A reference to another annotation is represented in the table as "[Annotation Id]"; for example, "[3]" represents a reference to annotation 3. Where the value of an feature is a sequence of items, these items are separated by commas. No special operations are provided in the current architecture for manipulating constituents. At a less esoteric level, annotations can be used to record the overall structure of documents, including in particular documents which have structured headers, as is shown in the third example (Table 5.3).

If the Addressee, Source, ... annotations are recorded when the document is indexed for retrieval, it will be possible to perform retrieval selectively on information in particular fields. Our final example (Table 5.4) involves an annotation which effectively modifies the document. The current architecture does not make any specific provision for the modification

Text				
Topster tackles 2 terrorbytes.				
^0...^5...^10..^15..^20..^25..				
Annotations				
Id	Type	SpanStart	Span End	Features
1	token	0	7	pos=NP correction=TIPSTER
2	token	8	15	pos=VBZ
3	token	16	17	pos=CD
4	token	18	29	pos=NNS correction=terabytes
5	token	29	30	

Table 5.4: Annotation modifying the document

of the original text. However, some allowance must be made for processes such as spelling correction. This information will be recorded as a correction feature on token annotations and possibly on name annotations:

5.4.3 Creating, Viewing and Editing Diverse Annotation Types

Note that annotation types should consist of a single word with no spaces. Otherwise they may not be recognised by other components such as JAPE transducers, and may create problems when annotations are saved as inline ('Save Preserving Format' in the context menu).

To view and edit annotation types, see Section 3.4. To add annotations of a new type, see Section 3.4.5. To add a new annotation schema, see Section 3.4.6.

5.5 Document Formats

The following document formats are supported by GATE:

- Plain Text
- HTML
- SGML
- XML
- RTF
- Email

- PDF (some documents)
- Microsoft Word (some documents)

By default GATE will try and identify the type of the document, then strip and convert any markup into GATE's annotation format. To disable this process, set the `markupAware` parameter on the document to `false`.

When reading a document of one of these types, GATE extracts the text between tags (where such exist) and create a GATE annotation filled as follows:

The name of the tag will constitute the annotation's type, all the tags attributes will materialize in the annotation's features and the annotation will span over the text covered by the tag. A few exceptions of this rule apply for the RTF, Email and Plain Text formats, which will be described later in the input section of these formats.

The text between tags is extracted and appended to the GATE document's content and all annotations created from tags will be placed into a GATE annotation set named 'Original markups'.

Example:

If the markup is like this:

```
<aTagName attrib1="value1" attrib2="value2" attrib3="value3"> A
piece of text</aTagName>
```

then the annotation created by GATE will look like:

```
annotation.type = "aTagName";
annotation.fm = {attrib1=value1;attrib2=value2;attrib3=value3};
annotation.start = startNode;
annotation.end = endNode;
```

The `startNode` and `endNode` are created from offsets referring the beginning and the end of 'A piece of text' in the document's content.

The documents supported by GATE have to be in one of the encodings accepted by Java. The most popular is the 'UTF-8' encoding which is also the most storage efficient one for UNICODE. If, when loading a document in GATE the *encoding* parameter is set to ''(the empty string), then the default encoding of the platform will be used.

5.5.1 Detecting the Right Reader

In order to successfully apply the document creation algorithm described above, GATE needs to detect the proper reader to use for each document format. If the user knows in advance what kind of document they are loading then they can specify the MIME type (e.g. *text/html*) using the init parameter `mimeType`, and GATE will respect this. If an explicit type is not given, GATE attempts to determine the type by other means, taking into consideration (where possible) the information provided by three sources:

- Document's extension
- The web server's content type
- Magic numbers detection

The first represents the extension of a file like (*xml,htm,html,txt,sgm,rtf, etc*), the second represents the HTTP information sent by a web server regarding the content type of the document being send by it (*text/html; text/xml, etc*), and the third one represents certain sequences of chars which are ultimately number sequences. GATE is capable of supporting multimedia documents, if the right reader is added to the framework. Sometimes, multimedia documents are identified by a signature consisting in a sequence of numbers. Inside GATE they are called magic numbers. For textual documents, certain char sequences form such magic numbers. Examples of magic numbers sequences will be provided in the Input section of each format supported by GATE.

All those tests are applied to each document read, and after that, a voting mechanism decides what is the best reader to associate with the document. There is a degree of priority for all those tests. The document's extension test has the highest priority. If the system is in doubt which reader to choose, then the one associated with document's extension will be selected. The next higher priority is given to the web server's content type and the third one is given to the magic numbers detection. However, any two tests that identify the same mime type, will have the highest priority in deciding the reader that will be used. The web server test is not always successful as there might be documents that are loaded from a local file system, and the magic number detection test is not always applicable. In the next paragraphs we will see how those tests are performed and what is the general mechanism behind reader detection.

The method that detects the proper reader is a static one, and it belongs to the `gate.DocumentFormat` class. It uses the information stored in the maps filled by the `init()` method of each reader. This method comes with three signatures:

```
1  static public DocumentFormat getDocumentFormat( gate.Document
2  aGateDocument, URL url)
3
4  static public DocumentFormat getDocumentFormat(gate.Document
5  aGateDocument, String fileSuffix)
```

```

6
7  static public DocumentFormat getDocumentFormat (gate.Document
8  aGateDocument, MimeType mimeType)

```

The first two methods try to detect the right MimeType for the GATE document, and after that, they call the third one to return the reader associate with a MimeType. Of course, if an explicit mimeType parameter was specified, GATE calls the third form of the method directly, passing the specified type. GATE uses the implementation from ‘http://jigsaw.w3.org’ for mime types.

The magic numbers test is performed using the information form magic2mimeTypeMap map. Each key from this map, is searched in the first bufferSize (the default value is 2048) chars of text. The method that does this is called runMagicNumbers(InputStreamReader aReader) and it belongs to DocumentFormat class. More details about it can be found in the GATE API documentation.

In order to activate a reader to perform the unpacking, the creole definition of a GATE document defines a parameter called ‘markupAware’ initialized with a default value of **true**. This parameter, forces GATE to detect a proper reader for the document being read. If no reader is found, the document’s content is load and presented to the user, just like any other text editor (this for textual documents).

The next subsections investigates particularities for each format and will describe the file extensions registered with each document format.

5.5.2 XML

Input

GATE permits the processing of any XML document and offers support for XML namespaces. It benefits the power of Apache’s Xerces parser and also makes use of Sun’s JAXP layer. Changing the XML parser in GATE can be achieved by simply replacing the value of a Java system property (‘javax.xml.parsers.SAXParserFactory’).

GATE will accept any well formed XML document as input. Although it has the possibility to validate XML documents against DTDs it does not do so because the validating procedure is time consuming and in many cases it issues messages that are annoying for the user.

There is an open problem with the general approach of reading XML, HTML and SGML documents in GATE. As we previously said, the text covered by tags/elements is appended to the GATE document content and a GATE annotation refers to this particular span of text. When appending, in cases such as ‘end.</P><P>Start’ it might happen that the ending word of the previous annotation is concatenated with the beginning phrase of the annotation currently being created, resulting in a garbage input for GATE processing resources that

operate at the text surface.

Let's take another example in order to better understand the problem:

```
<title>This is a title</title><p>This is a paragraph</p><a  
href="#link">Here is an useful link</a>
```

When the markup is transformed to annotations, it is likely that the text from the document's content will be as follows:

```
This is a titleThis is a paragraphHere is an useful link
```

The annotations created will refer the right parts of the texts but for the GATE's processing resources like (tokenizer, gazetter, etc) which work on this text, this will be a major disaster. Therefore, in order to prevent this problem from happening, GATE checks if it's likely to join words and if this happens then it inserts a space between those words. So, the text will look like this after loaded in GATE Developer:

```
This is a title This is a paragraph Here is an useful link
```

There are cases when these words are meant to be joined, but they are rare. This is why it's an open problem.

The extensions associate with the XML reader are:

- xml
- xhtm
- xhtml

The web server content type associate with xml documents is: *text/xml*.

The magic numbers test searches inside the document for the XML(`<?xml version="1.0"`) signature. It is also able to detect if the XML document uses the semantics described in the GATE document format DTD (see 5.5.2 below) or uses other semantics.

Output

GATE is capable of ensuring persistence for its resources. The types of persistent storage used for Language Resources are:

- Java serialization;
- XML serialization.

We describe the latter case here.

XML persistence doesn't necessarily preserve all the objects belonging to the annotations, documents or corpora. Their features can be of all kinds of objects, with various layers of nesting. For example, *lists containing lists containing maps, etc.* Serializing these arbitrary data types in XML is not a simple task; GATE does the best it can, and supports native Java types such as Integers and Booleans, but where complex data types are used, information may be lost (the types will be converted into Strings). GATE provides a full serialization of certain types of features such as collections, strings and numbers. It is possible to serialize only those collections containing strings or numbers. The rest of other features are serialized using their string representation and when read back, they will be all strings instead of being the original objects. Consequences of this might be observed when performing evaluations (see Chapter 10).

When GATE outputs an XML document it may do so in one of two ways:

- When the original document that was imported into GATE was an XML document, GATE can dump that document back into XML (possibly with additional markup added);
- For all document formats, GATE can dump its internal representation of the document into XML.

In the former case, the XML output will be close to the original document. In the latter case, the format is a GATE-specific one which can be read back by the system to recreate all the information that GATE held internally for the document.

In order to understand why there are two types of XML serialization, one needs to understand the structure of a GATE document. GATE allows a graph of annotations that refer to parts of the text. Those annotations are grouped under annotation sets. Because of this structure, sometimes it is impossible to save a document as XML using tags that surround the text referred to by the annotation, because tags crossover situations could appear (XML is essentially a tree-based model of information, whereas GATE uses graphs). Therefore, in order to preserve all annotations in a GATE document, a custom type of XML document was developed.

The problem of crossover tags appears with GATE's second option (the preserve format one), which is implemented at the cost of losing certain annotations. The way it is applied in GATE is that it tries to restore the original markup and where it is possible, to add in the same manner annotations produced by GATE.

How to Access and Use the Two Forms of XML Serialization

Save as XML Option This option is available in GATE Developer in the pop-up menu associated with each language resource (document or corpus). Saving a corpus as XML is done by calling ‘Save as XML’ on each document of the corpus. This option saves all the annotations of a document together their features (applying the restrictions previously discussed), using the GateDocument.dtd :

```
<!ELEMENT GateDocument (GateDocumentFeatures,
    TextWithNodes, (AnnotationSet+))>
<!ELEMENT GateDocumentFeatures (Feature+)>
<!ELEMENT Feature (Name, Value)>
<!ELEMENT Name (\#PCDATA)>
<!ELEMENT Value (\#PCDATA)>
<!ELEMENT TextWithNodes (\#PCDATA | Node)*>
<!ELEMENT AnnotationSet (Annotation*)>
<!ATTLIST AnnotationSet Name CDATA \#IMPLIED>
<!ELEMENT Annotation (Feature*)>
<!ATTLIST Annotation Type CDATA \#REQUIRED
    StartNode CDATA \#REQUIRED
    EndNode CDATA \#REQUIRED>
<!ELEMENT Node EMPTY>
<!ATTLIST Node id CDATA \#REQUIRED>
```

The document is saved under a name chosen by the user and it may have any extension. However, the recommended extension would be ‘xml’.

Using GATE Embedded, this option is available by calling `gate.Document`’s `toXml()` method. This method returns a string which is the XML representation of the document on which the method was called.

Note: It is recommended that the string representation to be saved on the file system using the UTF-8 encoding, as the first line of the string is : `<?xml version="1.0" encoding="UTF-8"?>`

Example of such a GATE format document:

```
<?xml version="1.0" encoding="UTF-8" ?>
<GateDocument>

<!-- The document’s features-->

<GateDocumentFeatures>
<Feature>
  <Name className="java.lang.String">MimeType</Name>
  <Value className="java.lang.String">text/plain</Value>
</Feature>
```

```

<Feature>
  <Name className="java.lang.String">gate.SourceURL</Name>
  <Value className="java.lang.String">file:/G:/tmp/example.txt</Value>
</Feature>
</GateDocumentFeatures>

```

```

<!-- The document content area with serialized nodes -->

```

```

<TextWithNodes>
<Node id="0"/>A TEENAGER <Node
id="11"/>yesterday<Node id="20"/> accused his parents of cruelty
by feeding him a daily diet of chips which sent his weight
ballooning to 22st at the age of 12<Node id="146"/>.<Node
id="147"/>
</TextWithNodes>

```

```

<!-- The default annotation set -->

```

```

<AnnotationSet>
<Annotation Type="Date" StartNode="11"
EndNode="20">
<Feature>
  <Name className="java.lang.String">rule2</Name>
  <Value className="java.lang.String">DateOnlyFinal</Value>
</Feature> <Feature>
  <Name className="java.lang.String">rule1</Name>
  <Value className="java.lang.String">GazDateWords</Value>
</Feature> <Feature>
  <Name className="java.lang.String">kind</Name>
  <Value className="java.lang.String">date</Value>
</Feature> </Annotation> <Annotation Type="Sentence" StartNode="0"
EndNode="147"> </Annotation> <Annotation Type="Split"
StartNode="146" EndNode="147"> <Feature>
  <Name className="java.lang.String">kind</Name>
  <Value className="java.lang.String">internal</Value>
</Feature> </Annotation> <Annotation Type="Lookup" StartNode="11"
EndNode="20"> <Feature>
  <Name className="java.lang.String">majorType</Name>
  <Value className="java.lang.String">date_key</Value>
</Feature> </Annotation>
</AnnotationSet>

```

```

<!-- Named annotation set -->

```

```

<AnnotationSet Name="Original markups" >
  <Annotation
Type="paragraph" StartNode="0" EndNode="147"> </Annotation>

```

```
</AnnotationSet>  
</GateDocument>
```

Note: One must know that all features that are not collections containing numbers or strings or that are not numbers or strings are discarded. With this option, GATE does not preserve those features it cannot restore back.

The Preserve Format Option This option is available in GATE Developer from the popup menu of the annotations table. If no annotation in this table is selected, then the option will restore the document's original markup. If certain annotations are selected, then the option will attempt to restore the original markup and insert all the selected ones. When an annotation violates the crossed over condition, that annotation is discarded and a message is issued.

This option makes it possible to generate an XML document with tags surrounding the annotation's referenced text and features saved as attributes. All features which are collections, strings or numbers are saved, and the others are discarded. However, when read back, only the attributes under the GATE namespace (see below) are reconstructed back differently to the others. That is because GATE does not store in the XML document the information about the features class and for collections the class of the items. So, when read back, all features will become strings, except those under the GATE namespace.

One will notice that all generated tags have an attribute called 'gateId' under the namespace 'http://www.gate.ac.uk'. The attribute is used when the document is read back in GATE, in order to restore the annotation's old ID. This feature is needed because it works in close cooperation with another attribute under the same namespace, called 'matches'. This attribute indicates annotations/tags that refer the same entity¹. They are under this namespace because GATE is sensitive to them and treats them differently to all other elements with their attributes which fall under the general reading algorithm described at the beginning of this section.

The 'gateId' under GATE namespace is used to create an annotation which has as ID the value indicated by this attribute. The 'matches' attribute is used to create an ArrayList in which the items will be Integers, representing the ID of annotations that the current one matches.

Example:

If the text being processed is as follows:

```
<Person gate:gateId="23">John</Person> and <Person  
gate:gateId="25" gate:matches="23;25;30">John Major</Person> are  
the same person.
```

¹It's not an XML entity but a information extraction named entity

What GATE does when it parses this text is it creates two annotations:

```
a1.type = "Person"
a1.ID = Integer(23)
a1.start = <the start offset of
John>
a1.end = <the end offset of John>
a1.featureMap = {}

a2.type = "Person"
a2.ID = Integer(25)
a2.start = <the start offset
of John Major>
a2.end = <the end offset of John Major>
a2.featureMap = {matches=[Integer(23); Integer(25); Integer(30)]}
```

Under GATE Embedded, this option is available by calling `gate.Document's toXml(Set aSetContainingAnnotations)` method. This method returns a string which is the XML representation of the document on which the method was called. If called with **null** as a parameter, then the method will attempt to restore only the original markup. If the parameter is a set that contains annotations, then each annotation is tested against the crossover restriction, and for those found to violate it, a warning will be issued and they will be discarded.

In the next subsections we will show how this option applies to the other formats supported by GATE.

5.5.3 HTML

Input

HTML documents are parsed by GATE using the NekoHTML parser. The documents are read and created in GATE the same way as the XML documents.

The extensions associate with the HTML reader are:

- htm
- html

The web server content type associate with html documents is: *text/html*.

The magic numbers test searches inside the document for the HTML(<html) signature. There are certain HTML documents that do not contain the HTML tag, so the magical numbers test might not hold.

There is a certain degree of customization for HTML documents in that GATE introduces new lines into the document's text content in order to obtain a readable form. The annotations will refer the pieces of text as described in the original document but there will be a few extra new line characters inserted.

After reading H1, H2, H3, H4, H5, H6, TR, CENTER, LI, BR and DIV tags, GATE will introduce a new line (NL) char into the text. After a TITLE tag it will introduce two NLs. With P tags, GATE will introduce one NL at the beginning of the paragraph and one at the end of the paragraph. All newly added NLs are not considered to be part of the text contained by the tag.

Output

The 'Save as XML' option works exactly the same for all GATE's documents so there is no particular observation to be made for the HTML formats.

When attempting to preserve the original markup formatting, GATE will generate the document in xhtml. The html document will look the same with any browser after processed by GATE but it will be in another syntax.

5.5.4 SGML

Input

The SGML support in GATE is fairly light as there is no freely available Java SGML parser. GATE uses a light converter attempting to transform the input SGML file into a well formed XML. Because it does not make use of a DTD, the conversion might not be always good. It is advisable to perform a SGML2XML conversion outside the system(using some other specialized tools) before using the SGML document inside GATE.

The extensions associate with the SGML reader are:

- sgm
- sgml

The web server content type associate with xml documents is : *text/sgml*.

There is no magic numbers test for SGML.

Output

When attempting to preserve the original markup formatting, GATE will generate the document as XML because the real input of a SGML document inside GATE is an XML one.

5.5.5 Plain text

Input

When reading a plain text document, GATE attempts to detect its paragraphs and add ‘paragraph’ annotations to the document’s ‘Original markups’ annotation set. It does that by detecting two consecutive NLs. The procedure works for both UNIX like or DOS like text files.

Example:

If the plain text read is as follows:

Paragraph 1. This text belongs to the first paragraph.

Paragraph 2. This text belongs to the second paragraph

then two ‘paragraph’ type annotation will be created in the ‘Original markups’ annotation set (referring the first and second paragraphs) with an empty feature map.

The extensions associate with the plain text reader are:

- txt
- text

The web server content type associate with plain text documents is: *text/plain*.

There is no magic numbers test for plain text.

Output

When attempting to preserve the original markup formatting, GATE will dump XML markup that surrounds the text refereed.

The procedure described above applies both for plain text and RTF documents.

5.5.6 RTF

Input

Accessing RTF documents is performed by using the Java's RTF editor kit. It only extracts the document's text content from the RTF document.

The extension associate with the RTF reader is '*rtf*'.

The web server content type associate with xml documents is : *text/rtf*.

The magic numbers test searches for `{\\rtf1`.

Output

Same as the plain tex output.

5.5.7 Email

Input

GATE is able to read email messages packed in one document (UNIX mailbox format). It detects multiple messages inside such documents and for each message it creates annotations for all the fields composing an e-mail, like date, from, to, subject, etc. The message's body is analyzed and a paragraph detection is performed (just like in the plain text case) . All annotation created have as type the name of the e-mail's fields and they are placed in the Original markup annotation set.

Example:

From someone@zzz.zzz.zzz Wed Sep 6 10:35:50 2000

Date: Wed, 6 Sep2000 10:35:49 +0100 (BST)

From: forename1 surname2 <someone1@yyy.yyy.xxx>

To: forename2 surname2 <someone2@ddd.dddd.dd.dd>

Subject: A subject

Message-ID: <Pine.SOL.3.91.1000906103251.26010A-100000@servername>

MIME-Version: 1.0

Content-Type: TEXT/PLAIN; charset=US-ASCII

This text belongs to the e-mail body....

This is a paragraph in the body of the e-mail

This is another paragraph.

GATE attempts to detect lines such as ‘*From someone@zzz.zzz.zzz Wed Sep 6 10:35:50 2000*’ in the e-mail text. Those lines separate e-mail messages contained in one file. After that, for each field in the e-mail message annotations are created as follows:

The annotation type will be the name of the field, the feature map will be empty and the annotation will span from the end of the field until the end of the line containing the e-mail field.

Example:

```
a1.type = "date" a1 spans between the two ^ ^. Date:^ Wed,
6Sep2000 10:35:49 +0100 (BST)^
```

```
a2.type = "from"; a2 spans between the two ^ ^. From:^ forename1
surname2 <someone1@yyy.yyy.xxx>^
```

The extensions associated with the email reader are:

- eml
- email
- mail

The web server content type associate with plain text documents is: *text/email*.

The magic numbers test searches for keywords like *Subject;*,etc.

Output

Same as plain text output.

5.6 XML Input/Output

Support for input from and output to XML is described in Section 5.5.2. In short:

- GATE will read any well-formed XML document (it does not attempt to validate XML documents). Markup will by default be converted into native GATE format.
- GATE will write back into XML in one of two ways:
 1. Preserving the original format and adding selected markup (for example to add the results of some language analysis process to the document).
 2. In GATE's own XML serialisation format, which encodes all the data in a GATE Document (as far as this is possible within a tree-structured paradigm – for 100% non-lossy data storage use GATE's RDBMS or binary serialisation facilities – see Section 4.5).

When using GATE Embedded, object representations of XML documents such as DOM or jDOM, or query and transformation languages such as X-Path or XSLT, may be used in parallel with GATE's own Document representation (`gate.Document`) without conflicts.

Chapter 6

ANNIE: a Nearly-New Information Extraction System

And so the time had passed predictably and soberly enough in work and routine chores, and the events of the previous night from first to last had faded; and only now that both their days' work was over, the child asleep and no further disturbance anticipated, did the shadowy figures from the masked ball, the melancholy stranger and the dominoes in red, revive; and those trivial encounters became magically and painfully interfused with the treacherous illusion of missed opportunities. Innocent yet ominous questions and vague ambiguous answers passed to and fro between them; and, as neither of them doubted the other's absolute candour, both felt the need for mild revenge. They exaggerated the extent to which their masked partners had attracted them, made fun of the jealous stirrings the other revealed, and lied dismissively about their own. Yet this light banter about the trivial adventures of the previous night led to more serious discussion of those hidden, scarcely admitted desires which are apt to raise dark and perilous storms even in the purest, most transparent soul; and they talked about those secret regions for which they felt hardly any longing, yet towards which the irrational wings of fate might one day drive them, if only in their dreams. For however much they might belong to one another heart and soul, they knew last night was not the first time they had been stirred by a whiff of freedom, danger and adventure.

Dream Story, Arthur Schnitzler, 1926 (pp. 4-5).

GATE was originally developed in the context of Information Extraction (IE) R&D, and IE systems in many languages and shapes and sizes have been created using GATE with the IE components that have been distributed with it (see [Maynard *et al.* 00] for descriptions of some of these projects).¹

¹The principal architects of the IE systems in GATE version 1 were Robert Gaizauskas and Kevin Humphreys. This work lives on in the LaSIE system. (A derivative of LaSIE was distributed with GATE

GATE is distributed with an IE system called ANNIE, A Nearly-New IE system (developed by Hamish Cunningham, Valentin Tablan, Diana Maynard, Kalina Bontcheva, Marin Dimitrov and others). ANNIE relies on finite state algorithms and the JAPE language (see Chapter 8).

ANNIE components form a pipeline which appears in figure 6.1. ANNIE components are

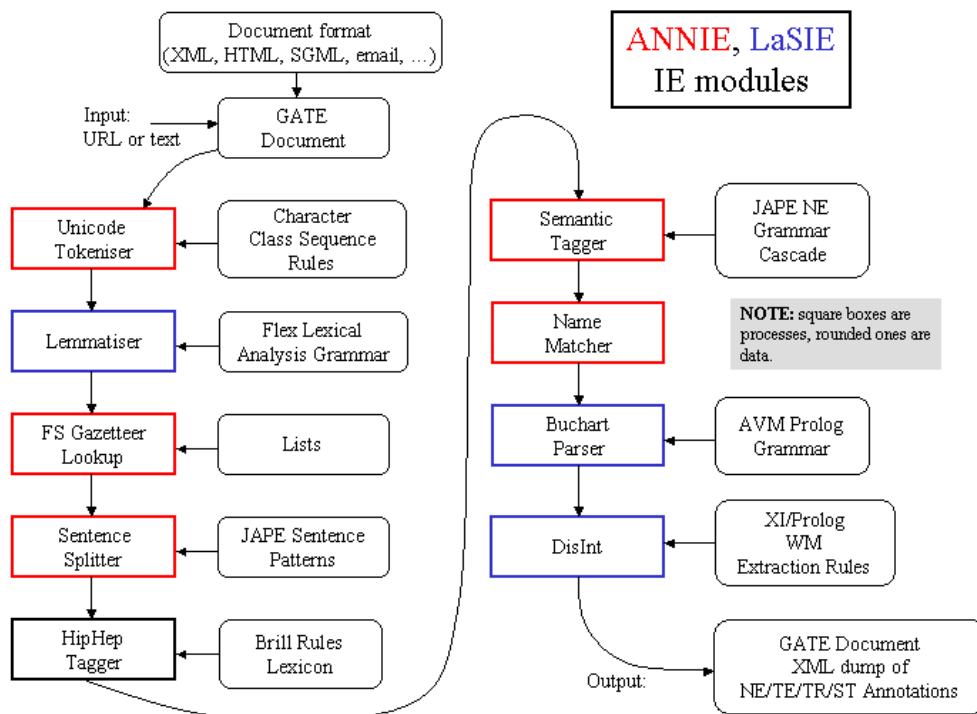


Figure 6.1: ANNIE and LaSIE

included with GATE (though the linguistic resources they rely on are generally more simple than the ones we use in-house). The rest of this chapter describes these components.

6.1 Document Reset

The document reset resource enables the document to be reset to its original state, by removing all the annotation sets and their contents, apart from the one containing the document format analysis (Original Markups). An optional parameter, `keepOriginalMarkupsAS`, allows users to decide whether to keep the Original Markups AS or not while resetting the document. The parameter `annotationTypes` can be used to specify a list of annotation types to remove from all the sets instead of the whole sets.

version 1 under the name VIE, a Vanilla IE system.)

Alternatively, if the parameter `setsToRemove` is not empty, the other parameters except `annotationTypes` are ignored and only the annotation sets specified in this list will be removed. If `annotationTypes` is also specified, only those annotation types in the specified sets are removed. In order to specify that you want to reset the default annotation set, just click the "Add" button without entering a name – this will add `<null>` which denotes the default annotation set. This resource is normally added to the beginning of an application, so that a document is reset before an application is rerun on that document.

6.2 Tokeniser

The tokeniser splits the text into very simple tokens such as numbers, punctuation and words of different types. For example, we distinguish between words in uppercase and lowercase, and between certain types of punctuation. The aim is to limit the work of the tokeniser to maximise efficiency, and enable greater flexibility by placing the burden on the grammar rules, which are more adaptable.

6.2.1 Tokeniser Rules

A rule has a left hand side (LHS) and a right hand side (RHS). The LHS is a regular expression which has to be matched on the input; the RHS describes the annotations to be added to the AnnotationSet. The LHS is separated from the RHS by '>'. The following operators can be used on the LHS:

```
| (or)
* (0 or more occurrences)
? (0 or 1 occurrences)
+ (1 or more occurrences)
```

The RHS uses ';' as a separator, and has the following format:

```
{LHS} > {Annotation type};{attribute1}={value1};...;{attribute
n}={value n}
```

Details about the primitive constructs available are given in the tokeniser file (`DefaultTokeniser.Rules`).

The following tokeniser rule is for a word beginning with a single capital letter:

```
'UPPERCASE_LETTER' 'LOWERCASE_LETTER'* >
  Token;orth=upperInitial;kind=word;
```

It states that the sequence must begin with an uppercase letter, followed by zero or more lowercase letters. This sequence will then be annotated as type 'Token'. The attribute 'orth' (orthography) has the value 'upperInitial'; the attribute 'kind' has the value 'word'.

6.2.2 Token Types

In the default set of rules, the following kinds of Token and SpaceToken are possible:

Word

A word is defined as any set of contiguous upper or lowercase letters, including a hyphen (but no other forms of punctuation). A word also has the attribute 'orth', for which four values are defined:

- upperInitial - initial letter is uppercase, rest are lowercase
- allCaps - all uppercase letters
- lowerCase - all lowercase letters
- mixedCaps - any mixture of upper and lowercase letters not included in the above categories

Number

A number is defined as any combination of consecutive digits. There are no subdivisions of numbers.

Symbol

Two types of symbol are defined: currency symbol (e.g. '\$', '£') and symbol (e.g. '&', '^'). These are represented by any number of consecutive currency or other symbols (respectively).

Punctuation

Three types of punctuation are defined: start_punctuation (e.g. '('), end_punctuation (e.g. ')'), and other_punctuation (e.g. ':'). Each punctuation symbol is a separate token.

SpaceToken

White spaces are divided into two types of SpaceToken - space and control - according to whether they are pure space characters or control characters. Any contiguous (and homogeneous) set of space or control characters is defined as a SpaceToken.

The above description applies to the default tokeniser. However, alternative tokenisers can be created if necessary. The choice of tokeniser is then determined at the time of text processing.

6.2.3 English Tokeniser

The English Tokeniser is a processing resource that comprises a normal tokeniser and a JAPE transducer (see Chapter 8). The transducer has the role of adapting the generic output of the tokeniser to the requirements of the English part-of-speech tagger. One such adaptation is the joining together in one token of constructs like “ ’30s”, “ ’Cause”, “ ’em”, “ ’N”, “ ’S”, “ ’s”, “ ’T”, “ ’d”, “ ’ll”, “ ’m”, “ ’re”, “ ’til”, “ ve”, etc. Another task of the JAPE transducer is to convert negative constructs like “don’t” from three tokens (“don”, “ ’ “ and “t”) into two tokens (“do” and “n’t”).

The English Tokeniser should always be used on English texts that need to be processed afterwards by the POS Tagger.

6.3 Gazetteer

The role of the gazetteer is to identify entity names in the text based on lists. The ANNIE gazetteer is described here, and also covered in Chapter 13 in Section 13.2.

The gazetteer lists used are plain text files, with one entry per line. Each list represents a set of names, such as names of cities, organisations, days of the week, etc.

Below is a small section of the list for units of currency:

```
Ecu
European Currency Units
FFr
Fr
German mark
German marks
New Taiwan dollar
New Taiwan dollars
NT dollar
```

NT dollars

An index file (`lists.def`) is used to access these lists; for each list, a major type is specified and, optionally, a minor type ². In the example below, the first column refers to the list name, the second column to the major type, and the third to the minor type. These lists are compiled into finite state machines. Any text tokens that are matched by these machines will be annotated with features specifying the major and minor types. Grammar rules then specify the types to be identified in particular circumstances. Each gazetteer list should reside in the same directory as the index file.

```
currency_prefix.lst:currency_unit:pre_amount
currency_unit.lst:currency_unit:post_amount
date.lst:date:specific
day.lst:date:day
```

So, for example, if a specific day needs to be identified, the minor type ‘day’ should be specified in the grammar, in order to match only information about specific days; if any kind of date needs to be identified, the major type ‘date’ should be specified, to enable tokens annotated with any information about dates to be identified. More information about this can be found in the following section.

In addition, the gazetteer allows arbitrary feature values to be associated with particular entries in a single list. ANNIE does not use this capability, but to enable it for your own gazetteers, set the optional `gazetteerFeatureSeparator` parameter to a single character (or an escape sequence such as `\t` or `\uNNNN`) when creating a gazetteer. In this mode, each line in a `.lst` file can have feature values specified, for example, with the following entry in the index file:

```
software_company.lst:company:software
```

the following `software_company.lst`:

```
Red Hat&stockSymbol=RHAT
Apple Computer&abbrev=Apple&stockSymbol=AAPL
Microsoft&abbrev=MS&stockSymbol=MSFT
```

and `gazetteerFeatureSeparator` set to `&`, the gazetteer will annotate Red Hat as a Lookup with features `majorType=company`, `minorType=software` and `stockSymbol=RHAT`. Note that you do not have to provide the same features for every line in the file, in particular it is possible to provide extra features for some lines in the list but not others.

²it is also possible to include a language in the same way, where lists for different languages are used, though ANNIE is only concerned with monolingual recognition

Here is a full list of the parameters used by the Default Gazetteer:

Init-time parameters

listsURL A URL pointing to the index file (usually lists.def) that contains the list of pattern lists.

encoding The character encoding to be used while reading the pattern lists.

gazetteerFeatureSeparator The character used to add arbitrary features to gazetteer entries. See above for an example.

caseSensitive Should the gazetteer be case sensitive during matching.

Run-time parameters

document The document to be processed.

annotationSetName The name for annotation set where the resulting Lookup annotations will be created.

wholeWordsOnly Should the gazetteer only match whole words? If set to true, a string segment in the input document will only be matched if it is bordered by characters that are not letters, non spacing marks, or combining spacing marks (as identified by the Unicode standard).

longestMatchOnly Should the gazetteer only match the longest possible string starting from any position. This parameter is only relevant when the list of lookups contains proper prefixes of other entries (e.g when both ‘Dell’ and ‘Dell Europe’ are in the lists). The default behaviour (when this parameter is set to **true**) is to only match the longest entry, ‘Dell Europe’ in this example. This is the default GATE gazetteer behaviour since version 2.0. Setting this parameter to **false** will cause the gazetteer to match all possible prefixes.

6.4 Sentence Splitter

The **sentence splitter** is a cascade of finite-state transducers which segments the text into sentences. This module is required for the tagger. The splitter uses a gazetteer list of abbreviations to help distinguish sentence-marking full stops from other kinds.

Each sentence is annotated with the type `Sentence`. Each sentence break (such as a full stop) is also given a ‘Split’ annotation. This has several possible types: ‘.’, ‘punctuation’, ‘CR’ (a line break) or ‘multi’ (a series of punctuation marks such as ‘?!?’).

The sentence splitter is domain and application-independent.

There is an alternative ruleset for the Sentence Splitter which considers newlines and carriage returns differently. In general this version should be used when a new line on the page indicates a new sentence). To use this alternative version, simply load the `main-single-nl.jape` from the default location instead of `main.jape` (the default file) when asked to select the location of the grammar file to be used.

6.5 RegEx Sentence Splitter

The RegEx sentence splitter is an alternative to the standard ANNIE Sentence Splitter. Its main aim is to address some performance issues identified in the JAPE-based splitter, mainly do to with improving the execution time and robustness, especially when faced with irregular input.

As its name suggests, the RegEx splitter is based on regular expressions, using the default Java implementation.

The new splitter is configured by three files containing (Java style, see <http://java.sun.com/j2se/1.5.0/docs/api/java/util/regex/Pattern.html>) regular expressions, one regex per line. The three different files encode patterns for:

internal splits sentence splits that are part of the sentence, such as sentence ending punctuation;

external splits sentence splits that are NOT part of the sentence, such as 2 consecutive new lines;

non splits text fragments that might be seen as splits but they should be ignored (such as full stops occurring inside abbreviations).

The new splitter comes with an initial set of patterns that try to emulate the behaviour of the original splitter (apart from the situations where the original one was obviously wrong, like not allowing sentences to start with a number).

Here is a full list of the parameters used by the RegEx Sentence Splitter:

Init-time parameters

encoding The character encoding to be used while reading the pattern lists.

externalSplitListURL URL for the file containing the list of external split patterns;

internalSplitListURL URL for the file containing the list of internal split patterns;

nonSplitListURL URL for the file containing the list of non split patterns;

Run-time parameters

document The document to be processed.

outputASName The name for annotation set where the resulting **Split** and **Sentence** annotations will be created.

6.6 Part of Speech Tagger

The **tagger** [Hepple 00] is a modified version of the Brill tagger, which produces a part-of-speech tag as an annotation on each word or symbol. The list of tags used is given in Appendix G. The tagger uses a default lexicon and ruleset (the result of training on a large corpus taken from the Wall Street Journal). Both of these can be modified manually if necessary. Two additional lexicons exist - one for texts in all uppercase (`lexicon_cap`), and one for texts in all lowercase (`lexicon_lower`). To use these, the default lexicon should be replaced with the appropriate lexicon at load time. The default ruleset should still be used in this case.

The ANNIE Part-of-Speech tagger requires the following parameters.

- `encoding` - encoding to be used for reading rules and lexicons (init-time)
- `lexiconURL` - The URL for the lexicon file (init-time)
- `rulesURL` - The URL for the ruleset file (init-time)
- `document` - The document to be processed (run-time)
- `inputASName` - The name of the annotation set used for input (run-time)
- `outputASName` - The name of the annotation set used for output (run-time). This is an optional parameter. If user does not provide any value, new annotations are created under the default annotation set.
- `baseTokenAnnotationType` - The name of the annotation type that refers to Tokens in a document (run-time, default = Token)
- `baseSentenceAnnotationType` - The name of the annotation type that refers to Sentences in a document (run-time, default = Sentences)
- `outputAnnotationType` - POS tags are added as category features on the annotations of type 'outputAnnotationType' (run-time, default = Token)

- `failOnMissingInputAnnotations` - if set to false, the PR will not fail with an `ExecutionException` if no input Annotations are found and instead only log a single warning message per session and a debug message per document that has no input annotations (run-time, default = true).

If - (`inputASName == outputASName`) AND (`outputAnnotationType == baseTokenAnnotationType`)

then - New features are added on existing annotations of type ‘`baseTokenAnnotationType`’.

otherwise - Tagger searches for the annotation of type ‘`outputAnnotationType`’ under the ‘`outputASName`’ annotation set that has the same offsets as that of the annotation with type ‘`baseTokenAnnotationType`’. If it succeeds, it adds new feature on a found annotation, and otherwise, it creates a new annotation of type ‘`outputAnnotationType`’ under the ‘`outputASName`’ annotation set.

6.7 Semantic Tagger

ANNIE’s semantic tagger is based on the JAPE language – see Chapter 8. It contains rules which act on annotations assigned in earlier phases, in order to produce outputs of annotated entities.

6.8 Orthographic Coreference (OrthoMatcher)

(Note: this component was previously known as a ‘`NameMatcher`’.)

The Orthomatcher module adds identity relations between named entities found by the semantic tagger, in order to perform coreference. It does not find new named entities as such, but it may assign a type to an unclassified proper name, using the type of a matching name.

The matching rules are only invoked if the names being compared are both of the same type, i.e. both already tagged as (say) organisations, or if one of them is classified as ‘unknown’. This prevents a previously classified name from being recategorised.

6.8.1 GATE Interface

Input – entity annotations, with an `id` attribute.

Output – matches attributes added to the existing entity annotations.

6.8.2 Resources

A lookup table of aliases is used to record non-matching strings which represent the same entity, e.g. ‘IBM’ and ‘Big Blue’, ‘Coca-Cola’ and ‘Coke’. There is also a table of spurious matches, i.e. matching strings which do not represent the same entity, e.g. ‘BT Wireless’ and ‘BT Cellnet’ (which are two different organizations). The list of tables to be used is a load time parameter of the orthomatcher: a default list is set but can be changed as necessary.

6.8.3 Processing

The wrapper builds an array of the strings, types and IDs of all `name` annotations, which is then passed to a string comparison function for pairwise comparisons of all entries.

6.9 Pronominal Coreference

The pronominal coreference module performs anaphora resolution using the JAPE grammar formalism. Note that this module is not automatically loaded with the other ANNIE modules, but can be loaded separately as a Processing Resource. The main module consists of three submodules:

- quoted text module
- pleonastic it module
- pronominal resolution module

The first two modules are helper submodules for the pronominal one, because they do not perform anything related to coreference resolution except the location of quoted fragments and pleonastic it occurrences in text. They generate temporary annotations which are used by the pronominal submodule (such temporary annotations are removed later).

The main coreference module can operate successfully only if all ANNIE modules were already executed. The module depends on the following annotations created from the respective ANNIE modules:

- Token (English Tokenizer)
- Sentence (Sentence Splitter)
- Split (Sentence Splitter)
- Location (NE Transducer, OrthoMatcher)

- Person (NE Transducer, OrthoMatcher)
- Organization (NE Transducer, OrthoMatcher)

For each pronoun (anaphor) the coreference module generates an annotation of type ‘Coreference’ containing two features:

- antecedent offset - this is the offset of the starting node for the annotation (entity) which is proposed as the antecedent, or null if no antecedent can be proposed.
- matches - this is a list of annotation IDs that comprise the coreference chain comprising this anaphor/antecedent pair.

6.9.1 Quoted Speech Submodule

The quoted speech submodule identifies quoted fragments in the text being analysed. The identified fragments are used by the pronominal coreference submodule for the proper resolution of pronouns such as I, me, my, etc. which appear in quoted speech fragments. The module produces ‘Quoted Text’ annotations.

The submodule itself is a JAPE transducer which loads a JAPE grammar and builds an FSM over it. The FSM is intended to match the quoted fragments and generate appropriate annotations that will be used later by the pronominal module.

The JAPE grammar consists of only four rules, which create temporary annotations for all punctuation marks that may enclose quoted speech, such as ”, ’, ‘, etc. These rules then try to identify fragments enclosed by such punctuation. Finally all temporary annotations generated during the processing, except the ones of type ‘Quoted Text’, are removed (because no other module will need them later).

6.9.2 Pleonastic It Submodule

The pleonastic it submodule matches pleonastic occurrences of ‘it’. Similar to the quoted speech submodule, it is a JAPE transducer operating with a grammar containing patterns that match the most commonly observed pleonastic it constructs.

6.9.3 Pronominal Resolution Submodule

The main functionality of the coreference resolution module is in the pronominal resolution submodule. This uses the result from the execution of the quoted speech and pleonastic it submodules. The module works according to the following algorithm:

- Preprocess the current document. This step locates the annotations that the submodule need (such as Sentence, Token, Person, etc.) and prepares the appropriate data structures for them.
- For each pronoun do the following:
 - inspect the proper appropriate context for all candidate antecedents for this kind of pronoun;
 - choose the best antecedent (if any);
- Create the coreference chains from the individual anaphor/antecedent pairs and the coreference information supplied by the OrthoMatcher (this step is performed from the main coreference module).

6.9.4 Detailed Description of the Algorithm

Full details of the pronominal coreference algorithm are as follows.

Preprocessing

The preprocessing task includes the following subtasks:

- Identifying the sentences in the document being processed. The sentences are identified with the help of the Sentence annotations generated from the Sentence Splitter. For each sentence a data structure is prepared that contains three lists. The lists contain the annotations for the person/organization/location named entities appearing in the sentence. The named entities in the sentence are identified with the help of the Person, Location and Organization annotations that are already generated from the Named Entity Transducer and the OrthoMatcher.
- The gender of each person in the sentence is identified and stored in a global data structure. It is possible that the gender information is missing for some entities - for example if only the person family name is observed then the Named Entity transducer will be unable to deduce the gender. In such cases the list with the matching entities generated by the OrthoMatcher is inspected and if some of the orthographic matches contains gender information it is assigned to the entity being processed.
- The identified pleonastic it occurrences are stored in a separate list. The ‘Pleonastic It’ annotations generated from the pleonastic submodule are used for the task.
- For each quoted text fragment, identified by the quoted text submodule, a special structure is created that contains the persons and the 3rd person singular pronouns such as ‘he’ and ‘she’ that appear in the sentence containing the quoted text, but not in the quoted text span (i.e. the ones preceding and succeeding the quote).

Pronoun Resolution

This task includes the following subtasks:

Retrieving all the pronouns in the document. Pronouns are represented as annotations of type ‘Token’ with feature ‘category’ having value ‘PRP\$’ or ‘PRP’. The former classifies possessive adjectives such as *my*, *your*, etc. and the latter classifies personal, reflexive etc. pronouns. The two types of pronouns are combined in one list and sorted according to their offset in the text.

For each pronoun in the list the following actions are performed:

- If the pronoun is ‘it’, then the module performs a check to determine if this is a pleonastic occurrence. If it is, then no further attempt for resolution is made.
- The proper context is determined. The context size is expressed in the number of sentences it will contain. The context always includes the current sentence (the one containing the pronoun), the preceding sentence and zero or more preceding sentences.
- Depending on the type of pronoun, a set of candidate antecedents is proposed. The candidate set includes the named entities that are compatible with this pronoun. For example if the current pronoun is *she* then only the Person annotations with ‘gender’ feature equal to ‘female’ or ‘unknown’ will be considered as candidates.
- From all candidates, one is chosen according to evaluation criteria specific for the pronoun.

Coreference Chain Generation

This step is actually performed by the main module. After executing each of the submodules on the current document, the coreference module follows the steps:

- Retrieves the anaphor/antecedent pairs generated from them.
- For each pair, the orthographic matches (if any) of the antecedent entity is retrieved and then extended with the anaphor of the pair (i.e. the pronoun). The result is the coreference chain for the entity. The coreference chain contains the IDs of the annotations (entities) that co-refer.
- A new Coreference annotation is created for each chain. The annotation contains a single feature ‘matches’ whose value is the coreference chain (the list with IDs). The annotations are exported in a pre-specified annotation set.

The resolution of she, her, her\$, he, him, his, herself and himself are similar because an analysis of a corpus showed that these pronouns are related to their antecedents in a similar manner. The characteristics of the resolution process are:

- Context inspected is not very big - cases where the antecedent is found more than 3 sentences back from the anaphor are rare.
- Recency factor is heavily used - the candidate antecedents that appear closer to the anaphor in the text are scored better.
- Anaphora have higher priority than cataphora. If there is an anaphoric candidate and a cataphoric one, then the anaphoric one is preferred, even if the recency factor scores the cataphoric candidate better.

The resolution process performs the following steps:

- Inspect the context of the anaphor for candidate antecedents. Every Person annotation is considered to be a candidate. Cases where she/her refers to inanimate entity (ship for example) are not handled.
- For each candidate perform a gender compatibility check - only candidates having 'gender' feature equal to 'unknown' or compatible with the pronoun are considered for further evaluation.
- Evaluate each candidate with the best candidate so far. If the two candidates are anaphoric for the pronoun then choose the one that appears closer. The same holds for the case where the two candidates are cataphoric relative to the pronoun. If one is anaphoric and the other is cataphoric then choose the former, even if the latter appears closer to the pronoun.

Resolution of 'it', 'its', 'itself'

This set of pronouns also shares many common characteristics. The resolution process contains certain differences with the one for the previous set of pronouns. Successful resolution for it, its, itself is more difficult because of the following factors:

- There is no gender compatibility restriction. In the case in which there are several candidates in the context, the gender compatibility restriction is very useful for rejecting some of the candidates. When no such restriction exists, and with the lack of any syntactic or ontological information about the entities in the context, the recency factor plays the major role in choosing the best antecedent.
- The number of nominal antecedents (i.e. entities that are not referred by name) is much higher compared to the number of such antecedents for she, he, etc. In this case trying to find an antecedent only amongst named entities degrades the precision a lot.

Resolution of ‘I’, ‘me’, ‘my’, ‘myself’

Resolution of these pronouns is dependent on the work of the quoted speech submodule. One important difference from the resolution process of other pronouns is that the context is not measured in sentences but depends solely on the quote span. Another difference is that the context is not contiguous - the quoted fragment itself is excluded from the context, because it is unlikely that an antecedent for I, me, etc. appears there. The context itself consists of:

- the part of the sentence where the quoted fragment originates, that is not contained in the quote - i.e. the text prior to the quote;
- the part of the sentence where the quoted fragment ends, that is not contained in the quote - i.e. the text following the quote;
- the part of the sentence preceding the sentence where the quote originates, which is not included in other quote.

It is worth noting that contrary to other pronouns, the antecedent for I, me, my and myself is most often cataphoric or if anaphoric it is not in the same sentence with the quoted fragment.

The resolution algorithm consists of the following steps:

- Locate the quoted fragment description that contains the pronoun. If the pronoun is not contained in any fragment then return without proposing an antecedent.
- Inspect the context for the quoted fragment (as defined above) for candidate antecedents. Candidates are considered annotations of type Pronoun or annotations of type Token with features category = ‘PRP’, string = ‘she’ or category = ‘PRP’, string = ‘he’.
- Try to locate a candidate in the text succeeding the quoted fragment (first pattern). If more than one candidate is present, choose the closest to the end of the quote. If a candidate is found then propose it as antecedent and exit.
- Try to locate a candidate in the text preceding the quoted fragment (third pattern). Choose the closest one to the beginning of the quote. If found then set as antecedent and exit.
- Try to locate antecedents in the unquoted part of the sentence preceding the sentence where the quote starts (second pattern). Give preference to the one closest to the end of the quote (if any) in the preceding sentence or closest to the sentence beginning.

6.10 A Walk-Through Example

Let us take an example of a 3-stage procedure using the tokeniser, gazetteer and named-entity grammar. Suppose we wish to recognise the phrase ‘800,000 US dollars’ as an entity of type ‘Number’, with the feature ‘money’.

First of all, we give an example of a grammar rule (and corresponding macros) for money, which would recognise this type of pattern.

```
Macro: MILLION_BILLION
({Token.string == "m"}|
{Token.string == "million"}|
{Token.string == "b"}|
{Token.string == "billion"}
)

Macro: AMOUNT_NUMBER
({Token.kind == number}
(({{Token.string == ","}|
  {Token.string == "."})
{Token.kind == number})*
({SpaceToken.kind == space})?
(MILLION_BILLION)?)
)

Rule: Money1
// e.g. 30 pounds
(
  (AMOUNT_NUMBER)
  (SpaceToken.kind == space)?
  ({Lookup.majorType == currency_unit})
)
:money -->
:money.Number = {kind = "money", rule = "Money1"}
```

6.10.1 Step 1 - Tokenisation

The tokeniser separates this phrase into the following tokens. In general, a word is comprised of any number of letters of either case, including a hyphen, but nothing else; a number is composed of any sequence of digits; punctuation is recognised individually (each character is a separate token), and any number of consecutive spaces and/or control characters are recognised as a single spacetoken.

```
Token, string = '800', kind = number, length = 3
Token, string = ',', kind = punctuation, length = 1
Token, string = '000', kind = number, length = 3
SpaceToken, string = ' ', kind = space, length = 1
Token, string = 'US', kind = word, length = 2, orth = allCaps
SpaceToken, string = ' ', kind = space, length = 1
Token, string = 'dollars', kind = word, length = 7, orth = lowercase
```

6.10.2 Step 2 - List Lookup

The gazetteer lists are then searched to find all occurrences of matching words in the text. It finds the following match for the string 'US dollars':

```
Lookup, minorType = post_amount, majorType = currency_unit
```

6.10.3 Step 3 - Grammar Rules

The grammar rule for money is then invoked. The macro MILLION_BILLION recognises any of the strings 'm', 'million', 'b', 'billion'. Since none of these exist in the text, it passes onto the next macro. The AMOUNT_NUMBER macro recognises a number, optionally followed by any number of sequences of the form 'dot or comma plus number', followed by an optional space and an optional MILLION_BILLION. In this case, '800,000' will be recognised. Finally, the rule Money1 is invoked. This recognises the string identified by the AMOUNT_NUMBER macro, followed by an optional space, followed by a unit of currency (as determined by the gazetteer). In this case, 'US dollars' has been identified as a currency unit, so the rule Money1 recognises the entire string '800,000 US dollars'. Following the rule, it will be annotated as a Number entity of type Money:

```
Number, kind = money, rule = Money1
```

Part II

GATE for Advanced Users

Chapter 7

GATE Embedded

7.1 Quick Start with GATE Embedded

Embedding GATE-based language processing in other applications using GATE Embedded (the GATE API) is straightforward:

- add `$GATE_HOME/bin/gate.jar` and the JAR files in `$GATE_HOME/lib` to the Java CLASSPATH (`$GATE_HOME` is the GATE root directory)
- tell Java that the GATE Unicode Kit is an extension:
`-Djava.ext.dirs=$GATE_HOME/lib/ext`
N.B. This is only necessary for GUI applications that need to support Unicode text input; other applications such as command line or web applications don't generally need GUK.
- initialise GATE with `gate.Gate.init()`;
- program to the framework API.

For example, this code will create the ANNIE extraction system:

```
1  // initialise the GATE library
2  Gate.init();
3
4  // load ANNIE as an application from a gapp file
5  SerialAnalyserController controller = (SerialAnalyserController)
6      PersistenceManager.loadObjectFromFile(new File(new File(
7          Gate.getPluginsHome(), ANNIEConstants.PLUGIN_DIR),
8          ANNIEConstants.DEFAULT_FILE));
```

If you want to use resources from any plugins, you need to load the plugins before calling `createResource`:

```

1  Gate.init();
2
3  // need Tools plugin for the Morphological analyser
4  Gate.getCreoleRegister().registerDirectories(
5      new File(Gate.getPluginsHome(), "Tools").toURL()
6  );
7
8  ...
9
10 ProcessingResource morpher = (ProcessingResource)
11     Factory.createResource("gate.creole.morph.Morph");

```

Instead of creating your processing resources individually using the `Factory`, you can create your application in GATE Developer, save it using the ‘save application state’ option (see Section 3.8.3), and then load the saved state from your code. This will automatically reload any plugins that were loaded when the state was saved, you do not need to load them manually.

```

1  Gate.init();
2
3  CorpusController controller = (CorpusController)
4      PersistenceManager.loadObjectFromFile(new File("savedState.xgapp"));
5
6  // loadObjectFromUrl is also available

```

There are many examples of using GATE Embedded available at <http://gate.ac.uk/wiki/code-repository/>.

7.2 Resource Management in GATE Embedded

As outlined earlier, GATE defines three different types of resources:

Language Resources : (LRs) entities that hold linguistic data.

Processing Resources : (PRs) entities that process data.

Visual Resources : (VRs) components used for building graphical interfaces.

These resources are collectively named **CREOLE**¹ resources.

All CREOLE resources have some associated meta-data in the form of an entry in a special XML file named `creole.xml`. The most important role of that meta-data is to specify the set

¹CREOLE stands for Collection of REusable Objects for Language Engineering

of parameters that a resource understands, which of them are required and which not, if they have default values and what those are. The valid parameters for a resource are described in the resource's section of its `creole.xml` file or in Java annotations on the resource class – see Section 4.7.

All resource types have creation-time parameters that are used during the initialisation phase. Processing Resources also have run-time parameters that get used during execution (see Section 7.5 for more details).

Controllers are used to define GATE applications and have the role of controlling the execution flow (see Section 7.6 for more details).

This section describes how to create and delete CREOLE resources as objects in a running Java virtual machine. This process involves using GATE's Factory class², and, in the case of LRs, may also involve using a DataStore.

CREOLE resources are Java Beans; creation of a resource object involves using a default constructor, then setting parameters on the bean, then calling an `init()` method. The Factory takes care of all this, makes sure that the GATE Developer GUI is told about what is happening (when GUI components exist at runtime), and also takes care of restoring LRs from DataStores. **A programmer using GATE Embedded should never call the constructor of a resource: always use the Factory!**

Creating a resource involves providing the following information:

- **fully qualified class name** for the resource. This is the only **required** value. For all the rest, defaults will be used if actual values are not provided.
- values for the **creation time parameters**.[†]
- initial values for **resource features**.[†] For an explanation on features see Section 7.4.2.
- a **name** for the new resource;

[†] Parameters and features need to be provided in the form of a GATE Feature Map which is essentially a java Map (`java.util.Map`) implementation, see Section 7.4.2 for more details on Feature Maps.

Creating a resource via the Factory involves passing values for any create-time parameters that require setting to the Factory's `createResource` method. If no parameters are passed, the defaults are used. So, for example, the following code creates a default ANNIE part-of-speech tagger:

```
1 Gate.getCreoleRegister().registerDirectories(new File(
2   Gate.getPluginsHome(), ANNIEConstants.PLUGIN_DIR).toURI().toURL());
3 FeatureMap params = Factory.newFeatureMap(); //empty map: default params
```

²Fully qualified name: `gate.Factory`

```

4 ProcessingResource tagger = (ProcessingResource)
5   Factory.createResource("gate.creole.POSTagger", params);

```

Note that if the resource created here had any parameters that were both mandatory and had no default value, the `createResource` call would throw an exception. In this case, all the information needed to create a tagger is available in default values given in the tagger's XML definition (in `plugins/ANNIE/creole.xml`):

```

<RESOURCE>
  <NAME>ANNIE POS Tagger</NAME>
  <COMMENT>Mark Hepple's Brill-style POS tagger</COMMENT>
  <CLASS>gate.creole.POSTagger</CLASS>
  <PARAMETER NAME="document"
    COMMENT="The document to be processed"
    RUNTIME="true">gate.Document</PARAMETER>
  . . . .
  <PARAMETER NAME="rulesURL" DEFAULT="resources/heptag/ruleset"
    COMMENT="The URL for the ruleset file"
    OPTIONAL="true">java.net.URL</PARAMETER>
</RESOURCE>

```

Here the two parameters shown are either 'runtime' parameters, which are set before a PR is executed, or have a default value (in this case the default rules file is distributed with GATE itself).

When creating a Document, however, the URL of the source for the document must be provided³. For example:

```

1 URL u = new URL("http://gate.ac.uk/hamish/");
2 FeatureMap params = Factory.newFeatureMap();
3 params.put("sourceUrl", u);
4 Document doc = (Document)
5   Factory.createResource("gate.corpora.DocumentImpl", params);

```

Note that the document created here is transient: when you quit the JVM the document will no longer exist. If you want the document to be persistent, you need to store it in a `DataStore` (see Section 7.4.5).

Apart from `createResource()` methods with different signatures, `Factory` also provides some shortcuts for common operations, listed in table 7.1.

GATE maintains various data structures that allow the retrieval of loaded resources. When a resource is no longer required, it needs to be removed from those structures in order to remove all references to it, thus making it a candidate for garbage collection. This is achieved using

³Alternatively a string giving the document source may be provided.

Method	Purpose
<code>newFeatureMap()</code>	Creates a new Feature Map (as used in the example above).
<code>newDocument(String content)</code>	Creates a new GATE Document starting from a String value that will be used to generate the document content.
<code>newDocument(URL sourceUrl)</code>	Creates a new GATE Document using the text pointed by an URL to generate the document content.
<code>newDocument(URL sourceUrl, String encoding)</code>	Same as above but allows the specification of an encoding to be used while downloading the document content.
<code>newCorpus(String name)</code>	creates a new GATE Corpus with a specified name.

Table 7.1: Factory Operations

the `deleteResource(Resource res)` method on `Factory`.

Simply removing all references to a resource from the user code will **NOT** be enough to make the resource collect-able. Not calling `Factory.deleteResource()` **will** lead to memory leaks!

7.3 Using CREOLE Plugins

As shown in the examples above, in order to use a CREOLE resource the relevant CREOLE plugin must be loaded. Processing Resources, Visual Resources and Language Resources other than Document, Corpus and DataStore all require that the appropriate plugin is first loaded. When using Document, Corpus or DataStore, you do not need to first load a plugin. The following API calls listed in table 7.2 are relevant to working with CREOLE plugins.

7.4 Language Resources

This section describes the implementation of documents and corpora in GATE.

Class gate.Gate	
Method	Purpose
public static void addKnownPlugin (URL pluginURL)	adds the plugin to the list of known plugins.
public static void removeKnownPlugin (URL pluginURL)	tells the system to ‘forget’ about one previously known directory. If the specified directory was loaded, it will be unloaded as well - i.e. all the metadata relating to resources defined by this directory will be removed from memory.
public static void addAutoloadPlugin (URL pluginUrl)	adds a new directory to the list of plugins that are loaded automatically at start-up.
public static void removeAutoloadPlugin (URL pluginURL)	tells the system to remove a plugin URL from the list of plugins that are loaded automatically at system start-up. This will be reflected in the user’s configuration data file.
Class gate.CreoleRegister	
public void registerDirectories (URL directoryUrl)	loads a new CREOLE directory. The new plugin is added to the list of known plugins if not already there.
public void removeDirectory (URL directory)	unloads a loaded CREOLE plugin.

Table 7.2: Calls Relevant to CREOLE Plugins

7.4.1 GATE Documents

Documents are modelled as content plus annotations (see Section 7.4.4) plus features (see Section 7.4.2).

The content of a document can be any implementation of the `gate.DocumentContent` interface; the features are <attribute, value> pairs stored a Feature Map. Attributes are String values while the values can be any Java object.

The annotations are grouped in sets (see section 7.4.3). A document has a default (anonymous) annotations set and any number of named annotations sets.

Documents are defined by the `gate.Document` interface and there is also a provided implementation:

`gate.corpora.DocumentImpl` : transient document. Can be stored persistently through Java serialisation.

Main Document functions are presented in table 7.3.

7.4.2 Feature Maps

All CREOLE resources as well as the *Controllers* and the annotations can have attached meta-data in the form of *Feature Maps*.

A Feature Map is a Java Map (i.e. it implements the `java.util.Map` interface) and holds <attribute-name, attribute-value> pairs. The attribute names are Strings while the values can be any Java Objects.

The use of non-*Serializable* objects as values is strongly discouraged.

Feature Maps are created using the `gate.Factory.newFeatureMap()` method.

The actual implementation for FeatureMaps is provided by the `gate.util.SimpleFeatureMapImpl` class.

Objects that have features in GATE implement the `gate.util.FeatureBearer` interface which has only the two accessor methods for the object features: `FeatureMap getFeatures()` and `void setFeatures(FeatureMap features)`.

Content Manipulation	
Method	Purpose
DocumentContent getContent()	Gets the Document content.
void edit (Long start, Long end, DocumentContent replacement)	Modifies the Document content.
void setContent (DocumentContent newContent)	Replaces the entire content.
Annotations Manipulation	
Method	Purpose
public AnnotationSet getAnnotations()	Returns the default annotation set.
public AnnotationSet getAnnotations (String name)	Returns a <i>named</i> annotation set.
public Map getNamedAnnotationSets()	Returns <i>all</i> the named annotation sets.
void removeAnnotationSet (String name)	Removes a named annotation set.
Input Output	
String toXml()	Serialises the Document in XML format.
String toXml (Set aSourceAnnotationSet, boolean includeFeatures)	Generates XML from a set of annotations only, trying to preserve the original format of the file used to create the document.

Table 7.3: gate.Document methods.

Getting a particular feature from an object

```

1 Object obj;
2 String featureName = "length";
3 if(obj instanceof FeatureBearer){
4     FeatureMap features = ((FeatureBearer)obj).getFeatures();
5     Object value = (features == null) ? null :
6                     features.get(featureName);
7 }

```

7.4.3 Annotation Sets

A GATE document can have one or more annotation layers — an anonymous one, (also called *default*), and as many *named* ones as necessary.

An annotation layer is organised as a *Directed Acyclic Graph (DAG)* on which the nodes are particular locations — *anchors*— in the document content and the arcs are made out of annotations reaching from the location indicated by the start node to the one pointed by the end node (see Figure 7.1 for an illustration). Because of the *graph* metaphor, the annotation layers are also called *annotation graphs*. In terms of Java objects, the annotation layers are represented using the *Set* paradigm as defined by the collections library and they are hence named *annotation sets*. The terms of annotation layer, *graph* and *set* are interchangeable and refer to the same concept when used in this book.




Figure 7.1: The Annotation Graph model.

An annotation set holds a number of annotations and maintains a series of indices in order to provide fast access to the contained annotations.

The GATE Annotation Sets are defined by the `gate.AnnotationSet` interface and there is a default implementation provided:

```

gate.annotation.AnnotationSetImpl annotation set implementation used by transient
documents.

```

The annotation sets are created by the document as required. The first time a particular annotation set is requested from a document it will be transparently created if it doesn't exist.

Annotations Manipulation	
Method	Purpose
Integer add (Long start, Long end, String type, FeatureMap features)	Creates a new annotation between two offsets, adds it to this set and returns its id.
Integer add (Node start, Node end, String type, FeatureMap features)	Creates a new annotation between two nodes, adds it to this set and returns its id.
boolean remove (Object o)	Removes an annotation from this set.
Nodes	
Method	Purpose
Node firstNode ()	Gets the node with the smallest offset.
Node lastNode ()	Gets the node with the largest offset.
Node nextNode (Node node)	Get the first node that is relevant for this annotation set and which has the offset larger than the one of the node provided.
Set implementation	
Iterator iterator ()	
int size ()	

Table 7.4: `gate.AnnotationSet` methods (general purpose).

Searching	
<code>AnnotationSet get(Long offset)</code>	Select annotations by offset. This returns the set of annotations whose start node is the least such that it is less than or equal to offset. If a positional index doesn't exist it is created. If there are no nodes at or beyond the offset parameter then it will return null.
<code>AnnotationSet get(Long startOffset, Long endOffset)</code>	Select annotations by offset. This returns the set of annotations that overlap totally or partially with the interval defined by the two provided offsets. The result will include all the annotations that either: <ul style="list-style-type: none"> • start before the start offset and end strictly after it • start at a position between the start and the end offsets
<code>AnnotationSet get(String type)</code>	Returns all annotations of the specified type.
<code>AnnotationSet get(Set types)</code>	Returns all annotations of the specified types.
<code>AnnotationSet get(String type, FeatureMap constraints)</code>	Selects annotations by type and features.
<code>Set getAllTypes()</code>	Gets a set of <code>java.lang.String</code> objects representing all the annotation types present in this annotation set.

Table 7.5: `gate.AnnotationSet` methods (searching).

Tables 7.4 and 7.5 list the most used Annotation Set functions.

Iterating from left to right over all annotations of a given type

```

1 AnnotationSet annSet = ...;
2 String type = "Person";
3 //Get all person annotations
4 AnnotationSet persSet = annSet.get(type);
5 //Sort the annotations
6 List persList = new ArrayList(persSet);
7 Collections.sort(persList, new gate.util.OffsetComparator());
8 //Iterate
9 Iterator persIter = persList.iterator();
10 while(persIter.hasNext()){
11     ...
12 }

```

7.4.4 Annotations

An **annotation**, is a form of meta-data attached to a particular section of document content. The connection between the annotation and the content it refers to is made by means of two pointers that represent the start and end locations of the covered content. An annotation must also have a type (or a name) which is used to create classes of similar annotations, usually linked together by their semantics.

An Annotation is defined by:

start node a location in the document content defined by an offset.

end node a location in the document content defined by an offset.

type a String value.

features (see Section 7.4.2).

ID an Integer value. All annotations IDs are unique inside an annotation set.

In GATE Embedded, annotations are defined by the `gate.Annotation` interface and implemented by the `gate.annotation.AnnotationImpl` class. Annotations exist only as members of annotation sets (see Section 7.4.3) and they should not be directly created by means of a constructor. Their creation should always be delegated to the containing annotation set.

7.4.5 GATE Corpora

A corpus in GATE is a Java List (i.e. an implementation of `java.util.List`) of documents. GATE corpora are defined by the `gate.Corpus` interface and the following implementations

are available:

`gate.corpora.CorporaImpl` used for transient corpora.

`gate.corpora.SerialCorporaImpl` used for persistent corpora that are stored in a serial datastore (i.e. as a directory in a file system).

Apart from implementation for the standard List methods, a Corpus also implements the methods in table 7.6.

Method	Purpose
<code>String getDocumentName(int index)</code>	Gets the name of a document in this corpus.
<code>List getDocumentNames()</code>	Gets the names of all the documents in this corpus.
<code>void populate(URL directory, FileFilter filter, String encoding, boolean recurseDirectories)</code>	Fills this corpus with documents created on the fly from selected files in a directory. Uses a <code>FileFilter</code> to select which files will be used and which will be ignored. A simple file filter based on extensions is provided in the Gate distribution (<code>gate.util.ExtensionFileFilter</code>).

Table 7.6: `gate.Corpora` methods.

Creating a corpus from all XML files in a directory

```

1 Corpus corpus = Factory.newCorpus("My XML Files");
2 File directory = ...;
3 ExtensionFileFilter filter = new ExtensionFileFilter("XML files", "xml");
4 URL url = directory.toURL();
5 corpus.populate(url, filter, null, false);

```

Using a DataStore

Assuming that you have a `DataStore` already open called `myDataStore`, this code will ask the datastore to take over persistence of your document, and to synchronise the memory representation of the document with the disk storage:

```

Document persistentDoc = myDataStore.adopt(doc, mySecurity);
myDataStore.sync(persistentDoc);

```

When you want to restore a document (or other LR) from a datastore, you make the same `createResource` call to the `Factory` as for the creation of a transient resource, but this time you tell it the datastore the resource came from, and the ID of the resource in that datastore:

```

1  URL u = .....; // URL of a serial datastore directory
2  SerialDataStore sds = new SerialDataStore(u.toString());
3  sds.open();
4
5  // getLrIds returns a list of LR Ids, so we get the first one
6  Object lrId = sds.getLrIds("gate.corpora.DocumentImpl").get(0);
7
8  // we need to tell the factory about the LR's ID in the data
9  // store, and about which datastore it is in – we do this
10 // via a feature map:
11 FeatureMap features = Factory.newFeatureMap();
12 features.put(DataStore.LR_ID_FEATURE_NAME, lrId);
13 features.put(DataStore.DATASTORE_FEATURE_NAME, sds);
14
15 // read the document back
16 Document doc = (Document)
17     Factory.createResource("gate.corpora.DocumentImpl", features);

```

7.5 Processing Resources

Processing Resources (**PRs**) represent entities that are primarily algorithmic, such as parsers, generators or ngram modellers.

They are created using the GATE Factory in manner similar the Language Resources. Besides the creation-time parameters they also have a set of run-time parameters that are set by the system just before executing them.

Analysers are a particular type of processing resources in the sense that they always have a **document** and a **corpus** among their run-time parameters.

The most used methods for Processing Resources are presented in table 7.7

7.6 Controllers

Controllers are used to create GATE applications. A Controller handles a set of Processing Resources and can execute them following a particular strategy. GATE provides a series of serial controllers (i.e. controllers that run their PRs in sequence):

gate.creole.SerialController: a serial controller that takes any kind of PRs.

gate.creole.SerialAnalyserController: a serial controller that only accepts Language Analysers as member PRs.

Method	Purpose
<code>void setParameterValue(String paramaterName, Object parameterValue)</code>	Sets the value for a specified parameter. method inherited from <code>gate.Resource</code>
<code>void setParameterValues(FeatureMap parameters)</code>	Sets the values for more parameters in one step. method inherited from <code>gate.Resource</code>
<code>Object getParameterValue(String paramaterName)</code>	Gets the value of a named parameter of this resource. method inherited from <code>gate.Resource</code>
<code>Resource init()</code>	Initialise this resource, and return it. method inherited from <code>gate.Resource</code>
<code>void reInit()</code>	Reinitialises the processing resource. After calling this method the resource should be in the state it is after calling <code>init</code> . If the resource depends on external resources (such as rules files) then the resource will re-read those resources. If the data used to create the resource has changed since the resource has been created then the resource will change too after calling <code>reInit()</code> .
<code>void execute()</code>	Starts the execution of this Processing Resource.
<code>void interrupt()</code>	Notifies this PR that it should stop its execution as soon as possible.
<code>boolean isInterrupted()</code>	Checks whether this PR has been interrupted since the last time its <code>Executable.execute()</code> method was called.

Table 7.7: `gate.ProcessingResource` methods.

`gate.creole.ConditionalSerialController`: a serial controller that accepts all types of PRs and that allows the inclusion or exclusion of member PRs from the execution chain according to certain run-time conditions (currently features on the document being processed are used).

`gate.creole.ConditionalSerialAnalyserController`: a serial controller that only accepts Language Analysers and that allows the conditional run of member PRs.

Creating an ANNIE application and running it over a corpus

```

1  // load the ANNIE plugin
2  Gate.getCreoleRegister().registerDirectories(new File(
3    Gate.getPluginsHome(), "ANNIE").toURI().toURL());
4
5  // create a serial analyser controller to run ANNIE with
6  SerialAnalyserController annieController =
7    (SerialAnalyserController) Factory.createResource(
8      "gate.creole.SerialAnalyserController",
9      Factory.newFeatureMap(),
10     Factory.newFeatureMap(), "ANNIE");
11
12 // load each PR as defined in ANNIEConstants
13 for(int i = 0; i < ANNIEConstants.PR_NAMES.length; i++) {
14   // use default parameters
15   FeatureMap params = Factory.newFeatureMap();
16   ProcessingResource pr = (ProcessingResource)
17     Factory.createResource(ANNIEConstants.PR_NAMES[i],
18                           params);
19   // add the PR to the pipeline controller
20   annieController.add(pr);
21 } // for each ANNIE PR
22
23 // Tell ANNIE's controller about the corpus you want to run on
24 Corpus corpus = ...;
25 annieController.setCorpus(corpus);
26 // Run ANNIE
27 annieController.execute();

```

7.7 Duplicating a Resource

Sometimes, particularly in a multi-threaded application, it is useful to be able to create an independent copy of an existing PR, controller or LR. The obvious way to do this is to call `createResource` again, passing the same class name, parameters, features and name, and for many resources this will do the right thing. However there are some resources for which this may be insufficient (e.g. controllers, which also need to duplicate their PRs), unsafe (if a PR uses temporary files, for instance), or simply inefficient. For example for a large gazetteer this would involve loading a second copy of the lists into memory and compiling them into a second identical state machine representation, but a much more efficient way to

achieve the same behaviour would be to use a `SharedDefaultGazetteer` (see section 13.11), which can re-use the existing state machine.

The GATE Factory provides a `duplicate` method which takes an existing resource instance and creates and returns an independent copy of the resource. By default it uses the algorithm described above, extracting the parameter values from the template resource and calling `createResource` to create a duplicate. However, if a particular resource type knows of a better way to duplicate itself it can implement the `CustomDuplication` interface, and provide its own `duplicate` method which the factory will use instead of performing the default duplication algorithm. A caller who needs to duplicate an existing resource can simply call `Factory.duplicate` to obtain a copy, which will be constructed in the appropriate way depending on the resource type.

Note that the duplicate object returned by `Factory.duplicate` will *not necessarily* be of the same class as the original object. However the contract of `Factory.duplicate` specifies that where the original object implements any of a list of core GATE interfaces, the duplicate can be assumed to implement the same ones – if you duplicate a `DefaultGazetteer` the result may not be an instance of `DefaultGazetteer` but it is guaranteed to implement the `Gazetteer` interface.

Full details of how to implement a custom duplicate method in your own resource type can be found in the JavaDoc documentation for the `CustomDuplication` interface and the `Factory.duplicate` method.

7.8 Persistent Applications

GATE Embedded allows the persistent storage of applications in a format based on XML serialisation. This is particularly useful for applications management and distribution. A developer can save the state of an application when he/she stops working on its design and continue developing it in a next session. When the application reaches maturity it can be deployed to the client site using the same method.

When an application (i.e. a *Controller*) is saved, GATE will actually only save the values for the parameters used to create the Processing Resources that are contained in the application. When the application is reloaded, all the PRs will be re-created using the saved parameters.

Many PRs use external resources (files) to define their behaviour and, in most cases, these files are identified using URLs. During the saving process, all the URLs are converted relative URLs based on the location of the application file. This way, if the resources are packaged together with the application file, the entire application can be reliably moved to a different location.

API access to application saving and loading is provided by means of two static methods on the `gate.util.persistence.PersistenceManager` class, listed in table 7.8.

Method	Purpose
<code>public static void saveObjectToFile(Object obj, File file)</code>	Saves the data needed to re-create the provided GATE object to the specified file. The Object provided can be any type of Language or Processing Resource or a Controller. The procedures may work for other types of objects as well (e.g. it supports most Collection types).
<code>public static Object loadObjectFromFile(File file)</code>	Parses the file specified (which needs to be a file created by the above method) and creates the necessary object(s) as specified by the data in the file. Returns the root of the object tree.

Table 7.8: Application Saving and Loading

Saving and loading a GATE application

```

1  //Where to save the application?
2  File file = ...;
3  //What to save?
4  Controller theApplication = ...;
5
6  //save
7  gate.util.persistence.PersistenceManager.
8      saveObjectToFile(theApplication, file);
9  //delete the application
10 Factory.deleteResource(theApplication);
11 theApplication = null;
12
13 [...]
14 //load the application back
15 theApplication = gate.util.persistence.PersistenceManager.
16     loadObjectFromFile(file);

```

7.9 Ontologies

Starting from GATE version 3.1, support for ontologies has been added. Ontologies are nominally Language Resources but are quite different from documents and corpora and are detailed in chapter 14.

Classes related to ontologies are to be found in the `gate.creole.ontology` package and its sub-packages. The top level package defines an abstract API for working with ontologies while the sub-packages contain concrete implementations. A client program should only use

the classes and methods defined in the API and never any of the classes or methods from the implementation packages.

The entry point to the ontology API is the `gate.creole.ontology.Ontology` interface which is the base interface for all concrete implementations. It provides methods for accessing the class hierarchy, listing the instances and the properties.

Ontology implementations are available through plugins. Before an ontology language resource can be created using the `gate.Factory` and before any of the classes and methods in the API can be used, one of the implementing ontology plugins must be loaded. For details see chapter 14.

7.10 Creating a New Annotation Schema

An annotation schema (see Section 3.4.6) can be brought inside GATE through the `creole.xml` file. By using the `AUTOINSTANCE` element, one can create instances of resources defined in `creole.xml`. The `gate.creole.AnnotationSchema` (which is the Java representation of an annotation schema file) initializes with some predefined annotation definitions (annotation schemas) as specified by the GATE team.

Example from GATE's internal `creole.xml` (in `src/gate/resources/creole`):

```
<!-- Annotation schema -->
<RESOURCE>
  <NAME>Annotation schema</NAME>
  <CLASS>gate.creole.AnnotationSchema</CLASS>
  <COMMENT>An annotation type and its features</COMMENT>
  <PARAMETER NAME="xmlFileUrl" COMMENT="The url to the definition file"
    SUFFIXES="xml;xsd">java.net.URL</PARAMETER>
  <AUTOINSTANCE>
    <PARAM NAME ="xmlFileUrl" VALUE="schema/AddressSchema.xml" />
  </AUTOINSTANCE>
  <AUTOINSTANCE>
    <PARAM NAME ="xmlFileUrl" VALUE="schema/DateSchema.xml" />
  </AUTOINSTANCE>
  <AUTOINSTANCE>
    <PARAM NAME ="xmlFileUrl" VALUE="schema/FacilitySchema.xml" />
  </AUTOINSTANCE>
  <!-- etc. -->
</RESOURCE>
```

In order to create a `gate.creole.AnnotationSchema` object from a schema annotation file, one must use the `gate.Factory` class;

```
1 FeatureMap params = new FeatureMap();\n2 param.put("xmlFileUrl",annotSchemaFile.toURL());\n3 AnnotationSchema annotSchema = \n4 Factory.createResource("gate.creole.AnnotationSchema", params);
```

Note: All the elements and their values must be written in lower case, as XML is defined as case sensitive and the parser used for XML Schema inside GATE searches is case sensitive.

In order to be able to write XML Schema definitions, the ones defined in GATE (resources/creole/schema) can be used as a model, or the user can have a look at <http://www.w3.org/2000/10/XMLSchema> for a proper description of the semantics of the elements used.

Some examples of annotation schemas are given in Section 5.4.1.

7.11 Creating a New CREOLE Resource

To create a new resource you need to:

- write a Java class that implements GATE's beans model;
- compile the class, and any others that it uses, into a Java Archive (JAR) file;
- write some XML configuration data for the new resource;
- tell GATE the URL of the new JAR and XML files.

GATE Developer helps you with this process by creating a set of directories and files that implement a basic resource, including a Java code file and a Makefile. This process is called 'bootstrapping'.

For example, let's create a new component called GoldFish, which will be a Processing Resource that looks for all instances of the word 'fish' in a document and adds an annotation of type 'GoldFish'.

First start GATE Developer (see Section 2.2). From the 'Tools' menu select 'BootStrap Wizard', which will pop up the dialogue in figure 7.2. The meaning of the data entry fields:

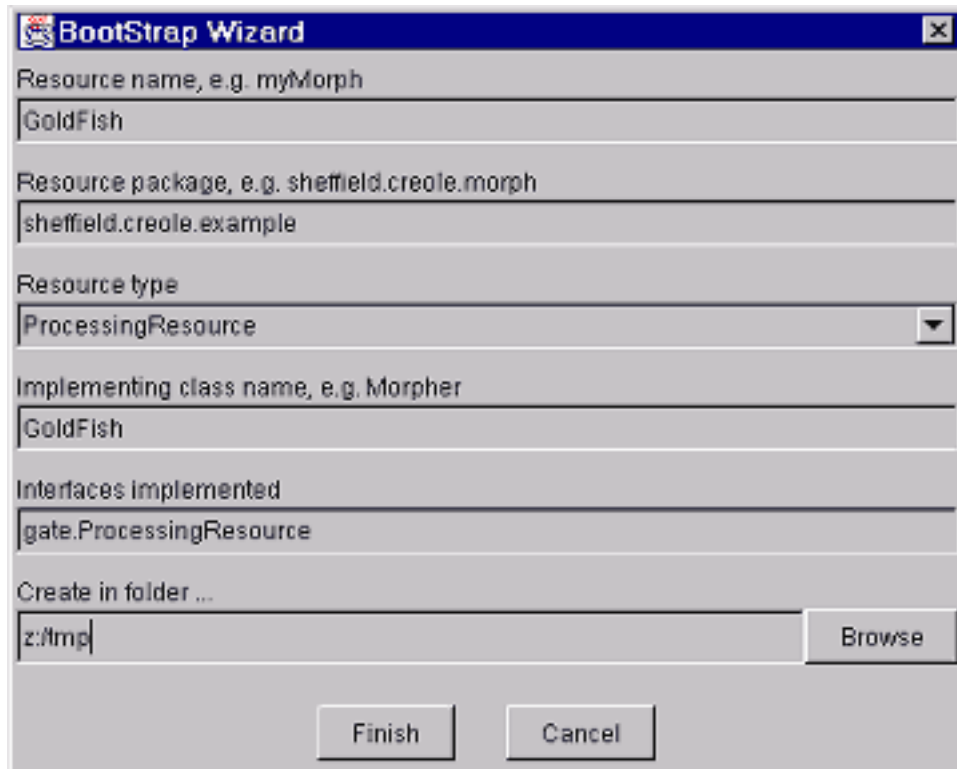


Figure 7.2: BootStrap Wizard Dialogue

- The ‘resource name’ will be displayed when GATE Developer loads the resource, and will be the name of the directory the resource lives in. For our example: **GoldFish**.
- ‘Resource package’ is the Java package that the class representing the resource will be created in. For our example: **sheffield.creole.example**.
- ‘Resource type’ must be one of Language, Processing or Visual Resource. In this case we’re going to process documents (and add annotations to them), so we select **ProcessingResource**.
- ‘Implementing class name’ is the name of the Java class that represents the resource. For our example: **GoldFish**.
- The ‘interfaces implemented’ field allows you to add other interfaces (e.g. `gate.creole.ControllerAwarePR`⁴) that you would like your new resource to implement. In this case we just leave the default (which is to implement the `gate.ProcessingResource` interface).
- The last field selects the directory that you want the new resource created in. For our example: **z:/tmp**.

⁴See Section 4.4.

Now we need to compile the class and package it into a JAR file. The bootstrap wizard creates an Ant build file that makes this very easy – so long as you have Ant set up properly, you can simply run

```
ant jar
```

This will compile the Java source code and package the resulting classes into `GoldFish.jar`. If you don't have your own copy of Ant, you can use the one bundled with GATE – suppose your GATE is installed at `/opt/gate-5.0-snapshot`, then you can use `/opt/gate-5.0-snapshot/bin/ant jar` to build.

You can now load this resource into GATE; see Section 3.6. The default Java code that was created for our GoldFish resource looks like this:

```

1  /*
2  *  GoldFish.java
3  *
4  *  You should probably put a copyright notice here. Why not use the
5  *  GNU licence? (See http://www.gnu.org/.)
6  *
7  *  hamish, 26/9/2001
8  *
9  *  $Id: howto.tex,v 1.130 2006/10/23 12:56:37 ian Exp $
10 /*
11
12 package sheffield.creole.example;
13
14 import java.util.*;
15 import gate.*;
16 import gate.creole.*;
17 import gate.util.*;
18
19 /**
20 * This class is the implementation of the resource GOLDFISH.
21 */
22 @CreoleResource(name = "GoldFish",
23                 comment = "Add a descriptive comment about this resource")
24 public class GoldFish extends AbstractProcessingResource
25     implements ProcessingResource {
26
27
28 } // class GoldFish

```

The default XML configuration for GoldFish looks like this:

```

<!-- creole.xml GoldFish -->
<!-- hamish, 26/9/2001 -->
<!-- $Id: howto.tex,v 1.130 2006/10/23 12:56:37 ian Exp $ -->

<CREOLE-DIRECTORY>

```

```
<JAR SCAN="true">GoldFish.jar</JAR>  
</CREOLE-DIRECTORY>
```

The directory structure containing these files is shown in figure 7.3. `GoldFish.java` lives

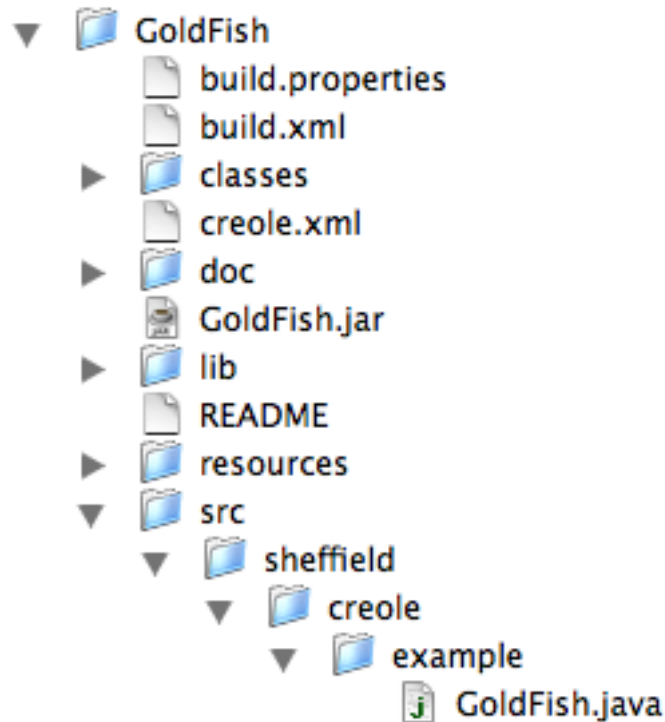


Figure 7.3: BootStrap directory tree

in the `src/sheffield/creole/example` directory. `creole.xml` and `build.xml` are in the top `GoldFish` directory. The `lib` directory is for libraries; the `classes` directory is where Java class files are placed; the `doc` directory is for documentation. These last two, plus `GoldFish.jar` are created by Ant.

This process has the advantage that it creates a complete source tree and build structure for the component, and the disadvantage that it creates a complete source tree and build structure for the component. If you already have a source tree, you will need to chop out the bits you need from the new tree (in this case `GoldFish.java` and `creole.xml`) and copy it into your existing one.

See the example code at <http://gate.ac.uk/wiki/code-repository/>.

7.12 Adding Support for a New Document Format

In order to add a new document format, one needs to extend the `gate.DocumentFormat` class and to implement an abstract method called:

```
1 public void unpackMarkup(Document doc) throws
2   DocumentFormatException
```

This method is supposed to implement the functionality of each format reader and to create annotations on the document. Finally the document's old content will be replaced with a new one containing only the text between markups.

If one needs to add a new textual reader will extend the `gate.corpora.TextualDocumentFormat` and override the `unpackMarkup(doc)` method.

This class needs to be implemented under the Java bean specifications because it will be instantiated by GATE using `Factory.createResource()` method.

The `init()` method that one needs to add and implement is very important because in here the reader defines its means to be selected successfully by GATE. What one needs to do is to add some specific information into certain static maps defined in `DocumentFormat` class, that will be used at reader detection time.

After that, a definition of the reader will be placed into the one's `creole.xml` file and the reader will be available to GATE.

We present for the rest of the section a complete three step example of adding such a reader. The reader we describe in here is an XML reader.

Step 1

Create a new class called `XmlDocumentFormat` that extends `gate.corpora.TextualDocumentFormat`.

Step 2

Implement the `unpackMarkup(Document doc)` which performs the required functionality for the reader. Add XML detection means in `init()` method:

```
1 public Resource init() throws ResourceInstantiationException{
2   // Register XML mime type
3   MimeType mime = new MimeType("text","xml");
4   // Register the class handler for this mime type
5   mimeType2ClassHandlerMap.put(mime.getType()+ "/" + mime.getSubtype(),
6                               this);
7   // Register the mime type with mime string
8   mimeType2mimeTypeMap.put(mime.getType() + "/" + mime.getSubtype(),
9                             mime);
10  // Register file suffixes for this mime type
11  suffixes2mimeTypeMap.put("xml",mime);
```



```

12  suffixes2mimeTypeMap.put("xhtm",mime);
13  suffixes2mimeTypeMap.put("xhtml",mime);
14  // Register magic numbers for this mime type
15  magic2mimeTypeMap.put("<?xml",mime);
16  // Set the mimeType for this language resource
17  setMimeType(mime);
18  return this;
19 } // init()

```

More details about the information from those maps can be found in Section 5.5.1

Step 3

Add the following creole definition in the creole.xml document.

```

<RESOURCE>
  <NAME>My XML Document Format</NAME>
  <CLASS>mypackage.XmlDocumentFormat</CLASS>
  <AUTOINSTANCE/>
  <PRIVATE/>
</RESOURCE>

```

More information on the operation of GATE's document format analysers may be found in Section 5.5.

7.13 Using GATE Embedded in a Multithreaded Environment

GATE Embedded can be used in multithreaded applications, so long as you observe a few restrictions. First, you must initialise GATE by calling `Gate.init()` *exactly once* in your application, typically in the application startup phase before any concurrent processing threads are started.

Secondly, you must not make calls that affect the global state of GATE (e.g. loading or unloading plugins) in more than one thread at a time. Again, you would typically load all the plugins your application requires at initialisation time. It is safe to create *instances* of resources in multiple threads concurrently.

Thirdly, it is important to note that individual GATE processing resources, language resources and controllers are by design *not* thread safe – it is not possible to use a single instance of a controller/PR/LR in multiple threads at the same time – but for a well written resource it should be possible to use several different instances of the same resource at once, each in a different thread. When writing your own resource classes you should bear the following in mind, to ensure that your resource will be useable in this way.

- Avoid static data. Where possible, you should avoid using static fields in your class, and you should try and take all configuration data via the CREOLE parameters you declare in your creole.xml file. System properties may be appropriate for truly static configuration, such as the location of an external executable, but even then it is generally better to stick to CREOLE parameters – a user may wish to use two different instances of your PR, each talking to a different executable.
- Read parameters at the correct time. Init-time parameters should be read in the `init()` (and `reInit()`) method, and for processing resources runtime parameters should be read at each `execute()`.
- Use temporary files correctly. If your resource makes use of external temporary files you should create them using `File.createTempFile()` at `init` or `execute` time, as appropriate. Do not use hardcoded file names for temporary files.
- If there are objects that can be shared between different instances of your resource, make sure these objects are accessed either read-only, or in a thread-safe way. In particular you must be very careful if your resource can take other resource instances as `init` or runtime parameters (e.g. the Flexible Gazetteer, Section 13.7).

Of course, if you are writing a PR that is simply a wrapper around an external library that imposes these kinds of limitations there is only so much you can do. If your resource cannot be made safe you should *document this fact clearly*.

All the standard ANNIE PRs are safe when independent instances are used in different threads concurrently, as are the standard transient document, transient corpus and controller classes. A typical pattern of development for a multithreaded GATE-based application is:

- Develop your GATE processing pipeline in GATE Developer.
- Save your pipeline as a `.gapp` file.
- In your application's initialisation phase, load n copies of the pipeline using `PersistenceManager.loadObjectFromFile()` (see the Javadoc documentation for details), or load the pipeline once and then make copies of it using `Factory.duplicate` as described in section 7.7, and either give one copy to each thread or store them in a pool (e.g. a `LinkedList`).
- When you need to process a text, get one copy of the pipeline from the pool, and return it to the pool when you have finished processing.

Alternatively you can use the Spring Framework as described in the next section to handle the pooling for you.

7.14 Using GATE Embedded within a Spring Application

GATE Embedded provides helper classes to allow GATE resources to be created and managed by the Spring framework. For Spring 2.0 or later, GATE Embedded provides a custom namespace handler that makes them extremely easy to use. To use this namespace, put the following declarations in your bean definition file:

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:gate="http://gate.ac.uk/ns/spring"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://gate.ac.uk/ns/spring
    http://gate.ac.uk/ns/spring.xsd">
```

You can have Spring initialise GATE:

```
<gate:init gate-home="WEB-INF" user-config-file="WEB-INF/user.xml">
  <gate:preload-plugins>
    <value>WEB-INF/ANNIE</value>
    <value>http://example.org/gate-plugin</value>
  </gate:preload-plugins>
</gate:init>
```

To create a GATE resource, use the `<gate:resource>` element.

```
<gate:resource id="sharedOntology" scope="singleton"
  resource-class="gate.creole.ontology.owlim.OWLIMOntologyLR">
  <gate:parameters>
    <entry key="rdfXmlURL">
      <value type="org.springframework.core.io.Resource"
        >WEB-INF/ontology.rdf</value>
    </entry>
  </gate:parameters>
</gate:resource>
```

If you are familiar with Spring you will see that `<gate:parameters>` uses the same format as the standard `<map>` element, but values whose type is a Spring `Resource` will be converted to URLs before being passed to the GATE resource.

You can load a GATE saved application with

```

<gate:saved-application location="WEB-INF/application.gapp" scope="prototype">
  <gate:customisers>
    <gate:set-parameter pr-name="custom transducer" name="ontology"
      ref="sharedOntology" />
  </gate:customisers>
</gate:saved-application>

```

‘Customisers’ are used to customise the application after it is loaded. In the example above, we load a singleton copy of an ontology which is then shared between all the separate instances of the (prototype) application. The `<gate:set-parameter>` customiser accepts all the same ways to provide a value as the standard Spring `<property>` element (a “value” or “ref” attribute, or a sub-element - `<value>`, `<list>`, `<bean>`, `<gate:resource>` ...).

The `<gate:add-pr>` customiser provides support for the case where most of the application is in a saved state, but we want to create one or two extra PRs with Spring (maybe to inject other Spring beans as init parameters) and add them to the pipeline.

```

<gate:saved-application ...>
  <gate:customisers>
    <gate:add-pr add-before="OrthoMatcher" ref="myPr" />
  </gate:customisers>
</gate:saved-application>

```

By default, the `<gate:add-pr>` customiser adds the target PR at the end of the pipeline, but an `add-before` or `add-after` attribute can be used to specify the name of a PR before (or after) which this PR should be placed. Alternatively, an `index` attribute places the PR at a specific (0-based) index into the pipeline. The PR to add can be specified either as a ‘ref’ attribute, or with a nested `<bean>` or `<gate:resource>` element.

7.14.1 Duplication in Spring

The above example defines the `<gate:application>` as a prototype-scoped bean, which means the saved application state will be loaded afresh each time the bean is fetched from the bean factory (either explicitly using `getBean` or implicitly when it is injected as a dependency of another bean). However in many cases it is better to load the application once and then duplicate it as required (as described in section 7.7), as this allows resources to optimise their memory usage, for example by sharing a single in-memory representation of a large gazetteer list between several instances of the gazetteer PR. This approach is supported by the `<gate:duplicate>` tag.

```

<gate:duplicate id="theApp">
  <gate:saved-application location="/WEB-INF/application.xgapp" />
</gate:duplicate>

```

The `<gate:duplicate>` tag acts like a prototype bean definition, in that each time it is fetched or injected it will call `Factory.duplicate` to create a new duplicate of its template resource (declared as a nested element or referenced by the `template-ref` attribute). However the tag also keeps track of all the duplicate instances it has returned over its lifetime, and will ensure they are released (using `Factory.deleteResource`) when the Spring context is shut down.

The `<gate:duplicate>` tag also supports customisers, which will be applied to the newly-created *duplicate* resource before it is returned. This is subtly different from applying the customisers to the template resource itself, which would cause them to be applied once to the *original* resource before it is first duplicated.

Finally, `<gate:duplicate>` takes an optional boolean attribute `return-template`. If set to false (or omitted, as this is the default behaviour), the tag always returns a duplicate — the original template resource is used only as a template and is not made available for use. If set to true, the first time the bean defined by the tag is injected or fetched, the original template resource is returned. Subsequent uses of the tag will return duplicates. Generally speaking, it is only safe to set `return-template="true"` when there are no customisers, and when the duplicates will all be created up-front before any of them are used. If the duplicates will be created asynchronously (e.g. with a dynamically expanding pool, see below) then it is possible that, for example, a template application may be duplicated in one thread whilst it is being executed by another thread, which may lead to unpredictable behaviour.

7.14.2 Spring pooling

In a multithreaded application it is vital that individual GATE resources are not used in more than one thread at the same time. Because of this, multithreaded applications that use GATE Embedded often need to use some form of pooling to provide thread-safe access to GATE components. This can be managed by hand, but the Spring framework has built-in tools to support transparent pooling of Spring-managed beans. Spring can create a pool of identical objects, then expose a single “proxy” object (offering the same interface) for use by clients. Each method call on the proxy object will be routed to an available member of the pool in such a way as to guarantee that each member of the pool is accessed by no more than one thread at a time.

Since the pooling is handled at the level of method calls, this approach is not used to create a pool of GATE resources directly — making use of a GATE PR typically involves a sequence of method calls (at least `setDocument(doc)`, `execute()` and `setDocument(null)`), and creating a pooling proxy for the resource may result in these calls going to different members of the pool. Instead the typical use of this technique is to define a helper object with a single method that internally calls the GATE API methods in the correct sequence, and then create a pool of these helpers. The interface `gate.util.DocumentProcessor` and its associated implementation `gate.util.LanguageAnalyserDocumentProcessor` are useful for this. The `DocumentProcessor` interface defines a `processDocument` method that takes a GATE

document and performs some processing on it. `LanguageAnalyserDocumentProcessor` implements this interface using a GATE `LanguageAnalyser` (such as a saved “corpus pipeline” application) to do the processing. A pool of `LanguageAnalyserDocumentProcessor` instances can be exposed through a proxy which can then be called from several threads.

The machinery to implement this is all built into Spring, but the configuration typically required to enable it is quite fiddly, involving at least three co-operating bean definitions. Since the technique is so useful with GATE Embedded, GATE provides a special syntax to configure pooling in a simple way. Given the `<gate:duplicate id="theApp">` definition from the previous section we can create a `DocumentProcessor` proxy that can handle up to five concurrent requests as follows:

```
<bean id="processor"
      class="gate.util.LanguageAnalyserDocumentProcessor">
  <property name="analyser" ref="theApp" />
  <gate:pooled-proxy max-size="5" />
</bean>
```

The `<gate:pooled-proxy>` element decorates a singleton bean definition. It converts the original definition to prototype scope and replaces it with a singleton proxy delegating to a pool of instances of the prototype bean. The pool parameters are controlled by attributes of the `<gate:pooled-proxy>` element, the most important ones being:

max-size The maximum size of the pool. If more than this number of threads try to call methods on the proxy at the same time, the others will (by default) block until an object is returned to the pool.

initial-size The default behaviour of Spring’s pooling tools is to create instances in the pool on demand (up to the max-size). This attribute instead causes **initial-size** instances to be created up-front and added to the pool when it is first created.

when-exhausted-action-name What to do when the pool is exhausted (i.e. there are already **max-size** concurrent calls in progress and another one arrives). Should be set to one of `WHEN_EXHAUSTED_BLOCK` (the default, meaning block the excess requests until an object becomes free), `WHEN_EXHAUSTED_GROW` (create a new object anyway, even though this pushes the pool beyond **max-size**) or `WHEN_EXHAUSTED_FAIL` (cause the excess calls to fail with an exception).

Many more options are available, corresponding to the properties of the Spring `CommonsPoolTargetSource` class. These allow you, for example, to configure a pool that dynamically grows and shrinks as necessary, releasing objects that have been idle for a set amount of time. See the JavaDoc documentation of `CommonsPoolTargetSource` (and the documentation for Apache `commons-pool`) for full details.

Note that the `<gate:pooled-proxy>` technique is not tied to GATE in any way, it is simply an easy way to configure standard Spring beans and can be used with any bean that needs to be pooled, not just objects that make use of GATE.

7.14.3 Further reading

These custom elements all define various factory beans. For full details, see the JavaDocs for `gate.util.spring` (the factory beans) and `gate.util.spring.xml` (the `gate:` namespace handler). The main Spring framework API documentation is the best place to look for more detail on the pooling facilities provided by Spring AOP.

Note: the former approach using factory methods of the `gate.util.spring.SpringFactory` class will still work, but should be considered deprecated in favour of the new factory beans.

7.15 Using GATE Embedded within a Tomcat Web Application

Embedding GATE in a Tomcat web application involves several steps.

1. Put the necessary JAR files (`gate.jar` and all or most of the jars in `gate/lib`) in your `webapp/WEB-INF/lib`.
2. Put the plugins that your application depends on in a suitable location (e.g. `webapp/WEB-INF/plugins`).
3. Create suitable `gate.xml` configuration files for your environment.
4. Set the appropriate paths in your application before calling `Gate.init()`.

This process is detailed in the following sections.

7.15.1 Recommended Directory Structure

You will need to create a number of other files in your web application to allow GATE to work:

- Site and user `gate.xml` config files - we highly recommend defining these specifically for the web application, rather than relying on the default files on your application server.
- The plugins your application requires.

In this guide, we assume the following layout:

```
webapp/  
  WEB-INF/  
    gate.xml  
    user-gate.xml  
  plugins/  
    ANNIE/  
    etc.
```

7.15.2 Configuration Files

Your `gate.xml` (the ‘site-wide configuration file’) should be as simple as possible:

```
<?xml version="1.0" encoding="UTF-8" ?>  
<GATE>  
  <GATECONFIG Save_options_on_exit="false"  
              Save_session_on_exit="false" />  
</GATE>
```

Similarly, keep the `user-gate.xml` (the ‘user config file’) simple:

```
<?xml version="1.0" encoding="UTF-8" ?>  
<GATE>  
  <GATECONFIG Known_plugin_path=";"  
              Load_plugin_path=";" />  
</GATE>
```

This way, you can control exactly which plugins are loaded in your webapp code.

7.15.3 Initialization Code

Given the directory structure shown above, you can initialize GATE in your web application like this:

```
1 // imports  
2 ...  
3 public class MyServlet extends HttpServlet {  
4   private static boolean gateInited = false;  
5  
6   public void init() throws ServletException {  
7     if(!gateInited) {
```



```

8     try {
9         ServletContext ctx = getServletContext();
10
11         // use /path/to/your/webapp/WEB-INF as gate.home
12         File gateHome = new File(ctx.getRealPath("/WEB-INF"));
13
14         Gate.setGateHome(gateHome);
15         // thus webapp/WEB-INF/plugins is the plugins directory, and
16         // webapp/WEB-INF/gate.xml is the site config file.
17
18         // Use webapp/WEB-INF/user-gate.xml as the user config file,
19         // to avoid confusion with your own user config.
20         Gate.setUserConfigFile(new File(gateHome, "user-gate.xml"));
21
22         Gate.init();
23         // load plugins, for example...
24         Gate.getCreoleRegister().registerDirectories(
25             ctx.getResource("/WEB-INF/plugins/ANNIE"));
26
27         gateInited = true;
28     }
29     catch(Exception ex) {
30         throw new ServletException("Exception initialising GATE",
31             ex);
32     }
33 }
34 }
35 }

```

Once initialized, you can create GATE resources using the Factory in the usual way (for example, see Section 7.1 for an example of how to create an ANNIE application). You should also read Section 7.13 for important notes on using GATE Embedded in a multithreaded application.

Instead of an initialization servlet you could also consider doing your initialization in a `ServletContextListener`, or using Spring (see Section 7.14).

7.16 Groovy for GATE

Groovy is a dynamic programming language based on Java. Groovy is not used in the core GATE distribution, so to enable the Groovy features in GATE you must first load the Groovy plugin. Loading this plugin:

- provides access to the Groovy scripting console (configured with some extensions for GATE) from the GATE Developer “Tools” menu.
- provides a PR to run a Groovy script over documents.

- enhances a number of core GATE classes with additional convenience methods that can be used from any Groovy code including the console, the script PR, and any Groovy class that uses the GATE Embedded API.

This section describes these features in detail, but assumes that the reader already has some knowledge of the Groovy language. If you are not already familiar with Groovy you should read this section in conjunction with Groovy’s own documentation at <http://groovy.codehaus.org/>.

7.16.1 Groovy Scripting Console for GATE

Loading the Groovy plugin in GATE Developer will provide a “Groovy Console” item in the Tools/Groovy Tools menu. This menu item opens the standard Groovy console window (<http://groovy.codehaus.org/Groovy+Console>).

To help scripting GATE in Groovy, the console is pre-configured to import all classes from the `gate`, `gate.annotation`, `gate.util`, `gate.jape` and `gate.creole.ontology` packages of the core GATE API⁵. This means you can refer to classes and interfaces such as `Factory`, `AnnotationSet`, `Gate`, etc. without needing to prefix them with a package name. In addition, the following (read-only) variable bindings are pre-defined in the Groovy Console.

- **corpora**: a list of loaded corpora LRs (Corpus)
- **docs**: a list of all loaded document LRs (DocumentImpl)
- **prs**: a list of all loaded PRs
- **apps**: a list of all loaded Applications (AbstractController)

These variables are automatically updated as resources are created and deleted in GATE.

Here’s an example script. It finds all documents with a feature “annotator” set to “fred”, and puts them in a new corpus called “fredsDocs”.

```

1 Factory.newCorpus("fredsDocs").addAll(
2     docs.findAll{
3         it.features.annotator == "fred"
4     }
5 )

```

You can find other examples (and add your own) in the Groovy script repository on the GATE Wiki: <http://gate.ac.uk/wiki/groovy-recipes/>.

⁵These are the same classes that are imported by default for use in Java code on the right hand side of JAPE rules.

Why won't the 'Groovy executing' dialog go away? Sometimes, when you execute a Groovy script through the console, a dialog will appear, saying "Groovy is executing. Please wait". The dialog fails to go away even when the script has ended, and cannot be closed by clicking the "Interrupt" button. You can, however, continue to use the Groovy Console, and the dialog will usually go away next time you run a script. This is not a GATE problem: it is a Groovy problem.

7.16.2 Groovy scripting PR

The Groovy scripting PR enables you to load and execute Groovy scripts as part of a GATE application pipeline. The Groovy scripting PR is made available when you load the Groovy plugin via the plugin manager.

Parameters

The Groovy scripting PR has a single initialisation parameter

- **scriptURL**: the path to a valid Groovy script

It has three runtime parameters

- **inputASName**: an optional annotation set intended to be used as input by the PR (but note that the PR has access to all annotation sets)
- **outputASName**: an optional annotation set intended to be used as output by the PR (but note that the PR has access to all annotation sets)
- **scriptParams**: optional parameters for the script. In a creole.xml file, these should be specified as key=value pairs, each pair separated by a comma. For example: 'name=fred,type=person'. In the GATE GUI, these are specified via a dialog.

Script bindings

As with the Groovy console described above, and with JAPE right-hand-side Java code, Groovy scripts run by the scripting PR implicitly import all classes from the `gate`, `gate.annotation`, `gate.util`, `gate.jape` and `gate.creole.ontology` packages of the core GATE API. The Groovy scripting PR also makes available the following bindings, which you can use in your scripts:

- **doc**: the current document (`DocumentImpl`)

- **content**: the string content of the current document
- **inputAS**: the annotation set specified by inputASName in the PRs runtime parameters
- **outputAS**: the annotation set specified by outputASName in the PRs runtime parameters

Note that inputAS and outputAS are intended to be used as input and output Annotation-Sets. This is, however, a convention: there is nothing to stop a script writing to or reading from any AnnotationSet.

Passing parameters to the script

In addition to the above bindings, one further binding is available to the script:

- **scriptParams**: a FeatureMap with keys and values as specified by the scriptParams runtime parameter

For example, if you were to create a scriptParams runtime parameter for your PR, with the keys and values: 'name=fred,type=person', then the values could be retrieved in your script via scriptParams.name and scriptParams.type

Examples

The plugin directory Groovy/resources/scripts contains some example scripts. Below is the code for a naive regular expression PR.

```

1
2 matcher = content =~ scriptParams.regex
3 while(matcher.find())
4     outputAS.add(matcher.start(),
5                 matcher.end(),
6                 scriptParams.type,
7                 Factory.newFeatureMap())

```

The script needs to have the runtime parameter scriptParams set with keys and values as follows:

- **regex**: the Groovy regular expression that you want to match e.g. `[^\s]*ing`
- **type**: the type of the annotation to create for each regex match, e.g. **regexMatch**

When the PR is run over a document, the script will first make a matcher over the document content for the regular expression given by the **regex** parameter. It will iterate over all matches for this regular expression, adding a new annotation for each, with a type as given by the **type** parameter.

7.16.3 Utility methods

Loading the Groovy plugin adds some additional methods to several of the core GATE API classes and interfaces using the Groovy “mixin” mechanism. Any Groovy code that runs after the plugin has been loaded can make use of these additional methods, including snippets run in the Groovy console, scripts run using the Script PR, and any other Groovy code that uses the GATE Embedded API.

The methods that are injected come from two classes. The `gate.Utills` class (part of the core GATE API in `gate.jar`) defines a number of static methods that can be used to simplify common tasks such as getting the string covered by an annotation or annotation set, finding the start or end offset of an annotation (or set), etc. These methods do not use any Groovy-specific types, so they are usable from pure Java code in the usual way as well as being mixed in for use in Groovy. Additionally, the class `gate.groovy.GateGroovyMethods` (part of the Groovy plugin) provides methods that use Groovy types such as closures and ranges.

The added methods include:

- Unified access to the start and end offsets of an `Annotation`, `AnnotationSet` or `Document`: e.g. `someAnnotation.start()` or `anAnnotationSet.end()`
- Simple access to the `DocumentContent` or string covered by an annotation or annotation set: `document.stringFor(anAnnotation)`, `document.contentFor(annotationSet)`
- Simple access to the length of an annotation or document, either as an `int` (`annotation.length()`) or a `long` (`annotation.lengthLong()`).
- A method to construct a `FeatureMap` from any map, to support constructions like `def params = [sourceUrl:'http://gate.ac.uk', encoding:'UTF-8'].toFeatureMap()`
- A method to convert an annotation set into a `List` of annotations in the order they appear in the document, for iteration in a predictable order: `annSet.inDocumentOrder().collect { it.type }`
- The `each`, `eachWithIndex` and `collect` methods for a corpus have been redefined to properly load and unload documents if the corpus is stored in a datastore.
- Various `getAt` methods to support constructions like `annotationSet["Token"]` (get all `Token` annotations from the set), `annotationSet[15..20]` (get all annotations between offsets 15 and 20), `documentContent[0..10]` (get the document content between offsets 0 and 10).

- A `withResource` method for any resource, which calls a closure with the resource passed as a parameter, and ensures that the resource is properly deleted when the closure completes (analogous to the default Groovy method `InputStream.withStream`).

For full details, see the source code or javadoc documentation for these two classes.

7.17 Saving Config Data to `gate.xml`

Arbitrary feature/value data items can be saved to the user's `gate.xml` file via the following API calls:

To get the config data: `Map configData = Gate.getUserConfig()`.

To add config data simply put pairs into the map: `configData.put("my new config key", "value");`.

To write the config data back to the XML file: `Gate.writeUserConfig();`

Note that new config data will simply override old values, where the keys are the same. In this way defaults can be set up by putting their values in the main `gate.xml` file, or the site `gate.xml` file; they can then be overridden by the user's `gate.xml` file.

7.18 Annotation merging through the API

If we have annotations about the same subject on the same document from different annotators, we may need to merge those annotations to form a unified annotation. Two approaches for merging annotations are implemented in the API, via static methods in the class `gate.util.AnnotationMerging`.

The two methods have very similar input and output parameters. Each of the methods takes an array of annotation sets, which should be the same annotation type on the same document from different annotators, as input. A single feature can also be specified as a parameter (or given as *null* if no feature is to be specified).

The output is a map, the key of which is one merged annotation and the value of which represents the annotators (in terms of the indices of the array of annotation sets) who support the annotation. The methods also have a boolean input parameter to indicate whether or not the annotations from different annotators are based on the same set of instances, which can be determined by the static method `public boolean isSameInstancesForAnnotators(AnnotationSet[] annsA)` in the class `gate.util.IaaCalculation`. One instance corresponds to all the annotations with the same span. If the annotation sets are based on the

same set of instances, the merging methods will ensure that the merged annotations are on the same set of instances.

The two methods corresponding to those described for the Annotation Merging plugin described in Section 19.11. They are:

- The Method *public static void mergeAnnotation(AnnotationSet[] annsArr, String nameFeat, HashMap<Annotation,String>mergeAnns, int numMinK, boolean isTheSameInstances)* merges the annotations stored in the array *annsArr*. The merged annotation is put into the map *mergeAnns*, with a key of the merged annotation and value of a string containing the indices of elements in the annotation set array *annsArr* which contain that annotation. *NumMinK* specifies the minimal number of the annotators supporting one merged annotation. The boolean parameter *isTheSameInstances* indicate if or not those annotation sets for merging are based on the same instances.
- Method *public static void mergeAnnotationMajority(AnnotationSet[] annsArr, String nameFeat, HashMap<Annotation, String>mergeAnns, boolean isTheSameInstances)* selects the annotations which the majority of the annotators agree on. The meanings of parameters are the same as those in the above method.

Chapter 8

JAPE: Regular Expressions over Annotations

If Osama bin Laden did not exist, it would be necessary to invent him. For the past four years, his name has been invoked whenever a US president has sought to increase the defence budget or wriggle out of arms control treaties. He has been used to justify even President Bush's missile defence programme, though neither he nor his associates are known to possess anything approaching ballistic missile technology. Now he has become the personification of evil required to launch a crusade for good: the face behind the faceless terror.

The closer you look, the weaker the case against Bin Laden becomes. While the terrorists who inflicted Tuesday's dreadful wound may have been inspired by him, there is, as yet, no evidence that they were instructed by him. Bin Laden's presumed guilt appears to rest on the supposition that he is the sort of man who would have done it. But his culpability is irrelevant: his usefulness to western governments lies in his power to terrify. When billions of pounds of military spending are at stake, rogue states and terrorist warlords become assets precisely because they are liabilities.

The need for dissent, George Monbiot, The Guardian, Tuesday September 18, 2001.

JAPE is a Java Annotation Patterns Engine. JAPE provides finite state transduction over annotations based on regular expressions. JAPE is a version of CPSL – Common Pattern Specification Language¹. This chapter introduces JAPE, and outlines the functionality available. (You can find an excellent tutorial [here](#); thanks to Dhaval Thakker, Taha Osmin and Phil Lakin).

¹A good description of the original version of this language is in Doug Appelt's *TextPro* manual. Doug was a great help to us in implementing JAPE. Thanks Doug!

JAPE allows you to recognise regular expressions in annotations on documents. Hang on, there's something wrong here: a regular language can only describe sets of strings, not graphs, and GATE's model of annotations is based on graphs. Hmm. Another way of saying this: typically, regular expressions are applied to character strings, a simple linear sequence of items, but here we are applying them to a much more complex data structure. The result is that in certain cases the matching process is non-deterministic (i.e. the results are dependent on random factors like the addresses at which data is stored in the virtual machine): when there is structure in the graph being matched that requires more than the power of a regular automaton to recognise, JAPE chooses an alternative arbitrarily. However, this is not the bad news that it seems to be, as it turns out that in many useful cases the data stored in annotation graphs in GATE (and other language processing systems) can be regarded as simple sequences, and matched deterministically with regular expressions.

A JAPE grammar consists of a set of phases, each of which consists of a set of pattern/action rules. The phases run sequentially and constitute a cascade of finite state transducers over annotations. The left-hand-side (LHS) of the rules consist of an annotation pattern description. The right-hand-side (RHS) consists of annotation manipulation statements. Annotations matched on the LHS of a rule may be referred to on the RHS by means of labels that are attached to pattern elements. Consider the following example:

```
Phase: Jobtitle
Input: Lookup
Options: control = appelt debug = true

Rule: Jobtitle1
(
  {Lookup.majorType == jobtitle}
  (
    {Lookup.majorType == jobtitle}
  )?
)
:jobtitle
-->
  :jobtitle.JobTitle = {rule = "JobTitle1"}
```

The LHS is the part preceding the '-->' and the RHS is the part following it. The LHS specifies a pattern to be matched to the annotated GATE document, whereas the RHS specifies what is to be done to the matched text. In this example, we have a rule entitled 'Jobtitle1', which will match text annotated with a 'Lookup' annotation with a 'majorType' feature of 'jobtitle', followed optionally by further text annotated as a 'Lookup' with 'majorType' of 'jobtitle'. Once this rule has matched a sequence of text, the entire sequence is allocated a

label by the rule, and in this case, the label is ‘jobtitle’. On the RHS, we refer to this span of text using the label given in the LHS; ‘jobtitle’. We say that this text is to be given an annotation of type ‘JobTitle’ and a ‘rule’ feature set to ‘JobTitle1’.

We began the JAPE grammar by giving it a phase name, e.g. ‘Phase: Jobtitle’. JAPE grammars can be cascaded, and so each grammar is considered to be a ‘phase’ (see Section 8.5). The phase name makes up part of the Java class name for the compiled RHS actions. Because of this, it must contain alphanumeric characters and underscores only, and cannot start with a number.

We also provide a list of the annotation types we will use in the grammar. In this case, we say ‘Input: Lookup’ because the only annotation type we use on the LHS are Lookup annotations. If no annotations are defined, all annotations will be matched.

Then, several options are set:

- Control; in this case, ‘appelt’. This defines the method of rule matching (see Section 8.4)
- Debug. When set to true, if the grammar is running in Appelt mode and there is more than one possible match, the conflicts will be displayed on the standard output.

A wide range of functionality can be used with JAPE, making it a very powerful system. Section 8.1 gives an overview of some common LHS tasks. Section 8.2 talks about the various operators available for use on the LHS. After that, Section 8.3 outlines RHS functionality. Section 8.4 talks about priority and Section 8.5 talks about phases. Section 8.6 talks about using Java code on the RHS, which is the main way of increasing the power of the RHS. We conclude the chapter with some miscellaneous JAPE-related topics of interest.

8.1 The Left-Hand Side

The LHS of a JAPE grammar aims to match the text span to be annotated, whilst avoiding undesirable matches. There are various tools available to enable you to do this. This section outlines how you would approach various common tasks on the LHS of your JAPE grammar.

8.1.1 Matching a Simple Text String

To match a simple text string, you need to refer to a feature on an annotation that contains the string; for example,

```
{Token.string == "of"}
```

The following grammar shows a sequence of strings being matched. Bracketing, along with the ‘or’ operator, is used to define how the strings should come together:

```
Phase: UrlPre
Input:  Token SpaceToken
Options: control = appelt

Rule: Urlpre

( (({Token.string == "http"} |
   {Token.string == "ftp"})
  {Token.string == ":"}
  {Token.string == "/"})
  {Token.string == "/"})
) |
({Token.string == "www"}
 {Token.string == "."}
)
):urlpre
-->
:urlpre.UrlPre = {rule = "UrlPre"}
```

Alternatively you could use the ‘string’ metaproperty. See Section 8.1.4 for an example of using metaproperties.

8.1.2 Matching Entire Annotation Types

You can specify the presence of an annotation previously assigned from a gazetteer, tokeniser, or other module. For example, the following will match a Lookup annotation:

```
{Lookup}
```

The following will match if there is *not* a Lookup annotation at this location:

```
{!Lookup}
```

The following rule shows several different annotation types being matched. We also see a string being matched, and again, the use of the ‘or’ operator:

```

Rule: Known
Priority: 100
(
  {Location}|
  {Person}|
  {Date}|
  {Organization}|
  {Address}|
  {Money} |
  {Percent}|
  {Token.string == "Dear"}|
  {JobTitle}|
  {Lookup}
):known
-->
{}

```

8.1.3 Using Attributes and Values

You can specify the attributes (and values) of an annotation to be matched. Several operators are supported; see Section 8.2 for full details:

- `{Token.kind == "number"}`, `{Token.length != 4}` - equality and inequality.
- `{Token.string > "aardvark"}`, `{Token.length < 10}` - comparison operators. `>=` and `<=` are also supported.
- `{Token.string =~ "[Dd]ogs"}`, `{Token.string !~ "(?i)hello"}` - regular expression. `==~` and `!=~` are also provided, for whole-string matching.
- `{X contains Y}` and `{X within Y}` for checking annotations within the context of other annotations.

In the following rule, the ‘category’ feature of the ‘Token’ annotation is used, along with the ‘equals’ operator:

```

Rule: Unknown
Priority: 50
(
  {Token.category == NNP}
)
:unknown
-->
:unknown.Unknown = {kind = "PN", rule = Unknown}

```

8.1.4 Using Meta-Properties

In addition to referencing annotation features, JAPE allows access to other ‘meta-properties’ of an annotation. This is done by using an ‘@’ symbol rather than a ‘.’ symbol after the annotation type name. The three meta-properties that are built in are:

- `length` - returns the spanning length of the annotation.
- `string` - returns the string spanned by the annotation in the document.
- `cleanString` - Like `string`, but with extra white space stripped out. (i.e. ‘\s+’ goes to a single space and leading or trailing white space is removed).

```
{X@length > "5"}:label-->:label.New = {}
```

Only the value of ‘meta-property’ `@string` can be accessed from a non-java RHS of a rule:

```
{X@length > "5"}:label-->:label.New = {somefeat = :label.X@string }
```

The ‘meta-properties’ `@length` and `@cleanString` are not available in the RHS but may be added in the future.

8.1.5 Multiple Pattern/Action Pairs

It is also possible to have more than one pattern and corresponding action, as shown in the rule below. On the LHS, each pattern is enclosed in a set of round brackets and has a unique label; on the RHS, each label is associated with an action. In this example, the `Lookup` annotation is labelled ‘`jobtitle`’ and is given the new annotation `JobTitle`; the `TempPerson` annotation is labelled ‘`person`’ and is given the new annotation ‘`Person`’.

```
Rule: PersonJobTitle
```

```
Priority: 20
```

```
(
  {Lookup.majorType == jobtitle}
):jobtitle
(
  {TempPerson}
):person
-->
:jobtitle.JobTitle = {rule = "PersonJobTitle"},
:person.Person = {kind = "personName", rule = "PersonJobTitle"}
```

Similarly, labelled patterns can be nested, as in the example below, where the whole pattern is annotated as Person, but within the pattern, the jobtitle is annotated as JobTitle.

```
Rule: PersonJobTitle2
Priority: 20

(
  (
    {Lookup.majorType == jobtitle}
  ):jobtitle
  {TempPerson}
):person
-->
  :jobtitle.JobTitle = {rule = "PersonJobTitle"},
  :person.Person = {kind = "personName", rule = "PersonJobTitle"}
```

8.1.6 LHS Macros

Macros allow you to create a definition that can then be used multiple times in your JAPE rules. In the following JAPE grammar, we have a cascade of macros used. The macro ‘AMOUNT_NUMBER’ makes use of the macros ‘MILLION_BILLION’ and ‘NUMBER_WORDS’, and the rule ‘MoneyCurrencyUnit’ then makes use of ‘AMOUNT_NUMBER’:

```
Phase: Number
Input: Token Lookup
Options: control = appelt

Macro: MILLION_BILLION
({Token.string == "m"}|
{Token.string == "million"}|
{Token.string == "b"}|
{Token.string == "billion"}|
{Token.string == "bn"}|
{Token.string == "k"}|
{Token.string == "K"}
)

Macro: NUMBER_WORDS
(
  (({Lookup.majorType == number}
    {Token.string == "-"})?
```

```

)*
  {Lookup.majorType == number}
  {Token.string == "and"}
)*
({Lookup.majorType == number}
  ({Token.string == "-"})?
)*
  {Lookup.majorType == number}
)

```

Macro: AMOUNT_NUMBER

```

(({Token.kind == number}
  (({Token.string == ","}|
    {Token.string == "."}
  )
  {Token.kind == number}
)*
|
(NUMBER_WORDS)
)
(MILLION_BILLION)?
)

```

Rule: MoneyCurrencyUnit

```

(
  (AMOUNT_NUMBER)
  ({Lookup.majorType == currency_unit})
)
:number -->
  :number.Money = {kind = "number", rule = "MoneyCurrencyUnit"}

```

8.1.7 Using Context

Context can be dealt with in the grammar rules in the following way. The pattern to be annotated is always enclosed by a set of round brackets. If preceding context is to be included in the rule, this is placed before this set of brackets. This context is described in exactly the same way as the pattern to be matched. If context following the pattern needs to be included, it is placed after the label given to the annotation. Context is used where a pattern should only be recognised if it occurs in a certain situation, but the context itself does not form part of the pattern to be annotated.

For example, the following rule for Time (assuming an appropriate macro for ‘year’) would mean that a year would only be recognised if it occurs preceded by the words ‘in’ or ‘by’:

Rule: YearContext1

```
{Token.string == "in"}|
  {Token.string == "by"}
)
(YEAR)
:date -->
  :date.Timex = {kind = "date", rule = "YearContext1"}
```

Similarly, the following rule (assuming an appropriate macro for ‘email’) would mean that an email address would only be recognised if it occurred inside angled brackets (which would not themselves form part of the entity):

Rule: Emailaddress1

```
{Token.string == '<'}
(
  (EMAIL)
)
:email
({Token.string == '>'})
-->
  :email.Address= {kind = "email", rule = "Emailaddress1"}
```

Also, it is possible to specify the constraint that one annotation must start at the same place as another. For example:

Rule: SurnameStartingWithDe

```
(
  {Token.string == "de",
   Lookup.majorType == "name",
   Lookup.minorType == "surname"}
):de
-->
  :de.Surname = {prefix = "de"}
```

This rule would match anywhere where a Token with string ‘de’ and a Lookup with majorType ‘name’ and minorType ‘surname’ start at the same offset in the text. Both the Lookup and Token annotations would be included in the `:de` binding, so the Surname annotation generated would span the longer of the two. Constraints on the same annotation type must be satisfied by a single annotation, so in this example there must be a single Lookup matching both the major and minor types – the rule would not match if there were two different lookups at the same location, one of them satisfying each constraint.

It is important to remember that context is consumed by the rule, so it cannot be reused in another rule within the same phase. So, for example, right context cannot be used as left context for another rule.

8.1.8 Multi-Constraint Statements

In the examples we have seen so far, most statements have contained only one constraint. For example, in this statement, the ‘category’ of ‘Token’ must equal ‘NNP’:

```
Rule: Unknown
Priority: 50
(
  {Token.category == NNP}
)
:unknown
-->
:unknown.Unknown = {kind = "PN", rule = Unknown}
```

However, it is equally acceptable to have multiple constraints in a statement. In this example, the ‘majorType’ of ‘Lookup’ must be ‘name’ **and** the ‘minorType’ must be ‘surname’:

```
Rule: Surname
(
  {Lookup.majorType == "name",
   Lookup.minorType == "surname"}
):surname
-->
:surname.Surname = {}
```

As we saw in Section 8.1.7, the constraints may refer to different annotations. In this example, in addition to the constraints on the ‘majorType’ and ‘minorType’ of ‘Lookup’, we also have a constraint on the ‘string’ of ‘Token’:

```
Rule: SurnameStartingWithDe
(
  {Token.string == "de",
   Lookup.majorType == "name",
   Lookup.minorType == "surname"}
):de
-->
:de.Surname = {prefix = "de"}
```

8.1.9 Negation

All the examples in the preceding sections involve constraints that require the presence of certain annotations to match. JAPE also supports ‘negative’ constraints which specify the *absence* of annotations. A negative constraint is signalled in the grammar by a ‘!’ character.

Negative constraints are generally used in combination with positive ones to constrain the locations at which the positive constraint can match. For example:

```
Rule: PossibleName
(
  {Token.orth == "upperInitial", !Lookup}
):name
-->
:name.PossibleName = {}
```

This rule would match any uppercase-initial Token, but only where there is no Lookup annotation starting at the same location. The general rule is that a negative constraint matches at any location where the corresponding positive constraint would *not* match. Negative constraints do not contribute any annotations to the bindings - in the example above, the `:name` binding would contain only the Token annotation. The exception to this is when a negative constraint is used alone, without any positive constraints in the combination. In this case it binds *all* the annotations at the match position that do not match the constraint. Thus, `{!Lookup}` would bind all the annotations starting at this location except Lookups. In most cases, negative constraints should only be used in combination with positive ones.

Any constraint can be negated, for example:

```
Rule: SurnameNotStartingWithDe
(
  {Surname, !Token.string =~ "[Dd]e"}
):name
-->
:name.NotDe = {}
```

This would match any Surname annotation that does not start at the same place as a Token with the string ‘de’ or ‘De’. Note that this is subtly different from `{Surname, Token.string ! =~ "[Dd]e"}`, as the second form requires a Token annotation to be present, whereas the first form (`!Token...`) will match if there is no Token annotation at all at this location.²

²In the Montreal transducer, the two forms were equivalent

Although JAPE provides an operator to look for the absence of a single annotation type, there is no support for a general negative operator to prevent a rule from firing if a particular *sequence* of annotations is found. One solution to this is to create a ‘negative rule’ which has higher priority than the matching ‘positive rule’. The style of matching must be Appelt for this to work. To create a negative rule, simply state on the LHS of the rule the pattern that should NOT be matched, and on the RHS do nothing. In this way, the positive rule cannot be fired if the negative pattern matches, and vice versa, which has the same end result as using a negative operator. A useful variation for developers is to create a dummy annotation on the RHS of the negative rule, rather than to do nothing, and to give the dummy annotation a rule feature. In this way, it is obvious that the negative rule has fired. Alternatively, use Java code on the RHS to print a message when the rule fires. An example of a matching negative and positive rule follows. Here, we want a rule which matches a surname followed by a comma and a set of initials. But we want to specify that the initials shouldn’t have the POS category PRP (personal pronoun). So we specify a negative rule that will fire if the PRP category exists, thereby preventing the positive rule from firing.

```
Rule: NotPersonReverse
Priority: 20
// we don't want to match 'Jones, I'
(
  {Token.category == NNP}
  {Token.string == ","}
  {Token.category == PRP}
)
:foo
-->
{}
```

```
Rule: PersonReverse
Priority: 5
// we want to match 'Jones, F.W.'
(
  {Token.category == NNP}
  {Token.string == ","}
  (INITIALS)?
)
:person -->
```

8.1.10 Escaping Special Characters

To specify a single or double quote as a string, precede it with a backslash, e.g.

```
{Token.string=="\""}}
```

will match a double quote. For other special characters, such as '\$', enclose it in double quotes, e.g.

```
{Token.category == "PRP\$"}
```

8.2 LHS Operators in Detail

This section gives more detail on the behaviour of the matching operators used on the left-hand side of JAPE rules.

8.2.1 Compositional Operators

Compositional operators are used to combine matching constructions in the manner intended. Union and Kleene operators are available, as is range notation.

Union and Kleene Operators

The following union and Kleene operators are available:

- | - or
- * - zero or more occurrences
- ? - zero or one occurrences
- + - one or more occurrences

In the following example, you can see the '|' and '?' operators being used:

```
Rule: LocOrganization
```

```
Priority: 50
```

```
(  
  ({Lookup.majorType == location} |  
   {Lookup.majorType == country_adj})  
  {Lookup.majorType == organization}
```

```
({Lookup.majorType == organization})?
)
:orgName -->
  :orgName.TempOrganization = {kind = "orgName", rule=LocOrganization}
```

Range Notation

A range notation can also be added. e.g.

```
({Token}) [1,3]
```

matches one to three Tokens in a row.

```
({Token.kind == number}) [3]
```

matches exactly 3 number Tokens in a row.

8.2.2 Matching Operators

Matching operators are used to specify how matching must take place between a specification and an annotation in the document. Equality ('==' and '!=') and comparison ('<', '<=', '>=' and '>') operators can be used, as can regular expression matching and contextual operators ('contains' and 'within').

Equality Operators

The equality operators are '==' and '!='. The basic operator in JAPE is equality. `{Lookup.majorType == "person"}` matches a Lookup annotation whose majorType feature has the value 'person'. Similarly `{Lookup.majorType != "person"}` would match any Lookup whose majorType feature does *not* have the value 'person'. If a feature is missing it is treated as if it had an empty string as its value, so this would also match a Lookup annotation that did not have a majorType feature at all.

Certain type coercions are performed:

- If the constraint's attribute is a string, it is compared with the annotation feature value using string equality (`String.equals()`).
- If the constraint's attribute is an integer it is treated as a `java.lang.Long`. If the annotation feature value is also a `Long`, or is a string that can be parsed as a `Long`, then it is compared using `Long.equals()`.

- If the constraint's attribute is a floating-point number it is treated as a `java.lang.Double`. If the annotation feature value is also a `Double`, or is a string that can be parsed as a `Double`, then it is compared using `Double.equals()`.
- If the constraint's attribute is `true` or `false` (without quotes) it is treated as a `java.lang.Boolean`. If the annotation feature value is also a `Boolean`, or is a string that can be parsed as a `Boolean`, then it is compared using `Boolean.equals()`.

The `!=` operator matches exactly when `==` doesn't.

Comparison Operators

The comparison operators are '`<`', '`<=`', '`>=`' and '`>`'. Comparison operators have their expected meanings, for example `{Token.length > 3}` matches a `Token` annotation whose length attribute is an integer greater than 3. The behaviour of the operators depends on the type of the constraint's attribute:

- If the constraint's attribute is a string it is compared with the annotation feature value using Unicode-lexicographic order (see `String.compareTo()`).
- If the constraint's attribute is an integer it is treated as a `java.lang.Long`. If the annotation feature value is also a `Long`, or is a string that can be parsed as a `Long`, then it is compared using `Long.compareTo()`.
- If the constraint's attribute is a floating-point number it is treated as a `java.lang.Double`. If the annotation feature value is also a `Double`, or is a string that can be parsed as a `Double`, then it is compared using `Double.compareTo()`.

Regular Expression Operators

The regular expression operators are '`=~`', '`==~`', '`!~`' and '`!=~`'. These operators match regular expressions. `{Token.string =~ "[Dd]ogs"}` matches a `Token` annotation whose string feature contains a substring that matches the regular expression `[Dd]ogs`, using `!~` would match if the feature value does not contain a substring that matches the regular expression. The `==~` and `!=~` operators are like `=~` and `!~` respectively, but require that the *whole* value match (or not match) the regular expression³. As with `==`, missing features are treated as if they had the empty string as their value, so the constraint `{Identifier.name ==~ "(?i)[aeiou]*"}` would match an `Identifier` annotation which does not have a name feature, as well as any whose name contains only vowels.

³This syntax will be familiar to Groovy users.

The matching uses the standard Java regular expression library, so full details of the pattern syntax can be found in the JavaDoc documentation for `java.util.regex.Pattern`. There are a few specific points to note:

- To enable flags such as case-insensitive matching you can use the *(?flags)* notation. See the Pattern JavaDocs for details.
- If you need to include a double quote character in a regular expression you must precede it with a backslash, otherwise JAPE will give a syntax error. Quoted strings in JAPE grammars also convert the sequences `\n`, `\r` and `\t` to the characters newline (U+000A), carriage return (U+000D) and tab (U+0009) respectively, but these characters can match literally in regular expressions so it does not make any difference to the result in most cases.⁴

Contextual Operators

The contextual Operators are ‘contains’ and ‘within’. These operators match annotations within the context of other annotations.

- `contains` - Written as `{X contains Y}`, returns true if an annotation of type X completely contains an annotation of type Y.
- `within` - Written as `{X within Y}`, returns true if an annotation of type X is completely covered by an annotation of type Y.

For either operator, the right-hand value (Y in the above examples) can be a full constraint itself. For example `{X contains {Y.foo==bar}}` is also accepted. The operators can be used in a multi-constraint statement (see Section 8.1.8) just like any of the traditional ones, so `{X.f1 != "something", X contains {Y.foo==bar}}` is valid.

Custom Operators

It is possible to add additional custom operators without modifying the JAPE language. There are new init-time parameters to Transducer so that additional annotation ‘meta-property’ accessors and custom operators can be referenced at runtime. To add a custom operator, write a class that implements `gate.jape.constraint.ConstraintPredicate`, make the class available to GATE (either by putting the class in a JAR file in the `lib` directory or

⁴However this does mean that it is not possible to include an n, r or t character after a backslash in a JAPE quoted string, or to have a backslash as the last character of your regular expression. Workarounds include placing the backslash in a character class (`[\\]`) or enabling the `(?x)` flag, which allows you to put whitespace between the backslash and the offending character without changing the meaning of the pattern.

by putting the class in a plugin and loading the plugin), and then list that class name for the Transducer's 'operators' property. Similarly, to add a custom 'meta-property' accessor, write a class that implements `gate.jape.constraint.AnnotationAccessor`, and then list that class name in the Transducer's 'annotationAccessors' property.

8.3 The Right-Hand Side

The RHS of the rule contains information about the annotation to be created/manipulated. Information about the text span to be annotated is transferred from the LHS of the rule using the label just described, and annotated with the entity type (which follows it). Finally, attributes and their corresponding values are added to the annotation. Alternatively, the RHS of the rule can contain Java code to create or manipulate annotations, see Section 8.6.

8.3.1 A Simple Example

In the simple example below, the pattern described will be awarded an annotation of type 'Enamex' (because it is an entity name). This annotation will have the attribute 'kind', with value 'location', and the attribute 'rule', with value 'GazLocation'. (The purpose of the 'rule' attribute is simply to ease the process of manual rule validation).

```
Rule: GazLocation
(
{Lookup.majorType == location}
)
:location -->
  :location.Enamex = {kind="location", rule=GazLocation}
```

8.3.2 Copying Feature Values from the LHS to the RHS

JAPE provides limited support for copying annotation feature values from the left to the right hand side of a rule, for example:

```
Rule: LocationType
(
  {Lookup.majorType == location}
):loc
-->
  :loc.Location = {rule = "LocationType", type = :loc.Lookup.minorType}
```

This will set the ‘type’ feature of the generated location to the value of the ‘minorType’ feature from the ‘Lookup’ annotation bound to the loc label. If the Lookup has no minorType, the Location will have no ‘type’ feature. The behaviour of `newFeat = :bind.Type.oldFeat` is:

- Find all the annotations of type `Type` from the left hand side binding `bind`.
- Find one of them that has a non-null value for its `oldFeat` feature (if there is more than one, which one is chosen is up to the JAPE implementation).
- If such a value exists, set the `newFeat` feature of our newly created annotation to this value.
- If no such non-null value exists, do not set the `newFeat` feature at all.

Notice that the behaviour is *deliberately underspecified* if there is more than one `Type` annotation in `bind`. If you need more control, or if you want to copy several feature values from the same left hand side annotation, you should consider using Java code on the right hand side of your rule (see Section 8.6).

8.3.3 RHS Macros

Macros, first introduced in the context of the left-hand side (Section 8.1.6) can also be used on the RHS of rules. In this case, the label (which matches the label on the LHS of the rule) should be included in the macro. Below we give an example of using a macro on the RHS:

```
Macro: UNDERScores_OKAY          // separate
:match                          // lines
{
    AnnotationSet matchedAnns = bindings.get("match");

    int begOffset = matchedAnns.firstNode().getOffset().intValue();
    int endOffset = matchedAnns.lastNode().getOffset().intValue();
    String mydocContent = doc.getContent().toString();
    String matchedString = mydocContent.substring(begOffset, endOffset);

    FeatureMap newFeatures = Factory.newFeatureMap();

    if(matchedString.equals("Spanish")) {
        newFeatures.put("myrule", "Lower");
    }
    else {
        newFeatures.put("myrule", "Upper");
    }
}
```

```

    }

    newFeatures.put("quality", "1");
    annotations.add(matchedAnns.firstNode(), matchedAnns.lastNode(),
                    "Spanish_mark", newFeatures);
}

```

Rule: Lower

```

(
  ({Token.string == "Spanish"})
:match)-->UNDERScores_OKAY // no label here, only macro name

```

Rule: Upper

```

(
  ({Token.string == "SPANISH"})
:match)-->UNDERScores_OKAY // no label here, only macro name

```

8.4 Use of Priority

Each grammar has one of 5 possible control styles: ‘brill’, ‘all’, ‘first’, ‘once’ and ‘appelt’. This is specified at the beginning of the grammar.

The Brill style means that when more than one rule matches the same region of the document, they are all fired. The result of this is that a segment of text could be allocated more than one entity type, and that no priority ordering is necessary. Brill will execute all matching rules starting from a given position and will advance and continue matching from the position in the document where the longest match finishes.

The ‘all’ style is similar to Brill, in that it will also execute all matching rules, but the matching will continue from the next offset to the current one.

For example, where [] are annotations of type Ann

```
[aaa[bbb]] [ccc[ddd]]
```

then a rule matching {Ann} and creating {Ann-2} for the same spans will generate:

```
BRILL: [aaabbb] [cccddd]
ALL: [aaa[bbb]] [ccc[ddd]]
```

With the ‘first’ style, a rule fires for the first match that’s found. This makes it inappropriate for rules that end in ‘+’ or ‘?’ or ‘*’. Once a match is found the rule is fired; it does not attempt to get a longer match (as the other two styles do).

With the ‘once’ style, once a rule has fired, the whole JAPE phase exits after the first match.

With the appelt style, only one rule can be fired for the same region of text, according to a set of priority rules. Priority operates in the following way.

1. From all the rules that match a region of the document starting at some point X, the one which matches the longest region is fired.
2. If more than one rule matches the same region, the one with the highest priority is fired
3. If there is more than one rule with the same priority, the one defined earlier in the grammar is fired.

An optional priority declaration is associated with each rule, which should be a positive integer. The higher the number, the greater the priority. By default (if the priority declaration is missing) all rules have the priority -1 (i.e. the lowest priority).

For example, the following two rules for location could potentially match the same text.

```
Rule: Location1
Priority: 25

(
  ({Lookup.majorType == loc_key, Lookup.minorType == pre}
   {SpaceToken})?
  {Lookup.majorType == location}
  ({SpaceToken}
   {Lookup.majorType == loc_key, Lookup.minorType == post})?
)
:locName -->
  :locName.Location = {kind = "location", rule = "Location1"}
```

```
Rule: GazLocation
Priority: 20
(
  ({Lookup.majorType == location}):location
)
--> :location.Name = {kind = "location", rule=GazLocation}
```

Assume we have the text ‘China sea’, that ‘China’ is defined in the gazetteer as ‘location’, and that sea is defined as a ‘loc_key’ of type ‘post’. In this case, rule Location1 would apply, because it matches a longer region of text starting at the same point (‘China sea’, as opposed to just ‘China’). Now assume we just have the text ‘China’. In this case, both rules could be fired, but the priority for Location1 is highest, so it will take precedence. In this case, since both rules produce the same annotation, so it is not so important which rule is fired, but this is not always the case.

One important point of which to be aware is that prioritisation only operates within a single grammar. Although we could make priority global by having all the rules in a single grammar, this is not ideal due to other considerations. Instead, we currently combine all the rules for each entity type in a single grammar. An index file (main.jape) is used to define which grammars should be used, and in which order they should be fired.

Note also that depending on the control style, firing a rule may ‘consume’ that part of the text, making it unavailable to be matched by other rules. This can be a problem for example if one rule uses context to make it more specific, and that context is then missed by later rules, having been consumed due to use of for example the ‘Brill’ control style. ‘All’, on the other hand, would allow it to be matched.

Using priority to resolve ambiguity

If the Appelt style of matching is selected, rule priority operates in the following way.

1. Length of rule – a rule matching a longer pattern will fire first.
2. Explicit priority declaration. Use the optional Priority function to assign a ranking. The higher the number, the higher the priority. If no priority is stated, the default is -1.
3. Order of rules. In the case where the above two factors do not distinguish between two rules, the order in which the rules are stated applies. Rules stated first have higher priority.

Because priority can only operate within a single grammar, this can be a problem for dealing with ambiguity issues. One solution to this is to create a temporary set of annotations in initial grammars, and then manipulate this temporary set in one or more later phases (for example, by converting temporary annotations from different phases into permanent annotations in a single final phase). See the default set of grammars for an example of this.

If two possible ways of matching are found for the same text string, a conflict can arise. Normally this is handled by the priority mechanism (test length, rule priority and finally rule precedence). If all these are equal, Jape will simply choose a match at random and fire it. This leads to non-deterministic behaviour, which should be avoided.

8.5 Using Phases Sequentially

A JAPE grammar consists of a set of sequential phases. The list of phases is specified (in the order in which they are to be run) in a file, conventionally named `main.jape`. When loading the grammar into GATE, it is only necessary to load this main file – the phases will then be loaded automatically. It is, however, possible to omit this main file, and just load the phases individually, but this is much more time-consuming. The grammar phases do not need to be located in the same directory as the main file, but if they are not, the relative path should be specified for each phase.

One of the main reasons for using a sequence of phases is that a pattern can only be used once in each phase, but it can be reused in a later phase. Combined with the fact that priority can only operate within a single grammar, this can be exploited to help deal with ambiguity issues.

The solution currently adopted is to write a grammar phase for each annotation type, or for each combination of similar annotation types, and to create temporary annotations. These temporary annotations are accessed by later grammar phases, and can be manipulated as necessary to resolve ambiguity or to merge consecutive annotations. The temporary annotations can either be removed later, or left and simply ignored.

Generally, annotations about which we are more certain are created earlier on. Annotations which are more dubious may be created temporarily, and then manipulated by later phases as more information becomes available.

An annotation generated in one phase can be referred to in a later phase, in exactly the same way as any other kind of annotation (by specifying the name of the annotation within curly braces). The features and values can be referred to or omitted, as with all other annotations. Make sure that if the `Input` specification is used in the grammar, that the annotation to be referred to is included in the list.

8.6 Using Java Code on the RHS

The RHS of a JAPE rule can consist of any Java code. This is useful for removing temporary annotations and for percolating and manipulating features from previous annotations. In the example below

The first rule below shows a rule which matches a first person name, e.g. ‘Fred’, and adds a gender feature depending on the value of the `minorType` from the gazetteer list in which the name was found. We first get the bindings associated with the person label (i.e. the `Lookup` annotation). We then create a new annotation called ‘`personAnn`’ which contains this annotation, and create a new `FeatureMap` to enable us to add features. Then we get the `minorType` features (and its value) from the `personAnn` annotation (in this case, the feature

will be ‘gender’ and the value will be ‘male’), and add this value to a new feature called ‘gender’. We create another feature ‘rule’ with value ‘FirstName’. Finally, we add all the features to a new annotation ‘FirstPerson’ which attaches to the same nodes as the original ‘person’ binding.

Note that inputAS and outputAS represent the input and output annotation set. Normally, these would be the same (by default when using ANNIE, these will be the ‘Default’ annotation set). Since the user is at liberty to change the input and output annotation sets in the parameters of the JAPE transducer at runtime, it cannot be guaranteed that the input and output annotation sets will be the same, and therefore we must specify the annotation set we are referring to.

Rule: FirstName

```
(
  {Lookup.majorType == person_first}
):person
-->
{
  AnnotationSet person = bindings.get("person");
  Annotation personAnn = person.iterator().next();
  FeatureMap features = Factory.newFeatureMap();
  features.put("gender", personAnn.getFeatures().get("minorType"));
  features.put("rule", "FirstName");
  outputAS.add(person.firstNode(), person.lastNode(), "FirstPerson",
  features);
}
```

The second rule (contained in a subsequent grammar phase) makes use of annotations produced by the first rule described above. Instead of percolating the minorType from the annotation produced by the gazetteer lookup, this time it percolates the feature from the annotation produced by the previous grammar rule. So here it gets the ‘gender’ feature value from the ‘FirstPerson’ annotation, and adds it to a new feature (again called ‘gender’ for convenience), which is added to the new annotation (in outputAS) ‘TempPerson’. At the end of this rule, the existing input annotations (from inputAS) are removed because they are no longer needed. Note that in the previous rule, the existing annotations were not removed, because it is possible they might be needed later on in another grammar phase.

Rule: GazPersonFirst

```
(
  {FirstPerson}
)
:person
```

```
-->
{
AnnotationSet person = bindings.get("person");
Annotation personAnn = person.iterator().next();
FeatureMap features = Factory.newFeatureMap();

features.put("gender", personAnn.getFeatures().get("gender"));
features.put("rule", "GazPersonFirst");
outputAS.add(person.firstChild(), person.lastChild(), "TempPerson",
features);
inputAS.removeAll(person);
}
```

You can combine Java blocks and normal assignments (separating each block or assignment from the next with a comma), so the above RHS could be more simply expressed as

```
-->
:person.TempPerson = { gender = :person.FirstPerson.gender,
                      rule = "GazPersonFirst" },
{
  inputAS.removeAll(bindings.get("person"));
}
```

8.6.1 A More Complex Example

The example below is more complicated, because both the title and the first name (if present) may have a gender feature. There is a possibility of conflict since some first names are ambiguous, or women are given male names (e.g. Charlie). Some titles are also ambiguous, such as 'Dr', in which case they are not marked with a gender feature. We therefore take the gender of the title in preference to the gender of the first name, if it is present. So, on the RHS, we first look for the gender of the title by getting all Title annotations which have a gender feature attached. If a gender feature is present, we add the value of this feature to a new gender feature on the Person annotation we are going to create. If no gender feature is present, we look for the gender of the first name by getting all firstPerson annotations which have a gender feature attached, and adding the value of this feature to a new gender feature on the Person annotation we are going to create. If there is no firstPerson annotation and the title has no gender information, then we simply create the Person annotation with no gender feature.

```
Rule: PersonTitle
Priority: 35
```



```

/* allows Mr. Jones, Mr Fred Jones etc. */

(
  (TITLE)
  (FIRSTNAME | FIRSTNAMEAMBIG | INITIALS2)*
  (PREFIX)?
  {Upper}
  ({Upper})?
  (PERSONENDING)?
)
:person -->
{
FeatureMap features = Factory.newFeatureMap();
AnnotationSet personSet = bindings.get("person");

// get all Title annotations that have a gender feature
HashSet fNameSet = new HashSet();
  fNameSet.add("gender");
  AnnotationSet personTitle = personSet.get("Title", fNameSet);

// if the gender feature exists
if (personTitle != null && personTitle.size()>0)
{
  Annotation personAnn = personTitle.iterator().next();
  features.put("gender", personAnn.getFeatures().get("gender"));
}
else
{
  // get all firstPerson annotations that have a gender feature
  AnnotationSet firstPerson = personSet.get("FirstPerson", fNameSet);

  if (firstPerson != null && firstPerson.size()>0)
  // create a new gender feature and add the value from firstPerson
  {
    Annotation personAnn = firstPerson.iterator().next();
    features.put("gender", personAnn.getFeatures().get("gender"));
  }
}

  // create some other features
  features.put("kind", "personName");
  features.put("rule", "PersonTitle");
  // create a Person annotation and add the features we've created
outputAS.add(personSet.firstChild(), personSet.lastChild(), "TempPerson",
features);

```

```
}

```

8.6.2 Adding a Feature to the Document

This is useful when using conditional controllers, where we only want to fire a particular resource under certain conditions. We first test the document to see whether it fulfils these conditions or not, and attach a feature to the document accordingly.

In the example below, we test whether the document contains an annotation of type ‘message’. In emails, there is often an annotation of this type (produced by the document format analysis when the document is loaded in GATE). Note that annotations produced by document format analysis are placed automatically in the ‘Original markups’ annotation set, so we must ensure that when running the processing resource containing this grammar that we specify the Original markups set as the input annotation set. It does not matter what we specify as the output annotation set, because the annotation we produce is going to be attached to the document and not to an output annotation set. In the example, if an annotation of type ‘message’ is found, we add the feature ‘genre’ with value ‘email’ to the document.

```
Rule: Email
Priority: 150

(
  {message}
)
-->
{
doc.getFeatures().put("genre", "email");
}
```

8.6.3 Finding the Tokens of a Matched Annotation

In this section we will demonstrate how by using Java on the right-hand side one can find all Token annotations that are covered by a matched annotation, e.g., a Person or an Organization. This is useful if one wants to transfer some information from the matched annotations to the tokens. For example, to add to the Tokens a feature indicating whether or not they are covered by a named entity annotation deduced by the rule-based system. This feature can then be given as a feature to a learning PR, e.g. the HMM. Similarly, one can add a feature to all tokens saying which rule in the rule based system did the match, the idea being that some rules might be more reliable than others. Finally, yet another useful feature might be the length of the coreference chain in which the matched entity is involved, if such exists.

The example below is one of the pre-processing JAPE grammars used by the HMM application. To inspect all JAPE grammars, see the `muse/applications/hmm` directory in the distribution.

Phase: NEInfo

Input: Token Organization Location Person

Options: control = appelt

Rule: NEInfo

Priority:100

```
({Organization} | {Person} | {Location}):entity
-->
{
  //get the annotation set
  AnnotationSet annSet = bindings.get("entity");

  //get the only annotation from the set
  Annotation entityAnn = annSet.iterator().next();

  AnnotationSet tokenAS = inputAS.get("Token",
    entityAnn.getStartNode().getOffset(),
    entityAnn.getEndNode().getOffset());
  List<Annotation> tokens = new ArrayList<Annotation>(tokenAS);
  //if no tokens to match, do nothing
  if (tokens.isEmpty())
    return;
  Collections.sort(tokens, new gate.util.OffsetComparator());

  Annotation curToken=null;
  for (int i=0; i < tokens.size(); i++) {
    curToken = tokens.get(i);
    String ruleInfo = (String) entityAnn.getFeatures().get("rule1");
    String NMRuleInfo = (String) entityAnn.getFeatures().get("NMRule");
    if ( ruleInfo != null) {
      curToken.getFeatures().put("rule_NE_kind", entityAnn.getType());
      curToken.getFeatures().put("NE_rule_id", ruleInfo);
    }
    else if (NMRuleInfo != null) {
      curToken.getFeatures().put("rule_NE_kind", entityAnn.getType());
      curToken.getFeatures().put("NE_rule_id", "orthomatcher");
    }
  }
}
```

```

    }
    else {
        curToken.getFeatures().put("rule_NE_kind", "None");
        curToken.getFeatures().put("NE_rule_id", "None");
    }
    List matchesList = (List) entityAnn.getFeatures().get("matches");
    if (matchesList != null) {
        if (matchesList.size() == 2)
            curToken.getFeatures().put("coref_chain_length", "2");
        else if (matchesList.size() > 2 && matchesList.size() < 5)
            curToken.getFeatures().put("coref_chain_length", "3-4");
        else
            curToken.getFeatures().put("coref_chain_length", "5-more");
    }
    else
        curToken.getFeatures().put("coref_chain_length", "0");
} //for
}

Rule: TokenNEInfo
Priority:10
({Token}):entity
-->
{
    //get the annotation set
    AnnotationSet annSet = bindings.get("entity");

    //get the only annotation from the set
    Annotation entityAnn = annSet.iterator().next();

    entityAnn.getFeatures().put("rule_NE_kind", "None");
    entityAnn.getFeatures().put("NE_rule_id", "None");
    entityAnn.getFeatures().put("coref_chain_length", "0");
}

```

8.6.4 Using Named Blocks

For the common case where a Java block refers just to the annotations from a single left-hand-side binding, JAPE provides a shorthand notation:

```
Rule: RemoveDoneFlag
```

```
(
  {Instance.flag == "done"}
):inst
-->
:inst{
  Annotation theInstance = instAnnots.iterator().next();
  theInstance.getFeatures().remove("flag");
}
```

This rule is equivalent to the following:

Rule: RemoveDoneFlag

```
(
  {Instance.flag == "done"}
):inst
-->
{
  AnnotationSet instAnnots = bindings.get("inst");
  if(instAnnots != null && instAnnots.size() != 0) {
    Annotation theInstance = instAnnots.iterator().next();
    theInstance.getFeatures().remove("flag");
  }
}
```

A label `:<label>` on a Java block creates a local variable `<label>Annots` within the Java block which is the `AnnotationSet` bound to the `<label>` label. Also, the Java code in the block is only executed if there is at least one annotation bound to the label, so you do not need to check this condition in your own code. Of course, if you need more flexibility, e.g. to perform some action in the case where the label is not bound, you will need to use an unlabelled block and perform the `bindings.get()` yourself.

8.6.5 Java RHS Overview

When a JAPE grammar is parsed, a Jape parser creates action classes for all Java RHSs in the grammar. (one action class per RHS) RHS Java code will be embedded as a body of the method `doit` and will work in context of this method. When a particular rule is fired, the method `doit` will be executed.

Method `doit` is specified by the interface `gate.jape.RhsAction`. Each action class implements this interface and is generated with the following template:

```

1  import java.io.*;
2  import java.util.*;
3  import gate.*;
4  import gate.jape.*;
5  import gate.creole.ontology.*;
6  import gate.annotation.*;
7  import gate.util.*;
8  class <AutogeneratedActionClassName>
9      implements java.io.Serializable, gate.jape.RhsAction {
10     public void doit(
11         gate.Document doc,
12         java.util.Map<java.lang.String, gate.AnnotationSet> bindings,
13         gate.AnnotationSet annotations,
14         gate.AnnotationSet inputAS,
15         gate.AnnotationSet outputAS,
16         gate.creole.ontology.Ontology ontology) throws JapeException {
17         // your RHS Java code will be embedded here ...
18     }
19 }

```

Method `doit` has the following parameters that can be used in RHS Java code:

- `gate.Document doc` - a document that is currently processed
- `java.util.Map<String, AnnotationSet> bindings` - a map of binding variables where a key is a (String) name of binding variable and value is (AnnotationSet) set of annotations corresponding to this binding variable⁵
- `gate.AnnotationSet annotations` - Do not use this (it's a synonym for `outputAS` that is still used in some grammars but is now deprecated).
- `gate.AnnotationSet inputAS` - input annotations
- `gate.AnnotationSet outputAS` - output annotations
- `gate.creole.ontology.Ontology ontology` - a GATE's transducer ontology

In your Java RHS you can use short names for all Java classes that are imported by the action class (plus Java classes from the packages that are imported by default according to JVM specification: `java.lang.*`, `java.math.*`). But you need to use fully qualified Java class names for all other classes. For example:

```

1  -->
2  {
3      // VALID line examples
4      AnnotationSet as = ...
5      InputStream is = ...

```

⁵Prior to GATE 5.2 this parameter was a plain Map without type parameters, which is why you will see a lot of now-unnecessary casts in existing JAPE grammars such as those in ANNIE.

```

6  java.util.logging.Logger myLogger =
7      java.util.logging.Logger.getLogger("JAPELogger");
8  java.sql.Statement stmt = ...
9
10 // INVALID line examples
11 Logger myLogger = Logger.getLogger("JapePhaseLogger");
12 Statement stmt = ...
13 }

```

In order to add additional Java `import` or `import static` statements to all Java RHS' of the rules in a JAPE grammar file, you can use the following code at the beginning of the JAPE file:

```

1  Imports: {
2  import java.util.logging.Logger;
3  import java.sql.*;
4  }

```

These import statements will be added to the default import statements for each action class generated for a RHS and the corresponding classes can be used in the RHS Java code without the need to use fully qualified names. A useful class to know about is `gate.Utils` (see the javadoc documentation for details), which provides static utility methods to simplify some common tasks that are frequently used in RHS Java code. Adding an `import static gate.Utils.*;` to the Imports block allows you to use these methods without any prefix, for example:

```

1  {
2  AnnotationSet lookups = bindings.get("lookup");
3  outputAS.add(start(lookups), end(lookups), "Person",
4      featureMap("text", stringFor(doc, lookups)));
5  }

```

You can do the same with your own utility classes — JAPE rules can import any class available to GATE, including classes defined in a plugin.

8.7 Optimising for Speed

The way in which grammars are designed can have a huge impact on the processing speed. Some simple tricks to keep the processing as fast as possible are:

- avoid the use of the `*` and `+` operators. Replace them with range queries where possible. For example, instead of

```
{Token}*
```

use

`(Token)[0,3]`

if you can predict that you won't need to recognise a string of Tokens longer than 3.

- avoid specifying unnecessary elements such as SpaceTokens where you can. To do this, use the Input specification at the beginning of the grammar to stipulate the annotations that need to be considered. If no Input specification is used, all annotations will be considered (so, for example, you cannot match two tokens separated by a space unless you specify the SpaceToken in the pattern). If, however, you specify Tokens but not SpaceTokens in the Input, SpaceTokens do not have to be mentioned in the pattern to be recognised. If, for example, there is only one rule in a phase that requires SpaceTokens to be specified, it may be judicious to move that rule to a separate phase where the SpaceToken can be specified as Input.
- avoid the shorthand syntax for copying feature values (`newFeat = :bind.Type.oldFeat`), particularly if you need to copy multiple features from the left to the right hand side of your rule.

8.8 Ontology Aware Grammar Transduction

GATE supports two different methods for ontology aware grammar transduction. Firstly it is possible to use the `ontology` feature both in grammars and annotations, while using the default transducer. Secondly it is possible to use an ontology aware transducer by passing an ontology language resource to one of the subsumes methods in `SimpleFeatureMapImpl`. This second strategy does not check for ontology features, which will make the writing of grammars easier, as there is no need to specify `ontology` when writing them. More information about the ontology-aware transducer can be found in Section 14.9.

8.9 Serializing JAPE Transducer

JAPE grammars are written as files with the extension `.jape`, which are parsed and compiled at run-time to execute them over the GATE document(s). Serialization of the JAPE Transducer adds the capability to serialize such grammar files and use them later to bootstrap new JAPE transducers, where they do not need the original JAPE grammar file. This allows people to distribute the serialized version of their grammars without disclosing the actual contents of their `jape` files. This is implemented as part of the JAPE Transducer PR. The following sections describe how to serialize and deserialize them.

8.9.1 How to Serialize?

Once an instance of a JAPE transducer is created, the option to serialize it appears in the context menu of that instance. The context menu can be activated by right clicking on the respective PR. Having done so, it asks for the file name where the serialized version of the respective JAPE grammar is stored.

8.9.2 How to Use the Serialized Grammar File?

The JAPE Transducer now also has an init-time parameter *binaryGrammarURL*, which appears as an optional parameter to the *grammarURL*. The User can use this parameter (i.e. *binaryGrammarURL*) to specify the serialized grammar file.

8.10 The JAPE Debugger

As of Version 5.1 the Jape debugger is not supported.

8.11 Notes for Montreal Transducer Users

In June 2008, the standard JAPE transducer implementation gained a number of features inspired by Luc Plamondon's 'Montreal Transducer', which was available as a GATE plugin for several years, and was made obsolete in Version 5.1. If you have existing Montreal Transducer grammars and want to update them to work with the standard JAPE implementation you should be aware of the following differences in behaviour:

- Quantifiers (*, + and ?) in the Montreal transducer are always greedy, but this is not necessarily the case in standard JAPE.
- The Montreal Transducer defines `{Type.feature != value}` to be the same as `{!Type.feature == value}` (and likewise the `!~` operator in terms of `=~`). In standard JAPE these constructs have different semantics. `{Type.feature != value}` will only match if there is a `Type` annotation whose `feature` feature does not have the given value, and if it matches it will bind the single `Type` annotation. `{!Type.feature == value}` will match if there is no `Type` annotation at a given place with this feature (including when there is no `Type` annotation at all), and if it matches it will bind every other annotation that starts at that location. If you have used `!=` in your Montreal grammars and want them to continue to behave the same way you must change them to use the prefix-! form instead (see Section 8.1.9).

- The `=~` operator in standard JAPE looks for regular expression matches anywhere within a feature value, whereas in the Montreal transducer it requires the whole string to match. To obtain the whole-string matching behaviour in standard JAPE, use the `==~` operator instead (see Section 8.2.2).

Chapter 9

ANNIC: ANNotations-In-Context

ANNIC (ANNotations-In-Context) is a full-featured annotation indexing and retrieval system. It is provided as part of an extension of the Serial Data-stores, called Searchable Serial Data-store (SSD).

ANNIC can index documents in any format supported by the GATE system (i.e., XML, HTML, RTF, e-mail, text, etc). Compared with other such query systems, it has additional features addressing issues such as extensive indexing of linguistic information associated with document content, independent of document format. It also allows indexing and extraction of information from overlapping annotations and features. Its advanced graphical user interface provides a graphical view of annotation markups over the text, along with an ability to build new queries interactively. In addition, ANNIC can be used as a first step in rule development for NLP systems as it enables the discovery and testing of patterns in corpora.

ANNIC is built on top of the Apache Lucene¹ – a high performance full-featured search engine implemented in Java, which supports indexing and search of large document collections. Our choice of IR engine is due to the customisability of Lucene. For more details on how Lucene was modified to meet the requirements of indexing and querying annotations, please refer to [Aswani *et al.* 05].

As explained earlier, SSD is an extension of the serial data-store. In addition to the persist location, SSD asks user to provide some more information (explained later) that it uses to index the documents. Once the SSD has been initiated, user can add/remove documents/-corpora to the SSD in a similar way it is done with other data-stores. When documents are added to the SSD, it automatically tries to index them. It updates the index whenever there is a change in any of the documents stored in the SSD and removes the document from the index if it is deleted from the SSD. Be warned that only the annotation sets, types and features initially provided during the SSD creation time, will be updated when adding/removing documents to the datastore.

¹<http://lucene.apache.org>

SSD has an advanced graphical interface that allows users to issue queries over the SSD. Below we explain the parameters required by SSD and how to instantiate it, how to use its graphical interface and how to use SSD programmatically.

9.1 Instantiating SSD

Steps:

1. In GATE Developer, right click on ‘Datastores’ and select ‘Create Datastore’.
2. From a drop-down list select ‘Lucene Based Searchable DataStore’.
3. Here, you will see an input window. Please provide these parameters:
 - (a) DataStore URL: Select an empty folder where the DS is created.
 - (b) Index Location: Select an empty folder. This is where the index will be created.
 - (c) Annotation Sets: Here, you can provide one or more annotation sets that you wish to index or exclude from being indexed. In order to be able to index the default annotation set, you must click on the edit list icon and add an empty field to the list. If there are no annotation sets provided, all the annotation sets in all documents are indexed.
 - (d) Base-Token Type: (e.g. Token or Key.Token) These are the basic tokens of any document. Your documents must have the annotations of Base-Token-Type in order to get indexed. These basic tokens are used for displaying contextual information while searching patterns in the corpus. In case of indexing more than one annotation set, user can specify the annotation set from which the tokens should be taken (e.g. Key.Token- annotations of type Token from the annotation set called Key). In case user does not provide any annotation set name (e.g. Token), the system searches in all the annotation sets to be indexed and the base-tokens from the first annotation set with the base token annotations are taken. Please note that the documents with no base-tokens are not indexed. However, if the ”create tokens automatically” option is selected, the SSD creates base-tokens automatically. Here, each string delimited with white space is considered as a token.
 - (e) Index Unit Type: (e.g. Sentence, Key.Sentence) This specifies the unit of Index. In other words, annotations lying within the boundaries of these annotations are indexed (e.g. in the case of “Sentences”, no annotations that are spanned across the boundaries of two sentences are considered for indexing). User can specify from which annotation set the index unit annotations should be considered. If user does not provide any annotation set, the SSD searches among all annotation sets for index units. If this field is left empty or SSD fails to locate index units, the entire document is considered as a single unit.




Figure 9.1: Searchable Serial Datastore Viewer.

(f) Features: Finally, users can specify the annotation types and features that should be indexed or excluded from being indexed. (e.g. SpaceToken and Split). If user wants to exclude only a specific feature of a specific annotation type, he/she can specify it using a '.' separator between the annotation type and its feature (e.g. Person.matches).

4. Click OK. If all parameters are OK, a new empty DS will be created.
5. Create an empty corpus and save it to the SSD.
6. Populate it with some documents. Each document added to the corpus and eventually to the SSD is indexed automatically. If the document does not have the required annotations, that document is skipped and not indexed.

9.2 Search GUI

9.2.1 Overview

Figure 9.1 shows the search GUI for a datastore. The top section contains a text area to write a query, lists to select the corpus and annotation set to search in, sliders to set the size of the results and context and icons to execute and clear the query.

The central section shows a graphical visualisation of stacked annotations and feature values for the result row selected in the bottom results table. You can also see the stack view configuration window where you define which annotation type and feature to display in the central section.

The bottom section contains the results table of the query, i.e. the text that matches the query with their left and right contexts. The bottom section contains also a tabbed pane of statistics.

9.2.2 Syntax of Queries

SSD enables you to formulate versatile queries using a subset of JAPE patterns. Below, we give the JAPE pattern clauses which can be used as SSD queries. Queries can also be a combination of one or more of the following pattern clauses.

1. String
2. {AnnotationType}
3. {AnnotationType == String}
4. {AnnotationType.feature == feature value}
5. {AnnotationType1, AnnotationType2.feature == featureValue}
6. {AnnotationType1.feature == featureValue, AnnotationType2.feature == featureValue}

JAPE patterns also support the | (OR) operator. For instance, {A} ({B} | {C}) is a pattern of two annotations where the first is an annotation of type A followed by the annotation of type either B or C.

ANNIC supports two operators, + and *, to specify the number of times a particular annotation or a sub pattern should appear in the main query pattern. Here, ({A})+n means one and up to n occurrences of annotation {A} and ({A})*n means zero or up to n occurrences of annotation {A}.

Below we explain the steps to search in SSD.

1. Double click on SSD. You will see an extra tab “Lucene DataStore Searcher”. Click on it to activate the searcher GUI.
2. Here you can specify a query to search in your SSD. The query here is a L.H.S. part of the JAPE grammar. Here are some examples:
 - (a) `{Person}` – This will return annotations of type Person from the SSD
 - (b) `{Token.string == “Microsoft”}` – This will return all occurrences of “Microsoft” from the SSD.
 - (c) `{Person}({Token})*2{Organization}` – Person followed by zero or up to two tokens followed by Organization.
 - (d) `{Token.orth==“upperInitial”, Organization}` – Token with feature orth with value set to “upperInitial” and which is also annotated as Organization.




Figure 9.2: Searchable Serial Datastore Viewer - Auto-completion.

9.2.3 Top Section

A text-area located in the top left part of the GUI is used to input a query. You can copy/cut/paste with `Control+C/X/V`, undo/redo your changes with `Control+Z/Y` as usual. To add a new line, use `Control+Enter` key combination.

Auto-completion shown on the figure 9.2 for annotation type is triggered when typing '{' or ',' and for feature when typing '.' after a valid annotation type. It shows only the annotation types and features related to the selected corpus and annotation set.

If you right-click on an expression it will automatically select the shortest valid enclosing brace and if you click on a selection it will propose you to add quantifiers for allowing the expression to appear zero, one or more times.

To execute the query, click on the magnifying glass icon, use `Enter` key or `Alt+Enter` key combination. To clear the query, click on the red X icon or use `Alt+Backspace` key combination.

It is possible to have more than one corpus, each containing a different set of documents, stored in a single data-store. ANNIC, by providing a drop down box with a list of stored corpora, also allows searching within a specific corpus. Similarly a document can have more than one annotation set indexed and therefore ANNIC also provides a drop down box with a list of indexed annotation sets for the selected corpus.

A large corpus can have many hits for a given query. This may take a long time to refresh the GUI and may create inconvenience while browsing through results. ANNIC therefore allows you to specify a number of results that you wish to retrieve at once and provides a way to iterate through next pages with the *Next Page of Results* button. Due to technical complexities, it is not possible to visit a previous page. To retrieve all the results at the same time, push the results slider to the right end.

9.2.4 Central Section

Annotation types and features to show can be configured from the stack view configuration window by clicking on the *Configure* button in the central section. You can also change the feature value to be displayed by double clicking on the annotation type name.

The central section shows coloured rectangles exactly below the spans of text where these annotations occur. If only an annotation type is displayed, the rectangle remains empty. When you hover the mouse over the rectangle, it shows all their features and values in a popup window. If an annotation type and a feature are displayed, the value of that feature is shown in the rectangle.

Shortcuts are expression that stand for an "AnnotationType.Feature" expression. For example, on the figure 9.1, the shortcut "POS" stands for the expression "Token.category".

When you double click on an annotation rectangle, the respective query expression is placed at the caret position in the query text area. If you have selected anything in the query text area, it gets replaced. You can also double click on a word on the first line to add it to the query.

9.2.5 Bottom Section

The table of results contains the text matched by the query, the contexts, the features from the matched annotations, the effective query, the document and annotation set names. You can sort a table column by clicking on its header.

You can remove a result from the results table or open the document containing it by right-clicking on a result in the results table.

ANNIC provides an *Export* button to export results into an HTML file. You can also select then copy/paste the table in your word processor or spreadsheet.

A statistics tabbed pane is displayed on the bottom-right. There is always a global statistics pane that list the count of the occurrences of all annotation types for the selected corpus and annotation set.

Statistics can be obtained for matched spans of the query in the results, with or without contexts, just by annotation type, an annotation type + feature or an annotation type + feature + value. A second pane contains the one item statistics that you can add by right-clicking on a non empty annotation rectangle or on the header of a row in the central section. You can sort a table column by clicking on its header.

9.3 Using SSD from GATE Embedded

9.3.1 How to instantiate a searchable datastore

```

1
2 // create an instance of datastore
3 LuceneDataStoreImpl ds = (LuceneDataStoreImpl)
4     Factory.createDataStore('gate.persist.LuceneDataStoreImpl',
5                             dsLocation);
6
7 // we need to set Indexer
8 Indexer indexer = new LuceneIndexer(new URL(indexLocation));
9
10 // set the parameters
11 Map parameters = new HashMap();
12
13 // specify the index url
14 parameters.put(Constants.INDEX_LOCATION_URL, new URL(indexLocation));
15
16 // specify the base token type
17 // and specify that the tokens should be created automatically
18 // if not found in the document
19 parameters.put(Constants.BASE_TOKEN_ANNOTATION_TYPE, 'Token');
20 parameters.put(Constants.CREATE_TOKENS_AUTOMATICALLY, new Boolean(true));
21
22 // specify the index unit type
23 parameters.put(Constants.INDEX_UNIT_ANNOTATION_TYPE, 'Sentence');
24
25 // specifying the annotation sets "Key" and "Default Annotation Set"
26 // to be indexed
27 List<String> setsToInclude = new ArrayList<String>();
28 setsToInclude.add("Key");
29 setsToInclude.add("<null>");
30 parameters.put(Constants.ANNOTATION_SETS_NAMES_TO_INCLUDE, setsToInclude);
31 parameters.put(Constants.ANNOTATION_SETS_NAMES_TO_EXCLUDE,
32     new ArrayList<String>());
33
34 // all features should be indexed
35 parameters.put(Constants.FEATURES_TO_INCLUDE, new ArrayList<String>());
36 parameters.put(Constants.FEATURES_TO_EXCLUDE, new ArrayList<String>());
37
38 // set the indexer
39 ds.setIndexer(indexer, parameters);
40
41 // set the searcher
42 ds.setSearcher(new LuceneSearcher());

```

9.3.2 How to search in this datastore

```

1
2 // obtain the searcher instance

```

```
3 Searcher searcher = ds.getSearcher();
4 Map parameters = new HashMap();
5
6 // obtain the url of index
7 String indexLocation =
8     new File(((URL) ds.getIndexer().getParameters()
9         .get(Constants.INDEX_LOCATION_URL)).getFile()).getAbsolutePath();
10 ArrayList indexLocations = new ArrayList();
11 indexLocations.add(indexLocation);
12
13
14 // corpus2SearchIn = mention corpus name that was indexed here.
15
16 // the annotation set to search in
17 String annotationSet2SearchIn = "Key";
18
19 // set the parameter
20 parameters.put(Constants.INDEX_LOCATIONS, indexLocations);
21 parameters.put(Constants.CORPUS_ID, corpus2SearchIn);
22 parameters.put(Constants.ANNOTATION_SET_ID, annotationSet);
23 parameters.put(Constants.CONTEXT_WINDOW, contextWindow);
24 parameters.put(Constants.NO_OF_PATTERNS, noOfPatterns);
25
26 // search
27 String query = "{Person}";
28 Hit[] hits = searcher.search(query, parameters);
```


Chapter 10

Performance Evaluation of Language Analysers

When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meager and unsatisfactory kind: it may be the beginning of knowledge, but you have scarcely in your thoughts advanced to the stage of science. (Kelvin)

Not everything that counts can be counted, and not everything that can be counted counts. (Einstein)

GATE provides a variety of tools for automatic evaluation. The Annotation Diff tool compares two annotation sets within a document. Corpus QA extends Annotation Diff to an entire corpus. The Corpus Benchmark tool also provides functionality for comparing annotation sets over an entire corpus. Additionally, two plugins cover similar functionality; one implements inter-annotator agreement, and the other, the balanced distance metric.

These tools are particularly useful not just as a final measure of performance, but as a tool to aid system development by tracking progress and evaluating the impact of changes as they are made. Applications include evaluating the success of a machine learning or language engineering application by comparing its results to a gold standard and also comparing annotations prepared by two human annotators to each other to ensure that the annotations are reliable.

This chapter begins by introducing the concepts and metrics relevant, before describing each of the tools in turn.

10.1 Metrics for Evaluation in Information Extraction

When we evaluate the performance of a processing resource such as tokeniser, POS tagger, or a whole application, we usually have a human-authored ‘gold standard’ against which to compare our software. However, it is not always easy or obvious what this gold standard should be, as different people may have different opinions about what is correct. Typically, we solve this problem by using more than one human annotator, and comparing their annotations. We do this by calculating inter-annotator agreement (IAA), also known as inter-rater reliability.

IAA can be used to assess how difficult a task is. This is based on the argument that if two humans cannot come to agreement on some annotation, it is unlikely that a computer could ever do the same annotation ‘correctly’. Thus, IAA can be used to find the ceiling for computer performance.

There are many possible metrics for reporting IAA, such as Cohen’s Kappa, prevalence, and bias [Eugenio & Glass 04]. Kappa is the best metric for IAA when all the annotators have identical exhaustive sets of questions on which they might agree or disagree. In other words, it is a classification task. This could be a task like ‘are these names male or female names’. However, sometimes there is disagreement about the set of questions, e.g. when the annotators themselves determine which text spans they ought to annotate, such as in named entity extraction. That could be a task like ‘read over this text and mark up all references to politics’. When annotators determine their own sets of questions, it is appropriate to use precision, recall, and F-measure to report IAA. Precision, recall and F-measure are also appropriate choices when assessing performance of an automated application against a trusted gold standard.

In this section, we will first introduce some relevant terms, before outlining Cohen’s Kappa and similar measures, in Section 10.1.2. We will then introduce precision, recall and F-measure in Section 10.1.3.

10.1.1 Annotation Relations

Before introducing the metrics we will use in this chapter, we will first outline the ways in which annotations can relate to each other. These ways of comparing annotations to each other are used to determine the counts that then go into calculating the metrics of interest. Consider a document with two annotation sets upon it. These annotation sets might for example be prepared by two human annotators, or alternatively, one set might be produced by an automated system and the other might be a trusted gold standard. We wish to assess the extent to which they agree. We begin by counting incidences of the following relations:

Coextensive Two annotations are coextensive if they hit the same span of text in a document. Basically, both their start and end offsets are equal.

Overlaps Two annotations overlap if they share a common span of text.

Compatible Two annotations are compatible if they are coextensive and if the features of one (usually the ones from the key) are included in the features of the other (usually the response).

Partially Compatible Two annotations are partially compatible if they overlap and if the features of one (usually the ones from the key) are included in the features of the other (response).

Missing This applies only to the key annotations. A key annotation is missing if either it is not coextensive or overlapping, or if one or more features are not included in the response annotation.

Spurious This applies only to the response annotations. A response annotation is spurious if either it is not coextensive or overlapping, or if one or more features from the key are not included in the response annotation.

10.1.2 Cohen's Kappa

The three commonly used IAA measures are *observed agreement*, *specific agreement*, and *Kappa* (κ) [Hripcsak & Heitjan 02]. Those measures can be calculated from a contingency table, which lists the numbers of instances of agreement and disagreement between two annotators on each category. To explain the IAA measures, a general contingency table for two categories *cat1* and *cat2* is shown in Table 10.1.

Table 10.1: Contingency table for two-category problem

	Annotator-2		
Annotator-1	cat1	cat2	marginal sum
cat1	a	b	a+b
cat2	c	d	c+d
marginal sum	a+c	b+d	a+b+c+d

Observed agreement is the portion of the instances on which the annotators agree. For the two annotators and two categories as shown in Table 10.1, it is defined as

$$A_o = \frac{a + d}{a + b + c + d} \quad (10.1)$$

The extension of the above formula to more than two categories is straightforward. The extension to more than two annotators is usually taken as the mean of the pair-wise agreements [Fleiss 75], which is the average agreement across all possible pairs of annotators. An alternative compares each annotator with the majority opinion of the others [Fleiss 75].

However, the observed agreement has two shortcomings. One is that a certain amount of agreement is expected by chance. The Kappa measure is a chance-corrected agreement. Another is that it sums up the agreement on all the categories, but the agreements on each category may differ. Hence the category specific agreement is needed.

Specific agreement quantifies the degree of agreement for each of the categories separately. For example, the specific agreement for the two categories list in Table 10.1 is the following, respectively,

$$A_{cat1} = \frac{2a}{2a + b + c}; \quad A_{cat2} = \frac{2d}{b + c + 2d} \quad (10.2)$$

Kappa is defined as the observed agreements A_o minus the agreement expected by chance A_e and is normalized as a number between -1 and 1.

$$\kappa = \frac{A_o - A_e}{1 - A_e} \quad (10.3)$$

$\kappa = 1$ means perfect agreements, $\kappa = 0$ means the agreement is equal to chance, $\kappa = -1$ means ‘perfect’ disagreement.

There are two different ways of computing the chance agreement A_e (for a detailed explanations about it see [Eugenio & Glass 04]; however, a quick outline will be given below). The Cohen’s Kappa is based on the individual distribution of each annotator, while the Siegel & Castellan’s Kappa is based on the assumption that all the annotators have the same distribution. The former is more informative than the latter and has been used widely.

Let us consider an example:

Table 10.2: Example contingency table for two-category problem

	Annotator-2		
Annotator-1	cat1	cat2	marginal sum
cat1	1	2	3
cat2	3	4	7
marginal sum	4	6	10

Cohen’s Kappa requires that the expected agreement be calculated as follows. Divide marginal sums by the total to get the portion of the instances that each annotator allocates to each category. Multiply annotator’s proportions together to get the likelihood of chance agreement, then total these figures. Table 10.3 gives a worked example.

The formula can easily be extended to more than two categories.

Siegel & Castellan’s Kappa is applicable for any number of annotators. Siegel & Castellan’s Kappa for two annotators is also known as Scott’s Pi (see [Lombard *et al.* 02]). It

Table 10.3: Calculating Expected Agreement for Cohen's Kappa

	Annotator-1	Annotator 2	Multiplied
cat1	3 / 10 = 0.3	4 / 10 = 0.4	0.12
cat2	7 / 10 = 0.7	6 / 10 = 0.6	0.42
Total			0.54

differs from Cohen's Kappa only in how the expected agreement is calculated. Table 10.4 shows a worked example. Annotator totals are added together and divided by the number of decisions to form joint proportions. These are then squared and totalled.

Table 10.4: Calculating Expected Agreement for Siegel & Castellan's Kappa (Scott's Pi)

	Ann-1	Ann-2	Sum	Joint Prop	JP-Squared
cat1	3	4	7	7/20	49/400=0.1225
cat2	7	6	13	13/20	169/400=0.4225
Total					218/400 = 0.545

The Kappa suffers from the prevalence problem which arises because imbalanced distribution of categories in the data increases A_e . The prevalence problem can be alleviated by reporting the positive and negative specified agreement on each category besides the Kappa [Hripcsak & Heitjan 02, Eugenio & Glass 04]. In addition, the so-called bias problem affects the Cohen's Kappa, but not S&C's. The bias problem arises as one annotator prefers one particular category more than another annotator. [Eugenio & Glass 04] advised to compute the S&C's Kappa and the specific agreements along with the Cohen's Kappa in order to handle these problems.

Despite the problem mentioned above, the Cohen's Kappa remains a popular IAA measure. Kappa can be used for more than two annotators based on pair-wise figures, e.g. the mean of all the pair-wise Kappa as an overall Kappa measure. The Cohen's Kappa can also be extended to the case of more than two annotators by using the following single formula [Davies & Fleiss 82]

$$\kappa = 1 - \frac{IJ^2 - \sum_i \sum_c Y_{ic}^2}{I(J(J-1) \sum_c (p_c(1-p_c)) + \sum_c \sum_j (p_{cj} - p_c)^2)} \quad (10.4)$$

Where I and J are the number of instances and annotators, respectively; Y_{ic} is the number of annotators who assigns the category c to the instance I ; p_{cj} is the probability of the annotator j assigning category c ; p_c is the probability of assigning category by all annotators (i.e. averaging p_{cj} over all annotators).

The Krippendorff's alpha, another variant of Kappa, differs only slightly from the S&C's Kappa on nominal category problem (see [Carletta 96, Eugenio & Glass 04]).

However, note that the Kappa (and the observed agreement) is not applicable to some tasks. Named entity annotation is one such task [Hripcsak & Rothschild 05]. In the named

entity annotation task, annotators are given some text and are asked to annotate some named entities (and possibly their categories) in the text. Different annotators may annotate different instances of the named entity. So, if one annotator annotates one named entity in the text but another annotator does not annotate it, then that named entity is a non-entity for the latter. However, generally the non-entity in the text is not a well-defined term, e.g. we don't know how many words should be contained in the non-entity. On the other hand, if we want to compute Kappa for named entity annotation, we need the non-entities. This is why people don't compute Kappa for the named entity task.

10.1.3 Precision, Recall, F-Measure

Much of the research in IE in the last decade has been connected with the MUC competitions, and so it is unsurprising that the MUC evaluation metrics of precision, recall and F-measure [Chinchor 92] also tend to be used, along with slight variations. These metrics have a very long-standing tradition in the field of IR [van Rijsbergen 79] (see also [Manning & Schütze 99, Frakes & Baeza-Yates 92]).

Precision measures the number of correctly identified items as a percentage of the number of items identified. In other words, it measures how many of the items that the system identified were actually correct, regardless of whether it also failed to retrieve correct items. The higher the precision, the better the system is at ensuring that what is identified is correct.

Error rate is the inverse of precision, and measures the number of incorrectly identified items as a percentage of the items identified. It is sometimes used as an alternative to precision.

Recall measures the number of correctly identified items as a percentage of the total number of correct items. In other words, it measures how many of the items that should have been identified actually were identified, regardless of how many spurious identifications were made. The higher the recall rate, the better the system is at not missing correct items.

Clearly, there must be a tradeoff between precision and recall, for a system can easily be made to achieve 100% precision by identifying nothing (and so making no mistakes in what it identifies), or 100% recall by identifying everything (and so not missing anything). The **F-measure** [van Rijsbergen 79] is often used in conjunction with Precision and Recall, as a weighted average of the two. **False positives** are a useful metric when dealing with a wide variety of text types, because it is not dependent on *relative document richness* in the same way that precision is. By this we mean the relative number of entities of each type to be found in a set of documents.

When comparing different systems on the same document set, relative document richness is unimportant, because it is equal for all systems. When comparing a single system's performance on different documents, however, it is much more crucial, because if a particular

document type has a significantly different number of any type of entity, the results for that entity type can become skewed. Compare the impact on precision of one error where the total number of correct entities = 1, and one error where the total = 100. Assuming the document length is the same, then the false positive score for each text, on the other hand, should be identical.

Common metrics for evaluation of IE systems are defined as follows:

$$Precision = \frac{Correct + 1/2Partial}{Correct + Spurious + Partial} \quad (10.5)$$

$$Recall = \frac{Correct + 1/2Partial}{Correct + Missing + Partial} \quad (10.6)$$

$$F - measure = \frac{(\beta^2 + 1)P * R}{(\beta^2 R) + P} \quad (10.7)$$

where β reflects the weighting of P vs. R. If β is set to 1, the two are weighted equally.

$$FalsePositive = \frac{Spurious}{c} \quad (10.8)$$

where c is some constant independent from document richness, e.g. the number of tokens or sentences in the document.

Note that we consider annotations to be partially correct if the entity type is correct and the spans are overlapping but not identical. Partially correct responses are normally allocated a half weight.

10.1.4 Macro and Micro Averaging

Where precision, recall and f-measure are calculated over a corpus, there are options in terms of how document statistics are combined.

- Micro averaging essentially treats the corpus as one large document. Correct, spurious and missing counts span the entire corpus, and precision, recall and f-measure are calculated accordingly.
- Macro averaging calculates precision, recall and f-measure on a per document basis, and then averages the results.

The method of choice depends on the priorities of the case in question. Macro averaging tends to increase the importance of shorter documents.

It is also possible to calculate a macro average across annotation types; that is to say, precision, recall and f-measure are calculated separately for each annotation type and the results then averaged.

10.2 The Annotation Diff Tool

The Annotation Diff tool enables two sets of annotations in one or two documents to be compared, in order either to compare a system-annotated text with a reference (hand-annotated) text, or to compare the output of two different versions of the system (or two different systems). For each annotation type, figures are generated for precision, recall, F-measure. Each of these can be calculated according to 3 different criteria - strict, lenient and average. The reason for this is to deal with partially correct responses in different ways.

- The Strict measure considers all partially correct responses as incorrect (spurious).
- The Lenient measure considers all partially correct responses as correct.
- The Average measure allocates a half weight to partially correct responses (i.e. it takes the average of strict and lenient).

It can be accessed both from GATE Developer and from GATE Embedded. Annotation Diff compares sets of annotations with the same type. When performing the comparison, the annotation offsets and their features will be taken into consideration. and after that, the comparison process is triggered. Figure 10.1 shows the Annotation Diff window.

All annotations from the key set are compared with the ones from the response set, and those found to have the same start and end offsets are displayed on the same line in the table. Then, the Annotation Diff evaluates if the features of each annotation from the response set subsume those features from the key set, as specified by the features names you provide.

In order to create a gold standard set from two sets you need to show the 'Adjudication' panel at the bottom. It will insert two checkboxes columns in the central table. Tick boxes in the 'K(ey)' and 'R(espone)' then input a Target set in the text field and use the 'Copy selection to target' button to copy all annotations selected to the target annotation set. There is a context menu for the checkboxes to tick them quickly.

To use the annotation diff tool, see Section 10.2.1. To compare more than two annotation sets, see Section 3.4.3.

Key doc: EP-1016726-A1.x... Key set: annotator1 Type: Measurement Weight: 1.00 Compare

Resp. doc: EP-1016726-A1.x... Resp. set: annotator2 Features: all some none 1.00

Start	End	Key	Features	K	R	Start	End	Re
26517	26532	2.and+0.25mg/ml	{rule=measurement.Me...fore=MgCl, conj=and}	<input type="checkbox"/>	<input type="checkbox"/>	26517	26532	2.and+0.25mg/ml
32763	32769	70-80%	{conj=, type=interv...urement.MeasureSpan}	<input checked="" type="checkbox"/>	<input type="checkbox"/>	32763	32769	70-80%
37992	38003	10'.and-15'	{conj=and, type=inte...urement.MeasureSpan}	<input type="checkbox"/>	<input type="checkbox"/>	37992	38003	10'.and-15'
25580	25593	32°C.and+39°C	{rule=measurement.Me...before=at, conj=and}	<input type="checkbox"/>	<input type="checkbox"/>	25580	25593	32°C.and+39°C
31576	31578	cm	{dimension=[length],...afepat;centi_metre}}	<input type="checkbox"/>	<input checked="" type="checkbox"/>	31576	31580	cm-2
25093	25097	cm-2	{dimension=[length],...afepat;centi_metre}}	<input type="checkbox"/>	<input type="checkbox"/>			
57200	57211	mm-2./field	{dimension=[length],...afepat;milli-meter}}	<input type="checkbox"/>	<input type="checkbox"/>			
40226	40227	3	{rule=measurement.Si...e, type=scalarValue}	<input type="checkbox"/>	<input type="checkbox"/>			
30461	30462	2	{rule=measurement.Me...n, type=scalarValue}	<input type="checkbox"/>	<input type="checkbox"/>			
35865	35866	'	{dimension=[length],...safepat;arc_minute}}	<input type="checkbox"/>	<input type="checkbox"/>			
45397	45398	5	{rule=measurement.Si...e, type=scalarValue}	<input type="checkbox"/>	<input type="checkbox"/>			
				<input type="checkbox"/>	<input checked="" type="checkbox"/>	33474	33480	per ml

Target set: consensus Copy selection to target set

2 annotations copied to consensus and 2 hidden Show document Export to HTML

Statistics Adjudication

Figure 10.1: Annotation diff window with the parameters at the top, the comparison table in the center and the adjudication panel at the bottom.

10.2.1 Performing Evaluation with the Annotation Diff Tool

The Annotation Diff tool is activated by selecting it from the Tools menu at the top of the GATE Developer window. It will appear in a new window. Select the key and response documents to be used (note that both must have been previously loaded into the system), the annotation sets to be used for each, and the annotation type to be compared.

Note that the tool automatically intersects all the annotation types from the selected key annotation set with all types from the response set.

On a separate note, you can perform a diff on the same document, between two different annotation sets. One annotation set could contain the key type and another could contain the response one.

After the type has been selected, the user is required to decide how the features will be compared. It is important to know that the tool compares them by analysing if features from the key set are contained in the response set. It checks for both the feature name and feature value to be the same.

There are three basic options to select:

- To take 'all' the features from the key set into consideration

Key doc: EP-1016726-A1.x... Key set: annotator1 Type: Measurement Weight: 1.00

Resp. doc: EP-1016726-A1.x... Resp. set: annotator2 Features: all some none

Start	End	Key	Features	=?	Start	End	Response
26517	26532	2.and+0.25mg/ml	{rule=measurement.Me...fore=MgCl, conj=and}	=	26517	26532	2.and+0.25mg/ml
32763	32769	70-80%	{conj=-, type=interv...urement.MeasureSpan}	=	32763	32769	70-80%
37992	38003	10'.and:15'	{conj=and, type=inte...urement.MeasureSpan}	=	37992	38003	10'.and:15'
25580	25593	32°C.and+39°C	{rule=measurement.Me...before=at, conj=and}	=	25580	25593	32°C.and+39°C
31576	31578	cm	{dimension=[length],...afepat;centi_metre}}	~	31576	31580	cm-2
25093	25097	cm-2	{dimension=[length],...afepat;centi_metre}}	~	25093	25095	cm
57200	57211	mm-2./field	{dimension=[length],...afepat;milli-meter}}	~	57200	57204	mm-2
40226	40227	3	{rule=measurement.Si...e, type=scalarValue}	-?			
30461	30462	2	...n, type=scalarValue}	-?			
35865	35866	4	...afepat;arc_minute}}	-?			
45397	45398	5	...e, type=scalarValue}	-?			
				?	24794	24796	37
				?	33474	33480	per-ml

Legend:
 l = new line, - = tab, . = space

Correct:	462	Recall	Precision	F-measure
Partially correct:	3	Strict: 0,99	0,95	0,97
Missing:	4	Lenient: 0,99	0,95	0,97
False positives:	23	Average: 0,99	0,95	0,97

2 annotations copied to consensus and 2 hidden

Buttons: Show document, Export to HTML

Navigation: Statistics, Adjudication

Figure 10.2: Annotation diff window with the parameters at the top, the comparison table in the center and the statistics panel at the bottom.

- To take only 'some' user selected features
- To take 'none' of the features from the key set.

The weight for the F-Measure can also be changed - by default it is set to 1.0 (i.e. to give precision and recall equal weight). Finally, click on 'Compare' to display the results. Note that the window may need to be resized manually, by dragging the window edges as appropriate).

In the main window, the key and response annotations will be displayed. They can be sorted by any category by clicking on the central column header: '='. The key and response annotations will be aligned if their indices are identical, and are color coded according to the legend displayed at the bottom.

Precision, recall, F-measure are also displayed below the annotation tables, each according to 3 criteria - strict, lenient and average. See Sections 10.2 and 10.1 for more details about the evaluation metrics.

The results can be saved to an HTML file by using the 'Export to HTML' button. This creates an HTML snapshot of what the Annotation Diff table shows at that moment. The columns and rows in the table will be shown in the same order, and the hidden columns will not appear in the HTML file. The colours will also be the same.

If you need more details or context you can use the button ‘Show document’ to display the document and the annotations selected in the annotation diff drop down lists and table.

10.3 Corpus Quality Assurance

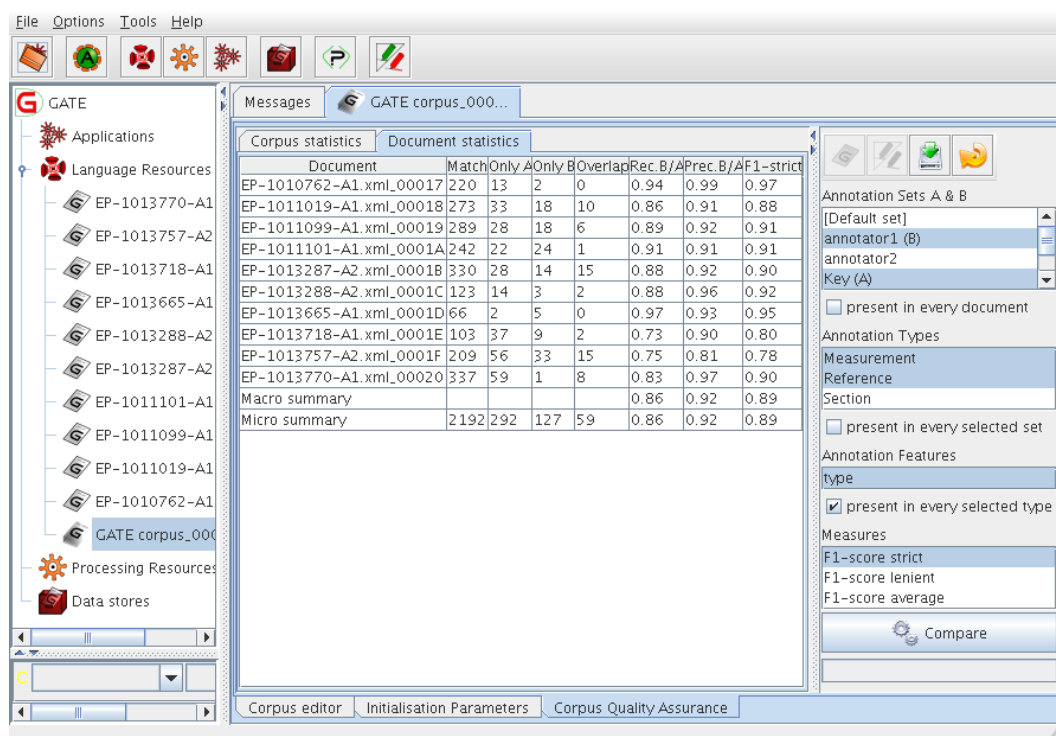


Figure 10.3: Corpus Quality Assurance showing the document statistics table

10.3.1 Description of the interface

A bottom tab in each corpus view is entitled ‘Corpus Quality Assurance’. This tab will allow you to calculate precision, recall and F-score between two annotation sets in a corpus without the need to load a plugin. It extends the Annotation Diff functionality to the entire corpus in a convenient interface.

The main part of the view consists of two tabs each containing a table. One tab is entitled ‘Corpus statistics’ and the other is entitled ‘Document statistics’.

To the right of the tabbed area is a configuration pane in which you can select the annotation sets you wish to compare, the annotation types you are interested in and the annotation features you wish to specify for use in the calculation if any.

You can also choose whether to calculate agreement on a strict or lenient basis or take the average of the two. (Recall that strict matching requires two annotations to have an identical span if they are to be considered a match, where lenient matching accepts a partial match; annotations are overlapping but not identical in span.)

At the top, four icons are for opening a document or Annotation Diff only when a row in the document statistics table is selected, exporting the two tables to an HTML file and reloading the list of sets, types and features when some documents have been modified in the corpus.

Corpus Quality Assurance works also with a corpus inside a datastore. Using a datastore is useful to minimise memory consumption when you have a big corpus.

10.3.2 Step by step usage

Begin by selecting the annotation sets you wish to compare in the top list in the configuration pane. Clicking on an annotation set labels it annotation set A (an '(A)' will appear beside it to indicate that this is your selection for annotation set A). Now click on another annotation set. This will be labelled annotation set B.

To change your selection, deselect an annotation set by clicking on it a second time. You can now choose another annotation set. Note that you do not need to hold the control key down to select the second annotation set. This list is configured to accept two (and no more than two) selections. If you wish, you may check the box 'present in every document' to reduce the annotation sets list to only those sets present in every document.

You may now choose the annotation types you are interested in. If you don't choose any then all will be used. If you wish, you may check the box 'present in every selected set' to reduce the annotation types list to only those present in every selected annotation set.

Optionally you can choose the annotation features you wish to include in the calculation. If you choose features, then for an annotation to be considered a match to another, their feature values must also match. If you select the box 'present in every selected type' the features list will be reduced to only those present in every type you selected.

The 'Measures' list allows you to choose whether to calculate strict or lenient figures or average the two. You may choose as many as you wish, and they will be included as columns in the table to the left.

Finally, click on the 'Compare' button to recalculate the tables. The figures that appear in the two tables (one per tab) are described below.

10.3.3 Details of the Corpus statistics table

In this table you will see that one row appears for every annotation type you chose. Columns give total counts for matching annotations ('Match'), annotations only present in annotation set A ('Only A'), annotations only present in annotation set B ('Only B') and annotations that overlapped ('Overlap').

Depending on whether one of your annotation sets is considered a gold standard, you might prefer to think of 'Only A' as missing and 'Only B' as spurious, or vice versa, but the Corpus Quality Assurance tool makes no assumptions about which if any annotation set is the gold standard. Where it is being used to calculate Inter Annotator Agreement there is no concept of a 'correct' set. However, in 'MUC' terms, 'Match' would be correct and 'Overlap' would be partial.

After these columns, three columns appear for every measure you chose to calculate. If you chose to calculate a strict F1, a recall, precision and F1 column will appear for the strict counts. If you chose to calculate a lenient F1, precision, recall and F1 columns will also appear for lenient counts.

In the corpus statistics table, calculations are done on a per type basis and include all documents in the calculation. Final rows in the table provide summaries; total counts are given along with a micro and a macro average.

Micro averaging treats the entire corpus as one big document where macro averaging, on this table, is the arithmetic mean of the per-type figures. See Section 10.1.4 for more detail on the distinction between a micro and a macro average.

10.3.4 Details of the Document statistics table

In this table you will see that one row appears for every document in the corpus. Columns give counts as in the corpus statistics table, but this time on a per-document basis.

As before, for every measure you choose to calculate, precision, recall and F1 columns will appear in the table.

Summary rows, again, give a macro average (arithmetic mean of the per-document measures) and micro average (identical to the figure in the corpus statistics table).

10.4 Corpus Benchmark Tool

Like the Corpus Quality Assurance functionality, the corpus benchmark tool enables evaluation to be carried out over a whole corpus rather than a single document. Unlike Corpus

QA, it uses matched corpora to achieve this, rather than comparing annotation sets within a corpus. It enables tracking of the system's performance over time. It provides more detailed information regarding the annotations that differ between versions of the corpus (e.g. annotations created by different versions of an application) than the Corpus QA tool does.

The basic idea with the tool is to evaluate an application with respect to a 'gold standard'. You have a 'marked' corpus containing the gold standard reference annotations; you have a 'clean' copy of the corpus that does not contain the annotations in question, and you have an application that creates the annotations in question. Now you can see how you are getting on, by comparing the result of running your application on 'clean' to the 'marked' annotations.

10.4.1 Preparing the Corpora for Use

You will need to prepare the following directory structure:

```
main directory (can have any name)
|
|__"clean" (directory containing unannotated documents in XML form)
|
|__"marked" (directory containing annotated documents in XML form)
|
|__"processed" (directory containing the datastore which is generated
                when you 'store corpus for future evaluation')
```

- **main:** you should have a main directory containing subdirectories for your matched corpora. It does not matter what this directory is called. This is the directory you will select when the program prompts, 'Please select a directory which contains the documents to be evaluated'.
- **clean:** Make a directory called 'clean' (case-sensitive), and in it, make a copy of your corpus that does not contain the annotations that your application creates (though it may contain other annotations). The corpus benchmark tool will apply your application to this corpus, so it is important that the annotations it creates are not already present in the corpus. You can create this corpus by copying your 'marked' corpus and deleting the annotations in question from it.
- **marked:** you should have a 'gold standard' copy of your corpus in a directory called 'marked' (case-sensitive), containing the annotations to which the program will compare those produced by your application. The idea of the corpus benchmark tool is to tell you how good your application performance is relative to this annotation set. The 'marked' corpus should contain exactly the same documents as the 'clean' set.

- **processed**: this directory contains a third version of the corpus. This directory will be created by the tool itself, when you run ‘store corpus for future evaluation’. We will explain how to do this in Section 10.4.3

10.4.2 Defining Properties

The properties of the corpus benchmark tool are defined in the file ‘corpus_tool.properties’, which should be located in the GATE home directory. GATE will tell you where it’s looking for the properties file in the ‘message’ panel when you run the Corpus Benchmark Tool. It is important to prepare this file before attempting to run the tool because there is no file present by default, so unless you prepare this file, the corpus benchmark tool will not work!

The following properties should be set:

- the precision/recall performance threshold for verbose mode, below which the annotation will be displayed in the results file. This enables problem annotations to be easily identified. By default this is set to 0.5;
- the name of the annotation set containing the human-marked annotations (annotSetName);
- the name of the annotation set containing the system-generated annotations (outputSetName);
- the annotation types to be considered (annotTypes);
- the feature values to be considered, if any (annotFeatures).

The default annotation set has to be represented by an empty string. The outputSetName and annotSetName must be different, and cannot both be the default annotation set. (If they are the same, then use the Annotation Set Transfer PR to change one of them.) If you omit any line (or just leave the value blank), that property reverts to default. For example, ‘annotSetName=’ is the same as leaving that line out.

An example file is shown below:

```
threshold=0.7
annotSetName=Key
outputSetName=ANNIE
annotTypes=Person;Organization;Location;Date;Address;Money
annotFeatures=type;gender
```

Here is another example:

```
threshold=0.6
annotSetName=Filtered
outputSetName=
annotTypes=Mention
annotFeatures=class
```

10.4.3 Running the Tool

To use the tool, first make sure the properties of the tool have been set correctly (see Section 10.4.2 for how to do this) and that the corpora and directory structure have been prepared as outlined in Section 10.4.1. Also, make sure that your application is **saved to file** (see Section 3.8.3). Then, from the ‘Tools’ menu, select ‘Corpus Benchmark’. You have four options:

1. Default Mode
2. Store Corpus for Future Evaluation
3. Human Marked Against Stored Processing Results
4. Human Marked Against Current Processing Results

We will describe these options in a different order to that in which they appear on the menu, to facilitate explanation.

Store Corpus for Future Evaluation populates the ‘processed’ directory with a datastore containing the result of running your application on the ‘clean’ corpus. If a ‘processed’ directory exists, the results will be placed there; if not, one will be created. This creates a record of the current application performance. You can rerun this operation any time to update the stored set.

Human Marked Against Stored Processing Results compares the stored ‘processed’ set with the ‘marked’ set. This mode assumes you have already run ‘Store corpus for future evaluation’. It performs a diff between the ‘marked’ directory and the ‘processed’ directory and prints out the metrics.

Human Marked Against Current Processing Results compares the ‘marked’ set with the result of running the application on the ‘clean’ corpus. It runs your application on the documents in the ‘clean’ directory creating a temporary annotated corpus and performs a diff with the documents in the ‘marked’ directory. After the metrics (recall, precision, etc.) are calculated and printed out, it deletes the temporary corpus.

Default Mode runs ‘Human Marked Against Current Processing Results’ and ‘Human Marked Against Stored Processing Results’ and compares the results of the two, showing

you where things have changed between versions. This is one of the main purposes of the benchmark tool; to show the difference in performance between different versions of your application.

Once the mode has been selected, the program prompts, 'Please select a directory which contains the documents to be evaluated'. Choose the main directory containing your corpus directories. (Do not select 'clean', 'marked', or 'processed'.) Then (except in 'Human marked against stored processing results' mode) you will be prompted to select the file containing your application (e.g. an .xgapp file).

The tool can be used either in verbose or non-verbose mode, by selecting or unselecting the verbose option from the menu. In verbose mode, for any precision/recall figure below the user's pre-defined threshold (stored in `corpus_tool.properties` file) the tool will show the the non-coextensive annotations (and their corresponding text) for that entity type, thereby enabling the user to see where problems are occurring.

10.4.4 The Results

Running the tool (either in 'Human marked against stored processing results', 'Human marked against current processing results' or 'Default' mode) produces an HTML file, in tabular form, which is output in the main GATE Developer messages window. This can then be pasted into a text editor and viewed in a web browser for easier viewing. See figure 10.4 for an example.

In each mode, the following statistics will be output:

1. Per-document figures, itemised by type: precision and recall, as well as detailed information about the differing annotations;
2. Summary by type ('Statistics'): correct, partially correct, missing and spurious totals, as well as whole corpus (micro-average) precision, recall and f-measure (F1), itemised by type;
3. Overall average figures: precision, recall and F1 calculated as a macro-average (arithmetic average) of the individual document precisions and recalls.

In 'Default' mode, information is also provided about whether the figures have increased or decreased in comparison with the 'Marked' corpus.

			PARTIALLY CORRECT ANNOTATIONS in the automatic texts: CC [2202,2210] England : [12,19] Commerce : [2618,2626] England : [2260,2267] Bank : [344,348]				
ft-airlines-27-jul-2001.xml							
Annotation Type	Precision	Recall	Annotations				
Mention	0.851063829787234	0.6896551724137931	MISSING ANNOTATIONS in the automatic texts: 1997 : [652,656] the Airline Group : [1669,1686] London Underground : [388,406] The Airline Group : [1266,1283] The Airline Group : [2412,2429] Swanwick : [2634,2642] 2000 : [2018,2022] The Airline Group : [938,955] The Airline Group : [1751,1768] Swanwick : [2376,2384] Monarch Airlines : [1052,1068] Virgin Atlantic : [45,60] Britannia Airways : [1030,1047] Labour : [634,640] London Area and Terminal Control Centre : [2029,2068] SPURIOUS ANNOTATIONS in the automatic texts: peak days : [2194,2203] PPP : [376,379] The Airline Group : [938,955] London : [2029,2035] PARTIALLY CORRECT ANNOTATIONS in the automatic texts: last year Nats handled more than 2m air traffic movements with volumes growing by 5 per cent in 2000 : [1922,2022] Abbey : [1367,1372] March : [1743,1748] March : [1486,1491] January : [2402,2409] next 10 : [2520,2527]				
ft-bmi-09-may-2001.xml							
Annotation Type	Precision	Recall	Annotations				
Mention	0.9791666666666666	0.94					
Statistics							
Annotation Type	Correct	Partially Correct	Missing	Spurious	Precision	Recall	F-Measure
Mention	434	16	34	8	0.9650655021834061	0.9132231404958677	0.9384288747346073
Overall average precision: 0.959158791298373 Overall average recall: 0.9108515466531439 Overall average fMeasure : 0.9332844120837758 Finished!							

Figure 10.4: Fragment of results from corpus benchmark tool

10.5 A Plugin Computing Inter-Annotator Agreement (IAA)

The interannotator agreement plugin, ‘Inter_Annotator_Agreement’, computes the F-measures, namely precision, recall and F1, suitable for named entity annotations (see Section 10.1.3), and agreement, Cohen’s kappa and Scott’s pi, suitable for text classification tasks (see Section 10.1.2). In the latter case, a confusion matrix is also provided. In this section we describe those measures and the output results from the plugin. But first we explain how to load the plugin, and the input to and the parameters of the plugin.

First you need to load the plugin named ‘Inter_Annotator_Agreement’ into GATE Developer using the tool *Manage CREOLE Plugins*, if it is not already loaded. Then you can create a PR for the plugin from the ‘IAA Computation’ in the existing PR list. After that you can put the PR into a *Corpus Pipeline* to use it.

The IAA Computation PR differs from the Corpus Benchmark Tool in the data preparation required. As in the Corpus Benchmark Tool, the idea is to compare annotation sets, for example, prepared by different annotators, but in the IAA Computation PR, these annotation sets should be on the same set of documents. Thus, one corpus is loaded into GATE on which the PR is run. Different annotation sets contain the annotations which will be

compared. These should (obviously) have different names.

It falls to the user to decide whether to use annotation type or an annotation feature as class; are two annotations considered to be in agreement because they have the same type and the same span? Or do you want to mark up your data with an annotation type such as ‘Mention’, thus defining the relevant annotations, then give it a ‘class’ feature, the value of which should be matched in order that they are considered to agree? This is a matter of convenience. For example, data from the Batch Learning PR (see Section 15.2) uses a single annotation type and a class feature. In other contexts, using annotation type might feel more natural; the annotation sets should agree about what is a ‘Person’, what is a ‘Date’ etc. It is also possible to mix the two, as you will see below.

The IAA plugin has two runtime parameters **annSetsForIaa** and **annTypesAndFeats** for specifying the annotation sets and the annotation types and features, respectively. Values should be separated by semicolons. For example, to specify annotation sets ‘Ann1’, ‘Ann2’ and ‘Ann3’ you should set the value of *annSetsForIaa* to ‘Ann1;Ann2;Ann3’. Note that more than two annotation sets are possible. Specify the value of *annTypesAndFeats* as ‘Per’ to compute the IAA for the three annotation sets on the annotation type *Per*. You can also specify more than one annotation type and separate them by ‘;’ too, and optionally specify an annotation feature for a type by attaching a ‘->’ followed by feature name to the end of the annotation name. For example, ‘Per->label;Org’ specifies two annotation types *Per* and *Org* and also a feature name *label* for the type *Per*. If you specify an annotation feature for an annotation type, then two annotations of the same type will be regarded as being different if they have different values of that feature, even if the two annotations occupy exactly the same position in the document. On the other hand, if you do not specify any annotation feature for an annotation type, then the two annotations of the type will be regarded as the same if they occupy the same position in the document.

The parameter **measureType** specifies the type of measure computed. There are two measure types; the *F-measure* (i.e. Precision, Recall and F1), and the *observed agreement and Cohen’s Kappa*. For classification tasks such as document or sentence classification, the observed agreement and Cohen’s Kappa is often used, though the F-measure is applicable too. In these tasks, the targets are already identified, and the task is merely to classify them correctly. However, for the named entity recognition task, only the F-measure is applicable. In such tasks, finding the ‘named entities’ (text to be annotated) is as much a part of the task as correctly labelling it. Observed agreement and Cohen’s kappa are not suitable in this case. See Section 10.1.2 for further discussion. The parameter has two values, *FMEASURE* and *AGREEMENTANDKAPPA*. The default value of the parameter is *FMEASURE*.

Another parameter **verbosity** specifies the verbosity level of the plugin’s output. Level 2 displays the most detailed output, including the IAA measures on each document and the macro-averaged results over all documents. Level 1 only displays the IAA measures averaged over all documents. Level 0 does not have any output. The default value of the parameter is 1. In the following we will explain the outputs in detail.

Yet another runtime parameter **bdmScoreFile** specifies the URL for a file containing the

BDM scores used for the BDM based IAA computation. The BDM score file should be produced by the BDM computation plugin, which is described in Section 10.6. The BDM-based IAA computation will be explained below. If the parameter is not assigned any value, or is assigned a file which is not a BDM score file, the PR will not compute the BDM based IAA.

10.5.1 IAA for Classification

IAA has been used mainly in classification tasks, where two or more annotators are given a set of instances and are asked to classify those instances into some pre-defined categories. IAA measures the agreements among the annotators on the class labels assigned to the instances by the annotators. Text classification tasks include document classification, sentence classification (e.g. opinionated sentence recognition), and token classification (e.g. POS tagging). The important point to note is that the evaluation set and gold standard set have exactly the same instances, but some instances in the two sets have different class labels. Identifying the instances is not part of the problem.

The three commonly used IAA measures are *observed agreement*, *specific agreement*, and *Kappa* (κ) [Hripcsak & Heitjan 02]. See Section 10.1.2 for the detailed explanations of those measures. If you select the value of the runtime parameter *measureType* as *AGREEMENTANDKAPPA*, the IAA plugin will compute and display those IAA measures for your classification task. Below, we will explain the output of the PR for the agreement and Kappa measures.

At the verbosity level 2, the output of the plugin is the most detailed. It first prints out a list of the names of the annotation sets used for IAA computation. In the rest of the results, the first annotation set is denoted as annotator 0, and the second annotation set is denoted as annotator 1, etc. Then the plugin outputs the IAA results for each document in the corpus.

For each document, it displays one annotation type and optionally an annotation feature if specified, and then the results for that type and that feature. Note that the IAA computations are based on the pairwise comparison of annotators. In other words, we compute the IAA for each pair of annotators. The first results for one document and one annotation type are the macro-averaged ones over all pairs of annotators, which have three numbers for the three types of IAA measures, namely *Observed agreement*, *Cohen's kappa* and *Scott's pi*. Then for each pair of annotators, it outputs the three types of measures, a confusion matrix (or contingency table), and the specific agreements for each label. The labels are obtained from the annotations of that particular type. For each annotation type, if a feature is specified, then the labels are the values of that feature. Please note that two terms may be added to the label list: one is the empty one obtained from those annotations which have the annotation feature but do not have a value for the feature; the other is 'Non-cat', corresponding to those annotations not having the feature at all. If no feature is specified, then two labels are used: 'Anns' corresponding to the annotations of that type, and 'Non-cat' corresponding to those annotations which are annotated by one annotator but are not

annotated by another annotator.

After displaying the results for each document, the plugin prints out the macro-averaged results over all documents. First, for each annotation type, it prints out the results for each pair of annotators, and the macro-averaged results over all pairs of annotators. Finally it prints out the macro-averaged results over all pairs of annotators, all types and all documents.

Please note that the classification problem can be evaluated using the F-measure too. If you want to evaluate a classification problem using the F-measure, you just need to set the run time parameter *measureType* to *FMEASURE*.

10.5.2 IAA For Named Entity Annotation

The commonly used IAA measures, such as kappa, have not been used in text mark-up tasks such as named entity recognition and information extraction, for reasons explained in Section 10.1.2 (also see [Hripcsak & Rothschild 05]). Instead, the F-measures, such as Precision, Recall, and F1, have been widely used in information extraction evaluations such as MUC, ACE and TERN for measuring IAA. This is because the computation of the F-measures does not need to know the number of non-entity examples. Another reason is that F-measures are commonly used for evaluating information extraction systems. Hence IAA F-measures can be directly compared with results from other systems published in the literature.

For computing F-measure between two annotation sets, one can use one annotation set as gold standard and another set as system's output and compute the F-measures such as Precision, Recall and F1. One can switch the roles of the two annotation sets. The Precision and Recall in the former case become Recall and Precision in the latter, respectively. But the F1 remains the same in both cases. For more than two annotators, we first compute F-measures between any two annotators and use the mean of the pair-wise F-measures as an overall measure.

The computation of the F-measures (e.g. Precision, Recall and F1) are shown in Section 10.1. As noted in [Hripcsak & Rothschild 05], the F1 computed for two annotators for one specific category is equivalent to the positive specific agreement of the category.

The outputs of the IAA plugins for named entity annotation are similar to those for classification. But the outputs are the F-measures, such as Precision, Recall and F1, instead of the agreements and Kappas. It first prints out the results for each document. For one document, it prints out the results for each annotation type, macro-averaged over all pairs of annotators, then the results for each pair of annotators. In the last part, the micro-averaged results over all documents are displayed. Note that the results are reported in both the strict measure and the lenient measure, as defined in Section 10.2.

Please note that, for computing the F-measures for the named entity annotations, the IAA plugin carries out the same computation as the *Corpus Benchmark tool*. The IAA plugin is

simpler than the Corpus benchmark tool in the sense that the former needs only one set of documents with two or more annotation sets, whereas the latter needs three sets of the same documents, one without any annotation, another with one annotation set, and the third one with another annotation set. Additionally, the IAA plugin can deal with more than two annotation sets but the Corpus benchmark tool can only deal with two annotation sets.

10.5.3 The BDM-Based IAA Scores

For a named entity recognition system, if the named entity's class labels are the names of concepts in some ontology (e.g. in the ontology-based information extraction), the system can be evaluated using the IAA measures based on the BDM scores. The BDM measures the closeness of two concepts in an ontology. If an entity is identified but is assigned a label which is close to but not the same as the true label, the system should obtain some credit for it, which the BDM-based metric can do. In contrast, the conventional named entity recognition measure does not take into account the closeness of two labels and does not give any credit to one identified entity with a wrong label, regardless of how close the assigned label is to the true label. For more explanation about BDM see Section 10.6.

In order to compute the BDM-based IAA, one has to assign the plugin's runtime parameter **bdmScoreFile** to the URL of a file containing the BDM scores. The file should be obtained by using the BDM computation plugin, which is described in Section 10.6. Currently the BDM-based IAA is only used for computing the F-measures for e.g. the entity recognition problem. Please note that the F-measures can also be used for evaluation of classification problem. The BDM is not used for computing other measures such as the *observed agreement* and *Kappa*, though it is possible to implement it. Therefore currently one has to select *FMEASURE* for the run time parameter *measureType* in order to use the BDM based IAA computation.

10.6 A Plugin Computing the BDM Scores for an Ontology

The BDM (balanced distance metric) measures the closeness of two concepts in an ontology or taxonomy [Maynard 05, Maynard *et al.* 06]. It is a real number between 0 and 1. The closer the two concepts are in an ontology, the greater their BDM score is. For detailed explanation about the BDM, see the papers [Maynard 05, Maynard *et al.* 06]. The BDM can be seen as an improved version of the learning accuracy [Cimiano *et al.* 03]. It is dependent on the length of the shortest path connecting the two concepts and also the deepness of the two concepts in ontology. It is also normalised with the size of ontology and also takes into account the concept density of the area containing the two involved concepts.

The BDM has been used to evaluate the ontology based information extraction (qOBIE)

system [Maynard *et al.* 06]. The OBIE identifies the instances for the concepts of an ontology. It's possible that an OBIE system identifies an instance successfully but does not assign it the correct concept. Instead it assigns the instance a concept being close to the correct one. For example, the entity 'London' is an instance of the concept *Capital*, and an OBIE system assigns it the concept *City* which is close to the concept *Capital* in some ontology. In that case the OBIE should obtain some credit according to the closeness of the two concepts. That is where the BDM can be used. The BDM has also been used to evaluate the hierarchical classification system [Li *et al.* 07b]. It can also be used for ontology learning and alignment.

The BDM computation plugin computes BDM score for each pair of concepts in an ontology. It has two run time parameters:

- **ontologyURL** – its value should be the URL of the ontology that one wants to compute the BDM scores for.
- **outputBDMFile** – its value is the URL of a file which will store the BDM scores computed.

The plugin has the name *Ontology_BDM_Computation* and the corresponding processing resource's name is *BDM Computation PR*. The PR can be put into a Pipeline. If it is put into a Corpus Pipeline, the corpus used should contain at least one document.

The BDM computation used the formula given in [Maynard *et al.* 06]. The resulting file specified by the runtime parameter *outputBDMFile* contains the BDM scores. It is a text file. The first line of the file gives some meta information such as the name of ontology used for BDM computation. From the second line of the file, each line corresponds to one pair of concepts. One line is like

```
key=Service, response=Object, bdm=0.6617647, msca=Object, cp=1, dpk=1, dpr=0,  
n0=2.0, n1=2.0, n2=2.8333333, bran=1.9565217
```

It first shows the names of the two concepts (one as *key* and another as *response*, and the BDM score, and then other parameters' values used for the computation. Note that, since the BDM is symmetric for the two concepts, the resulting file contains only one line for each pair. So if you want to look for the BDM score for one pair of concepts, you can choose one as key and another as response. If you cannot find the line for the pair, you have to change the order of two concepts and retrieve the file again.

Chapter 11

Profiling Processing Resources

11.1 Overview

This is a reporting tool for GATE processing resources. It reports the total time taken by processing resources and the time taken for each document to be processed by an application of type corpus pipeline.

GATE use log4j, a logging system, to write profiling informations in a file. The GATE profiling reporting tool uses the file generated by log4j and produces a report on the processing resources. It profiles JAPE grammars at the rule level, enabling the user precisely identify the performance bottlenecks. It also produces a report on the time taken to process each document to find problematic documents.

This initial code for the reporting tool was written by Intelius employees Andrew Borthwick and Chirag Viradiya and generously released under the LGPL licence to be part of GATE.

Processing elements of following pipelines	Time in seconds	% time taken
• ANNIE		
[-] ANNIE	67.233	100.0
[+] pr_ANNIE_NE_Transducer	47.631	70.8
[+] pr_ANNIE_English_Tokeniser	7.053	10.5
pr_ANNIE_OrthoMatcher	5.574	8.3
[+] pr_ANNIE_Sentence_Splitter	5.181	7.7
pr_ANNIE_POS_Tagger	0.903	1.3
All others	0.47	0.7
pr_ANNIE_Gazetteer	0.217	0.3
pr_Document_Reset_PR	0.204	0.3

Figure 11.1: Example of HTML profiling report for ANNIE

11.1.1 Features

- Ability to generate the following two reports
 - Report on processing resources. For each level of processing: application, processing resource (PR) and grammar rule, subtotaled at each level.
 - Report on documents processed. For some or all PR, sorted in decreasing processing time.
- Report on processing resources specific features
 - Sort order by time or by execution.
 - Show or hide processing elements which took 0 milliseconds.
 - Generate HTML report with a collapsible tree.
- Report on documents processed specific features
 - Limit the number of document to show from the most time consuming.
 - Filter the PR to display statistics for.
- Features common to both reports
 - Generate report as indented text or in HTML format.
 - Generate a report only on the log entries from the last logical run of GATE.
 - All processing times are reported in milliseconds and in terms of percentage (rounded to nearest 0.1%) of total time.
 - Command line interface and API.
 - Detect if the benchmark.txt file is modified while generating the report.

11.1.2 Limitations

Be aware that the profiling doesn't support non corpus pipeline as application type. There is indeed no interest in profiling a non corpus pipeline that works on one or no document at all. To get meaningful results you should run your corpus pipeline on at least 10 documents.

11.2 Graphical User Interface

The activation of the profiling and the creation of profiling reports are accessible from the 'Tools' menu in GATE with the submenu 'Profiling Reports'.

You can ‘Start Profiling Applications’ and ‘Stop Profiling Applications’ at any time. The logging is cumulative so if you want to get a new report you must use the ‘Clear Profiling History’ menu item when the profiling is stopped.

Two types of reports are available: ‘Report on Processing Resources’ and ‘Report on Documents Processed’. See the previous section for more information.

11.3 Command Line Interface

Report on processing resources Usage: `java gate.util.reporting.PRTimeReporter [Options]`

Options:

- i input file path (default: benchmark.txt in the user’s .gate directory¹)
- m print media - html/text (default: html)
- z suppressZeroTimeEntries - true/false (default: true)
- s sorting order - exec_order/time_taken (default: exec_order)
- o output file path (default: report.html/txt in the system temporary directory)
- l logical start (not set by default)
- h show help

Note that suppressZeroTimeEntries will be ignored if the sorting order is ‘time_taken’

Report on documents processed Usage: `java gate.util.reporting.DocTimeReporter [Options]`

Options:

- i input file path (default: benchmark.txt in the user’s .gate directory²)
- m print media - html/text (default: html)
- d number of docs, use -1 for all docs (default: 10 docs)
- p processing resource name to be matched (default: all_prs)

¹GATE versions up to 5.2 placed benchmark.txt in the execution directory.

²GATE versions up to 5.2 placed benchmark.txt in the execution directory.

- o output file path (default: report.html/txt in the system temporary directory)
- l logical start (not set by default)
- h show help

Examples

- Run report 1: Report on Total time taken by each processing element across corpus
 - java -cp "gate/bin:gate/lib/GnuGetOpt.jar" gate.util.reporting.PRTimeReporter -i benchmark.txt -o report.txt -m text
- Run report 2: Report on Time taken by document within given corpus.
 - java -cp "gate/bin:gate/lib/GnuGetOpt.jar" gate.util.reporting.DocTimeReporter -i benchmark.txt -o report.html -m html

11.4 Application Programming Interface

11.4.1 Log4j.properties

This is required to direct the profiling information to the benchmark.txt file. The benchmark.txt generated by GATE will be used as input for GATE profiling report tool as input.

- # File appender that outputs only benchmark messages
- log4j.appender.benchmarklog=org.apache.log4j.RollingFileAppender
- log4j.appender.benchmarklog.Threshold=DEBUG
- log4j.appender.benchmarklog.File=\$user.home/.gate/benchmark.txt
- log4j.appender.benchmarklog.MaxFileSize=5MB
- log4j.appender.benchmarklog.MaxBackupIndex=1
- log4j.appender.benchmarklog.layout=org.apache.log4j.PatternLayout
- log4j.appender.benchmarklog.layout.ConversionPattern=%m%n
- # Configure the Benchmark logger so that it only goes to the benchmark log file
- log4j.logger.gate.util.Benchmark=DEBUG, benchmarklog
- log4j.additivity.gate.util.Benchmark=false

11.4.2 Benchmark log format

The format of the benchmark file that logs the times is as follow:

```
timestamp START PR_name
timestamp duration benchmarkID class features
timestamp duration benchmarkID class features
...
```

with the timestamp being the difference, measured in milliseconds, between the current time and midnight, January 1, 1970 UTC.

Example:

```
1257269774770 START Sections_splitter
1257269774773 0 Sections_splitter.doc_EP-1026523-A1_xml_00008.documentLoaded
gate.creole.SerialAnalyserController
{corpusName=Corpus for EP-1026523-A1.xml_00008,
documentName=EP-1026523-A1.xml_00008}
...
```

11.4.3 Enabling profiling

There are two ways to enable profiling of the processing resources:

1. In `gate/build.properties`, add the line: `run.gate.enable.benchmark=true`
2. In your Java code, use the method: `Benchmark.setBenchmarkingEnabled(true)`

11.4.4 Reporting tool

Report on processing resources

1. Instantiate the Class `PRTIMEReporter`
 - (a) `PRTIMEReporter report = new PRTIMEReporter();`
2. Set the input benchmark file
 - (a) `File benchmarkFile = new File("benchmark.txt");`

- (b) `report.setBenchmarkFile(benchmarkFile);`
- 3. Set the output report file
 - (a) `File reportFile = new File("report.txt");` or
 - (b) `File reportFile = new File("report.html");`
 - (c) `report.setReportFile(reportFile);`
- 4. Set the output format: in html or text format (default: `MEDIA_HTML`)
 - (a) `report.setPrintMedia(PRTimeReporter.MEDIA_TEXT);` or
 - (b) `report.setPrintMedia(PRTimeReporter.MEDIA_HTML);`
- 5. Set the sorting order: Sort in order of execution or descending order of time taken (default: `EXEC_ORDER`)
 - (a) `report.setSortOrder(PRTimeReporter.SORT_TIME_TAKEN);` or
 - (b) `report.setSortOrder(PRTimeReporter.SORT_EXEC_ORDER);`
- 6. Set if suppress zero time entries: `True/False` (default: `True`). Parameter ignored if `SortOrder` specified is `'SORT_TIME_TAKEN'`
 - (a) `report.setSuppressZeroTimeEntries(true);`
- 7. Set the logical start: A string indicating the logical start to be operated upon for generating reports
 - (a) `report.setLogicalStart("InteliusPipelineStart");`
- 8. Generate the text/html report
 - (a) `report.executeReport();`

Report on documents processed

- 1. Instantiate the Class `DocTimeReporter`
 - (a) `DocTimeReporter report = new DocTimeReporter();`
- 2. Set the input benchmark file
 - (a) `File benchmarkFile = new File("benchmark.txt");`
 - (b) `report.setBenchmarkFile(benchmarkFile);`
- 3. Set the output report file
 - (a) `File reportFile = new File("report.txt");` or

- (b) `File reportFile = new File("report.html");`
- (c) `report.setReportFile(reportFile);`
- 4. Set the output format: Generate report in html or text format (default: MEDIA_HTML)
 - (a) `report.setPrintMedia(DocTimeReporter.MEDIA_TEXT);` or
 - (b) `report.setPrintMedia(DocTimeReporter.MEDIA_HTML);`
- 5. Set the maximum number of documents: Maximum number of documents to be displayed in the report (default: 10 docs)
 - (a) `report.setNoOfDocs(2); // 2 docs or`
 - (b) `report.setNoOfDocs(DocTimeReporter.ALL_DOCS); // All documents`
- 6. Set the PR matching regular expression: A PR name or a regular expression to filter the results (default: MATCH_ALL_PR_REGEX).
 - (a) `report.setSearchString("HTML"); // match ALL PRS having HTML as substring`
- 7. Set the logical start: A string indicating the logical start to be operated upon for generating reports
 - (a) `report.setLogicalStart("InteliusPipelineStart");`
- 8. Generate the text/html report
 - (a) `report.executeReport();`

Chapter 12

Developing GATE

This chapter describes ways of getting involved in and contributing to the GATE project. Sections 12.1 and 12.2 are good places to start. Sections 12.3 and 12.4 describe protocol and provide information for committers; we cover creating new plugins and updating this user guide. See Section 12.2 for information on becoming a committer.

12.1 Reporting Bugs and Requesting Features

The GATE bug tracker can be found on SourceForge, [here](#). When reporting bugs, please give as much detail as possible. Include the GATE version number and build number, the platform on which you observed the bug, and the version of Java you were using (1.5.0_15, 1.6.0_03, etc.). Include steps to reproduce the problem, and a full stack trace of any exceptions, including ‘Caused by ...’. You may wish to first check whether the bug is already fixed in the latest nightly build. You may also request new features.

12.2 Contributing Patches

Patches may be submitted on SourceForge. The best format for patches is an SVN diff against the latest subversion. The diff can be saved as a file and attached; it should not be pasted into the bug report. Note that we generally do not accept patches against earlier versions of GATE. Also, GATE is intended to be compatible with Java 5, so if you regularly develop using a later version of Java it is very important to compile and test your patches on Java 5. Patches that use Java-6-only features and do not compile and run on Java 5 will not be accepted.

If you intend to submit larger changes, you might prefer to become a committer! We welcome input to the development process of GATE. The code is hosted on SourceForge, providing

anonymous Subversion access (see Section 2.2.3). We're happy to give committer privileges to anyone with a track record of contributing good code to the project. We also make the current version available nightly on the ftp site.

12.3 Creating New Plugins

GATE provides a flexible structure where new resources can be plugged in very easily. There are three types of resources: Language Resource (LR), Processing Resource (PR) and Visual Resource (VR). In the following subsections we describe the necessary steps to write new PRs and VRs, and to add plugins to the nightly build. The guide on writing new LRs will be available soon.

12.3.1 Where to Keep Plugins in the GATE Hierarchy

Each new resource added as a plugin should contain its own subfolder under the `%GATE-HOME%/plugins` folder with an associated `creole.xml` file. A plugin can have one or more resources declared in its `creole.xml` file and/or using source-level annotations as described in section 4.7.

12.3.2 What to Call your Plugin

The plugins are many and the list is constantly expanding. The naming convention aims to impose order and group plugins in a readable manner. When naming new plugins, please adhere to the following guidelines;

- Words comprising plugin names should be capitalized and separated by underscores Like_So. This means that they will format nicely in GATE Developer. For example, 'Inter_Annotator_Agreement'.
- Plugin names should begin with the word that best describes their function. Practically, this means that words are often reversed from the usual order, for example, the Chemistry Tagger plugin should be called 'Tagger_Chemistry'. This means that for example parsers will group together alphabetically and thus will be easy to find when someone is looking for parsers. Before naming your plugin, look at the existing plugins and see where it might group well.

12.3.3 Writing a New PR

Class Definition

Below we show a template class definition, which can be used in order to write a new Processing Resource.

```

1
2 package example;
3
4 import gate.*;
5 import gate.creole.*;
6 import gate.creole.metadata.*;
7
8 /**
9  * Processing Resource. The @CreoleResource annotation marks this
10  * class as a GATE Resource, and gives the information GATE needs
11  * to configure the resource appropriately.
12  */
13 @CreoleResource(name = "Example PR",
14                 comment = "An example processing resource")
15 public class NewPlugin extends AbstractLanguageAnalyser {
16
17     /*
18     * this method gets called whenever an object of this
19     * class is created either from GATE Developer GUI or if
20     * initiated using Factory.createResource() method.
21     */
22     public Resource init() throws ResourceInstantiationException {
23         // here initialize all required variables, and may
24         // be throw an exception if the value for any of the
25         // mandatory parameters is not provided
26
27         if(this.rulesURL == null)
28             throw new ResourceInstantiationException("rules URL null");
29
30         return this;
31     }
32
33
34     /*
35     * this method should provide the actual functionality of the PR
36     * (from where the main execution begins). This method
37     * gets called when user click on the "RUN" button in the
38     * GATE Developer GUI's application window.
39     */
40     public void execute() throws ExecutionException {
41         // write code here
42     }
43
44     /* this method is called to reinitialize the resource */
45     public void reInit() throws ResourceInstantiationException {

```

```

46     // reinitialization code
47 }
48
49 /*
50  * There are two types of parameters
51  * 1. Init time parameters – values for these parameters need to be
52  * provided at the time of initializing a new resource and these
53  * values are not supposed to be changed.
54  * 2. Runtime parameters – values for these parameters are provided
55  * at the time of executing the PR. These are runtime parameters and
56  * can be changed before starting the execution
57  * (i.e. before you click on the "RUN" button in GATE Developer)
58  * A parameter myParam is specified by a pair of methods getMyParam
59  * and setMyParam (with the first letter of the parameter name
60  * capitalized in the normal Java Beans style), with the setter
61  * annotated with a @CreoleParameter annotation.
62  *
63  * for example to set a value for outputAnnotationSetName
64  */
65 String outputAnnotationSetName;
66
67 //getter and setter methods
68
69 /* get<parameter name with first letter Capital> */
70 public String getOutputAnnotationSetName() {
71     return outputAnnotationSetName;
72 }
73
74 /* The setter method is annotated to tell GATE that it defines an
75  * optional runtime parameter.
76  */
77 @Optional
78 @RunTime
79 @CreoleParameter(
80     comment = "name of the annotationSet used for output")
81 public void setOutputAnnotationSetName(String setName) {
82     this.outputAnnotationSetName = setName;
83 }
84
85 /** Init-time parameter */
86 URL rulesURL;
87
88 // getter and setter methods
89 public URL getRulesURL() {
90     return rulesFile;
91 }
92
93 /* This parameter is not annotated @RunTime or @Optional, so it is a
94  * required init-time parameter.
95  */
96 @CreoleParameter(
97     comment = "example of an inittime parameter",
98     defaultValue = "resources/morph/default.rul")

```



```
99     public void setRulesURL(URL rulesURL) {
100         this.rulesURL = rulesURL;
101     }
102 }
```

PR Creole Entry

The creole.xml file simply needs to tell GATE which JAR file to look in to find the PR.

```
<?xml version="1.0"?>
<CREOLE-DIRECTORY>
    <JAR SCAN="true">newplugin.jar</JAR>
</CREOLE-DIRECTORY>
```

Alternatively the configuration can be given in the XML file directly instead of using source annotations. Section 4.7 gives the full details.

Context Menu

Each resource (LR,PR) has some predefined actions associated with it. These actions appear in a context menu that appears in GATE Developer when the user right clicks on any of the resources. For example if the selected resource is a Processing Resource, there will be at least four actions available in its context menu: 1. Close 2. Hide 3. Rename and 4. Reinitialize. New actions in addition to the predefined actions can be added by implementing the *gate.gui.ActionsPublisher* interface in either the LR/PR itself or in any associated VR. Then the user has to implement the following method.

```
public List getActions() {
    return actions;
}
```

Here the variable *actions* should contain a list of instances of type *javax.swing.AbstractAction*. A string passed in the constructor of an *AbstractAction* object appears in the context menu. Adding a *null* element adds a separator in the menu.

Listeners

There are at least four important listeners which should be implemented in order to listen to the various relevant events happening in the background. These include:

- CreoleListener

Creole-register keeps information about instances of various resources and refreshes itself on new additions and deletions. In order to listen to these events, a class should implement the *gate.event.CreoleListener*. Implementing CreoleListener requires users to implement the following methods:

- public void resourceLoaded(CreoleEvent creoleEvent);
- public void resourceUnloaded(CreoleEvent creoleEvent);
- public void resourceRenamed(Resource resource, String oldName, String newName);
- public void datastoreOpened(CreoleEvent creoleEvent);
- public void datastoreCreated(CreoleEvent creoleEvent);
- public void datastoreClosed(CreoleEvent creoleEvent);

- DocumentListener

A traditional GATE document contains text and a set of annotationSets. To get notified about changes in any of these resources, a class should implement the *gate.event.DocumentListener*. This requires users to implement the following methods:

- public void contentEdited(DocumentEvent event);
- public void annotationSetAdded(DocumentEvent event);
- public void annotationSetRemoved(DocumentEvent event);

- AnnotationSetListener

As the name suggests, AnnotationSet is a set of annotations. To listen to the addition and deletion of annotations, a class should implement the *gate.event.AnnotationSetListener* and therefore the following methods:

- public void annotationAdded(AnnotationSetEvent event);
- public void annotationRemoved(AnnotationSetEvent event);

- AnnotationListener

Each annotation has a featureMap associated with it, which contains a set of feature names and their respective values. To listen to the changes in annotation, one needs to implement the *gate.event.AnnotationListener* and implement the following method:

- public void annotationUpdated(AnnotationEvent event);




Figure 12.1: GATE GUI

12.3.4 Writing a New VR

Each resource (PR and LR) can have its own associated visual resource. When double clicked, the resource's respective visual resource appears in GATE Developer. The GATE Developer GUI is divided into three visible parts (See Figure 12.1). One of them contains a tree that shows the loaded instances of resources. The one below this is used for various purposes - such as to display document features and that the execution is in progress. This part of the GUI is referred to as 'small'. The third and the largest part of the GUI is referred to as 'large'. One can specify which one of these two should be used for displaying a new visual resource in the creole.xml.

Class Definition

Below we show a template class definition, which can be used in order to write a new Visual Resource.

```

1 package example.gui;
2
3 import gate.*;
4 import gate.creole.*;
5 import gate.creole.metadata.*;
6
7 /*
8  * An example Visual Resource for the New Plugin
9  * Note that here we extends the AbstractVisualResource class.
10 * The @CreoleResource annotation associates this VR with the
11 * underlying PR type it displays.
12 */
13 @CreoleResource(name = "Visual resource for new plugin",

```

```

14         guiType = GuiType.LARGE,
15         resourceDisplayed = "example.NewPlugin",
16         mainViewer = true)
17 public class NewPluginVR extends AbstractVisualResource {
18
19     /*
20     * An Init method called when the GUI is initialized for
21     * the first time
22     */
23     public Resource init() {
24         // initialize GUI Components
25         return this;
26     }
27
28     /*
29     * Here target is the PR class to which this Visual Resource
30     * belongs. This method is called after the init() method.
31     */
32     public void setTarget(Object target) {
33         // check if the target is an instance of what you expected
34         // and initialize local data structures if required
35     }
36 }

```

Every document has its own document viewer associated with it. It comes with a single component that shows the text of the original document. GATE provides a way to attach new GUI plugins to the document viewer. For example AnnotationSet viewer, AnnotationList viewer and Co-Reference editor. These are the examples of DocumentViewer plugins shipped as part of the core GATE build. These plugins can be displayed either on the right or on top of the document viewer. They can also replace the text viewer in the center (See figure 12.1). A separate button is added at the top of the document viewer which can be pressed to display the GUI plugin.

Below we show a template class definition, which can be used to develop a new DocumentViewer plugin.

```

1
2  /*
3  * Note that the class needs to extends the AbstractDocumentView class
4  */
5  @CreoleResource
6  public class DocumentViewerPlugin extends AbstractDocumentView {
7
8      /* Implementers should override this method and use it for
9      * populating the GUI.
10     */
11     public void initGUI() {
12         // write code to initialize GUI
13     }
14
15     /* Returns the type of this view */
16     public int getType() {

```

```

17         // it can be any of the following constants
18         // from the gate.gui.docview.DocumentView
19         // CENTRAL, VERTICAL, HORIZONTAL
20     }
21
22     /* Returns the actual UI component this view represents. */
23     public Component getGUI() {
24         // return the top level GUI component
25     }
26
27     /* This method called whenever view becomes active.*/
28     public void registerHooks() {
29         // register listeners
30     }
31
32     /* This method called whenever view becomes inactive. */
33     public void unregisterHooks() {
34         // do nothing
35     }
36 }

```

12.3.5 Adding Plugins to the Nightly Build

If you add a new plugin and want it to be part of the build process, you should create a build.xml file with targets ‘build’, ‘test’, ‘distro.prepare’, ‘javadoc’ and ‘clean’. The build target should build the JAR file, test should run any unit tests, distro.prepare should clean up any intermediate files (e.g. the classes/ directory) and leave just what’s in Subversion, plus the compiled JAR file and javadocs. The clean target should clean up everything, including the compiled JAR and any generated sources, etc. You should also add your plugin to ‘plugins.to.build’ in the top-level build.xml to include it in the build. This is by design - not all the plugins have build files, and of the ones that do, not all are suitable for inclusion in the nightly build (viz. SUPPLE, Section 17.11).

Note that if you are currently building gate by doing ‘ant jar’, be aware that this does not build the plugins. Running just ‘ant’ or ‘ant all’ will do so.

12.4 Updating this User Guide

The GATE User Guide is maintained in the GATE subversion repository at SourceForge. If you are a developer at Sheffield you do not need to check out the userguide explicitly, as it will appear under the tao directory when you check out sale. For others, you can check it out as follows:

```
svn checkout https://svn.sourceforge.net/svnroot/gate/userguide/trunk userguide
```

The user guide is written in L^AT_EX and translated to PDF using `pdflatex` and to HTML using `tex4ht`. The main file that ties it all together is `tao_main.tex`, which defines the various macros used in the rest of the guide and `\inputs` the other `.tex` files, one per chapter.

12.4.1 Building the User Guide

You will need:

- A standard POSIX shell environment including GNU Make. On Windows this generally means Cygwin, on Mac OS X the XCode developer tools and on Unix the relevant packages from your distribution.
- A copy of the userguide sources (see above).
- A L^AT_EX installation, including `pdflatex` if you want to build the PDF version, and `tex4ht` if you want to build the HTML. MiKTeX should work for Windows, `texlive` (available in MacPorts) for Mac OS X, or your choice of package for Unix.
- The BibTeX database `big.bib`. It must be located in the directory **above** where you have checked out the userguide, i.e. if the guide sources are in `/home/bob/svn/userguide` then `big.bib` needs to go in `/home/bib/svn`. Sheffield developers will find that it is already in the right place, under `sale`, others will need to download it from <http://gate.ac.uk/sale/big.bib>.
- The file <http://gate.ac.uk/sale/utills.tex>.
- A bit of luck.

Once these are all assembled it *should* be a case of running `make` to perform the actual build. To build the PDF do `make tao.pdf`, for the one page HTML do `make index.html` and for the several pages HTML do `make split.html`.

The PDF build generally works without problems, but the HTML build is known to hang on some machines for no apparent reason. If this happens to you try again on a different machine.

12.4.2 Making Changes to the User Guide

To make changes to the guide simply edit the relevant `.tex` files, make sure the guide still builds (at least the PDF version), and check in your changes **to the source files only**. Please do not check in your own built copy of the guide, the official user guide builds are produced by a Hudson continuous integration server in Sheffield.

If you add a section or subsection you should use the `\sect` or `\subsection` commands rather than the normal LaTeX `\section` or `\subsection`. These shorthand commands take an optional first parameter, which is the label to use for the section and should follow the pattern of existing labels. The label is also set as an anchor in the HTML version of the guide. For example a new section for the ‘Fish’ plugin would go in `misc-creole.tex` with a heading of:

```
\sect[sec:misc-creole:fish]{The Fish Plugin}
```

and would have the persistent URL `http://gate.ac.uk/userguide/sec:misc-creole:fish`.

If your changes are to document a bug fix or a new (or removed) feature then you should also add an entry to the change log in `recent-changes.tex`. You should include a reference to the full documentation for your change, in the same way as the existing changelog entries do. You should find yourself adding to the changelog every time except where you are just tidying up or rewording existing documentation. Unlike in the other source files, if you add a section or subsection you should use the `\rcSect` or `\rcSubsect`. Recent changes appear both in the introduction and the appendix, so these commands enable nesting to be done appropriately.

Section/subsection labels should comprise ‘sec’ followed by the chapter label and a descriptive section identifier, each colon-separated. New chapter labels should begin ‘chap:’.

Try to avoid changing chapter/section/subsection labels where possible, as this may break links to the section. If you need to change a label, add it in the file ‘sections.map’. Entries in this file are formatted one per line, with the old section label followed by a tab followed by the new section label.

The quote marks used should be ‘ and ’.

Titles should be in title case (capitalise the first word, nouns, pronouns, verbs, adverbs and adjectives but not articles, conjunctions or prepositions). When referring to a numbered chapter, section, subsection, figure or table, capitalise it, e.g. ‘Section 3.1’. When merely using the words chapter, section, subsection, figure or table, e.g. ‘the next chapter’, do not capitalise them. Proper nouns should be capitalised (‘Java’, ‘Groovy’), as should strings where the capitalisation is significant, but not terms like ‘annotation set’ or ‘document’.

The user guide is rebuilt automatically whenever changes are checked in, so your change should appear in the online version of the guide within 20 or 30 minutes.

Part III

CREOLE Plugins

Chapter 13

Gazetteers

...neurobiologists still go on openly studying reflexes and looking under the hood, not huddling passively in the trenches. Many of them still keep wondering: how does the inner life arise? Ever puzzled, they oscillate between two major fictions: (1) The brain can be understood; (2) We will never come close. Meanwhile they keep pursuing brain mechanisms, partly from habit, partly out of faith. Their premise: The brain is the organ of the mind. Clearly, this three-pound lump of tissue is the source of our ‘insight information’ about our very being. Somewhere in it there might be a few hidden guidelines for better ways to lead our lives.

Zen and the Brain, James H. Austin, 1998 (p. 6).

13.1 Introduction to Gazetteers

A gazetteer consists of a set of lists containing names of entities such as cities, organisations, days of the week, etc. These lists are used to find occurrences of these names in text, e.g. for the task of named entity recognition. The word ‘gazetteer’ is often used interchangeably for both the set of entity lists and for the processing resource that makes use of those lists to find occurrences of the names in text.

When a gazetteer processing resource is run on a document, annotations of type Lookup are created for each matching string in the text. Gazetteers usually do not depend on Tokens or on any other annotation and instead find matches based on the textual content of the document. (the **Flexible Gazetteer**, described in section 13.7, being the exception to the rule). This means that an entry may span more than one word and may start or end within a word. If a gazetteer that directly works on text does respect word boundaries, the way how word boundaries are found might differ from the way the GATE tokeniser finds word boundaries. A Lookup annotation will only be created if the entire gazetteer entry is matched in the text. The details of how gazetteer entries match text depend on the gazetteer

processing resource and its parameters. In this chapter, we will cover several gazetteers.

13.2 ANNIE Gazetteer

The rest of this introductory section describes the **ANNIE Gazetteer** which is part of ANNIE and also described in section 6.3. The ANNIE gazetteer is part of and proved by the ANNIE plugin.

Each individual gazetteer list is a plain text file, with one entry per line.

Below is a section of the list for units of currency:

```
Ecu
European Currency Units
FFr
Fr
German mark
German marks
New Taiwan dollar
New Taiwan dollars
NT dollar
NT dollars
```

An index file (usually called `lists.def`) is used to describe all such gazetteer list files that belong together. Each gazetteer list should reside in the same directory as the index file.

The gazetteer index files describes for each list the major type and optionally, a minor type and a language, separated by colons. In the example below, the first column refers to the list name, the second column to the major type, and the third to the minor type. These lists are compiled into finite state machines. Any text strings matched by these machines will be annotated with features specifying the major and minor types.

```
currency_prefix.lst:currency_unit:pre_amount
currency_unit.lst:currency_unit:post_amount
date.lst:date:specific_date
day.lst:date:day
monthen.lst:date:month:en
monthde.lst:date:month:de
season.lst:date:season
```

The major and minor type as well as the language will be added as features to only Lookup annotation generated from a matching entry from the respective list. For example, if an entry

from the `currency_unit.lst` gazetteer list matches some text in a document, the gazetteer processing resource will generate a Lookup annotation spanning the matching text and assign the features `major="currency_unit"` and `minor="post_amount"` to that annotation.

Grammar rules (JAPE rules) can specify the types to be identified in particular circumstances. The major and minor types enable this identification to take place, by giving access to items stored in particular lists or combinations of lists.

For example, if a day needs to be identified, the minor type ‘day’ would be specified in the grammar, in order to match only information about specific days. If any kind of date needs to be identified, the major type ‘date’ would be specified. This might include weeks, months, years etc. as well as days of the week, and would give access to all the items stored in `day.lst`, `month.lst`, `season.lst`, and `date.lst` in the example shown.

13.2.1 Creating and Modifying Gazetteer Lists

Gazetteer lists can be modified using any text editor. Use of an editor that can edit Unicode UTF-8 files (e.g. the GATE Unicode editor) is advised, however, in order to ensure that the lists are stored as UTF-8, which will minimise any language encoding problems, particularly if e.g. accents, umlauts or characters from non-Latin scripts are present.

To create a new list, simply add an entry for that list to the definitions file and add the new list in the same directory as the existing lists.

After any modifications have been made, ensure that you reinitialise the gazetteer PR in GATE, if one is already loaded, before rerunning your application.

13.3 Gazetteer Visual Resource - GAZE

Gaze is a tool for editing the gazetteer lists, definitions and mapping to ontology. It is suitable for use both for Plain/Linear Gazetteers (Default and Hash Gazetteers) and Ontology-enabled Gazetteers (OntoGazetteer). The Gazetteer PR associated with the viewer is reinitialised every time a save operation is performed. Note that GAZE does not scale up to very large lists (we suggest not using it to view over 40,000 entries and not to copy inside more than 10,000 entries).

Gaze is part of and provided by the ANNIE plugin. To make it possible to visualize gazetteers with the Gaze visualizer, the ANNIE plugin must be loaded first. Double clicking on a gazetteer PR that uses a gazetteer definition (index) file will display the contents of the gazetteer in the main window. The first pane will display the definition file, while the right pane will display whichever gazetteer list has been selected from it.

A gazetteer list can be modified simply by typing in it. It can be saved by clicking the Save

button. When a list is saved, the whole gazetteer is automatically reinitialised (and will be ready for use in GATE immediately).

To edit the definition file, right click inside the pane and choose from the options (Inset, Edit, Remove). A pop-up menu will appear to guide you through the remaining process. Save the definition file by selecting Save. Again, the gazetteer will be reinitialised automatically.

13.3.1 Display Modes

The display mode depends on the type of gazetteer loaded in the VR. The mode in which Linear/Plain Gazetteers are loaded is called Linear/Plain Mode. In this mode, the Linear Definition is displayed in the left pane, and the Gazetteer List is displayed in the right pane. The Ontology/Extended mode is on when the displayed gazetteer is ontology-aware, which means that there exists a mapping between classes in the ontology and lists of phrases. Two more panes are displayed when in this mode. On the top in the left-most pane there is a tree view of the ontology hierarchy, and at the bottom the mapping definition is displayed. This section describes the Linear/Plain display mode, the Ontology/Extended mode is described in section 13.5.

Whenever a gazetteer PR that uses a gazetteer definition (index) file is loaded, the Gaze gazetteer visualisation will appear on double-click over the gazetteer in the Processing Resources branch of the Resources Tree.

13.3.2 Linear Definition Pane

This pane displays the nodes of the linear definition, and allows manipulation of the whole definition as a file, as well as the single nodes. Whenever a gazetteer list is modified, its node in the linear definition is coloured in red.

13.3.3 Linear Definition Toolbar

All the functionality explained in this section (New, Load, Save, Save As) is accessible also via File — Linear Definition in the menu bar of Gaze.

New – Pressing New invokes a file dialog where the location of the new definition is specified.

Load – Pressing Load invokes a file dialog, and after locating the new definition it is loaded by pressing Open.

Save – Pressing Save saves the definition to the location from which it has been read.

Save As – Pressing Save As allows another location to be chosen, and the definition saved

there.

13.3.4 Operations on Linear Definition Nodes

Double-click node – Double-clicking on a definition node forces the displaying of the gazetteer list of the node in the right-most pane of the viewer.

Insert – On right-click over a node and choosing Insert, a dialog is displayed, requesting List, Major Type, Minor Type and Languages. The mandatory fields are List and Major Type. After pressing OK, a new linear node is added to the definition.

Remove – On right-click over a node and choosing Remove, the selected linear node is removed from the definition.

Edit – On right-click over a node and choosing Edit a dialog is displayed allowing changes of the fields List, Major Type, Minor Type and Languages.

13.3.5 Gazetteer List Pane

The gazetteer list pane has a toolbar with similar to the linear definition's buttons (New, Load, Save, Save As). They work as predicted by their names and as explained in the Linear Definition Pane section, and are also accessible from File / Gazetteer List in the menu bar of Gaze. The only addition is Save All which saves all modified gazetteer lists. The editing of the gazetteer list is as simple as editing a text file. One could use Ctrl+A to select the whole list, Ctrl+C to copy the selected, Ctrl+V to paste it, Del to delete the selected text or a single character, etc.

13.3.6 Mapping Definition Pane

The mapping definition is displayed one mapping node per row. It consists of a gazetteer list, ontology URL, and class id. The content of the gazetteer list in the node is accessible through double-clicking. It is displayed in the Gazetteer List Pane. The toolbar allows the creation of a new definition (New), the loading of an existing one (Load), saving to the same or new location (Save/Save As). The functionality of the toolbar buttons is also available via File.

13.4 OntoGazetteer

The Ontogazetteer, or Hierarchical Gazetteer, is a processing resource which can associate the entities from a specific gazetteer list with a class in a GATE ontology language resource. The OntoGazetteer assigns classes rather than major or minor types, and is aware of mappings between lists and class IDs. The Gaze visual resource can display the lists, ontology mappings and the class hierarchy of the ontology for a OntoGazetteer processing resource and provides ways of editing these components.

13.5 Gaze Ontology Gazetteer Editor

This section describes the Gaze gazetteer editor when it displays an OntoGazetteer processing resource. The editor consists of two parts: one for the editing of the lists and the mapping of lists and one for editing the ontology. These two parts are described in the following subsections.

13.5.1 The Gaze Gazetteer List and Mapping Editor

This is a VR for editing the gazetteer lists, and mapping them to classes in an ontology. It provides load/store/edit for the lists, load/store/edit for the mapping information, loading of ontologies, load/store/edit for the linear definition file, and mapping of the lists file to the major type, minor type and language.

Left pane: A single ontology is visualized in the left pane of the VR. The mapping between a list and a class is displayed by showing the list as a subclass with a different icon. The mapping is specified by drag and drop from the linear definition pane (in the middle) and/or by right click menu.

Middle pane: The middle pane displays the nodes/lines in the linear definition file. By double clicking on a node the corresponding list is opened. Editing of the line/node is done by right clicking and choosing edit: a dialogue appears (lower part of the scheme) allowing the modification of the members of the node.

Right pane: In the right pane a single gazetteer list is displayed. It can be edited and parts of it can be cut/copied/pasted.

13.5.2 The Gaze Ontology Editor

Note: to edit ontologies within gate, the more recent ontology viewer editor provided by the `Ontology_Tools` which provides many more features can be used, see section 14.5.

This is a VR for editing the class hierarchy of an ontology. it provides storing to and loading from RDF/RDFS, and provides load/edit/store of the class hierarchy of an ontology.

Left pane: The various ontologies loaded are listed here. On double click or right click and edit from the menu the ontology is visualized in the Right pane.

Right pane: Besides the visualization of the class hierarchy of the ontology the following operations are allowed:

- expanding/collapsing parts of the ontology
- adding a class in the hierarchy: by right clicking on the intended parent of the new class and choosing add sub class.
- removing a class: via right clicking on the class and choosing remove.

As a result of this VR, the ontology definition file is affected/altered.

13.6 Hash Gazetteer

The Hash Gazetteer is a gazetter implemented by the OntoText Lab (<http://www.ontotext.com/>). Its implementation is based on simple lookup in several `java.util.HashMap` objects, and is inspired by the strange idea of Atanas Kiryakov, that searching in HashMaps will be faster than a search in a Finite State Machine (FSM). The Hash Gazetteer processing resource is part of the ANNIE plugin.

This gazetter processing resource is implemented in the following way: Every phrase i.e. every list entry is separated into several parts. The parts are determined by the whitespaces lying among them. e.g. the phrase : "form is emptiness" has three parts : "form", "is", and "emptiness". There is also a list of HashMaps: `mapsList` which has as many elements as the longest (in terms of 'count of parts') phrase in the lists. So the first part of a phrase is placed in the first map. The first part + space + second part is placed in the second map, etc. The full phrase is placed in the appropriate map, and a reference to a Lookup object is attached to it.

On first sight it seems that this algorithm is certainly much more memory-consuming than a finite state machine (FSM) with the parts of the phrases as transitions, but this is actually not so important since the average length of the phrases (in parts) in the lists is 1.1. On the other hand, one advantage of the algorithm is that, although unconventional, on average it takes four times less memory and works three times faster than an optimized FSM implementation.

13.6.1 Prerequisites

The phrases to be recognised should be listed in a set of files, one for each type of occurrence (as for the standard gazetteer).

The gazetteer is built with the information from a file that contains the set of lists (which are files as well) and the associated type for each list. The file defining the set of lists should have the following syntax: each list definition should be written on its own line and should contain:

- the file name (required)
- the major type (required)
- the minor type (optional)
- the language(s) (optional)

The elements of each definition are separated by ‘:’. The following is an example of a valid definition:

```
personmale.lst:person:male:english
```

Each file named in the lists definition file is just a list containing one entry per line.

When this gazetteer is run over some input text (a GATE document) it will generate annotations of type Lookup having the attributes specified in the definition file.

13.6.2 Parameters

The Hash Gazetteer processing resource allows the specification of the following parameters when it is created:

caseSensitive: this can be switched between `true` and `false` to indicate if matches should be done in a case-sensitive way.

encoding: the encoding of the gazetteer lists

listsURL: the URL of the list definitions (index) file, i.e. the file that contains the filenames, major types and optionally minor types and languages of all the list files.

There is one run-time parameter, **annotationSetName** that allows the specification of the annotation set in which the Lookup annotations will be created. If nothing is specified the default annotation set will be used.

13.7 Flexible Gazetteer

The Flexible Gazetteer provides users with the flexibility to choose their own customized input and an external Gazetteer. For example, the user might want to replace words in the text with their base forms (which is an output of the Morphological Analyser) or to segment a Chinese text (using the Chinese Tokeniser) before running the Gazetteer on the Chinese text.

The Flexible Gazetteer performs lookup over a document based on the values of an arbitrary feature of an arbitrary annotation type, by using an *externally provided* gazetteer. It is important to use an external gazetteer as this allows the use of any type of gazetteer (e.g. an Ontological gazetteer).

Input to the Flexible Gazetteer:

Runtime parameters:

- Document – the document to be processed
- **inputAnnotationSetName** The annotationSet where the Flexible Gazetteer should search for the AnnotationType.feature specified in the inputFeatureNames.
- **outputAnnotationSetName** The AnnotationSet where Lookup annotations should be placed.

Creation time parameters:

- **inputFeatureNames** – when selected, these feature values are used to replace the corresponding original text. A temporary document is created from the values of the specified features on the specified annotation types. For example: for Token.string the temporary document will have the same content as the original one but all the SpaceToken annotations will have been replaced by single spaces.
- **gazetteerInst** – the actual gazetteer instance, which should run over a temporary document. This generates the Lookup annotations with features. This must be an instance of `gate.creole.gazetteer.Gazetteer` which has already been created. All such instances will be shown in the dropdown menu for this parameter in GATE Developer.

Once the external gazetteer has annotated text with Lookup annotations, Lookup annotations on the temporary document are converted to Lookup annotations on the original document. Finally the temporary document is deleted.

13.8 Gazetteer List Collector

The gazetteer list collector collects occurrences of entities directly from a set of annotated training texts, and populates gazetteer lists with the entities. The entity types and structure of the gazetteer lists are defined as necessary by the user. Once the lists have been collected, a semantic grammar can be used to find the same entities in new texts.

An empty list must be created first for each annotation type, if no list exists already. The set of lists must be loaded into GATE before the PR can be run. If a list already exists, the list will simply be augmented with any new entries. The list collector will only collect one occurrence of each entry: it first checks that the entry is not present already before adding a new one.

There are 4 runtime parameters:

- `annotationTypes`: a list of the annotation types that should be collected
- `gazetteer`: the gazetteer where the results will be stored (this must be already loaded in GATE)
- `markupASname`: the annotation set from which the annotation types should be collected
- `theLanguage`: sets the language feature of the gazetteer lists to be created to the appropriate language (in the case where lists are collected for different languages)

Figure 13.1 shows a screenshot of a set of lists collected automatically for the Hindi language. It contains 4 lists: Person, Organisation, Location and a list of stopwords. Each list has a `majorType` whose value is the type of list, a `minorType` 'inferred' (since the lists have been inferred from the text), and the language 'Hindi'.

The list collector also has a facility to split the Person names that it collects into their individual tokens, so that it adds both the entire name to the list, and adds each of the tokens to the list (i.e. each of the first names, and the surname) as a separate entry. When the grammar annotates Persons, it can require them to be at least 2 tokens or 2 consecutive Person Lookups. In this way, new Person names can be recognised by combining a known first name with a known surname, even if they were not in the training corpus. Where only a single token is found that matches, an Unknown entity is generated, which can later be matched with an existing longer name via the orthomatcher component which performs orthographic coreference between named entities. This same procedure can also be used for other entity types. For example, parts of Organisation names can be combined together in different ways. The facility for splitting Person names is hardcoded in the file `gate/src/gate/creole/GazetteerListsCollector.java` and is commented.

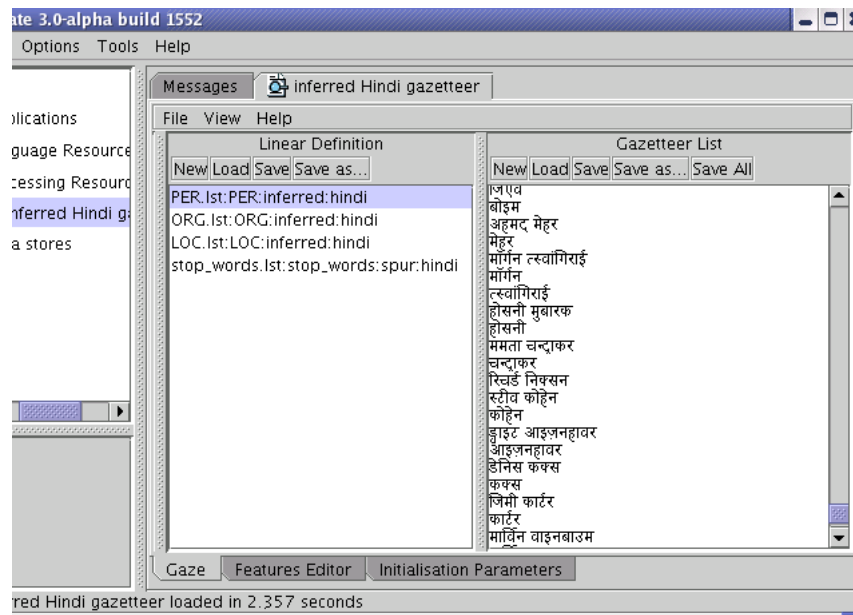


Figure 13.1: Lists collected automatically for Hindi

13.9 OntoRoot Gazetteer

OntoRoot Gazetteer is a type of a dynamically created gazetteer that is, in combination with few other generic GATE resources, capable of producing ontology-based annotations over the given content with regards to the given ontology. This gazetteer is a part of ‘Gazetteer_Ontology_Based’ plugin that has been developed as a part of the TAO project.

13.9.1 How Does it Work?

To produce ontology-based annotations i.e. annotations that link to the specific concepts or relations from the ontology, it is essential to pre-process the Ontology Resources (e.g., Classes, Instances, Properties) and extract their human-understandable lexicalisations.

As a precondition for extracting human-understandable content from the ontology, first a list of the following is being created:

- names of all ontology resources i.e. fragment identifiers ¹ and

¹An ontology resource is usually identified by an URI concatenated with a set of characters starting with ‘#’. This set of characters is called *fragment identifier*. For example, if the URI of a class representing *GATE POS Tagger* is: ‘http://gate.ac.uk/ns/gate-ontology#POSTagger’, the fragment identifier will be ‘POSTagger’.

- assigned property values for all ontology resources (e.g., label and datatype property values)

Each item from the list is further processed so that:

- any name containing dash ("-") or underline ("_") character(s) is processed so that each of these characters is replaced by a blank space. For example, `Project_Name` or `Project-Name` would become a `Project Name`.
- any name that is written in *camelCase* style is actually split into its constituent words, so that `ProjectName` becomes a `Project Name` (optional).
- any name that is a compound name such as 'POS Tagger for Spanish' is split so that both 'POS Tagger' and 'Tagger' are added to the list for processing. In this example, 'for' is a stop word, and any words after it are ignored (optional).

Each item from this list is analysed separately by the Onto Root Application (ORA) on execution (see figure 13.2). The Onto Root Application first tokenises each linguistic term, then assigns part-of-speech and lemma information to each token.

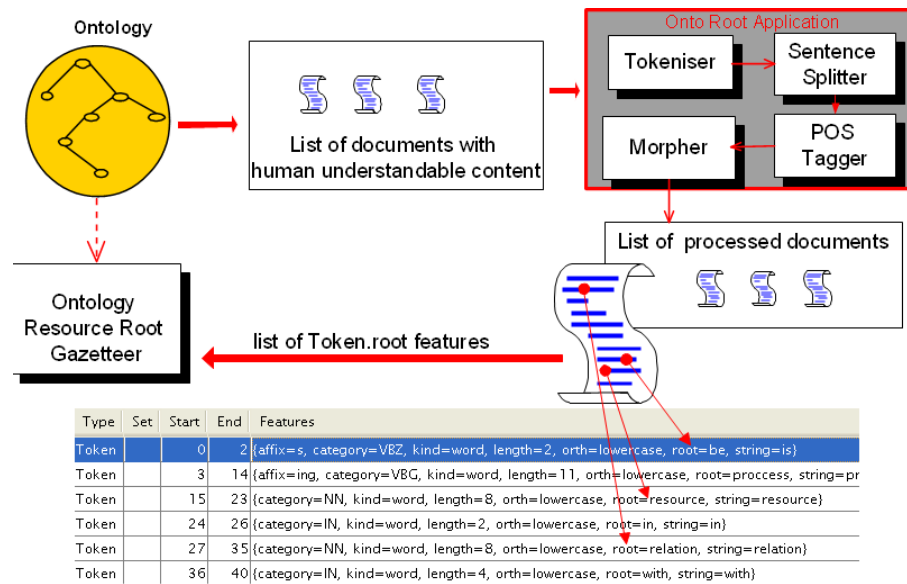


Figure 13.2: Building Ontology Resource Root (OntoRoot) Gazetteer from the Ontology

As a result of that pre-processing, each token in the terms will have additional feature named ‘root’, which contains the lemma as created by the morphological analyser. It is this lemma or a set of lemmas which are then added to the dynamic gazetteer list, created from the ontology.

For instance, if there is a resource with a short name (i.e., fragment identifier) *ProjectName*, without any assigned properties the created list before executing the OntoRoot gazetteer collection will contain the following strings:

- ‘*ProjectName*’,
- ‘*Project Name*’ after separating camelCased word and
- ‘*Name*’ after applying heuristic rules.

Each of the item from the list is then analysed separately and the results would be the same as the input strings, as all of entries are nouns given in singular form.

13.9.2 Initialisation of OntoRoot Gazetteer

To initialise the gazetteer there are few mandatory parameters:

- *Ontology* to be processed;

- *Tokeniser*, *POS Tagger* and *GATE Morphological Analyser* to be used during processing.

and few optional ones:

- *useResourceUri*, default is set to true - should this gazetteer analyse resource URIs or not;
- *considerProperties*, default is set to true - should this gazetteer consider properties or not;
- *propertiesToInclude* - checked only if *considerProperties* is set to true - this parameter contains the list of property names (URIs) to be included, comma separated;
- *propertiesToExclude* - checked only if *considerProperties* is set to true - this parameter contains the list of property names to be excluded, comma separated;
- *caseSensitive*, default set to be false -should this gazetteer differentiate on case;
- *separateCamelCasedWords*, default set to true - should this gazetteer separate emph-camelCased words, e.g. 'ProjectName' into 'Project Name';
- *considerHeuristicRules*, default set to false - should this gazetteer consider several heuristic rules or not. Rules include splitting the words containing spaces, and using prepositions as stop words; for example, if 'pos tagger for Spanish' would be analysed, 'for' would be considered as a stop word; heuristically derived would be 'pos tagger' and this would be further used to add 'pos tagger' to the gazetteer list, with a feature emphheuristic level set to be 0, and 'tagger' with emphheuristic level 1; at runtime lower heuristical level should be preferred. NOTE: setting *considerHeuristicRules* to true can cause a lot of noise for some ontologies and is likely to require implementing an additional filtering resource that will prefer the annotations with the lower heuristic level;

13.9.3 Simple steps to run OntoRoot Gazetteer

OntoRoot Gazetteer is a part of the Gazetteer_Ontology_Based plugin.

Easy way

For a quick start with the OntoRoot Gazetteer, consider running it from the GATE Developer (GATE GUI):

- Start GATE




Figure 13.3: Sample ontology-based annotation as a result of running OntoRoot Gazetteer. Feature *URI* refers to the URI of the ontology resource, while *type* identifies the type of the resource such as *class*, *instance*, *property*, or *datatypePropertyValue*

- Load a sample application from resources folder (exampleApp.xgapp). This will load *CAT App* application.
- Run *CAT App* application and open *query-doc* to see a set of Lookup annotations generated as a result (see Figure 13.3).

Hard way

OntoRoot Gazetteer can easily be set up to be used with any ontology. To generate a GATE application which demonstrates the use of the OntoRoot Gazetteer, follow these steps:

1. Start GATE
2. Load necessary plugins: Click on *Manage CREOLE plugins* and check the following:
 - Tools
 - Ontology
 - Ontology_Based_Gazetteer
 - Ontology_Tools (optional); this parameter is required in order to view ontology using the GATE Ontology Editor.
 - ANNIE.

Make sure that these plugins are loaded from GATE/plugins/[plugin_name] folder.

3. Load an ontology. Right click on *Language Resource*, and select the last option to create an *OWLIM Ontology LR*. Specify the format of the ontology, for example *rdxmlURL*,

and give the correct path to the ontology: either the absolute path on your local machine such as `c:/myOntology.owl` or the URL such as `http://gate.ac.uk/ns/gate-ontology`. Specify the *name* such as *myOntology* (this is optional).

4. Create Processing Resources: Right click on the *Processing Resource* and create the following PRs (with default parameters):
 - Document Reset PR
 - ANNIE English Tokeniser
 - ANNIE POS Tagger
 - GATE Morphological Analyser
 - RegEx Sentence Splitter (or ANNIE Sentence Splitter)
5. Create an *Onto Root Gazetteer* and set the init parameters. Mandatory ones are:
 - *Ontology*: select previously created myOntology;
 - *Tokeniser*: select previously created Tokeniser;
 - *POS Tagger*: select previously created POS Tagger;
 - *Morpher*: select previously created Morpher.

OntoRoot gazetteer is quite flexible in that it can be configured using the optional parameters. List of all parameters is detailed in Section 13.9.2.

When all parameters are set click OK. It can take some time to initialise OntoRoot Gazetteer. For example, loading GATE knowledge base from `http://gate.ac.uk/ns/gate-kb` takes around 6-15 seconds. Larger ontologies can take much longer.

6. Create another PR which is a Flexible Gazetteer. As init parameters it is mandatory to select previously created OntoRoot Gazetteer for *gazetteerInst*. For another parameter, *inputFeatureNames*, click on the button on the right and when prompted with a window, add 'Token.root' in the provided textbox, then click Add button. Click OK, give name to the new PR (optional) and then click OK.
7. Create an application. Right click on Application, then New Pipeline (or Corpus Pipeline). Add the following PRs to the application in this particular order:
 - Document Reset PR
 - RegEx Sentence Splitter (or ANNIE Sentence Splitter)
 - ANNIE English Tokeniser
 - ANNIE POS Tagger
 - GATE Morphological Analyser
 - Flexible Gazetteer

8. Create a document to process with the new application; for example, if the ontology was `http://gate.ac.uk/ns/gate-kb`, then the document could be the GATE home page: `http://gate.ac.uk`. Run application and then investigate the results further. All annotations are of type *Lookup*, with additional features that give details about the resources they are referring to in the given ontology.

13.10 Large KB Gazetteer

The large KB gazetteer provides support for ontology-aware NLP. You can load any ontology from RDF and then use the gazetteer to obtain lookup annotations that have both instance and class URI.

The large KB gazetteer is available as the plugin `Gazetteer_LKB`.

The current version of the large KB gazetteer does not use GATE ontology language resources. Instead, it uses its own mechanism to load and process ontologies. The current version is likely to change significantly in the near future.

The Large KB gazetteer grew from a component in the semantic search platform Ontotext KIM. The gazetteer is developed by people from the KIM team (see http://nmwiki.ontotext.com/lkb_gazetteer/team-list.html). You may find the name *kim* left in several places in the source code, documentation or source files.

13.10.1 Quick usage overview

- To use the Large KB gazetteer, set up your dictionary first. The dictionary is a folder with some configuration files. Use the samples at `GATE_HOME/plugins/Gazetteer_LKB/samples` as a guide or download a prebuilt dictionary from http://ontotext.com/kim/lkb_gazetteer/dictionaries.
- Load `GATE_HOME/plugins/Gazetteer_LKB` as a CREOLE plugin. See Section 3.5 for details.
- Create a new ‘Large KB Gazetteer’ processing resource (PR). Put the folder of the dictionary you created in the ‘dictionaryPath’ parameter. You can leave the rest of the parameters as defaults.
- Add the PR to your GATE application. The gazetteer doesn’t require a tokenizer or the output of any other processing resources.
- The gazetteer will create annotations with type ‘Lookup’ and two features; ‘inst’, which contains the URI of the ontology instance, and ‘class’ which contains the URI of the ontology class that instance belongs to.

13.10.2 Dictionary setup

The dictionary is a folder with some configuration files. You can find samples at *GATE_HOME/plugins/Gazetteer_LKB/samples*.

Setting up your own dictionary is easy. You need to define your RDF ontology and then specify a SPARQL or SERQL query that will retrieve a subset of that ontology as a dictionary.

config.ttl is a Turtle RDF file which configures a local RDF ontology or connection to a remote Sesame RDF database.

If you want to see examples of how to use local RDF files, please check *samples/dictionary_from_local_ontology/config.ttl*. The *Sesame repository configuration* section configures a local Ontotext SwiftOWLIM database that loads a list of RDF files. Simply create a list of your RDF files and reuse the rest of the configuration. The sample configuration support datasets with 10,000,000 triples with acceptable performance. For working with larger datasets, advanced users can substitute SwiftOWLIM with another Sesame RDF engine. In that case, make sure you add the necessary JARs to the list in *GATE_HOME/plugins/Gazetteer_LKB/creole.xml*. For example, Ontotext BigOWL is a Sesame RDF engine that can load billions of triples on desktop hardware.

Since any Sesame repository can be configured in *config.ttl*, the Large KB Gazetteer can extract dictionaries from **all** significant RDF databases. See the page on database compatibility for more information.

query.txt contains a SPARQL query. You can write any query you like, as long as its projection contains at least two columns in the following order: label and instance. As an option, you can also add a third column for the ontology class of the RDF entity. Below you can see a sample query, which creates a dictionary from the names and the unique identifiers of 10,000 entertainers in DbPedia.

```
PREFIX opencyc: <http://sw.opencyc.org/2008/06/10/concept/en/>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>

SELECT ?Name ?Person WHERE {
    ?Person a opencyc:Entertainer ; rdfs:label ?Name .
    FILTER (lang(?Name) = "en")
} LIMIT 10000
```

Try this query at the Linked Data Semantic Repository.

When you load the dictionary configuration in GATE for the first time, it creates a binary snapshot of the dictionary. Thereafter it will load only this binary snapshot. If the dictionary configuration is changed, the snapshot will be reinitialized automatically. For more information, please see the dictionary lifecycle specification.

13.10.3 Additional dictionary configuration

The `config.ttl` may contain additional dictionary configuration. Such configuration concerns only the initial loading of the dictionary from the RDF database. The options are still being determined and more will appear in future versions. They must be placed below the repository configuration section as attributes of a dictionary configuration. Here is a sample `config.ttl` file with additional configuration.

```
# Sesame configuration template for a (proxy for a) remote repository
#
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#>.
@prefix rep: <http://www.openrdf.org/config/repository#>.
@prefix hr: <http://www.openrdf.org/config/repository/http#>.
@prefix lkgb: <http://www.ontotext.com/lkb_gazetteer#>.

[] a rep:Repository ;
  rep:repositoryImpl [
    rep:repositoryType "openrdf:HTTPRepository" ;
    hr:repositoryURL <http://ldsr.ontotext.com/openrdf-sesame/repositories/owlim>
  ] ;
  rep:repositoryID "owlim" ;
  rdfs:label "LDSR" .

[] a lkgb:DictionaryConfiguration ;
  lkgb:caseSensitivity "CASE_INSENSITIVE" .
```

13.10.4 Processing Resource Configuration

The following options can be set when the gazetteer PR is initialized:

- `dictionaryPath`; the dictionary folder described above.
- `forceCaseSensitive`; whether the gazetteer should return case-sensitive matches regardless of the loaded dictionary.

13.10.5 Runtime configuration

- `annotationSetName` - The annotation set, which will receive the generated lookup annotations.
- `annotationLimit` - The maximum number of the generated annotations. NULL or 0 for no limit. Setting limit of the number of the created annotations will reduce the memory consumption of GATE on large documents. Note that GATE documents

consume gigabytes of memory if there are tens of thousands of annotations in the document. All PRs that create large number of annotations like the gazetteers and tokenizers may cause an Out Of Memory error on large texts. Setting that options limits the amount of memory that the gazetteer will use.

13.10.6 Semantic Enrichment PR

The Semantic Enrichment PR allows adding new data to semantic annotations by querying external RDF (Linked Data) repositories. It is a companion to the large KB gazetteer that showcases the usefulness of using Linked Data URI as identifiers.

Here a semantic annotation is an annotation that is linked to an RDF entity by having the URI of the entity in the 'inst' feature of the annotation. For all such annotation of a given type, this PR runs a SPARQL query against the defined repository and puts a comma-separated list of the values mentioned in the query output in the 'connections' feature of the same annotation.

There is a sample pipeline that features the Semantic Enrichment PR.

Parameters

- `inputASName`; the annotation set, which annotation will be processed.
- `server`; the URL of the Sesame 2 HTTP repository. Support for generic SPARQL endpoints can be implemented if required.
- `repositoryId`; the ID of the Sesame repository.
- `annotationTypes`; a list of types of annotation that will be processed.
- `query`; a SPARQL query pattern. The query will be processed like this - `String.format(query, uriFromAnnotation)`, so you can use parameters like `%s` or `%1$s`.
- `deleteOnNoRelations`; whether we want to delete the annotation that weren't enriched. Helps to clean up the input annotations.

13.11 The Shared Gazetteer for multithreaded processing

The `DefaultGazetteer` (and its subclasses such as the `OntoRootGazetteer`) compiles its gazetteer data into a finite state matcher at initialization time. For large gazetteers this FSM requires a considerable amount of memory. However, once the FSM has been built

then (as long as you do not modify it dynamically using Gaze) it is accessed in a read-only manner at runtime. For a multi-threaded application that requires several identical copies of its processing resources (see section 7.13), GATE provides a mechanism whereby a single compiled FSM can be shared between several gazetteer PRs that can then be executed concurrently in different threads, saving the memory that would otherwise be required to load the lists several times.

This feature is not available in the GATE Developer GUI, as it is only intended for use in embedded code. To make use of it, first create a single instance of the regular `DefaultGazetteer` or `OntoRootGazetteer`:

```
FeatureMap params = Factory.newFeatureMap();
params.put("listsUrl", listsDefLocation);
LanguageAnalyser mainGazetteer = (LanguageAnalyser)Factory.createResource(
    "gate.creole.gazetteer.DefaultGazetteer", params);
```

Then create any number of `SharedDefaultGazetteer` instances, passing this regular gazetteer as a parameter:

```
FeatureMap params = Factory.newFeatureMap();
params.put("bootstrapGazetteer", mainGazetteer);
LanguageAnalyser sharedGazetteer = (LanguageAnalyser)Factory.createResource(
    "gate.creole.gazetteer.SharedDefaultGazetteer", params);
```

The `SharedDefaultGazetteer` instance will re-use the FSM that was built by the `mainGazetteer` instead of loading its own.

Chapter 14

Working with Ontologies

GATE provides an API for modeling and manipulating ontologies and comes with two plugins that provide implementations for the API and several tools for editing ontologies and using ontologies for document annotation.

Ontologies in GATE are classified as language resources. In order to create an ontology language resource, the user must first *load one of the two plugins containing an ontology implementation*.

The following implementations and ontology related tools are provided as plugins:

- Plugin `Ontology_OWLIM2` provides an implementation that is fully backwards-compatible with the implementation that was part of GATE prior to version 5.1 (see Section 14.4).
- Plugin `Ontology` provides a modified and current implementation (see Section 14.3). Unless noted otherwise, all information in this chapter applies to this implementation.
- Plugin `Ontology_Tools` provides a simple graphical ontology editor (see Section 14.5) and OCAT, a tool for interactive ontology based document annotation (see Section 14.6). It also provides a gazetteer processing resource, `OntoGaz`, that allows the mapping of linear gazetteers to classes in an ontology (see Section 13.4).
- Plugin `Gazetteer_Ontology_Based` provides the ‘Onto Root Gazetteer’ for the automatic creating of a gazetteer from an ontology (see Section 13.9)
- Plugin `Ontology_BDM_Computation` can be used to compute BDM scores (see Section 10.6).
- Plugin `Gazetteer_LKB` provides a processing resource for creating annotations based on the contents of a large ontology.

GATE ontology support aims to simplify the use of ontologies both within the set of GATE tools and for programmers using the GATE ontology API. The GATE ontology API hides the details of the actual backend implementation and allows a simplified manipulation of ontologies by modeling ontology resources as easy-to-use Java objects. Ontologies can be loaded from and saved to various serialization formats.

The GATE ontology support roughly conforms to the representation, manipulation and inference that conforms to what is supported in OWL-Lite (see <http://www.w3.org/TR/owl-features/>). This means that a user can represent information in an ontology that conforms to OWL-Lite and that the GATE ontology model will provide inferred information equivalent to what an OWL-Lite reasoner would provide. The GATE ontology model makes an attempt to also to some extent provide useful information for ontologies that do not conform to OWL-Lite: RDFS, OWL-DL, OWL-Full or OWL2 ontologies can be loaded but GATE might ignore part of all contents of those ontologies, or might only provide part of, or incorrect inferred facts for such ontologies.

The GATE API tries to prevent clients from modifying an ontology that conforms to OWL-Lite to become OWL-DL or OWL-Full and also tries to prevent or warn about some of the most common errors that would make the ontology inconsistent. However, the current implementation is not able to prevent all such errors and has no way of finding out if an ontology conforms to OWL-Lite or is inconsistent.

14.1 Data Model for Ontologies

14.1.1 Hierarchies of Classes and Restrictions

Class hierarchy (or taxonomy) plays the central role in the ontology data model. This consists of a set of ontology classes (represented by `OCClass` objects in the ontology API) linked by `subClassOf`, `superClassOf` and `equivalentClassAs` relations. Each ontology class is identified by an URI (unless it is a restriction or an anonymous class, see below). The URI of each ontology resource must be unique.

Each class can have a set of superclasses and a set of subclasses; these are used to build the class hierarchy. The `subClassOf` and `superClassOf` relations are transitive and methods are provided by the API for calculating the transitive closure for each of these relations given a class. The transitive closure for the set of superclasses for a given class is a set containing all the superclasses of that class, as well as all the superclasses of its direct superclasses, and so on until no more are found. This calculation is finite, the upper bound being the set of all the classes in the ontology. A class that has no superclasses is called a *top class*. An ontology can have several top classes. Although the GATE ontology API can deal with cycles in the hierarchy graph, these can cause problems for processes using the API and probably indicate an error in the definition of the ontology. Also other components of GATE, like the ontology editor cannot deal with cyclic class structures and will terminate with an error. Care should

be taken to avoid such situations.

A pair of ontology classes can also have an `equivalentClassAs` relation, which indicates that the two classes are virtually the same and all their properties and instances should be shared.

A restriction (represented by `Restriction` objects in the GATE ontology API) is an anonymous class (i.e., the class is not identified by an URI/IRI) and is set on an object or a datatype property to restrict some instances of the specified domain of the property to have only certain values (also known as value constraint) or certain number of values (also known as cardinality restriction) for the property. Thus for each restriction there exists at least three triples in the repository. One that defines resource as a restriction, another one that indicates on which property the restriction is specified, and finally the third one that indicates what is the constraint set on the cardinality or value on the property. There are six types of restrictions:

1. *Cardinality* Restriction (`owl:cardinalityRestriction`): the only valid values for this restriction in OWL-Lite are 0 and 1. A cardinality restriction set to either 0 or 1 implies both a *MinCardinality* Restriction and a *MaxCardinality* Restriction set to the same value.
2. *MinCardinality* Restriction (`owl:minCardinalityRestriction`)
3. *MaxCardinality* Restriction (`owl:maxCardinalityRestriction`)
4. *HasValue* Restriction (`owl:hasValueRestriction`)
5. *AllValuesFrom* Restriction (`owl:allValuesFromRestriction`)
6. *SomeValuesFrom* Restriction (`owl:someValuesFromRestriction`)

Please visit the OWL Reference for more detailed information on restrictions.

14.1.2 Instances

Instances, also often called *individuals* are objects that belong to classes. Like named classes, each instance is identified by an URI. Each instance can belong to one or more classes and can have properties with values. Two instances can have the `sameInstanceAs` relation, which indicates that the property values assigned to both instances should be shared and that all the properties applicable to one instance are also valid for the other. In addition, there is a `differentInstanceAs` relation, which declares the instances as disjoint.

Instances are represented by `OInstance` objects in the API. API methods are provided for getting all the instances in an ontology, all the ones that belong to a given class, and all the property values for a given instance. There is also a method to retrieve a list of classes that the instance belongs to, using either transitive or direct closure.

14.1.3 Hierarchies of Properties

The last part of the data model is made up of hierarchies of properties that can be associated with objects in the ontology. The specification of the type of objects that properties apply to is done through the means of domains. Similarly, the types of values that a property can take are restricted through the definition of a range. A property with a domain that is an empty set can apply to instances of any type (i.e. there are no restrictions given). Like classes, properties can also have `superPropertyOf`, `subPropertyOf` and `equivalentPropertyAs` relations among them.

GATE supports the following property types:

1. Annotation Property:

An annotation property is associated with an ontology resource (i.e. a class, property or instance) and can have a *Literal* as value. A *Literal* is a Java object that can refer to the URI of any ontology resource or a string (`http://www.w3.org/2001/XMLSchema#string`) with the specified language or a data type (discussed below) with a compatible value. Two annotation properties can *not* be declared as equivalent. It is also not possible to specify a domain or range for an annotation property or a super or subproperty relation between two annotation properties. Five annotation properties, predefined by OWL, are made available to the user whenever a new ontology instance is created:

- `owl:versionInfo`,
- `rdfs:label`,
- `rdfs:comment`,
- `rdfs:seeAlso`, and
- `rdfs:isDefinedBy`.

In other words, even when the user creates an empty ontology, these annotation properties are created automatically and available to users.

2. Datatype Property:

A datatype property is associated with an ontology instance and can have a *Literal* value that is compatible with its data type. A data type can be one of the pre-defined data types in the GATE ontology API:

```
http://www.w3.org/2001/XMLSchema#boolean
http://www.w3.org/2001/XMLSchema#byte
http://www.w3.org/2001/XMLSchema#date
http://www.w3.org/2001/XMLSchema#decimal
http://www.w3.org/2001/XMLSchema#double
http://www.w3.org/2001/XMLSchema#duration
```

```
http://www.w3.org/2001/XMLSchema#float
http://www.w3.org/2001/XMLSchema#int
http://www.w3.org/2001/XMLSchema#integer
http://www.w3.org/2001/XMLSchema#long
http://www.w3.org/2001/XMLSchema#negativeInteger
http://www.w3.org/2001/XMLSchema#nonNegativeInteger
http://www.w3.org/2001/XMLSchema#nonPositiveInteger
http://www.w3.org/2001/XMLSchema#positiveInteger
http://www.w3.org/2001/XMLSchema#short
http://www.w3.org/2001/XMLSchema#string
http://www.w3.org/2001/XMLSchema#time
http://www.w3.org/2001/XMLSchema#unsignedByte
http://www.w3.org/2001/XMLSchema#unsignedInt
http://www.w3.org/2001/XMLSchema#unsignedLong
http://www.w3.org/2001/XMLSchema#unsignedShort
```

A set of ontology classes can be specified as a property's domain; in that case the property can be associated with the instance belonging to all of the classes specified in that domain only (the intersection of the set of domain classes).

Datatype properties can have other datatype properties as subproperties.

3. Object Property:

An object property is associated with an ontology instance and has an instance as value. A set of ontology classes can be specified as property's domain and range. Then the property can only be associated with the instances belonging to all of the classes specified as the domain. Similarly, only the instances that belong to all the classes specified in the range can be set as values.

Object properties can have other object properties as subproperties.

4. RDF Property:

RDF properties are more general than datatype or object properties. The GATE ontology API uses `RDFProperty` objects to hold datatype properties, object properties, annotation properties or actual RDF properties (`rdf:Property`).

Note: The use of `RDFProperty` objects for creating, or manipulating RDF properties is carried over from previous implementations for compatibility reasons but should be avoided.

All properties (except the annotation properties) can be marked as functional properties, which means that for a given instance in their domain, they can only take at most one value, i.e. they define a function in the algebraic sense. Properties inverse to functional properties are marked as *inverse functional*. If one likes ontology properties with algebraic relations, the semantics of these become apparent.

14.1.4 URIs

URIs are used to identify resources (instances, classes, properties) in an ontology. All URIs that identify classes, instances, or properties in an ontology must consist of two parts:

- a name part: this is the part after the last slash (#) or the first hash (#) in the URI. This part of the URI is often used as a shorthand name for the entity (e.g. in the ontology editor) and is often called a *fragment identifier*
- a namespace part: the part that precedes the name, including the trailing slash or hash character.

URIs uniquely identify resources: each resource can have at most one URI and each URI can be associated with at most one resource.

URIs are represented by `OURI` objects in the API. The `Ontology` object provides factory methods to create `OURIs` from a complete URI string or by appending a name to the default namespace of the ontology. However it is the responsibility of the caller to ensure that any strings that are passed to these factory methods do in fact represent valid URIs. GATE provides some helper methods in the `OUtils` class to help with encoding and decoding URI strings.

14.2 Ontology Event Model

An Ontology Event Model (OEM) is implemented and incorporated into the new GATE ontology API. Under the new OEM, events are fired when a resource is added, modified or deleted from the ontology.

An interface called `OntologyModificationListener` is created with five methods (see below) that need to be implemented by the listeners of ontology events.

```
public void resourcesRemoved(Ontology ontology, String[] resources);
```

This method is invoked whenever an ontology resource (a class, property or instance) is removed from the ontology. Deleting one resource can also result into the deletion of the other dependent resources. For example, deleting a class should also delete all its instances (more details on how deletion works are explained later). The second parameter, an array of strings, provides a list of URIs of resources deleted from the ontology.

```
public void resourceAdded(Ontology ontology, OResource resource);
```

This method is invoked whenever a new resource is added to the ontology. The parameters provide references to the ontology and the resource being added to it.

```
public void ontologyRelationChanged(Ontology ontology, OResource resource1,
                                   OResource resource2, int eventType);
```

This method is invoked whenever a relation between two resources (e.g. OClass and OClass, RDFProperty, RDFProperty, etc) is changed. Example events are addition or removal of a subclass or a subproperty, two classes or properties being set as equivalent or different and two instances being set as same or different. The first parameter is the reference to the ontology, the next two parameters are the resources being affected and the final parameter is the event type. Please refer to the list of events specified below for different types of events.

```
public void resourcePropertyValueChanged(Ontology ontology,
                                       OResource resource, RDFProperty
                                       property, Object value, int eventType)
```

This method is invoked whenever any property value is added or removed to a resource. The first parameter provides a reference to the ontology in which the event took place. The second provides a reference to the resource affected, the third parameter provides a reference to the property for which the value is added or removed, the fourth parameter is the actual value being set on the resource and the fifth parameter identifies the type of event.

```
public void ontologyReset(Ontology ontology)
```

This method is called whenever ontology is reset. In other words when all resources of the ontology are deleted using the ontology.cleanup method.

The OConstants class defines the static constants, listed below, for various event types.

```
public static final int OCLASS_ADDED_EVENT;
public static final int ANONYMOUS_CLASS_ADDED_EVENT;
public static final int CARDINALITY_RESTRICTION_ADDED_EVENT;
public static final int MIN_CARDINALITY_RESTRICTION_ADDED_EVENT;
public static final int MAX_CARDINALITY_RESTRICTION_ADDED_EVENT;
public static final int HAS_VALUE_RESTRICTION_ADDED_EVENT;
public static final int SOME_VALUES_FROM_RESTRICTION_ADDED_EVENT;
public static final int ALL_VALUES_FROM_RESTRICTION_ADDED_EVENT;
public static final int SUB_CLASS_ADDED_EVENT;
public static final int SUB_CLASS_REMOVED_EVENT;
public static final int EQUIVALENT_CLASS_EVENT;
```

```

public static final int ANNOTATION_PROPERTY_ADDED_EVENT;
public static final int DATATYPE_PROPERTY_ADDED_EVENT;
public static final int OBJECT_PROPERTY_ADDED_EVENT;
public static final int TRANSTIVE_PROPERTY_ADDED_EVENT;
public static final int SYMMETRIC_PROPERTY_ADDED_EVENT;
public static final int ANNOTATION_PROPERTY_VALUE_ADDED_EVENT;
public static final int DATATYPE_PROPERTY_VALUE_ADDED_EVENT;
public static final int OBJECT_PROPERTY_VALUE_ADDED_EVENT;
public static final int RDF_PROPERTY_VALUE_ADDED_EVENT;
public static final int ANNOTATION_PROPERTY_VALUE_REMOVED_EVENT;
public static final int DATATYPE_PROPERTY_VALUE_REMOVED_EVENT;
public static final int OBJECT_PROPERTY_VALUE_REMOVED_EVENT;
public static final int RDF_PROPERTY_VALUE_REMOVED_EVENT;
public static final int EQUIVALENT_PROPERTY_EVENT;
public static final int OINSTANCE_ADDED_EVENT;
public static final int DIFFERENT_INSTANCE_EVENT;
public static final int SAME_INSTANCE_EVENT;
public static final int RESOURCE_REMOVED_EVENT;
public static final int RESTRICTION_ON_PROPERTY_VALUE_CHANGED;
public static final int SUB_PROPERTY_ADDED_EVENT;
public static final int SUB_PROPERTY_REMOVED_EVENT;

```

An ontology is responsible for firing various ontology events. Object wishing to listen to the ontology events must implement the methods above and must be registered with the ontology using the following method.

```
addOntologyModificationListener(OntologyModificationListener oml);
```

The following method cancels the registration.

```
removeOntologyModificationListener(OntologyModificationListener oml);
```

14.2.1 What Happens when a Resource is Deleted?

Resources in an ontology are connected with each other. For example, one class can be a sub or superclass of another classes. A resource can have multiple properties attached to it. Taking these various relations into account, change in one resource can affect other resources in the ontology. Below we describe what happens (in terms of what does the GATE ontology API do) when a resource is deleted.

- When a class is deleted

- A list of all its super classes is obtained. For each class in this list, a list of its subclasses is obtained and the deleted class is removed from it.
 - All subclasses of the deleted class are removed from the ontology. A list of all its equivalent classes is obtained. For each class in this list, a list of its equivalent classes is obtained and the deleted class is removed from it.
 - All instances of the deleted class are removed from the ontology.
 - All properties are checked to see if they contain the deleted class as a member of their domain or range. If so, the respective property is also deleted from the ontology.
- When an instance is deleted
 - A list of all its same instances is obtained. For each instance in this list, a list of its same instances is obtained and the deleted instance is removed.
 - A list of all instances set as different from the deleted instance is obtained. For each instance in this list, a list of instances set as different from it is obtained and the deleted instance is removed.
 - All the instances of ontology are checked to see if any of their set properties have the deleted instance as value. If so, the respective set property is altered to remove the deleted instance.
 - When a property is deleted
 - A list of all its super properties is obtained. For each property in this list, a list of its sub properties is obtained and the deleted property is removed.
 - All sub properties of the deleted property are removed from the ontology.
 - A list of all its equivalent properties is obtained. For each property in this list, a list of its equivalent properties is obtained and the deleted property is removed.
 - All instances and resources of the ontology are checked to see if they have the deleted property set on them. If so the respective property is deleted.

14.3 The Ontology Plugin: Current Implementation

The plugin `Ontology` contains the current ontology API implementation. This implementation provides the additions and enhancements introduced into the GATE ontology API as of release 5.1. It is based on a backend that uses Sesame version 2 and OWLIM version 3.

Before any ontology-based functionality can be used, the plugin must be loaded into GATE. To do this in the GATE Developer GUI, select the ‘Manage CREOLE Plugins’ option from the ‘File’ menu and check the ‘Load now’ checkbox for the ‘Ontology’ plugin, then click OK. After this, the context menu for Language Resources will include the following ontology language resources:

- *OWLIMOntology*: this is the standard language resource to use in most situations. It allows the user to create a new ontology backed by files in a local directory and optionally load ontology data into it.
- *OWLIMOntology DEPRECATED*: this language resource has the same functionality as *OWLIMOntology* but uses the exactly same package and class name as the language resource in the plugin `Ontology_OWLIM2`. This LR is provided to allow an easier upgrade of existing pipelines to the new implementation but users should move the the *OWLIMOntology LR* as soon as possible.
- *ConnectSesameOntology*: This language resources allows the use of ontologies that are already stored in a Sesame2 repository which is either stored in a directory or accessible from a server. This is useful for quickly re-using a very large ontology that has been previously created as a persistent *OWLIMOntology* language resource.
- *CreateSesameOntology*: This language resource allows the user to create a new empty ontology by specifying the repository configuration for creating the sesame repository. **Note:** *This is for advanced uses only!*

Each of these language resources is explained in more detail in the following sections.

To make the plugin available to your GATE Embedded application, load the plugin prior to creating one of the ontology language resources using the following code:

```

1 // Find the directory for the Ontology plugin
2 File pluginHome =
3   new File(new File(Gate.getGateHome(), "plugins"), "Ontology");
4 // Load the plugin from that directory
5 Gate.getCreoleRegister().registerDirectories(pluginHome.toURI().toURL());

```

14.3.1 The OWLIMOntology Language Resource

The *OWLIMOntology* language resource is the main ontology language resource provided by the plugin and provides a similar functionality to the *OWLIMOntologyLR* language resource provided by the pre-5.1 implementation and provided by the `Ontology_OWLIM2` plugin from version 5.1 on. This language resource creates an in-memory store backed by files in a directory on the file system to hold the ontology data.

To create a new OWLIM Ontology resource, select ‘OWLIM Ontology’ from the right-click ‘New’ menu for language resources. A dialog as shown in Figure 14.1 appears with the following parameters to fill in or change:

- *Name* (optional): if no name is given, a default name will be generated, if an ontology is loaded from an URL, based on that URL, otherwise based on the language resource name.

- *baseURI* (optional): the URI to be used for resolving relative URI references in the ontology during loading.
- *dataDirectoryName* (optional): the name of an existing directory on the file system where the directory will be created that backs the ontology store. The name of the directory that will be created within the data directory will be `GATE_OWLIMOntology_` followed by a string representation of the system time. If this parameter is not specified, the value for system property `java.io.tmpdir` is used, if this is not set either an error is raised.
- *loadImports* (optional): either true or false. If set to false all ontology import specifications found in the loaded ontology are ignored. This parameter is ignored if no ontology is loaded when the language resource is created.
- *mappingsURL* (optional): the URL of a text file containing import mappings specifications. See below for a description of the mappings file. if no URL is specified, the GATE will interpret each import URI found as an URL and try to import the data from that URL.
- *persistent* (optional): true or false: if false, the directory created inside the data directory is removed when the language resource is closed, otherwise, that directory is kept. The `ConnectSesameOntology` language resource can be used at a later time to connect to such a directory and create an ontology language resource for it (see Section 14.3.2).
- *rdfXmlUrl* (optional): an URL specifying the location of an ontology in RDF/XML serialization format (see <http://www.w3.org/TR/rdf-syntax-grammar/>) from which to load initial ontology data from. The parameter name can be changed from `rdfXmlUrl` to `n3Url` to indicate N3 serialization format (see <http://www.w3.org/DesignIssues/Notation3.html>), to `ntriplesUrl` to indicate N-Triples format (see <http://www.w3.org/TR/2004/REC-rdf-testcases-20040210/#ntriples>), and to `turtleUrl` to indicate TURTLE serialization format (see <http://www.w3.org/TeamSubmission/turtle/>). If this is left blank, no ontology is loaded and an empty ontology language resource is created.




Figure 14.1: The New OWLIM Ontology Dialog

Note: you could create a language resource such as *OWLIM Ontology* from GATE Developer successfully, but you will not be able to browse/edit the ontology unless you loaded *Ontology Tools* plugin beforehand.

Additional ontology data can be loaded into an existing ontology language resource by selecting the ‘Load’ option from the language resource’s context menu. This will show the dialog shown in figure 14.2. The parameters in this dialog correspond to the parameters in the dialog for creating a new ontology with the addition of one new parameter: ‘load as import’. If this parameter is checked, the ontology data is loaded specifically as an ontology import. Ontology imports can be excluded from what is saved at a later time.

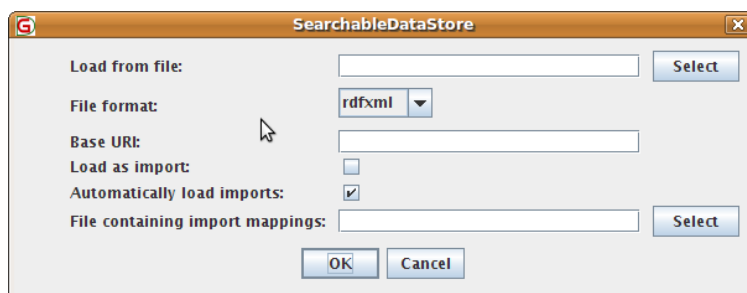


Figure 14.2: The Load Ontology Dialog

Figure 14.3 shows the ontology save dialog that is shown when the option ‘Save as...’ is selected from the language resource’s context menu. The parameter ‘include imports’ allows the user to specify if the data that has been loaded through imports should be included in the saved data or not.




Figure 14.3: The Save Ontology Dialog

14.3.2 The ConnectSesameOntology Language Resource

This ontology language resource can be created from either a directory on the local file system that holds an ontology backing store (as created in the ‘data directory’ for the ‘OWLIM Ontology’ language resource), or from a sesame repository on a server that holds an OWLIM ontology store.

This is very useful when using very large ontologies with GATE. Loading a very large ontology from a serialized format takes a significant amount of time because the file has to be deserialized and all implied facts have to get generated. Once an ontology has been loaded into a persisting `OWLIMOntology` language resource, the `ConnectSesameOntology` language resource can be used with the directory created to re-connect to the already de-serialized and inferred data much faster.

Figure 14.4 shows the dialog for creating a `ConnectSesameOntology` language resource.

- *repositoryID*: the name of the sesame repository holding the ontology store. For a backing store created with the ‘OWLIM Ontology’ language resource, this is always ‘owlim3’.
- *repositoryLocation*: the URL of the location where to find the repository holding the ontology store. For a backing store created with the ‘OWLIM Ontology’ language resource this is the directory that was created inside the data directory (the name of the directory starting with `GATE_OWLIMOntology_`). This can also contain any other name containing a Sesame2 repository or the URL of a Sesame2 server.

Note that this ontology language resource is only supported when connected with an OWLIM3 repository with the `owl-max` ruleset and `partialRDFS` optimizations disabled! Connecting to any other repository is experimental and for expert users only! Also note that connecting to a repository that is already in use by GATE or any other application is not supported and might result in unwanted or erroneous behavior!




Figure 14.4: The New `ConnectSesameOntology` Dialog

14.3.3 The `CreateSesameOntology` Language Resource

This ontology language resource can be directly created from a Sesame2 repository configuration file. This is an experimental language resource intended for expert users only. This can be used to create any kind of Sesame2 repository, but the only repository configuration supported by GATE and the GATE ontology API is an OWLIM repository with the `owl-max` ruleset and `partialRDFS` optimizations disabled. The dialog for creating this language resource is shown in Figure 14.5.




Figure 14.5: The New CreateSesameOntology Dialog

14.3.4 The OWLIM2 Backwards-Compatible Language Resource

This language resource is shown as “OWLIM Ontology DEPRECATED” in the ‘New Language Resource’ submenu from the ‘File’ menu. It provides the “OWLIM Ontology” language resource in a way that attempts maximum backwards-compatibility with the ontology language resource provided by prior versions or the `Ontology_OWLIM2` language resource. This means, the class name is identical to those language resources (`gate.creole.ontology.owlim.OWLIMOntologyLR`) and the parameters are made compatible. This means that the parameter `defaultNameSpace` is added as an alias for the parameter `baseURI` (also the methods `setPersistsLocation` and `getPersistLocation` are available for legacy Java code that expects them, but the persist location set that way is not actually used).

In addition, this language resource will still automatically add the resource name of a resource as the String value for the annotation property “label”.

14.4 The Ontology_OWLIM2 plugin: backwards-compatible implementation

14.4.1 The OWLIMOntologyLR Language Resource

This implementation is identical to the implementation that was part of GATE core before version 5.1. It is based on SwiftOWLIM version 2 and Sesame version 1.

In order to load an ontology in an OWLIM repository, the user has to provide certain configuration parameters. These include the name of the repository, the URL of the ontology, the default name space, the format of the ontology (RDF/XML, N3, NTriples and Turtle), the URLs or absolute locations of the other ontologies to be imported, their respective name spaces and so on. Ontology files, based on their format, are parsed and persisted in the NTriples format.

In order to utilize the power of OWLIM and the simplicity of GATE ontology API, GATE provides an implementation of the OWLIM Ontology. Its basic purpose is to hide all the complexities of OWLIM and Sesame and provide an easy to use API and interface to create, load, save and update ontologies. Based on certain parameters that the user provides when instantiating the ontology, a configuration file is dynamically generated to create a dummy repository in memory (unless persistence is specified).

When creating a new ontology, one can use an existing file to pre-populate it with data. If no such file is provided, an empty ontology is created. A detailed description for all the parameters that are available for new ontologies follows:

1. *defaultNameSpace* is the base URI to be used for all new items that are only mentioned using their local name. This can safely be left empty, in which case, while adding new resources to the ontology, users are asked to provide name spaces for each new resource.
2. As indicated earlier, OWLIM supports four different formats: RDF/XML, NTriples, Turtle and N3. According to the format of the ontology file, user should select one of the four URL options (*rdfXmlURL*, *ntriplesURL*, *turtleURL* and *n3URL* (*not supported yet*)) and provide a URL pointing to the ontology data.

Once an ontology is created, additional data can be loaded that will be merged with the existing information. This can be done by right-clicking on the ontology in the resources tree in GATE Developer and selecting 'Load ... data' where '...' is one of the supported formats.

Other options available are cleaning the ontology (deleting all the information from it) and saving it to a file in one of the supported formats.

Ontology can be saved in different formats (rdf/xml, ntriples, n3 and turtle) using the options provided in the context menu that can be invoked by right clicking on the instance of an ontology in GATE Developer. All the changes made to the ontology are logged and stored as an ontology feature. Users can also export these changes to a file by selecting the 'Save Ontology Event Log' option from the context menu. Similarly, users can also load the exported event log and apply the changes on a different ontology by using the 'Load Ontology Event Log' option. Any change made to the ontology can be described by a set of triples either added or deleted from the repository. For example, in GATE Embedded, addition of a new instance results into addition of two statements into the repository:

```
// Adding a new instance "Rec1" of type "Recognized"  
// Here + indicates the addition  
+ <http://proton.semanticweb.org/2005/04/protons#Rec1>  
<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>  
<http://proton.semanticweb.org/2005/04/protons#Recognized>
```

```
// Adding a label (annotation property) to the instance with value "Rec Instance"
+ <http://proton.semanticweb.org/2005/04/protons#Rec1>
<http://www.w3.org/2000/01/rdf-schema#label>
<Rec Instance>
<http://www.w3.org/2001/XMLSchema#string>
```

The event log therefore contains a list of such triples, the latest change being at the bottom of the change log. Each triple consists of a subject followed by a predicate followed by an object. Below we give an illustration explaining the syntax used for recording the changes.

```
// Adding a new instance "Rec1" of type "Recognized"
// Here + indicates the addition
+ <http://proton.semanticweb.org/2005/04/protons#Rec1>
<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>
<http://proton.semanticweb.org/2005/04/protons#Recognized>

// Adding a label (annotation property) to the instance with value "Rec Instance"
+ <http://proton.semanticweb.org/2005/04/protons#Rec1>
<http://www.w3.org/2000/01/rdf-schema#label>
<Rec Instance>
<http://www.w3.org/2001/XMLSchema#string>

// Adding a new class called TrustSubClass
+ <http://proton.semanticweb.org/2005/04/protons#TrustSubClass>
<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>
<http://www.w3.org/2002/07/owl#Class>

// TrustSubClass is a subclassOf the class Trusted
+ <http://proton.semanticweb.org/2005/04/protons#TrustSubClass>
<http://www.w3.org/2000/01/rdf-schema#subclassOf>
<http://proton.semanticweb.org/2005/04/protons#Trusted>

// Deleting a property called hasAlias and all relevant statements
// Here - indicates the deletion
// * indicates any value in place
- <http://proton.semanticweb.org/2005/04/protons#hasAlias> <*> <*>
- <*> <http://proton.semanticweb.org/2005/04/protons#hasAlias> <*>
- <*> <*> <http://proton.semanticweb.org/2005/04/protons#hasAlias>

// Deleting a label set on the instance Rec1
- <http://proton.semanticweb.org/2005/04/protons#Rec1>
```



```
<http://www.w3.org/2000/01/rdf-schema#label>  
<Rec Instance>  
<http://www.w3.org/2001/XMLSchema#string>  
  
// Reseting the entire ontology (Deleting all statements)  
- <*> <*> <*>
```

14.5 GATE Ontology Editor

GATE's ontology support also includes a viewer/editor that can be used within GATE Developer to navigate an ontology and quickly inspect the information relating to any of the objects defined in it—classes and restrictions, instances and their properties. Also, resources can be deleted and new resources can be added through the viewer.

Before the ontology editor can be used, one of the ontology implementation plugins must be loaded. In addition the `Ontology_Tools` must be loaded.

Note: To make it possible to show a loaded ontology in the ontology editor, the `Ontology_Tools` plugin must be loaded *before* the ontology language resource is created.

The viewer is divided into two areas. One on the left shows separate tabs for hierarchy of classes and instances and for (as of Gate 4) hierarchy of properties. The view on right hand side shows the details pertaining of the object currently selected in the other two.

First tab on the left view displays a tree which shows all the classes and restrictions defined in the ontology. The tree can have several root nodes—one for each top class in the ontology. The same tree also shows each instances for each class. **Note:** Instances that belong to several classes are shown as children of all the classes they belong to.

Second tab on the left view displays a tree of all the properties defined in the ontology. This tree can also have several root nodes—one for each top property in the ontology. Different types of properties are distinguished by using different icons.

Whenever an item is selected in the tree view, the right-hand view is populated with the details that are appropriate for the selected object. For an ontology class, the details include the brief information about the resource such as the URI of the selected class, type of the selected class etc., set of direct superclasses, the set of all superclasses using the transitive closure, the set of direct subclasses, the set of all the subclasses, the set of equivalent classes, the set of applicable property types, the set of property values set on the selected class, and the set of instances that belong to the selected class. For a restriction, in addition to the above information, it displays on which property the restriction is applicable to and what type of the restriction that is.

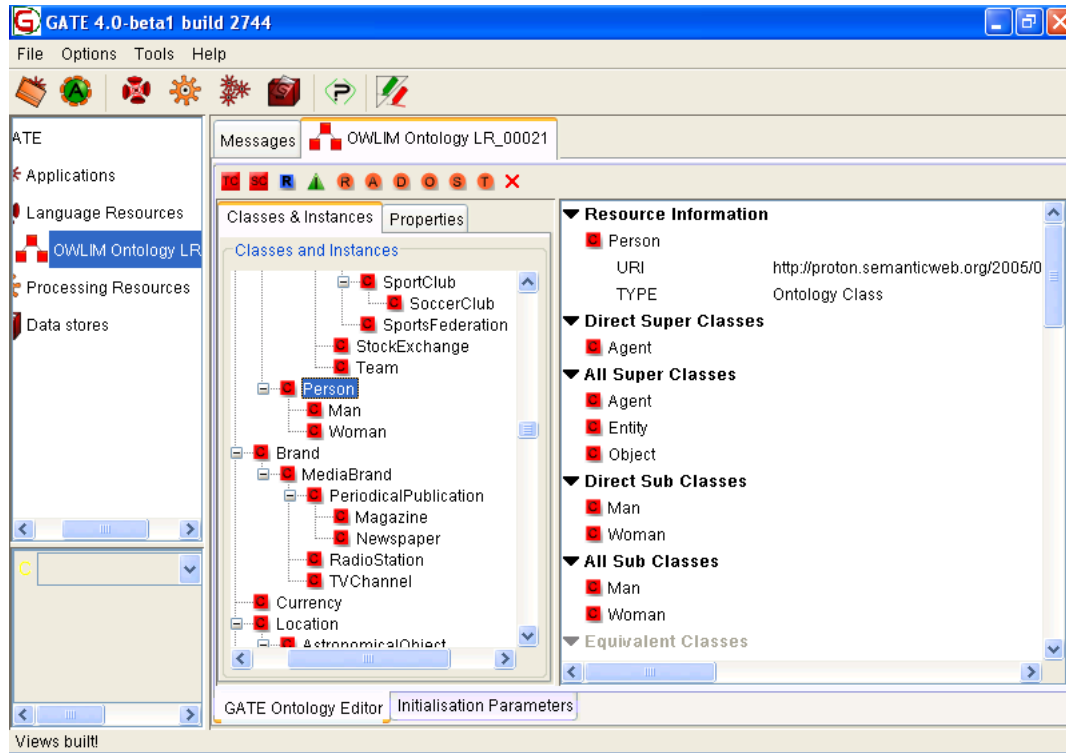


Figure 14.6: The GATE Ontology Viewer

For an instance, the details displayed include the brief information about the instance, set of direct types (the list of classes this instance is known to belong to), the set of all types this instance belongs to (through the transitive closure of the set of direct types), the set of same instances, the set of different instances and the values for all the properties that are set.

When a property is selected, different information is displayed in the right-hand view according to the property type. It includes the brief information about the property itself, set of direct superproperties, the set of all superproperties (obtained through the transitive closure), the set of direct subproperties, the set of all subproperties (obtained through the transitive closure), the set of equivalent properties, and domain and range information.

As mentioned in the description of the data model, properties are not directly linked to the classes, but rather define their domain of applicability through a set of domain restrictions. This means that the list of properties should not really be listed as a detail for class objects but only for instances. It is however quite useful to have an indication of the types of properties that could apply to instances of a given class. Because of the semantics of property domains, it is not possible to calculate precisely the list of applicable properties for a given class, but only an estimate of it. If a property for instance requires its domain instances to belong to two different classes then it cannot be known with certitude whether it is applicable to either of the two classes—it does not apply to all instances of any of those classes, but

only to those instances the two classes have in common. Because of this, such properties will not be listed as applicable to any class.

The information listed in the details pane is organised in sub-lists according to the type of the items. Each sub-list can be collapsed or expanded by clicking on the little triangular button next to the title. The ontology viewer is dynamic and will update the information displayed whenever the underlying ontology is changed through the API.

When you double click on any resource in the details table, the respective resource is selected in the class or in the property tree and the selected resource's details are shown in the details table. To change a property value, user can double click on a value of the property (second column) and the relevant window is shown where user is asked to provide a new value. Along with each property value, a button (with red X caption) is provided. If user wants to remove a property value he or she can click on the button and the property value is deleted.

A new toolbar has been added at the top of the ontology viewer, which contains the following buttons to add and delete ontology resources:

- Add new top class (TC)
- Add new subclass (SC)
- Add new instance (I)
- Add new restriction (R)
- Add new Annotation property (A)
- Add new Datatype property (D)
- Add new Object property (O)
- Add new Symmetric property (S)
- Add new Transitive property (T)
- Remove the selected resource(s) (X)
- Search
- Refresh ontology

The tree components allow the user to select more than one node, but the details table on the right-hand side of the GATE Developer GUI only shows the details of the first selected node. The buttons in the toolbar are enabled and disabled based on users' selection of nodes in the tree.

1. **Creating a new top class:**

A window appears which asks the user to provide details for its namespace (default name space if specified), and class name. If there is already a class with same name in ontology, GATE Developer shows an appropriate message.

2. **Creating a new subclass:**

A class can have multiple super classes. Therefore, selecting multiple classes in the ontology tree and then clicking on the 'SC' button, automatically considers the selected classes as the super classes. The user is then asked for details for its namespace and class name.

3. **Creating a new instance:**

An instance can belong to more than one class. Therefore, selecting multiple classes in the ontology tree and then clicking on the 'I' button, automatically considers the selected classes as the type of new instance. The user is then prompted to provide details such as namespace and instance name.

4. **Creating a new restriction:**

As described above, restriction is a type of an anonymous class and is specified on a property with a constraint set on either the number of values it can take or the type of value allowed for instances to have for that property. User can click on the blue 'R' square button which shows a window for creating a new restriction. User can select a type of restriction, property and a value constraint for the same. Please note that restrictions are considered as anonymous classes and therefore user does not have to specify any URI for the same but restrictions are named automatically by the system.

5. **Creating a new property:**

Editor allows creating five different types of properties:

- **Annotation property:** Since an annotation property cannot have any domain or range constraints, clicking on the new annotation property button brings up a dialog that asks the user for information such as the namespace and the annotation property name.
- **Datatype property:** A datatype property can have one or more ontology classes as its domain and one of the pre-defined datatypes as its range. Selecting one or more classes and clicking on the new Datatype property icon, brings up a window where the selected classes in the tree are taken as the property's domain. The user is then asked to provide information such as the namespace and the property name. A drop down box allows users to select one of the data types from the list.
- **Object, Symmetric and Transitive properties:** These properties can have one or more classes as their domain and range. For a symmetric property the domain and range are the same. Clicking on any of these options brings up a window where user is asked to provide information such as the namespace and the property

name. The user is also given two buttons to select one or more classes as values for domain and range.

6. Removing the selected resources:

All the selected nodes are removed when user clicks on the ‘X’ button. Please note that since ontology resources are related in various ways, deleting a resource can affect other resources in the ontology; for example, deleting a resource can cause other resources in the same ontology to be deleted too.

7. Searching in ontology:

The Search button allows users to search for resources in the ontology. A window pops up with an input text field that allows incremental searching. In other words, as user types in name of the resource, the drop-down list refreshes itself to contain only the resources that start with the typed string. Selecting one of the resources in this list and pressing OK, selects the appropriate resource in the editor. The Search function also allows selecting resources by the property values set on them.

8. Refresh Ontology

The refresh button reloads the ontology and updates the editor.

9. Setting properties on instances/classes:

Right-clicking on an instance brings up a menu that provides a list of properties that are inherited and applicable to its classes. Selecting a specific property from the menu allows the user to provide a value for that property. For example, if the property is an Object property, a new window appears which allows the user to select one or more instances which are compatible to the range of the selected property. The selected instances are then set as property values. For classes, all the properties (e.g. annotation and RDF properties) are listed on the menu.

10. Setting relations among resources:

Two or more classes, or two or more properties, can be set as equivalent; similarly two or more instances can be marked as the same. Right-clicking on a resource brings up a menu with an appropriate option (Equivalent Class for ontology classes, Same As Instance for instances and Equivalent Property for properties) which when clicked then brings up a window with a drop down box containing a list of resources that the user can select to specify them as equivalent or the same.

14.6 Ontology Annotation Tool

The Ontology Annotation Tool (OAT) is a GATE plugin available from the Ontology Tools plugin set, which enables a user to manually annotate a text with respect to one or more




Figure 14.7: Viewing Ontology-Based Annotations

ontologies. The required ontology must be selected from a pull-down list of available ontologies.

The OAT tool supports annotation with information about the ontology classes, instances and properties.

14.6.1 Viewing Annotated Text

Ontology-based annotations in the text can be viewed by selecting the desired classes or instances in the ontology tree in GATE Developer (see Figure 14.7). By default, when a class is selected, all of its sub-classes and instances are also automatically selected and their mentions are highlighted in the text. There is an option to disable this default behaviour (see Section 14.6.4).

Figure 14.7 shows the mentions of each class and instance in a different colour. These colours can be customised by the user by clicking on the class/instance names in the ontology tree. It is also possible to expand and collapse branches of the ontology.




Figure 14.8: Editing Existing Annotations

14.6.2 Editing Existing Annotations

In order to view the class/instance of a highlighted annotation in the text (e.g., United States - see Figure 14.8), hover the mouse over it and an edit dialogue will appear. It shows the current class or instance (Country in our example) and allows the user to delete it or change it. To delete an existing annotation, press the Delete button.

A class or instance can be changed by starting to type the name of the new class in the combo-box. Then it displays a list of available classes and instances, which start with the typed string. For example, if we want to change the type from Country to Location, we can type 'Lo' and all classes and instances which names start with Lo will be displayed. The more characters are typed, the fewer matching classes remain in the list. As soon as one sees the desired class in the list, it is chosen by clicking on it.

It is possible to apply the changes to all occurrences of the same string and the same previous class/instance, not just to the current one. This is useful when annotating long texts. The user needs to make sure that they still check the classes and instances of annotations further down in the text, in case the same string has a different meaning (e.g., bank as a building vs. bank as a river bank).

The edit dialogue also allows correcting annotation offset boundaries. In other words, user can expand or shrink the annotation offsets' boundaries by clicking on the relevant arrow




Figure 14.9: Add New Annotation

buttons.

OAT also allows users to assign property values as annotation features to the existing class and instance annotations. In the case of class annotation, all annotation properties from the ontology are displayed in the table. In the case of instance annotations, all properties from the ontology applicable to the selected instance are shown in the table. The table also shows existing features of the selected annotation. User can then add, delete or edit any value(s) of the selected feature. In the case of a property, user is allowed to provide an arbitrary number of values. User can, by clicking on the editList button, add, remove or edit any value to the property. In case of object properties, users are only allowed to select values from a pre-selected list of values (i.e. instances which satisfy the selected property's range constraints).

14.6.3 Adding New Annotations

New annotations can be added in two ways: using a dialogue (see Figure 14.9) or by selecting the text and clicking on the desired class or instance in the ontology tree.

When adding a new annotation using the dialogue, select a text and after a very short while, if the mouse is not moved, a dialogue will appear (see Figure 14.9). Start typing the name of the desired class or instance, until you see it listed in the combo-box, then select it with




Figure 14.10: Tool Options

the mouse. This operation is the same, as in changing the class/instance of an existing annotation. One has the option of applying this choice to the current selection only or to all mentions of the selected string in the current document (Apply to All check box).

User can also create an instance from the selected text. If user checks the 'create instance' checkbox prior to selecting the class, the selected text is annotated with the selected class and a new instance of the selected class (with the name equivalent to the selected text) is created (provided there isn't any existing instance available in the ontology with that name).

14.6.4 Options

There are several options that control the OAT behaviour (see Figure 14.10):

- **Disable child feature:** By default, when a class is selected, all of its sub-classes are also automatically selected and their mentions are highlighted in the text. This option disables that behaviour, so only mentions of the selected class are highlighted.
- **Delete confirmation:** By default, OAT deletes ontological information without asking for confirmation, when the delete button is pressed. However, if this leads to too many mistakes, it is possible to enable delete confirmations from this option.

- **Disable Case-Sensitive Feature:** When user decides to annotate all occurrences of the selected text ('apply to all' option) in the document and if the 'disable case-sensitive feature' is selected, the tool, when searching for the identical strings in the document text, ignores the case-sensitivity.
- **Setting up a filter to disable resources from the OAT GUI:** When user wants to annotate the text of a document with certain classes/instances of the ontology, s/he may disable the resources which s/he is not going to use. This option allows users to select a file which contains class or instance names, one per line. These names are case sensitive. After selecting a file, when user turns on the 'filter' check box, the resources specified in the filter file are disabled and removed from the annotation editor window. User can also add new resources to this list or remove some or all from the list by right clicking on the respective resource and by selecting the relevant option. Once modified, the 'save' button allows users to export this list to a file.
- **Annotation Set:** GATE stores information in annotation sets and OAT allows you to select which set to use as input and output.
- **Annotation Type:** By default, this is annotation of type Mention, but that can be changed to any other name. This option is required because OAT uses Gate annotations to store and read the ontological data. However, to do that, it needs a type (i.e. name) so ontology-based annotations can be distinguished easily from other annotations (e.g. tokens, gazetteer lookups).

14.7 Using the ontology API

The following code demonstrates how to use the GATE API to create an instance of the OWLIM Ontology language resource. This example shows how to use the current version of the API and ontology implementation.

For an example of using the old API and the backwards compatibility plugin, see 14.8.

```

1 // step 1: initialize GATE
2 if(!Gate.isInitialized()) { Gate.init(); }
3
4 // step 2: load the Ontology plugin that contains the implementation
5 File ontoHome = new File(Gate.getPluginsHome(), "Ontology");
6 Gate.getCreoleRegister().addDirectory(ontoHome.toURL());
7
8 // step 3: set the parameters
9 FeatureMap fm = Factory.newFeatureMap();
10 fm.put("rdfXmlURL", urlOfTheOntology);
11 fm.put("baseURI", theBaseURI);
12 fm.put("mappingsURL", urlOfTheMappingsFile);
13 // .. any other parameters
14
15 // step 4: finally create an instance of ontology

```

```

16  Ontology ontology = (Ontology)
17  Factory.createResource("gate.creole.ontology.impl.sesame.OWLIMOntology",
18                          fm);
19
20  // retrieving a list of top classes
21  Set<OClass> topClasses = ontology.getOClasses(true);
22
23  // for all top classes, printing their direct sub classes and print
24  // their URI or blank node ID in turtle format.
25  for(OClass c : topClasses) {
26      Set<OClass> dcs = c.getSubClasses(OConstants.DIRECT_CLOSURE);
27      for(OClass sClass : dcs) {
28          System.out.println(sClass.getONodeID().toTurtle());
29      }
30  }
31
32  // creating a new class from a full URI
33  OURI aURI1 = ontology.createOURI("http://sample.en/owlim#Organization");
34  OClass organizationClass = ontology.addOClass(aURI1);
35
36  // create a new class from a name and the default name space set for
37  // the ontology
38  OURI aURI2 = ontology.createOURIForName("someOtherName");
39  OClass someOtherClass = ontology.addOClass(aURI2);
40
41  // set the label for the class
42  someOtherClass.setLabel("some other name", OConstants.ENGLISH);
43
44  // creating a new Datatype property called name
45  // with domain set to Organization
46  // with datatype set to string
47  URI dURI = new URI("http://sample.en/owlim#Name", false);
48  Set<OClass> domain = new HashSet<OClass>();
49  domain.add(organizationClass);
50  DatatypeProperty dp =
51      ontology.addDatatypeProperty(dURI, domain, Datatype.getStringDataType());
52
53  // creating a new instance of class organization called IBM
54  OURI iURI = ontology.createOURI("http://sample.en/owlim#IBM");
55  OInstance ibm = Ontology.addOInstance(iURI, organizationClass);
56
57  // assigning a Datatype property, name to ibm
58  ibm.addDatatypePropertyValue(dp, new Literal("IBM Corporation",
59
60
61  // get all the set values of all Datatype properties on the instance ibm
62  Set<DatatypeProperty> dps = Ontology.getDatatypeProperties();
63  for(DatatypeProperty dp : dps) {
64      List<Literal> values = ibm.getDatatypePropertyValues(dp);
65      System.out.println("DP : "+dp.getOURI());
66      for (Literal l : values) {
67          System.out.println("Value : "+l.getValue());
68          System.out.println("Datatype : "+ l.getDataType().getXmlSchemaURI());

```

```

69  }
70  }
71
72  // export data to a file in Turtle format
73  BufferedWriter writer = new BufferedWriter(new FileWriter(someFile));
74  ontology.writeOntologyData(writer, OConstants.OntologyFormat.TURTLE);
75  writer.close();

```

14.8 Using the ontology API (old version)

The following code demonstrates how to use the GATE API to create an instance of the OWLIM Ontology language resource. *This example shows how to use the API with the backwards-compatibility plugin `Ontology_OWLIM2`*

For how to use the API with the current implementation plugin, see 14.7.

```

1  // step 1: initialize GATE
2  Gate.init();
3
4  // step 2: load the Ontology_Tools plugin
5  File ontoHome = new File(Gate.getPluginsHome(), "Ontology_Tools");
6  Gate.getCreoleRegister().addDirectory(ontoHome.toURL());
7
8  // step 3: set the parameters
9  FeatureMap fm = Factory.newFeatureMap();
10 fm.put("rdfXmlURL", url-of-the-ontology);
11
12 // step 4: finally create an instance of ontology
13 Ontology ontology = (Ontology)
14 Factory.createResource("gate.creole.ontology.owlim.OWLIMOntologyLR", fm);
15
16 // retrieving a list of top classes
17 Set<OClass> topClasses = ontology.getOClasses(true);
18
19 // for all top classes, printing their direct sub classes
20 Iterator<OClass> iter = topClasses.iterator();
21 while(iter.hasNext()) {
22     Set<OClass> dcs = iter.next().getSubClasses(OConstants.DIRECT_CLOSURE);
23     for(OClass aClass : dcs) {
24         System.out.println(aClass.getURI().toString());
25     }
26 }
27
28 // creating a new class
29 // false indicates that it is not an anonymous URI
30 URI aURI = new URI("http://sample.en/owlim#Organization", false);
31 OClass organizationClass = ontology.addOClass(aURI);
32
33 // creating a new Datatype property called name
34 // with domain set to Organization

```

```

35 // with datatype set to string
36 URI dURI = new URI("http://sample.en/owlim#Name", false);
37 Set<OClass> domain = new HashSet<OClass>();
38 domain.add(organizationClass);
39 DatatypeProperty dp = ontology.addDatatypeProperty(dURI, domain,
40                                                    Datatype.getStringDatatype());
41
42 // creating a new instance of class organization called IBM
43 URI iURI = new URI("http://sample.en/owlim#IBM", false);
44 OInstance ibm = Ontology.addOInstance(iURI, organizationClass);
45
46 // assigning a Datatype property, name to ibm
47 ibm.addDatatypePropertyValue(dp, new Literal("IBM Corporation",
48                                             dp.getDataType()));
49
50 // get all the set values of all Datatype properties on the instance ibm
51 Set<DatatypeProperty> dps = Ontology.getDatatypeProperties();
52 for(DatatypeProperty dp : dps) {
53     List<Literal> values = ibm.getDatatypePropertyValues(dp);
54     System.out.println("DP : "+dp.getURI().toString());
55     for (Literal l : values) {
56         System.out.println("Value : "+l.getValue());
57         System.out.println("Datatype : "
58                             + l.getDataType().getXmlSchemaURI().toString()); }
59 }
60
61 // export data to a file in the ntriples format
62 BufferedWriter writer = new BufferedWriter(new FileWriter(someFile));
63 String output = ontology.getOntologyData(
64     OConstants.ONTOLLOGY_FORMAT_NTRIPLES);
65 writer.write(output);
66 writer.flush();
67 writer.close();

```

14.9 Ontology-Aware JAPE Transducer

One of the GATE components that makes use of the ontology support is the JAPE transducer (see Chapter 8). Combining the power of ontologies with JAPE's pattern matching mechanisms can ease the creation of applications.

In order to use ontologies with JAPE, one needs to load an ontology in GATE before loading the JAPE transducer. Once the ontology is known to the system, it can be set as the value for the optional ontology parameter for the JAPE grammar. Doing so alters slightly the way the matching occurs when the grammar is executed. If a transducer is ontology-aware (i.e. it has a value set for the 'ontology' parameter) it will treat all occurrences of the feature named `class` differently from the other features of annotations. The values for the feature `class` on any type of annotation will be considered to be the names of classes belonging the ontology and the matching between two values will not be based on simple equality but rather

hierarchical compatibility. For example if the ontology contains a class named ‘Politician’, which is a sub class of the class ‘Person’, then a pattern of `{Entity.class == ‘Person’}` will successfully match an annotation of type `Entity` with a feature `class` having the value ‘Politician’. If the JAPE transducer were not ontology-aware, such a test would fail.

This behaviour allows a larger degree of generalisation when designing a set of rules. Rules that apply several types of entities mentioned in the text can be written using the most generic class they apply to and need not be repeated for each subtype of entity. One could have rules applying to `Locations` without needing to know whether a particular location happens to be a country or a city.

If a domain ontology is available at the time of building an application, using it in conjunction with the JAPE transducers can significantly simplify the set of grammars that need to be written.

The ontology does not normally affect actions on the right hand side of JAPE rules, but when Java is used on the right hand side, then the ontology becomes accessible via a local variable named `ontology`, which may be referenced from within the right-hand-side code.

In Java code, the `class` feature should be referenced using the static final variable, `LOOKUP_CLASS_FEATURE_NAME`, that is defined in `gate.creole.ANNIEConstants`.

14.10 Annotating Text with Ontological Information

The ontology-aware JAPE transducer enables the text to be linked to classes in an ontology by means of annotations. Essentially this means that each annotation can have a class and ontology feature. To add the relevant class feature to an annotation is very easy: simply add a feature ‘class’ with the classname as its value. To add the relevant ontology, use `ontology.getURL()`.

Below is a sample rule which looks for a location annotation and identifies it as a ‘Mention’ annotation with the class ‘Location’ and the ontology loaded with the ontology-aware JAPE transducer (via the runtime parameter of the transducer).

Rule: Location

```
({Location}):mention

-->
:mention{
  // create the ontology and class features
  FeatureMap features = Factory.newFeatureMap();
  features.put("ontology", ontology.getURL());
```

```

features.put("class", "Location");

// create the new annotation
try {
    annotations.add(mentionAnnots.firstNode().getOffset(),
        mentionAnnots.lastNode().getOffset(), "Mention", features);
}
catch(InvalidOffsetException e) {
    throw new JapeException(e);
}
}

```

14.11 Populating Ontologies

Another typical application that combines the use of ontologies with NLP techniques is finding mentions of entities in text. The scenario is that one has an existing ontology and wants to use Information Extraction to populate it with instances whenever entities belonging to classes in the ontology are mentioned in the input texts.

Let us assume we have an ontology and an IE application that marks the input text with annotations of type ‘Mention’ having a feature ‘class’ specifying the class of the entity mentioned. The task we are seeking to solve is to add instances in the ontology for every `Mention` annotation.

The example presented here is based on a JAPE rule that uses Java code on the action side in order to access directly the GATE ontology API:

```

Rule: FindEntities
({Mention}):mention
-->
:mention{
    //find the annotation matched by LHS
    //we know the annotation set returned
    //will always contain a single annotation
    Annotation mentionAnn = mentionAnnots.iterator().next();

    //find the class of the mention
    String className = (String)mentionAnn.getFeatures().
        get(gate.creole.ANNIEConstants.LOOKUP_CLASS_FEATURE_NAME);
    // should normalize class name and avoid invalid class names here!
    OClass aClass = ontology.getOClass(ontology.createOURIForName(className));
    if(aClass == null) {

```

```

    System.err.println("Error class \"" + className + "\" does not exist!");
    return;
}

//find the text covered by the annotation
String theMentionText = gate.Utills.stringFor(doc, mentionAnn);

// when creating a URI from text that came from a document you must take care
// to ensure that the name does not contain any characters that are illegal
// in a URI. The following method does this nicely for English but you may
// want to do your own normalization instead if you have non-English text.
String mentionName = OUtills.toResourceName(theMentionText);

// get the property to store mention texts for mention instances
DatatypeProperty prop =
    ontology.getDatatypeProperty(ontology.createOURIForName("mentionText"));

OURI mentionURI = ontology.createOURIForName(mentionName);
// if that mention instance does not already exist, add it
if (!ontology.containsOInstance(mentionURI)) {
    OInstance inst = ontology.addOInstance(mentionURI, aClass);
    // add the actual mention text to the instance
    try {
        inst.addDatatypePropertyValue(prop,
            new Literal(theMentionText, OConstants.ENGLISH));
    }
    catch(InvalidValueException e) {
        throw new JapeException(e);
    }
}
}
}

```

This will match each annotation of type `Mention` in the input and assign it to a label `'mention'`. That label is then used in the right hand side to find the annotation that was matched by the pattern (lines 5–10); the value for the `class` feature of the annotation is used to identify the ontological class name (lines 12–14); and the annotation span is used to extract the text covered in the document (lines 16–26). Once all these pieces of information are available, the addition to the ontology can be done. First the right class in the ontology is identified using the class name (lines 28–37) and then a new instance for that class is created (lines 38–50).

Beside JAPE, another tool that could play a part in this application is the Ontological Gazetteer, see Section 13.4, which can be useful in bootstrapping the IE application that finds entity mentions.

The solution presented here is purely pedagogical as it does not address many issues that would be encountered in a real life application solving the same problem. For instance, it is naïve to assume that the name for the entity would be exactly the text found in the document. In many cases entities have several aliases – for example the same person name can be written in a variety of forms depending on whether titles, first names, or initials are used. A process of name normalisation would probably need to be employed in order to make sure that the same entity, regardless of the textual form it is mentioned in, will always be linked to the same ontology instance.

For a detailed description of the GATE ontology API, please consult the JavaDoc documentation.

14.12 Ontology API and Implementation Changes

This section describes the changes in the API and the implementation made in GATE Developer version 5.1. The most important change is that *the implementation of the ontology API has been removed from the GATE core and is now being made available as plugins*. Currently the plugin `Ontology_OWLIM2` provides the implementation that was present in the GATE core previously and the plugin `Ontology` provides a new and upgraded implementation that also implements some new features that were added to the API. The `Ontology_OWLIM2` plugin is intended to provide maximum backwards compatibility but will not be developed further and be phased out in the future, while the `Ontology` plugin provides the current actively developed implementation.

Before any ontology-related function can be used in GATE, one of the ontology implementation plugins must be loaded.

14.12.1 Differences between the implementation plugins

The implementation provided in plugin `Ontology_OWLIM2` is based on Sesame version 1 and OWLIM version 2, while the changed implementation provided in plugin `Ontology` is based on Sesame version 2 and OWLIM version 3.

The plugin `Ontology` provides the ontology language resource `OWLIM Ontology` with new and changed parameters. In addition, there are two language resources for advanced users, `Create Sesame Ontology` and `Connect Sesame Ontology`. Finally the new implementation provides the language resource `OWLIM Ontology DEPRECATED` to make the move from the old to the new implementation easier: this language resource has the same name, parameters and Java package as the language resource `OWLIMOntologyLR` in backwards-compatibility plugin `Ontology_OWLIM2`. This allows to test existing pipelines and applications with the new implementation without the necessity to adapt the names of the language resource or parameters.

The implementation in plugin `Ontology` makes various attempts to reduce the amount of memory needed to load an ontology. This will allow to load significantly larger ontologies into GATE. This comes at the price of some methods needing more time than before, as the implementation does not cache all ontology entities in GATE's memory any more.

The new implementation does not provide access to any implementation detail anymore, the method `getSesameRepository` will therefore throw an exception. The return type of this method in the old implementation has been changed to `Object` to remove the dependency on a Sesame class in the GATE API.

14.12.2 Changes in the Ontology API

The class `gate.creole.ontology.URI` has been deprecated. Instead, the ontology API client must use objects that implement the new `OURI`, `OBnodeID` or `ONodeID` interfaces. An API client can only directly create `OURI` objects and must use the ontology factory methods `createOURI`, `createOURIFroName` or `generateOURI` to create such objects.

Also, the intended way how ontologies are modeled has been changed: the API tries to prevent a user from adding anything to an ontology that would make an ontology that conforms to OWL-Lite go beyond that sublanguage (and e.g. become OWL-Full). However, if an ontology already is not conforming to OWL-Lite, the API tries to make as much information visible to the client as possible. That means for instance that RDF classes will be included in the list of classes returned by method `getOClasses`, but there is no support for adding RDF classes to an ontology. Similarly, all methods that already existed which would allow to add entities to an ontology that do not conform to OWL-Lite have been deprecated.

Most methods that use a constant from class `OConstants` which is defined as a byte value have been deprecated and replaced by methods that use enums that replace the byte constants instead. (however, the byte constants used for literal string languages are still used).

The API now supports the handling of ontology imports more flexibly. Ontology imports are internally kept in a named graph that is different from the named graph data from loaded ontologies is kept in. Imported ontology data is still visible to the ontology API but can be ignored when storing (serializing) an ontology. The ontology API now also allows to explicitly resolve ontology imports and it allows the specification of mappings between import URIs and URLs of either local files or substitute web URLs. The import map can also specify patterns for substituting import URIs with replacement URLs (or ignoring them altogether).

The default namespace URI is now set automatically from the ontology if possible and the API allows getting and setting the ontology URI.

The ontology API now offers methods for getting an iterator when accessing some ontology resources, e.g. when getting all classes in the ontology. This helps to prevent the excessive

use of memory when retrieving a large number of such resources from a large ontology.

Ontology objects do not internally store copies of all ontology resources in hash maps any more. This means that re-fetching ontology resources will be a slower operation and old methods that rely on this mechanism are either deprecated (`getOResourcesByName`, `getOResourceByName`) or do not work at all any more (`getOResourceFromMap`, `addOResourceToMap`, `removeOResourceFromMap`).

Chapter 15

Machine Learning

This chapter presents machine learning PRs available in GATE. Currently, two PRs are available:

- The **Batch Learning PR** (in the **Learning** plugin) is GATE's most comprehensive and developed machine learning offering. It is specifically targetted at NLP tasks including text classification, chunk learning (e.g. for named entity recognition) and relation learning. It integrates LibSVM for improved speed, along with the PAUM algorithm, offering very competitive performance and speed. It also offers a Weka interface. It is documented in Section 15.2.
- The **Machine Learning PR** (in the **Machine_Learning** plugin) is GATE's older machine learning offering. It offers wrappers for Maxent, Weka and SVM Light. It is documented in Section 15.3.

The rest of the chapter is organised as follows. Section 15.1 introduces machine learning in general, focusing on the terminology used and the meaning of the terms within GATE. We then move on to describe the two Machine Learning processing resources, beginning with the Batch Learning PR in Section 15.2. Section 15.2.1 describes all the configuration settings of the Batch Learning PR one by one; i.e. all the elements in the configuration file for setting the Batch Learning PR (the learning algorithm to be used and the options for learning) and defining the NLP features for the problem. Section 15.2.2 presents three case studies with example configuration files for the three types of NLP learning problems. Section 15.2.3 lists the steps involved in using the Batch Learning PR. Finally, Section 15.2.4 explains the outputs of the Batch Learning PR for the four usage modes; namely training, application, evaluation and producing feature files only, and in particular, the format of the feature files and label list file produced by the Batch Learning PR. Section 15.3 outlines the original Machine Learning PR in GATE.

15.1 ML Generalities

There are two main types of ML; supervised learning and unsupervised learning. Supervised learning is more effective and much more widely used in NLP. Classification is a particular example of supervised learning, in which the set of training examples is split into multiple subsets (classes) and the algorithm attempts to distribute new examples into the existing classes. This is the type of ML that is used in GATE, and all further references to ML actually refer to classification.

An ML algorithm ‘learns’ about a phenomenon by looking at a set of occurrences of that phenomenon that are used as examples. Based on these, a model is built that can be used to predict characteristics of future (unseen) examples of the phenomenon.

An ML implementation has two modes of functioning: training and application. The training phase consists of building a model (e.g. a statistical model, a decision tree, a rule set, etc.) from a dataset of already classified instances. During application, the model built during training is used to classify new instances.

Machine Learning in NLP falls broadly into three categories of task type; text classification, chunk recognition, and relation extraction

- **Text classification** classifies text into pre-defined categories. The process can be equally well applied at the document, sentence or token level. Typical examples of text classification might be document classification, opinionated sentence recognition, POS tagging of tokens and word sense disambiguation.
- **Chunk recognition** often consists of two steps. First, it identifies the chunks of interest in the text. It then assigns a label or labels to these chunks. However some problems comprise simply the first step; identifying the relevant chunks. Examples of chunk recognition include named entity recognition (and more generally, information extraction), NP chunking and Chinese word segmentation.
- **Relation extraction** determines whether or not a pair of terms in the text has some type(s) of pre-defined relations. Two examples are named entity relation extraction and co-reference resolution.

Typically, the three types of NLP learning use different linguistic features and feature representations. For example, it has been recognised that for text classification the so-called *tf – idf* representation of n-grams is very effective (e.g. with SVM). For chunk recognition, identifying the start token and the end token of the chunk by using the linguistic features of the token itself and the surrounding tokens is effective and efficient. Relation extraction benefits from both the linguistic features from each of the two terms involved in the relation and the features of the two terms combined.

The rest of this section explains some basic definitions in ML and their specification in the ML plugin.

15.1.1 Some Definitions

- **instance:** an example of the studied phenomenon. An ML algorithm learns a model from a set of known instances, called a (training) dataset. It can then apply the learned model to another (application) dataset.
- **attribute:** a characteristic of the instances. Each instance is defined by the values of its attributes. The set of possible attributes is well defined and is the same for all instances in the training and application datasets. ‘Feature’ is also often used. However, in this context, this can cause confusion with GATE annotation features.
- **class:** an attribute for which the values are available in the training dataset for learning, but which are not present in the application dataset. ML is used to find the value of this attribute in the application dataset.

15.1.2 GATE-Specific Interpretation of the Above Definitions

- **instance:** an annotation. In order to use ML in GATE, users will need to choose the type of annotations used as instances. Token annotations are a good candidate for many NLP learning tasks such as information extraction and POS tagging, but any type of annotation could be used (e.g. things that were found by a previously run JAPE grammar, such as sentence annotations and document annotations for sentence and document classification respectively).
- **attribute:** an attribute is the value of a named feature of a particular annotation type, which can either (partially) cover the instance annotation considered or another instance annotation which is related to the instance annotation considered. The value of the attribute can refer to the current instance or to an instance either situated at a specified location relative to the current instance or having special relation with the current instance.
- **class:** any attribute referring to the current instance can be marked as class attribute.

15.2 Batch Learning PR

This section describes the newest machine learning PR in GATE. The implementation focuses on the three main types of learning in NLP, namely chunk recognition (e.g. named entity recognition), text classification and relation extraction. The implementation for chunk recognition is based on our work using support vector machines (SVM) for information extraction [Li *et al.* 05a]. The text classification is based on our work on opinionated sentence classification and patent document classification (see [Li *et al.* 07c] and [Li *et al.* 07d], respectively). The relation extraction is based on our work on named entity relation extraction [Wang *et al.* 06].

The Batch Learning PR, given a set of documents, can also produce feature files, containing linguistic features and feature vectors, and labels if there are any in the documents. It can also produce document-term matrices and n-gram based language models. Feature files are in text format and can be used outside of GATE. Hence, users can use GATE-produced feature files off-line, for their own purpose, e.g. evaluating new learning algorithms.

The PR also provides facilities for active learning, based on support vector machines (SVM), mainly ranking the unlabelled documents according to the confidence scores of the current SVM models for those documents.

The primary learning algorithm implemented is SVM, which has achieved state of the art performances for many NLP learning tasks. The training of SVM uses a Java version of the SVM package LibSVM [CC001]. Application of SVM is implemented by ourselves. The PAUM (Perceptron Algorithm with Uneven Margins) is also included [Li *et al.* 02], and on our test datasets has consistently produced a performance to rival the SVM with much reduced training times. Moreover, the ML implementation provides an interface to the open-source machine learning package Weka [Witten & Frank 99], and can use machine learning algorithms implemented in Weka. Three widely-used learning algorithms are available in the current implementation: Naive Bayes, KNN and the C4.5 decision tree algorithm.

Access to ML implementations is provided in GATE by the ‘Batch Learning PR’ (in the ‘learning’ plugin). The PR handles training and application of an ML model, evaluation of learning on GATE documents, producing feature files and ranking documents for Active Learning. It also makes it possible to view the primal forms of a linear SVM. This PR is a Language Analyser so it can be used in all default types of GATE controllers.

In order to use the Batch Learning processing resource, the user has to do three things. First, the user has to annotate some training documents with the labels that s/he wants the learning system to annotate in new documents. Those label annotations should be GATE annotations. Secondly, the user may need to pre-process the documents to obtain linguistic features for the learning. Again, these features should be in the form of GATE annotations. GATE’s plugin ANNIE might be helpful for producing the linguistic features. Other resources such as the NP Chunker and parser may also be helpful. By providing the machine learning algorithm with more and better information on which to base learning, chances of a good result are increased, so this preprocessing stage is important. Finally the user has to create a configuration file for setting the ML PR, e.g. selecting the learning algorithm and defining the linguistic features used in learning. Three example configuration files are presented in this section; it might be helpful to take one of them as a starting point and modify it.

15.2.1 Batch Learning PR Configuration File Settings

In order to allow for more flexibility, all configuration parameters for the PR are set through one external XML file, except for the learning mode, which is selected through normal PR

parameterisation. The XML file contains both the configuration parameters of the Batch Learning PR itself and of the linguistic data (namely the definitions of the instance and attributes) used by the Batch Learning PR. The XML file is specified when creating a new Batch Learning PR.

The parent directory of the XML configuration file becomes the working directory. A subdirectory in the working directory, named ‘savedFiles’, will be created (if it does not already exist). All the files produced by the Batch Learning PR, including the NLP features files, label list file, feature vector file and learned model file, will be stored in that subdirectory. A log file recording the learning session is also created in this directory.

Below, we first describe the parameters of the Batch Learning PR. Then we explain those settings specified in the configuration file.

PR Parameters: Settings not Specified in the Configuration File

For the sake of convenience, a few settings are not specified in the configuration file. Instead the user should specify them as initialization or run-time parameters of the PR, as in other PRs.

- **URL (or path and name) of the configuration file.** The user is required to give the URL of the configuration file when creating the PR. The configuration file should be in XML format with the extension name *.xml*. It contains most of learning settings and will be explained in detail in the next subsection.
- **Corpus.** This is a run-time parameter, meaning that the user should specify it after creating the PR, and may change it between runs. The corpus contains the documents that the PR will use as learning data (training or application). For application, the documents should include all the annotations specified in the configuration file, except the class attribute. The annotations for class attribute should be available in the documents used for training or evaluation.
- **inputASName** is the annotation set containing the annotations for the linguistic features to be used and the class labels.
- **outputASName** is the annotation set in which the results of applying the models will be put. Note that it should be set the same as the *inputASName* when doing the evaluation (i.e. setting the *learningMode* as ‘EVALUATION’).
- **learningMode** is a run-time parameter. It can be set as one of the following values, ‘TRAINING’, ‘APPLICATION’, ‘EVALUATION’, ‘ProduceFeatureFilesOnly’, ‘MITRAINING’, ‘VIEWPRIMALFORMMODELS’ and ‘RankingDocsForAL’. The default learning mode is ‘TRAINING’.

- In **TRAINING** mode, the PR learns from the data provided and saves the models into a file called ‘learnedModels.save’ under the sub-directory ‘savedFiles’ of the working directory.
- If the user wants to apply the learned model to the data, s/he should select **APPLICATION** mode. In application mode, the PR reads the learned model from the file ‘learnedModels.save’ in the subdirectory ‘savedFiles’ and then applies the model to the data.
- In **EVALUATION** mode, the PR will do k-fold or hold-out test set evaluation on the corpus provided (the method of the evaluation is specified in the configuration file, see below), and output the evaluation results to the messages window of GATE Developer, or standard out when using GATE Embedded, and into the log file. When using evaluation mode, please make sure that the *outputASName* is set to the same annotation set as the *inputASName*.
- If the user only wants to produce feature data and feature vectors but does not want to train or apply a model, s/he may select the **ProduceFeatureFilesOnly** mode. The feature files that the PR produces will be explained in detail in Section 15.2.4.
- In **MITRAINING** (mixed initiative training) mode, the training data are appended to the end of any existing feature file. In contrast, in training mode, the training data created in the current session overwrite any existing feature file. Consequently, mixed initiative training mode uses both the training data obtained in this session and the data that existed in the feature file before starting the session. Hence, training mode is for batch learning, while mixed initiative training mode can be used for on-line (or adaptive, or mixed-initiative) learning. There is one parameter for mixed initiative training mode specifying the minimal number of newly added documents before starting the learning procedure to update the learned model. The parameter can be defined in the configuration file.
- **VIEWPRIMALFORMMODELS** mode is used for displaying the most salient NLP features in the learned models. In the current implementation, the mode is only valid with the linear SVM model, in which the most salient NLP features correspond to the biggest (absolute values of) weights in the weight vector. In the configuration file one can specify two parameters to determine the number of displayed NLP features for positive and negative weights. Note that if e.g. the number for negative weight is set as 0, then no NLP feature is displayed for negative weights.
- **RankingDocsForAL** applies the current learned SVM models (in the sub-directory ‘savedFiles’) to the feature vectors stored in the file ‘fvsDataSelecting.save’ in the sub-directory ‘savedFiles’ and ranks the documents according to the margins of the examples in one document to the SVM models. The ranked list of documents will be put into the file ‘ALRankedDocs.save’.

In most cases it is not safe to run more than one instance of the batch learning PR with the same working directory at the same time, because the PR needs to update the model (in

TRAINING, MITRAINING or EVALUATION mode) or other data files. It *is* safe to run multiple instances at once provided they are all in *APPLICATION* mode¹.

Order of document processing In the usual case, in a GATE corpus pipeline application, documents are processed one at a time, and each PR is applied in turn to the document, processing it fully, before moving on to the next document. The Batch Learning PR breaks from this rule. ML training algorithms, including SVM, typically run as a batch process over a training set, and require all the data to be fully prepared and passed to the algorithm in one go. This means that in training (or evaluation) mode, the Batch Learning PR will wait for all the documents to be processed and will then run as a single operation at the end. Therefore, the Batch Learning PR needs to be positioned *last* in the pipeline. Post-processing cannot be done within the pipeline after the Batch Learning PR. Where further processing needs to be done, this should take the form of a separate application, and be applied to the data afterwards.

There is an exception to the above, however. In application mode, the situation is slightly different, since the ML model has already been created, and the PR only applies it to the data. This can be done on a document by document basis, in the manner of a normal PR. However, although it can be done document by document, there may be advantages in terms of efficiency to grouping documents into batches before applying the algorithm. A parameter in the configuration file, *BATCH-APP-INTERVAL*, described later, allows the user to specify the size of such batches, and by default this is set to 1; in other words, by default, the Batch Learning PR in application mode behaves like a normal PR and processes each document separately. There may be substantial efficiency gains to be had through increasing this parameter (although higher values require more memory consumption), but if the Batch Learning PR is applied in application mode *and* the parameter *BATCH-APP-INTERVAL* is set to 1, the PR can be treated like any other, and other PRs may be positioned after it in a pipeline.

Settings in the Batch Learning PR XML Configuration File

The root element of the XML configuration file needs to be called ‘ML-CONFIG’, and it must contain two basic elements; *DATASET* and *ENGINE*, and optionally other settings. In the following, we first describe the optional settings, then the *ENGINE* element, and finally the *DATASET* element. In the next section, some examples of the XML configuration file are given for illustration. Please also refer to the configuration files in the test directory (i.e. `plugins/learning/test/` under the main gate directory) for more examples.

¹This is only true for GATE 5.2 or later; in earlier versions *all* modes were unsafe for multiple instances of the PR.

Optional Settings in the Configuration File The Batch Learning PR provides a variety of optional settings, which facilitate different tasks. Every optional setting has a default value; if an optional setting is not specified in the configuration file, the Batch Learning PR will adopt its default value. Each of the following optional settings can be set as an element in the XML configuration file.

- **SURROUND** should be set to ‘true’ if the user wants the Batch Learning PR to learn chunks by identifying the start token and the end token of the chunk. This approach to chunk learning, for example, named entity recognition, where a span of several tokens is to be identified, often produces better results than trying to learn every token in the chunk. For classification problems and relation extraction, set its value as ‘false’. This element appears in the configuration file as:

```
<SURROUND VALUE='X'/>
```

where the variable *X* has two possible values: ‘true’ or ‘false’. The default value is ‘false’.

- **FILTERING** relates to SVM training. Where the ratio of positive examples to negative examples is low, i.e. the instances belonging in the class are much outweighed by instances outside of the class (e.g. ‘one against others’ is used, see *multiClassification2Binary* below) SVMs can run into difficulties. The positive examples may be swamped by outlying negative examples. The ML plugin provides functionality developed through research (e.g. [Li & Bontcheva 08]) to assist in such cases. One example is the *FILTERING* parameter. The filtering functionality performs initial SVM training, then removes negative examples on the basis of their position relative to the separator. It then retrains on the smaller dataset. Typically, negative instances close to the boundary are removed. Note that this two-step process takes longer than simple training. However, the second training step will be quicker than the first, as it is performed on a somewhat reduced dataset. If the item *dis* is set as ‘near’, the PR selects and removes those negative examples which are closest to the SVM hyper-plane. If it is set as ‘far’, those negative examples that are furthest from the SVM hyper-plane are removed. The value of the item *ratio* determines what proportion of negative examples will be filtered out. This element appears in the configuration file as:

```
< FILTERING ratio='X' dis='Y'/>
```

where *X* represents a number between 0 and 1 and *Y* can be set as ‘near’ or ‘far’. If the filtering element is not present in the configuration file, or the value of *ratio* is set as 0.0, the PR does not perform filtering. The default value of *ratio* is 0.0. The default value of *dis* is ‘far’.

- **EVALUATION** As outlined above, if the learning mode parameter *learningMode* is set to ‘EVALUATION’, the PR will perform evaluation of the ML model; it will split the documents in the corpus into two parts, the training dataset and the test dataset, learn a model from the training dataset, apply the model to the testing dataset, and finally compare the annotations assigned by the model on the test set with the true annotations and output measures of success (e.g. F-measure). The evaluation element specifies the method of splitting the corpus. The item *method* determines which method

to use for evaluation. Currently two commonly used methods are implemented, namely *k-fold cross-validation* and *hold-out test*. In *k-fold cross-validation* the PR segments the corpus into *k* partitions of equal size, and uses each of the partitions in turn as a test set, with all the remaining documents as a training set. For *hold-out test*, the system randomly selects some documents as testing data and uses all other documents as training data. The value of the item *runs* specifies the number ‘*k*’ for *k-fold cross-validation*. The value of the item *ratio* specifies the ratio of the data used for training in the *hold-out test* method. The element in the configuration file appears as so:

```
<EVALUATION method="X" runs="Y" ratio="Z"/>
```

where the variable *X* has two possible values ‘*kfold*’ and ‘*holdout*’, *Y* is a positive integer, and *Z* is a float number between 0 and 1. The default value of *method* is ‘*holdout*’. The default value of *runs* is ‘1’. The default value of *ratio* is ‘0.66’.

- **multiClassification2Binary**. Certain machine learning algorithms, including SVM, are designed to operate on two class problems; they find a separator between two groups of instances. In order to use such algorithms to classify items into a larger number of classes, the problem has to be converted into a series of ‘binary’ (two class) problems. The ML plugin implements two common methods for converting a multi-class problem into several binary problems, namely *one against others* and *one against another*. The two methods may have slightly different names in other publications, but the principle is the same. Suppose we have a multi-class classification problem with *n* classes. For the *one against others* method, one binary classification problem is derived for each of the *n* classes. Examples belonging to the class in question are considered to be positive examples and all other examples in the training set are negative examples. In contrast, for the *one against another* method, one binary classification problem is derived for each pair (*c1*, *c2*) of the *n* classes. Training examples belonging to the class *c1* are the positive examples and those belonging to the other class, *c2*, are the negative examples. The user can select one of the two methods by specifying the value of the item *method* of the element. The element appears as so:

```
<multiClassification2Binary method="X" thread-pool-size="N"/>
```

where the variable *X* has two values, ‘*one-vs-others*’ and ‘*one-vs-another*’. Note that depending on the sample size, the two methods may differ greatly in their speed of execution. The default method is the *one-vs-others* method. If the configuration file does not have the element or the item *method* is missed, then the PR will use the *one-vs-others* method. Since the derived binary classifiers are independent it is possible to learn several of them in parallel. The ‘*thread-pool-size*’ attribute gives the number of threads that will be used to learn and apply the binary classifiers. If omitted, a single thread will be used to process all the classifiers in sequence.

- **thresholdProbabilityBoundary** sets a confidence threshold on start and end tokens for chunk learning. It is used in post-processing the learning results. Only those boundary tokens in which the confidence level is above the threshold are selected as candidates for the entities. The element in configuration file appears as so:

```
<PARAMETER name="thresholdProbabilityBoundary" value="X"/>
```

The value *X* is between 0 and 1. The default value is 0.4.

- **thresholdProbabilityEntity** sets a confidence threshold on chunks (which is the multiplication of the probabilities of the start token and end token of the chunk) for chunk learning. Only those entities in which the confidence level is above the threshold are selected as candidates of the entities. The element in configuration file appears as so:

```
<PARAMETER name="thresholdProbabilityEntity" value="X"/>
```

The value X is between 0 and 1. The default value is 0.2.
- The threshold parameter **thresholdProbabilityClassification** is the confidence threshold for classification (e.g. text classification and relation extraction tasks. In contrast, the above two probabilities are for the chunking recognition task.) The corresponding element in configuration file appears as so:

```
<PARAMETER name="thresholdProbabilityClassification" value="X"/>
```

The value X is between 0 and 1. The default value is 0.5.
- **IS-LABEL-UPDATABLE** is a Boolean parameter. If its value is set to ‘true’, the label list is updated from the labels in the training data. Otherwise, a pre-defined label list will be used and cannot be updated from the training data. The configuration element appears as so:

```
<IS-LABEL-UPDATABLE value="X"/>
```

The value X is ‘true’ or ‘false’. The default value is ‘true’.
- **IS-NLPFEATURELIST-UPDATABLE** is a Boolean parameter. If its value is set to ‘true’, the NLP feature list is updated from the features in the training or application data. Otherwise, a pre-defined NLP feature list will be used and cannot be updated. The configuration element appears as so:

```
<IS-NLPFEATURELIST-UPDATABLE value="X"/>
```

The value X is ‘true’ or ‘false’. The default value is ‘true’.
- The parameter **VERBOSITY** specifies the verbosity level of the output of the system, both to the message window of GATE Developer (or standard out when using GATE Embedded) and into the log file. Currently there are three verbosity levels. Level 0 only allows the output of warning messages. Level 1 outputs some important setting information and the results for evaluation mode. Level 2 is used for debugging purposes. The element in the configuration file appears as so:

```
<VERBOSITY level="X"/>
```

The value X can be set as 0, 1 or 2. The default value is 1.
- **MI-TRAINING-INTERVAL** specifies the minimal number of newly added documents needed to trigger retraining the model. This parameter is used in *MITRAINING*. The number is specified by the value of the feature ‘num’ as so:

```
<MI-TRAINING-INTERVAL num="X"/>
```

The default value of X is 1.
- **BATCH-APP-INTERVAL** is used in application mode, and specifies the number of documents to be collected and passed as a batch for classification. Please refer to

Section 15.2.1 for a detailed explanation of this option. The corresponding element in the configuration file is:

```
<BATCH-APP-INTERVAL num="X"/>
```

The default value of X is 1.

- **DISPLAY-NLPFEATURES-LINEARSVM** relates to ‘VIEWPRIMALFORM-MODELS’ mode. In this mode, the most significant features are displayed for each class. For more information about this mode see Section 15.2.1. Two numbers are specified; the number of positively weighted features to display and the number of negatively weighted features to display. It has the following form in the configuration file;

```
<DISPLAY-NLPFEATURES-LINEARSVM numP="X" numN="Y"/>
```

where X and Y represent the numbers of positively and negatively weighted features to display, respectively. The default values of X and Y are 10 and 0.

- **ACTIVELEARNING** specifies the settings for active learning. Active learning ranks documents based on the average of a sample of ML annotation confidence scores. A larger sample gives a more accurate ranking but takes longer to calculate. The option has the following form:

```
<ACTIVELEARNING numExamplesPerDoc='X'/>
```

where X represents the number of examples per document used to obtain the confidence score with respect to the learned model. The default value of *numExamplesPerDoc* is 3.

The ENGINE Element The *ENGINE* element specifies which ML algorithm will be used, and also allows the options to be set for that algorithm.

For SVM learning, the user can choose one of two learning engines. We will discuss the two SVM learning engines below. Note that only linear and polynomial kernels are supported. This is despite the fact that the original SVM packages implemented other types of kernel. Linear and polynomial kernels are popular in natural language learning, and other types of kernel are rarely used. However, if you want to experiment with other types of kernel, you can do so by first running the Batch Learning PR in GATE to produce the training and testing data, then using the data with the SVM implementation outside of GATE.

The configuration files in the test directory (i.e. `plugins/learning/test/` under the main gate directory) contain examples for setting the learning engine.

The *ENGINE* element in the configuration file is specified as follows:

```
<ENGINE nickname='X' implementationName='Y' options='Z'/>
```

It has three items:

- **nickname** can be the name of the learning algorithm or whatever the user wants it to

be.

- **implementationName** refers to the implementation of the particular learning algorithm that the user wants to use. Its value should be one of the following:
 - **SVMLibSvmJava**, the binary classification SVM algorithm implemented in the Java version of the SVM package LibSVM.
 - **SVMExec**, a binary SVM implementation of your choice, potentially in a language other than Java, run as a separate process outside of GATE. Currently it can use the *SVM^{light}* SVM package²; see the XML file in the GATE distribution (at `gate/plugins/learning/test/chunklearning/engines-svm-svmlight.xml`) for an example of how to specify the learning engine to be used. The learning engines *SVMExec* and *SVMLibSvmJava* should produce the same results in theory but may get slightly different results in practice due to implementational differences. *SVMLibSvmJava* tends to be faster than *SVMExec* for smaller training sets. There may be cases where it is an advantage to run SVM as a separate process however, in which case, *SVMExec* would be preferable.
 - **PAUM**, the Perceptron with uneven margins, a simple and fast classification learning algorithm. (For details about the learning algorithm PAUM, see [Li *et al.* 02]).
 - **PAUMExec**, a binary PAUM implementation of your choice, potentially in a language other than Java, run as a separate process outside of GATE. The relationship between the *PAUM* and *PAUMExec* is similar to that of *SVMLibSvmJava* and *SVMExec*. You may download and use an implementation in C from <http://www.dcs.shef.ac.uk/~yaoyong/paum/paum-learning.zip>. See the XML file in the GATE distribution (at `gate/plugins/learning/test/chunklearning/engines-paum-exec.xml`) for an example of how to specify the learning engine to be used.
 - **NaiveBayesWeka**, the Naive Bayes learning algorithm implemented in Weka.
 - **KNNWeka**, the K nearest neighbour (KNN) algorithm implemented in Weka.
 - **C4.5Weka**, the decision tree algorithm C4.5 implemented in Weka.
- **Options**: the value of this item, which is dependent on the particular learning algorithm, will be passed verbatim to the ML engine used. Where an option is absent, defaults for that engine will be used.
 - The options for *SVMLibSvmJava* are similar to those for LibSVM but with the exception that since *SVMLibSvmJava* implements the uneven margins SVM algorithms described in [Li & Shawe-Taylor 03], it takes the uneven margins parameter as an option. *SVMLibSvmJava* options are as follows:
 - * **-s svm_type**; whether the SVM should be binary or multiclass. Default value is 0. Since only binary is supported, the option should be set to 0 or excluded.

²The SVM package *SVM^{light}* can be downloaded from <http://svmlight.joachims.org/>.

- * **-t kernel_type**; 0 for a linear kernel or 1 for a polynomial kernel. Default value is 0. Note that the current implementation does not support other kernel types such as radial and sigmoid function.
 - * **-d degree**; the degree in polynomial kernel, e.g. 2 for quadratic kernel. Default value is 3.
 - * **-c cost**; the cost parameter C in the SVM. Default value is 1. This parameter determines the cost associated with allowing training errors ('soft margins'). Allowing some points to be misclassified by the SVM may produce a more generalizable result.
 - * **-m cachesize**; the cache memory size in MB (default 100).
 - * **-tau value**; setting the value of uneven margins parameter of the SVM. $\tau = 1$ corresponds to the standard SVM. If the training data has just a small number of positive examples and a large number of negative examples, setting the parameter τ to a value less than 1 (e.g. $\tau = 0.4$) often results in better F-measure than the standard SVM (see [Li & Shawe-Taylor 03]).
- The options for *SVMExec*, using *SVM^{light}*, are similar to those for using *SVM^{light}* directly for training. Options set the type of kernel, the parameters in the kernel function, the cost parameter, the memory used, etc. The parameter tau is also included, to set the uneven margins parameter, as explained above. The last two terms in the parameter options are the training data file and the model file. An example of the options for SVMExec might be '-c 0.7 -t 0 -m 100 -v 0 -tau 0.6 /yaoyong/software/svm-light/data_svm.dat /yaoyong/software/svm-light/model_svm.dat', meaning that the learner uses a linear kernel, the uneven margins parameter is set as 0.6, and two data files /yaoyong/software/svm-light/data_svm.dat and /yaoyong/software/svm-light/model_svm.dat for writing and reading data. Note that both the data files specified here are temporary files, which are used only by the svm-light training program, can be in anywhere in your computer, and are independent of the data files produced by the GATE learning plugin. SVMExec also takes a further argument, **executableTraining**, which specifies the SVM learning program svm_learn.exe in the *SVM^{light}*. For example, executableTraining='/yaoyong/software/svm-light/svm_learn.exe' specifies one particular svm_learn.exe obtained from the package *SVM^{light}*.
 - The *PAUM* engine has three options; '-p' for the positive margin, '-n' for the negative margin, and '-optB' for the modification of the bias term. For example, options='-p 50 -n 5 -optB 0.3' means $\tau_+ = 50$, $\tau_- = 5$ and $b = b + 0.3$ in the PAUM algorithm.
 - The KNN algorithm has one option; the number of neighbours used. It is set via '-k X'. The default value is 1.
 - There are no options for Naive Bayes and C4.5 algorithms.

The DATASET Element The *DATASET* element defines the type of annotation to be used as training instance and the set of attributes that characterise the instances. The

INSTANCE-TYPE sub-element is used to select the annotation type to be used for instances. There will be one training instance for every one of the instance annotations in the corpus. For example, if *INSTANCE-TYPE* has ‘Token’ as its value, there will be one training instance in the document per token. This also means that the *positions* (see below) are defined in relation to tokens. *INSTANCE-TYPE* can be seen as the basic unit to be taken into account for machine learning. The attributes of the instance are defined by a sequence of *ATTRIBUTE*, *ATTRIBUTE_REL* or *ATTRIBUTELIST* elements.

Different NLP learning tasks may have different instance types and use different kinds of attribute elements. Chunking recognition often uses the token as instance type and the linguistic features of ‘Token’ and other annotations as features. Text classification’s instance type is the text unit for classification, e.g. the whole document, or sentence, or token. If classifying for example a sentence, n-grams (see below) are often a good feature representation for many statistical learning algorithms. For relation extraction, the instance type is a pair of terms that may be related, and the features come from not only the linguistic features of each of the two terms but also those related to both terms taken together.

The *DATASET* element should define an *INSTANCE-TYPE* sub-element, it should define an *ATTRIBUTE* sub-element or an *ATTRIBUTE_REL* sub-element as class, and it should define some linguistic feature related sub-elements (‘linguistic feature’ or ‘NLP feature’ is used here to distinguish features or attributes used for machine learning from features in the sense of a feature of a GATE annotation). All the annotation types involved in the dataset definition should be in the same annotation set. Each of the sub-elements defining the linguistic features (attributes) should contain an element defining the annotation *TYPE* to be used and an element defining the *FEATURE* of the annotation type to use. For instance, *TYPE* might be ‘Person’ and *FEATURE* might be ‘gender’. For an *ATTRIBUTE* sub-element, if you do not specify *FEATURE*, the entire sub-element will be ignored. Therefore, if an annotation type you want to use does not have any annotation features, you should add an annotation feature to it and assign the same value to the feature for all annotations of that type. Note that if blank spaces are contained in the values of the annotation features, they will be replaced by the character ‘_’ in each occurrence. So it is advisable that the values of the annotation features used, in particular for the class label, do not contain any blank space.

Below, we explain all the sub-elements one by one. Please also refer to the example configuration files presented in next section. Note that each sub-element should have a unique name, if it requires a name, unless we explicitly state otherwise.

- The **INSTANCE-TYPE** sub-element is defined as `<INSTANCE-TYPE>X</INSTANCE-TYPE>` where X is the annotation type used as instance unit for learning, for example ‘Token’. For relation extraction, the user should also specify the two arguments of the relation, as so:

`<INSTANCE-ARG1>A</INSTANCE-ARG1>`

`<INSTANCE-ARG2>B</INSTANCE-ARG2>`

The values of A and B should be identifiers for the first and second terms of the relation,

respectively. These names will be used later in the configuration file. An example can be found at `/gate/plugins/learning/test/relation-learning/engines-svm.xml`.

- An **ATTRIBUTE** element has the following sub-elements:
 - **NAME**; the name of the attribute. Its value should not end with ‘gram’, since this is reserved for n-gram features as mentioned below. This attribute name will appear in output files, so it is useful to give a descriptive name.
 - **SEMTYPE**; type of the attribute value. It can be ‘NOMINAL’ or ‘NUMERIC’. Currently only nominal is supported.
 - **TYPE**; the annotation type used to extract the attribute.
 - **FEATURE**; the value of the attribute will be the value of the named feature on the annotation of the specified type.
 - **POSITION**; the position of the instance annotation to be used for extracting the feature relative to the current instance annotation. 0 refers to the current instance annotation, -1 refers to the preceding instance annotation, 1 refers to the following one and so forth. Recall that we defined *INSTANCE-TYPE* at the start of the *DATASET* element. This type might for example be ‘Token’. In the current *ATTRIBUTE* element we are defining an annotation type to use to get the feature from, separate and possibly different from the *INSTANCE-TYPE*. For example, we might be interested in the ‘majorType’ of a ‘Lookup’. By specifying -1, we would be saying, move to the preceding ‘Token’ and then try to extract the ‘majorType’ of the ‘Lookup’ on that token. The default value of the parameter is 0. Note that if our *INSTANCE-TYPE* were to be for example a named entity annotation comprising multiple tokens, and we wanted to extract a feature on the ‘Token’ annotation, then all the tokens within it would be considered to be in the zero position relative to the current instance annotation, and the current implementation would simply pick the first. (Useful in this case might be the *NGRAM* attribute type, described later, which can be used to extract features for each member of a multi-token annotation.) In the current implementation, features are weighted according to their distance from the current instance annotation. In other words, features which are further removed from the current instance annotation are given reduced importance. The component value in the feature vector for one attribute feature is 1 if the attribute’s position p is 0. Otherwise its value is $1.0/|p|$.
 - **<CLASS/>**: an empty element used to mark the class attribute. There can only be one attribute marked as class in a dataset definition. The attribute, as described above, has specified *TYPE* and *FEATURE*; the features of the type are the class labels. Since only one attribute can be marked as class, it may be necessary to preprocess your data to put all class labels into a feature of one type of annotation, e.g. you might create a ‘Mention’ annotation, with the feature ‘Class’, which is set to the class name.

- The **ATTRIBUTE** element is similar to *ATTRIBUTE* except that it has no *POSITION* sub-element but instead a *RANGE* element. This will be converted into several attributes with position ranging from the value of ‘from’ to the value of ‘to’. It defines a ‘context window’ containing several consecutive examples. The *ATTRIBUTE* should be preferred when defining a context window for features, because not only it can avoid the duplication of *ATTRIBUTE* elements, but also because processing is speeded up (see the discussion for the element *WINDOWSIZE* below).
- The **WINDOWSIZE** element specifies the size of the context window. This will override the context window size defined in every *ATTRIBUTE*. If the *WINDOWSIZE* element is not present in the configuration file, the window size defined in each element *ATTRIBUTE* will be used; otherwise, the window size specified by this element will be used for each *ATTRIBUTE* if it contains one *ATTRIBUTE* at position 0 (otherwise the *ATTRIBUTE* will be ignored). This element can be used for speeding up the process of extracting the feature vectors from the documents. The element has two features specifying the length of left and right sides of context window. It has the following form:


```
<WINDOWSIZE windowSizeLeft="X" windowSizeRight="Y"/>
```

 where X and Y represent the the length of left and right sides of context window, respectively. For example, if X = 2 and Y = 1, then the context window will be from the position -2 to 1 (e.g. from the second token in the left through the current token to the first token in the right).
- An **GRAM** feature is used for characterising an instance annotation in terms of constituent sequences of subsumed feature annotations. It is essentially a reversal of the *ATTRIBUTE* principle; where *ATTRIBUTE* uses a sequence *surrounding* an instance in order to classify the instance, *GRAM* uses sequences *within* the instance as features. It simply creates a series of attributes that constitute a sliding window across the entire of the current instance annotation. For example, *INSTANCE-TYPE* might be sentences, in sentence classification, and the *GRAM* attribute specification could be used for example to create a series of unigram features for the sentence, effectively a ‘bag of words’ representation. Conventionally, one would use the string of the token, or perhaps its lemma, as the feature for the *GRAM*; however, it is possible to specify multiple features of choice, as shown below.
 - **NAME**; name of the n-gram. Its value should end with ‘gram’.
 - **NUMBER**; the ‘n’ of the n-gram, with value 1 for unigram, and 2 for bigram, etc.
 - **CONSNUM**; several features can be used to generate n-grams. For example, n-grams of token strings could be used as well as n-grams of lemmas. Where *CONSNUM* is ‘k’, the *GRAM* element should have ‘k’ **CONS-X** sub-elements, where X= 1, ..., k. Each *CONS-X* element has one *TYPE* sub-element and one *FEATURE* sub-element, which define feature to be used for that term to create n-grams.

- The **WEIGHT** sub-element specifies a weight for the n-gram feature. The n-gram part of the feature vector for one instance is normalised, thus having a default value of 1.0. If the user wants to adjust the contributions of the n-gram to the whole feature vector, s/he can do so by setting the **WEIGHT** parameter. For example, if the user is doing sentence classification and s/he uses two features; the unigram of tokens in a sentence and the length of the sentence, by default the entire of the *NGRAM* attribute specification is given only the same importance as the sentence length feature. In order to experiment with increasing the importance of the n-gram element, the user can set the weight sub-element of the n-gram element with a number bigger than 1.0 (like 10.0). Then every component of the n-gram part of the feature vector would be multiplied by the parameter.
- The **ValueTypeNgram** element specifies the type of value used in the n-gram. Currently it can take one of the three types; ‘binary’, ‘tf’, and ‘tf-idf’, which are explained in Section 15.2.4. The value is specified by the **X** in


```
<ValueTypeNgram>X</ValueTypeNgram>
```

X = 1 for binary, = 2 for tf, and = 3 for tf-idf. The default value is 3.
- The **FEATURES-ARG1** element defines the features related to the first argument of the relation for relation learning. It should include one *ARG* sub-element referring to the GATE annotation of the argument (see below for a detailed explanation). It may include other sub-elements, such as *ATTRIBUTE*, *ATTRIBUTELIST* and/or *NGRAM*, to define the linguistic features related to the argument. Features pertaining particularly to one or the other argument of a relation should be defined in *FEATURES-ARG1* or *FEATURES-ARG2* as appropriate. Features relating to both arguments should be defined using an *ATTRIBUTE_REL*.
- The **FEATURES-ARG2** element defines the features related to the second argument of relation. Like the element *FEATURES-ARG1*, it should include one *ARG* sub-element. It may also include other sub-elements. The *ARG* sub-element in the *FEATURES-ARG2* should have a unique name which is different from the name for the *ARG* sub-element in the *FEATURES-ARG1*. However, other sub-elements may have the same name as corresponding ones in the *FEATURES-ARG1*, if they refer to the same annotation type and feature in the text.
- The **ARG** element is used in both *FEATURES-ARG1* and *FEATURES-ARG2*. It specifies the annotation corresponding to one argument of a relation. It has four sub-elements, as follows;
 - **NAME**; a unique name for the argument (e.g. ‘ARG1’).
 - **SEMTYPE**; the type of the arg value. This can be ‘NOMINAL’ or ‘NUMERIC’. Currently only nominal is implemented.
 - **TYPE**; the annotation type for the argument.
 - **FEATURE**; the value of the named feature on the annotation of specified type is the identifier of the argument. Only if the value of the feature is same as the value

of the feature specified in the sub-element $\langle \text{INSTANCE-ARG1} \rangle \mathbf{A} \langle / \text{INSTANCE-ARG1} \rangle$ (or $\langle \text{INSTANCE-ARG2} \rangle \mathbf{B} \langle / \text{INSTANCE-ARG2} \rangle$), the argument is regarded as one argument of the relation instance considered.

- **ATTRIBUTE_REL** element is similar to the *ATTRIBUTE* element. However, it does not have the *POSITION* sub-element, and it has two other sub-elements, *ARG1* and *ARG2*, relating to the two argument features of the (relation) instance type. In other words, if and only if the value \mathbf{X} in the sub-element $\langle \text{ARG1} \rangle \mathbf{X} \langle / \text{ARG1} \rangle$ is same as the value \mathbf{A} in the first argument instance $\langle \text{INSTANCE-ARG1} \rangle \mathbf{A} \langle / \text{INSTANCE-ARG1} \rangle$ and the value \mathbf{Y} in the sub-element $\langle \text{ARG2} \rangle \mathbf{Y} \langle / \text{ARG2} \rangle$ is same as the value \mathbf{B} in the second argument instance $\langle \text{INSTANCE-ARG2} \rangle \mathbf{B} \langle / \text{INSTANCE-ARG2} \rangle$ is the feature defined in this *ATTRIBUTE_REL* sub-element assigned to the instance considered. For relation learning, an *ATTRIBUTE_REL* is denoted as the class attribute by including $\langle \text{CLASS} \rangle$.

15.2.2 Case Studies for the Three Learning Types

The following are three illustrated examples of configuration files for information extraction, sentence classification and relation extraction. Note that the configuration file is in the XML format, and should be stored in a file with the ‘.xml’ extension.

Information Extraction

The first example is for information extraction. The corpus is prepared with annotations providing class information as well as the features to be used. Class information is provided in the form of a single annotation type, ‘Mention’, which contains a feature ‘class’. Within the class feature is the name of the class of the textual chunk. Other annotations in the dataset include ‘Token’ and ‘Lookup’ annotations as provided by ANNIE. All of these annotations are in the same annotation set, the name of which will be passed as a runtime parameter.

The configuration file is given below. The optional settings are in the first part. It first specifies surround mode as ‘true’; we will find the chunks that correspond to our entities by using machine learning to locate the start and end of the chunks. Then it specifies the filtering settings. Since we are going to use SVM in this problem, we can filter our data to remove some of the negative instances that can cause problems if they are too dominant. The *ratio*’s value is ‘0.1’ and the *dis*’s value is ‘near’, meaning that an initial SVM learning step will be executed and the 10% of negative examples which are closest to the learned SVM hyper-plane will be removed in the filtering stage, before the final learning is executed. The threshold probabilities for the boundary tokens and information entity are set as ‘0.4’ and ‘0.2’, respectively; boundary tokens found with a lower confidence than the threshold will be rejected. The threshold probability for classification is also set as ‘0.5’; this, however, will not be used in this case since we are doing chunk learning with *surround* mode set as

‘true’. The parameter will be ignored. *multiClassification2Binary* is set as ‘one-vs-others’, meaning that the ML API will convert the multi-class classification problem into a series of binary classification problems using the *one against others* approach. In evaluation mode, ‘2-fold’ cross-validation will be used, dividing the corpus into two equal parts and running two training/test cycles with each part as the training data.

The second part is the sub-element *ENGINE*, specifying the learning algorithm. The PR will use the *LibSVM* SVM implementation. The options determine that it will use the linear kernel with the cost *C* as 0.7 and the cache memory as 100M. Additionally it will use uneven margins, with τ as 0.4.

The last part is the *DATASET* sub-element, defining the linguistic features used. It first specifies the ‘Token’ annotation as instance type. The first *ATTRIBUTE* allows the token’s string as a feature of an instance. The range from ‘-5’ to ‘5’ means that the strings of the current token instance as well as its five preceding tokens and its five ensuing tokens will be used as features for the current token instance. The next two attribute lists define features based on the tokens’ capitalisation information and types. The *ATTRIBUTE* named ‘Gaz’ uses as attributes the values of the feature ‘majorType’ of the annotation type ‘Lookup’. The final *ATTRIBUTE* feature defines the class attribute; it has the sub-element *CLASS*. The values of the feature ‘class’ of the annotation type ‘Mention’ are the class labels.

```
<?xml version="1.0"?>
<ML-CONFIG>
  <SURROUND value="true"/>
  <FILTERING ratio="0.1" dis="near"/>
  <PARAMETER name="thresholdProbabilityEntity" value="0.2"/>
  <PARAMETER name="thresholdProbabilityBoundary" value="0.4"/>
  <PARAMETER name="thresholdProbabilityClassification" value="0.5"/>
  <multiClassification2Binary method="one-vs-others"/>
  <EVALUATION method="kfold" runs="2"/>
  <ENGINE nickname="SVM" implementationName="SVMLibSvmJava"
    options=" -c 0.7 -t 0 -m 100 -tau 0.4  "/>
<DATASET>
  <INSTANCE-TYPE>Token</INSTANCE-TYPE>
  <ATTRIBUTE<LIST>
    <NAME>Form</NAME>
    <SEMTYPE>NOMINAL</SEMTYPE>
    <TYPE>Token</TYPE>
    <FEATURE>string</FEATURE>
    <RANGE from="-5" to="5"/>
  </ATTRIBUTE<LIST>
  <ATTRIBUTE<LIST>
    <NAME>Orthography</NAME>
```

```

    <SEMTYPE>NOMINAL</SEMTYPE>
    <TYPE>Token</TYPE>
    <FEATURE>orth</FEATURE>
    <RANGE from="-5" to="5"/>
</ATTRIBUTELIST>
<ATTRIBUTELIST>
  <NAME>Tokenkind</NAME>
  <SEMTYPE>NOMINAL</SEMTYPE>
  <TYPE>Token</TYPE>
  <FEATURE>kind</FEATURE>
  <RANGE from="-5" to="5"/>
</ATTRIBUTELIST>
<ATTRIBUTELIST>
  <NAME>Gaz</NAME>
  <SEMTYPE>NOMINAL</SEMTYPE>
  <TYPE>Lookup</TYPE>
  <FEATURE>majorType</FEATURE>
  <RANGE from="-5" to="5"/>
</ATTRIBUTELIST>
<ATTRIBUTE>
  <NAME>Class</NAME>
  <SEMTYPE>NOMINAL</SEMTYPE>
  <TYPE>Mention</TYPE>
  <FEATURE>class</FEATURE>
  <POSITION>0</POSITION>
  <CLASS/>
</ATTRIBUTE>
</DATASET>
</ML-CONFIG>

```

Sentence Classification

We will now consider the case of sentence classification. The corpus in this example is annotated with ‘Sentence’ annotations, which contain the feature ‘sent_size’, as well as the class of the sentence. Furthermore, ‘Token’ annotations are applied, having features ‘category’ and ‘root’. As before, all annotations are in the same set, and the annotation set name will be passed to the PR at run time.

Below is an example configuration file. It first specifies surround mode as ‘false’, because it is a text classification problem; we are interested in classifying single instances rather than chunks of instances. Our targets of interest, sentences, have already been found (unlike in the information extraction example, where identifying the limits of the entity was part of the

problem). The next two options allow the label list and the NLP feature list to be updated from the training data when retraining. It also specifies probability thresholds for entity and entity boundary. Note that these two specifications will not be used in this case. However, their presence is not problematic; they will simply be ignored. The probability threshold for classification is set as '0.5'. This will be used to decide which classifications to accept and which to reject as being too unlikely. (Altering this parameter can trade off precision against recall and vice versa.) The evaluation will use the hold-out test method. It will randomly select 66% of the documents from the corpus for training, and the other 34% documents will be used for testing. It will run the evaluation twice, and average the results over the two runs. Note that it does not specify the method of converting a multi-class classification problem into several binary class problem, meaning that it will adopt the default (namely one against all others).

The configuration file specifies KNN (K-Nearest Neighbour) as the learning algorithm. It also specifies the number of neighbours used as 5. Of course other learning algorithms can be used as well. For example, the *ENGINE* element in the previous example, which specifies SVM as learning algorithm, can be put into this configuration file to replace the current one.

In the *DATASET* element, the annotation 'Sentence' is used as instance type. Two kinds of linguistic features are defined; one is *NGRAM* and the other is *ATTRIBUTE*. The n-gram is based on the annotation 'Token'. It is a unigram, as its *NUMBER* element has the value 1. This means that a 'bag of words' feature will be formed from the tokens comprising the sentence. It is based on the two features, 'root' and 'category', of the annotation 'Token'. This introduces a new aspect to the n-gram. The n-gram feature comprises counts of the unigrams appearing in the sentence. For example, if the sentence were 'the man walked the dog', the unigram feature would contain the information that 'the' appeared twice, and 'man', 'walked' and 'dog' appeared once. However, since our n-gram has two features, 'root' and 'category', two tokens will be considered the same term if and only if they have the same 'root' feature and the same 'category' feature. The weight of the ngram is set as 10.0, meaning its contribution is ten times that of the contribution of the other feature, the sentence length. The feature 'sent_size' of the annotation 'Sentence' is given as an *ATTRIBUTE* feature. Finally the values of the feature 'class' of the annotation 'Sentence' are nominated as the class labels.

```
<?xml version="1.0"?>
<ML-CONFIG>
  <SURROUND value="false"/>
  <IS-LABEL-UPDATABLE value="true"/>
  <IS-NLPFEATURELIST-UPDATABLE value="true"/>
  <PARAMETER name="thresholdProbabilityEntity" value="0.2"/>
  <PARAMETER name="thresholdProbabilityBoundary" value="0.42"/>
  <PARAMETER name="thresholdProbabilityClassification" value="0.5"/>
  <EVALUATION method="holdout" runs="2" ratio="0.66"/>
  <ENGINE nickname="KNN" implementationName="KNNWeka" options = " -k 5 "/>
```

```

<DATASET>
  <INSTANCE-TYPE>Sentence</INSTANCE-TYPE>
  <NGRAM>
    <NAME>Sent1gram</NAME>
    <NUMBER>1</NUMBER>
    <CONSNUM>2</CONSNUM>
    <CONS-1>
      <TYPE>Token</TYPE>
      <FEATURE>root</FEATURE>
    </CONS-1>
    <CONS-2>
      <TYPE>Token</TYPE>
      <FEATURE>category</FEATURE>
    </CONS-2>
    <WEIGHT>10.0</WEIGHT>
  </NGRAM>
  <ATTRIBUTE>
    <NAME>Class</NAME>
    <SEMTYPE>NOMINAL</SEMTYPE>
    <TYPE>Sentence</TYPE>
    <FEATURE>sent_size</FEATURE>
    <POSITION>0</POSITION>
  </ATTRIBUTE>
  <ATTRIBUTE>
    <NAME>Class</NAME>
    <SEMTYPE>NOMINAL</SEMTYPE>
    <TYPE>Sentence</TYPE>
    <FEATURE>class</FEATURE>
    <POSITION>0</POSITION>
    <CLASS/>
  </ATTRIBUTE>
</DATASET>
</ML-CONFIG>

```

Relation Extraction

The last example is for relation extraction. The relation extraction support in the PR is based on the work described in [Wang *et al.* 06].

Two concepts are key in a relation extraction corpus. Entities are the things that may be related, and relations describe the relationship between the entities if any. In our example,

entities are pre-identified, and the task is to identify the relationships between them. The corpus for this example is annotated with the following:

- ‘ACEEntity’ annotations indicate the entities of interest in the corpus.
- ‘RE_INS’ annotations form the instances, and there is an instance for every pair of ‘ACEEntities’ within a sentence. ‘RE_INS’ annotations span the entire of the text between and including their ‘ACEEntity’ annotations. For example, ‘the commander of Israeli troops’ might be a potential relationship between a person, ‘the commander’, and an entity, ‘Israeli troops’. Its ‘RE_INS’ annotation covers the entire of this text. It contains ‘arg1’ and ‘arg2’ features containing the numerical identifiers of the two ‘ACEEntities’ to which it pertains. These numerical identifiers match the ‘MENTION_ID’ feature of the ‘ACEEntity’ annotation.
- ‘ACERelation’ annotations indicate the relations we wish to learn, and also span the entire of the text involved in the relationship. They include the features ‘MENTION_ARG1’ and ‘MENTION_ARG2’, which, again, contain the numerical identifier found in the ‘MENTION_ID’ feature of the ‘ACEEntity’ annotations, as well as ‘Relation_type’, indicating the type of the relation.
- Various ANNIE-style annotations are also included.

Our task is to select the ‘RE_INS’ instances that match the ‘ACERelations’. You will see that throughout the configuration file, annotation types are specified in conjunction with argument identifiers. This is because we need to ensure that the annotation in question pertains to the right entities. Therefore, argument identifiers are used to constrain the match.

The configuration file does not specify any optional settings, meaning that it uses all the default values for those settings (see Section 15.2.1 for the default values of all possible settings).

- it sets the *surround mode* as ‘false’;
- both the label list and NLP feature list are updatable;
- the probability threshold for classification is set as 0.5;
- it uses ‘one against others’ for converting multi-class problem into binary class problems for SVM learning;
- for evaluation it uses hold-out testing with a ratio of 0.66 and only one run.

The configuration file specifies the learning algorithm as the Naive Bayes method implemented in Weka. However, other learning algorithms could equally well be used.

We begin by defining ‘RE_INS’ as the instance type. Next, we provide the numeric identifiers of each argument of the relationship by specifying elements *INSTANCE-ARG1* and *INSTANCE-ARG2* as the feature names ‘arg1’ and ‘arg2’ respectively. This indicates that the argument identifiers of the instances can be found in the ‘arg1’ and ‘arg2’ features of the ‘RE_INS’ annotations.

Attributes might pertain to the entire relation or they might pertain to one or other argument within the relation. We are going to begin by defining the features specific to each argument of the relation. Recall that our ‘RE_INS’ annotations have as arguments two ‘ACEEntity’ annotations, and that these are identified by their ‘MENTION_ID’ being the same as the ‘arg1’ or ‘arg2’ features of the ‘RE_INS’. It is from these ‘ACEEntity’ annotations that we wish to obtain argument-specific features. *FEATURES-ARG1* and *FEATURES-ARG2* elements begin by specifying which annotation we are referring to. We use the *ARG* element to explain this. We are interested in annotations of type ‘ACEEntity’, and their ‘MENTION_ID’ must match ‘arg1’ or ‘arg2’ of ‘RE_INS’ as appropriate. Having identified precisely which ‘ACEEntity’ we are interested in we can go on to give argument-specific features; in this case, unigrams of the ‘Token’ feature ‘string’.

We now wish to define features pertaining to the entire relation. We indicate that the ‘t12’ feature of ‘RE_INS’ annotations is to be used (this feature contains type information derived from ‘ACEEntity’). Again, rather than just specifying the ‘RE_INS’ annotation, we also indicate that the ‘arg1’ and ‘arg2’ feature values must match the argument identifiers of the instance, as defined in the *INSTANCE-ARG1* and *INSTANCE-ARG2* elements at the beginning. This ensures that we are taking our features from the correct annotation.

Finally, we define the class attribute. We indicate that the class attribute is contained in the ‘Relation_type’ feature of the ‘ACERelation’ annotation. The ‘ACERelation’ annotation type has features ‘MENTION_ARG1’ and ‘MENTION_ARG2’, indicating its arguments. Again, we use the elements *ARG1* and *ARG2* to indicate that it is these features that must be matched to the arguments of the instance if that instance is to be considered a positive example of the class.

```
<?xml version="1.0"?>
<ML-CONFIG>
  <ENGINE nickname="NB" implementationName="NaiveBayesWeka"/>
  <DATASET>
    <INSTANCE-TYPE>RE_INS</INSTANCE-TYPE>
    <INSTANCE-ARG1>arg1</INSTANCE-ARG1>
    <INSTANCE-ARG2>arg2</INSTANCE-ARG2>
    <FEATURES-ARG1>
      <ARG>
        <NAME>ARG1</NAME>
        <SEMTYPE>NOMINAL</SEMTYPE>
        <TYPE>ACEEntity</TYPE>
        <FEATURE>MENTION_ID</FEATURE>
      </ARG>
    </FEATURES-ARG1>
  </DATASET>
</ML-CONFIG>
```

```

</ARG>
<ATTRIBUTE>
  <NAME>Form</NAME>
  <SEMTYPE>NOMINAL</SEMTYPE>
  <TYPE>Token</TYPE>
  <FEATURE>string</FEATURE>
  <POSITION>0</POSITION>
</ATTRIBUTE>
</FEATURES-ARG1>
<FEATURES-ARG2>
  <ARG>
    <NAME>ARG2</NAME>
    <SEMTYPE>NOMINAL</SEMTYPE>
    <TYPE>ACEEntity</TYPE>
    <FEATURE>MENTION_ID</FEATURE>
  </ARG>
  <ATTRIBUTE>
    <NAME>Form</NAME>
    <SEMTYPE>NOMINAL</SEMTYPE>
    <TYPE>Token</TYPE>
    <FEATURE>string</FEATURE>
    <POSITION>0</POSITION>
  </ATTRIBUTE>
</FEATURES-ARG2>
<ATTRIBUTE_REL>
  <NAME>EntityCom1</NAME>
  <SEMTYPE>NOMINAL</SEMTYPE>
  <TYPE>RE_INS</TYPE>
  <ARG1>arg1</ARG1>
  <ARG2>arg2</ARG2>
  <FEATURE>t12</FEATURE>
</ATTRIBUTE_REL>
<ATTRIBUTE_REL>
  <NAME>Class</NAME>
  <SEMTYPE>NOMINAL</SEMTYPE>
  <TYPE>ACERelation</TYPE>
  <ARG1>MENTION_ARG1</ARG1>
  <ARG2>MENTION_ARG2</ARG2>
  <FEATURE>Relation_type</FEATURE>
  <CLASS/>
</ATTRIBUTE_REL>
</DATASET>
</ML-CONFIG>

```

15.2.3 How to Use the Batch Learning PR in GATE Developer

The Batch Learning PR implements the procedure of using supervised machine learning for NLP, which generally has two steps; training and application. The training step learns models from labelled data. The application step applies the learned models to the unlabelled data in order to add labels. Therefore, in order to use supervised ML for NLP, one should have some labelled data, which can be obtained either by manually annotating documents or from other resources. One also needs to determine which linguistic features are to be used in training. (The same features should be used in the application as well.) In this implementation, all machine learning attributes are GATE annotation features. Finally, one should determine which learning algorithm will be used.

Based on the general procedure outlined above, we explain how to use the Batch Learning PR step by step below:

1. Annotate some documents with labels that you want to learn. The labels should be represented by the values of a feature of a GATE annotation type (not the annotation type itself).
2. Determine the linguistic features that you want the PR to use for learning.
3. Annotate the documents (training and application) with the desired features. ANNIE can be useful in this regard. Other PRs such as GATE morphological analyser and the parsers may produce useful features as well. You may need to write some JAPE scripts to produce the features you want.
4. Create an XML configuration file for your learning problem. The file should contain one *DATASET* element specifying the NLP features used, one *ENGINE* element specifying the learning algorithm, and some optional settings as necessary. (**Tip:** it may be easier to copy one of the configuration files presented above and modify it for your problem than to write a configuration file from scratch.)
5. Load the training documents containing the required annotations representing the linguistic features and the class label, and put them into a corpus. All linguistic features and the class feature should be in the same annotation set. (The Annotation Set Transfer PR in the ‘Tools’ plugin can be useful here.)
6. Load the Batch Learning PR into GATE Developer. First you need load the plugin named ‘learning’ using the tool *Manage CREOLE Plugins*. Then you can create a new ‘Batch Learning PR’. You will need to provide the configuration file as an initialization parameter. After that you can put the PR into a *Corpus Pipeline* application to use it. Add the corpus containing the training documents to the application too. Set the inputASName to the annotation set containing the annotations for linguistic features and class labels.

7. Set the run-time parameter *learningMode* to ‘TRAINING’ to learn a model from the training data, or set *learningMode* to ‘EVALUATION’ to do evaluation on the training data and get figures indicating the success of the learning. When using evaluation mode, make sure that the *outputASName* is the same as the *inputASName*. (**Tip:** it may save time if you first try evaluation mode on a small number of documents to make sure that the ML PR works well on your problem and outputs reasonable results before training on the large data.)
8. If you want to apply the learned model to new documents, load those new documents into GATE and pre-process them in the same way as the training documents, to ensure that the same features are present. (Class labels need not be present, of course.) Then set *learningMode* to ‘APPLICATION’ and run the PR on this corpus. The application results, namely the new annotations containing the class labels, will be added into the annotation set specified by the *outputASName*.
9. If you just want the feature files produced by the system and do not want to do any learning or application, select the learning mode ‘ProduceFeatureFilesOnly’.

15.2.4 Output of the Batch Learning PR

The Batch Learning PR outputs several different kinds of information. Firstly, it outputs information about the learning settings. This information will be printed in the Messages Window of the GATE Developer (or standard out if using GATE Embedded) and also into the log file ‘logFileForNLPLearning.save’. The amount of information displayed can be determined via the *VERBOSITY* parameter in the configuration file. The main output of the learning system is different for different usage modes. In training mode the system produces the learned models. In application mode it annotates the documents using the learned models. In evaluation mode it displays the evaluation results. Finally, in ‘ProduceFeatureFilesOnly’ mode, it produces feature files for the current corpus. Below, we explain the outputs for different learning modes.

Note that all the files produced by the Batch Learning PR, including the log file, are placed in the sub-directory ‘savedFiles’ of the ML working directory. The ML working directory is the directory containing the configuration file.

Training results

When the Batch Learning PR is used in training mode, its main output is the learned model, stored in a file named ‘learnedModels.save’. For the SVM algorithm, the learned model file is a text file. For the learning algorithms implemented in Weka, the model file is a binary file. The output also includes the feature files described in Section 15.2.4.

Application Results

The main application result is the annotations added to the documents. Those annotations are the results of applying the ML model to the documents. In the configuration file, the annotation type and feature of the class labels are specified; class labels must be the value of a feature of an annotation type. In application mode, those annotation types are created in the new documents, and the feature specified will hold the class label. An additional feature will also be included on the specified annotation type; ‘prob’ will hold the confidence level for the annotation.

Evaluation Results

The Batch Learning PR outputs the evaluation results for each run and also the averaged results over all runs. For each run, it first prints a message about the names of the documents in training and testing corpora respectively. Then it displays the evaluation results of this run; first the results for each class label and then the micro-averaged results over all labels. For each label, it presents the name of the label, the number of instances belonging to the label in the training data and results on the test data; the numbers of correct, partially correct, spurious and missing instances in the testing data, and the precision, recall and F1, calculated using correct only (strict) and correct plus partial (lenient). The F-measure results are obtained using the *AnnotationDiff Tool* which is described in Chapter 10. Finally, the system presents the means of the results of all runs for each label and the micro-averaged results.

Feature Files

The Batch Learning PR is able to produce several feature files. These feature files could be used for evaluating learning algorithms not implemented in this plugin. We describe the formats of those feature files below. Note that all the data files described below can be obtained by setting the run time parameter *learningMode* to ‘ProduceFeatureFilesOnly’, but some may be produced as part of other learning modes.

The **NLP feature file**, named *NLPFeatureData.save*, contains the NLP features of the instances defined in the configuration file. Below is an example of the first few lines of an NLP feature file for information extraction:

```
Class(es) Form(-1) Form(0) Form(1) Ortho(-1) Ortho(0) Ortho(1)
0 ft-airlines-27-jul-2001.xml 512
1 Number_BB _NA[-1] _Form_Seven _Form_UK[1] _NA[-1] _Ortho_upperInitial
  _Ortho_allCaps[1]
1 Country_BB _Form_Seven[-1] _Form_UK _Form_airlines[1] _Ortho_upperInitial[-1]
  _Ortho_allCaps _Ortho_lowercase[1]
```



```

0 _Form_UK[-1] _Form_airlines _Form_including[1] _Ortho_allCaps[-1]
  _Ortho_lowercase _Ortho_lowercase[1]
0 _Form_airlines[-1] _Form_including _Form_British[1] _Ortho_lowercase[-1]
  _Ortho_lowercase _Ortho_upperInitial[1]
1 Airline_BB _Form_including[-1] _Form_British _Form_Airways[1]
  _Ortho_lowercase[-1] _Ortho_upperInitial _Ortho_upperInitial[1]
1 Airline _Form_British[-1] _Form_Airways _Form_[1], _Ortho_upperInitial[-1]
  _Ortho_upperInitial _NA[1]
0 _Form_Airways[-1] _Form_, _Form_Virgin[1] _Ortho_upperInitial[-1] _NA
  _Ortho_upperInitial[1]

```

The first line of the NLP feature file lists the names of all features used. These names are the names the user gave to their features in the configuration file. The number in the parenthesis following a feature name indicates the position of the feature. For example, ‘Form(-1)’ means the Form feature of the token which is immediately before the current token, and ‘Form(0)’ means the Form feature of the current token. The NLP features for all instances are listed for one document before moving on to the next. For each document, the first line shows the index of the document, the document’s name and the number of instances in the document, as shown in the second line above. After that, each line corresponds to an instance in the document, in their order of appearance. The first item on the line is a number n , representing the number of class labels of the instance. Then, the following n items are the labels. If the current instance is the first instance of an entity, its corresponding label has a suffix ‘_BB’. The other items following the label item(s) are the NLP features of the instance, in the order listed in the first line of the file. Each NLP feature contains the feature’s name and value, separated by ‘.’. At the end of one NLP feature, there may be an integer in square brackets, which represents the position of the feature relative to the current instance. If there is no square-bracketed integer at the end of one NLP feature, then the feature is at the position 0.

The **Feature vector file** has the file name ‘featureVectorsData.save’, and stores the feature vector in sparse format for each instance. The first few lines of the feature vector file corresponding to the NLP feature file shown above are as follows:

```

0 512 ft-airlines-27-jul-2001.xml
1 2 1 2 439:1.0 761:1.0 100300:1.0 100763:1.0
2 2 3 4 300:1.0 763:1.0 50439:1.0 50761:1.0 100440:1.0 100762:1.0
3 0 440:1.0 762:1.0 50300:1.0 50763:1.0 100441:1.0 100762:1.0
4 0 441:1.0 762:1.0 50440:1.0 50762:1.0 100020:1.0 100761:1.0
5 1 5 20:1.0 761:1.0 50441:1.0 50762:1.0 100442:1.0 100761:1.0
6 1 6 442:1.0 761:1.0 50020:1.0 50761:1.0 100066:1.0
7 0 66:1.0 50442:1.0 50761:1.0 100443:1.0 100761:1.0

```

The feature vectors are also listed for each document in sequence. For each document, the first line shows the index of the document, the number of instances in the document and the

document's name. Each of the following lines is for each of the instances in the document. The first item in the line is the index of the instance in the document. The second item is a number n , representing the number of labels the instance has. The following n items are indices representing the class labels.

For text classification and relation learning, the label's index comes directly from the label list file, described below. For chunk learning, the label's index presented in the feature vector file is a bit more complicated. If an instance (e.g. token) is the first one of a chunk with label k , then the instance has as the label's index $2 * k - 1$, as shown in the fifth instance. If it is the last instance of the chunk, it has the label's index as $2 * k$, as shown in the sixth instance. If the instance is both the first one and the last one of the chunk (namely the chunk consists of one instance), it has two label indices, $2 * k - 1$ and $2 * k$, as shown in the first and second instances.

The items following the label(s) are the non-zero components of the feature vector. Each component is represented by two numbers separated by ':'. The first number is the dimension (position) of the component in the feature vector, and the second one is the value of the component.

The **Label list file** has the name 'LabelsList.save', and stores a list of labels and their indices. The following is a part of a label list. Each line shows one label name and its index in the label list.

```
Airline 3
Bank 13
CalendarMonth 11
CalendarYear 10
Company 6
Continent 8
Country 2
CountryCapital 15
Date 21
DayOfWeek 4
```

The **NLP feature list** has the name 'NLPFeaturesList.save', and contains a list of NLP features and their indices in the list. The following are the first few lines of an NLP feature list file.

```
totalNumDocs=14915
_EntityType_Date 13 1731
_EntityType_Location 170 1081
_EntityType_Money 523 3774
_EntityType_Organization 12 2387
_EntityType_Person 191 421
```

```

_EntityType_Unknown 76 218
_Form_ ' 112 775
_Form_ \$ 527 74
_Form_ ' 508 37
_Form_ 's 63 731
_Form_ ( 526 111

```

The first line of the file shows the number of instances from which the NLP features were collected. The number of instances will be used for computing of the *idf* (inverse document frequency) in document or sentence classification. The following lines are for the NLP features. Each line is for one unique feature. The first item in the line represents the NLP feature, which is a combination of the feature's name defined in the configuration file and the value of the feature. The second item is a positive integer representing the index of the feature in the list. The last item is the number of times that the feature occurs, which is needed for computing the *idf*.

The **N-grams (or language model) file** has the name 'NgramList.save', and can only be produced by setting the learning mode to 'ProduceFeatureFilesOnly'. In order to produce n-gram data, the user may use a very simple configuration file, i.e. it need only contain the *DATASET* element, and the data element need contain only an *NGRAM* element to specify the type of n-gram and the *INSTANCE-TYPE* element to define the annotation type from which the n-gram data are created (e.g. sentence). The *NGRAM* element in configuration file specifies what type of n-grams the PR produces (see Section 15.2.1 for the explanation of the n-gram definition). For example, if you specify a bigram based on the string form of 'Token', you will obtain a list of bigrams from the corpus you used. The following are the first lines of a bigram list based on the token annotation's 'string' feature, and was calculated over 3 documents.

```

## The following 2-gram were obtained from 3 documents or examples
Aug<>, 3
Female<>; 3
Human<>; 3
2004<>Aug 3
;<>Female 3
.<>The 3
of<>a 3
)<>: 3
,<>and 3
to<>be 3
;<>Human 3

```

The two terms of the bigram are separated by '<>'. The number following one n-gram is the number of occurrences of that n-gram in the corpus. The n-gram list is ordered according

to the number of occurrences of the n-gram terms. The most frequent terms in the corpus are therefore at the start of the list.

The n-gram data produced can be based on any features of annotations available in the documents. Hence it can not only produce the conventional n-gram data based on the token's form or lemma, but also n-grams based on e.g. the token's POS, or a combination of the token's POS and form, or any feature of the 'sentence' annotation (see Section 15.2.1 for how to define different types of n-gram).

The **Document-term matrix file** has the name 'documentByTermMatrix.save', and can only be produced by setting the learning mode to 'ProduceFeatureFilesOnly'. The document-term matrix presents the weights of terms appearing in each document (see Section 19.4 for more explanation). Currently three types of weight are implemented; binary, term frequency (tf) and tf-idf. The binary weight is simply 1 if the term appears in document and 0 if it does not. *tf* (term frequency) refers to the number of occurrences of one term in a document. *tf-idf* is popular in information retrieval and text mining. It is a multiplication of term frequency and inverse document frequency. Inverse document frequency is calculated as follows:

$$idf_i = \log \frac{|D|}{|\{d_j : t_i \in d_j\}|}$$

where $|D|$ is the total number of documents in the corpus, and $|\{d_j : t_i \in d_j\}|$ is the number of documents in which the term t_i appears. The type of weight is specified by the sub-element *ValueTypeNgram* in the *DATASET* element in configuration file (see Section 15.2.1).

Like the n-gram data, in order to produce the document-term matrix, the user may use a very simple configuration file, i.e. it need only contain the *DATASET* element, and the data element need only contain two elements; the *INSTANCE-TYPE* element, to define the annotation type from which the terms are counted, and an *NGRAM* element to specify the type of n-gram. As mentioned previously, the element *ValueTypeNgram* specifies the type of value used in the matrix. If it is not present, the default type *tf-idf* will be used. The conventional document-term matrix can be produced using a unigram based on the token's form or lemma and the instance type covering the whole document. In other words, *INSTANCE-TYPE* is set to an annotation type such as for example 'body', which covers the entire document, and the n-gram definition then specifies the 'string' feature of the 'Token' annotation type.

The following was extracted from the beginning of a document-term matrix file, produced using unigrams of the token's form. It presents a part of the matrix of terms and their term frequency values in the document named '27.xml'. Each term and its term frequency are separated by ':'. The terms are in alphabetic order.

```
0 Documentname="27.xml", has 1 parts: ":2 (:6 ):6 ,:14 -:1 ..:16 /:1
124:1 2004:1 22:1 29:1 330:1 54:1 8:2 ::5 ;:11 Abstract:1 Adaptation:1
```

Adult:1 Atopic:2 Attachment:3 Aug:1 Bindungsicherheit:1 Cross-:1
 Dermatitis:2 English:1 F-SOZU:1 Female:1 Human:1 In:1 Index:1
 Insecure:1 Interpersonal:1 Irrespective:1 It:1 K-:1 Lebensqualität:1
 Life:1 Male:1 NSI:2 Neurodermitis:2 OT:1 Original:1 Patients:1
 Psychological:1 Psychologie:1 Psychosomatik:1 Psychotherapie:1
 Quality:1 Questionnaire:1 RSQ:1 Relations:1 Relationship:1 SCORAD:1
 Scales:1 Sectional:1 Securely:1 Severity:2 Skindex-:1 Social:1
 Studies:1 Suffering:1 Support:1 The:1 Title:1 We:3 [:1]:1 a:4
 absence:1 affection:1 along:2 amount:1 an:1 and:9 as:1 assessed:1
 association:2 atopic:5 attached:7

A list of names of documents processed can also be obtained. The file has the name ‘docsName.save’, and only can be produced by setting the learning mode to ‘ProduceFeatureFilesOnly’. It contains the names of all the documents processed. The first line shows the number of documents in the list. Then, each line lists one document’s name. The first lines of an example file are shown below:

```
##totalDocs=3
ft-bank-of-england-02-aug-2001.xml
ft-airtours-08-aug-2001.xml
ft-airlines-27-jul-2001.xml
```

A list of names of the selected documents for active learning purposes can also be produced. The file has the name ‘ALSelectedDocs.save’. It is a text file. It is produced in ‘ProduceFeatureFilesOnly’ mode. The file contains the names of documents which have been selected for annotating and training in the active learning process. It is used by the ‘RankingDocsForAL’ learning mode to exclude those selected documents from the ranked documents for active learning purposes. When one or more documents are selected for annotating and training, their names should be put into this file, one line per document.

A list of names of ranked documents for active learning purposes; the file has the name ‘ALRankedDocs.save’, and is produced in ‘RankingDocsForAL’ mode. The file contains the list of names of the documents ranked for active learning, according to their usefulness for learning. Those in the front of the list are the most useful documents for learning. The first line in the file shows the total number of documents in the list. Each of other lines in the file lists one document and the averaged confidence score for classifying the document. An example of the file is shown below:

```
##numDocsRanked=3
ft-airlines-27-jul-2001.xml_000201 8.61744
ft-bank-of-england-02-aug-2001.xml_000221 8.672693
ft-airtours-08-aug-2001.xml_000211 9.82562
```

15.2.5 Using the Batch Learning PR from the API

Using the Batch Learning PR from the API is a simple matter if you have some familiarity with GATE Embedded. Chapter 7 provides a more comprehensive introduction to programming with GATE Embedded, and should be consulted for any general points. There is also a complete example program on the code examples page.

The following snippet shows creating a pipeline application, with a corpus, then creating a batch learning PR and adding it to the application. The location of the configuration file and the mode in which the PR is to be run are added to the PR. The application is then run. ‘corpus’ is a GATE corpus that you have previously set up. (To learn more about creating a corpus from GATE Embedded, see chapter 7 or the example at the code examples page.)

```

1 File configFile = new File("/home/you/ml_config.xml"); //Wherever it is
2 RunMode mode = RunMode.EVALUATION; //or TRAINING, or APPLICATION ..
3
4 //Make a pipeline and add the corpus
5 FeatureMap pfm = Factory.newFeatureMap();
6 pfm.put("corpus", corpus);
7 gate.creole.SerialAnalyserController pipeline =
8     (gate.creole.SerialAnalyserController)
9     gate.Factory.createResource("gate.creole.SerialAnalyserController", pfm);
10
11 //Set up the PR and add it to the pipeline.
12 //As with using the PR from GATE Developer, it needs a config file
13 //and a mode.
14 FeatureMap fm = Factory.newFeatureMap();
15 fm.put("configFileURL", configFile.toURI().toURL());
16 fm.put("learningMode", mode);
17 gate.learning.LearningAPIMain learner =
18     (gate.learning.LearningAPIMain)
19     gate.Factory.createResource("gate.learning.LearningAPIMain", fm);
20 pipeline.add(learner);
21
22 //Run it!
23 pipeline.execute();

```

Having run the PR in EVALUATION mode, you can access the results programmatically:

```

1 EvaluationBasedOnDocs ev = learner.getEvaluation();
2 System.out.println(
3     ev.macroMeasuresOfResults.precision + "," +
4     ev.macroMeasuresOfResults.recall + "," +
5     ev.macroMeasuresOfResults.f1 + "," +
6     ev.macroMeasuresOfResults.precisionLenient + "," +
7     ev.macroMeasuresOfResults.recallLenient + "," +
8     ev.macroMeasuresOfResults.f1Lenient + "\n");

```

15.3 Machine Learning PR

The ‘Machine Learning PR’ is GATE’s earlier machine learning PR. It handles both the training and application of ML model on GATE documents. This PR is a Language Analyser so it can be used in all default types of GATE controllers. It can be found in the ‘Machine_Learning’ plugin.

In order to allow for more flexibility, all the configuration parameters for the Machine Learning PR are set through an external XML file and not through the normal PR parameterisation. The root element of the file needs to be called ‘ML-CONFIG’ and it contains two elements: ‘DATASET’ and ‘ENGINE’. An example XML configuration file is given in Section 15.3.6.

15.3.1 The DATASET Element

The DATASET element defines the type of annotation to be used as instance and the set of attributes that characterise all the instances.

An ‘INSTANCE-TYPE’ element is used to select the annotation type to be used for instances, and the attributes are defined by a sequence of ‘ATTRIBUTE’ elements.

For example, if an ‘INSTANCE-TYPE’ has a ‘Token’ for value, there will one instance in the dataset per ‘Token’. This also means that the **positions** (see below) are defined in relation to Tokens. The ‘INSTANCE-TYPE’ can be seen as the smallest unit to be taken into account for the Machine Learning.

An ATTRIBUTE element has the following sub-elements:

- **NAME**: the name of the attribute
- **TYPE**: the annotation type used to extract the attribute.
- **FEATURE** (optional): if present, the value of the attribute will be the value of the named feature on the annotation of specified type.
- **POSITION**: the position of the annotation used to extract the feature relative to the current instance annotation.
- **VALUES**(optional): includes a list of VALUE elements.
- **<CLASS/>**: an empty element used to mark the class attribute. There can only be one attribute marked as class in a dataset definition.

The **VALUES** being defined as XML entities, the characters <, > and & must be replaced by <, > and &. It is recommended to write the XML configuration file in UTF-8 in order that uncommon characters are correctly parsed.

Semantically, there are three types of attributes:

- **nominal attributes:** both type and features are defined and a list of allowed values is provided;
- **numeric:** both type and features are defined but no list of allowed values is provided; it is assumed that the feature can be converted to a number (a double value).
- **boolean:** no feature or list of values is provided; the attribute will take one of the 'true' or 'false' values based on the presence (or absence) of the specified annotation type at the required position.

Figure 15.1 gives some examples of what the values of specified attributes would be in a situation when 'Token' annotations are used as instances.




Figure 15.1: Sample attributes and their values

An `ATTRIBUTE` element is similar to `ATTRIBUTELIST` except that it has no `POSITION` sub-element but a `RANGE` element. This will be converted into several `ATTRIBUTELIST` with position ranging from the value of the attribute 'from' to the value of the attribute 'to'. This can be used in order to avoid the duplication of `ATTRIBUTE` elements.

15.3.2 The ENGINE Element

The ENGINE element defines which particular ML implementation will be used, and allows the setting of options for that particular implementation.

The ENGINE element has three sub-elements:

- **WRAPPER:** defines the class name for the ML implementation (or implementation wrapper). The specified class needs to extend `gate.creole.ml.MLEngine`.
- **BATCH-MODE-CLASSIFICATION:** this element is optional. If present (as an empty element `<BATCH-MODE-CLASSIFICATION />`), the training instances will be passed to the engine in a single batch. If absent, the instances are passed to the engine one at a time. Not every engine supports this option, but for those that do, it can greatly improve performance.
- **OPTIONS:** the contents of the OPTIONS element will be passed verbatim to the ML engine used.

15.3.3 The WEKA Wrapper

The PR provides a wrapper for the WEKA ML Library (<http://www.cs.waikato.ac.nz/ml/weka/>) in the form of the `gate.creole.ml.weka.Wrapper` class.

Options for the WEKA Wrapper

The WEKA wrapper accepts the following options:

- **CLASSIFIER:** the class name for the classifier to be used.
- **CLASSIFIER-OPTIONS:** the options string as required for the classifier.
- **CONFIDENCE-THRESHOLD:** a double value. If the classifier can provide a probability distribution rather than a simple classification then all possible classifications that have a probability value larger or equal to the confidence threshold will be considered.
- **DATASET-FILE:** location of the weka arff file. This item is not mandatory, it is possible to specify the file using the saving option on the GUI.

Training an ML Model with the WEKA Wrapper

The Machine Learning PR has a Boolean runtime parameter named "training". When the value of this parameter is set to true, the PR will collect a dataset of instances from the documents on which it is run. If the classifier used is an updatable classifier then the ML model will be built while collecting the dataset. If the selected classifier is not updatable, then the model will be built the first time a classification is attempted.

Training a model consists of designing a definition file for the ML PR, and creating an application containing a Machine Learning PR. When the application is run over a corpus, the dataset (and the model if possible) is built.

Applying a Learnt Model

Using the same PR, set the 'training' parameter to false and run your application.

Depending on the type of the attribute that is marked as class, different actions will be performed when a classification occurs:

- if the attribute is boolean, a new annotation of the specified type will be created with no features;
- if the attribute is nominal or numeric, a new annotation of the specified type will be created with the feature named in the attribute definition having the value predicted by the classifier.

Once a model is learnt, it can be saved and reloaded at a later time. The WEKA wrapper also provides an operation for saving only the dataset in the ARFF format, which can be used for experiments in the WEKA interface. This could be useful for determining the best algorithm to be used and the optimal options for the selected algorithm.

15.3.4 The MAXENT Wrapper

GATE also provides a wrapper for the Open NLP MAXENT library (<http://maxent.sourceforge.net/about.html>). The MAXENT library provides an implementation of the maximum entropy learning algorithm, and can be accessed using the `gate.creole.ml.maxent.MaxentWrapper` class.

The MAXENT library requires all attributes except for the class attribute to be boolean, and that the class attribute be boolean or nominal. (It should be noted that, within maximum entropy terminology, the class attribute is called the 'outcome'.) Because the MAXENT library does not provide a specific format for data sets, there is no facility to save or load

data sets separately from the model, but if there should be a need to do this, the WEKA wrapper can be used to collect the data.

Training a MAXENT model follows the same general procedure as for WEKA models, but the following difference should be noted. MAXENT models are not updateable, so the model will always be created and trained the first time a classification is attempted. The training of the model might take a considerable amount of time, depending on the amount of training data and the parameters of the model.

Options for the MAXENT Wrapper

- **CUT-OFF:** MAXENT features will only be included in the model if they occur at least this many times. (The default value of this parameter is zero.)
- **ITERATIONS:** The number of times the training procedure should iterate when finding the model's parameters (default is 10). In general no more than about 100 iterations should be needed to train a model, and it is recommended that less are used during development to allow for shorter training times.
- **CONFIDENCE-THRESHOLD:** Same as for the WEKA wrapper (see above). However, if this parameter is not set, or is set to zero, the model will not use a confidence threshold, but will simply return the most likely classification.
- **SMOOTHING:** Use smoothing when training the model. Smoothing can improve the accuracy of the learned models, but it will result in longer training times, and training will use more memory. The size of the learned models will also be larger. Generally smoothing will only improve performance for those models trained from small data sets with a few outcomes. With larger data sets with lots of outcomes, it may make performance worse.
- **SMOOTHING-OBSERVATION:** When using smoothing, this will specify the number of times that trainer will imagine that it has seen features which it did not see (default value is 0.1).
- **VERBOSE:** If selected, this will cause the classifier to output more details of its operation during execution.

15.3.5 The SVM Light Wrapper

The PR provides a wrapper for the SVM Light ML system (<http://svmlight.joachims.org>). SVM Light is a support vector machine implementation, written in C, which is provided as a set of command line programs. The wrapper takes care of the mundane work of converting the data structures between GATE and SVM Light formats, and calls the command line

programs in the right sequence, passing the data back and forth in temporary files. The `<WRAPPER>` value for this engine is `gate.creole.ml.svmlight.SVMLightWrapper`.

The SVM Light binaries themselves are not distributed with GATE – you should download the version for your platform from <http://svmlight.joachims.org> and place `svm_learn` and `svm_classify` on your path.

Classifying documents using the `SVMLightWrapper` is a two phase procedure. In its first phase, `SVMWrapper` collects data from the pre-annotated documents and builds the SVM model using the collected data to classify the unseen documents in its second phase. Below we describe briefly an example of classifying the start time of the seminar in a corpus of email announcing seminars and provide more details later in the section.

Figure 15.2 explains step by step the process of collecting training data for the SVM classifier. GATE documents, which are pre-annotated with the annotations of type *Class* and feature *type='stime'*, are used as the training data. In order to build the SVM model, we require start and end annotations for each *stime* annotation. We use a pre-processor JAPE transduction script to mark the *sTimeStart* and *sTimeEnd* annotations on *stime* annotations. Following this step, the Machine Learning PR (`SVMLightWrapper`) with training mode set to true collects the training data from all training documents. A GATE corpus pipeline, given a set of documents and PRs to execute on them, executes all PRs one by one, only on one document at a time. Unless provided in a separate pipeline, it makes it impossible to send all training data (i.e. collected from all documents) altogether to the `SVMWrapper` using the same pipeline to build the SVM model. This results in the model not being built at the time of collecting training data. The state of the `SVMWrapper` can be saved to an external file once the training data is collected.




Figure 15.2: Flow diagram explaining the SVM training data collection

Before classifying any unseen document, SVM requires the SVM model to be available. In the absence of an up-to-date SVM model, `SVMWrapper` builds a new one using a command line `SVM_learn` utility and the training data collected from the training corpus. In other words, the first SVM model is built when a user tries to classify the first document. At this point the user has an option to save the model somewhere. This is to enable reloading of the model prior to classifying other documents and to avoid rebuilding of the SVM model everytime the user classifies a new set of documents. Once the model becomes available, `SVMWrapper` classifies the unseen documents which creates new *sTimeStart* and *sTimeEnd* annotations over the text. Finally, a post-processor JAPE transduction script is used to combine them into the *sTime* annotation. Figure 15.3 explains this process.

The wrapper allows support vector machines to be created which do either boolean classi-




Figure 15.3: Flow diagram explaining document classifying process

fication or regression (estimation of numeric parameters), and so the class attribute can be boolean or numeric. Additionally, when learning a classifier, SVM Light supports *transduction*, whereby additional examples can be presented during training which do not have the value of the class attribute marked. Presenting such examples can, in some circumstances, greatly improve the performance of the classifier. To make use of this, the class attribute can be a three value nominal, in which case the first value specified for that nominal in the configuration file will be interpreted as *true*, the second as *false* and the third as *unknown*. Transduction will be used with any instances for which this attribute is set to the *unknown* value. It is also possible to use a two value nominal as the class attribute, in which case it will simply be interpreted as *true* or *false*.

The other attributes can be boolean, numeric or nominal, or any combination of these. If an attribute is nominal, each value of that attribute maps to a separate SVM Light feature. Each of these SVM Light features will be given the value 1 when the nominal attribute has the corresponding value, and will be omitted otherwise. If the value of the nominal is not specified in the configuration file or there is no value for an instance, then no feature will be added.

An extension to the basic functionality of SVM Light is that each attribute can receive a weighting. These weighting can be specified in the configuration file by adding `<WEIGHTING>` tags to the parts of the XML file specifying each attribute. The weighting for the attribute must be specified as a numeric value, and be placed between an opening `<WEIGHTING>` tag and a closing `</WEIGHTING>` one. Giving an attribute a greater weighting, will cause it to play a greater role in learning the model and classifying data. This is achieved by multiplying the value of the attribute by the weighting before creating the training or test data that is passed to SVM Light. Any attribute left without an explicitly specified weighting is given a

default weighting of one. Support for these weightings is contained in the Machine Learning PR itself, and so is available to other wrappers, though at time of writing only the SVM Light wrapper makes use of weightings.

As with the MAXENT wrapper, SVM Light models are not updateable, so the model will be trained at the first classification attempt. The SVM Light wrapper supports `<BATCH-MODE-CLASSIFICATION />`, which should be used unless you have a very good reason not to.

The SVM Light wrapper allows both data sets and models to be loaded and saved to files in the same formats as those used by SVM Light when it is run from the command line. When a model is saved, a file will be created which contains information about the state of the SVM Light Wrapper, and which is needed to restore it when the model is loaded again. This file does not, however, contain any information about the SVM Light model itself. If an SVM Light model exists at the time of saving, and that model is up to date with respect to the current state of the training data, then it will be saved as a separate file, with the same name as the file containing information about the state of the wrapper, but with `.NativePart` appended to the filename. These files are in the standard SVM Light model format, and can be used with SVM Light when it is run from the command line. When a model is reloaded by GATE, both of these files must be available, and in the same directory, otherwise an error will result. However, if an up to date trained model does not exist at the time the model is saved, then only one file will be created upon saving, and only that file is required when the model is reloaded. So long as at least one training instance exists, it is possible to bring the model up to date at any point simply by classifying one or more instances (i.e. running the model with the *training* parameter set to false).

Options for the SVM Light Engine

Only one `<OPTIONS>` subelement is currently supported:

- `<CLASSIFIER-OPTIONS>` a string of options to be passed to `svm_learn` on the command line. The only difference is that the user should not specify whether regression or classification is to be used, as the wrapper will detect this automatically, based on the type of the class attribute, and set the option accordingly.

15.3.6 Example Configuration File

```
<?xml version="1.0" encoding="UTF-8"?>
<ML-CONFIG>
  <DATASET>
    <!-- The type of annotation used as instance -->
    <INSTANCE-TYPE>Token</INSTANCE-TYPE>
  <ATTRIBUTE>
    <!-- The name given to the attribute -->
    <NAME>Lookup(0)</NAME>
```

```

    <!-- The type of annotation used as attribute -->
    <TYPE>Lookup</TYPE>
    <!-- The position relative to the instance annotation -->
    <POSITION>0</POSITION>
</ATTRIBUTE>

<ATTRIBUTE>
    <!-- The name given to the attribute -->
    <NAME>Lookup_MT(-1)</NAME>
    <!-- The type of annotation used as attribute -->
    <TYPE>Lookup</TYPE>
    <!-- Optional: the feature name for the feature used to extract values
    for the attribute -->
    <FEATURE>majorType</FEATURE>

    <!-- The position relative to the instance annotation -->
    <POSITION>-1</POSITION>
    <!-- The list of permitted values.
    if present, marks a nominal attribute;
    if absent, the attribute is numeric (double) -->
    <VALUES>
        <!-- One permitted value -->
        <VALUE>address</VALUE>
        <VALUE>cdg</VALUE>
        <VALUE>country_adj</VALUE>
        <VALUE>currency_unit</VALUE>
        <VALUE>date</VALUE>
        <VALUE>date_key</VALUE>
        <VALUE>date_unit</VALUE>
        <VALUE>facility</VALUE>
        <VALUE>facility_key</VALUE>
        <VALUE>facility_key_ext</VALUE>
        <VALUE>govern_key</VALUE>
        <VALUE>greeting</VALUE>
        <VALUE>ident_key</VALUE>
        <VALUE>jobtitle</VALUE>
        <VALUE>loc_general_key</VALUE>
        <VALUE>loc_key</VALUE>
        <VALUE>location</VALUE>
        <VALUE>number</VALUE>
        <VALUE>org_base</VALUE>
        <VALUE>org_ending</VALUE>
        <VALUE>org_key</VALUE>
        <VALUE>org_pre</VALUE>
        <VALUE>organization</VALUE>
        <VALUE>organization_noun</VALUE>
        <VALUE>percent</VALUE>
        <VALUE>person_ending</VALUE>
        <VALUE>person_first</VALUE>
        <VALUE>person_full</VALUE>
        <VALUE>phone_prefix</VALUE>
        <VALUE>sport</VALUE>
        <VALUE>spur</VALUE>
        <VALUE>spur_ident</VALUE>
        <VALUE>stop</VALUE>
        <VALUE>surname</VALUE>
        <VALUE>time</VALUE>
        <VALUE>time_modifier</VALUE>
        <VALUE>time_unit</VALUE>
        <VALUE>title</VALUE>
        <VALUE>year</VALUE>
    </VALUES>
    <!-- Optional: if present marks the attribute used as CLASS
    Only one attribute can be marked as class -->
</ATTRIBUTE>

```

```

<ATTRIBUTE>
  <!-- The name given to the attribute -->
  <NAME>Lookup_MT(0)</NAME>
  <!-- The type of annotation used as attribute -->
  <TYPE>Lookup</TYPE>
  <!-- Optional: the feature name for the feature used to extract values
  for the attribute -->
  <FEATURE>majorType</FEATURE>

  <!-- The position relative to the instance annotation -->
  <POSITION>0</POSITION>
  <!-- The list of permitted values.
  if present, marks a nominal attribute;
  if absent, the attribute is numeric (double)      -->
  <VALUES>
    <!-- One permitted value -->
    <VALUE>address</VALUE>
    <VALUE>cdg</VALUE>
    <VALUE>country_adj</VALUE>
    <VALUE>currency_unit</VALUE>
    <VALUE>date</VALUE>
    <VALUE>date_key</VALUE>
    <VALUE>date_unit</VALUE>
    <VALUE>facility</VALUE>
    <VALUE>facility_key</VALUE>
    <VALUE>facility_key_ext</VALUE>
    <VALUE>govern_key</VALUE>
    <VALUE>greeting</VALUE>
    <VALUE>ident_key</VALUE>
    <VALUE>jobtitle</VALUE>
    <VALUE>loc_general_key</VALUE>
    <VALUE>loc_key</VALUE>
    <VALUE>location</VALUE>
    <VALUE>number</VALUE>
    <VALUE>org_base</VALUE>
    <VALUE>org_ending</VALUE>
    <VALUE>org_key</VALUE>
    <VALUE>org_pre</VALUE>
    <VALUE>organization</VALUE>
    <VALUE>organization_noun</VALUE>
    <VALUE>percent</VALUE>
    <VALUE>person_ending</VALUE>
    <VALUE>person_first</VALUE>
    <VALUE>person_full</VALUE>
    <VALUE>phone_prefix</VALUE>
    <VALUE>sport</VALUE>
    <VALUE>spur</VALUE>
    <VALUE>spur_ident</VALUE>
    <VALUE>stop</VALUE>
    <VALUE>surname</VALUE>
    <VALUE>time</VALUE>
    <VALUE>time_modifier</VALUE>
    <VALUE>time_unit</VALUE>
    <VALUE>title</VALUE>
    <VALUE>year</VALUE>
  </VALUES>
  <!-- Optional: if present marks the attribute used as CLASS
  Only one attribute can be marked as class -->
</ATTRIBUTE>

<ATTRIBUTE>
  <!-- The name given to the attribute -->
  <NAME>Lookup_MT(1)</NAME>
  <!-- The type of annotation used as attribute -->
  <TYPE>Lookup</TYPE>
  <!-- Optional: the feature name for the feature used to extract values
  for the attribute -->

```



```

<FEATURE>majorType</FEATURE>

<!-- The position relative to the instance annotation -->
<POSITION>1</POSITION>

<!-- The list of permitted values.
if present, marks a nominal attribute;
if absent, the attribute is numeric (double) -->
<VALUES>
  <!-- One permitted value -->
  <VALUE>address</VALUE>
  <VALUE>cdg</VALUE>
  <VALUE>country_adj</VALUE>
  <VALUE>currency_unit</VALUE>
  <VALUE>date</VALUE>
  <VALUE>date_key</VALUE>
  <VALUE>date_unit</VALUE>
  <VALUE>facility</VALUE>
  <VALUE>facility_key</VALUE>
  <VALUE>facility_key_ext</VALUE>
  <VALUE>govern_key</VALUE>
  <VALUE>greeting</VALUE>
  <VALUE>ident_key</VALUE>
  <VALUE>jobtitle</VALUE>
  <VALUE>loc_general_key</VALUE>
  <VALUE>loc_key</VALUE>
  <VALUE>location</VALUE>
  <VALUE>number</VALUE>
  <VALUE>org_base</VALUE>
  <VALUE>org_ending</VALUE>
  <VALUE>org_key</VALUE>
  <VALUE>org_pre</VALUE>
  <VALUE>organization</VALUE>
  <VALUE>organization_noun</VALUE>
  <VALUE>percent</VALUE>
  <VALUE>person_ending</VALUE>
  <VALUE>person_first</VALUE>
  <VALUE>person_full</VALUE>
  <VALUE>phone_prefix</VALUE>
  <VALUE>sport</VALUE>
  <VALUE>spur</VALUE>
  <VALUE>spur_ident</VALUE>
  <VALUE>stop</VALUE>
  <VALUE>surname</VALUE>
  <VALUE>time</VALUE>
  <VALUE>time_modifier</VALUE>
  <VALUE>time_unit</VALUE>
  <VALUE>title</VALUE>
  <VALUE>year</VALUE>
</VALUES>
<!-- Optional: if present marks the attribute used as CLASS
Only one attribute can be marked as class -->
</ATTRIBUTE>

<ATTRIBUTE>
  <!-- The name given to the attribute -->
  <NAME>POS_category(-1)</NAME>
  <!-- The type of annotation used as attribute -->
  <TYPE>Token</TYPE>
  <!-- Optional: the feature name for the feature used to extract values
for the attribute -->
  <FEATURE>category</FEATURE>

  <!-- The position relative to the instance annotation -->
  <POSITION>-1</POSITION>

  <!-- The list of permitted values.

```

```

if present, marks a nominal attribute;
if absent, the attribute is numeric (double)    -->
<VALUES>
  <!-- One permitted value -->
  <VALUE>NN</VALUE>
  <VALUE>NNP</VALUE>
  <VALUE>NNPS</VALUE>
  <VALUE>NNS</VALUE>
  <VALUE>NP</VALUE>
  <VALUE>NPS</VALUE>
  <VALUE>JJ</VALUE>
  <VALUE>JJR</VALUE>
  <VALUE>JJS</VALUE>
  <VALUE>JJSS</VALUE>
  <VALUE>RB</VALUE>
  <VALUE>RBR</VALUE>
  <VALUE>RBS</VALUE>
  <VALUE>VB</VALUE>
  <VALUE>VBD</VALUE>
  <VALUE>VBG</VALUE>
  <VALUE>VBN</VALUE>
  <VALUE>VBP</VALUE>
  <VALUE>VBZ</VALUE>
  <VALUE>FW</VALUE>
  <VALUE>CD</VALUE>
  <VALUE>CC</VALUE>
  <VALUE>DT</VALUE>
  <VALUE>EX</VALUE>
  <VALUE>IN</VALUE>
  <VALUE>LS</VALUE>
  <VALUE>MD</VALUE>
  <VALUE>PDT</VALUE>
  <VALUE>POS</VALUE>
  <VALUE>PP</VALUE>
  <VALUE>PRP</VALUE>
  <VALUE>PRP$</VALUE>
  <VALUE>PRPR$</VALUE>
  <VALUE>RP</VALUE>
  <VALUE>TO</VALUE>
  <VALUE>UH</VALUE>
  <VALUE>WDT</VALUE>
  <VALUE>WP</VALUE>
  <VALUE>WP$</VALUE>
  <VALUE>WRB</VALUE>
  <VALUE>SYM</VALUE>
  <VALUE>\ "</VALUE>
  <VALUE>#</VALUE>
  <VALUE>$</VALUE>
  <VALUE>'</VALUE>
  <VALUE>(</VALUE>
  <VALUE>)</VALUE>
  <VALUE>,</VALUE>
  <VALUE>--</VALUE>
  <VALUE>-LRB-</VALUE>
  <VALUE>.</VALUE>
  <VALUE>' '</VALUE>
  <VALUE>:</VALUE>
  <VALUE>::</VALUE>
  <VALUE>'</VALUE>
</VALUES>
  <!-- Optional: if present marks the attribute used as CLASS
  Only one attribute can be marked as class -->
</ATTRIBUTE>

<ATTRIBUTE>
  <!-- The name given to the attribute -->
  <NAME>POS_category(0)</NAME>

```

```

<!-- The type of annotation used as attribute -->
<TYPE>Token</TYPE>
<!-- Optional: the feature name for the feature used to extract values
for the attribute -->
<FEATURE>category</FEATURE>

<!-- The position relative to the instance annotation -->
<POSITION>0</POSITION>

<!-- The list of permitted values.
if present, marks a nominal attribute;
if absent, the attribute is numeric (double) -->
<VALUES>
  <!-- One permitted value -->
  <VALUE>NN</VALUE>
  <VALUE>NNP</VALUE>
  <VALUE>NNPS</VALUE>
  <VALUE>NNS</VALUE>
  <VALUE>NP</VALUE>
  <VALUE>NPS</VALUE>
  <VALUE>JJ</VALUE>
  <VALUE>JJR</VALUE>
  <VALUE>JJS</VALUE>
  <VALUE>JJSS</VALUE>
  <VALUE>RB</VALUE>
  <VALUE>RBR</VALUE>
  <VALUE>RBS</VALUE>
  <VALUE>VB</VALUE>
  <VALUE>VBD</VALUE>
  <VALUE>VBG</VALUE>
  <VALUE>VBN</VALUE>
  <VALUE>VBP</VALUE>
  <VALUE>VBZ</VALUE>
  <VALUE>FW</VALUE>
  <VALUE>CD</VALUE>
  <VALUE>CC</VALUE>
  <VALUE>DT</VALUE>
  <VALUE>EX</VALUE>
  <VALUE>IN</VALUE>
  <VALUE>LS</VALUE>
  <VALUE>MD</VALUE>
  <VALUE>PDT</VALUE>
  <VALUE>POS</VALUE>
  <VALUE>PP</VALUE>
  <VALUE>PRP</VALUE>
  <VALUE>PRP$</VALUE>
  <VALUE>PRPR$</VALUE>
  <VALUE>RP</VALUE>
  <VALUE>TO</VALUE>
  <VALUE>UH</VALUE>
  <VALUE>WDT</VALUE>
  <VALUE>WP</VALUE>
  <VALUE>WP$</VALUE>
  <VALUE>WRB</VALUE>
  <VALUE>SYM</VALUE>
  <VALUE>"</VALUE>
  <VALUE>#</VALUE>
  <VALUE>$</VALUE>
  <VALUE>'</VALUE>
  <VALUE>(</VALUE>
  <VALUE>)</VALUE>
  <VALUE>,</VALUE>
  <VALUE>--</VALUE>
  <VALUE>-LRB-</VALUE>
  <VALUE>.</VALUE>
  <VALUE>' '</VALUE>
  <VALUE>:</VALUE>

```

```

    <VALUE>::</VALUE>
    <VALUE>'</VALUE>
</VALUES>
<!-- Optional: if present marks the attribute used as CLASS
Only one attribute can be marked as class -->
</ATTRIBUTE>

<ATTRIBUTE>
  <!-- The name given to the attribute -->
  <NAME>POS_category(1)</NAME>
  <!-- The type of annotation used as attribute -->
  <TYPE>Token</TYPE>
  <!-- Optional: the feature name for the feature used to extract values
for the attribute -->
  <FEATURE>category</FEATURE>

  <!-- The position relative to the instance annotation -->
  <POSITION>1</POSITION>

  <!-- The list of permitted values.
if present, marks a nominal attribute;
if absent, the attribute is numeric (double) -->
  <VALUES>
    <!-- One permitted value -->
    <VALUE>NN</VALUE>
    <VALUE>NNP</VALUE>
    <VALUE>NNPS</VALUE>
    <VALUE>NNS</VALUE>
    <VALUE>NP</VALUE>
    <VALUE>NPS</VALUE>
    <VALUE>JJ</VALUE>
    <VALUE>JJR</VALUE>
    <VALUE>JJS</VALUE>
    <VALUE>JJSS</VALUE>
    <VALUE>RB</VALUE>
    <VALUE>RBR</VALUE>
    <VALUE>RBS</VALUE>
    <VALUE>VB</VALUE>
    <VALUE>VBD</VALUE>
    <VALUE>VBG</VALUE>
    <VALUE>VBN</VALUE>
    <VALUE>VBP</VALUE>
    <VALUE>VBZ</VALUE>
    <VALUE>FW</VALUE>
    <VALUE>CD</VALUE>
    <VALUE>CC</VALUE>
    <VALUE>DT</VALUE>
    <VALUE>EX</VALUE>
    <VALUE>IN</VALUE>
    <VALUE>LS</VALUE>
    <VALUE>MD</VALUE>
    <VALUE>PDT</VALUE>
    <VALUE>POS</VALUE>
    <VALUE>PP</VALUE>
    <VALUE>PRP</VALUE>
    <VALUE>PRP$</VALUE>
    <VALUE>PRPR$</VALUE>
    <VALUE>RP</VALUE>
    <VALUE>TO</VALUE>
    <VALUE>UH</VALUE>
    <VALUE>WDT</VALUE>
    <VALUE>WP</VALUE>
    <VALUE>WP$</VALUE>
    <VALUE>WRB</VALUE>
    <VALUE>SYM</VALUE>
    <VALUE>\"</VALUE>
    <VALUE>#</VALUE>

```

```

    <VALUE>${</VALUE>
    <VALUE>'</VALUE>
    <VALUE>(</VALUE>
    <VALUE>)</VALUE>
    <VALUE>,</VALUE>
    <VALUE>--</VALUE>
    <VALUE>-LRB-</VALUE>
    <VALUE>.</VALUE>
    <VALUE>'</VALUE>
    <VALUE>:</VALUE>
    <VALUE>::</VALUE>
    <VALUE>'</VALUE>
  </VALUES>
  <!-- Optional: if present marks the attribute used as CLASS
  Only one attribute can be marked as class -->
</ATTRIBUTE>

<ATTRIBUTE>
  <!-- The name given to the attribute -->
  <NAME>Entity(0)</NAME>
  <!-- The type of annotation used as attribute -->
  <TYPE>Entity</TYPE>
  <!-- The position relative to the instance annotation -->
  <POSITION>0</POSITION>

  <CLASS/>
  <!-- Optional: if present marks the attribute used as CLASS
  Only one attribute can be marked as class -->
</ATTRIBUTE>

</DATASET>

<ENGINE>
  <WRAPPER>gate.creole.ml.weka.Wrapper</WRAPPER>
  <OPTIONS>
    <CLASSIFIER OPTIONS="-S -C 0.25 -B -M 2">weka.classifiers.trees.J48</CLASSIFIER>
    <CONFIDENCE-THRESHOLD>0.85</CONFIDENCE-THRESHOLD>
  </OPTIONS>
</ENGINE>
</ML-CONFIG>

```


Chapter 16

Tools for Alignment Tasks

16.1 Introduction

This chapter introduces a new plugin called ‘Alignment’ that comprises of tools to perform text alignment at various level (e.g word, phrase, sentence etc). It allows users to integrate other tools that can be useful for speeding up the alignment process.

Text alignment can be achieved at a document, section, paragraph, sentence and a word level. Given two parallel corpora, where the first corpus contains documents in a source language and the other in a target language, the first task is to find out the parallel documents and align them at the document level. For these tasks one would need to refer to more than one document at the same time. Hence, a need arises for Processing Resources (PRs) which can accept more than one document as parameters. For example given two documents, a source and a target, a Sentence Alignment PR would need to refer to both of them to identify which sentence of the source document aligns with which sentence of the target document. However, the problem occurs when such a PR is part of a corpus pipeline. In a corpus pipeline, only one document from the selected corpus at a time is set on the member PRs. Once the PRs have completed their execution, the next document in the corpus is taken and set on the member PRs. Thus it is not possible to use a corpus pipeline and at the same time supply for than one document to the underlying PRs.

16.2 The Tools

We have introduced a few new resources in GATE that allows processing parallel data. These include resources such as CompoundDocument, CompositeDocument, and a new AlignmentEditor to name a few. Below we describe these components. Please note that all these resources are distributed as part of the ‘Alignment’ plugin and therefore the users should load the plugin first in order to use these resources.

16.2.1 Compound Document

A new Language Resource (LR), called `CompoundDocument`, is introduced which is a collection of documents and allow various documents to be grouped together under a single document. The `CompoundDocument` allows adding more documents to it and removing them if required. It implements the `gate.Document` interface allowing users to carry out all operations that can be done on a normal gate document. For example, if a PR such as Sentence Aligner needs access to two documents (e.g. source and target documents), these documents can be grouped under a single compound document and supplied to the Sentence Alignment PR.

To instantiate `CompoundDocument` user needs to provide the following parameters.

- `encoding` - encoding of the member documents. All document members must have the same encoding (e.g. Unicode, UTF-8, UTF-16).
- `collectRepositioningInfo` - this parameter indicates whether the underlying documents should collect the repositioning information in case the contents of these documents change.
- `preserveOriginalContent` - if the original content of the underlying documents should be preserved.
- `documentIDs` - users need to provide a unique ID for each document member. These ids are used to locate the appropriate documents.
- `sourceUrl` - given a URL of one of the member documents, the instance of `CompoundDocument` searches for other members in the same folder based on the ids provided in the `documentIDs` parameter. Following document name conventions are followed to search other member documents:
 - `FileName.id.extension` (filename followed by id followed the extension and all of these separated by a ‘.’ (dot)).
 - For example if user provides three document IDs (e.g. ‘en’, ‘hi’ and ‘gu’) and selects a file with name ‘File.en.xml’, the `CompoundDocument` will search for rest of the documents (i.e. ‘File.hi.xml’ and ‘File.gu.xml’). The file name (i.e. ‘File’) and the extension (i.e. ‘xml’) remain common for all three members of the compound document.

Figure 16.1 shows a snapshot for instantiating a compound document from GATE Developer.

Compound document provides various methods that help in accessing their individual members.

```
public Document getDocument(String docid);
```





Figure 16.1: Compound Document

The following method returns a map of documents where the key is a document ID and the value is its respective document.

```
public Map getDocuments();
```

Please note that only one member document in a compound document can have focus set on it. Then all the standard document methods of `gate.Document` interface apply to the document with focus set on it. For example, if there are two documents, 'hi' and 'en', and the focus is set on the document 'hi' then the `getAnnotations()` method will return a default annotation set of the 'hi' document. One can use the following method to switch the focus of a compound document to a different document:

```
public void setCurrentDocument(String documentID);
public Document getCurrentDocument();
```

As explained above, new documents can be added to or removed from the compound document using the following method:

```
public void addDocument(String documentID, Document document);
public void removeDocument(String documentID);
```

The following code snippet demonstrates how to create a new compound document using GATE Embedded:

```
1
2 // step 1: initialize GATE
3 Gate.init();
4
5 // step 2: load the Alignment plugin
6 File alignmentHome = new File(Gate.getPluginsHome(), "Alignment");
7 Gate.getCreoleRegister().addDirectory(alignmentHome.toURL());
8
9 // step 3: set the parameters
10 FeatureMap fm = Factory.newFeatureMap();
11
12 // for example you want to create a compound document for
13 // File.id1.xml and File.id2.xml
14 List docIDs = new ArrayList();
15 docIDs.add("id1");
16 docIDs.add("id2");
17 fm.put("documentIDs", docIDs);
18 fm.put("sourceUrl", new URL("file:///url/to/File.id1.xml"));
19
20 // step 4: finally create an instance of compound document
21 Document aDocument = (gate.compound.CompoundDocument)
22     Factory.createResource("gate.compound.impl.CompoundDocumentImpl", fm);
```

16.2.2 Compound Document Editor

The compound document editor is a visual resource (VR) associated with the compound document. The VR contains several tabs - each representing a different member of the compound document. All standard functionalities such as GATE document editor, with all its add-on plugins such as AnnotationSetView, AnnotationsList, coreference editor etc., are available to be used with each individual member.

Figure 16.2 shows a compound document editor with English and Hindi documents as members of the compound document.

16.2.3 Composite Document

The composite document allows users to merge the texts of member documents and keep the merged text linked with their respective member documents. In other words, if users make any change to the composite document (e.g. add new annotations or remove any existing annotations), the relevant effect is made to their respective documents.




Figure 16.2: Compound Document Editor

A PR called `CombineMembersPR` allows creation of a new composite document. It asks for a class name that implements the `CombiningMethod` interface. The `CombiningMethod` tells the `CombineMembersPR` how to combine texts and create a new composite document.

For example, a default implementation of the `CombiningMethod`, called `DefaultCombiningMethod`, takes the following parameters and puts the text of the compound document's members into a new composite document.

```
unitAnnotationType=Sentence
inputASName=Key
copyUnderlyingAnnotations=true;
```

The first parameter tells the combining method that it is the 'Sentence' annotation type whose text needs to be merged and it should be taken from the 'Key' annotation set (second parameter) and finally all the underlying annotations of every Sentence annotation must be copied in the composite document.

If there are two members of a compound document (e.g. ‘hi’ and ‘en’), given the above parameters, the combining method finds out all the annotations of type Sentence from each document and sorts them in ascending order, and one annotation from each document is put one after another in a composite document. This operation continues until all the annotations have been traversed.

Document en	Document hi
Sen1	Shi1
Sen2	Shi2
Sen3	Shi3

```
Document Composite
Sen1
Shi1
Sen2
Shi2
Sen3
Shi3
```

The composite document also maintains a mapping of text offsets such that if someone adds a new annotation to or removes any annotation from the composite document, they are added to or removed from their respective documents. Finally the newly created composite document becomes a member of the same compound document.

16.2.4 DeleteMembersPR

This PR allows deletion of a specific member of the compound document. It takes a parameter called ‘documentID’ and deletes a document with this name.

16.2.5 SwitchMembersPR

As described above, only one member of the compound document can have focus set on it. PRs trying to use the `getDocument()` method get a pointer to the compound document; however all the other methods of the compound document give access to the information of the document member with the focus set on it. So if user wants to process a particular member of the compound document with some PRs, s/he should use the `SwitchMembersPR` that takes one parameter called `documentID` and sets focus to the document with that specific id.

16.2.6 Saving as XML

Calling the `toXml()` method on a compound document returns the XML representation of the member which has focus. However, GATE Developer provides an option to save all member documents in different files. This option appears in the options menu when the user right-clicks on the compound document. The user is asked to provide a name for the directory in which all the members of the compound document will be saved in separate files.

It is also possible to save all members of the compound document in a single XML file. The option, 'Save in a single XML Document', also appears in the options menu. After saving it in a single XML document, the user can use the option 'Compound Document from XML' to load the document back into GATE Developer.

16.2.7 Alignment Editor

A new Visual Resource (VR) called `AlignmentEditor` has been implemented (Figure 16.3 and 16.4) and is attached with every compound document. As the name suggest, the purpose of the `AlignmentEditor` is to allow users to align texts from different members of the compound document at the section, paragraph, sentence and word level. It provides a user-friendly interface in which to perform manual text alignment.

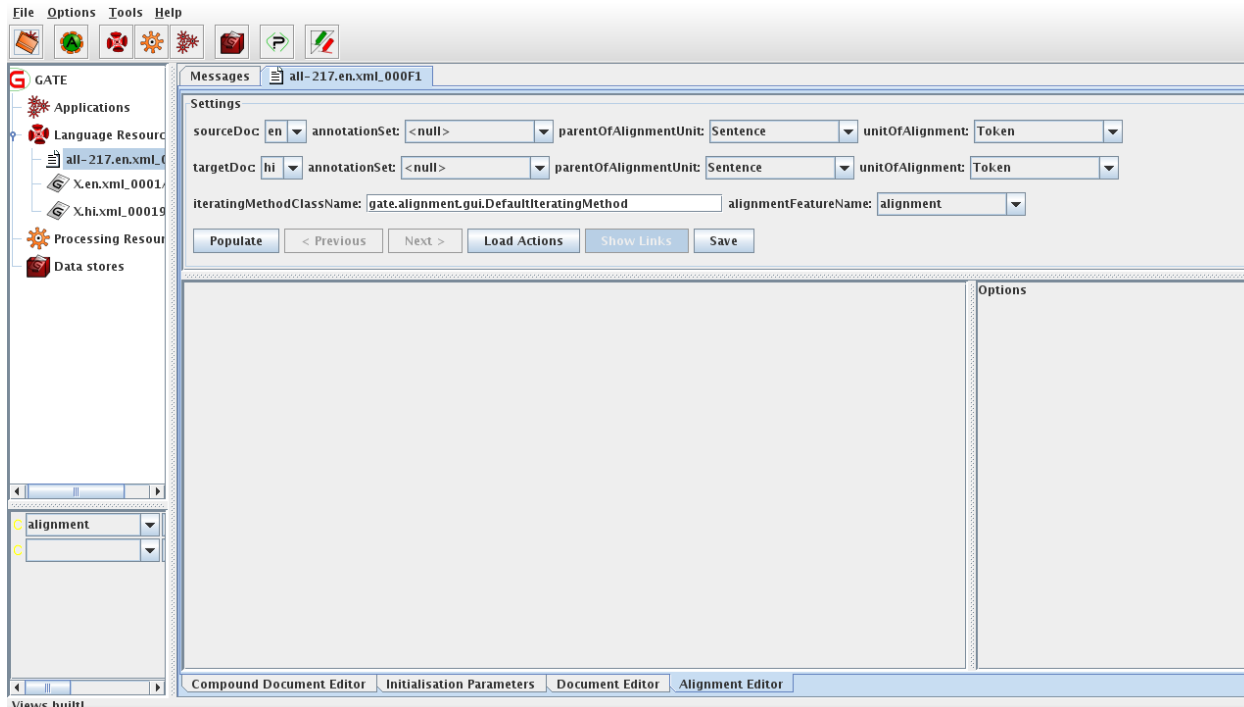


Figure 16.3: Alignment Editor

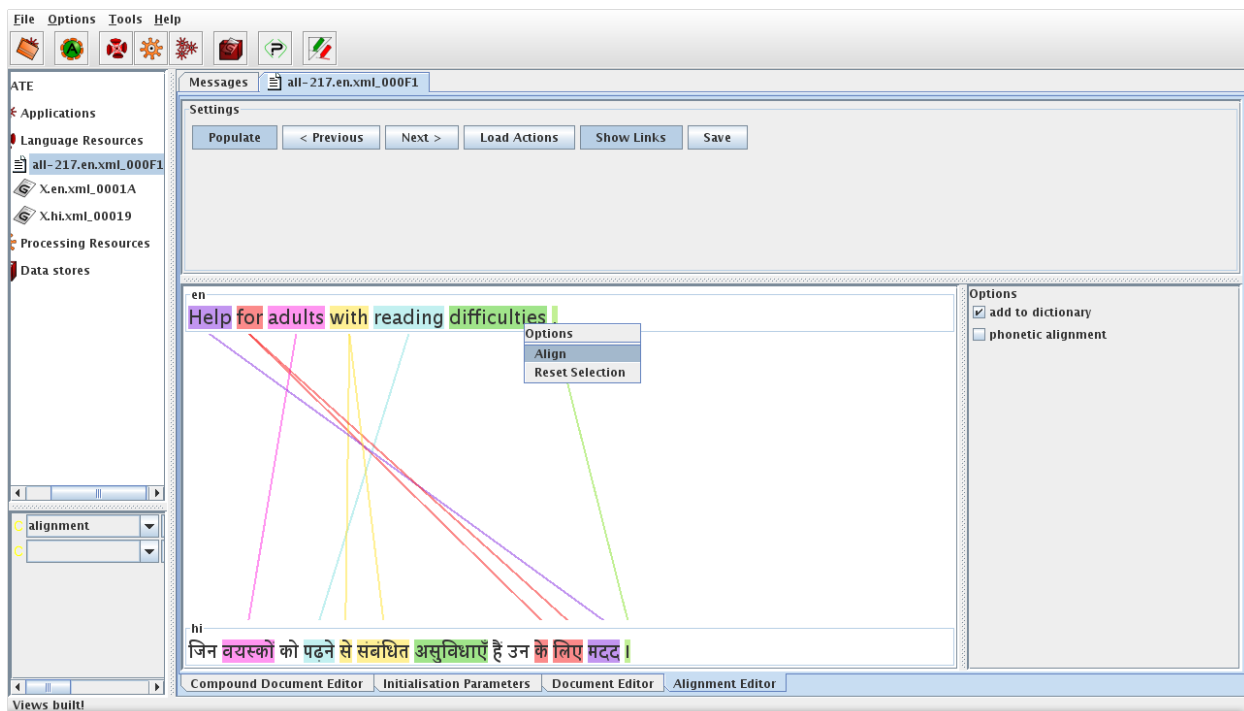


Figure 16.4: Alignment Editor with an example Pair

The user is asked to provide certain parameters based on the requirements of the task. Once the task has been set up, the user is shown some text and is asked to align it. An instance of the `gate.alignment.Alignment` class is created and stored as a document feature on the compound document. This object is then used for storing all the alignment information such as which annotation is aligned with which annotations of what document and so on.

The parameters needed for setting up a new alignment task are as follows:

- **Source and Target Documents:** The user is asked to choose one of the members of the compound document as a source document and an another one as a target document.
- **Annotation Sets:** Users are also asked to choose relevant annotation sets in both the source and the target documents that need to be aligned.
- **Unit Of Alignment:** This is the annotation type that users want to perform alignment on. For example, if users want to align the text at a word level, they will need to process their documents with a tokenizer (e.g. ANNIE English Tokenizer) to generate tokens and provide Token as a unit of alignment.
- **Parent Of Unit Of Alignment:** Generally, if performing a word alignment task, people consider a pair of aligned sentences, one in a source language and the other one in the target language. Thus, the 'Sentence' is a parent of unit of alignment. In other

words, users should also process their documents with a sentence splitter (e.g. AN-NIE Sentence Splitter) that identifies the boundaries of the sentences and creates an annotation for each sentence in the text.

- **Iterating Method:** If the parent of the unit of alignment is Sentence, it is essential to know the order in which sentences from the source and the target documents should be paired together. For example, one could simply specify to pair one sentence from the source document with one sentence from the target document in the order in which they appear in their documents. However, it is possible that the sentences in both documents are not in the correct order or one sentence from the source document refers to more than one sentence in the target document or vice versa. To make it customizable, this parameter allows users to specify a class that implements the `gate.alignment.gui.IteratingMethod` interface. The implementing class is seen as an iterator with a `next()` method that returns an object of `gate.alignment.gui.Pair`, one at a time. One such default implementation, `gate.alignment.gui.DefaultIteratingMethod`, is provided that takes two annotations (of type parent of unit of alignment), one from each, the source and the target documents, in their order of appearance, and forms a pair.
- **Alignment Feature Name:** Information about the alignment (i.e. which annotation is aligned with what annotation) is stored as a document feature. Using this parameter, user can specify the name of the feature that should be used to store the alignment information.

Document en	Document hi
Sen1	Shi1
Sen2	Shi2
Sen3	Shi3

Given a compound document with two members (en and hi) as shown above, if the user selects 'en' as a source document, 'hi' as a target document, 'Key' as an input annotation set, 'Sentence' as a parent of unit of alignment, 'Token' as a value for unit of alignment and 'gate.alignment.gui.DefaultIteratingMethod' as an iterating method, pairs will be created in the following manner.

```
Pair1 Sen1 Shi1
Pair2 Sen2 Shi2
Pair3 Sen3 Sen3
```

Each of these pairs is shown one at a time. If the user clicks on the next button, the next pair of sentences is shown. Similarly clicking on the previous button brings up the previous pair. In each of these sentences, the individual tokens are highlighted with a default colour (to mark the boundary for each unit of alignment). In order to align one or more units in

the source language with one or more units in the target language, the user needs to select them by clicking on them individually. Clicking on units highlights them with an identical colour. Right-clicking on any of the selected units brings up a menu with ‘Align’ and ‘Reset Selection’ options. The user can select ‘Align’ to align the selected units or can select the ‘Reset Selection’ option to reset the selection. If the annotations are unaligned, they are highlighted with the same color and a link (a line with the same color) between them is shown. In order to unalign them, the user needs to right-click on the aligned annotation and click on the ‘Remove Alignment’ option. If the annotation is part of a one-to-one alignment, both the annotations (i.e. the source and the target annotations) are unaligned. However, if there is another annotation in the same pair and the same document that is aligned with the same annotations in the target document, the annotation on which the user right-clicks is taken out of the alignment leaving the rest of the annotations still aligned.

Currently there is no implementation provided to export this alignment information, but one could easily write a PR that reads the information and export it to his/her desired format.

Advanced Features

The editor also allows adding more actions to the editor. There are in total three different types of actions:

- `PreDisplayAction`
- `AlignmentAction`
- `FinishAlignmentAction`

When users click on the next or previous button, the editor obtains a pair to be shown in the editor. Before it is displayed, the editor calls the registered instances of the `PreDisplayAction` and passes them the pair object. This could be helpful in allowing a pair to be preprocessed before it is displayed in the editor. For example, a wrapper could be written for a word alignment algorithm that identifies word alignments in the given sentence pair. More information on the methods of the `PreDisplayAction` interface can be found in the javadoc.

In the case of the word alignment scenario, when a sentence pair is displayed, users can align new words and delete existing ones if needed. This can be achieved by clicking on the relevant buttons in the options menu. All buttons that appear in the options menu are instances of the `AlignmentAction`. As explained earlier, ‘Align’, ‘Reset Selection’ and ‘Remove Alignments’ are the three default buttons that are available to users. The editor also has an ‘options tab’ where users are allowed to add new actions. Users wishing to add new options to this tab or to the options menu need to provide their own implementations of the `AlignmentAction` interface. Below we list some of the methods of the `AlignmentAction` interface.

- `public boolean invokeForAlignedAnnotation()`
- `public boolean invokeForHighlightedUnalignedAnnotation()`
- `public boolean invokeForUnhighlightedUnalignedAnnotation()`
- `public boolean invokeWithAlignAction()`
- `public boolean invokeWithRemoveAction()`
- `getCaption()`

Users may make the visibility of these buttons conditional. For example, the ‘Align’ button appears only when users select unaligned units. The ‘Remove Alignments’ button appears only when users right-click on any of the aligned units. This can be controlled with the help of first three methods as specified above. For example the method ‘`invokeForAlignedAnnotation()`’ indicates that the button should only appear when users right-click on any of the unaligned units.

It is also possible that users might want to perform additional tasks when they click on any of the ‘Align’ or the ‘Remove Alignment’ buttons. For example, users can build a dictionary with new entries while aligning word pairs. In this case, an additional task of adding new entries to the dictionary can be performed when the ‘Align’ button is clicked. On the other hand, not all entries that users align should be included in the dictionary. For the ones which aligners think should go in dictionary, they might want to ask the editor to add them explicitly. All these issues can be controlled by returning appropriate values for the last two methods of the `AlignmentAction` interface (i.e. `invokeWithAlignAction()` and `invokeWithRemoveAction()`).

It is important to note that the new option is added either to the options tab or to the options menu. Users wishing to add it as a button to the options menu must return ‘false’ for the `invokeWithAlignAction()` and `invokeWithRemoveAction()` methods. Users wishing to add it to the options tab, must return true for at least one of these two methods. In case of the latter, the `getCaption()` method is used for obtaining a string that is used for creating a checkbox which is then added to the options tab. When users click on the ‘Align’ or ‘Remove Alignment’ button, the editor also calls the respective actions for the checked checkboxes.

Last of the three types of actions is `FinishedAlignmentAction`. Before users click on the ‘next’ button, they are asked if the pair they were aligning has been aligned completely; in other words, if there is any alignment unit left that still needs to be aligned. If the alignment is complete, the registered instances of the `FinishedAlignmentAction` interface are called. This could be helpful in writing an alignment exporter that takes an aligned pair as input and exports it in an appropriate format.

How to register actions? Having implemented various actions, users need to register them with the alignment editor. In order to do so, users can click on the ‘Load Actions’ button. The user is prompted to provide a configuration file. A configuration file is a simple text file with fully-qualified class names specified in it. After the class name, users can specify any necessary parameters (delimited by a comma sign) that they wish to pass to respective actions classes when they are initialized. Below, we give an example of such an entry in the actions configuration file.

```
#use the class DictionaryBuilder and pass the ‘/user-home/dictionary.txt’ and  
‘root’ as two parameters to the init method of the class.  
gate.alignment.actions.DictionaryBuilder,/user-home/dictionary.txt,root
```

16.2.8 Section-by-Section Processing

In this section, we describe a component that allows processing documents section-by-section. Processing documents this way is useful for many reasons:

For example, a patent document has several different sections but user is interested in processing only the ‘claims’ section or the ‘technical details section’. This is also useful for processing a large document where processing it as a single document is not possible and the only alternative is to divide it in several small documents to process them independently. However, doing so would need another process that merges all the small documents and their annotations back into the original document. On the other hand, a webpage may contain profiles of different people. If the document has more than one person with similar names, running the ‘Orthomatcher PR’ on such a document would produce incorrect coreference chains.

All such problems can be solved by using a PR called ‘Segment Processing PR’. This PR is distributed as part of the ‘Alignment’ plugin. User needs to provide the following four parameters to run this PR.

1. **document**: This is the document to be processed.
2. **controller**: This is a corpus controller that needs to be used for processing the segments of the document.
3. **segmentAnnotationType**: Sections of the documents (that need to be processed) should be annotated with some annotation type and the type of such annotation should be provided as the value to this parameter.
4. **inputASName**: This is the name of the annotation set that contains the segment annotations.

Given these parameters, each span in the document that is annotated as the type specified by the `segmentAnnotationType` is processed independently.

Given a corpus of publications, if you just want to process the abstract section with the ANNIE application, please follow the following steps. It is assumed that the boundaries of abstracts in all these publications are already identified. If not, you would have to do some processing to identify them prior to using the following steps. In the following example, we assume that the abstract boundaries have been annotated as ‘Abstract’ annotations and stored under the ‘Original markups’ annotation set.

Steps:

1. Create a new corpus and populate it with a set of publications that you would like to process with ANNIE.
2. Load the ANNIE application.
3. Load the ‘Alignment’ plugin.
4. Create an instance of the ‘Segment Processing PR’ by selecting it from the list of processing resources.
5. Create a corpus pipeline.
6. Add the ‘Segment Processing PR’ into the pipeline and provide the following parameters:
 - (a) Provide the corpus with publication documents in it as a parameter to the corpus controller.
 - (b) Select the ‘ANNIE’ controller for the ‘controller’ parameter.
 - (c) Type ‘Abstract’ in the ‘segmentAnnotationType’ parameter.
 - (d) Type ‘Original markups’ in the ‘inputASName’ parameter.
7. Run the application.

Now, you should see that the ANNIE application has only processed the text in each document that was annotated as ‘Abstract’.

Chapter 17

Parsers and Taggers

17.1 Verb Group Chunker

The rule-based verb chunker is based on a number of grammars of English [Cobuild 99, Azar 89]. We have developed 68 rules for the identification of non recursive verb groups. The rules cover finite ('is investigating'), non-finite ('to investigate'), participles ('investigated'), and special verb constructs ('is going to investigate'). All the forms may include adverbials and negatives. The rules have been implemented in JAPE. The finite state analyser produces an annotation of type 'VG' with features and values that encode syntactic information ('type', 'tense', 'voice', 'neg', etc.). The rules use the output of the POS tagger as well as information about the identity of the tokens (e.g. the token 'might' is used to identify modals).

The grammar for verb group identification can be loaded as a Jape grammar into the GATE architecture and can be used in any application: the module is domain independent.

17.2 Noun Phrase Chunker

The NP Chunker application is a Java implementation of the Ramshaw and Marcus BaseNP chunker (in fact the files in the resources directory are taken straight from their original distribution) which attempts to insert brackets marking noun phrases in text which have been marked with POS tags in the same format as the output of Eric Brill's transformational tagger. The output from this version should be identical to the output of the original C++/Perl version released by Ramshaw and Marcus.

For more information about baseNP structures and the use of transformation-based learning to derive them, see [Ramshaw & Marcus 95].

17.2.1 Differences from the Original

The major difference is the assumption is made that if a POS tag is not in the mapping file then it is tagged as 'I'. The original version simply failed if an unknown POS tag was encountered. When using the GATE wrapper the chunk tag can be changed from 'I' to any other legal tag (B or O) by setting the unknownTag parameter.

17.2.2 Using the Chunker

The Chunker requires the Creole plugin 'Parser_NP_Chunking' to be loaded. The two load-time parameters are simply urls pointing at the POS tag dictionary and the rules file, which should be set automatically. There are five runtime parameters which should be set prior to executing the chunker.

- `annotationName`: name of the annotation the chunker should create to identify noun phrases in the text.
- `inputASName`: The chunker requires certain types of annotations (e.g. Tokens with part of speech tags) for identifying noun chunks. This parameter tells the chunker which annotation set to use to obtain such annotations from.
- `outputASName`: This is where the results (i.e. new noun chunk annotations will be stored).
- `posFeature`: Name of the feature that holds POS tag information. '
- `unknownTag`: it works as specified in the previous section.

The chunker requires the following PRs to have been run first: tokeniser, sentence splitter, POS tagger.

17.3 Tree Tagger

The TreeTagger plugin ('Tagger_TreeTagger') has been deprecated in favour of the Tagger Framework plugin ('Tagger_Framework'). The Tagger Framework is a generic wrapper around a number of different taggers including the TreeTagger. The TreeTagger plugin will be removed from future versions of GATE and so you should consider updating your applications to use the Tagger Framework, see Section 17.4 for full details.

The TreeTagger is a language-independent part-of-speech tagger, which currently supports English, French, German, Spanish, Italian and Bulgarian (although the latter two are not

available in GATE). It is integrated with GATE using a GATE CREOLE wrapper, originally designed by the CLaC lab (Computational Linguistics at Concordia), Concordia University, Montreal (<http://www.cs.concordia.ca/research/researchgroups/clac/>).

The GATE wrapper calls TreeTagger as an external program, passing gate Tokens as input, and adding two new features to them, which hold the features as described below:

- Features of the TreeTaggerToken:
 - i) category: the part-of-speech tag of the token;
 - ii) lemma: the lemma of the token
- Runtime parameters:
 - i) document: the document to be processed
 - ii) treeTaggerBinary: a URL indicating the location of a (language-specific) GATE TreeTagger wrapper shell script. Note that the scripts used by GATE are different from the original TreeTagger scripts (in cmd), since the latter perform their own tokenisation, whereas the GATE scripts rely on Token annotations as they have been computed by a Tokeniser component. The GATE scripts reside in plugins/Tagger_TreeTagger/resources. Currently available are command scripts for German, French, and Spanish.
 - iii) encoding: The character encoding to use when passing data to and from the tagger. This must be **ISO-8859-1** to work with the standard TreeTagger distribution – do not change it unless you know what you are doing.
 - iv) failOnUnmappableChar: What to do if a character is encountered in the document which cannot be represented in the selected encoding. If the parameter is **true** (the default), unmappable characters cause the wrapper to throw an exception and fail. If set to **false**, unmappable characters are replaced by question marks when the document is passed to the tagger. This is useful if your documents are largely OK but contain the odd character from outside the Latin-1 range.
- Requirement: The TreeTagger, which is available from <http://www.ims.uni-stuttgart.de/projekte/corplex/TreeTagger/DecisionTreeTagger.html>, must be correctly installed on the same machine as GATE. It must be installed in a directory that does not contain any spaces in its path, otherwise the scripts will fail. Once the TreeTagger is installed, the first two lines of the shell script may need to be modified to indicate the installed location of the bin and lib directories of the tagger, as shown below:

```
# THESE VARIABLES HAVE TO BE SET:  
BIN=/usr/local/clactools/TreeTagger/bin  
LIB=/usr/local/clactools/TreeTagger/lib
```

The TreeTagger plugin works on any platform that supports the tree tagger tool, including Linux, Mac OS X and Windows, but the GATE-specific scripts require a POSIX-style Bourne shell with the `gawk`, `tr` and `grep` commands, plus Perl for the Spanish tagger. For Windows this means that you will need to install the appropriate parts of the Cygwin environment from <http://www.cygwin.com> and set the system property `treetagger.sh.path` to contain the path to your `sh.exe` (typically `C:\cygwin\bin\sh.exe`). If this property is set, the TreeTagger plugin runs the shell given in the property and passes the tagger script as its first argument; without the property, the plugin will attempt to run the shell script directly, which fails on Windows with a cryptic ‘error=193’. For GATE Developer, put the following line in `build.properties` (see Section 2.3, and note the extra backslash before each backslash and colon in the path):

```
run.treetagger.sh.path: C:\\cygwin\\bin\\sh.exe
```

Figure 17.1 shows a screenshot of a French document processed with the TreeTagger.

17.3.1 POS Tags

For English the POS tagset is a slightly modified version of the Penn Treebank tagset, where the second letter of the tags for verbs distinguishes between ‘be’ verbs (B), ‘have’ verbs (H) and other verbs (V).

The tagsets for French, German, Italian, Spanish and Bulgarian can be found in the original TreeTagger documentation at <http://www.ims.uni-stuttgart.de/projekte/complex/TreeTagger/DecisionTreeTagger.html>.

17.4 TaggerFramework

The Tagger Framework is an extension of work originally developed in order to provide support for the TreeTagger plugin within GATE. Rather than focusing on providing support for a single external tagger this plugin provides a generic wrapper that can easily be customised (no Java code is required) to incorporate many different taggers within GATE.

The plugin currently provides example applications (see `plugins/Tagger_Framework/resources`) for the following taggers: GENIA (a biomedical tagger), Hunpos (providing support for English and Hungarian), TreeTagger (supporting German, French, Spanish and Italian as well as English), and the Stanford Tagger (supporting English, German and Arabic).

The basic idea behind this plugin is to allow the use of many external taggers. Providing such a generic wrapper requires a few assumptions. Firstly we assume that the external tagger will read from a file and that the contents of this file will be one annotation per line

(i.e. one token or sentence per line). Secondly we assume that the tagger will write its response to stdout and that it will also be based on one annotation per line – although there is no assumption that the input and output annotation types are the same.

- Initialization Parameters

- `preProcessURL`: The URL of a JAPE grammar that should be run over each document before running the tagger.
- `postProcessURL`: The URL of a JAPE grammar that should be run over each document after running the tagger. This can be used, for example, to add chunk annotations using IOB tags output by the tagger and stored as features on Token annotations.

- Runtime Parameters

- `debug`: if set to `true` then a whole heap of useful information will be printed to the messages tab as the tagger runs. Defaults to `false`.
- `encoding`: this must be set to the encoding that the tagger expects the input/output files to use. If this is incorrectly set is highly likely that either the tagger will fail or the results will be meaningless. Defaults to `ISO-8859-1` as this seems to be the most commonly required encoding.
- `failOnUnmappableCharacter`: What to do if a character is encountered in the document which cannot be represented in the selected encoding. If the parameter is `true` (the default), unmappable characters cause the wrapper to throw an exception and fail. If set to `false`, unmappable characters are replaced by question marks when the document is passed to the tagger. This is useful if your documents are largely OK but contain the odd character from outside the Latin-1 range.
- `regex`: this should be a Java regular expression that matches a single line in the output from the tagger. Capturing groups should be used to define the sections of the expression which match the useful output.
- `featureMapping`: this is a mapping from feature name to capturing group in the regular expression. Each feature will be added to the output annotations with a value equal to the specified capturing group. For example, the `TreeTagger` uses a regular expression `(.+)\t(.+)\t(.+)` to capture the three column output. This is then combined with the feature mapping `{string=1, category=2, lemma=3}` to add the appropriate feature/values to the output annotations.
- `inputASName`: the name of the annotation set which should be used for input. If not specified the default (i.e. un-named) annotation set will be used.
- `inputAnnotationType`: the name of the annotation used as input to the tagger. This will usually be `Token`. Note that the input annotations must contain a `string` feature which will be used as input to the tagger. Tokens usually have this feature but if, for example, you wish to use `Sentence` as the input annotation

then you will need to add the `string` feature. JAPE grammars for doing this are provided in `plugins/Tagger_Framework/resources`.

- `outputASName`: the name of the annotation set which should be used for output. If not specified the default (i.e. un-named) annotation set will be used.
- `outputAnnotationType`: the name of the annotation to be provided as output. This is usually `Token`.
- `taggerBinary`: a URL indicating the location of the external tagger. This is usually a shell script which may perform extra processing before executing the tagger. The `plugins/Tagger_Framework/resources` directory contains example scripts (where needed) for the supported taggers. These scripts may need editing (for example, to set the installation directory of the tagger) before they can be used.
- `taggerDir`: the directory from which the tagger must be executed. This can be left unspecified.
- `taggerFlags`: an ordered set of flags that should be passed to the tagger as command line options
- `updateAnnotations`: If set to `true` then the plugin will attempt to update existing output annotations. This can fail if the output from the tagger and the existing annotations are created differently (i.e. the tagger does it's own tokenization). Setting this option to `false` will make the plugin create new output annotations, removing any existing ones, to prevent the two sets getting out of sync. This is also useful when the tagger is domain specific and may do a better job than GATE. For example, the GENIA tagger is better at tokenizing biomedical text than the ANNIE tokenizer. Defaults to `true`.

By default the `GenericTagger` PR simply tries to execute the `taggerBinary` using the normal `Java Runtime.exec()` mechanism. This works fine on Unix-style platforms such as Linux or Mac OS X, but on Windows it will only work if the `taggerBinary` is a `.exe` file. Attempting to invoke other types of program fails on Windows with a rather cryptic “error=193”.

To support other types of tagger programs such as shell scripts or Perl scripts, the `GenericTagger` PR supports a Java system property `shell.path`. If this property is set then instead of invoking the `taggerBinary` directly the PR will invoke the program specified by `shell.path` and pass the tagger binary as the first command-line parameter.

If the tagger program is a shell script then you will need to install the appropriate interpreter, such as `sh.exe` from the cygwin tools, and set the `shell.path` system property to point to `sh.exe`. For GATE Developer you can do this by adding the following line to `build.properties` (see Section 2.3, and note the extra backslash before each backslash and colon in the path):

```
run.shell.path: C:\\cygwin\\bin\\sh.exe
```

Similarly, for Perl or Python scripts you should install a suitable interpreter and set `shell.path` to point to that.

You can also run taggers that are invoked using a Windows batch file (`.bat`). To use a batch file you do not need to use the `shell.path` system property, but instead set the `taggerBinary` runtime parameter to point to `C:\WINDOWS\system32\cmd.exe` and set the first two `taggerFlags` entries to `"/c"` and the Windows-style path to the tagger batch file (e.g. `C:\MyTagger\runTagger.bat`). This will cause the PR to run `cmd.exe /c runTagger.bat` which is the way to run batch files from Java.

17.5 Chemistry Tagger

This GATE module is designed to tag a number of chemistry items in running text. Currently the tagger tags compound formulas (e.g. SO_2 , H_2O , H_2SO_4 ...) ions (e.g. Fe^{3+} , Cl^-) and element names and symbols (e.g. Sodium and Na). Limited support for compound names is also provided (e.g. sulphur dioxide) but only when followed by a compound formula (in parenthesis or commas).

17.5.1 Using the Tagger

The Tagger requires the Creole plugin ‘`Tagger_Chemistry`’ to be loaded. It requires the following PRs to have been run first: `tokeniser` and `sentence splitter` (the annotation set containing the `Tokens` and `Sentences` can be set using the `annotationSetName` runtime parameter). There are four init parameters giving the locations of the two gazetteer list definitions, the element mapping file and the JAPE grammar used by the tagger (in previous versions of the tagger these files were fixed and loaded from inside the `ChemTagger.jar` file). Unless you know what you are doing you should accept the default values.

The annotations added to documents are ‘`ChemicalCompound`’, ‘`ChemicalIon`’ and ‘`ChemicalElement`’ (currently they are always placed in the default annotation set). By default ‘`ChemicalElement`’ annotations are removed if they make up part of a larger compound or ion annotation. This behaviour can be changed by setting the `removeElements` parameter to `false` so that all recognised chemical elements are annotated.

17.6 ABNER

ABNER is A Biomedical Named Entity Recogniser. It uses machine learning (linear-chain conditional random fields, CRFs) to find entities such as genes, cell types, and DNA in text. Full details of ABNER can be found at <http://pages.cs.wisc.edu/~bsettles/abner/>

The ABNER plugin, called ‘Tagger_Abner’, contains a single PR, called AbnerTagger, which wraps ABNER. To use AbnerTagger, first load the Tagger_Abner plugin through the plugins console, and then create a new AbnerTagger PR in the usual way. The AbnerTagger PR has no loadtime parameters (apart from Name). It does not require any other PRs to be run prior to execution.

The AbnerTagger has two runtime parameters:

- **abnerMode** The Abner model that will be used for tagging. The plugin can use one of two previously trained machine learning models for tagging text, as provided by Abner:
 - **BIOCREATIVE** trained on the BioCreative corpus
 - **NLPBA** trained on the NLPBA corpus
- **outputASName** The name of the output annotation set to which AbnerTagger output will be written.

The AbnerTagger creates annotations of type ‘Tagger’ with a feature and value ‘source=abner’. Each annotation may also have features of ‘class’ and ‘type’ set by Abner to values such as:

- Protein
- DNA
- RNA
- Cell Line
- Cell Type
- Gene

Abner does support training of models on other data, but this functionality is not, however, supported by the GATE wrapper.

For further details please refer to the Abner documentation at <http://pages.cs.wisc.edu/~bsettles/abner/>

17.7 Stemmer

The stemmer plugin, ‘Stemmer_Snowball’, consists of a set of stemmers PRs for the following 11 European languages: Danish, Dutch, English, Finnish, French, German, Italian, Norwegian, Portuguese, Russian, Spanish and Swedish. These take the form of wrappers for the Snowball stemmers freely available from <http://snowball.tartarus.org>. Each Token is annotated with a new feature ‘stem’, with the stem for that word as its value. The stemmers should be run as other PRs, on a document that has been tokenised.

There are three runtime parameters which should be set prior to executing the stemmer on a document.

- `annotationType`: This is the type of annotations that represent tokens in the document. Default value is set to ‘Token’.
- `annotationFeature`: This is the name of a feature that contains tokens’ strings. The stemmer uses value of this feature as a string to be stemmed. Default value is set to ‘string’.
- `annotationSetName`: This is where the stemmer expects the annotations of type as specified in the `annotationType` parameter to be.

17.7.1 Algorithms

The stemmers are based on the Porter stemmer for English [Porter 80], with rules implemented in Snowball e.g.

```
define Step_1a as
( [substring] among (
  'sses' (<-'ss')
  'ies' (<-'i')
  'ss' () 's' (delete)
)
```

17.8 GATE Morphological Analyzer

The Morphological Analyser PR can be found in the Tools plugin. It takes as input a tokenized GATE document. Considering one token and its part of speech tag, one at a time, it identifies its lemma and an affix. These values are then added as features on the Token annotation. Morpher is based on certain regular expression rules. These rules were

originally implemented by *Kevin Humphreys* in GATE1 in a programming language called *Flex*. Morpher has a capability to interpret these rules with an extension of allowing users to add new rules or modify the existing ones based on their requirements. In order to allow these operations with as little effort as possible, we changed the way these rules are written. More information on how to write these rules is explained later in Section 17.8.1.

Two types of parameters, Init-time and run-time, are required to instantiate and execute the PR.

- **rulesFile** (Init-time) The rule file has several regular expression patterns. Each pattern has two parts, L.H.S. and R.H.S. L.H.S. defines the regular expression and R.H.S. the function name to be called when the pattern matches with the word under consideration. Please see 17.8.1 for more information on rule file.
- **caseSensitive** (init-time) By default, all tokens under consideration are converted into lowercase to identify their lemma and affix. If the user selects *caseSensitive* to be *true*, words are no longer converted into lowercase.
- **document** (run-time) Here the document must be an instance of a GATE document.
- **affixFeatureName** (run-time) Name of the feature that should hold the affix value.
- **rootFeatureName** (run-time) Name of the feature that should hold the root value.
- **annotationSetName** (run-time) Name of the annotationSet that contains Tokens.
- **considerPOSTag** (run-time) Each rule in the rule file has a separate tag, which specifies which rule to consider with what part-of-speech tag. If this option is set to false, all rules are considered and matched with all words. This option is very useful. For example if the word under consideration is "singing". "singing" can be used as a noun as well as a verb. In the case where it is identified as a verb, the lemma of the same would be "sing" and the affix "ing", but otherwise there would not be any affix.
- **failOnMissingInputAnnotations** (run-time) If set to true (the default) the PR will terminate with an Exception if none of the required input Annotations are found in a document. If set to false the PR will not terminate and instead log a single warning message per session and a debug message per document that has no input annotations.

17.8.1 Rule File

GATE provides a default rule file, called *default.rul*, which is available under the *gate/plugins/Tools/morph/resources* directory. The rule file has two sections.

1. Variables
2. Rules

Variables

The user can define various types of variables under the section *defineVars*. These variables can be used as part of the regular expressions in rules. There are three types of variables:

1. Range With this type of variable, the user can specify the range of characters. e.g. A ==> [-a-z0-9]
2. Set With this type of variable, user can also specify a set of characters, where one character at a time from this set is used as a value for the given variable. When this variable is used in any regular expression, all values are tried one by one to generate the string which is compared with the contents of the document. e.g. A ==> [abcdqurs09123]
3. Strings Where in the two types explained above, variables can hold only one character from the given set or range at a time, this allows specifying strings as possibilities for the variable. e.g. A ==> 'bb' OR 'cc' OR 'dd'

Rules

All rules are declared under the section *defineRules*. Every rule has two parts, LHS and RHS. The LHS specifies the regular expression and the RHS the function to be called when the LHS matches with the given word. '==>' is used as delimiter between the LHS and RHS.

The LHS has the following syntax:

< " *"—"verb"—"noun" >< *regularexpression* >.

User can specify which rule to be considered when the word is identified as 'verb' or 'noun'. '*' indicates that the rule should be considered for all part-of-speech tags. If the part-of-speech should be used to decide if the rule should be considered or not can be enabled or disabled by setting the value of *considerPOSTags* option. Combination of any string along with any of the variables declared under the *defineVars* section and also the Kleene operators, '+' and '*', can be used to generate the regular expressions. Below we give few examples of L.H.S. expressions.

- <verb>"bias"
- <verb>"canvas"{ESEDING} "ESEDING" is a variable defined under the *defineVars* section. Note: variables are enclosed with "{" and "}".
- <noun>({A}* "metre") "A" is a variable followed by the Kleene operator "*", which means "A" can occur zero or more times.

- $\langle \text{noun} \rangle (\{A\}^+ \text{itis})$ "A" is a variable followed by the Kleene operator "+", which means "A" can occur one or more times.
- $\langle * \rangle$ "aches" " $\langle * \rangle$ " indicates that the rule should be considered for all part-of-speech tags.

On the RHS of the rule, the user has to specify one of the functions from those listed below. These rules are hard-coded in the Morph PR in GATE and are invoked if the regular expression on the LHS matches with any particular word.

- $\text{stem}(n, \text{string}, \text{affix})$ Here,
 - n = number of characters to be truncated from the end of the string.
 - string = the string that should be concatenated after the word to produce the root.
 - affix = affix of the word
- $\text{irreg_stem}(\text{root}, \text{affix})$ Here,
 - root = root of the word
 - affix = affix of the word
 - $\text{null_stem}()$ This means words are themselves the base forms and should not be analyzed.
- $\text{semi_reg_stem}(n, \text{string})$ semir_reg_stem function is used with the regular expressions that end with any of the {EDING} or {ESEDING} variables defined under the variable section. If the regular expression matches with the given word, this function is invoked, which returns the value of variable (i.e. {EDING} or {ESEDING}) as an affix. To find a lemma of the word, it removes the n characters from the back of the word and adds the string at the end of the word.

17.9 MiniPar Parser

MiniPar is a shallow parser. In its shipped version, it takes one sentence as an input and determines the dependency relationships between the words of a sentence. It parses the sentence and brings out the information such as:

- the lemma of the word;
- the part of speech of the word;
- the head modified by this word;

- name of the dependency relationship between this word and the head;
- the lemma of the head.

In the version of MiniPar integrated in GATE ('Parser_Minipar' plugin), it generates annotations of type 'DepTreeNode' and the annotations of type '[relation]' that exists between the head and the child node. The document is required to have annotations of type 'Sentence', where each annotation consists of a string of the sentence.

Minipar takes one sentence at a time as an input and generates the tokens of type 'DepTreeNode'. Later it assigns relation between these tokens. Each DepTreeNode consists of feature called 'word': this is the actual text of the word.

For each and every annotation of type '[Rel]', where 'Rel' is obj, pred etc. This is the name of the dependency relationship between the child word and the head word (see Section 17.9.5). Every '[Rel]' annotation is assigned four features:

- **child_word**: this is the text of the child annotation;
- **child_id**: IDs of the annotations which modify the current word (if any).
- **head_word**: this is the text of the head annotation;
- **head_id**: ID of the annotation modified by the child word (if any);

Messages GATE document_0001E

Annotation Sets Annotations Co-reference Editor Text

Type	Set	Start	End	Features
Token		0	2	{category=DET:ART, kind=word, lemma=un, length=2, orth=un}
Token		3	9	{category=NOM, kind=word, lemma=soldat, length=6, orth=soldat}
Token		10	19	{category=ADJ, kind=word, lemma=américain, length=9, orth=américain}
Token		20	23	{category=VER:pper, kind=word, lemma=tuer, length=3, orth=tuer}
Token		23	24	{category=PUN, kind=punctuation, lemma=., length=1, orth=.
Token		25	29	{category=NUM, kind=word, lemma=deux, length=4, orth=deux}
Token		30	37	{category=NOM, kind=word, lemma=blessé, length=7, orth=blessé}
Token		38	42	{category=PRP, kind=word, lemma=dans, length=4, orth=dans}
Token		43	46	{category=DET:ART, kind=word, lemma=un, length=3, orth=un}

185 Annotations (0 selected)

Un soldat américain tué, deux blessés dans une attaque en Irak
AFP | 07.10.04 | 05h45

Un soldat américain a été tué et deux autres blessés mercredi soir, lors de l'attaque de leur convoi près du bastion sunnite rebelle de Falloujah, à l'ouest de Bagdad, a annoncé jeudi l'armée américaine dans un communiqué. Un "engin explosif de type inconnu" a frappé le convoi à environ 21H45 (18H45 GMT), indique le communiqué. Les trois soldats ont été emmenés à l'hôpital militaire où l'un est mort de ses blessures une heure après son admission. L'un des blessés est grièvement atteint, tandis que l'état de l'autre s'est

Figure 17.1: a TreeTagger processed document

Figure 17.2 shows a MiniPar annotated document in GATE Developer.

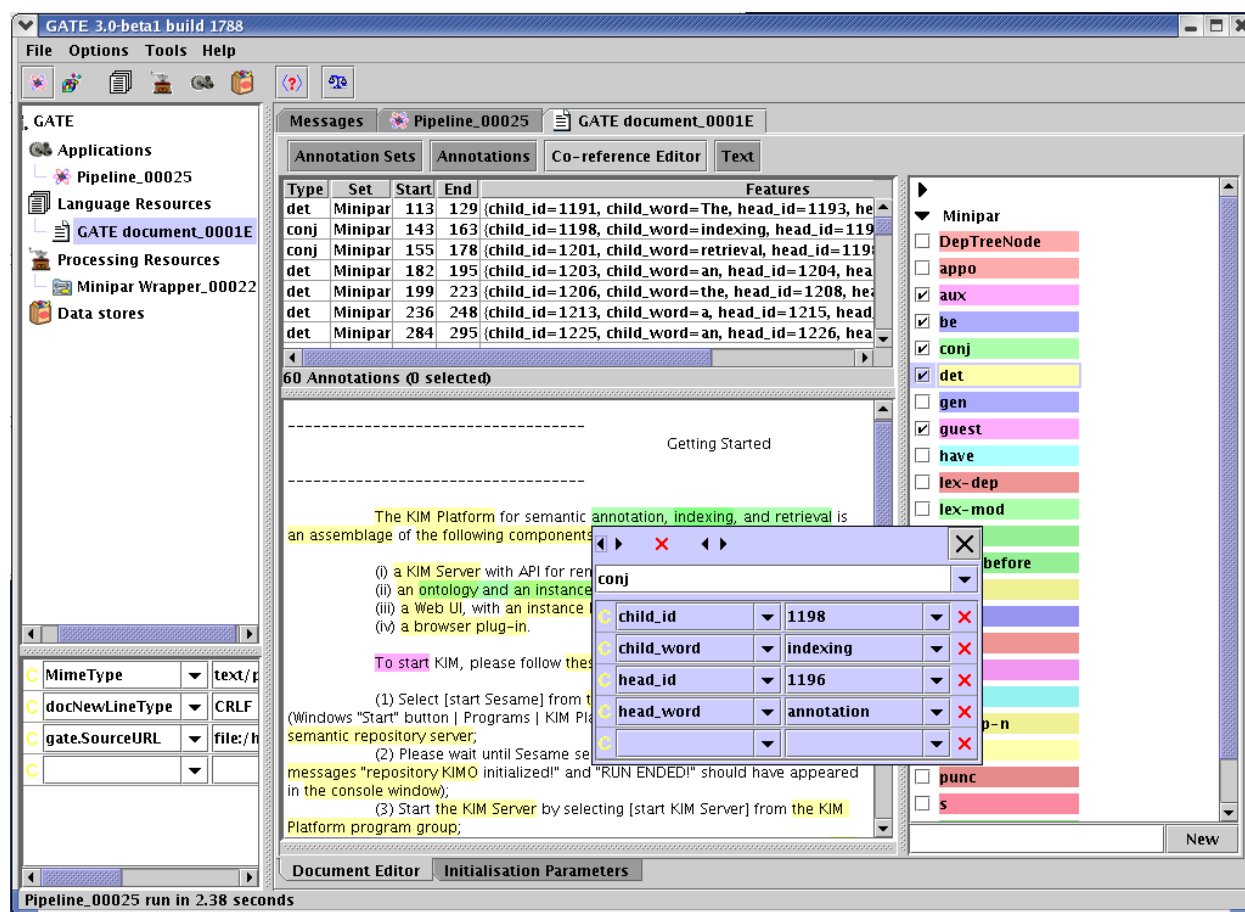


Figure 17.2: a MiniPar annotated document

17.9.1 Platform Supported

MiniPar in GATE is supported for the Linux and Windows operating systems. Trying to instantiate this PR on any other OS will generate the ResourceInstantiationException.

17.9.2 Resources

MiniPar in GATE is shipped with four basic resources:

- **MiniparWrapper.jar**: this is a JAVA Wrapper for MiniPar;
- **creole.XML**: this defines the required parameters for MiniPar Wrapper;

- **minipar.linux**: this is a modified version of pdemo.cpp.
- **minipar-windows.exe** : this is a modified version of pdemo.cpp compiled to work on windows.

17.9.3 Parameters

The MiniPar wrapper takes six parameters:

- **annotationTypeName**: new annotations are created with this type, default is "DepTreeNode";
- **annotationInputSetName**: annotations of Sentence type are provided as an input to MiniPar and are taken from the given annotationSet;
- **annotationOutputSetName**: All annotations created by Minipar Wrapper are stored under the given annotationOutputSet;
- **document**: the GATE document to process;
- **miniparBinary**: location of the MiniPar Binary file (i.e. either minipar.linux or minipar-windows.exe. These files are available under gate/plugins/minipar/ directory);
- **miniparDataDir**: location of the 'data' directory under the installation directory of MINIPAR. default is "%MINIPAR_HOME%/data".

17.9.4 Prerequisites

The MiniPar wrapper requires the MiniPar library to be available on the underlying Linux/Windows machine. It can be downloaded from the MiniPar homepage.

17.9.5 Grammatical Relationships

```

appo    "ACME president, --appo-> P.W. Buckman"
aux     "should <-aux-- resign"
be      "is <-be-- sleeping"
c       "that <-c-- John loves Mary"
comp1   first complement
det     "the <-det '-- hat"
gen     "Jane's <-gen-- uncle"
i       the relationship between a C clause and its I clause
inv-aux inverted auxiliary: "Will <-inv-aux-- you stop it?"

```

```

inv-be      inverted be: "Is <-inv-be-- she sleeping"
inv-have    inverted have: "Have <-inv-have-- you slept"
mod the relationship between a word and its adjunct modifier
pnmod      post nominal modifier
p-spec      specifier of prepositional phrases
pcomp-c     clausal complement of prepositions
pcomp-n     nominal complement of prepositions
post       post determiner
pre        pre determiner
pred       predicate of a clause
rel        relative clause
vrel       passive verb modifier of nouns
wha, whn, whp: wh-elements at C-spec positions
obj        object of verbs
obj2       second object of ditransitive verbs
subj       subject of verbs
s         surface subjec

```

17.10 RASP Parser

RASP (Robust Accurate Statistical Parsing) is a robust parsing system for English, developed by the Natural Language and Computational Linguistics group at the University of Sussex.

This plugin, ‘Parser_RASP’, developed by DigitalPebble, provides four wrapper PRs that call the RASP modules as external programs, as well as a JAPE component that translates the output of the ANNIE POS Tagger (Section 6.6).

RASP2 Tokenizer This PR requires **Sentence** annotations and creates **Token** annotations with a **string** feature. Note that sentence-splitting must be carried out before tokenization; the the RegEx Sentence Splitter (see Section 6.5) is suitable for this. (Alternatively, you can use the ANNIE Tokenizer (Section 6.2) and then the ANNIE Sentence Splitter (Section 6.4); their output is compatible with the other PRs in this plugin).

RASP2 POS Tagger This requires **Token** annotations and creates **WordForm** annotations with **pos**, **probability**, and **string** features.

RASP2 Morphological Analyser This requires **WordForm** annotations (from the POS Tagger) and adds **lemma** and **suffix** features.

RASP2 Parser This requires the preceding annotation types and creates multiple **Dependency** annotations to represent a parse of each sentence.

RASP POS Converter This PR requires `Token` annotations with a `category` feature as produced by the ANNIE POS Tagger (see Section 6.6 and creates `WordForm` annotations in the RASP Format. The ANNIE POS Tagger and this Converter can together be used as a substitute for the RASP2 POS Tagger.

Here are some examples of corpus pipelines that can be correctly constructed with these PRs.

1. RegEx Sentence Splitter
2. RASP2 Tokenizer
3. RASP2 POS Tagger
4. RASP2 Morphological Analyser
5. RASP2 Parser

1. RegEx Sentence Splitter
2. RASP2 Tokenizer
3. ANNIE POS Tagger
4. RASP POS Converter
5. RASP2 Morphological Analyser
6. RASP2 Parser

1. ANNIE Tokenizer
2. ANNIE Sentence Splitter
3. RASP2 POS Tagger
4. RASP2 Morphological Analyser
5. RASP2 Parser

1. ANNIE Tokenizer
2. ANNIE Sentence Splitter
3. ANNIE POS Tagger
4. RASP POS Converter

5. RASP2 Morphological Analyser

6. RASP2 Parser

Further documentation is included in the directory `gate/plugins/Parser_RASP/doc/`.

The RASP package, which provides the external programs, is available from the RASP web page.

RASP is only supported for Linux operating systems. Trying to run it on any other operating systems will generate an exception with the message: ‘The RASP cannot be run on any other operating systems except Linux.’

It must be correctly installed on the same machine as GATE, and must be installed in a directory whose path does not contain any spaces (this is a requirement of the RASP scripts as well as the wrapper). Before trying to run scripts for the first time, edit `rasp.sh` and `rasp_parse.sh` to set the correct value for the shell variable `RASP`, which should be the file system pathname where you have installed the RASP tools (for example, `RASP=/opt/RASP` or `RASP=/usr/local/RASP`. You will need to enter the same path for the initialization parameter `raspHome` for the POS Tagger, Morphological Analyser, and Parser PRs.

(On some systems the `arch` command used in the scripts is not available; a work-around is to comment that line out and add `arch='ix86_linux'`, for example.)

(The previous version of the RASP plugin can now be found in `plugins/Obsolete/rasp`.)

17.11 SUPPLE Parser

SUPPLE is a bottom-up parser that constructs syntax trees and logical forms for English sentences. The parser is complete in the sense that every analysis licensed by the grammar is produced. In the current version only the ‘best’ parse is selected at the end of the parsing process. The English grammar is implemented as an attribute-value context free grammar which consists of subgrammars for noun phrases (NP), verb phrases (VP), prepositional phrases (PP), relative phrases (R) and sentences (S). The semantics associated with each grammar rule allow the parser to produce logical forms composed of unary predicates to denote entities and events (e.g., *chase(e1)*, *run(e2)*) and binary predicates for properties (e.g. *lsubj(e1,e2)*). Constants (e.g., *e1*, *e2*) are used to represent entity and event identifiers. The GATE SUPPLE Wrapper stores syntactic information produced by the parser in the gate document in the form of `parse` annotations containing a bracketed representation of the parse; and `semantics` annotations that contains the logical forms produced by the parser. It also produces `SyntaxTreeNode` annotations that allow viewing of the parse tree for a sentence (see Section 17.11.4).

17.11.1 Requirements

The SUPPLE parser is written in Prolog, so you will need a Prolog interpreter to run the parser. A copy of PrologCafe (<http://kaminari.scitec.kobe-u.ac.jp/PrologCafe/>), a pure Java Prolog implementation, is provided in the distribution. This should work on any platform but it is not particularly fast. SUPPLE also supports the open-source SWI Prolog (<http://www.swi-prolog.org>) and the commercially licenced SICStus prolog (<http://www.sics.se/sicstus>, SUPPLE supports versions 3 and 4), which are available for Windows, Mac OS X, Linux and other Unix variants. For anything more than the simplest cases we recommend installing one of these instead of using PrologCafe.

17.11.2 Building SUPPLE

The SUPPLE plugin must be compiled before it can be used, so you will require a suitable Java SDK (GATE itself requires only the JRE to run). To build SUPPLE, first edit the file `build.xml` in the `Parser_SUPPLE` directory under `plugins`, and adjust the user-configurable options at the top of the file to match your environment. In particular, if you are using SWI or SICStus Prolog, you will need to change the `swi.executable` or `sicstus.executable` property to the correct name for your system. Once this is done, you can build the plugin by opening a command prompt or shell, going to the `Parser_SUPPLE` directory and running:

```
../../bin/ant swi
```

(on Windows, use `..\..\bin\ant`). For PrologCafe or SICStus, replace `swi` with `plcafe` or `sicstus` as appropriate.

17.11.3 Running the Parser in GATE

In order to parse a document you will need to construct an application that has:

- tokeniser
- splitter
- POS-tagger
- Morphology
- SUPPLE Parser with parameters
 - mapping file (`config/mapping.config`)
 - feature table file (`config/feature_table.config`)

parser file (supple.plcafe or supple.sicstus or supple.swi)

prolog implementation (shef.nlp.supple.prolog.PrologCafe, shef.nlp.supple.prolog.SICStusProlog3, shef.nlp.supple.prolog.SICStusProlog4, shef.nlp.supple.prolog.SWIProlog or shef.nlp.supple.prolog.SWIJavaProlog¹).

You can take a look at build.xml to see examples of invocation for the different implementations.

Note that prior to GATE 3.1, the parser file parameter was of type `java.io.File`. From 3.1 it is of type `java.net.URL`. If you have a saved application (.gapp file) from before GATE 3.1 which includes SUPPLE it will need to be updated to work with the new version. Instructions on how to do this can be found in the README file in the SUPPLE plugin directory.

17.11.4 Viewing the Parse Tree

GATE Developer provides a syntax tree viewer in the `Tools` plugin which can display the parse tree generated by SUPPLE for a sentence. To use the tree viewer, be sure that the `Tools` plugin is loaded, then open a document in GATE Developer that has been processed with SUPPLE and view its `Sentence` annotations. Right-click on the relevant `Sentence` annotation in the annotations table and select 'Edit with syntax tree viewer'. This viewer can also be used with the constituency output of the Stanford Parser PR (Section 17.12).

17.11.5 System Properties

The `SICStusProlog` (3 and 4) and `SWIProlog` implementations work by calling the native prolog executable, passing data back and forth in temporary files. The location of the prolog executable is specified by a system property:

- for `SICStus`: `supple.sicstus.executable` - default is to look for `sicstus.exe` (Windows) or `sicstus` (other platforms) on the PATH.
- for `SWI`: `supple.swi.executable` - default is to look for `plcon.exe` (Windows) or `swipl` (other platforms) on the PATH.

If your prolog is installed under a different name, you should specify the correct name in the relevant system property. For example, when installed from the source distribution, the Unix version of SWI prolog is typically installed as `pl`, most binary packages install it as

¹shef.nlp.supple.prolog.SICStusProlog exists for backwards compatibility and behaves the same as `SICStusProlog3`.

swipl, though some use the name `swi-prolog`. You can also use the properties to specify the full path to prolog (e.g. `/opt/swi-prolog/bin/pl`) if it is not on your default PATH.

For details of how to pass system properties to GATE, see the end of Section 2.3.

17.11.6 Configuration Files

Two files are used to pass information from GATE to the SUPPLE parser: the *mapping* file and the *feature table* file.

Mapping File

The mapping file specifies how annotations produced using GATE are to be passed to the parser. The file is composed of a number of pairs of lines, the first line in a pair specifies a GATE annotation we want to pass to the parser. It includes the AnnotationSet (or default), the AnnotationType, and a number of features and values that depend on the AnnotationType. The second line of the pair specifies how to encode the GATE annotation in a SUPPLE syntactic category, this line also includes a number of features and values. As an example consider the mapping:

```
Gate;AnnotationType=Token;category=DT;string=&S
SUPPLE;category=dt;m_root=&S;s_form=&S
```

It specifies how a determinant ('DT') will be translated into a category 'dt' for the parser. The construct '&S' is used to represent a variable that will be instantiated to the appropriate value during the mapping process. More specifically a token like 'The' recognised as a DT by the POS-tagging will be mapped into the following category:

```
dt(s_form:'The',m_root:'The',m_affix:'_',text:'_').
```

As another example consider the mapping:

```
Gate;AnnotationType=Lookup;majorType=person_first;minorType=female;string=&S
SUPPLE;category=list_np;s_form=&S;ne_tag=person;ne_type=person_first;gender=female
```

It specified that an annotation of type 'Lookup' in GATE is mapped into a category 'list_np' with specific features and values. More specifically a token like 'Mary' identified in GATE as a Lookup will be mapped into the following SUPPLE category:

```
list_np(s_form:'Mary',m_root:'_',m_affix:'_',
text:'_',ne_tag:'person',ne_type:'person_first',gender:'female').
```

Feature Table

The feature table file specifies SUPPLE ‘lexical’ categories and its features. As an example an entry in this file is:

```
n;s_form;m_root;m_affix;text;person;number
```

which specifies which features and in which order a noun category should be written. In this case:

```
n(s_form:...,m_root:...,m_affix:...,text:...,person:...,number:....).
```

17.11.7 Parser and Grammar

The parser builds a semantic representation compositionally, and a ‘best parse’ algorithm is applied to each final chart, providing a partial parse if no complete sentence span can be constructed. The parser uses a feature valued grammar. Each **Category** entry has the form:

```
Category(Feature1:Value1,...,FeatureN:ValueN)
```

where the number and type of features is dependent on the category type (see Section 5.1). All categories will have the features **s_form** (surface form) and **m_root** (morphological root); nominal and verbal categories will also have **person** and **number** features; verbal categories will also have **tense** and **vform** features; and adjectival categories will have a **degree** feature. The **list_np** category has the same features as other nominal categories plus **ne_tag** and **ne_type**.

Syntactic rules are specified in Prolog with the predicate *rule(LHS,RHS)* where *LHS* is a syntactic category and *RHS* is a list of syntactic categories. A rule such as *BNP_HEAD* \Rightarrow *N* (‘a basic noun phrase head is composed of a noun’) is written as follows:

```
rule(bnp_head(sem:E^[[R,E],[number,E,N]],number:N),
[n(m_root:R,number:N)]).
```

where the feature ‘sem’ is used to construct the semantics while the parser processes input, and E, R, and N are variables to be instantiated during parsing.

The full grammar of this distribution can be found in the `prolog/grammar` directory, the file `load.pl` specifies which grammars are used by the parser. The grammars are compiled when the system is built and the compiled version is used for parsing.

17.11.8 Mapping Named Entities

SUPPLE has a prolog grammar which deals with named entities, the only information required is the Lookup annotations produced by Gate, which are specified in the mapping file. However, you may want to pass named entities identified with your own Jape grammars in GATE. This can be done using a special syntactic category provided with this distribution. The category `sem_cat` is used as a bridge between Gate named entities and the SUPPLE grammar. An example of how to use it (provided in the mapping file) is:

```
Gate;AnnotationType=Date;string=&S
SUPPLE;category=sem_cat;type=Date;text=&S;kind=date;name=&S
```

which maps a named entity 'Date' into a syntactic category 'sem_cat'. A grammar file called `semantic_rules.pl` is provided to map `sem_cat` into the appropriate syntactic category expected by the phrasal rules. The following rule for example:

```
rule(ne_np(s_form:F,sem:X^[[name,X,NAME],[KIND,X]]),[
sem_cat(s_form:F,text:TEXT,type:'Date',kind:KIND,name:NAME)]).
```

is used to parse a 'Date' into a named entity in SUPPLE which in turn will be parsed into a noun phrase.

17.11.9 Upgrading from BuChart to SUPPLE

In theory upgrading from BuChart to SUPPLE should be relatively straightforward. Basically any instance of BuChart needs to be replaced by SUPPLE. Specific changes which must be made are:

- The compiled parser files are now `supple.swi`, `supple.sicstus`, or `supple.plcafe`
- The GATE wrapper parameter `buchartFile` is now `SUPPLEFile`, and it is now of type `java.net.URL` rather than `java.io.File`. Details of how to compensate for this in existing saved applications are given in the SUPPLE README file.
- The Prolog wrappers now start `shef.nlp.supple.prolog` instead of `shef.nlp.buchart.prolog`
- The `mapping.conf` file now has lines starting `SUPPLE;` instead of `Buchart;`

- Most importantly the main wrapper class is now called `nlp.shef.supple.SUPPLE`

Making these changes to existing code should be trivial and allow application to benefit from future improvements to SUPPLE.

17.12 Stanford Parser

The Stanford Parser is a probabilistic parsing system implemented in Java by Stanford University's Natural Language Processing Group. Data files are available from Stanford for parsing Arabic, Chinese, English, and German.

This plugin, 'Parser.Stanford', developed by the GATE team, provides a PR (`gate.stanford.Parser`) that acts as a wrapper around the Stanford Parser (version 1.6.1) and translates GATE annotations to and from the data structures of the parser itself. The plugin is supplied with the unmodified `jar` file and one English data file obtained from Stanford. Stanford's software itself is subject to the full GPL.

The parser itself can be trained on other corpora and languages, as documented on the website, but this plugin does not provide a means of doing so. Trained data files are not compatible between different versions of the parser; in particular, note that you need version 1.6.1 data files for GATE builds numbered above 3120 (when we upgraded the plugin to Stanford version 1.6.1 on 22 January 2009) but version 1.6 files for earlier versions, including *Release 5.0 beta 1*.

Creating multiple instances of this PR in the same JVM with different trained data files does not work—the PRs can be instantiated, but runtime errors will almost certainly occur.

17.12.1 Input Requirements

Documents to be processed by the Parser PR must already have **Sentence** and **Token** annotations, such as those produced by either ANNIE Sentence Splitter (Sections 6.4 and 6.5) and the ANNIE English Tokeniser (Section 6.2).

If the `reusePosTags` parameter is true, then the **Token** annotations must have `category` features with compatible POS tags. The tags produced by the ANNIE POS Tagger are compatible with Stanford's parser data files for English (which also use the Penn treebank tagset).

17.12.2 Initialization Parameters

parserFile the path to the trained data file; the default value points to the English data file² included with the GATE distribution. You can also use other files downloaded from the Stanford Parser website or produced by training the parser.

mappingFile the optional path to a mapping file: a flat, two-column file which the wrapper can use to ‘translate’ tags. A sample file is included.³ By default this value is `null` and mapping is ignored.

tlppClass an implementation of `TreebankLangParserParams`, used by the parser itself to extract the dependency relations from the constituency structures. The default value is compatible with the English data file supplied. Please refer to the Stanford NLP Group’s documentation and the parser’s javadoc for a further explanation.

17.12.3 Runtime Parameters

annotationSetName the name of the `annotationSet` used for input (`Token` and `Sentence` annotations) and output (`SyntaxTreeNode` and `Dependency` annotations, and `category` and `dependencies` features added to `Tokens`).

debug a boolean value which controls the verbosity of the wrapper’s output.

reusePosTags if true, the wrapper will read `category` features (produced by an earlier POS-tagging PR) from the `Token` annotations and force the parser to use them.

useMapping if this is true and a mapping file was loaded when the PR was initialized, the POS and syntactic tags produced by the parser will be translated using that file. If no mapping file was loaded, this parameter is ignored.

The following boolean parameters switch on and off the various types of output that the parser can produce. Any or all of them can be true, but if all are false the PR will simply print a warning to save time (instead of running the parser).

addPosTags if this is true, the wrapper will add `category` features to the `Token` annotations.

addConstituentAnnotations if true, the wrapper will mark the syntactic constituents with `SyntaxTreeNode` annotations that are compatible with the Syntax Tree Viewer (see Section 17.11.4).

addDependencyAnnotations if true, the wrapper will add `Dependency` annotations to indicate the dependency relations in the sentence.

²`resources/englishPCFG.ser.gz`

³`resources/english-tag-map.txt`

addDependencyFeatures if true, the wrapper will add **dependencies** features to the **Token** annotations to indicate the dependency relations in the sentence.

The parser will derive the dependency structures only if either or both of the dependency output options is enabled, so if you do not need the dependency analysis, you can disable both of them and the PR will run faster.

Two sample GATE applications for English are included in the `plugins/Parser.Stanford` directory: `sample_parser_en.gapp` runs the Regex Sentence Splitter and ANNIE Tokenizer and then this PR to annotate constituency and dependency structures, whereas `sample_pos+parser_en.gapp` also runs the ANNIE POS Tagger and makes the parser re-use its POS tags.

17.13 OpenCalais, LingPipe and OpenNLP

Further parsing and tagging functionality can be found in plugins described in Chapter 19: OpenCalais (Section 19.14), LingPipe (Section 19.15) and OpenNLP (Section 19.16).

Chapter 18

Combining GATE and UIMA

UIMA (Unstructured Information Management Architecture) is a platform for natural language processing, originally developed by IBM but now maintained by the Apache Software Foundation. It has many similarities to the GATE architecture – it represents documents as text plus annotations, and allows users to define pipelines of *analysis engines* that manipulate the document (or *Common Analysis Structure* in UIMA terminology) in much the same way as processing resources do in GATE. The Apache UIMA SDK provides support for building analysis components in Java and C++ and running them either locally on one machine, or deploying them as services that can be accessed remotely. The SDK is available for download from <http://incubator.apache.org/uima/>.

Clearly, it would be useful to be able to include UIMA components in GATE applications and vice-versa, letting GATE users take advantage of UIMA's flexible deployment options and UIMA users access JAPE and the many useful plugins already available in GATE. This chapter describes the interoperability layer provided as part of GATE to support this. The UIMA-GATE interoperability layer is based on Apache UIMA 2.2.2. GATE 5.0 and earlier included an implementation based on version 1.2.3 of the pre-Apache IBM UIMA SDK.

The rest of this chapter assumes that you have at least a basic understanding of core UIMA concepts, such as *type systems*, *primitive* and *aggregate analysis engines* (AEs), *feature structures*, the format of AE XML descriptors, etc. It will probably be helpful to refer to the relevant sections of the UIMA SDK User's Guide and Reference (supplied with the SDK) alongside this document.

There are two main parts to the interoperability layer:

1. A wrapper to allow a UIMA Analysis Engine (AE), whether primitive or aggregate, to be used within GATE as a Processing Resource (PR).
2. A wrapper to allow a GATE processing pipeline (specifically a `CorpusController`) to be used within UIMA as an AE.

The two components operate in very similar ways. Given a document in the source form (either a GATE Document or a UIMA CAS), a document in the target form is created with a copy of the source document's text. Some of the annotations from the source are transferred to the target, according to a mapping defined by the user, and the target component is then run. Finally, some of the annotations on the updated target document are then transferred back to the source, according to the user-defined mapping.

The rest of this document describes this process in more detail. Section 18.1 describes the GATE AE wrapper, and Section 18.2 describes the UIMA CorpusController wrapper.

18.1 Embedding a UIMA AE in GATE

Embedding a UIMA analysis engine in a GATE application is a two step process. First, you must construct a *mapping descriptor* XML file to define how to map annotations between the UIMA CAS and the GATE Document. This mapping file, along with the analysis engine descriptor, is used to instantiate an *AnalysisEnginePR* which calls the analysis engine on an appropriately initialized CAS. Examples of all the XML files discussed in this section are available in `examples/conf` under the UIMA plugin directory.

18.1.1 Mapping File Format

Figure 18.1 shows the structure of a mapping descriptor. The `inputs` section defines how annotations on the GATE document are transferred to the UIMA CAS. The `outputs` section defines how annotations which have been added, updated and removed by the AE are transferred back to the GATE document.

Input Definitions

Each input definition takes the following form:

```
<uimaAnnotation type="uima.Type" gateType="GATEType" indexed="true|false">
  <feature name="..." kind="string|int|float|fs">
    <!-- element defining the feature value goes here -->
  </feature>
  ...
</uimaAnnotation>
```

When a document is processed, this will create one UIMA annotation of type `uima.Type` in the CAS for each GATE annotation of type `GATEType` in the input annotation set, covering

```
<uimaGateMapping>
  <inputs>
    <uimaAnnotation type="..." gateType="..." indexed="true|false">
      <feature name="..." kind="string|int|float|fs">
        <!-- element defining the feature value goes here -->
      </feature>
      ...
    </uimaAnnotation>
  </inputs>

  <outputs>
    <added>
      <gateAnnotation type="..." uimaType="...">
        <feature name="...">
          <!-- element defining the feature value goes here -->
        </feature>
        ...
      </gateAnnotation>
    </added>

    <updated>
      ...
    </updated>

    <removed>
      ...
    </removed>
  </outputs>
</uimaGateMapping>
```

Figure 18.1: Structure of a mapping descriptor for an AE in GATE

the same offsets in the text. If `indexed` is `true`, GATE will keep a record of which GATE annotation gave rise to which UIMA annotation. If you wish to be able to track updates to this annotation's features and transfer the updated values back into GATE, you must specify `indexed="true"`. The `indexed` attribute defaults to `false` if omitted.

Each contained `feature` element will cause the corresponding feature to be set on the generated annotation. UIMA features can be string, integer or float valued, or can be a reference to another feature structure, and this must be specified in the `kind` attribute. The feature's value is specified using a nested element, but exactly how this value is handled is determined by the `kind`.

There are various options for setting feature values:

- `<string value="fixed string" />` The simplest case - a fixed Java String.
- `<docFeatureValue name="featureName" />` The value of the given named feature of the current GATE document.
- `<gateAnnotFeatureValue name="featureName" />` The value of a given feature on the current GATE annotation (i.e. the one on which the offsets of the UIMA annotation are based).
- `<featureStructure type="uima.fs.Type">...</featureStructure>` A feature structure of the given type. The `featureStructure` element can itself contain `feature` elements recursively.

The value is assigned to the feature according to the feature's `kind`:

string The value object's `toString()` method is called, and the resulting String is set as the string value of the feature.

int If the value object is a subclass of `java.lang.Number`, its `intValue()` method is called, and the result is set as the integer value of the feature. If the value object is not a `Number`, it is `toString()`ed, and the resulting String is parsed using `Integer.parseInt()`. If this succeeds, the integer result is used, if it fails the feature is set to zero.

float As for `int`, except that `Numbers` are converted by calling `floatValue()`, and non-`Numbers` are parsed using `Float.parseFloat()`.

fs The value object is assumed to be a `FeatureStructure`, and is used as-is. A `ClassCastException` will result if the value object is not a `FeatureStructure`.

In particular, `<featureStructure>` value elements should only be used with features of kind `fs`. While nothing will stop you using them with `string` features, the result will probably not be what you expected.

Output Definitions

The output definitions take a similar form. There are three groups:

added Annotations which have been added by the AE, and for which corresponding new annotations are to be created in the GATE document.

updated Annotations that were created by an input definition (with `indexed="true"`) whose feature values have been modified by the AE, and these values are to be transferred back to the original GATE annotations.

removed Annotations that were created by an input definition (with `indexed="true"`) which have been removed from the CAS¹ and whose source annotations are to be removed from the GATE document.

The definition elements for these three types all take the same form:

```
<gateAnnotation type="GATEType" uimaType="uima.Type">
  <feature name="featureName">
    <!-- element defining the feature value goes here -->
  </feature>
  ...
</gateAnnotation>
```

For **added** annotations, this has the mirror-image effect to the input definition – for each UIMA annotation of the given type, create a GATE annotation at the same offsets and set its feature values as specified by **feature** elements. For a **gateAnnotation** the **feature** elements do not have a **kind**, as features in GATE can have arbitrary Objects as values. The possible feature value elements for a **gateAnnotation** are:

- `<string value="fixed string" />` A fixed string, as before.
- `<uimaFSFeatureValue name="uima.Type:FeatureName" kind="string|int|float" />`
The value of the given feature of the current UIMA annotation. The feature name must be specified in fully-qualified form, including the type on which it is defined. The **kind** is used in a similar way as in input definitions:
 - string** The Java `String` object returned as the string value of the feature is used.
 - int** An `Integer` object is created from the integer value of the feature.
 - float** A `Float` object is created from the float value of the feature.

¹Strictly speaking, removed from the annotation index, as feature structures cannot be removed from the CAS entirely.

fs The UIMA `FeatureStructure` object is returned. Since `FeatureStructure` objects are not guaranteed to be valid once the CAS has been cleared, a downstream GATE component must extract the relevant information from the feature structure before the next document is processed. You have been warned.

Feature names in `uimaFSFeatureValue` must be qualified with their type name, as the feature may have been defined on a supertype of the feature's own type, rather than the type itself. For example, consider the following:

```
<gateAnnotation type="Entity" uimaType="com.example.Entity">
  <feature name="type">
    <uimaFSFeatureValue name="com.example.Entity:Type" kind="string" />
  </feature>
  <feature name="startOffset">
    <uimaFSFeatureValue name="uima.tcas.Annotation:begin" kind="int" />
  </feature>
</gateAnnotation>
```

For **updated** annotations, there must have been an input definition with `indexed="true"` with the same GATE and UIMA types. In this case, for each GATE annotation of the appropriate type, the UIMA annotation that was created from it is found in the CAS. The feature definitions are then used as in the **added** case, but here, the feature values are set on the *original* GATE annotation, rather than on a newly created annotation.

For **removed** annotations, the feature definitions are ignored, and the annotation is removed from GATE if the UIMA annotation which it gave rise to has been removed from the UIMA annotation index.

A Complete Example

Figure 18.2 shows a complete example mapping descriptor for a simple UIMA AE that takes tokens as input and adds a feature to each token giving the number of lower case letters in the token's string.² In this case the UIMA feature that holds the number of lower case letters is called `LowerCaseLetters`, but the GATE feature is called `numLower`. This demonstrates that the feature names do not need to agree, so long as a mapping between them can be defined.

²The Java code implementing this AE is in the `examples` directory of the UIMA plugin. The AE descriptor and mapping file are in `examples/conf`.

```

<uimaGateMapping>
  <inputs>
    <uimaAnnotation type="gate.uima.cas.Token" gateType="Token" indexed="true">
      <feature name="String" kind="string">
        <gateAnnotFeatureValue name="string" />
      </feature>
    </uimaAnnotation>
  </inputs>
  <outputs>
    <updated>
      <gateAnnotation type="Token" uimaType="gate.uima.cas.Token">
        <feature name="numLower">
          <uimaFSFeatureValue name="gate.uima.cas.Token:LowerCaseLetters"
            kind="int" />
        </feature>
      </gateAnnotation>
    </updated>
  </outputs>
</uimaGateMapping>

```

Figure 18.2: An example mapping descriptor

18.1.2 The UIMA Component Descriptor

As well as the mapping file, you must provide the UIMA component descriptor that defines how to access the AE that is to be called. This could be a primitive or aggregate analysis engine descriptor, or a URI specifier giving the location of a remote Vinci or SOAP service. It is up to the developer to ensure that the types and features used in the mapping descriptor are compatible with the type system and capabilities of the AE, or a runtime error is likely to occur.

18.1.3 Using the AnalysisEnginePR

To use a UIMA AE in GATE Developer, load the UIMA plugin and create a ‘UIMA Analysis Engine’ processing resource. If using the GATE Embedded, rather than GATE Developer, the class name is `gate.uima.AnalysisEnginePR`. The processing resource expects two parameters:

analysisEngineDescriptor The URL of the UIMA analysis engine descriptor (or URI specifier, for a remote AE service). This must be a `file:` URL, as UIMA needs a file path against which to resolve imports.

mappingDescriptor The URL of the mapping descriptor file. This may be any kind of URL (`file:`, `http:`, `Class.getResource()`, `ServletContext.getResource()`, etc.)

Any errors processing either of the descriptor files will cause an exception to be thrown. Once instantiated, you can add the PR to a pipeline in the usual way. `AnalysisEnginePR` implements `LanguageAnalyser`, so can be used in any of the standard GATE pipeline types.

The PR takes the following runtime parameter (in addition to the `document` parameter which is set automatically by a `CorpusController`):

annotationSetName The annotation set to process. Any input mappings take annotations from this set, and any output mappings place their new annotations in this set (`added` outputs) or update the input annotations in this set (`updated` or `removed`). If not specified, the default (unnamed) annotation set is used.

The Annotator implementation must be available for GATE to load. For an annotator written in Java, this means that the JAR file containing the annotator class (and any other classes it depends on) must be present in the GATE classloader. The easiest way to achieve this is to put the JAR file or files in a new directory, and create a `creole.xml` file in the same directory to reference the JARs:

```
<CREOLE-DIRECTORY>
  <JAR>my-annotator.jar</JAR>
  <JAR>classes-it-uses.jar</JAR>
</CREOLE-DIRECTORY>
```

This directory should then be loaded in GATE as a CREOLE plugin. Note that, due to the complex mechanics of classloaders in Java, putting your JARs in GATE's `lib` directory will *not* work.

For annotators written in C++ you need to ensure that the C++ enabler libraries (available separately from <http://incubator.apache.org/uima/>) and the shared library containing your annotator are in a directory which is on the `PATH` (Windows) or `LD_LIBRARY_PATH` (Linux) when GATE is run.

18.2 Embedding a GATE CorpusController in UIMA

The process of embedding a GATE controller in a UIMA application is more or less the mirror image of the process detailed in the previous section. Again, the developer must supply a mapping descriptor defining how to map between UIMA and GATE annotations, and pass this, plus the GATE controller definition, to an AE which performs the translation and calls the GATE controller.

18.2.1 Mapping File Format

The mapping descriptor format is virtually identical to that described in Section 18.1.1, except that the input definitions are `<gateAnnotation>` elements and the output definitions are `<uimaAnnotation>` elements. The input and output definition elements support an extra attribute, `annotationSetName`, which allows inputs to be taken from, and outputs to be placed in, different annotation sets. For example, the following hypothetical example maps `com.example.Person` annotations into the default set and `com.example.html.Anchor` annotations to ‘a’ tags in the ‘Original markups’ set.

```
<inputs>
  <gateAnnotation type="Person" uimaType="com.example.Person">
    <feature name="kind">
      <uimaFSFeatureValue name="com.example.Person:Kind" kind="string"/>
    </feature>
  </gateAnnotation>

  <gateAnnotation type="a" annotationSetName="Original markups"
    uimaType="com.example.html.Anchor">
    <feature name="href">
      <uimaFSFeatureValue name="com.example.html.Anchor:href" kind="string" />
    </feature>
  </gateAnnotation>
</inputs>
```

Figure 18.3 shows a mapping descriptor for an application that takes tokens and sentences produced by some UIMA component and runs the GATE part of speech tagger to tag them with Penn TreeBank POS tags.³ In the example, no features are copied from the UIMA tokens, but they are still `indexed="true"` as the POS feature must be copied back from GATE.

18.2.2 The GATE Application Definition

The GATE application to embed is given as a standard ‘.gapp file’, as produced by saving the state of an application in the GATE GUI. The .gapp file encodes the information necessary to load the correct plugins and create the various CREOLE components that make up the application. The .gapp file must be fully specified and able to be executed with no user intervention other than pressing the Go button. In particular, all runtime parameters must be set to their correct values before saving the application state. Also, since paths to things like CREOLE plugin directories, resource files, etc. are stored relative to the .gapp file’s

³The .gapp file implementing this example is in the `test/conf` directory under the UIMA plugin, along with the mapping file and the AE descriptor that will run it.

```

<uimaGateMapping>
  <inputs>
    <gateAnnotation type="Token"
                    uimaType="com.ibm.uima.examples.tokenizer.Token"
                    indexed="true" />
    <gateAnnotation type="Sentence"
                    uimaType="com.ibm.uima.examples.tokenizer.Sentence" />
  </inputs>
  <outputs>
    <updated>
      <uimaAnnotation type="com.ibm.uima.examples.tokenizer.Token"
                      gateType="Token">
        <feature name="POS" kind="string">
          <gateAnnotFeatureValue name="category" />
        </feature>
      </uimaAnnotation>
    </updated>
  </outputs>
</uimaGateMapping>

```

Figure 18.3: An example mapping descriptor for the GATE POS tagger

location, you must not move the `.gapp` file to a different directory unless you can keep all the CREOLE plugins it depends on at the same relative locations. The ‘Export for Teamware’ option (section 3.8.4) may help you here.

18.2.3 Configuring the GATEApplicationAnnotator

`GATEApplicationAnnotator` is the UIMA annotator that handles mapping the CAS into a GATE document and back again and calling the GATE controller. There is a template AE descriptor XML file for the annotator provided in the `conf` directory. Most of the template file can be used unchanged, but you will need to modify the type system definition and input/output capabilities to match the types and features used in your mapping descriptor. If the mapping descriptor references a type or feature that is not defined in the type system, a runtime error will occur.

The annotator requires two external resources:

GateApplication The `.gapp` file containing the saved application state.

MappingDescriptor The mapping descriptor XML file.

These must be bound to suitable URLs, either by editing the `resourceManagerConfiguration`

section of the primitive descriptor, or by supplying the binding in an aggregate descriptor that includes the `GATEApplicationAnnotator` as one of its delegates.

In addition, you may need to set the following Java system properties:

uima.gate.configdir The path to the GATE config directory. This defaults to `gate-config` in the same directory as `uima-gate.jar`.

uima.gate.siteconfig The location of the sitewide `gate.xml` configuration file. This defaults to `gate.uima.configdir/site-gate.xml`.

uima.gate.userconfig The location of the user-specific `gate.xml` configuration file. This defaults to `gate.uima.configdir/user-gate.xml`.

The default config files are deliberately simplified from the standard versions supplied with GATE, in particular they do not load any plugins automatically (not even ANNIE). All the plugins used by your application are specified in the `.gapp` file, and will be loaded when the application is loaded, so it is best to avoid loading any others from `gate.xml`, to avoid problems such as two different versions of the same plugin being loaded from different locations.

Classpath Notes

In addition to the usual UIMA library JAR files, `GATEApplicationAnnotator` requires a number of JAR files from the GATE distribution in order to function. In the first instance, you should include `gate.jar` from GATE's `bin` directory, and also all the JAR files from GATE's `lib` directory on the classpath. If you use the supplied Ant build file, `ant documentanalyser` will run the document analyser with this classpath. Depending on exactly which GATE plugins your application uses, you may be able to exclude some of the `lib` JAR files (for example, you will not need Weka if you do not use the machine learning plugin), but it is safest to start with them all. GATE will load plugin JAR files through its own classloader, so these do not need to be on the classpath.

Chapter 19

More (CREOLE) Plugins

For the previous reader was none other than myself. I had already read this book long ago.

The old sickness has me in its grip again: amnesia in litteris, the total loss of literary memory. I am overcome by a wave of resignation at the vanity of all striving for knowledge, all striving of any kind. Why read at all? Why read this book a second time, since I know that very soon not even a shadow of a recollection will remain of it? Why do anything at all, when all things fall apart? Why live, when one must die? And I clap the lovely book shut, stand up, and slink back, vanquished, demolished, to place it again among the mass of anonymous and forgotten volumes lined up on the shelf.

...

But perhaps - I think, to console myself - perhaps reading (like life) is not a matter of being shunted on to some track or abruptly off it. Maybe reading is an act by which consciousness is changed in such an imperceptible manner that the reader is not even aware of it. The reader suffering from amnesia in litteris is most definitely changed by his reading, but without noticing it, because as he reads, those critical faculties of his brain that could tell him that change is occurring are changing as well. And for one who is himself a writer, the sickness may conceivably be a blessing, indeed a necessary precondition, since it protects him against that crippling awe which every great work of literature creates, and because it allows him to sustain a wholly uncomplicated relationship to plagiarism, without which nothing original can be created.

Three Stories and a Reflection, Patrick Suskind, 1995 (pp. 82, 86).

This chapter describes additional CREOLE resources which do not form part of ANNIE, and have not been covered in previous chapters.

19.1 Language Plugins

There are plugins available for processing the following languages: French, German, Spanish, Italian, Chinese, Arabic, Romanian, Hindi and Cebuano. Some of the applications are quite basic and just contain some useful processing resources to get you started when developing a full application. Others (Cebuano and Hindi) are more like toy systems built as part of an exercise in language portability.

Note that if you wish to use individual language processing resources without loading the whole application, you will need to load the relevant plugin for that language in most cases. The plugins all follow the same kind of format. Load the plugin using the plugin manager in GATE Developer, and the relevant resources will be available in the Processing Resources set.

Some plugins just contain a list of resources which can be added ad hoc to other applications. For example, the Italian plugin simply contains a lexicon which can be used to replace the English lexicon in the default English POS tagger: this will provide a reasonable basic POS tagger for Italian.

In most cases you will also find a directory in the relevant plugin directory called data which contains some sample texts (in some cases, these are annotated with NEs).

19.1.1 French Plugin

The French plugin contains two applications for NE recognition: one which includes the TreeTagger for POS tagging in French (`french+tagger.gapp`), and one which does not (`french.gapp`). Simply load the application required from the `plugins/Lang_French` directory. You do not need to load the plugin itself from the GATE Developer's Plugin Management Console. Note that the TreeTagger must first be installed and set up correctly (see Section 17.3 for details). Check that the runtime parameters are set correctly for your TreeTagger in your application. The applications both contain resources for tokenisation, sentence splitting, gazetteer lookup, NE recognition (via JAPE grammars) and orthographic coreference. Note that they are not intended to produce high quality results, they are simply a starting point for a developer working on French. Some sample texts are contained in the `plugins/Lang_French/data` directory.

19.1.2 German Plugin

The German plugin contains two applications for NE recognition: one which includes the TreeTagger for POS tagging in German (`german+tagger.gapp`), and one which does not (`german.gapp`). Simply load the application required from the `plugins/Lang_German/resources` directory. You do not need to load the plugin itself from the GATE Developer's Plugin

Management Console. Note that the TreeTagger must first be installed and set up correctly (see Section 17.3 for details). Check that the runtime parameters are set correctly for your TreeTagger in your application. The applications both contain resources for tokenisation, sentence splitting, gazetteer lookup, compound analysis, NE recognition (via JAPE grammars) and orthographic coreference. Some sample texts are contained in the plugins/Lang_German/data directory. We are grateful to Fabio Ciravegna and the Dot.KOM project for use of some of the components for the German plugin.

19.1.3 Romanian Plugin

The Romanian plugin contains an application for Romanian NE recognition (`romanian.gapp`). Simply load the application from the `plugins/Lang_Romanian/resources` directory. You do not need to load the plugin itself from the GATE Developer's Plugin Management Console. The application contains resources for tokenisation, gazetteer lookup, NE recognition (via JAPE grammars) and orthographic coreference. Some sample texts are contained in the `plugins/romanian/corpus` directory.

19.1.4 Arabic Plugin

The Arabic plugin contains a simple application for Arabic NE recognition (`arabic.gapp`). Simply load the application from the `plugins/Lang_Arabic/resources` directory. You do not need to load the plugin itself from the GATE Developer's Plugin Management Console. The application contains resources for tokenisation, gazetteer lookup, NE recognition (via JAPE grammars) and orthographic coreference. Note that there are two types of gazetteer used in this application: one which was derived automatically from training data (Arabic inferred gazetteer), and one which was created manually. Note that there are some other applications included which perform quite specific tasks (but can generally be ignored). For example, `arabic-for-bbn.gapp` and `arabic-for-muse.gapp` make use of a very specific set of training data and convert the result to a special format. There is also an application to collect new gazetteer lists from training data (`arabic_lists_collector.gapp`). For details of the gazetteer list collector please see Section 13.8.

19.1.5 Chinese Plugin

The Chinese plugin contains two components: a simple application for Chinese NE recognition (`chinese.gapp`) and a component called "Chinese Segmenter".

In order to use the former, simply load the application from the `plugins/Lang_Chinese/resources` directory. You do not need to load the plugin itself from the GATE Developer's Plugin Management Console. The application contains resources for tokenisation, gazetteer lookup, NE

recognition (via JAPE grammars) and orthographic coreference. The application makes use of some gazetteer lists (and a grammar to process them) derived automatically from training data, as well as regular hand-crafted gazetteer lists. There are also applications (`listscollector.gapp`, `adj_collector.gapp` and `nounperson_collector.gapp`) to create such lists, and various other application to perform special tasks such as coreference evaluation (`coreference_eval.gapp`) and converting the output to a different format (`ace-to-muse.gapp`).

For details on the Chinese Segmenter please see Section 19.12.

19.1.6 Hindi Plugin

The Hindi plugin (`'Lang_Hindi'`) contains a set of resources for basic Hindi NE recognition which mirror the ANNIE resources but are customised to the Hindi language. You need to have the ANNIE plugin loaded first in order to load any of these PRs. With the Hindi, you can create an application similar to ANNIE but replacing the ANNIE PRs with the default PRs from the plugin.

19.2 Flexible Exporter

The Flexible Exporter enables the user to save a document (or corpus) in its original format with added annotations. The user can select the name of the annotation set from which these annotations are to be found, which annotations from this set are to be included, whether features are to be included, and various renaming options such as renaming the annotations and the file.

At load time, the following parameters can be set for the flexible exporter:

- `includeFeatures` - if set to true, features are included with the annotations exported; if false (the default status), they are not.
- `useSuffixForDumpFiles` - if set to true (the default status), the output files have the suffix defined in `suffixForDumpFiles`; if false, no suffix is defined, and the output file simply overwrites the existing file (but see the `outputFileUrl` runtime parameter for an alternative).
- `suffixForDumpFiles` - this defines the suffix if `useSuffixForDumpFiles` is set to true. By default the suffix is `.gate`.
- `useStandOffXML` - if true then the format will be the GATE XML format that separates nodes and annotations inside the file which allows overlapping annotations to be saved.

The following runtime parameters can also be set (after the file has been selected for the application):

- `annotationSetName` - this enables the user to specify the name of the annotation set which contains the annotations to be exported. If no annotation set is defined, it will use the Default annotation set.
- `annotationTypes` - this contains a list of the annotations to be exported. By default it is set to Person, Location and Date.
- `dumpTypes` - this contains a list of names for the exported annotations. If the annotation name is to remain the same, this list should be identical to the list in `annotationTypes`. The list of annotation names must be in the same order as the corresponding annotation types in `annotationTypes`.
- `outputDirectoryUrl` - this enables the user to specify the export directory where the file is exported with its original name and an extension (provided as a parameter) appended at the end of filename. Note that you can also save a whole corpus in one go. If not provided, use the temporary directory.

19.3 Annotation Set Transfer

The Annotation Set Transfer allows copying or moving annotations to a new annotation set if they lie between the beginning and the end of an annotation of a particular type (the covering annotation). For example, this can be used when a user only wants to run a processing resource over a specific part of a document, such as the Body of an HTML document. The user specifies the name of the annotation set and the annotation which covers the part of the document they wish to transfer, and the name of the new annotation set. All the other annotations corresponding to the matched text will be transferred to the new annotation set. For example, we might wish to perform named entity recognition on the body of an HTML text, but not on the headers. After tokenising and performing gazetteer lookup on the whole text, we would use the Annotation Set Transfer to transfer those annotations (created by the tokeniser and gazetteer) into a new annotation set, and then run the remaining NE resources, such as the semantic tagger and coreference modules, on them.

The Annotation Set Transfer has no loadtime parameters. It has the following runtime parameters:

- `inputASName` - this defines the annotation set from which annotations will be transferred (copied or moved). If nothing is specified, the Default annotation set will be used.

- **outputASName** - this defines the annotation set to which the annotations will be transferred. This default value for this parameter is 'Filtered'. If it is left blank the Default annotation set will be used.
- **tagASName** - this defines the annotation set which contains the annotation covering the relevant part of the document to be transferred. This default value for this parameter is 'Original markups'. If it is left blank the Default annotation set will be used.
- **textTagName** - this defines the type of the annotation covering the annotations to be transferred. The default value for this parameter is 'BODY'. If this is left blank, then all annotations from the inputASName annotation set will be transferred. If more than one covering annotation is found, the annotation covered by each of them will be transferred. If no covering annotation is found, the processing depends on the copyAllUnlessFound parameter (see below).
- **copyAnnotations** - this specifies whether the annotations should be moved or copied. The default value **false** will move annotations, removing them from the inputASName annotation set. If set to **true** the annotations will be copied.
- **transferAllUnlessFound** - this specifies what should happen if no covering annotation is found. The default value is **true**. In this case, all annotations will be copied or moved (depending on the setting of parameter **copyAnnotations**) if no covering annotation is found. If set to **false**, no annotation will be copied or moved.

For example, suppose we wish to perform named entity recognition on only the text covered by the BODY annotation from the Original Markups annotation set in an HTML document. We have to run the gazetteer and tokeniser on the entire document, because since these resources do not depend on any other annotations, we cannot specify an input annotation set for them to use. We therefore transfer these annotations to a new annotation set (Filtered) and then perform the NE recognition over these annotations, by specifying this annotation set as the input annotation set for all the following resources. In this example, we would set the following parameters (assuming that the annotations from the tokenise and gazetteer are initially placed in the Default annotation set).

- **inputASName**: Default
- **outputASName**: Filtered
- **tagASName**: Original markups
- **textTagName**: BODY
- **copyAnnotations**: true or false (depending on whether we want to keep the Token and Lookup annotations in the Default annotation set)
- **copyAllUnlessFound**: true

The AST PR makes a shallow copy of the feature map for each transferred annotation, i.e. it creates a new feature map containing the same keys and values as the original. It does *not* clone the feature values themselves, so if your annotations have a feature whose value is a collection and you need to make a deep copy of the collection value then you will not be able to use the AST PR to do this. Similarly if you are copying annotations and *do* in fact want to share the same feature map between the source and target annotations then the AST PR is not appropriate. In these sorts of cases a JAPE grammar or Groovy script would be a better choice.

19.4 Information Retrieval in GATE

GATE comes with a full-featured Information Retrieval (IR) subsystem that allows queries to be performed against GATE corpora. This combination of IE and IR means that documents can be retrieved from the corpora not only based on their textual content but also according to their features or annotations. For example, a search over the Person annotations for ‘Bush’ will return documents with higher relevance, compared to a search in the content for the string ‘bush’. The current implementation is based on the most popular open source full-text search engine - Lucene (available at <http://jakarta.apache.org/lucene/>) but other implementations may be added in the future.

An Information Retrieval system is most often considered a system that accepts as input a set of documents (corpus) and a query (combination of search terms) and returns as input only those documents from the corpus which are considered as relevant according to the query. Usually, in addition to the documents, a proper relevance measure (score) is returned for each document. There exist many relevance metrics, but usually documents which are considered more relevant, according to the query, are scored higher.

Figure 19.1 shows the results from running a query against an indexed corpus in GATE.

	$term_1$	$term_2$	$term_k$
doc_1	$w_{1,1}$	$w_{1,2}$	$w_{1,k}$
doc_2	$w_{2,1}$	$w_{2,2}$	$w_{2,k}$
...
...
doc_n	$w_{n,1}$	$w_{n,2}$	$w_{n,k}$

Table 19.1: An information retrieval document-term matrix




Figure 19.1: Documents with scores, returned from a search over a corpus

Information Retrieval systems usually perform some preprocessing on the input corpus in order to create the document-term matrix for the corpus. A document-term matrix is usually presented as in Table 19.1, where doc_i is a document from the corpus, $term_j$ is a word that is considered as important and representative for the document and $w_{i,j}$ is the weight assigned to the term in the document. There are many ways to define the term weight functions, but most often it depends on the term frequency in the document and in the whole corpus (i.e. the local and the global frequency). Note that the machine learning plugin described in Chapter 15 can produce such document-term matrix (for detailed description of the matrix produced, see Section 15.2.4).

Note that not all of the words appearing in the document are considered terms. There are many words (called ‘stop-words’) which are ignored, since they are observed too often and are not representative enough. Such words are articles, conjunctions, etc. During the preprocessing phase which identifies such words, usually a form of stemming is performed in order to minimize the number of terms and to improve the retrieval recall. Various forms of the same word (e.g. ‘play’, ‘playing’ and ‘played’) are considered identical and multiple

occurrences of the same term (probably ‘play’) will be observed.

It is recommended that the user reads the relevant Information Retrieval literature for a detailed explanation of stop words, stemming and term weighting.

IR systems, in a way similar to IE systems, are evaluated with the help of the precision and recall measures (see Section 10.1 for more details).

19.4.1 Using the IR Functionality in GATE

In order to run queries against a corpus, the latter should be ‘indexed’. The indexing process first processes the documents in order to identify the terms and their weights (stemming is performed too) and then creates the proper structures on the local file system. These file structures contain indexes that will be used by Lucene (the underlying IR engine) for the retrieval.

Once the corpus is indexed, queries may be run against it. Subsequently the index may be removed and then the structures on the local file system are removed too. Once the index is removed, queries cannot be run against the corpus.

Indexing the Corpus

In order to index a corpus, the latter should be stored in a serial datastore. In other words, the IR functionality is unavailable for corpora that are transient or stored in a RDBMS datastores (though support for the latter may be added in the future).

To index the corpus, follow these steps:

- Select the corpus from the resource tree (top-left pane) and from the context menu (right button click) choose ‘Index Corpus’. A dialogue appears that allows you to specify the index properties.
- In the index properties dialogue, specify the underlying IR system to be used (only Lucene is supported at present), the directory that will contain the index structures, and the set of properties that will be indexed such as document features, content, etc (the same properties will be indexed for each document in the corpus).
- Once the corpus is indexed, you may start running queries against it. Note that the directory specified for the index data should exist and be empty. Otherwise an error will occur during the index creation.




Figure 19.2: Indexing a corpus by specifying the index location and indexed features (and content)

Querying the Corpus

To query the corpus, follow these steps:

- Create a SearchPR processing resource. All the parameters of SearchPR are runtime so they are set later.
- Create a pipeline application containing the SearchPR.
- Set the following SearchPR parameters:
 - The corpus that will be queried.
 - The query that will be executed.
 - The maximum number of documents returned.

A query looks like the following:

```
{+/-}field1:term1 {+/-}field2:term2 ? {+/-}fieldN:termN
```

where `field` is the name of a index field, such as the one specified at index creation (the document content field is `body`) and `term` is a term that should appear in the field.

For example the query:

```
+body:government +author:CNN
```

will inspect the document content for the term ‘government’ (together with variations such as ‘governments’ etc.) and the index field named ‘author’ for the term ‘CNN’. The ‘author’ field is specified at index creation time, and is either a document feature or another document property.

- After the SearchPR is initialized, running the application executes the specified query over the specified corpus.
- Finally, the results are displayed (see fig.1) after a double-click on the SearchPR processing resource.

Removing the Index

An index for a corpus may be removed at any time from the ‘Remove Index’ option of the context menu for the indexed corpus (right button click).

19.4.2 Using the IR API

The IR API within GATE Embedded makes it possible for corpora to be indexed, queried and results returned from any Java application, without using GATE Developer. The following sample indexes a corpus, runs a query against it and then removes the index.

```

1
2 // open a serial datastore
3 SerialDataStore sds =
4 Factory.openDataStore("gate.persist.SerialDataStore",
5 "/tmp/datastore1");
6 sds.open();
7
8 //set an AUTHOR feature for the test document
9 Document doc0 = Factory.newDocument(new URL("/tmp/documents/doc0.html"));
10 doc0.getFeatures().put("author", "John Smith");
11
12 Corpus corp0 = Factory.newCorpus("TestCorpus");
13 corp0.add(doc0);
14
15 //store the corpus in the serial datastore
16 Corpus serialCorpus = (Corpus) sds.adopt(corp0, null);
17 sds.sync(serialCorpus);
18
19 //index the corpus – the content and the AUTHOR feature
20
21 IndexedCorpus indexedCorpus = (IndexedCorpus) serialCorpus;
22
23 DefaultIndexDefinition did = new DefaultIndexDefinition();

```

```

24 did.setIrEngineClassName(
25     gate.creole.ir.lucene.LuceneIREngine.class.getName());
26 did.setIndexLocation("/tmp/index1");
27 did.addField(new IndexField("content",
28     new DocumentContentReader(), false));
29 did.addField(new IndexField("author", null, false));
30 indexedCorpus.setIndexDefinition(did);
31
32 indexedCorpus.getIndexManager().createIndex();
33 //the corpus is now indexed
34
35 //search the corpus
36 Search search = new LuceneSearch();
37 search.setCorpus(ic);
38
39 QueryResultList res = search.search("+content:government +author:John");
40
41 //get the results
42 Iterator it = res.getQueryResults();
43 while (it.hasNext()) {
44     QueryResult qr = (QueryResult) it.next();
45     System.out.println("DOCUMENT_ID=" + qr.getDocumentID()
46         + ", score=" + qr.getScore());
47 }

```

19.5 Websphinx Web Crawler

The plugin ‘Web_Crawler_Websphinx’ enables GATE to build a corpus from a web crawl. The crawler itself is Websphinx. This is a JAVA based multi-threaded web crawler that can be customized for any application.

N.B. If you are using this plugin via an IDE, you may need to make sure that the websphinx.jar file is on the IDE’s classpath, or add to the IDE’s lib directory.

The basic idea is to be able to specify a source URL and a depth to build the initial corpus upon which further processing could be done. The PR itself provides a number of helpful features to set various parameters of the crawl.

19.5.1 Using the Crawler PR

In order to use the processing resource you first need to load the plugin using the plugin manager. Then load the crawler from the list of processing resources. User needs to create a corpus in which he or she wants to store crawled documents. In order to use the crawler, create a simple pipeline (note: do not create a corpus pipeline) and add the crawl PR to the pipeline.

Once the crawl PR is created there will be a number of parameters that can be set based on the PR required (see also Figure 19.3).

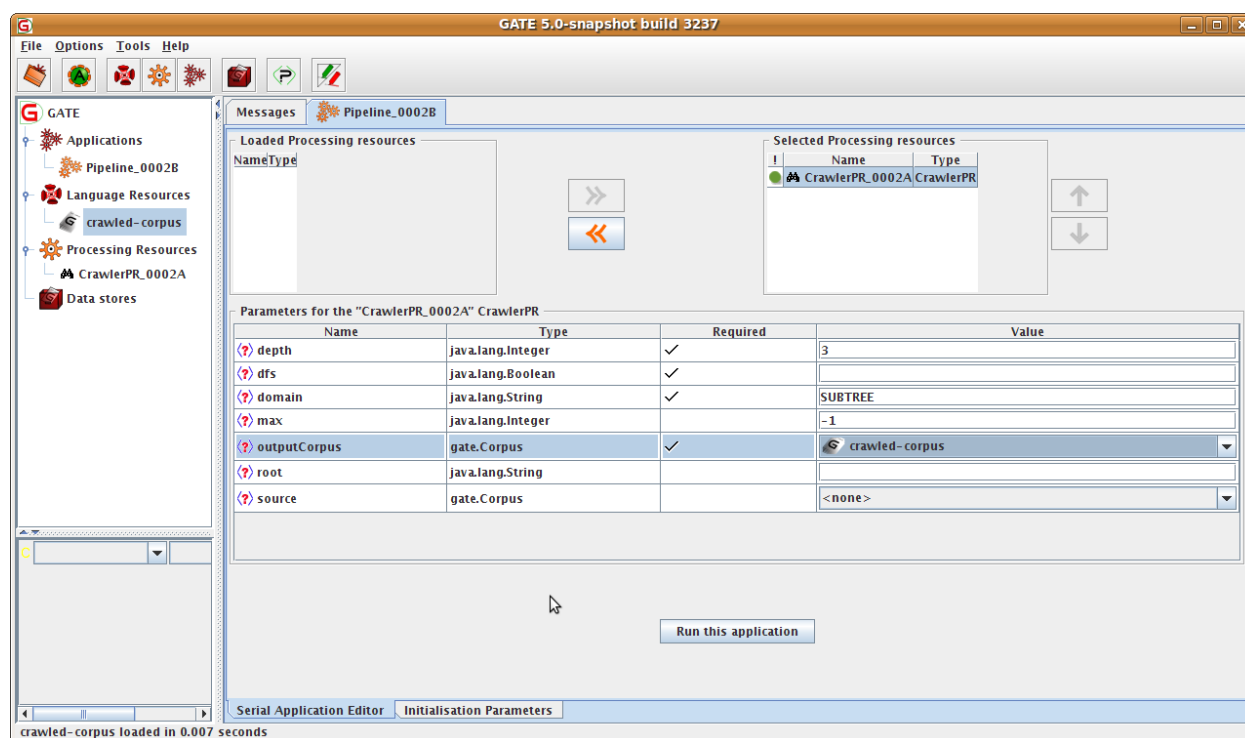


Figure 19.3: Crawler parameters

- depth: the depth to which the crawl should proceed.
- dfs / bfs: dfs if true bfs if false
 - Dfs : the crawler uses the depth first strategy for the crawl.
 - * Visits the nodes in dfs order until the specified depth limit is reached.
 - Bfs: the crawler used the breadth first strategy for the crawl.
 - * Visits the nodes on bfs order until the specified depth limit is reached.
- domain
 - SUBTREE: Crawler visits only the descendents of the page specified as the root for the crawl.
 - WEB: Crawler visits all the pages on the web.
 - SERVER: Crawler visits only the pages that are present on the server where the root page is located.
- max number of pages to be fetched

- outputCorpus an instance of Corpus to be used for storing crawled web pages
- root the starting URL to be used for the crawl to begin
- source is the corpus to be used that contains the documents from which the crawl must begin. Source is useful when the documents are fetched first from the Google PR and then need to be crawled to expand the web graph further. At any time either the source or the root needs to be set.

Once the parameters are set, the crawl can be run and the documents fetched are added to the specified corpus. Figure 19.4 shows the crawled pages added to the corpus.

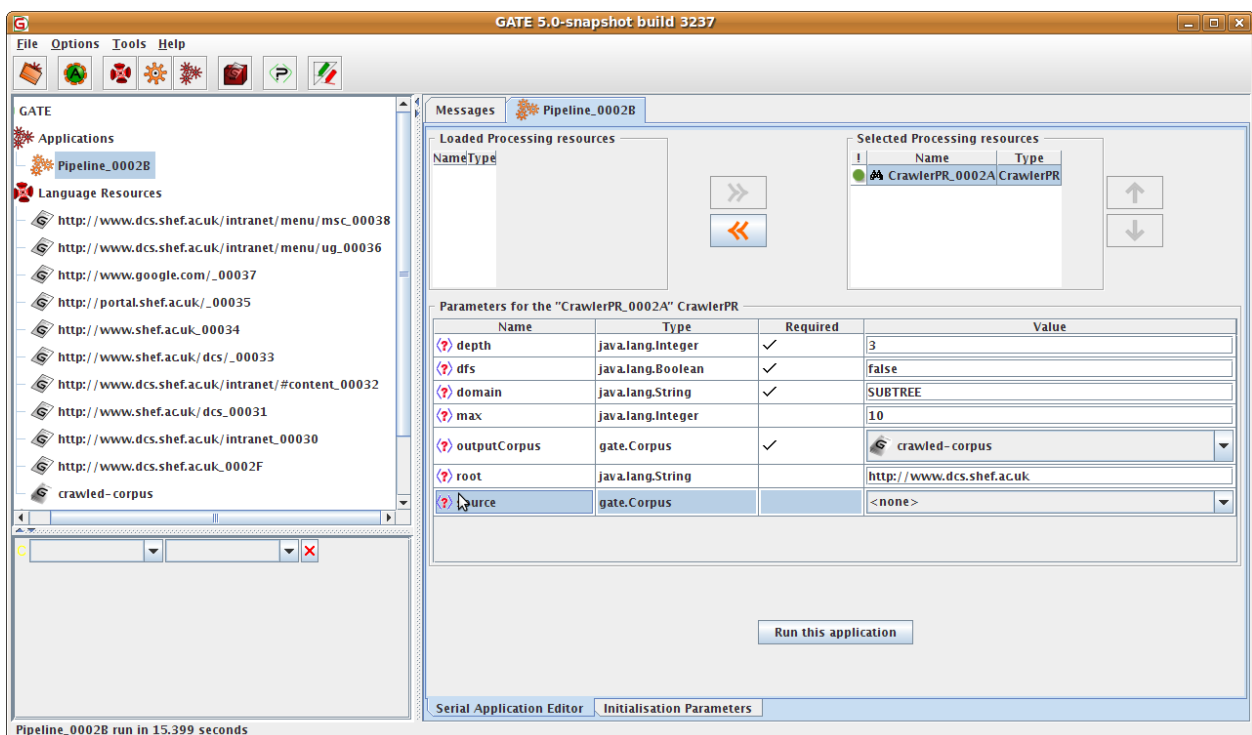


Figure 19.4: Crawled pages added to the corpus

N.B. Note that you must use a simple Pipeline, and not a Corpus Pipeline. If you wish to process the crawled documents, you must build a second Corpus Pipeline. Note that from GATE Version 5.1, you could combine the two pipelines as follows. Build a simple Pipeline containing your Web Crawler. Let's call this Pipeline P_s . Let's say you set the Corpus on your crawler to C . Now build your processing corpus pipeline, which we will call P_c . Put the original pipeline P_s as its first PR. Set its corpus to be the corpus as before, C .

19.6 Google Plugin

This plugin is no longer operational because the functionality, provided by Google, on which it depends, is no longer available.

19.7 Yahoo Plugin

The Yahoo API is now integrated with GATE, and can be used as a PR-based plugin. This plugin, 'Web_Search_Yahoo', allows the user to query Yahoo and build a document corpus that contains the search results returned by Yahoo for the query. For more information about the Yahoo API please refer to <http://developer.yahoo.com/search/>. In order to use the Yahoo PR, you need to obtain an application ID.

The Yahoo PR can be used for a number of different application scenarios. For example, one use case is where a user wants to find the different named entities that can be associated with a particular individual. In this example, the user could build a collection of documents by querying Yahoo with the individual's name and then running ANNIE over the collection. This would annotate the results and show the different Organization, Location and other entities that are associated with the query.

19.7.1 Using the YahooPR

In order to use the PR, you first need to load the plugin using the GATE Developer plugin manager. Once the PR is loaded, it can be initialized by creating an instance of a new PR. Here you need to specify the Yahoo Application ID. Please use the license key assigned to you by registering with Yahoo.

Once the Yahoo PR is initialized, it can be placed in a pipeline or a conditional pipeline application. This pipeline would contain the instance of the Yahoo PR just initialized as above. There are a number of parameters to be set at runtime:

- **corpus**: The corpus used by the plugin to add or append documents from the Web.
- **corpusAppendMode**: If set to `true`, will append documents to the corpus. If set to `false`, will remove preexisting documents from the corpus, before adding the documents newly fetched by the PR
- **limit**: A limit on the results returned by the search. Default set to 10.
- **pagesToExclude**: This is an optional parameter. It is a list with URLs not to be included in the search.

- query: The query sent to Yahoo. It is in the format accepted by Yahoo.

Once the required parameters are set we can run the pipeline. This will then download all the URLs in the results and create a document for each. These documents would be added to the corpus.

19.8 WordNet in GATE

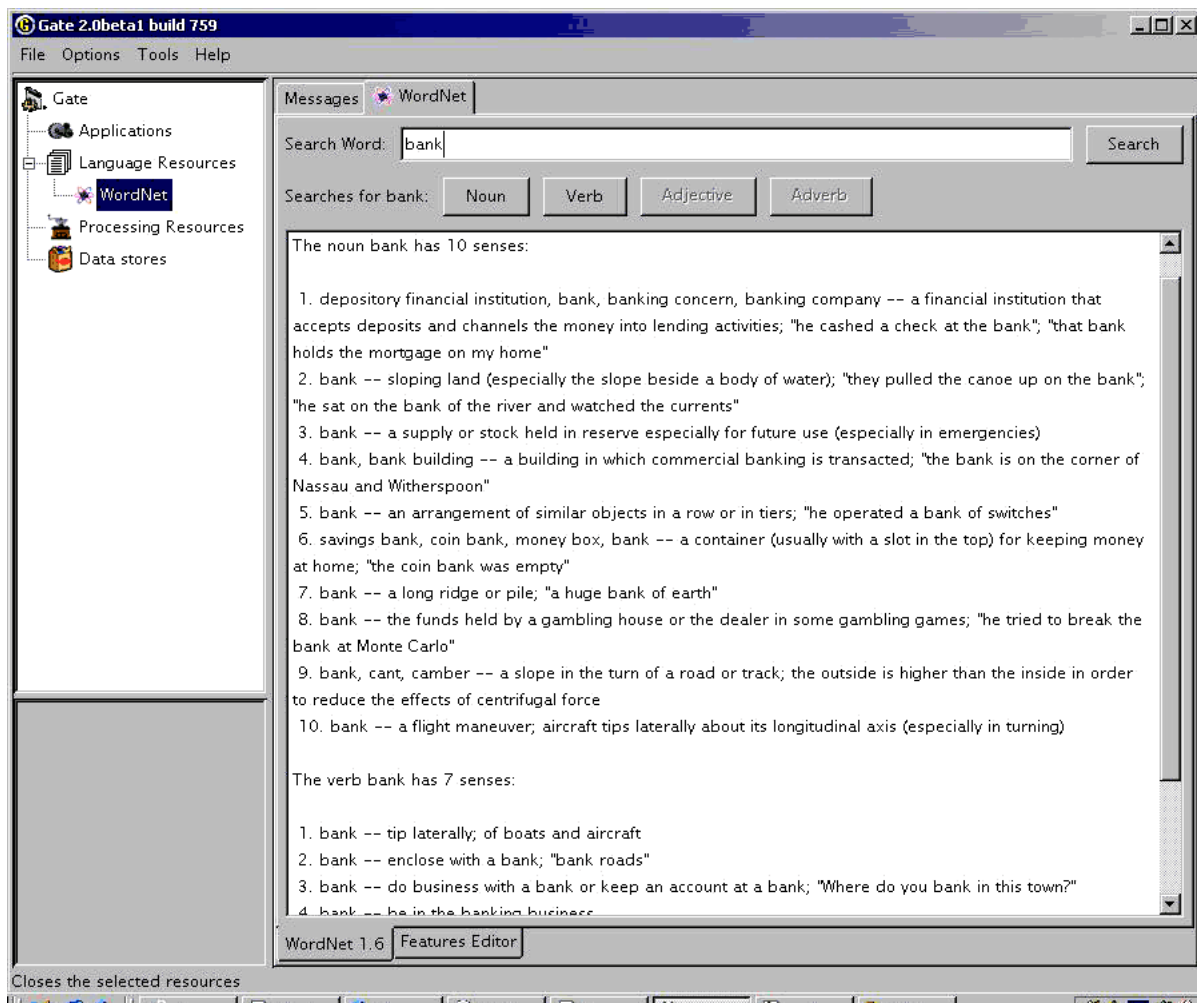


Figure 19.5: WordNet in GATE – results for ‘bank’

At present GATE supports only WordNet 1.6, so in order to use WordNet in GATE, you must first install WordNet 1.6 on your computer. WordNet is available at <http://wordnet.princeton.edu/>. The next step is to configure GATE to work with your local WordNet installation. Since GATE relies on the Java WordNet Library (JWNL) for

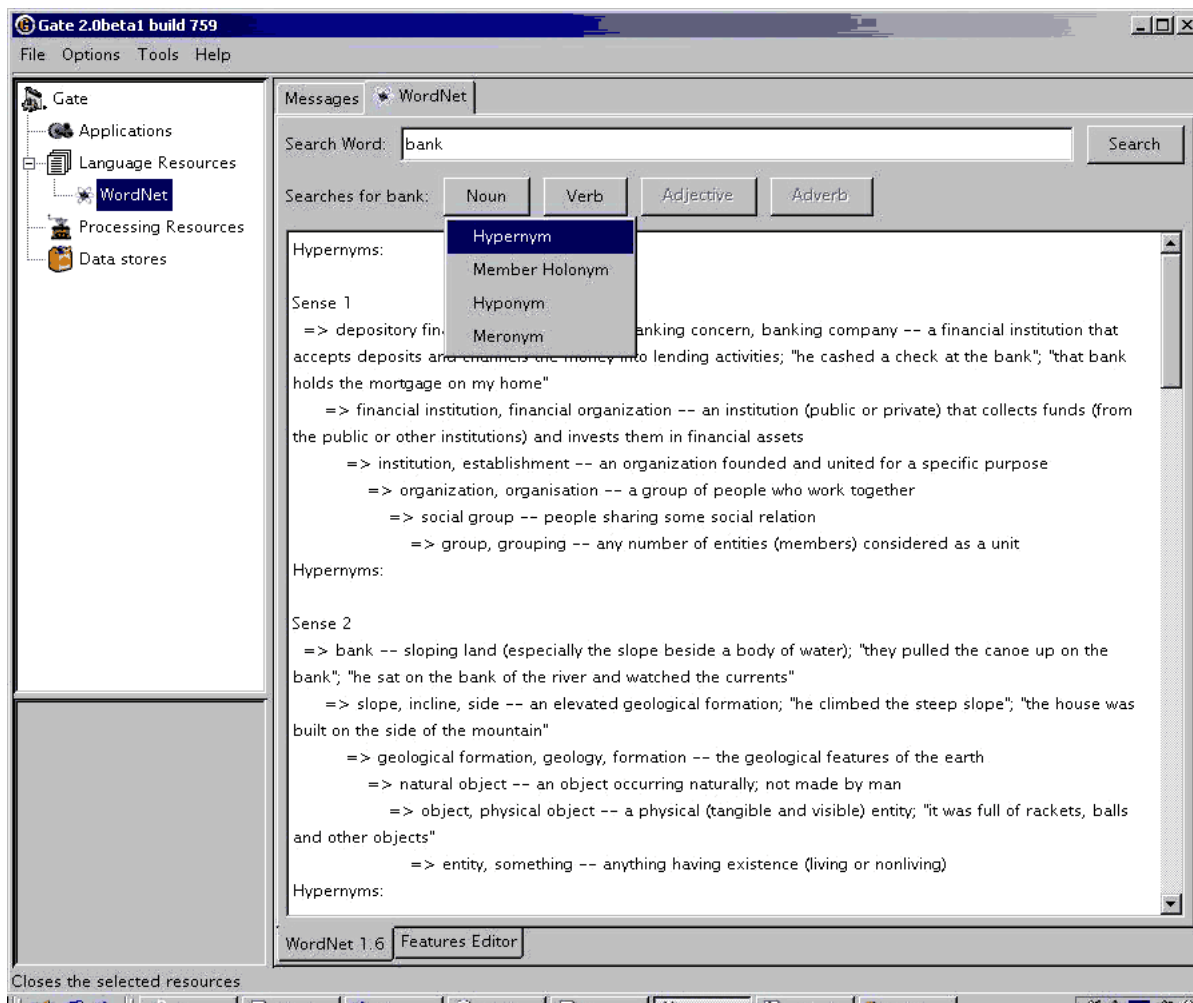


Figure 19.6: WordNet in GATE

WordNet access, this step consists of providing one special xml file that is used internally by JWNL. This file describes the location of your local copy of the WordNet 1.6 index files. An example of this wn-config.xml file is shown below:

```

<?xml version="1.0" encoding="UTF-8"?>

<jwnl_properties language="en">
  <version publisher="Princeton" number="1.6" language="en"/>
  <dictionary class="net.didion.jwnl.dictionary.FileBackedDictionary">
    <param
      name="morphological_processor"
      value="net.didion.jwnl.dictionary.DefaultMorphologicalProcessor"/>
    <param
      name="file_manager"
      value="net.didion.jwnl.dictionary.file_manager.FileManagerImpl">
      <param
        name="file_type"
        value="net.didion.jwnl.princeton.file.PrincetonRandomAccessDictionaryFile"/>
      <param
        name="dictionary_path"
        value="e:\wn16\dict"/>
      </param>
    </dictionary>
    <dictionary_element_factory
      class="net.didion.jwnl.princeton.data.PrincetonWN16DictionaryElementFactory"/>
    <resource class="PrincetonResource"/>
  </jwnl_properties>

```

All you have to do is to replace the value of the `dictionary_path` parameter to point to your local installation of WordNet 1.6.

After configuring GATE to use WordNet, you can start using the built-in WordNet browser or API. In GATE Developer, load the WordNet plugin via the Plugin Management Console. Then load WordNet by selecting it from the set of available language resources. Set the value of the parameter to the path of the xml properties file which describes the WordNet location (`wn-config`).

Once Word Net is loaded in GATE Developer, the well-known interface of WordNet will appear. You can search Word Net by typing a word in the box next to to the label ‘Search-Word’ and then pressing ‘Search’. All the senses of the word will be displayed in the window below. Buttons for the possible parts of speech for this word will also be activated at this point. For instance, for the word ‘play’, the buttons ‘Noun’, ‘Verb’ and ‘Adjective’ are activated. Pressing one of these buttons will activate a menu with hyponyms, hypernyms, meronyms for nouns or verb groups, and cause for verbs, etc. Selecting an item from the menu will display the results in the window below.

More information about WordNet can be found at <http://wordnet.princeton.edu/>

More information about the JWNL library can be found at <http://sourceforge.net/projects/jwordnet>

An example of using the WordNet API in GATE is available on the GATE examples page at <http://gate.ac.uk/GateExamples/doc/index.html>.

19.8.1 The WordNet API

GATE Embedded offers a set of classes that can be used to access the WordNet 1.6 Lexical Base. The implementation of the GATE API for WordNet is based on Java WordNet Library (JWNL). There are just a few basic classes, as shown in Figure 19.7. Details about the properties and methods of the interfaces/classes comprising the API can be obtained from the JavaDoc. Below is a brief overview of the interfaces:

- **WordNet**: the main WordNet class. Provides methods for getting the synsets of a lemma, for accessing the unique beginners, etc.
- **Word**: offers access to the word's lemma and senses
- **WordSense**: gives access to the synset, the word, POS and lexical relations.
- **Synset**: gives access to the word senses (synonyms) in the synset, the semantic relations, POS etc.
- **Verb**: gives access to the verb frames (not working properly at present)
- **Adjective**: gives access to the adj. position (attributive, predicative, etc.).
- **Relation**: abstract relation such as type, symbol, inverse relation, set of POS tags, etc. to which it is applicable.
- **LexicalRelation**
- **SemanticRelation**
- **VerbFrame**

19.9 Kea - Automatic Keyphrase Detection

Kea is a tool for automatic detection of key phrases developed at the University of Waikato in New Zealand. The home page of the project can be found at <http://www.nzdl.org/Kea/>.




Figure 19.7: The Wordnet API

This user guide section only deals with the aspects relating to the integration of Kea in GATE. For the inner workings of Kea, please visit the Kea web site and/or contact its authors.

In order to use Kea in GATE Developer, the ‘Keyphrase_Extraction_Algorithm’ plugin needs to be loaded using the plugins management console. After doing that, two new resource types are available for creation: the ‘KEA Keyphrase Extractor’ (a processing resource) and the ‘KEA Corpus Importer’ (a visual resource associated with the PR).

19.9.1 Using the ‘KEA Keyphrase Extractor’ PR

Kea is based on machine learning and it needs to be trained before it can be used to extract keyphrases. In order to do this, a corpus is required where the documents are annotated with keyphrases. Corpora in the Kea format (where the text and keyphrases are in separate files with the same name but different extensions) can be imported into GATE using the ‘KEA Corpus Importer’ tool. The usage of this tool is presented in a subsection below.

Once an annotated corpus is obtained, the ‘KEA Keyphrase Extractor’ PR can be used to build a model:

1. load a ‘KEA Keyphrase Extractor’
2. create a new ‘Corpus Pipeline’ controller.

Name	Type	Required	Value
document	gate.Document	✓	<none>
inputAS	java.lang.String		
outputAS	java.lang.String		
minPhraseLength	java.lang.Integer	✓	1
minNumOccur	java.lang.Integer	✓	2
maxPhraseLength	java.lang.Integer	✓	3
phrasesToExtract	java.lang.Integer	✓	5
keyphraseAnnotationType	java.lang.String	✓	Keyphrase
disallowInternalPeriods	java.lang.Boolean	✓	true
trainingMode	java.lang.Boolean	✓	true
useKFrequency	java.lang.Boolean	✓	true

Figure 19.8: Parameters used by the Kea PR

3. set the corpus for the controller
4. set the ‘trainingMode’ parameter for the PR to ‘true’
5. run the application.

After these steps, the Kea PR contains a trained model. This can be used immediately by switching the ‘trainingMode’ parameter to ‘false’ and running the PR over the documents that need to be annotated with keyphrases. Another possibility is to save the model for later use, by right-clicking on the PR name in the right hand side tree and choosing the ‘Save model’ option.

When a previously built model is available, the training procedure does not need to be repeated, the existing model can be loaded in memory by selecting the ‘Load model’ option in the PR’s context menu.

The Kea PR uses several parameters as seen in Figure 19.8:

document The document to be processed.

inputAS The input annotation set. This parameter is only relevant when the PR is running in training mode and it specifies the annotation set containing the keyphrase annotations.

outputAS The output annotation set. This parameter is only relevant when the PR is running in application mode (i.e. when the ‘trainingMode’ parameter is set to false) and it specifies the annotation set where the generated keyphrase annotations will be saved.

minPhraseLength the minimum length (in number of words) for a keyphrase.

minNumOccur the minimum number of occurrences of a phrase for it to be a keyphrase.

maxPhraseLength the maximum length of a keyphrase.

phrasesToExtract how many different keyphrases should be generated.

keyphraseAnnotationType the type of annotations used for keyphrases.

dissallowInternalPeriods should internal periods be disallowed.

trainingMode if ‘true’ the PR is running in training mode; otherwise it is running in application mode.

useKFrequency should the K-frequency be used.

19.9.2 Using Kea Corpora

The authors of Kea provide on the project web page a few manually annotated corpora that can be used for training Kea. In order to do this from within GATE, these corpora need to be converted to the format used in GATE (i.e. GATE documents with annotations). This is possible using the ‘KEA Corpus Importer’ tool which is available as a visual resource associated with the Kea PR. The importer tool can be made visible by double-clicking on the Kea PR’s name in the resources tree and then selecting the ‘KEA Corpus Importer’ tab, see Figure 19.9.

The tool will read files from a given directory, converting the text ones into GATE documents and the ones containing keyphrases into annotations over the documents.

The user needs to specify a few values:

Source Directory the directory containing the text and key files. This can be typed in or selected by pressing the folder button next to the text field.

Extension for text files the extension used for text fields (by default .txt).

Extension for keyphrase files the extension for the files listing keyphrases.

Encoding for input files the encoding to be used when reading the files.

Corpus name the name for the GATE corpus that will be created.




Figure 19.9: Options for the ‘KEA Corpus Importer’

Output annotation set the name for the annotation set that will contain the keyphrases read from the input files.

Keyphrase annotation type the type for the generated annotations.

19.10 Ontotext JapeC Compiler

Note: the JapeC compiler does not currently support the new JAPE language features introduced in July–September 2008. If you need to use negation, the @length and @string accessors, the contextual operators within and contains, or any comparison operators other than ==, then you will need to use the standard JAPE transducer instead of JapeC.

JapeC is an alternative implementation of the JAPE language which works by compiling JAPE grammars into Java code. Compared to the standard implementation, these compiled grammars can be several times faster to run. At Ontotext, a modified version of the ANNIE sentence splitter using compiled grammars has been found to run up to five times as fast as the standard version. The compiler can be invoked manually from the command line, or used through the ‘Ontotext JapeC Compiler’ PR in the *Jape_Compiler* plugin.

The ‘Ontotext JapeC Transducer’ (`com.ontotext.gate.japec.JapeCTransducer`) is a processing resource that is designed to be an alternative to the original Jape Transducer. You can simply replace `gate.creole.Transducer` with `com.ontotext.gate.japec.JapeCTransducer` in your

gate application and it should work as expected.

The Japec transducer takes the same parameters as the standard JAPE transducer:

grammarURL the URL from which the grammar is to be loaded. Note that the Japec Transducer will *only* work on `file:` URLs. Also, the alternative `binaryGrammarURL` parameter of the standard transducer is not supported.

encoding the character encoding used to load the grammars.

ontology the ontology used for ontolog-aware transduction.

Its runtime parameters are likewise the same as those of the standard transducer:

document the document to process.

inputASName name of the AnnotationSet from which input annotations to the transducer are read.

outputASName name of the AnnotationSet to which output annotations from the transducer are written.

The Japec compiler itself is written in Haskell. Compiled binaries are provided for Windows, Linux (x86) and Mac OS X (PowerPC), so no Haskell interpreter is required to run Japec on these platforms. For other platforms, or if you make changes to the compiler source code, you can build the compiler yourself using the Ant build file in the `Jape_Compiler` plugin directory. You will need to install the latest version of the Glasgow Haskell Compiler¹ and associated libraries. The japec compiler can then be built by running:

```
../../bin/ant japec.clean japec
```

from the `Jape_Compiler` plugin directory.

19.11 Annotation Merging Plugin

If we have annotations about the same subject on the same document from different annotators, we may need to merge the annotations.

This plugin implements two approaches for annotation merging.

MajorityVoting takes a parameter `numMinK` and selects the annotation on which at least `numMinK` annotators agree. If two or more merged annotations have the same span, then

¹GHC version 6.4.1 was used to build the supplied binaries for Windows, Linux and Mac

the annotation with the most supporters is kept and other annotations with the same span are discarded.

MergingByAnnotatorNum selects one annotation from those annotations with the same span, which the majority of the annotators support. Note that if one annotator did not create the annotation with the particular span, we count it as one non-support of the annotation with the span. If it turns out that the majority of the annotators did not support the annotation with that span, then no annotation with the span would be put into the merged annotations.

The annotation merging methods are available via the Annotation Merging plugin. The plugin can be used as a PR in a pipeline or corpus pipeline. To use the PR, each document in the pipeline or the corpus pipeline should have the annotation sets for merging. The annotation merging PR has no loading parameters but has several run-time parameters, explained further below.

The annotation merging methods are implemented in the GATE API, and are available in GATE Embedded as described in Section 7.18.

Parameters

- *annSetOutput*: the annotation set in the current document for storing the merged annotations. You should not use an existing annotation set, as the contents may be deleted or overwritten.
- *annSetsForMerging*: the annotation sets in the document for merging. It is an optional parameter. If it is not assigned with any value, the annotation sets for merging would be all the annotation sets in the document except the default annotation set. If specified, it is a sequence of the names of the annotation sets for merging, separated by ‘;’. For example, the value ‘a-1;a-2;a-3’ represents three annotation set, ‘a-1’, ‘a-2’ and ‘a-3’.
- *annTypeAndFeats*: the annotation types in the annotation set for merging. It is an optional parameter. It specifies the annotation types in the annotation sets for merging. For each type specified, it may also specify an annotation feature of the type. The parameter is a sequence of names of annotation types, separated by ‘;’. A single annotation feature can be specified immediately following the annotation type’s name, separated by ‘->’ in the sequence. For example, the value ‘SENT->senRel;OPINION_OPR;OPINION_SRC->type’ specifies three annotation types, ‘SENT’, ‘OPINION_OPR’ and ‘OPINION_SRC’ and specifies the annotation feature ‘senRel’ and ‘type’ for the two types SENT and OPINION_SRC, respectively but does not specify any feature for the type OPINION_OPR. If the *annTypeAndFeats* parameter is not set, the annotation types for merging are all the types in the annotation sets for merging, and no annotation feature for each type is specified.
- *keepSourceForMergedAnnotations*: should source annotations be kept in the *annSetsForMerging* annotation sets when merged? True by default.

- *mergingMethod*: specifies the method used for merging. Possible values are *MajorityVoting* and *MergingByAnnotatorNum*, referring to the two merging methods described above, respectively.
- *minimalAnnNum*: specifies the minimal number of annotators who agree on one annotation in order to put the annotation into merged set, which is needed by the merging method *MergingByAnnotatorNum*. If the value of the parameter is smaller than 1, the parameter is taken as 1. If the value is bigger than total number of annotation sets for merging, it is taken to be total number of annotation sets. If no value is assigned, a default value of 1 is used. Note that the parameter does not have any effect on the other merging method *MajorityVoting*.

19.12 Chinese Word Segmentation

Unlike English, Chinese text does not have a symbol (or delimiter) such as blank space to explicitly separate a word from the surrounding words. Therefore, for automatic Chinese text processing, we may need a system to recognise the words in Chinese text, a problem known as Chinese word segmentation. The plugin described in this section performs the task of Chinese word segmentation. It is based on our work using the Perceptron learning algorithm for the Chinese word segmentation task of the Sighan 2005². [Li *et al.* 05c]. Our Perceptron based system has achieved very good performance in the Sighan-05 task.

The plugin is called *Lang_Chinese* and is available in the GATE distribution. The corresponding processing resource's name is *Chinese Segmenter PR*. Once you load the PR into GATE, you may put it into a *Pipeline* application. Note that it does not process a corpus of documents, but a directory of documents provided as a parameter (see description of parameters below). The plugin can be used to learn a model from segmented Chinese text as training data. It can also use the learned model to segment Chinese text. The plugin can use different learning algorithms to learn different models. It can deal with different character encodings for Chinese text, such as UTF-8, GB2312 or BIG5. These options can be selected by setting the run-time parameters of the plugin.

The plugin has five run-time parameters, which are described in the following.

- **learningAlg** is a String variable, which specifies the learning algorithm used for producing the model. Currently it has two values, *PAUM* and *SVM*, representing the two popular learning algorithms Perceptron and SVM, respectively. The default value is *PAUM*.

Generally speaking, SVM may perform better than Perceptron, in particular for small training sets. On the other hand, Perceptron's learning is much faster than SVM's. Hence, if you have a small training set, you may want to use SVM to obtain a better

²See <http://www.sighan.org/bakeoff2005/> for the Sighan-05 task

model. However, if you have a big training set which is typical for the Chinese word segmentation task, you may want to use Perceptron for learning, because the SVM's learning may take too long time. In addition, using a big training set, the performance of the Perceptron model is quite similar to that of the SVM model. See [Li *et al.* 05c] for the experimental comparison of SVM and Perceptron on Chinese word segmentation.

- **learningMode** determines the two modes of using the plugin, either learning a model from training data or applying a learned model to segment Chinese text. Accordingly it has two values, *SEGMENTING* and *LEARNING*. The default value is *SEGMENTING*, meaning segmenting the Chinese text.

Note that you first need to learn a model and then you can use the learned model to segment the text. Several models using the training data used in the Sighan-05 Bakeoff are available for this plugin, which you can use to segment your Chinese text. More descriptions about the provided models will be given below.

- **modelURL** specifies an URL referring to a directory containing the model. If the plugin is in the *LEARNING* runmode, the model learned will be put into the directory. If it is in the *SEGMENTING* runmode, the plugin will use the model stored in the directory to segment the text. The models learned from the Sighan-05 bakeoff training data will be discussed below.
- **textCode** specifies the encoding of the text used. For example it can be UTF-8, BIG5, GB2312 or any other encoding for Chinese text. Note that, when you segment some Chinese text using a learned model, the Chinese text should use the same encoding as the one used by the training text for obtaining the model.
- **textFilesURL** specifies an URL referring to a directory containing the Chinese documents. All the documents contained in this directory (but not those documents contained in its sub-directory if there is any) will be used as input data. In the *LEARNING* runmode, those documents contain the segmented Chinese text as training data. In the *SEGMENTING* runmode, the text in those documents will be segmented. The segmented text will be stored in the corresponding documents in the sub-directory called *segmented*.

The following PAUM models are distributed with plugins and are available as compressed zip files under the `plugins/Lang_Chinese/resources/models` directory. Please unzip them to use. In detail, those models were learned using the PAUM learning algorithm from the corpora provided by Sighan-05 bakeoff task.

- the PAUM model learned from PKU training data, using the PAUM learning algorithm and the *UTF-8* encoding, is available as `model-paum-pku-utf8.zip`.
- the PAUM model learned from PKU training data, using the PAUM learning algorithm and the *GB2312* encoding, is available as `model-paum-pku-gb.zip`.

- the PAUM model learned from AS training data, using the PAUM learning algorithm and the *UTF-8* encoding, is available as `model-as-utf8.zip`.
- the PAUM model learned from AS training data, using the PAUM learning algorithm and the *BIG5* encoding, is available as `model-as-big5.zip`.

As you can see, those models were learned using different training data and different Chinese text encodings of the same training data. The PKU training data are news articles published in mainland China and use simplified Chinese, while the AS training data are news articles published in Taiwan and use traditional Chinese. If your text are in simplified Chinese, you can use the models trained by the PKU data. If your text are in traditional Chinese, you need to use the models trained by the AS data. If your data are in GB2312 encoding or any compatible encoding, you need use the model trained by the corpus in GB2312 encoding.

Note that the segmented Chinese text (either used as training data or produced by this plugin) use the blank space to separate a word from its surrounding words. Hence, if your data are in Unicode such as UTF-8, you can use the *GATE Unicode Tokeniser* to process the segmented text to add the Token annotations into your text to represent the Chinese words. Once you get the annotations for all the Chinese words, you can perform further processing such as POS tagging and named entity recognition.

19.13 Copying Annotations between Documents

Sometimes a document has two copies, each of which was annotated by different annotators for the same task. We may want to copy the annotations in one copy to the other copy of the document. This could be in order to use less resources, or so that we can process them with some other plugin, such as annotation merging or IAA. The **Copy_Annots_Between_Docs** plugin does exactly this.

The plugin is available with the GATE distribution. When loading the plugin into GATE, it is represented as a processing resource, **Copy Anns to Another Doc PR**. You need to put the PR into a *Corpus Pipeline* to use it. The plugin does not have any initialisation parameters. It has several run-time parameters, which specify the annotations to be copied, the source documents and target documents. In detail, the run-time parameters are:

- **sourceFilesURL** specifies a directory in which the source documents are in. The source documents must be GATE xml documents. The plugin copies the annotations from these source documents to target documents.
- **inputASName** specifies the name of the annotation set in the source documents. Whole annotations or parts of annotations in the annotation set will be copied.

- **annotationTypes** specifies one or more annotation types in the annotation set *inputASName* which will be copied into target documents. If no value is given, the plugin will copy all annotations in the annotation set.
- **outputASName** specifies the name of the annotation set in the target documents, into which the annotations will be copied. If there is no such annotation set in the target documents, the annotation set will be created automatically.

The **Corpus** parameter of the *Corpus Pipeline* application containing the plugin specifies a corpus which contains the target documents. Given one (target) document in the corpus, the plugin tries to find a source document in the source directory specified by the parameter *sourceFilesURL*, according to the similarity of the names of the source and target documents. The similarity of two file names is calculated by comparing the two strings of names from the start to the end of the strings. Two names have greater similarity if they share more characters from the beginning of the strings. For example, suppose two target documents have the names *aabcc.xml* and *abcab.xml* and three source files have names *abacc.xml*, *abcbb.xml* and *aacc.xml*, respectively. Then the target document *aabcc.xml* has the corresponding source document *aacc.xml*, and *abcab.xml* has the corresponding source document *abcbb.xml*.

19.14 OpenCalais Plugin

OpenCalais provides a web service for semantic annotation of text. The user submits a document to the web service, which returns entity and relations annotations in RDF, JSON or some other format. Typically, users integrate OpenCalais annotation of their web pages to provide additional links and ‘semantic functionality’. OpenCalais can be found at <http://www.opencalais.com>

The GATE OpenCalais PR submits a GATE document to the OpenCalais web service, and adds the annotations from the OpenCalais response as GATE annotations in the GATE document. It therefore provides OpenCalais semantic annotation functionality within GATE, for use by other PRs.

The PR only supports OpenCalais entities, not relations - although this should be straightforward for a competent Java programmer to add. Each OpenCalais entity is represented in GATE as an **OpenCalais** annotation, with features as given in the OpenCalais documentation.

The PR can be loaded with the CREOLE plugin manager dialog, from the creole directory in the gate distribution, *gate/plugins/Tagger.OpenCalais*. In order to use the PR, you will need to have an OpenCalais account, and request an OpenCalais service key. You can do this from the OpenCalais web site at <http://www.opencalais.com>. Provide your service key as an initialisation parameter when you create a new OpenCalais PR in GATE. OpenCalais

make restrictions on the number of requests you can make to their web service. See the OpenCalais web page for details.

Initialisation parameters are:

- **openCalaisURL** This is the URL of the OpenCalais REST service, and should not need to be changed - unless OpenCalais moves it!
- **licenseID** Your OpenCalais service key. This has to be requested from OpenCalais and is specific to you.

Various runtime parameters are available from the OpenCalais API, and are named the same as in that API. See the OpenCalais documentation for further details.

19.15 LingPipe Plugin

LingPipe is a suite of Java libraries for the linguistic analysis of human language³. We have provided a plugin called ‘LingPipe’ with wrappers for some of the resources available in the LingPipe library. In order to use these resources, please load the ‘LingPipe’ plugin. Currently, we have integrated the following five processing resources.

- LingPipe Tokenizer PR
- LingPipe Sentence Splitter PR
- LingPipe POS Tagger PR
- LingPipe NER PR
- LingPipe Language Identifier PR

Please note that most of the resources in the LingPipe library allow learning of new models. However, in this version of the GATE plugin for LingPipe, we have only integrated the application functionality. You will need to learn new models with Lingpipe outside of GATE. We have provided some example models under the ‘resources’ folder which were downloaded from LingPipe’s website. For more information on licensing issues related to the use of these models, please refer to the licensing terms under the LingPipe plugin directory.

The LingPipe system can be loaded from the GATE GUI by simply selecting the ‘Load LingPipe System’ menu item under the ‘File’ menu. This is similar to loading the ANNIE application with default values.

³see <http://alias-i.com/lingpipe/>

19.15.1 LingPipe Tokenizer PR

As the name suggests this PR tokenizes document text and identifies the boundaries of tokens. Each token is annotated with an annotation of type 'Token'. Every annotation has a feature called 'length' that gives a length of the word in number of characters. There are no initialization parameters for this PR. The user needs to provide the name of the annotation set where the PR should output Token annotations.

19.15.2 LingPipe Sentence Splitter PR

As the name suggests, this PR splits document text in sentences. It identifies sentence boundaries and annotates each sentence with an annotation of type 'Sentence'. There are no initialization parameters for this PR. The user needs to provide name of the annotation set where the PR should output Sentence annotations.

19.15.3 LingPipe POS Tagger PR

The LingPipe POS Tagger PR is useful for tagging individual tokens with their respective part of speech tags. This PR requires a model which it then uses to tag the tokens. An example model is provided under the 'resources' folder of this plugin. It must be provided at initialization time. It is a prerequisites of this PR that the document is processed with a Tokenizer and a Sentence Splitter. In other words, it expects annotations of type 'Token' and 'Sentence' to be available in the document. The Tokeniser and Sentence Splitter can be any such PR from GATE, and need not be the Lingpipe variants. This PR adds a feature called 'category' on each token.

Below we list the runtime parameters for this PR.

- `inputASName`: This is the name of the annotation set with 'Token' and 'Sentence' annotations in it.
- `applicationMode`: The POS tagger can be applied on the text in three different modes.
 1. `FIRSTBEST`: In this case, the POS tagger suggests one tag for each token that is best according to its calculations.
 2. `CONFIDENCE`: This is same as the `FIRSTBEST` except that in this case it also adds a feature 'score' with the actual calculated score for the tag that is assigned to the token.
 3. `NBEST`: In this case, the POS tagger suggests n best tags for each token. The default value is set to 5. In other words, it suggests five best tags for each token.

19.15.4 LingPipe NER PR

The LingPipe NER PR is used for named entity recognition. The PR recognizes entities such as Persons, Organizations and Locations in the text. This PR requires a model which it then uses to classify text as different entity types. An example model is provided under the ‘resources’ folder of this plugin. It must be provided at initialization time. Similar to other PRs, this PR expects users to provide name of the annotation set where the PR should output annotations.

19.15.5 LingPipe Language Identifier PR

As the name suggests, this PR is useful for identifying the language of a document. This PR requires a model file which it then uses to identify the language of the document. An example model is provided under the ‘resources’ folder of this plugin. It must be provided at initialization time. Unlike other PRs which produce annotations, this PR adds a document feature. The name of the document feature can be specified as a runtime parameter. More information on how many languages are supported by the PR can be found on the following url: <http://alias-i.com/lingpipe/demos/tutorial/langid/read-me.html>

19.16 OpenNLP Plugin

OpenNLP provides java-based tools for sentence detection, tokenization, pos-tagging, chunking and parsing, named-entity detection, and coreference. See the OpenNLP website for details: <http://opennlp.sourceforge.net/>. The tools use the Maxent machine learning package. See <http://maxent.sourceforge.net/> for details.

In order to use these tools as GATE processing resources, load the ‘OpenNLP’ plugin via the Plugin Management Console. Alternatively, the OpenNLP system can be loaded from the GATE GUI by simply selecting the ‘Load OpenNLP System’ menu item under the ‘File’ menu. The OpenNLP PRs will be loaded, together with a pre-configured corpus pipeline application containing these PRs. This is similar to loading the ANNIE application with default values.

We have integrated the following five processing resources:

- OpenNlpTokenizer Tokenizer PR
- OpenNlpSentenceSplit Sentence Splitter PR
- OpenNlpPOS POS Tagger PR
- OpenNlpNameFinder NER PR

- OpenNlpChunker Chunker PR

Note that in general, these PRs can be mixed with other PRs of similar types. For example, you could create a pipeline that uses the OpenNLP Tokenizer, and the ANNIE POS Tagger. You may occasionally have problems with some combinations. Notes on compatibility and PR prerequisites are given for each PR in Section 19.16.2.

Note also that some of the OpenNLP tools use quite large machine learning models, which the PRs need to load into memory. You may find that you have to give additional memory to GATE in order to use the OpenNLP PRs comfortably. See Section 2.7.5 for an example of how to do this.

Below, we describe the parameters common to all of the OpenNLP PRs. This is followed by a section which gives brief details of each PR. For more details on each, see the OpenNLP website, <http://opennlp.sourceforge.net/>.

19.16.1 Parameters common to all PRs

Load-time Parameters

All OpenNLP PRs have a ‘model’ parameter, which takes a URL. The URL should reference a valid Maxent model, or in the case of the Name Finder a directory containing a set of models for the different types of name sought. Default models can be found in the ‘models/english’ directory. In addition, the OpenNlpPOS POS Tagger PR has a ‘dictionary’ parameter, which also takes a URL, and a ‘dictionaryEncoding’ parameter giving the character encoding of the dictionary file. The default can be found in the ‘models/english’ directory.

For details of training new models (outside of the GATE framework), see Section 19.16.3

Run-time Parameters

The OpenNLP PRs have runtime parameters to specify the annotation sets they should use for input and/or output. These are detailed below in the description of each PR, but all PRs will use the default unnamed annotation set unless told otherwise.

19.16.2 OpenNLP PRs

OpenNlpTokenizer - Tokenizer PR

This PR adds Token annotations to the annotation set specified by the `annotationSetName` parameter.

This PR does not require any other PR to be run beforehand. It creates annotations of type `Token`, with a feature and value `'source=openNLP'` and a string feature that takes the underlying string as its value.

OpenNlpSentenceSplit - Sentence Splitter PR

This PR adds Sentence annotations to the annotation set specified by the `annotationSetName` parameter.

This PR does not require any other PR to be run beforehand. It creates annotations of type `Sentence`, with a feature and value `'source=openNLP'`. If the OpenNLP sentence splitter returns no output (i.e. fails to find any sentences), then GATE will add a single sentence annotation from offset 0 to the end of the document. This is to prevent downstream PRs that require sentence annotations from failing.

OpenNlpPOS - POS Tagger PR

This PR adds a feature for Part Of Speech to `Token` annotations.

This PR requires `Sentence` and `Token` annotations to be present in the annotation set specified by its `annotationSetName` parameter before it will work. These `Sentence` and `Token` annotations do not have to be from another OpenNLP PR. They could, for example, be from the ANNIE PRs. This PR adds a `'category'` feature to each `Token`, with the predicted Part Of Speech as value.

OpenNlpNameFinder - NER PR

This PR finds standard Named Entities, adding them as annotations named for each named Entity type.

This PR requires `Sentence` and `Token` annotations to be present in the annotation set specified by its `inputASName` parameter before it will work. These `Sentence` and `Token` annotations do not have to be from another OpenNLP PR. They could, for example, be from the ANNIE PRs. You may find, however, that not all pairings of `Tokenizer` and `Sentence Splitter` will work successfully.

The `Token` annotations do not need to have a `'category'` feature. In other words, you do not need to run a POS Tagger before using this PR.

This PR creates annotations for each named entity in the annotation set specified by the `outputASName` parameter, with a feature and value `'source=openNLP'`. It also adds a feature of `'type'` to each annotation, with the same values as the annotation type. Types include:

- person
- organization
- location
- date
- time
- percentage

For full details of all types, see the OpenNLP website, <http://opennlp.sourceforge.net/>.

OpenNlpChunker - Chunker PR

This PR finds noun, verb, and other chunks, adding their position as features Token annotations.

This PR requires Sentence and Token annotations to be present in the annotation set specified by the `annotationSetName` parameter before it will work. The Token annotations need to have a 'category' feature. In other words, you also need to run a POS Tagger before using this PR. The Sentence and Token annotations (and 'category' POS tag features) do not have to be from another OpenNLP PR. They could, for example, be from the ANNIE PRs.

This PR creates features of type 'chunk' for each token. The value of this feature define whether the token is at the beginning of a chunk, inside a chunk, or outside a chunk, using the standard BIO model. Example values and their interpretations are:

- **B-NP** Token is the first (beginning) token of a Noun Phrase
- **I-NP** Token is inside a Noun Phrase
- **B-VP** Token is the first (beginning) token of a Verb Phrase
- **I-VP** Token is inside a Verb Phrase
- **O** Token is outside any phrase
- **B-PP** Token is the first (beginning) token of a Prepositional Phrase
- **B-ADVP** Token is the first (beginning) token of an Adverbial Phrase

For full details of all chunk values, see the OpenNLP website, <http://opennlp.sourceforge.net/>.

19.16.3 Training new models

Within the OpenNLP framework, new models can be trained for each of the tools. By default, the GATE PRs use the standard Maxent models for English which can be found in the plugin's 'models/english' directory. The models are copies of those in the "models" module in the OpenNLP CVS repository at <http://opennlp.cvs.sourceforge.net>. If you need to train a different model for an OpenNLP PR, you will have to do this outside of GATE, and then use the file URL of your new model as a value for the 'model' parameter of the PR. For details on how to train models, see the Maxent website <http://maxent.sourceforge.net/>.

19.17 Inter Annotator Agreement

The IAA plugin, "Inter_Annotator_Agreement", computes interannotator agreement measures for various tasks. For named entity annotations, it computes the F-measures, namely Precision, Recall and F1, for two or more annotation sets. For text classification tasks, it computes Cohen's kappa and some other IAA measures which are more suitable than the F-measures for the task. This plugin is fully documented in Section 10.5. Chapter 10 introduces various measures of interannotator agreement and describes a range of tools provided in GATE for calculating them.

19.18 Balanced Distance Metric Computation

The BDM (balanced distance metric) measures the closeness of two concepts in an ontology or taxonomy [Maynard 05, Maynard *et al.* 06]. It is a real number between 0 and 1. The closer the two concepts are in an ontology, the greater their BDM score is. The plugin, "Ontology_BDM_Computation", is described more fully in Section 10.6.

19.19 Schema Annotation Editor

The plugin 'Schema_Annotation_Editor' constrains the annotation editor to permitted types. See Section 3.4.6 for more information.

Appendix A

Change Log

This chapter lists major changes to GATE in roughly chronological order by release. Changes in the documentation are also referenced here.

A.1 Version 5.2.1 (May 2010)

This is a bugfix release to resolve several bugs that were reported shortly after the release of version 5.2:

- Fixed some bugs with the automatic “create instance” feature in OAT (the ontology annotation tool) when used with the new `Ontology` plugin.
- Added validation to datatype property values of the *date*, *time* and *datetime* types.
- Fixed a bug with `Gazetteer_LKB` that prevented it working when the `dictionaryPath` contained spaces.
- Added a utility class to handle common cases of encoding URIs for use in ontologies, and fixed the example code to show how to make use of this. See chapter 14 for details.
- The annotation set transfer PR now copies the feature map of each annotation it transfers, rather than re-using the same `FeatureMap` (this means that when used to copy annotations rather than move them, the copied annotation is independent from the original and modifying the features of one does not modify the other). See section 19.3 for details.
- The Log4J log files are now created by default in the `.gate` directory under the user’s home directory, rather than being created in the current directory when GATE starts, to be more friendly when GATE is installed in a shared location where the user does not have write permission.

This release also fixes some shortcomings in the Groovy support added by 5.2, in particular:

- The `corpora` variable in the console now includes persistent corpora (loaded from a datastore) as well as transient corpora.
- The subscript notation for annotation sets works with long values as well as ints, so `someAS[annotation.start()..annotation.end()]` works as expected.

A.2 Version 5.2 (April 2010)

A.2.1 JAPE and JAPE-related

Introduced a utility class `gate.Utils` containing static utility methods for frequently-used idioms such as getting the string covered by an annotation, finding the start and end offsets of annotations and sets, etc. This class is particularly useful on the right hand side of JAPE rules (section 8.6.5).

Added type parameters to the `bindings` map available on the RHS of JAPE rules, so you can now do `AnnotationSet as = bindings.get("label")` without a cast (see section 8.6.5).

Fixed a bug with JAPE's handling of features called "class" in non-ontology-aware mode. Previously JAPE would always match such features using an equality test, even if a different operator was used in the grammar, i.e. `{SomeType.class != "foo"}` was matched as `{SomeType.class == "foo"}`. The correct operator is now used. Note that this does not affect the ontology-aware behaviour: when an ontology parameter is specified, "class" features are always matched using ontology subsumption.

Custom JAPE operators and annotation accessors can now be loaded from plugins as well as from the `lib` directory (see section 8.2.2).

A.2.2 Other Changes

Added a mechanism to allow plugins to contribute menu items to the "Tools" menu in GATE Developer. See section 4.8 for details.

Enhanced Groovy support in GATE: the Groovy console and Groovy Script PR (in the Groovy plugin) now import many GATE classes by default, and a number of utility methods are mixed in to some of the core GATE API classes to make them more natural to use in Groovy. See section 7.16 for details.

Modified the batch learning PR (in the `Learning` plugin) to make it safe to use several instances in *APPLICATION* mode with the same configuration file and the same learned

model at the same time (e.g. in a multithreaded application). The other modes (including training and evaluation) are unchanged, and thus are still *not* safe for use in this way. Also fixed a bug that prevented APPLICATION mode from working anywhere other than as the last PR in a pipeline when running over a corpus in a datastore.

Introduced a simple way to create duplicate copies of an existing resource instance, with a way for individual resource types to override the default duplication algorithm if they know a better way to deal with duplicating themselves. See section 7.7.

Enhanced the Spring support in GATE to provide easy access to the new duplication API, and to simplify the configuration of the built-in Spring pooling mechanisms when writing multi-threaded Spring-based applications. See section 7.14.

The GAPP packager Ant task now respects the ordering of mapping hints, with earlier hints taking precedence over later ones (see section E.2.3).

Bug fix in the UIMA plugin from Roland Cornelissen - `AnalysisEnginePR` now properly shuts down the wrapped `AnalysisEngine` when the PR is deleted.

Patch from Matt Nathan to allow several instances of a gazetteer PR in an embedded application to share a single copy of their internal data structures, saving considerable memory compared to loading several complete copies of the same gazetteer lists (see section 13.11).

A.3 Version 5.1 (December 2009)

Version 5.1 is a major increment with lots of new features and integration of a number of important systems from 3rd parties (e.g. LingPipe, OpenNLP, OpenCalais, a revised UIMA connector). We've stuck with the 5 series (instead of jumping to 6.0) because the core remains stable and backwards compatible.

Other highlights include:

- an entirely new ontology API from Johann Petrak of OFAI (the old one is still available but as a plugin)
- new benchmarking facilities for JAPE from Andrew Borthwick and colleagues at Intelius
- new quality assurance tools from Thomas Heitz and colleagues at Ontotext and Sheffield
- a generic tagger integration framework from René Witte of Concordia University
- several new code contributions from Ontotext, including a large knowledge-based gazetteer and various plugin wrappers from Marin Nozchev, Georgi Georgiev and colleagues

- a revised and reordered user guide, amalgamated with the programmers' guide and other materials
- Groovy support, application composition, section-by-section processing and lots of other bits and pieces

A.3.1 New Features

LingPipe Support

LingPipe is a suite of Java libraries for the linguistic analysis of human language. We have provided a plugin called 'LingPipe' with wrappers for some of the resources available in the LingPipe library. For more details, see the section 19.15.

OpenNLP Support

OpenNLP provides tools for sentence detection, tokenization, pos-tagging, chunking and parsing, named-entity detection, and coreference. The tools use Maximum Entropy modelling. We have provided a plugin called 'OpenNLP' with wrappers for some of the resources available in the OpenNLP Tools library. For more details, see section 19.16.

OpenCalais Support

We added a new PR called 'OpenCalais PR'. This will process a document through the OpenCalais service, and add OpenCalais entity annotations to the document. For more details, see Section 19.14.

Ontology API

The ontology API (package `gate.creole.ontology`) has been changed, the existing ontology implementation based on Sesame1 and OWLIM2 (package `gate.creole.ontology.owlim`) has been moved into the plugin `Ontology_OWLIM2`. An upgraded implementation based on Sesame2 and OWLIM3 that also provides a number of new features has been added as plugin `Ontology`. See Section 14.12 for a detailed description of all changes.

Benchmarking Improvements

A number of improvements to the benchmarking support in GATE. JAPE transducers now log the time spent in individual phases of a multi-phase grammar and by individual rules

within each phase. Other PRs that use JAPE grammars internally (the pronominal coreferencer, English tokeniser) log the time taken by their internal transducers. A reporting tool, called ‘Profiling Reports’ under the ‘Tools’ menu makes summary information easily available. For more details, see chapter 11.

GUI improvements

To deal with quality assurance of annotations, one component has been updated and two new components have been added. The annotation diff tool has a new mode to copy annotations to a consensus set, see section 10.2.1. An annotation stack view has been added in the document editor and it allows to copy annotations to a consensus set, see section 3.4.3. A corpus view has been added for all corpus to get statistics like precision, recall and F-measure, see section 10.3.

An annotation stack view has been added in the document editor to make easier to see overlapping annotations, see section 3.4.3.

ABNER Support

ABNER is A Biomedical Named Entity Recogniser, for finding entities such as genes in text. We have provided a plugin called ‘AbnerTagger’ with a wrapper for ABNER. For more details, see section 17.6.

Generic Tagger Support

A new plugin has been added to provide an easy route to integrate taggers with GATE. The Tagger_Framework plugin provides examples of incorporating a number of external taggers which should serve as a starting point for using other taggers. See Section 17.4 for more details.

Section-by-Section Processing

We have added a new PR called ‘Segment Processing PR’. As the name suggests this PR allows processing individual segments of a document independently of one other. For more details, please look at the section 16.2.8.

Application Composition

The `gate.Controller` implementations provided with the main GATE distribution now also implement the `gate.ProcessingResource` interface. This means that an application can now contain another application as one of its components.

Groovy Support

Groovy is a dynamic programming language based on Java. You can now use it as a scripting language for GATE, via the Groovy Console. For more details, see Section 7.16.

A.3.2 JAPE improvements

GATE now produces a warning when any Java right-hand-sides in JAPE rules make use of the deprecated `annotations` parameter. All bundled JAPE grammars have been updated to use the replacement `inputAS` and `outputAS` parameters as appropriate.

The new `Imports:` statement at the beginning of a JAPE grammar file can now be used to make additional Java import statements available to the Java RHS code, see 8.6.5.

The JAPE debugger has been removed. Debugging of JAPE has been made easier as stack traces now refer to the JAPE source file and line numbers instead of the generated Java source code.

The Montreal Transducer has been made obsolete.

A.3.3 Other improvements and bug fixes

Plugin names have been rationalised. Mappings exist so that existing applications will continue to work, but the new names should be used in the future. Plugin name mappings are given in Appendix B. Also, the `Segmenter_Chinese` plugin (used to be known as `chineseSegmenter` plugin) is now part of the `Lang_Chinese` plugin.

The User Guide has been amalgamated with the Programmer's Guide; all material can now be found in the User Guide. The 'How-To' chapter has been converted into separate chapters for installation, GATE Developer and GATE Embedded. Other material has been relocated to the appropriate specialist chapter.

Made Mac OS launcher 64-bit compatible. See section 2.2.1 for details.

The UIMA integration layer (Chapter 18) has been upgraded to work with Apache UIMA 2.2.2.

Oracle and PostgreSQL are no longer supported.

The MIAKT Natural Language Generation plugin has been removed.

The Minorthird plugin has been removed. Minorthird has changed significantly since this plugin was written. We will consider writing an up-to-date Minorthird plugin in the future.

A new gazetteer, Large KB Gazetteer (in the plugin 'Gazetteer_LKB') has been added, see Section 13.10 for details.

gate.creole.tokeniser.chinesetokeniser.ChineseTokeniser and related resources under the plugins/ANNIE/tokeniser/chinesetokeniser folder have been removed. Please refer to the Lang-Chinese plugin for resources related to the Chinese language in GATE.

Added an `isInitialised()` method to `gate.Gate()`.

Added a parameter to the chemistry tagger PR (section 17.5) to allow it to operate on annotation sets other than the default one.

Plus many more smaller bugfixes...

A.4 Version 5.0 (May 2009)

Note: *existing users – if you delete your user configuration file for any reason you will find that GATE Developer no longer loads the ANNIE plugin by default. You will need to manually select 'load always' in the plugin manager to get the old behaviour.*

A.4.1 Major New Features

JAPE Language Improvements

Several new extensions to the JAPE language to support more flexible pattern matching. Full details are in Chapter 8 but briefly:

- Negative constraints, that prevent a rule from matching if certain other annotations are present (Section 8.1.9).
- Additional matching operators for feature values, so you can now look for `{Token.length < 5}`, `{Lookup.minorType != "ignore"}`, etc. as well as simple equality (Section 8.2.2).

- ‘Meta-property’ accessors, see Section 8.1.4 to permit access to the string covered by an annotation, the length of the annotation, etc., e.g. `{Lookup@length > 4}`.
- Contextual operators, allowing you to search for one annotation contained within (or containing) another, e.g. `{Sentence contains {Lookup.majorType == "location"}}` (see Section 8.2.2).
- Additional Kleene operator for ranges, e.g. `({Token}) [2,5]` matches between 2 and 5 consecutive tokens, see Section 8.2.1.
- Additional operators can be added via runtime configuration (see Section 8.2.2).

Some of these extensions are similar to, but not the same as, those provided by the Montreal Transducer plugin. If you are already familiar with the Montreal Transducer, you should first look at Section 8.11 which summarises the differences.

Resource Configuration via Java 5 Annotations

Introduced an alternative style for supplying resource configuration information via Java 5 annotations rather than in `creole.xml`. The previous approach is still fully supported as well, and the two styles can be freely mixed. See Section 4.7 for full details.

Ontology-Based Gazetteer

Added a new plugin ‘Gazetteer_Ontology_Based’, which contains OntoRoot Gazetteer – a dynamically created gazetteer which is, in combination with few other generic resources, capable of producing ontology-aware annotations over the given content with regards to the given ontology. For more details see Section 13.9.

Inter-Annotator Agreement and Merging

New plugins to support tasks involving several annotators working on the same annotation task on the same documents. The plugin ‘Inter_Annotator_Agreement’ (Section 10.5) computes inter-annotator agreement scores between the annotators, the ‘Copy_Annots_Between_Docs’ plugin (Section 19.13) copies annotations from several parallel documents into a single master document, and the ‘Annotation_Merging’ plugin (Section 19.11) merges annotations from multiple annotators into a single ‘consensus’ annotation set.

Packaging Self-Contained Applications for GATE Teamware

Added a mechanism to assemble a saved GATE application along with all the resource files it uses into a single self-contained package to run on another machine (e.g. as a service in GATE Teamware). This is available as a menu option (Section 3.8.4) which will work for most common cases, but for complex cases you can use the underlying Ant task described in Section E.2.

GUI Improvements

- A new schema-driven tool to streamline manual annotation tasks (see Section 3.4.6).
- Context-sensitive help on elements in the resource tree and when pressing F1 key. Search in mailing list from the Help menu. Help is displayed in your browser or in a Java browser if you don't have one.
- Improved search function inside documents with a regular expression builder. Search and replace annotation function in all annotation editors.
- Remember for each resource type the last path used when loading/saving a resource.
- Remember the last annotations selected in the annotation set view when you shift click on the annotation set view button.
- Improved context menu and when possible added drag and drop in: resource tree, annotation set view, annotation list view, corpus view, controller view. Context menu key can be now used if you have Java 1.6.
- New dialog box for error messages with user oriented messages, optional display of the configuration and proposing some useful actions. This will progressively replace the old stack trace dump into the message panel which is still here for the moment but should be hide by default in the future.
- Add read-only document mode that can be enable from the Options menu.
- Add a selection filter in the status bar of the annotations list table to easily select rows based on the text you enter.
- Add the last five applications loaded/saved in the context menu of the language resources in the resources tree.
- Display more informations on what going's on in the waiting dialog box when running an application. The goal is to improve it to get a global progress bar and estimated time.

A.4.2 Other New Features and Improvements

- New parser plugins: A new plugin for the Stanford Parser (see Section 17.12) and a rewritten plugin for the RASP NLP tools (Section 17.10).
- A new sentence splitter, based on regular expressions, has been added to the ANNIE plugin. More details in Section 6.5.
- ‘Real-time’ corpus controller (Section 4.4), which terminates processing of a document if it takes longer than a configurable timeout..
- Major update to Annie OrthoMatcher coreference engine. Now correctly matches the sequence ‘David Jones ... David ... David Smith ... David’ as referring to two people. Also handles nicknames (David = Dave) via a new nickname list. Added optional parameter ‘highPrecisionOrgs’, which if set to true turns off riskier org matching rules. Many misc. bug fixes.
- Improved alignment editor (Chapter 16) with several advanced features and an API for adding your own actions to the editor.
- A new plugin for Chinese word segmentation, which is based on our work using machine learning algorithms for the Sighan-05 Chinese word segmentation task. It can learn a model from manually segmented text, and apply a learned model to segment Chinese text. In addition several learned models are available with the plugin, which can be used to segment text. For details about the plugin and those learned models see Section 19.12.
- New features in the ML API to produce an n-gram based language model from a corpus and a so-called ‘document-term matrix’ (see Section 19.4). Also introduced features to support active learning, a new learning algorithm (PAUM) and various optimisations including the ability to use an external executable for SVM training. Full details in Chapter 15.
- A new plugin to compute BDM scores for an ontology. The BDM score can be used to evaluate ontology based information extraction and classification. For details about the plugin see Section 10.6.
- Added new ‘getCovering’ method to AnnotationSet. This method returns annotations that completely span the provided range. An optional annotation type parameter can be provided to further limit the returned set.
- Complete redesign of ANNIE GUI. More details in Section 9.

A.4.3 Specific Bug Fixes

- HTML document format parser: several bugs fixed, including a null pointer exception if the document contained certain characters illegal in HTML (#1754749). Also, the

HTML parser now respects the ‘Add space on markup unpack’ configuration option – previously it would always add space, even if the option was set to false.

- Fixed a severe performance bug in the Annie Pronominal Coreferencer resulting in a 50X speed improvement.
- JAPE did not always correctly handle the case when the input and output annotation sets for a transducer were different. This has now been fixed.
- ‘Save Preserving Format’ was not correctly escaping ampersands and less than signs when two HTML entities are close together. Only the first one was replaced: A & B & C was output as A & B & C instead of A & B & C. This has now been fixed, and the fix is also valid for the flexible exporter but only if the standoff annotations parameter is set to false.

Plus many more minor bug fixes

A.5 Version 4.0 (July 2007)

A.5.1 Major New Features

ANNIC

ANNotations In Context: a full-featured annotation indexing and retrieval system designed to support corpus querying and JAPE rule authoring. It is provided as part of an extension of the Serial Datastores, called Searchable Serial Datastore (SSD). See Section 9 for more details.

New Machine Learning API

A brand new machine learning layer specifically targetted at NLP tasks including text classification, chunk learning (e.g. for named entity recognition) and relation learning. See Chapter 15 for more details.

Ontology API

A new ontology API, based on OWL In Memory (OWLIM), which offers a better API, revised ontology event model and an improved ontology editor to name but few. See Chapter 14 for more details.

OCAT

Ontology-based Corpus Annotation Tool to help annotators to manually annotate documents using ontologies. For more details please see Section 14.6.

Alignment Tools

A new set of components (e.g. `CompoundDocument`, `AlignmentEditor` etc.) that help in building alignment tools and in carrying out cross-document processing. See Chapter 16 for more details.

New HTML Parser

A new HTML document format parser, based on Andy Clark's `NekoHTML`. This parser is much better than the old one at handling modern HTML and XHTML constructs, JavaScript blocks, etc., though the old parser is still available for existing applications that depend on its behaviour.

Java 5.0 Support

GATE now requires Java 5.0 or later to compile and run. This brings a number of benefits:

- Java 5.0 syntax is now available on the right hand side of JAPE rules with the default Eclipse compiler. See Section D.5 for details.
- `enum` types are now supported for resource parameters. see Section 7.11 for details on defining the parameters of a resource.
- `AnnotationSet` and the `CreoleRegister` take advantage of generic types. The `AnnotationSet` interface is now an extension of `Set<Annotation>` rather than just `Set`, which should make for cleaner and more type-safe code when programming to the API, and the `CreoleRegister` now uses parameterized types, which are backwards-compatible but provide better type-safety for new code.

A.5.2 Other New Features and Improvements

- Hiding the view for a particular resource (by right clicking on its tab and selecting 'Hide this view') will now completely close the associated viewers and dispose them. Re-selecting the same resource at a later time will lead to re-creating the necessary viewers and displaying them. This has two advantages: firstly it offers a mechanism

for disposing views that are not needed any more without actually closing the resource and secondly it provides a way to refresh the view of a resource in the situations where it becomes corrupted.

- The DataStore viewer now allows multiple selections. This lets users load or delete an arbitrarily large number of resources in one operation.
- The Corpus editor has been completely overhauled. It now allows re-ordering of documents as well as sorting the document list by either index or document name.
- Support has been added for resource parameters of type `gate.FeatureMap`, and it is also possible to specify a default value for parameters whose type is `Collection`, `List` or `Set`. See Section 7.3 for details.
- (Feature Request #1446642) After several requests, a mechanism has been added to allow overriding of GATE's document format detection routine. A new creation-time parameter `contentType` has been added to the standard document implementation, which forces a document to be interpreted as a specific MIME type and prevents the usual detection based on file name extension and other information. See Section 5.5.1 for details.
- A capability has been added to specify arbitrary sets of additional features on individual gazetteer entries. These features are passed forward into the Lookup annotations generated by the gazetteer. See Section 6.3 for details.
- As an alternative to the Google plugin, a new plugin called *yahoo* has been added to allow users to submit their query to the Yahoo search engine and to load the found pages as GATE documents. See Section 19.7 for more details.
- It is now easier to run a corpus pipeline over a single document in the GATE Developer GUI – documents now provide a right-click menu item to create a singleton corpus containing just this document. See Section 3.3 for details.
- A new interface has been added that lets PRs receive notification at the start and end of execution of their containing controller. This is useful for PRs that need to do cleanup or other processing after a whole corpus has been processed. See Section 4.4 for details.
- The GATE Developer GUI does not call `System.exit()` any more when it is closed. Instead an effort is made to stop all active threads and to release all GUI resources, which leads to the JVM exiting gracefully. This is particularly useful when GATE is embedded in other systems as closing the main GATE window will not kill the JVM process any more.
- The set of AnnotationSchemas that used to be included in the core `gate.jar` and loaded as builtins have now been moved to the ANNIE plugin. When the plugin is loaded, the default annotation schemas are instantiated automatically and are available when doing manual annotation.

- There is now support in creole.xml files for automatically creating instances of a resource that are hidden (i.e. do not show in the GUI). One example of this can be seen in the creole.xml file of the ANNIE plugin where the default annotation schemas are defined.
- A couple of helper classes have been added to assist in using GATE within a Spring application. Section 7.14 explains the details.
- Improvements have been made to the thread-safety of some internal components, which mean that it is now safe to create resources in multiple threads (though it is not safe to use the same resource instance in more than one thread). This is a big advantage when using GATE in a multithreaded environment, such as a web application. See Section 7.13 for details.
- Plugins can now provide custom icons for their PRs and LRs in the plugin JAR file. See Section 7.11 for details.
- It is now possible to override the default location for the saved session file using a system property. See Section 2.3 for details.
- The TreeTagger plugin ('Tagger_TreeTagger') supports a system property to specify the location of the shell interpreter used for the tagger shell script. In combination with Cygwin this makes it much easier to use the tagger on Windows. See Section 17.3 for details.
- The Buchart plugin has been removed. It is superseded by SUPPLE, and instructions on how to upgrade your applications from Buchart to SUPPLE are given in Section 17.11. The probability finder plugin has also been removed, as it is no longer maintained.
- The bootstrap wizard now creates a basic plugin that builds with Ant. Since a Unix-style make command is no longer required this means that the generated plugin will build on Windows without needing Cygwin or MinGW.
- The GATE source code has moved from CVS into Subversion. See Section 2.2.3 for details of how to check out the code from the new repository.
- An optional parameter, keepOriginalMarkupsAS, has been added to the DocumentReset PR which allows users to decide whether to keep the Original Markups AS or not while resetting the document. See Section 6.1 for more details.

A.5.3 Bug Fixes and Optimizations

- The Morphological Analyser has been optimized. A new FSM based, although with minor alteration to the basic FSM algorithm, has been implemented to optimize the

Morphological Analyser. The previous profiling figures show that the morpher when integrated with ANNIE application used to take upto 60% of the overall processing time. The optimized version only takes 7.6% of the total processing time. See Section 17.8 for more details on the morpher.

- The ANNIE Sentence Splitter was optimised. The new version is about twice as fast as the previous one. The actual speed increase varies widely depending on the nature of the document.
- The implementation of the *OrthoMatcher* component has been improved. This resources takes significantly less time on large documents.
- The implementation of `AnnotationSets` has been improved. GATE now requires up to 40% less memory to run and is also 20% faster on average. The `get` methods of `AnnotationSet` return instances of `ImmutableAnnotationSet`. Any attempt at modifying the content of these objects will trigger an `Exception`. An empty `ImmutableAnnotationSet` is returned instead of `null`.
- The Chemistry tagger (Section 17.5) has been updated with a number of bugfixes and improvements.
- The Document user interface has been optimised to deal better with large bursts of events which tend to occur when the document that is currently displayed gets modified. The main advantages brought by this new implementation are:
 - The document UI refreshes faster than before.
 - The presence of the GUI for a document induces a smaller performance penalty than it used to. Due to a better threading implementation, machines benefiting from multiple CPUs (e.g. dual CPU, dual core or hyperthreading machines) should only see a negligible increase in processing time when a document is displayed compared to the situations where the document view is not shown. In the previous version, displaying a document while it was processed used to increase execution time by an order of magnitude.
 - The GUI is more responsive now when a large number of annotations are displayed, hidden or deleted.
 - The strange exceptions that used to occur occasionally while working with the document GUI should not happen any more.

And as always there are many smaller bugfixes too numerous to list here...

A.6 Version 3.1 (April 2006)

A.6.1 Major New Features

Support for UIMA

UIMA (<http://www.research.ibm.com/UIMA/>) is a language processing framework developed by IBM. UIMA and GATE share some functionality but are complementary in most respects. GATE now provides an interoperability layer to allow UIMA applications to include GATE components in their processing and vice-versa. For full information, see Chapter18.

New Ontology API

The ontology layer has been rewritten in order to provide an abstraction layer between the model representation and the tools used for input and output of the various representation formats. An implementation that uses Jena 2 (<http://jena.sourceforge.net/ontology>) for reading and writing OWL and RDF(S) is provided.

Ontotext Japec Compiler

Japec is a compiler for JAPE grammars developed by Ontotext Lab. It has some limitations compared to the standard JAPE transducer implementation, but can run JAPE grammars up to five times as fast. By default, GATE still uses the stable JAPE implementation, but if you want to experiment with Japec, see Section 19.10.

A.6.2 Other New Features and Improvements

- Addition of a new JAPE matching style ‘all’. This is similar to Brill, but once all rules from a given start point have matched, the matching will continue from the next offset to the current one, rather than from the position in the document where the longest match finishes. More details can be found in Section 8.4.
- Limited support for loading PDF and Microsoft Word document formats. Only the text is extracted from the documents, no formatting information is preserved.
- The Buchart parser has been deprecated and replaced by a new plugin called SUPPLE - the Sheffield University Prolog Parser for Language Engineering. Full details, including information on how to move your application from Buchart to SUPPLE, is in Section 17.11.

- The Hepple POS Tagger is now open-source. The source code has been included in the GATE Developer/Embedded distribution, under `src/hepple/postag`. More information about the POS Tagger can be found in Section 6.6.
- Minipar is now supported on Windows. `minipar-windows.exe`, a modified version of `pdemo.cpp` is added under the `gate/plugins/Parser_Minipar` directory to allow users to run Minipar on windows platform. While using Minipar on Windows, this binary should be provided as a value for `miniparBinary` parameter. For full information on Minipar in GATE, see Section 17.9.
- The `XmlGateFormat` writer (Save As Xml from GATE Developer GUI, `gate.Document.toXml()` from GATE Embedded API) and reader have been modified to write and read GATE annotation IDs. For backward compatibility reasons the old reader has been kept. This change fixes a bug which manifested in the following situation: If a GATE document had annotations carrying features of which values were numbers representing other GATE annotation IDs, after a save and a reload of the document to and from XML, the former values of the features could have become invalid by pointing to other annotations. By saving and restoring the GATE annotation ID, the former consistency of the GATE document is maintained. For more information, see Section 5.5.2.
- The NP chunker and chemistry tagger plugins have been updated. Mark Greenwood has relicenced them under the LGPL, so their source code has been moved into the GATE Developer/Embedded distribution. See Sections 17.2 and 17.5 for details.
- The Tree Tagger wrapper has been updated with an option to be less strict when characters that cannot be represented in the tagger's encoding are encountered in the document. Details are in Section 17.3.
- JAPE Transducers can be serialized into binary files. The option to load serialized version of JAPE Transducer (an init-time parameter `binaryGrammarURL`) is also implemented which can be used as an alternative to the parameter `grammarURL`. More information can be found in Section 8.9.
- On Mac OS, GATE Developer now behaves more 'naturally'. The application menu items and keyboard shortcuts for *About* and *Preferences* now do what you would expect, and exiting GATE Developer with command-Q or the *Quit* menu item properly saves your options and current session.
- Updated versions of Weka(3.4.6) and Maxent(2.4.0).
- Optimisation in `gate.creole.ml`: the conversion of `AnnotationSet` into ML examples is now faster.
- It is now possible to create your own implementation of `Annotation`, and have GATE use this instead of the default implementation. See `AnnotationFactory` and `AnnotationSetImpl` in the `gate.annotation` package for details.

A.6.3 Bug Fixes

- The Tree Tagger wrapper has been updated in order to run under Windows. See 17.3.
- The SUPPLE parser has been made more user-friendly. It now produces more helpful error messages if things go wrong. Note that you will need to update any saved applications that include SUPPLE to work with this version - see Section 17.11 for details.
- Miscellaneous fixes in the Ontotext JapeC compiler.
- Optimization : the creation of a Document is much faster.
- Google plugin: The optional pagesToExclude parameter was causing a NullPointerException when left empty at run time. Full details about the plugin functionality can be found in Section 19.6.
- Minipar, SUPPLE, TreeTagger: These plugins that call external processes have been fixed to cope better with path names that contain spaces. Note that some of the external tools themselves still have problems handling spaces in file names, but these are beyond our control to fix. If you want to use any of these plugins, be sure to read the documentation to see if they have any such restrictions.
- When using a non-default location for GATE configuration files, the configuration data is saved back to the correct location when GATE exits. Previously the default locations were always used.
- Jape Debugger: ConcurrentModificationException in JAPE debugger. The JAPE debugger was generating a ConcurrentModificationException during an attempt to run ANNIE. There is no exception when running without the debugger enabled. As result of fixing one unnecessary and incorrect callback to debugger was removed from SinglePhaseTransducer class.
- Plus many other small bugfixes...

A.7 January 2005

Release of version 3.

New plugins for processing in various languages (see 19.1). These are not full IE systems but are designed as starting points for further development (French, German, Spanish, etc.), or as sample or toy applications (Cebuano, Hindi, etc.).

Other new plugins:

- Chemistry Tagger 17.5

- Montreal Transducer (since retired)
- RASP Parser 17.10
- MiniPar 17.9
- Buchart Parser 17.11
- MinorThird (Version 5.1: removed)
- NP Chunker 17.2
- Stemmer 17.7
- TreeTagger 17.3
- Probability Finder
- Crawler 19.5
- Google PR 19.6

Support for SVM Light, a support vector machine implementation, has been added to the machine learning plugin ‘Learning’ (see section 15.3.5).

A.8 December 2004

GATE no longer depends on the Sun Java compiler to run, which means it will now work on any Java runtime environment of at least version 1.4. JAPE grammars are now compiled using the Eclipse JDT Java compiler by default.

A welcome side-effect of this change is that it is now much easier to integrate GATE-based processing into web applications in Tomcat. See Section 7.15 for details.

A.9 September 2004

GATE applications are now saved in XML format using the XStream library, rather than by using native java serialization. On loading an application, GATE will automatically detect whether it is in the old or the new format, and so applications in both formats can be loaded. However, older versions of GATE will be unable to load applications saved in the XML format. (A `java.io.StreamCorruptedException: invalid stream header exception` will occur.) It is possible to get new versions of GATE to use the old format by setting a flag in the source code. (See the `Gate.java` file for details.) This change has been made because it allows the details of an application to be viewed and edited in a text editor, which is sometimes easier than loading the application into GATE.

A.10 Version 3 Beta 1 (August 2004)

Version 3 incorporates a lot of new functionality and some reorganisation of existing components.

Note that Beta 1 is feature-complete but needs further debugging (please send us bug reports!).

Highlights include: completely rewritten document viewer/editor; extensive ontology support; a new plugin management system; separate .jar files and a Tomcat classloading fix; lots more CREOLE components (and some more to come soon).

Almost all the changes are backwards-compatible; some recent classes have been renamed (particularly the ontologies support classes) and a few events added (see below); datastores created by version 3 will probably not read properly in version 2. If you have problems use the mailing list and we'll help you fix your code!

The gorey details:

- Anonymous CVS is now available. See Section 2.2.3 for details.
- CREOLE repositories and the components they contain are now managed as plugins. You can select the plugins the system knows about (and add new ones) by going to 'Manage CREOLE Plugins' on the file menu.
- The `gate.jar` file no longer contains all the subsidiary libraries and CREOLE component resources. This makes it easier to replace library versions and/or not load them when not required (libraries used by CREOLE builtins will now not be loaded unless you ask for them from the plugins manager console).
- ANNIE and other bundled components now have their resource files (e.g. pattern files, gazetteer lists) in a separate directory in the distribution – `gate/plugins`.
- Some testing with Sun's JDK 1.5 pre-releases has been done and no problems reported.
- The `gate://` URL system used to load CREOLE and ANNIE resources in past releases is no longer needed. This means that loading in systems like Tomcat is now much easier.
- MAC OS X is now properly supported by the installed and the runtime.
- An Ontology-based Corpus Annotation Tool (OCAT) has been implemented as a plugin. Documentation of its functionality is in Section 14.6.
- The NLG Lexical tools from the MIAKT project have now been released.
- The Features viewer/editor has been completely updated – see Section 3.4.5 for details.

- The Document editor has been completely rewritten – see Section 3.2 for more information.
- The datastore viewer is now a full-size VR – see Section 3.8.2 for more information.

A.11 July 2004

GATE documents now fire events when the document content is edited. This was added in order to support the new facility of editing documents from the GUI. This change will break backwards compatibility by requiring all `DocumentListener` implementations to implement a new method:

```
public void contentEdited(DocumentEvent e);
```

A.12 June 2004

A new algorithm has been implemented for the `AnnotationDiff` function. A new, more usable, GUI is included, and an ‘Export to HTML’ option added. More details about the `AnnotationDiff` tool are in Section 10.2.1.

A new build process, based on ANT (<http://ant.apache.org/>) is now available. The old build process, based on make, is now unsupported. See Section 2.5 for details of the new build process.

A Jape Debugger from Ontos AG has been integrated. You can turn integration ON with command line option ‘-j’. If you run GATE Developer with this option, the new menu item for Jape Debugger GUI will appear in the Tools menu. The default value of integration is OFF. We are currently awaiting documentation for this.

NOTE! Keep in mind there is `ClassCastException` if you try to debug `ConditionalCorpusPipeline`. Jape Debugger is designed for `Corpus Pipeline` only. The Ontos code needs to be changed to allow debugging of `ConditionalCorpusPipeline`.

A.13 April 2004

There are now two alternative strategies for ontology-aware grammar transduction:

- using the `[ontology]` feature both in grammars and annotations; with the default `Transducer`.

- using the ontology aware transducer – passing an ontology LR to a new subsume method in the SimpleFeatureMapImpl. the latter strategy does not check for ontology features (this will make the writing of grammars easier – no need to specify ontology).

The changes are in:

- SinglePhaseTransducer (always call subsume with ontology – if null then the ordinary subsumption takes place)
- SimpleFeatureMapImpl (new subsume method using an ontology LR)

More information about the ontology-aware transducer can be found in Section 14.9.

A morphological analyser PR has been added. This finds the root and affix values of a token and adds them as features to that token.

A flexible gazetteer PR has been added. This performs lookup over a document based on the values of an arbitrary feature of an arbitrary annotation type, by using an externally provided gazetteer. See 13.7 for details.

A.14 March 2004

Support was added for the MAXENT machine learning library. (See 15.3.4 for details.)

A.15 Version 2.2 – August 2003

Note that GATE 2.2 works with JDK 1.4.0 or above. Version 1.4.2 is recommended, and is the one included with the latest installers.

GATE has been adapted to work with Postgres 7.3. The compatibility with PostgreSQL 7.2 has been preserved.

Note that as of Version 5.1 PostgreSQL is no longer supported.

New library version – Lucene 1.3 (rc1)

A bug in gate.util.Javac has been fixed in order to account for situations when String literals require an encoding different from the platform default.

Temporary .java files used to compile JAPE RHS actions are now saved using UTF-8 and the ‘-encoding UTF-8’ option is passed to the javac compiler.

A custom tools.jar is no longer necessary

Minor changes have been made to the look and feel of GATE Developer to improve its appearance with JDK 1.4.2

Some bug fixes (087, 088, 089, 090, 091, 092, 093, 095, 096 – see <http://gate.ac.uk/gate/doc/bugs.html> for more details).

A.16 Version 2.1 – February 2003

Integration of Machine Learning PR and WEKA wrapper (see Section 15.3).

Addition of DAML+OIL exporter.

Integration of WordNet (see Section 19.8).

The syntax tree viewer has been updated to fix some bugs.

A.17 June 2002

Conditional versions of the controllers are now available (see Section 3.7.2). These allow processing resources to be run conditionally on document features.

PostgreSQL Datastores are now supported.

These store data into a PostgreSQL RDBMS.

(As of Version 5.1 PostgreSQL is no longer supported.)

Addition of OntoGazetteer (see Section 13.4), an interface which makes ontologies visible within GATE Developer, and supports basic methods for hierarchy management and traversal.

Integration of Protégé, so that people with developed Protégé ontologies can use them within GATE.

Addition of IR facilities in GATE (see Section 19.4).

Modification of the corpus benchmark tool (see Section 10.4.3), which now takes an application as a parameter.

See also for details of other recent bug fixes.

Appendix B

Version 5.1 Plugins Name Map

In version 5.1 we attempted to impose order on chaos by further defining the plugin naming convention (see Section 12.3.2) and renaming those existing plugins that did not conform to it. Below, you will find a mapping of old plugin names to new.

Old Name	New Name
abner	Tagger_Abner
alignment	Alignment
annotationMerging	Annotation_Merging
arabic	Lang_Arabic
bdmComputation	Ontology_BDM_Computation
cebuano	Lang_Cebuano
Chemistry_Tagger	Tagger_Chemistry
chinese	Lang_Chinese
chineseSegmenter	Lang_Chinese
copyAS2AnoDoc	Copy_Annots_Between_Docs
crawl	Web_Crawler_Websphinx
french	Lang_French
german	Lang_German
google	Web_Search_Google
hindi	Lang_Hindi
iaaPlugin	Inter_Annotator_Agreement
italian	Lang_Italian
Kea	Keyphrase_Extraction_Algorithm
learning	Learning
lkb_gazetteer	Gazetteer_LKB
Minipar	Parser_Minipar
NP_Chunking	Tagger_NP_Chunking
Ontology_Based_Gazetteer	Gazetteer_Ontology_Based
OpenCalais	Tagger_OpenCalais
openNLP	OpenNLP
rasp	Parser_RASP
romanian	Lang_Romanian
Stanford	Parser_Stanford
Stemmer	Stemmer_Snowball
SUPPLE	Parser_SUPPLE
TaggerFramework	Tagger_Framework
TreeTagger	Tagger_TreeTagger
uima	UIMA
yahoo	Web_Search_Yahoo

Appendix C

Design Notes

Why has the pleasure of slowness disappeared? Ah, where have they gone, the amblers of yesteryear? Where have they gone, those loafing heroes of folk song, those vagabonds who roam from one mill to another and bed down under the stars? Have they vanished along with footpaths, with grasslands and clearings, with nature? There is a Czech proverb that describes their easy indolence by a metaphor: ‘they are gazing at God’s windows.’ A person gazing at God’s windows is not bored; he is happy. In our world, indolence has turned into having nothing to do, which is a completely different thing: a person with nothing to do is frustrated, bored, is constantly searching for an activity he lacks.

Slowness, Milan Kundera, 1995 (pp. 4-5).

GATE is a backplane into which specialised Java Beans plug. These beans are loose-coupled with respect to each other - they communicate entirely by means of the GATE framework. Inter-component communication is handled by model components - LanguageResources, and events.

Components are defined by conformance to various interfaces (e.g. LanguageResource), ensuring separation of interface and implementation.

The reason for adding to the normal bean initialisation mech is that LRs, PRs and VRs all have characteristic parameterisation phases; the GATE resources/components model makes explicit these phases.

C.1 Patterns

GATE is structured around a number of what we might call principles, or patterns, or alternatively, clever ideas stolen from better minds than mine. These patterns are:

- modelling most things as extensible sets of components (cf. Section C.1.1);
- separating components into model, view, or controller (cf. Section C.1.2) types;
- hiding implementation behind interfaces (cf. Section C.1.3).

Four interfaces in the top-level package describe the GATE view of components: `Resource`, `ProcessingResource`, `LanguageResource` and `VisualResource`.

C.1.1 Components

Architectural Principle

Wherever users of the architecture may wish to extend the set of a particular type of entity, those types should be expressed as components.

Another way to express this is to say that the architecture is based on *agents*. I've avoided this in the past because of an association between this term and the idea of bits of code moving around between machines of their own volition. I take this to be somewhat pointless, and probably the result of an anthropomorphic obsession with mobility as a correlate of intelligence. If we drop this connotation, however, we can say that GATE is an agent-based architecture. If we want to, that is.

Framework Expression

Many of the classes in the framework are components, by which we mean classes that conform to an interface with certain standard properties. In our case these properties are based on the Java Beans component architecture, with the addition of component metadata, automated loading and standardised storage, threading and distribution.

All components inherit from `Resource`, via one of the three sub-interfaces `LanguageResource` (LR), `VisualResource` (VR) or `ProcessingResource` (PR). `VisualResources` (VRs) are straight-forward – they represent visualisation and editing components that participate in GUIs – but the distinction between language and processing resources merits further discussion.

Like other software, LE programs consist of data and algorithms. The current orthodoxy in software development is to model both data and algorithms together, as *objects*¹. Systems that adopt the new approach are referred to as Object-Oriented (OO), and there are good reasons to believe that OO software is easier to build and maintain than other varieties [Booch 94, Yourdon 96].

¹Older development methods like Jackson Structured Design [Jackson 75] or Structured Analysis [Yourdon 89] kept them largely separate.

In the domain of human language processing R&D, however, the terminology is a little more complex. Language data, in various forms, is of such significance in the field that it is frequently worked on independently of the algorithms that process it. For example: a treebank² can be developed independently of the parsers that may later be trained from it; a thesaurus can be developed independently of the query expansion or sense tagging mechanisms that may later come to use it. This type of data has come to have its own term, *Language Resources* (LRs) [LREC-1 98], covering many data sources, from lexicons to corpora.

In recognition of this distinction, we will adopt the following terminology:

Language Resource (LR): refers to data-only resources such as lexicons, corpora, thesauri or ontologies. Some LRs come with software (e.g. Wordnet has both a user query interface and C and Prolog APIs), but where this is only a means of accessing the underlying data we will still define such resources as LRs.

Processing Resource (PR): refers to resources whose character is principally programmatic or algorithmic, such as lemmatisers, generators, translators, parsers or speech recognisers. For example, a part-of-speech tagger is best characterised by reference to the process it performs on text. PRs typically *include* LRs, e.g. a tagger often has a lexicon; a word sense disambiguator uses a dictionary or thesaurus.

Additional terminology worthy of note in this context: *language data* refers to LRs which are at their core examples of language in practice, or ‘performance data’, e.g. corpora of texts or speech recordings (possibly including added descriptive information as markup); *data about language* refers to LRs which are purely descriptive, such as a grammar or lexicon.

PRs can be viewed as algorithms that map between different types of LR, and which typically use LRs in the mapping process. An MT engine, for example, maps a monolingual corpus into a multilingual aligned corpus using lexicons, grammars, etc.³

Further support for the PR/LR terminology may be gleaned from the argument in favour of declarative data structures for grammars, knowledge bases, etc. This argument was current in the late 1980s and early 1990s [Gazdar & Mellish 89], partly as a response to what has been seen as the overly procedural nature of previous techniques such as augmented transition networks. Declarative structures represent a separation between data about language and the algorithms that use the data to perform language processing tasks; a similar separation to that used in GATE.

Adopting the PR/LR distinction is a matter of conforming to established domain practice and terminology. It does not imply that we cannot model the domain (or build software to support it) in an Object-Oriented manner; indeed the models in GATE are themselves Object-Oriented.

²A corpus of texts annotated with syntactic analyses.

³This point is due to Wim Peters.

C.1.2 Model, view, controller

According to Buschmann et al (Pattern-Oriented Software Architecture, 1996), the Model-View-Controller (MVC) pattern

...divides an interactive application into three components. The model contains the core functionality and data. Views display information to the user. Controllers handle user input. Views and controllers together comprise the user interface. A change-propagation mechanism ensures consistency between the user interface and the model. [p.125]

A variant of MVC, the Document-View pattern,

...relaxes the separation of view and controller... The View component of Document-View combines the responsibilities of controller and view in MVC, and implements the user interface of the system.

A benefit of both arrangements is that

...loose coupling of the document and view components enables multiple simultaneous synchronized but different views of the same document.

Geary (Graphic Java 2, 3rd Edtn., 1999) gives a slightly different view:

MVC separates applications into three types of objects:

- **Models:** Maintain data and provide data accessor methods
- **Views:** Paint a visual representation of some or all of a model's data
- **Controllers:** Handle events ... By encapsulating what other architectures intertwine, MVC applications are much more flexible and reusable than their traditional counterparts.

[pp. 71, 75]

Swing, the Java user interface framework, uses

a specialised version of the classic MVC meant to support pluggable look and feel instead of applications in general. [p. 75]

GATE may be regarded as an MVC architecture in two ways:

- directly, because we use the Swing toolkit for the GUIs;
- by analogy, where LR's are models, VR's are views and PR's are controllers. Of these, the latter sits least easily with the MVC scheme, as PR's may indeed be controllers but may also not be.

C.1.3 Interfaces

Architectural Principle

The implementation of types should generally be hidden from the clients of the architecture.

Framework Expression

With a few exceptions (such as for utility classes), clients of the framework work with the `gate.*` package. This package is mostly composed of interface definitions. Instantiations of these interfaces are obtained via the `Factory` class.

The subsidiary packages of GATE provide the implementations of the `gate.*` interfaces that are accessed via the factory. They themselves avoid directly constructing classes from other packages (with a few exceptions, such as JAPE's need for unattached annotation sets). Instead they use the factory.

C.2 Exception Handling

When and how to use exceptions? Borrowing from Bill Venners, here are some **guidelines** (with examples):

1. Exceptions exist to refer problem conditions up the call stack to a level at which they may be dealt with. "If your method encounters an abnormal condition *that it can't handle*, it should throw an exception." If the method can handle the problem rationally, it should catch the exception and deal with it.

Example:

If the creation of a resource such as a document requires a URL as a parameter, the method that does the creation needs to construct the URL and read from it. If there is an exception during this process, the GATE method should abort by throwing its own exception. The exception will be dealt with higher up the food chain, e.g. by asking the user to input another URL, or by aborting a batch script.

2. All GATE exceptions should inherit from `gate.util.GateException` (a descendant of `java.lang.Exception`, hence a checked exception) or `gate.util.GateRuntimeException` (a descendant of `java.lang.RuntimeException`, hence an unchecked exception). This rule means that clients of GATE code can catch all sorts of exceptions thrown by the system with only two catch statements. (This rule may be broken by methods that are not public, so long as their callers catch the non-GATE exceptions and deal with them or convert them to `GateException/GateRuntimeException`.) Almost **all** exceptions thrown by GATE should be checked exceptions: the point of an exception is that clients of your code get to know about it, so use a checked exception to make the compiler force them to deal with it. Except:

Example:

With reference to the previous example, a problem using the URL will be signalled by something like an `UnknownHostException` or an `IOException`. These should be caught and re-thrown as descendants of `GateException`.

3. In a situation where an exceptional condition is an indication of a bug in the GATE library, or in the implementation of some other library, then it is permissible to throw an unchecked exception.

Example:

If a method is creating annotations on a document, and before creating the annotations it checks that their start and end points are valid ranges in relation to the content of the document (i.e. they fall within the offset space of the document, and the end is after the start), then if the method receives an `InvalidOffsetException` from the `AnnotationSet.add` call, something is seriously wrong. In such cases it may be best to throw a `GateRuntimeException`.

4. Where you are inheriting from a non-GATE class and therefore have the exception signatures fixed for you, you may add a new exception deriving from a non-GATE class.

Example:

The SAX XML parser API uses `SaxException`. Implementing a SAX parser for a document type involves overriding methods that throw this exception. Where you want to have a subtype for some problem which is specific to GATE processing, you could use `GateSaxException` which extends `SaxException`.

5. Test code is different: in the JUnit test cases it is fine just to declare that each method throws `Exception` and leave it at that. The JUnit test runner will pick up the exceptions and report them to you. Test methods should, however, try and ensure that the exceptions thrown are meaningful. For example, avoid null pointer exceptions in the test code itself, e.g. by using `assertNonNull`.

Example:

```

1  public void testComments() throws Exception {
2      ResourceData docRd = (ResourceData) reg.get("gate.Document");
3      assertNotNull("testComments: couldn't find document res data", docRd);
4      String comment = docRd.getComment();
5      assert(
6          "testComments: incorrect or missing COMMENT on document",
7          comment != null && comment.equals("GATE document")
8      );
9  } // testComments()

```

See also the testing notes.

6. "Throw a different exception type for each abnormal condition." You can go too far on this one - a hundred exception types per package would certainly be too much - but in general you should create a new exception type for each different sort of problem you encounter.

Example:

The `gate.creole` package has a `ResourceInstantiationException` - this deals with all problems to do with creating resources. We could have had `"ResourceUrlProblem"` and `"ResourceParameterProblem"` but that would probably have ended up with too many. On the other hand, just throwing everything as `GateException` is too coarse (Hamish take note!).

7. Put exceptions in the package that they're thrown from (unless they're used in many packages, in which case they can go in `gate.util`). This makes it easier to find them in the documentation and prevents name clashes.

Example:

`gate.jape.ParserException` is correctly placed; if it was in `gate.util` it might clash with, for example, `gate.xml.ParserException` if there was such.

Appendix D

JAPE: Implementation

The annual Diagram prize for the oddest book title of the year has been awarded to Gerard Forlin's Butterworths Corporate Manslaughter Service, a hefty law tome providing guidance and analysis on corporate liability for deaths in the workplace.

The book, not published until January, was up against five other shortlisted titles: Fancy Coffins to Make Yourself; The Flat-Footed Flies of Europe; Lightweight Sandwich Construction; Tea Bag Folding; and The Art and Craft of Pounding Flowers: No Paint, No Ink, Just a Hammer! The shortlist was thrown open to readers of the literary trade magazine The Bookseller, who chose the winner by voting on the magazine's website. Butterworths Corporate Manslaughter Service, a snip at 375, emerged as the overall victor with 35% of the vote.

The Diagram prize has been a regular on the award circuit since 1978, when Proceedings of the Second International Workshop on Nude Mice carried off the inaugural award. Since then, titles such as American Bottom Archaeology and last year's winner, High-Performance Stiffened Structures (an engineering publication), have received unwonted publicity through the prize. This year's winner is perhaps most notable for its lack of *entendre*.

Manslaughter Service kills off competition in battle of strange titles, Emma Yates, The Guardian, November 30, 2001.

This appendix gives implementation details and formal definitions of the JAPE annotation patterns language. Section D.1 gives a more formal definition of the JAPE grammar, and some examples of its use. Section D.2 describes JAPE's relation to CPSL. Section D.3 describes the initialisation of a JAPE grammar, Section D.4 talks about the execution of JAPE grammars, and the final section explains how to switch the Java compiler used for JAPE.

D.1 Formal Description of the JAPE Grammar

JAPE is similar to CPSL (a Common Pattern Specification Language, developed in the TIPSTER programme by Doug Appelt and others), with a few exceptions. Figure D.1 gives a BNF (Backus-Naur Format) description of the grammar.

An example rule LHS:

```
Rule: KiloAmount
( ({Token.kind == "containsDigitAndComma"}):number
  {Token.string == "kilograms"} ):whole
```

A basic constraint specification appears between curly braces, and gives a conjunction of annotation/attribute/value specifiers which have to match at a particular point in the annotation graph. A complex constraint specification appears within round brackets, and may be bound to a label with the ‘:’ operator; the label then becomes available in the RHS for access to the annotations matched by the complex constraint. Complex constraints can also have Kleene operators (*, +, ?) applied to them. A sequence of constraints represents a sequential conjunction; disjunction is represented by separating constraints with ‘|’.

Converted to the format accepted by the JavaCC LL parser generator, the most significant fragment of the CPSL grammar (as described by Appelt, based on an original specification from a TIPSTER working group chaired by Boyan Onyshkevych) goes like this:

```
constraintGroup -->
    (patternElement)+ ("|" (patternElement)+ )*

patternElement -->
    "{" constraint ("," constraint)* "}"
|   "(" constraintGroup ")" (kleeneOp)? (binding)?
```

Here the first line of `patternElement` is a basic constraint, the second a complex one.

```

MultiPhaseTransducer ::=
  ( <multiphase> <ident> )?
  ( ( SinglePhaseTransducer )+ | ( <phases> ( <ident> )+ ) )
  <EOF>
SinglePhaseTransducer ::=
  <phase> <ident> ( <input> ( <ident> )* )?
  ( <option> ( <ident> <assign> <ident> )* )?
  ( ( Rule ) | MacroDef )*
Rule ::=
  <rule> <ident> ( <priority> <integer> )?
  LeftHandSide "-->" RightHandSide
MacroDef ::=
  <macro> <ident> ( PatternElement | Action )
LeftHandSide ::=
  ConstraintGroup
ConstraintGroup ::=
  ( PatternElement )+ ( <bar> ( PatternElement )+ )*
PatternElement ::=
  ( <ident> | BasicPatternElement | ComplexPatternElement )
BasicPatternElement ::=
  ( ( <leftBrace> Constraint ( <comma> Constraint )* <rightBrace> )
  | ( <string> ) )
ComplexPatternElement ::=
  <leftBracket> ConstraintGroup <rightBracket>
  ( <kleeneOp> )? ( <colon> ( <ident> | <integer> ) )?
Constraint ::=
  ( <pling> )? <ident> ( <period> <ident> <equals> AttrVal )?
AttrVal ::=
  ( <string> | <ident> | <integer> | <floatingPoint> | <bool> )
RightHandSide ::=
  Action ( <comma> Action )*
Action ::=
  ( NamedJavaBlock | AnonymousJavaBlock | AssignmentExpression | <ident> )
NamedJavaBlock ::=
  <colon> <ident> <leftBrace> ConsumeBlock
AnonymousJavaBlock ::=
  <leftBrace> ConsumeBlock
AssignmentExpression ::=
  ( <colon> | <colonplus> ) <ident> <period> <ident>
  <assign> <leftBrace> (
    <ident> <assign>
    ( AttrVal | ( <colon> <ident> <period> <ident> <period> <ident> ) )
    ( <comma> )?
  )* <rightBrace>
ConsumeBlock ::=
  Java code

```

Figure D.1: BNF of JAPE's grammar

An example of a complete rule:

```
Rule: NumbersAndUnit
( ( {Token.kind == "number"} )+:numbers {Token.kind == "unit"} )
-->
:numbers.Name = { rule = "NumbersAndUnit" }
```

This says ‘match sequences of numbers followed by a unit; create a Name annotation across the span of the numbers, and attribute rule with value NumbersAndUnit’.

D.2 Relation to CPSL

We *differ from the CPSL spec* in various ways:

1. No pre- or post-fix context is allowed on the LHS.
2. No function calls on the LHS.
3. No string shorthand on the LHS.
4. We have multiple rule application algorithms (see Section 8.4).
5. Expressions relating to labels unbound on the LHS are not evaluated on the RHS. (In TextPro they evaluate to ‘false’.)
6. JAPE allows arbitrary Java code on the RHS.
7. JAPE has a different macro syntax, and allows macros for both the RHS and LHS.
8. JAPE grammars are compiled and can be stored as serialised Java objects.

Apart from this, it is a full implementation of CPSL, and the formal power of the languages is the same (except that a JAPE RHS can delete annotations, which straight CPSL cannot). The rule LHS is a regular language over annotations; the rule RHS can perform arbitrary transformations on annotations, but the RHS is only fired *after* the LHS been evaluated, and the effects of a rule application can only be referenced after the phase in which it occurs, so the recognition power is no more than regular.

D.3 Initialisation of a JAPE Grammar

When a JAPE grammar is loaded in GATE, each phase is converted into a finite state machine, a process that has several stages. Each rule is treated as a regular expression using annotation-based patterns as input symbols. A JAPE phase is a disjunction of rules, so it is also a regular expression. The first stage of building the associated FSM for a JAPE phase is the construction of a non-deterministic finite-state automaton, following the algorithm described in [Aho *et al.* 86].

Additional to standard regular expressions, JAPE rules also contain bindings (labels associated to pattern segments). These are intended to be associated to the matched input symbols (i.e. annotations) during the matching process, and are used while executing the actions caused by the rule firing. Upon creating the equivalent FSM for a given JAPE rule, bindings are associated with the FSM transitions. This changes the semantics of a transition – besides moving the state machine into a new current state, a transition may also *bind* the consumed annotation(s) with one or more labels.

In order to optimise the execution time during matching (at the expense of storage space), NFAs are usually converted to *Deterministic Finite State Automata* (DFAs) using e.g. the *subset algorithm* [Aho *et al.* 86]. In the case of JAPE this transformation is not possible due to the binding labels: two or more transitions from the NFA that match the same annotation pattern cannot be compacted into a single transition in the DFA if they have different bindings. Because of this, JAPE grammars are represented as non-deterministic finite state machines. A partial optimisation that eliminates the ϵ -transitions from the NFA is however performed.

The actions represented on the right hand side of JAPE rules are converted to compiled Java classes and are associated with final states in the FSM. The final in-memory representation of a JAPE grammar thus consists of a non-deterministic finite state machine, with transitions that use annotation-based patterns as input symbols, additionally marked with bindings information and for which the final states are associated with actions.

Starting from the following two JAPE rules:

```
Rule: PersonPrefix
(
  ({Token})+ {Person}
):pers
--> {...}
```

```
Rule: OrganisationPrefix (
  ({Token})+ {Organisation}
):org --> {...}
```

the associated NFA is constructed, as illustrated in Figure D.2. Note that due to the fact that

the final states are associated with different actions, they cannot be joined into a single one and are kept separate. This automaton is then optimised by eliminating the ϵ -transitions, resulting in the NFA presented in Figure D.3. For the sake of simplicity, the annotation patterns used are the most basic ones, depending solely on annotation type. In the graphical representation, the transitions are marked with the type of annotation that they match and the associated binding in square brackets.




Figure D.2: Example of a non-deterministic finite state machine compiled from JAPE rules.

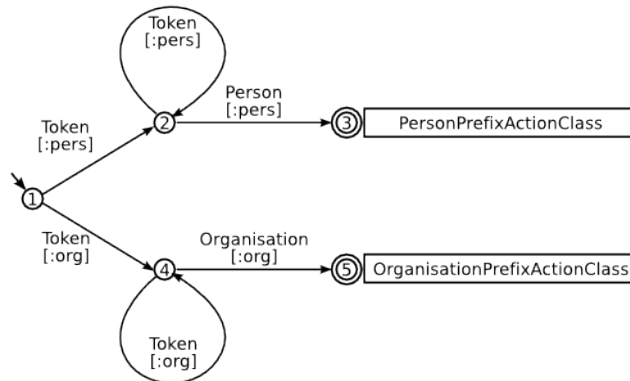


Figure D.3: Example of a finite state machine compiled from JAPE rules, with ϵ -transitions removed.

It can be observed in Figure D.3 that there are two transitions starting from state 1 (leading to states 2, respectively 4) that both consume annotations of type `Token`, thus even the optimised finite state machine is still non-deterministic.

Once a JAPE grammar is converted to the equivalent finite state automaton, the initialisation phase is complete.

D.4 Execution of JAPE Grammars

The execution of a JAPE grammar can be described in simple terms as finding a path through an annotation graph where all the annotations traversed form a sequence that is accepted by the finite state machine built during the initialisation phase. The actual process is somewhat more complex than that, as it also needs to take into account the various matching modes, the filtering of input annotation types, to deal with the assignment of matched annotation to bindings, and to manage the execution of actions whenever successful matches occur.

Executing a JAPE grammar involves simulating the execution of a non-deterministic finite state automaton (NFA) while using an annotation graph as input. At each step we start from a document position (initially zero) and a finite state machine in a given state (initially the start state). Annotations found at the given document position are compared with the restrictions encoded in the NFA transitions; if they match, the annotations are *consumed* and the state machine moves to a new state. Ambiguities are possible at each step both in terms of input (several matching annotations can start at the same offset) and in terms of available NFA transitions (the state machine is non-deterministic, so multiple transitions with the same restrictions can be present). When such ambiguities are encountered, the current state machine is cloned to create as many copies as necessary, and each such copy continues the matching process independently. The JAPE executor thus needs to keep track of a family of state machines that are running in parallel – henceforth we shall call these FSM instances.

Whenever one of the active FSM instances is moved to a new state, a test is performed to check if the new state is a final one. If that is the case, the FSM instance is said to be in an *accepting* state, and a copy of its state is saved for later usage.

When none of the active FSM instances can advance any further, the stored accepting FSM instances are used to execute JAPE actions, according to the declared matching style of the current grammar.

A high-level view¹ of the algorithm used during the execution of a JAPE grammar is presented in Listing D.1, in a Java-inspired pseudo-code.

The next paragraphs contain some more detailed comments, indexed using the line numbers in the listing:

line 1 The annotations present in the document are filtered according to the `Input` declaration in the JAPE code, if one was present. This causes the JAPE executor to completely ignore annotations that are not listed as valid input.

lines 2–4 The matching process is initialised by setting the document position to 0, and creating empty lists of active and accepting FSM instances.

¹The view of the algorithm presented here is greatly simplified, for the sake of clarity. The actual implementation consists of a few thousand lines of Java code.

Listing D.1: JAPE matching algorithm

```

1 processInputfilters();
2 currentDocPosition = 0;
3 activeFSMInstances = new List<FSMInstance>();
4 acceptingFSMInstances = new List<FSMInstance>();
5 while(currentDocPosition < document.length()){
6     //create an initial FSM instance, starting from
7     //the current document position
8     activeFSMInstances.add(
9         new FSMInstance(currentDocPosition));
10    //advance all FSM instances,
11    //until no further advance is possible
12    while(!activeFSMInstances.isEmpty()){
13        FSMInstance aFSM = activeFSMInstances.remove(0);
14        //advance aFSM, consuming annotations, linking used
15        //annotations to binding labels, as required;
16
17        //create cloned copies as necessary and add them to
18        //activeFSMInstances;
19
20        //save any accepting state of aFSM
21        //into acceptingFSMInstances;
22    }
23    if(!acceptingFSMInstances.isEmpty()){
24        //execute the action(s)
25    }
26
27    //move to the new document position, in accordance
28    //with the matching style.
29 }

```

lines 5–29 The matching continues until all the document text is exhausted.

line 8 Each step starts from the current document position with a single FSM instance.

lines 12–22 While there are still active FSM instances, they are advanced as far as possible. Whenever ambiguities are encountered, cloned copies are created and added to the list of active FSM instances. Whenever an FSM instance reaches a final state during its advancing, a copy of its state is saved to the list of accepting FSM instances.

lines 23–25 This segment of code is reached when there are no more active FSM instances – all active instances were advanced as far as possible and either saved to the accepting list (if they reached a final state during that process) or simply discarded (if they could advance no further but they still have not reached a final state).

At this point, any successful matches that occurred need to be acted upon, so the list of accepting FSM instances is inspected. If there are any, their associated actions are now executed, according to the desired matching style. For instance if the matching style used is `Appelt`, then only the accepting FSM instance that has covered the most

input will be executed; conversely, if the matching style is `Brill`, then all accepting FSM instances will have their actions executed, etc.

line 27 When this point is reached, all possible matches from the current document position were found and the required action executed. The next step is to move to the next starting position in the document, and re-start the matching process from there. Depending on the matching style selected, the new document position is either the *oldPosition* + 1, in the case of `All`, or *matchingEndPosition* + 1 in all other cases.

D.5 Using a Different Java Compiler

GATE allows you to choose which Java compiler is used to compile the action classes generated from JAPE rules. The preferred compiler is specified by the `Compiler.type` option in `gate.xml`. At present the supported values are:

Sun The Java compiler supplied with the JDK. Although the option is called `Sun`, it supports any JDK that supplies `com.sun.tools.javac.Main` in a standard location, including the IBM JDK (all platforms) and the Apple JDK for Mac OS X.

Eclipse The Eclipse compiler, from the Java Development Tools of the Eclipse project². Currently we use the compiler from Eclipse 3.2, which supports Java 5.0.

By default, the Eclipse compiler is used. It compiles faster than the Sun compiler, and loads dependencies via the GATE ClassLoader, which means that Java code on the right hand side of JAPE rules can refer to classes that were loaded from a plugin JAR file. The Sun compiler can only load classes from the system classpath, so it will not work if GATE is loaded from a subsidiary classloader, e.g. a Tomcat web application. You should generally use the Eclipse compiler unless you have a compelling reason not to.

Support for other compilers can be added, but this is not documented here - if you're in a position to do this, you won't mind reading the source code...

²<http://www.eclipse.org/jdt>

Appendix E

Ant Tasks for GATE

This chapter describes the Ant tasks provided by GATE that you can use in your own build files. The tasks require Ant 1.7 or later.

E.1 Declaring the Tasks

To use the GATE Ant tasks in your build file you must include the following `<typedef>` (where `${gate.home}` is the location of your GATE installation):

```
<typedef resource="gate/util/ant/antlib.xml">
  <classpath>
    <pathelement location="${gate.home}/bin/gate.jar" />
    <fileset dir="${gate.home}/lib" includes="*.jar" />
  </classpath>
</typedef>
```

If you have problems with library conflicts you should be able to reduce the JAR files included from the lib directory to just `jdom`, `xstream` and `jaxen` (plus `stax-api` and `wstx-lgpl` if you are running on Java 5, but these are not required on Java 6).

E.2 The packagegapp task - bundling an application with its dependencies

E.2.1 Introduction

GATE saved application states (GAPP files) are an XML representation of the state of a GATE application. One of the features of a GAPP file is that it holds references to the external resource files used by the application as paths relative to the location of the GAPP file itself (or relative to the location of the GATE home directory where appropriate). This is useful in many cases but if you want to package up a copy of an application to send to a third party or to use in a web application, etc., then you need to be very careful to save the file in a directory above *all* its resources, and package the resources up with the GAPP file at the same relative paths. If the application refers to resources outside its own file tree (i.e. with relative paths that include `..`) then you must either maintain this structure or manually edit the XML to move the resource references around and copy the files to the right places to match. This can be quite tedious and error-prone...

The `packagegapp` Ant task aims to automate this process. It extracts all the relative paths from a GAPP file, writes a modified version of the file with these paths rewritten to point to locations below the new GAPP file location (i.e. with no `..` path segments) and copies the referenced files to their rewritten locations. The result is a directory structure that can be easily packaged into a zip file or similar and moved around as a self-contained unit.

This Ant task is the underlying driver for the ‘Export for Teamware’ option described in Section 3.8.4. Export for teamware does the equivalent of:

```
<packagegapp src="sourceFile.gapp"
             destfile="{tempdir}/application.xgapp"
             copyPlugins="yes"
             copyResourceDirs="yes"
             onUnresolved="recover" />
```

followed by packaging the temporary directory into a zip file. These options are explained in detail below.

The `packagegapp` task requires Ant 1.7 or later.

E.2.2 Basic Usage

In many cases, the following simple invocation will do what you want:

```
<packagegapp src="original.xgapp"
```



```
gatehome="/path/to/GATE"  
destfile="package/target.xgapp" />
```

Note that the parent directory of the `destfile` (in this case `package`) must already exist. It will not be created automatically. The value for the `gatehome` attribute should be the path to your GATE installation (the directory containing `build.xml`, the `bin`, `lib` and `plugins` directories, etc.). If you know that the `gapp` file you want to package does not reference any resources relative to the GATE home directory¹ then this attribute may be omitted.

This will perform the following steps:

1. Read in the `original.xgapp` file and extract all the relative paths it contains.
2. For each plugin referred to by a relative path, `foo/bar/MyPlugin`, rewrite the plugin location to be `plugins/MyPlugin` (relative to the location of the `destfile`).
3. For each resource file referred to by the `gapp`, see if it lives under the original location of one of the plugins moved in the previous step. If so, rewrite its location relative to the new location of the plugin.
4. If there are any relative resource paths that are not accounted for by the above rule (i.e. they do not live inside a referenced plugin), the build fails (see Section E.2.3 for how to change this behaviour).
5. Write out the modified GAPP to the `destfile`.
6. Recursively copy the whole content of each of the plugins from step 2 to their new locations².

This means that the all the relative paths in the new GAPP file (`package/target.xgapp`) will point to `plugins/Something`. You can now bundle up the whole `package` directory and take it elsewhere.

E.2.3 Handling Non-Plugin Resources

By default, the task only handles relative resource paths that point within one of the plugins that the GAPP refers to. However, many applications refer to resources that live outside the plugin directories, for example custom JAPE grammars, gazetteer lists, etc. The task provides two approaches to support this: it can handle the unresolved references automatically, or you can provide your own ‘hints’ to augment the default plugin-based ones.

¹You can check this by searching for the string “`gatehome$`” in the XML

²This is done with an Ant copy task and so is subject to the normal `defaultexcludes`

Resolving Unresolved Resources

By default, the build will fail if there are any relative paths that cannot be accounted for by the plugins (or the explicit hints, see Section E.2.3). However, this is configurable using the `onUnresolved` attribute, which can take the following values:

fail (default) the build fails if an unresolved relative path is found.

absolute unresolved relative paths are left pointing to the same location as in the original file, but as an *absolute* rather than a relative URL. The same file will be used even if you move the GAPP file to a different directory. This option is useful if the resource in question is visible at the same absolute location on the machine where you will be putting the packaged file (for example a very large dictionary or ontology held on a network share).

recover attempt to recover gracefully (see below).

With `onUnresolved="recover"`, unresolved resources are relocated to a directory named `application-resources` under the target GAPP file location. Resources in the same original directory are copied to the same subdirectory of `application-resources`, files from different original directories are copied to different subdirectories. Typically, for a resource whose original location was `.../myresources/grammar/clever.jape` the target location would be `application-resources/grammar/clever.jape` but if the application also referred to (say) `.../otherresources/grammar/clean.jape` then this would be mapped into `application-resources/grammar-1` to avoid a name clash.

Example:

```
<packagegapp src="original.xgapp" destfile="package/target.xgapp"
  onUnresolved="recover" />
```

Providing Mapping Hints

By default, the task knows how to handle resources that live inside plugins. You can think of this as a ‘hint’ `/foo/bar/MyPlugin -> plugins/MyPlugin`, saying that whenever the mapper finds a resource path of the form `/foo/bar/MyPlugin/X`, it will relocate it to `plugins/MyPlugin/X` relative to the output GAPP file. You can specify your own hints which will be used the same way.

```
<packagegapp src="original.xgapp" destfile="package/target.xgapp">
  <hint from="{user.home}/my-app-v1" to="resources/my-app" />
  <hint from="/share/data/bigfiles" absolute="yes" />
</packagegapp>
```

In this example, `~/my-app-v1/grammar/main.jape` would be mapped to `resources/my-app/grammar/` (as always, relative to the output GAPP file). You can also hint that certain resources should be converted to absolute paths rather than being packaged with the application, using `absolute="yes"`. The `from` and `to` values refer to directories - you cannot hint a single file, nor put two files from the same original directory into different directories in the packaged GAPP.

Explicit hints override the default plugin-based hints. For example given the hint `from="${gate.home}/plugins/ANNIE/resources" to="resources/ANNIE"`, resources within the ANNIE plugin would be mapped into `resources/ANNIE`, but the plugin `creole.xml` itself would still be mapped into `plugins/ANNIE`.

As well as providing the hints inline in the build file you can also read them from a file in the normal Java Properties format³, using

```
<hint file="hints.properties" />
```

The keys in the property file are the `from` paths (in this case, relative paths are resolved against the project base directory, as with the `location` attribute of a `property` task) and the values are the `to` paths relative to the output file location.

The order of the `<hint>` elements is significant – if more than one hint could apply to the same resource file, the one defined first is used. For example, given the hints

```
<hint from="${gate.home}/plugins/ANNIE/resources/tokeniser" to="tokeniser" />
<hint from="${gate.home}/plugins/ANNIE/resources" to="annie" />
```

the resource `plugins/ANNIE/resources/tokeniser/DefaultTokeniser.rules` would be mapped into the `tokeniser` directory, but if the hints were reversed it would instead be mapped into `annie/tokeniser`. Note, however, that this does not necessarily extend to hints loaded from property files, as the order in which hints from a single property file are applied is not specified. Given

```
<hint file="file1.proeperties" />
<hint file="file2.properties" />
```

the relative precedence of two hints from `file1` is not fixed, but it is the case that all hints in `file1` will be applied before those in `file2`.

³the `hint` tag supports all the attributes of the standard Ant `property` tag so can load the hints from a file on disk or from a resource in a JAR file

E.2.4 Streamlining your Plugins

By default, the task will recursively copy the whole content of every plugin into the target directory. In most cases this is OK but it may be the case that your plugins contain many extraneous resources that are not used by your application. In this case you can specify `copyPlugins="no"`:

```
<packagegapp src="original.xgapp" destfile="package/target.xgapp"
             copyPlugins="no" />
```

In this mode, the packager task will copy only the following files from each plugin:

- `creole.xml`
- any JAR files referenced from `<JAR>` elements in `creole.xml`

In addition it will of course copy any files *directly* referenced by the GAPP, but not files referenced indirectly (the classic examples being `.lst` files used by a gazetteer `.def`, or the individual phases of a multiphase JAPE grammar) or files that are referenced by the `creole.xml` itself as `AUTOINSTANCE` parameters (e.g. the annotation schemas in ANNIE). You will need to name these extra files explicitly as extra resources (see the next section).

E.2.5 Bundling Extra Resources

Apart from plugins (when you don't use `copyPlugins="no"`), the only files copied into the target directory are those that are referenced directly from the GAPP file. This is often but not always sufficient, for example if your application contains a multiphase JAPE transducer then `packagegapp` will include the main JAPE file but not the individual phase files. The task provides two ways to include extra files in the package:

- If you set the attribute `copyResourceDirs="yes"` on the `packagegapp` task then whenever the task packages a referenced resource file it will also recursively include the whole contents of the directory containing that file in the output package. You probably don't want to use this option if you have resource files in a directory shared with other files (e.g. your home directory...).
- To include specific extra resources you can use an `<extraresourcespath>` (see below).

The `<extraresourcespath>` allows you to specify specific extra files that should be included in the package:

```

<packagegapp src="original.xgapp" destfile="package/target.xgapp">
  <extraresourcespath>
    <pathelement location="{user.home}/common-files/README" />
    <fileset dir="{user.home}/my-app-v1" includes="grammar/*.jape" />
  </extraresourcespath>
</packagegapp>

```

As the name suggests, this is a path-like structure and supports all the usual elements and attributes of an Ant `<path>`, including multiple nested `fileset`, `filelist`, `pathelement` and other path elements. For specific types of indirect references, there are helper elements that can be included under `extraresourcespath`. Currently the only one of these is `gazetteerlists`, which takes the path to a gazetteer definition file and returns the set of `.lst` files the definition uses:

```

<gazetteerlists definition="my/resources/lists.def" encoding="UTF-8" />

```

Other helpers (e.g. for multiphase JAPE) may be implemented in future.

You can also refer to a path defined elsewhere in the usual way:

```

<path id="extra.files">
  ...
</path>

<packagegapp ...>
  <extraresourcespath refid="extra.files" />
</packagegapp>

```

Resources declared in the `extraresourcespath` and directories included using `copyResourceDirs` are treated exactly the same as resources that are referenced by the GAPP file - their target locations in the package are determined by the mapping hints, default plugin-based hints, and the `onUnresolved` setting as above. If you want to put extra resource files at specific locations in the package tree, independent of the mapping hints mechanism, you should do this with a separate `<copy>` task after the `<packagegapp>` task has done its work.

E.3 The `expandcreoles` Task - Merging Annotation-Driven Config into `creole.xml`

The `expandcreoles` task processes a number of `creole.xml` files from plugins, processes any `@CreoleResource` and `@CreoleParameter` annotations on the declared resource classes,

and merges this configuration with the original XML configuration into a new copy of the `creole.xml`. It is not necessary to do this in the normal use of GATE, and this task is documented here simply for completeness. It is intended simply for use with non-GATE tools that can process the `creole.xml` file format to extract information about plugins (the prime use case for this is to generate the GATE plugins information page automatically from the plugin definitions).

The typical usage of this task (taken from the GATE `build.xml`) is:

```
<expandcreoles todir="build/plugins" gatehome="${basedir}">
  <fileset dir="plugins" includes="*/creole.xml" />
</expandcreoles>
```

This will initialise GATE with the given `GATE_HOME` directory, then read each file from the nested fileset, parse it as a `creole.xml`, expand it from any annotation configuration, and write it out to a file under `build/plugins`. Each output file will be generated at the same location relative to the `todir` as the original file was relative to the `dir` of its `fileset`.

Appendix F

Named-Entity State Machine Patterns

There are, it seems to me, two basic reasons why minds aren't computers... The first... is that human beings are organisms. Because of this we have all sorts of needs - for food, shelter, clothing, sex etc - and capacities - for locomotion, manipulation, articulate speech etc, and so on - to which there are no real analogies in computers. These needs and capacities underlie and interact with our mental activities. This is important, not simply because we can't understand how humans behave except in the light of these needs and capacities, but because any historical explanation of how human mental life developed can only do so by looking at how this process interacted with the evolution of these needs and capacities in successive species of hominids.

...

The second reason... is that... brains don't work like computers.

Minds, Machines and Evolution, Alex Callinicos, 1997 (ISJ 74, p.103).

This chapter describes the individual grammars used in GATE for Named Entity Recognition, and how they are combined together. It relates to the default NE grammar for ANNIE, but should also provide guidelines for those adapting or creating new grammars. For documentation about specific grammars other than this core set, use this document in combination with the comments in the relevant grammar files. chapter 8 also provides information about designing new grammar rules and tips for ensuring maximum processing speed.

F.1 Main.jape

This file contains a list of the grammars to be used, in the correct processing order. The ordering of the grammars is crucial, because they are processed in series, and later grammars

may depend on annotations produced by earlier grammars.

The default grammar consists of the following phases:

- first.jape
- firstname.jape
- name.jape
- name_post.jape
- date_pre.jape
- date.jape
- reldate.jape
- number.jape
- address.jape
- url.jape
- identifier.jape
- jobtitle.jape
- final.jape
- unknown.jape
- name_context.jape
- org_context.jape
- loc_context.jape
- clean.jape

F.2 first.jape

This grammar must always be processed first. It can contain any general macros needed for the whole grammar set. This should consist of a macro defining how space and control characters are to be processed (and may consequently be different for each grammar set, depending on the text type). Because this is defined first of all, it is not necessary to restate this in later grammars. This has a big advantage – it means that default grammars can be used for specialised grammar sets, without having to be adapted to deal with e.g. different

treatment of spaces and control characters. In this way, only the `first.jape` file needs to be changed for each grammar set, rather than every individual grammar.

The `first.jape` grammar also has a dummy rule in. This is never intended to fire – it is simply added because every grammar set must contain rules, but there are no specific rules we wish to add here. Even if the rule were to match the pattern defined, it is designed not to produce any output (due to the empty RHS).

F.3 `firstname.jape`

This grammar contains rules to identify first names and titles via the gazetteer lists. It adds a gender feature where appropriate from the gazetteer list. This gender feature is used later in order to improve co-reference between names and pronouns. The grammar creates separate annotations of type `FirstPerson` and `Title`.

F.4 `name.jape`

This grammar contains initial rules for organization, location and person entities. These rules all create temporary annotations, some of which will be discarded later, but the majority of which will be converted into final annotations in later grammars. Rules beginning with ‘Not’ are negative rules – this means that we detect something and give it a special annotation (or no annotation at all) in order to prevent it being recognised as a name. This is because we have no negative operator (we have ‘=’ but not ‘!=’).

F.4.1 `Person`

We first define macros for initials, first names, surnames, and endings. We then use these to recognise combinations of first names from the previous phase, and surnames from their POS tags or case information. Persons get marked with the annotation ‘TempPerson’. We also percolate feature information about the gender from the previous annotations if known.

F.4.2 `Location`

The rules for Location are fairly straightforward, but we define them in this grammar so that any ambiguity can be resolved at the top level. Locations are often combined with other entity types, such as Organisations. This is dealt with by annotating the two entity types separately, and then combining them in a later phase. Locations are recognised mainly by

gazetteer lookup, using not only lists of known places, but also key words such as mountain, lake, river, city etc. Locations are annotated as TempLocation in this phase.

F.4.3 Organization

Organizations tend to be defined either by straight lookup from the gazetteer lists, or, for the majority, by a combination of POS or case information and key words such as ‘company’, ‘bank’, ‘Services’ ‘Ltd.’ etc. Many organizations are also identified by contextual information in the later phase org_context.jape. In this phase, organizations are annotated as TempOrganization.

F.4.4 Ambiguities

Some ambiguities are resolved immediately in this grammar, while others are left until later phases. For example, a Christian name followed by a possible Location is resolved by default to a person rather than a Location (e.g. ‘Ken London’). On the other hand, a Christian name followed by a possible organisation ending is resolved to an Organisation (e.g. ‘Alexandra Pottery’), though this is a slightly less sure rule.

F.4.5 Contextual information

Although most of the rules involving contextual information are invoked in a much later phase, there are a few which are invoked here, such as ‘X joined Y’ where X is annotated as a Person and Y as an Organization. This is so that both annotations types can be handled at once.

F.5 name_post.jape

This grammar runs after the name grammar to fix some erroneous annotations that may have been created. Of course, a more elegant solution would be not to create the problem in the first instance, but this is a workaround. For example, if the surname of a Person contains certain stop words, e.g. ‘Mary And’ then only the first name should be recognised as a Person. However, it might be that the firstname is also an Organization (and has been tagged with TempOrganization already), e.g. ‘U.N.’ If this is the case, then the annotation is left untouched, because this is correct.

F.6 date_pre.jape

This grammar precedes the date phase, because it includes extra context to prevent dates being recognised erroneously in the middle of longer expressions. It mainly treats the case where an expression is already tagged as a Person, but could also be tagged as a date (e.g. 16th Jan).

F.7 date.jape

This grammar contains the base rules for recognising times and dates. Given the complexity of potential patterns representing such expressions, there are a large number of rules and macros.

Although times and dates can be mutually ambiguous, we try to distinguish between them as early as possible. Dates, times and years are generally tagged separately (as TempDate, TempTime and TempYear respectively) and then recombined to form a final Date annotation in a later phase. This is because dates, times and years can be combined together in many different ways, and also because there can be much ambiguity between the three. For example, 1312 could be a time or a year, while 9-10 could be a span of time or date, or a fixed time or date.

F.8 reldate.jape

This grammar handles relative rather than absolute date and time sequences, such as ‘yesterday morning’, ‘2 hours ago’, ‘the first 9 months of the financial year’ etc. It uses mainly explicit key words such as ‘ago’ and items from the gazetteer lists.

F.9 number.jape

This grammar covers rules concerning money and percentages. The rules are fairly straightforward, using keywords from the gazetteer lists, and there is little ambiguity here, except for example where ‘Pound’ can be money or weight, or where there is no explicit currency denominator.

F.10 address.jape

Rules for Address cover ip addresses, phone and fax numbers, and postal addresses. In general, these are not highly ambiguous, and can be covered with simple pattern matching, although phone numbers can require use of contextual information. Currently only UK formats are really handled, though handling of foreign zipcodes and phone number formats is envisaged in future. The annotations produced are of type Email, Phone etc. and are then replaced in a later phase with final Address annotations with ‘phone’ etc. as features.

F.11 url.jape

Rules for email addresses and Urls are in a separate grammar from the other address types, for the simple reason that SpaceTokens need to be identified for these rules to operate, whereas this is not necessary for the other Address types. For speed of processing, we place them in separate grammars so that SpaceTokens can be eliminated from the Input when they are not required.

F.12 identifier.jape

This grammar identifies ‘Identifiers’ which basically means any combination of numbers and letters acting as an ID, reference number etc. not recognised as any other entity type.

F.13 jobtitle.jape

This grammar simply identifies Jobtitles from the gazetteer lists, and adds a JobTitle annotation, which is used in later phases to aid recognition of other entity types such as Person and Organization. It may then be discarded in the Clean phase if not required as a final annotation type.

F.14 final.jape

This grammar uses the temporary annotations previously assigned in the earlier phases, and converts them into final annotations. The reason for this is that we need to be able to resolve ambiguities between different entity types, so we need to have all the different entity types handled in a single grammar somewhere. Ambiguities can be resolved using prioritisation

techniques. Also, we may need to combine previously annotated elements, such as dates and times, into a single entity.

The rules in this grammar use Java code on the RHS to remove the existing temporary annotations, and replace them with new annotations. This is because we want to retain the features associated with the temporary annotations. For example, we might need to keep track of whether a person is male or female, or whether a location is a city or country. It also enables us to keep track of which rules have been used, for debugging purposes.

For the sake of obfuscation, although this phase is called final, it is not the final phase!

F.15 unknown.jape

This short grammar finds proper nouns not previously recognised, and gives them an Unknown annotation. This is then used by the namematcher – if an Unknown annotation can be matched with a previously categorised entity, its annotation is changed to that of the matched entity. Any remaining Unknown annotations are useful for debugging purposes, and can also be used as input for additional grammars or processing resources.

F.16 name_context.jape

This grammar looks for Unknown annotations occurring in certain contexts which indicate they might belong to Person. This is a typical example of a grammar that would benefit from learning or automatic context generation, because useful contexts are (a) hard to find manually and may require large volumes of training data, and (b) often very domain-specific. In this core grammar, we confine the use of contexts to fairly general uses, since this grammar should not be domain-dependent.

F.17 org_context.jape

This grammar operates on a similar principle to name_context.jape. It is slightly oriented towards business texts, so does not quite fulfil the generality criteria of the previous grammar. It does, however, provide some insight into more detailed use of contexts.[i/pj](#)

F.18 `loc_context.jape`

This grammar also operates in a similar manner to the preceding two, using general context such as coordinated pairs of locations, and hyponymic types of information.

F.19 `clean.jape`

This grammar comes last of all, and simply aims to clean up (remove) some of the temporary annotations that may not have been deleted along the way.

Appendix G

Part-of-Speech Tags used in the Hepple Tagger

CC - coordinating conjunction: ‘and’, ‘but’, ‘nor’, ‘or’, ‘yet’, plus, minus, less, times (multiplication), over (division). Also ‘for’ (because) and ‘so’ (i.e., ‘so that’).

CD - cardinal number

DT - determiner: Articles including ‘a’, ‘an’, ‘every’, ‘no’, ‘the’, ‘another’, ‘any’, ‘some’, ‘those’.

EX - existential there: Unstressed ‘there’ that triggers inversion of the inflected verb and the logical subject; ‘There was a party in progress’.

FW - foreign word

IN - preposition or subordinating conjunction

JJ - adjective: Hyphenated compounds that are used as modifiers; happy-go-lucky.

JJR - adjective - comparative: Adjectives with the comparative ending ‘-er’ and a comparative meaning. Sometimes ‘more’ and ‘less’.

JJS - adjective - superlative: Adjectives with the superlative ending ‘-est’ (and ‘worst’). Sometimes ‘most’ and ‘least’.

JJSS - -unknown-, but probably a variant of JJS

-LRB- - -unknown-

LS - list item marker: Numbers and letters used as identifiers of items in a list.

MD - modal: All verbs that don’t take an ‘-s’ ending in the third person singular present: ‘can’, ‘could’, ‘dare’, ‘may’, ‘might’, ‘must’, ‘ought’, ‘shall’, ‘should’, ‘will’, ‘would’.

NN - noun - singular or mass

NNP - proper noun - singular: All words in names usually are capitalized but titles might not be.

NNPS - proper noun - plural: All words in names usually are capitalized but titles might not be.

NNS - noun - plural

NP - proper noun - singular

NPS - proper noun - plural

PDT - predeterminer: Determiner like elements preceding an article or possessive pronoun; 'all/PDT his marbles', 'quite/PDT a mess'.

POS - possessive ending: Nouns ending in 's' or ''.

PP - personal pronoun

PRPR\$ - unknown-, but probably possessive pronoun

PRP - unknown-, but probably possessive pronoun

PRP\$ - unknown, but probably possessive pronoun, such as 'my', 'your', 'his', 'his', 'its', 'one's', 'our', and 'their'.

RB - adverb: most words ending in '-ly'. Also 'quite', 'too', 'very', 'enough', 'indeed', 'not', '-n't', and 'never'.

RBR - adverb - comparative: adverbs ending with '-er' with a comparative meaning.

RBS - adverb - superlative

RP - particle: Mostly monosyllabic words that also double as directional adverbs.

STAART - start state marker (used internally)

SYM - symbol: technical symbols or expressions that aren't English words.

TO - literal to

UH - interjection: Such as 'my', 'oh', 'please', 'uh', 'well', 'yes'.

VBD - verb - past tense: includes conditional form of the verb 'to be'; 'If I were/VBD rich...'

VBG - verb - gerund or present participle

VCN - verb - past participle

VBP - verb - non-3rd person singular present

VB - verb - base form: subsumes imperatives, infinitives and subjunctives.

VBZ - verb - 3rd person singular present

WDT - wh-determiner

WP\$ - possessive wh-pronoun: includes 'whose'

WP - wh-pronoun: includes 'what', 'who', and 'whom'.

WRB - wh-adverb: includes 'how', 'where', 'why'. Includes 'when' when used in a temporal sense.

:: - literal colon

, - literal comma

\$ - literal dollar sign

- - literal double-dash

- literal double quotes

- literal grave

(- literal left parenthesis

. - literal period

- literal pound sign

) - literal right parenthesis

- literal single quote or apostrophe

References

[Agatonovic *et al.* 08]

M. Agatonovic, N. Aswani, K. Bontcheva, H. Cunningham, T. Heitz, Y. Li, I. Roberts, and V. Tablan. Large-scale, parallel automatic patent annotation. In *Proc. of 1st International CIKM Workshop on Patent Information Retrieval - PaIR'08*, Napa Valley, California, USA, October 30 2008.

[Aho *et al.* 86]

A. V. Aho, R. Sethi, and J. D. Ullman. **Compilers Principles, Techniques, and Tools**. Addison-Wesley, Reading, Massachusetts, 1986.

[Aswani *et al.* 05]

N. Aswani, V. Tablan, K. Bontcheva, and H. Cunningham. Indexing and Querying Linguistic Metadata and Document Content. In *Proceedings of Fifth International Conference on Recent Advances in Natural Language Processing (RANLP2005)*, Borovets, Bulgaria, 2005.

[Aswani *et al.* 06]

N. Aswani, K. Bontcheva, and H. Cunningham. Mining information for instance unification. In *5th International Semantic Web Conference (ISWC2006)*, Athens, Georgia, USA, 2006.

[Azar 89]

S. Azar. *Understanding and Using English Grammar*. Prentice Hall Regents, 1989.

[Baker *et al.* 02]

P. Baker, A. Hardie, T. McEnery, H. Cunningham, and R. Gaizauskas. EMILLE, A 67-Million Word Corpus of Indic Languages: Data Collection, Mark-up and Harmonisation. In *Proceedings of 3rd Language Resources and Evaluation Conference (LREC'2002)*, pages 819–825, 2002.

[Bird & Liberman 99]

S. Bird and M. Liberman. A Formal Framework for Linguistic Annotation. Technical Report MS-CIS-99-01, Department of Computer and Information Science, University of Pennsylvania, 1999. <http://xxx.lanl.gov/abs/cs.CL/9903003>.

[Bontcheva & Sabou 06]

K. Bontcheva and M. Sabou. Learning Ontologies from Software Artifacts: Exploring

and Combining Multiple Sources. In *Workshop on Semantic Web Enabled Software Engineering (SWESE)*, Athens, G.A., USA, November 2006.

[Bontcheva 04]

K. Bontcheva. Open-source Tools for Creation, Maintenance, and Storage of Lexical Resources for Language Generation from Ontologies. In *Proceedings of 4th Language Resources and Evaluation Conference (LREC'04)*, 2004.

[Bontcheva 05]

K. Bontcheva. Generating Tailored Textual Summaries from Ontologies. In *Second European Semantic Web Conference (ESWC'2005)*, 2005.

[Bontcheva et al. 00]

K. Bontcheva, H. Brugman, A. Russel, P. Wittenburg, and H. Cunningham. An Experiment in Unifying Audio-Visual and Textual Infrastructures for Language Processing R&D. In *Proceedings of the Workshop on Using Toolsets and Architectures To Build NLP Systems at COLING-2000*, Luxembourg, 2000. <http://gate.ac.uk/>.

[Bontcheva et al. 02a]

K. Bontcheva, H. Cunningham, V. Tablan, D. Maynard, and O. Hamza. Using GATE as an Environment for Teaching NLP. In *Proceedings of the ACL Workshop on Effective Tools and Methodologies in Teaching NLP*, 2002. <http://gate.ac.uk/sale/ac102/gate4teaching.pdf>.

[Bontcheva et al. 02b]

K. Bontcheva, H. Cunningham, V. Tablan, D. Maynard, and H. Saggion. Developing Reusable and Robust Language Processing Components for Information Systems using GATE. In *Proceedings of the 3rd International Workshop on Natural Language and Information Systems (NLIS'2002)*, Aix-en-Provence, France, 2002. IEEE Computer Society Press. <http://gate.ac.uk/sale/nlis/nlis.ps>.

[Bontcheva et al. 02c]

K. Bontcheva, M. Dimitrov, D. Maynard, V. Tablan, and H. Cunningham. Shallow Methods for Named Entity Coreference Resolution. In *Chaînes de références et résolveurs d'anaphores, workshop TALN 2002*, Nancy, France, 2002. <http://gate.ac.uk/sale/taln02/taln-ws-coref.pdf>.

[Bontcheva et al. 03]

K. Bontcheva, A. Kiryakov, H. Cunningham, B. Popov, and M. Dimitrov. Semantic web enabled, open source language technology. In *EACL workshop on Language Technology and the Semantic Web: NLP and XML*, Budapest, Hungary, 2003. <http://gate.ac.uk/sale/eac103-semweb/bontcheva-etal-final.pdf>.

[Bontcheva et al. 04]

K. Bontcheva, V. Tablan, D. Maynard, and H. Cunningham. Evolving GATE to Meet New Challenges in Language Engineering. *Natural Language Engineering*, 10(3/4):349—373, 2004.

- [Bontcheva *et al.* 06a]
K. Bontcheva, H. Cunningham, A. Kiryakov, and V. Tablan. Semantic Annotation and Human Language Technology. In J. Davies, R. Studer, and P. Warren, editors, *Semantic Web Technology: Trends and Research*. John Wiley and Sons, 2006.
- [Bontcheva *et al.* 06b]
K. Bontcheva, J. Davies, A. Duke, T. Glover, N. Kings, and I. Thurlow. Semantic Information Access. In J. Davies, R. Studer, and P. Warren, editors, *Semantic Web Technologies*. John Wiley and Sons, 2006.
- [Bontcheva *et al.* 09]
K. Bontcheva, B. Davis, A. Funk, Y. Li, and T. Wang. Human Language Technologies. In J. Davies, M. Grobelnik, and D. Mladenic, editors, *Semantic Knowledge Management*, pages 37–49. 2009.
- [Booch 94]
G. Booch. *Object-Oriented Analysis and Design 2nd Edn.* Benjamin/Cummings, 1994.
- [Brugman *et al.* 99]
H. Brugman, K. Bontcheva, P. Wittenburg, and H. Cunningham. Integrating Multimedia and Textual Software Architectures for Language Technology. Technical report MPI-TG-99-1, Max-Planck Institute for Psycholinguistics, Nijmegen, Netherlands, 1999.
- [Carletta 96]
J. Carletta. Assessing agreement on classification tasks: the Kappa statistic. *Computational Linguistics*, 22(2):249–254, 1996.
- [CC001]
LIBSVM: a library for support vector machines, 2001. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [Chinchor 92]
N. Chinchor. MUC-4 Evaluation Metrics. In *Proceedings of the Fourth Message Understanding Conference*, pages 22–29, 1992.
- [Cimiano *et al.* 03]
P. Cimiano, S. Staab, and J. Tane. Automatic Acquisition of Taxonomies from Text: FCA meets NLP. In *Proceedings of the ECML/PKDD Workshop on Adaptive Text Extraction and Mining*, pages 10–17, Cavtat-Dubrovnik, Croatia, 2003.
- [Cobuild 99]
C. Cobuild, editor. *English Grammar*. Harper Collins, 1999.
- [Cunningham & Bontcheva 05]
H. Cunningham and K. Bontcheva. Computational Language Systems, Architectures. *Encyclopedia of Language and Linguistics, 2nd Edition*, pages 733–752, 2005.

- [Cunningham & Scott 04a]
H. Cunningham and D. Scott. Introduction to the Special Issue on Software Architecture for Language Engineering. *Natural Language Engineering*, 2004. <http://gate.ac.uk/sale/jnle-sale/intro/intro-main.pdf>.
- [Cunningham & Scott 04b]
H. Cunningham and D. Scott, editors. *Special Issue of Natural Language Engineering on Software Architecture for Language Engineering*. Cambridge University Press, 2004.
- [Cunningham 94]
H. Cunningham. Support Software for Language Engineering Research. Technical Report 94/05, Centre for Computational Linguistics, UMIST, Manchester, 1994.
- [Cunningham 99a]
H. Cunningham. A Definition and Short History of Language Engineering. *Journal of Natural Language Engineering*, 5(1):1–16, 1999.
- [Cunningham 99b]
H. Cunningham. JAPE: a Java Annotation Patterns Engine. Research Memorandum CS-99-06, Department of Computer Science, University of Sheffield, May 1999.
- [Cunningham 00]
H. Cunningham. *Software Architecture for Language Engineering*. Unpublished PhD thesis, University of Sheffield, 2000. <http://gate.ac.uk/sale/thesis/>.
- [Cunningham 02]
H. Cunningham. GATE, a General Architecture for Text Engineering. *Computers and the Humanities*, 36:223–254, 2002.
- [Cunningham 05]
H. Cunningham. Information Extraction, Automatic. *Encyclopedia of Language and Linguistics, 2nd Edition*, pages 665–677, 2005.
- [Cunningham et al. 94]
H. Cunningham, M. Freeman, and W. Black. Software Reuse, Object-Oriented Frameworks and Natural Language Processing. In *New Methods in Language Processing (NeMLaP-1), September 1994*, Manchester, 1994. (Re-published in book form 1997 by UCL Press).
- [Cunningham et al. 95]
H. Cunningham, R. Gaizauskas, and Y. Wilks. A General Architecture for Text Engineering (GATE) – a new approach to Language Engineering R&D. Technical Report CS-95-21, Department of Computer Science, University of Sheffield, 1995. <http://xxx.lanl.gov/abs/cs.CL/9601009>.

- [Cunningham *et al.* 96a]
H. Cunningham, K. Humphreys, R. Gaizauskas, and M. Stower. CREOLE Developer's Manual. Technical report, Department of Computer Science, University of Sheffield, 1996. <http://www.dcs.shef.ac.uk/nlp/gate>.
- [Cunningham *et al.* 96b]
H. Cunningham, K. Humphreys, R. Gaizauskas, and Y. Wilks. TIPSTER-Compatible Projects at Sheffield. In *Advances in Text Processing, TIPSTER Program Phase II*. DARPA, Morgan Kaufmann, California, 1996.
- [Cunningham *et al.* 96c]
H. Cunningham, Y. Wilks, and R. Gaizauskas. GATE – a General Architecture for Text Engineering. In *Proceedings of the 16th Conference on Computational Linguistics (COLING-96)*, Copenhagen, August 1996. ftp://ftp.dcs.shef.ac.uk/home/hamish/auto_papers/Cun96b.ps.
- [Cunningham *et al.* 96d]
H. Cunningham, Y. Wilks, and R. Gaizauskas. Software Infrastructure for Language Engineering. In *Proceedings of the AISB Workshop on Language Engineering for Document Analysis and Recognition*, Brighton, U.K., April 1996.
- [Cunningham *et al.* 96e]
H. Cunningham, Y. Wilks, and R. Gaizauskas. New Methods, Current Trends and Software Infrastructure for NLP. In *Proceedings of the Conference on New Methods in Natural Language Processing (NeMLaP-2)*, Bilkent University, Turkey, September 1996. ftp://ftp.dcs.shef.ac.uk/home/hamish/auto_papers/Cun96c.ps.
- [Cunningham *et al.* 97a]
H. Cunningham, K. Humphreys, R. Gaizauskas, and Y. Wilks. GATE – a TIPSTER-based General Architecture for Text Engineering. In *Proceedings of the TIPSTER Text Program (Phase III) 6 Month Workshop*. DARPA, Morgan Kaufmann, California, May 1997. ftp://ftp.dcs.shef.ac.uk/home/hamish/auto_papers/Cun97e.ps.
- [Cunningham *et al.* 97b]
H. Cunningham, K. Humphreys, R. Gaizauskas, and Y. Wilks. Software Infrastructure for Natural Language Processing. In *Proceedings of the 5th Conference on Applied Natural Language Processing (ANLP-97)*, March 1997. ftp://ftp.dcs.shef.ac.uk/home/hamish/auto_papers/Cun97a.ps.gz.
- [Cunningham *et al.* 98a]
H. Cunningham, W. Peters, C. McCauley, K. Bontcheva, and Y. Wilks. A Level Playing Field for Language Resource Evaluation. In *Workshop on Distributing and Accessing Lexical Resources at Conference on Language Resources Evaluation, Granada, Spain*, 1998. <http://www.dcs.shef.ac.uk/hamish/dalr>.
- [Cunningham *et al.* 98b]
H. Cunningham, M. Stevenson, and Y. Wilks. Implementing a Sense Tagger within a

General Architecture for Language Engineering. In *Proceedings of the Third Conference on New Methods in Language Engineering (NeMLaP-3)*, pages 59–72, Sydney, Australia, 1998.

[Cunningham *et al.* 99]

H. Cunningham, R. Gaizauskas, K. Humphreys, and Y. Wilks. Experience with a Language Engineering Architecture: Three Years of GATE. In *Proceedings of the AISB'99 Workshop on Reference Architectures and Data Standards for NLP*, Edinburgh, April 1999. The Society for the Study of Artificial Intelligence and Simulation of Behaviour. <http://www.dcs.shef.ac.uk/~hamish/GateAisb99.html>.

[Cunningham *et al.* 00a]

H. Cunningham, K. Bontcheva, W. Peters, and Y. Wilks. Uniform language resource access and distribution in the context of a General Architecture for Text Engineering (GATE). In *Proceedings of the Workshop on Ontologies and Language Resources (OntoLex'2000)*, Sozopol, Bulgaria, September 2000. <http://gate.ac.uk/sale/ontolex/ontolex.ps>.

[Cunningham *et al.* 00b]

H. Cunningham, K. Bontcheva, V. Tablan, and Y. Wilks. Software Infrastructure for Language Resources: a Taxonomy of Previous Work and a Requirements Analysis. In *Proceedings of the 2nd International Conference on Language Resources and Evaluation (LREC-2)*, Athens, 2000. <http://gate.ac.uk/>.

[Cunningham *et al.* 00c]

H. Cunningham, D. Maynard, K. Bontcheva, V. Tablan, and Y. Wilks. Experience of using GATE for NLP R&D. In *Proceedings of the Workshop on Using Toolsets and Architectures To Build NLP Systems at COLING-2000*, Luxembourg, 2000. <http://gate.ac.uk/>.

[Cunningham *et al.* 00d]

H. Cunningham, D. Maynard, and V. Tablan. JAPE: a Java Annotation Patterns Engine (Second Edition). Research Memorandum CS-00-10, Department of Computer Science, University of Sheffield, November 2000.

[Cunningham *et al.* 02]

H. Cunningham, D. Maynard, K. Bontcheva, and V. Tablan. GATE: A Framework and Graphical Development Environment for Robust NLP Tools and Applications. In *Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistics (ACL'02)*, 2002.

[Cunningham *et al.* 03]

H. Cunningham, V. Tablan, K. Bontcheva, and M. Dimitrov. Language Engineering Tools for Collaborative Corpus Annotation. In *Proceedings of Corpus Linguistics 2003*, Lancaster, UK, 2003. <http://gate.ac.uk/sale/cl03/distrib-ollie-cl03.doc>.

[Damljjanovic & Bontcheva 08]

D. Damljjanovic and K. Bontcheva. Enhanced Semantic Access to Software Artefacts. In *Workshop on Semantic Web Enabled Software Engineering (SWESE)*, Karlsruhe, Germany, October 2008.

[Damljjanovic *et al.* 08]

D. Damljjanovic, V. Tablan, and K. Bontcheva. A text-based query interface to owl ontologies. In *6th Language Resources and Evaluation Conference (LREC)*, Marrakech, Morocco, May 2008. ELRA.

[Damljjanovic *et al.* 09]

D. Damljjanovic, F. Amardeilh, and K. Bontcheva. CA Manager Framework: Creating Customised Workflows for Ontology Population and Semantic Annotation. In *Proceedings of The Fifth International Conference on Knowledge Capture (KCAP'09)*, California, USA, September 2009.

[Davies & Fleiss 82]

M. Davies and J. Fleiss. Measuring Agreement for Multinomial Data. *Biometrics*, 38:1047–1051, 1982.

[Davis *et al.* 06]

B. Davis, S. Handschuh, H. Cunningham, and V. Tablan. Further use of Controlled Natural Language for Semantic Annotation of Wikis. In *Proceedings of the 1st Semantic Authoring and Annotation Workshop at ISWC2006*, Athens, Georgia, USA, November 2006.

[Della Valle *et al.* 08]

E. Della Valle, D. Cerizza, I. Celino, A. Turati, H. Lausen, N. Steinmetz, M. Erdmann, and A. Funk. Realizing Service-Finder: Web service discovery at web scale. In *European Semantic Technology Conference (ESTC)*, Vienna, September 2008.

[Dimitrov 02a]

M. Dimitrov. *A Light-weight Approach to Coreference Resolution for Named Entities in Text*. MSc Thesis, University of Sofia, Bulgaria, 2002. <http://www.ontotext.com/ie/thesis-m.pdf>.

[Dimitrov 02b]

M. Dimitrov. *A Light-weight Approach to Coreference Resolution for Named Entities in Text*. MSc Thesis, University of Sofia, Bulgaria, 2002. <http://www.ontotext.com/ie/thesis-m.pdf>.

[Dimitrov *et al.* 02]

M. Dimitrov, K. Bontcheva, H. Cunningham, and D. Maynard. A Light-weight Approach to Coreference Resolution for Named Entities in Text. In *Proceedings of the Fourth Discourse Anaphora and Anaphor Resolution Colloquium (DAARC)*, Lisbon, 2002.

[Dimitrov *et al.* 04]

M. Dimitrov, K. Bontcheva, H. Cunningham, and D. Maynard. A Light-weight Approach to Coreference Resolution for Named Entities in Text. In A. Branco, T. McEnery, and R. Mitkov, editors, *Anaphora Processing: Linguistic, Cognitive and Computational Modelling*. John Benjamins, 2004.

[Dowman *et al.* 05a]

M. Dowman, V. Tablan, H. Cunningham, and B. Popov. Content augmentation for mixed-mode news broadcasts. In *Proceedings of the 3rd European Conference on Interactive Television: User Centred ITV Systems, Programmes and Applications*, Aalborg University, Denmark, 2005. <http://gate.ac.uk/sale/euro-itv-2005/content-augmentation-for-mixed-mode-news-broadcast>

[Dowman *et al.* 05b]

M. Dowman, V. Tablan, H. Cunningham, and B. Popov. Web-assisted annotation, semantic indexing and search of television and radio news. In *Proceedings of the 14th International World Wide Web Conference*, Chiba, Japan, 2005. <http://gate.ac.uk/sale/www05/web-assisted-annotation.pdf>.

[Dowman *et al.* 05c]

M. Dowman, V. Tablan, H. Cunningham, C. Ursu, and B. Popov. Semantically enhanced television news through web and video integration. In *Second European Semantic Web Conference (ESWC'2005)*, 2005.

[Eugenio & Glass 04]

B. D. Eugenio and M. Glass. The kappa statistic: a second look. *Computational Linguistics*, 1(30), 2004. (squib).

[Fleiss 75]

J. L. Fleiss. Measuring agreement between two judges on the presence or absence of a trait. *Biometrics*, 31:651–659, 1975.

[Frakes & Baeza-Yates 92]

W. Frakes and R. Baeza-Yates, editors. *Information retrieval, data structures and algorithms*. Prentice Hall, New York, Englewood Cliffs, N.J., 1992.

[Funk *et al.* 07a]

A. Funk, D. Maynard, H. Saggion, and K. Bontcheva. Ontological integration of information extracted from multiple sources. In *Multi-source Multilingual Information Extraction and Summarization (MMIES) workshop at Recent Advances in Natural Language Processing (RANLP07)*, Borovets, Bulgaria, September 2007.

[Funk *et al.* 07b]

A. Funk, V. Tablan, K. Bontcheva, H. Cunningham, B. Davis, and S. Handschuh. CLOnE: Controlled Language for Ontology Editing. In *Proceedings of the 6th International Semantic Web Conference (ISWC 2007)*, Busan, Korea, November 2007.

- [Gaizauskas *et al.* 95]
R. Gaizauskas, T. Wakao, K. Humphreys, H. Cunningham, and Y. Wilks. Description of the LaSIE system as used for MUC-6. In *Proceedings of the Sixth Message Understanding Conference (MUC-6)*. Morgan Kaufmann, California, 1995.
- [Gaizauskas *et al.* 96a]
R. Gaizauskas, P. Rodgers, H. Cunningham, and K. Humphreys. GATE User Guide. <http://www.dcs.shef.ac.uk/nlp/gate>, 1996.
- [Gaizauskas *et al.* 96b]
R. Gaizauskas, H. Cunningham, Y. Wilks, P. Rodgers, and K. Humphreys. GATE – an Environment to Support Research and Development in Natural Language Engineering. In *Proceedings of the 8th IEEE International Conference on Tools with Artificial Intelligence (ICTAI-96)*, Toulouse, France, October 1996. <ftp://ftp.dcs.shef.ac.uk/home/robertg/ictai96.ps>.
- [Gaizauskas *et al.* 03]
R. Gaizauskas, M. A. Greenwood, M. Hepple, I. Roberts, H. Saggion, and M. Sargaison. The University of Sheffield’s TREC 2003 Q&A Experiments. In *In Proceedings of the 12th Text REtrieval Conference*, 2003.
- [Gaizauskas *et al.* 04]
R. Gaizauskas, M. A. Greenwood, M. Hepple, I. Roberts, H. Saggion, and M. Sargaison. The University of Sheffield’s TREC 2004 Q&A Experiments. In *In Proceedings of the 13th Text REtrieval Conference*, 2004.
- [Gaizauskas *et al.* 05]
R. Gaizauskas, M. Greenwood, M. Hepple, H. Harkema, H. Saggion, and A. Sanka. The University of Sheffield’s TREC 2005 Q&A Experiments. In *In Proceedings of the 11th Text REtrieval Conference*, 2005.
- [Gambäck & Olsson 00]
B. Gambäck and F. Olsson. Experiences of Language Engineering Algorithm Reuse. In *Second International Conference on Language Resources and Evaluation (LREC)*, pages 155–160, Athens, Greece, 2000.
- [Gazdar & Mellish 89]
G. Gazdar and C. Mellish. *Natural Language Processing in Prolog*. Addison-Wesley, Reading, MA, 1989.
- [Greenwood *et al.* 02]
M. Greenwood, I. Roberts, and R. Gaizauskas. The University of Sheffield’s TREC 2002 Q&A Experiments. In *In Proceedings of the 11th Text REtrieval Conference*, 2002.

[Grishman 97]

R. Grishman. TIPSTER Architecture Design Document Version 2.3. Technical report, DARPA, 1997. http://www.itl.nist.gov/div894/894.02/related_projects/-tipster/.

[Hepple 00]

M. Hepple. Independence and commitment: Assumptions for rapid training and execution of rule-based POS taggers. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics (ACL-2000)*, Hong Kong, October 2000.

[Hripcsak & Heitjan 02]

G. Hripcsak and D. Heitjan. Measuring agreement in medical informatics reliability studies. *Journal of Biomedical Informatics*, 35:99–110, 2002.

[Hripcsak & Rothschild 05]

G. Hripcsak and A. S. Rothschild. Agreement, the F-measure, and Reliability in Information Retrieval. *Journal of the American Medical Informatics Association*, 12(3):296–298, 2005.

[Humphreys *et al.* 96]

K. Humphreys, R. Gaizauskas, H. Cunningham, and S. Azzam. CREOLE Module Specifications. <http://www.dcs.shef.ac.uk/nlp/gate/>, 1996.

[Humphreys *et al.* 98]

K. Humphreys, R. Gaizauskas, S. Azzam, C. Huyck, B. Mitchell, H. Cunningham, and Y. Wilks. Description of the LaSIE system as used for MUC-7. In *Proceedings of the Seventh Message Understanding Conference (MUC-7)*. http://www.itl.nist.gov/iaui/894.02/related_projects/muc/index.html, 1998.

[Humphreys *et al.* 99]

K. Humphreys, R. Gaizauskas, M. Hepple, and M. Sanderson. The University of Sheffield TREC-8 Q&A System. In *In Proceedings of the 8th Text REtrieval Conference*, 1999.

[Jackson 75]

M. Jackson. *Principles of Program Design*. Academic Press, London, 1975.

[Kiryakov 03]

A. Kiryakov. Ontology and Reasoning in MUMIS: Towards the Semantic Web. Technical Report CS-03-03, Department of Computer Science, University of Sheffield, 2003. <http://gate.ac.uk/gate/doc/papers.html>.

[Laclavik & Maynard 09]

M. Laclavik and D. Maynard. Motivating intelligent email in business: an investigation into current trends for email processing and communication research. In *Proceedings of Workshop on Emails in e-Commerce and Enterprise Context, 11th IEEE Conference on Commerce and Enterprise Computing*, Vienna, Austria, 2009.

- [Lal & Ruger 02]
P. Lal and S. Ruger. Extract-based summarization with simplification. In *Proceedings of the ACL 2002 Automatic Summarization / DUC 2002 Workshop*, 2002. <http://www.doc.ic.ac.uk/~srueger/pr-p.lal-2002/duc02-final.pdf>.
- [Lal 02]
P. Lal. Text summarisation. Unpublished M.Sc. thesis, Imperial College, London, 2002.
- [Li & Bontcheva 08]
Y. Li and K. Bontcheva. Adapting support vector machines for f-term-based classification of patents. *ACM Transactions on Asian Language Information Processing*, 7(2):7:1–7:19, 2008.
- [Li & Cunningham 08]
Y. Li and H. Cunningham. Geometric and Quantum Methods for Information Retrieval. *SIGIR Forum*, 42(2):22–32, 2008.
- [Li & Shawe-Taylor 03]
Y. Li and J. Shawe-Taylor. The SVM with Uneven Margins and Chinese Document Categorization. In *Proceedings of The 17th Pacific Asia Conference on Language, Information and Computation (PACLIC17)*, Singapore, Oct. 2003.
- [Li & Shawe-Taylor 06]
Y. Li and J. Shawe-Taylor. Using KCCA for Japanese-English Cross-language Information Retrieval and Document Classification. *Journal of Intelligent Information Systems*, 27(2):117–133, 2006.
- [Li & Shawe-Taylor 07]
Y. Li and J. Shawe-Taylor. Advanced Learning Algorithms for Cross-language Patent Retrieval and Classification. *Information Processing and Management*, 43(5):1183–1199, 2007.
- [Li *et al.* 02]
Y. Li, H. Zaragoza, R. Herbrich, J. Shawe-Taylor, and J. Kandola. The Perceptron Algorithm with Uneven Margins. In *Proceedings of the 9th International Conference on Machine Learning (ICML-2002)*, pages 379–386, 2002.
- [Li *et al.* 04]
Y. Li, K. Bontcheva, and H. Cunningham. An SVM Based Learning Algorithm for Information Extraction. Machine Learning Workshop, Sheffield, 2004. <http://gate.ac.uk/sale/ml-ws04/mlw2004.pdf>.
- [Li *et al.* 05a]
Y. Li, K. Bontcheva, and H. Cunningham. SVM Based Learning System For Information Extraction. In M. N. J. Winkler and N. Lawrence, editors, *Deterministic and Statistical Methods in Machine Learning*, LNAI 3635, pages 319–339. Springer Verlag, 2005.

[Li *et al.* 05b]

Y. Li, K. Bontcheva, and H. Cunningham. Using Uneven Margins SVM and Perceptron for Information Extraction. In *Proceedings of Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, 2005.

[Li *et al.* 05c]

Y. Li, C. Miao, K. Bontcheva, and H. Cunningham. Perceptron Learning for Chinese Word Segmentation. In *Proceedings of Fourth SIGHAN Workshop on Chinese Language processing (Sighan-05)*, pages 154–157, Korea, 2005.

[Li *et al.* 07a]

Y. Li, K. Bontcheva, and H. Cunningham. Hierarchical, Perceptron-like Learning for Ontology Based Information Extraction. In *16th International World Wide Web Conference (WWW2007)*, pages 777–786, May 2007.

[Li *et al.* 07b]

Y. Li, K. Bontcheva, and H. Cunningham. Cost Sensitive Evaluation Measures for F-term Patent Classification. In *The First International Workshop on Evaluating Information Access (EVIA 2007)*, pages 44–53, May 2007.

[Li *et al.* 07c]

Y. Li, K. Bontcheva, and H. Cunningham. Experiments of opinion analysis on the corpora MPQA and NTCIR-6. In *Proceedings of the Sixth NTCIR Workshop Meeting on Evaluation of Information Access Technologies: Information Retrieval, Question Answering and Cross-Lingual Information Access*, pages 323–329, May 2007.

[Li *et al.* 07d]

Y. Li, K. Bontcheva, and H. Cunningham. SVM Based Learning System for F-term Patent Classification. In *Proceedings of the Sixth NTCIR Workshop Meeting on Evaluation of Information Access Technologies: Information Retrieval, Question Answering and Cross-Lingual Information Access*, pages 396–402, May 2007.

[Li *et al.* 09]

Y. Li, K. Bontcheva, and H. Cunningham. Adapting SVM for Data Sparseness and Imbalance: A Case Study on Information Extraction. *Natural Language Engineering*, 15(2):241–271, 2009.

[Lombard *et al.* 02]

M. Lombard, J. Snyder-Duch, and C. C. Bracken. Content analysis in mass communication: Assessment and reporting of intercoder reliability. *Human Communication Research*, 28:587–604, 2002.

[LREC-1 98]

Conference on Language Resources Evaluation (LREC-1), Granada, Spain, 1998.

[LREC-2 00]

Second Conference on Language Resources Evaluation (LREC-2), Athens, 2000.

[Manning & Schütze 99]

C. Manning and H. Schütze. *Foundations of Statistical Natural Language Processing*. MIT press, Cambridge, MA, 1999. Supporting materials available at <http://www.sultry.arts.usyd.edu.au/fsnlp/>.

[Manov *et al.* 03]

D. Manov, A. Kiryakov, B. Popov, K. Bontcheva, and D. Maynard. Experiments with geographic knowledge for information extraction. In *Workshop on Analysis of Geographic References, HLT/NAACL'03*, Edmonton, Canada, 2003. <http://gate.ac.uk/sale/hlt03/paper03.pdf>.

[Maynard 05]

D. Maynard. Benchmarking ontology-based annotation tools for the semantic web. In *UK e-Science Programme All Hands Meeting (AHM2005) Workshop on Text Mining, e-Research and Grid-enabled Language Technology*, Nottingham, UK, 2005.

[Maynard 08]

D. Maynard. Benchmarking textual annotation tools for the semantic web. In *Proc. of 6th International Conference on Language Resources and Evaluation (LREC)*, Marrakech, Morocco, 2008.

[Maynard *et al.* 00]

D. Maynard, H. Cunningham, K. Bontcheva, R. Catizone, G. Demetriou, R. Gaizauskas, O. Hamza, M. Hepple, P. Herring, B. Mitchell, M. Oakes, W. Peters, A. Setzer, M. Stevenson, V. Tablan, C. Ursu, and Y. Wilks. A Survey of Uses of GATE. Technical Report CS-00-06, Department of Computer Science, University of Sheffield, 2000.

[Maynard *et al.* 01]

D. Maynard, V. Tablan, C. Ursu, H. Cunningham, and Y. Wilks. Named Entity Recognition from Diverse Text Types. In *Recent Advances in Natural Language Processing 2001 Conference*, pages 257–274, Tzigorv Chark, Bulgaria, 2001.

[Maynard *et al.* 02a]

D. Maynard, K. Bontcheva, H. Saggion, H. Cunningham, and O. Hamza. Using a Text Engineering Framework to Build an Extendable and Portable IE-based Summarisation System. In *Proceedings of the ACL Workshop on Text Summarisation*, pages 19–26, Philadelphina, Pennsylvania, 2002. ACM.

[Maynard *et al.* 02b]

D. Maynard, H. Cunningham, K. Bontcheva, and M. Dimitrov. Adapting a robust multi-genre NE system for automatic content extraction. In *Proceedings of the 10th International Conference on Artificial Intelligence: Methodology, Systems, Applications (AIMSA '02)*, Varna, Bulgaria, Sep 2002.

[Maynard *et al.* 02c]

D. Maynard, H. Cunningham, K. Bontcheva, and M. Dimitrov. Adapting A Robust

Multi-Genre NE System for Automatic Content Extraction. In *Proceedings of the Tenth International Conference on Artificial Intelligence: Methodology, Systems, Applications (AIMSA 2002)*, 2002.

[Maynard *et al.* 02d]

D. Maynard, H. Cunningham, and R. Gaizauskas. Named entity recognition at sheffield university. In H. Holmboe, editor, *Nordic Language Technology – Arbog for Nordisk Sprogteknologisk Forskningsprogram 2002-2004*, pages 141–145. Museum Tusulanums Forlag, 2002.

[Maynard *et al.* 02e]

D. Maynard, V. Tablan, H. Cunningham, C. Ursu, H. Saggion, K. Bontcheva, and Y. Wilks. Architectural Elements of Language Engineering Robustness. *Journal of Natural Language Engineering – Special Issue on Robust Methods in Analysis of Natural Language Data*, 8(2/3):257–274, 2002.

[Maynard *et al.* 03a]

D. Maynard, K. Bontcheva, and H. Cunningham. From information extraction to content extraction. Submitted to EACL’2003, 2003.

[Maynard *et al.* 03b]

D. Maynard, K. Bontcheva, and H. Cunningham. Towards a semantic extraction of named entities. In G. Angelova, K. Bontcheva, R. Mitkov, N. Nicolov, and N. Nikolov, editors, *Proceedings of Recent Advances in Natural Language Processing (RANLP’03)*, pages 255–261, Borovets, Bulgaria, Sep 2003. <http://gate.ac.uk/sale/ranlp03/ranlp03.pdf>.

[Maynard *et al.* 03c]

D. Maynard, K. Bontcheva, and H. Cunningham. Towards a semantic extraction of Named Entities. In *Recent Advances in Natural Language Processing*, Bulgaria, 2003.

[Maynard *et al.* 03d]

D. Maynard, V. Tablan, K. Bontcheva, and H. Cunningham. Rapid customisation of an Information Extraction system for surprise languages. *Special issue of ACM Transactions on Asian Language Information Processing: Rapid Development of Language Capabilities: The Surprise Languages*, 2:295–300, 2003.

[Maynard *et al.* 03e]

D. Maynard, V. Tablan, and H. Cunningham. NE recognition without training data on a language you don’t speak. In *ACL Workshop on Multilingual and Mixed-language Named Entity Recognition: Combining Statistical and Symbolic Models*, Sapporo, Japan, 2003.

[Maynard *et al.* 04a]

D. Maynard, K. Bontcheva, and H. Cunningham. Automatic Language-Independent Induction of Gazetteer Lists. In *Proceedings of 4th Language Resources and Evaluation Conference (LREC’04)*, Lisbon, Portugal, 2004. ELRA.

- [Maynard *et al.* 04b]
D. Maynard, H. Cunningham, A. Kourakis, and A. Kokossis. Ontology-Based Information Extraction in hTechSight. In *First European Semantic Web Symposium (ESWS 2004)*, Heraklion, Crete, 2004.
- [Maynard *et al.* 04c]
D. Maynard, M. Yankova, N. Aswani, and H. Cunningham. Automatic Creation and Monitoring of Semantic Metadata in a Dynamic Knowledge Portal. In *Proceedings of the 11th International Conference on Artificial Intelligence: Methodology, Systems, Applications (AIMSA 2004)*, Varna, Bulgaria, 2004.
- [Maynard *et al.* 06]
D. Maynard, W. Peters, and Y. Li. Metrics for evaluation of ontology-based information extraction. In *WWW 2006 Workshop on Evaluation of Ontologies for the Web (EON)*, Edinburgh, Scotland, 2006.
- [Maynard *et al.* 07a]
D. Maynard, W. Peters, M. d’Aquin, and M. Sabou. Change management for metadata evolution. In *ESWC International Workshop on Ontology Dynamics (IWOD)*, Innsbruck, Austria, June 2007.
- [Maynard *et al.* 07b]
D. Maynard, H. Saggion, M. Yankova, K. Bontcheva, and W. Peters. Natural Language Technology for Information Integration in Business Intelligence. In *10th International Conference on Business Information Systems (BIS-07)*, Poznan, Poland, 25-27 April 2007.
- [Maynard *et al.* 08a]
D. Maynard, W. Peters, and Y. Li. Evaluating evaluation metrics for ontology-based applications: Infinite reflection. In *Proc. of 6th International Conference on Language Resources and Evaluation (LREC)*, Marrakech, Morocco, 2008.
- [Maynard *et al.* 08b]
D. Maynard, Y. Li, and W. Peters. NLP Techniques for Term Extraction and Ontology Population. In P. Buitelaar and P. Cimiano, editors, *Bridging the Gap between Text and Knowledge - Selected Contributions to Ontology Learning and Population from Text*. IOS Press, 2008.
- [McEnery *et al.* 00]
A. McEnery, P. Baker, R. Gaizauskas, and H. Cunningham. EMILLE: Building a Corpus of South Asian Languages. *Vivek, A Quarterly in Artificial Intelligence*, 13(3):23–32, 2000.
- [Pastra *et al.* 02]
K. Pastra, D. Maynard, H. Cunningham, O. Hamza, and Y. Wilks. How feasible is the reuse of grammars for named entity recognition? In *Proceedings of the 3rd Language Resources and Evaluation Conference*, 2002. <http://gate.ac.uk/sale/lrec2002/reusability.ps>.

[Peters *et al.* 98]

W. Peters, H. Cunningham, C. McCauley, K. Bontcheva, and Y. Wilks. Uniform Language Resource Access and Distribution. In *Workshop on Distributing and Accessing Lexical Resources at Conference on Language Resources Evaluation*, Granada, Spain, 1998.

[Polajnar *et al.* 05]

T. Polajnar, V. Tablan, and H. Cunningham. User-friendly ontology authoring using a controlled language. Technical Report CS Report No. CS-05-10, University of Sheffield, Sheffield, UK, 2005.

[Porter 80]

M. Porter. An algorithm for suffix stripping. *Program*, 14(3):130–137, 1980.

[Ramshaw & Marcus 95]

L. Ramshaw and M. Marcus. Text Chunking Using Transformation-Based Learning. In *Proceedings of the Third ACL Workshop on Very Large Corpora*, 1995.

[Saggion & Gaizauskas 04a]

H. Saggion and R. Gaizauskas. Mining on-line sources for definition knowledge. In *Proceedings of the 17th FLAIRS 2004*, Miami Beach, Florida, USA, May 17-19 2004. AAAI.

[Saggion & Gaizauskas 04b]

H. Saggion and R. Gaizauskas. Multi-document summarization by cluster/profile relevance and redundancy removal. In *Proceedings of the Document Understanding Conference 2004*. NIST, 2004.

[Saggion & Gaizauskas 05]

H. Saggion and R. Gaizauskas. Experiments on statistical and pattern-based biographical summarization. In *Proceedings of EPIA 2005*, pages 611–621, 2005.

[Saggion 04]

H. Saggion. Identifying definitions in text collections for question answering. lrec. In *Proceedings of Language Resources and Evaluation Conference*. ELDA, 2004.

[Saggion 06]

H. Saggion. Multilingual Multidocument Summarization Tools and Evaluation. In *Proceedings of LREC 2006*, 2006.

[Saggion 07]

H. Saggion. Shef: Semantic tagging and summarization techniques applied to cross-document coreference. In *Proceedings of SemEval 2007, Association for Computational Linguistics*, pages 292–295, June 2007.

[Saggion *et al.* 02a]

H. Saggion, H. Cunningham, K. Bontcheva, D. Maynard, C. Ursu, O. Hamza, and Y. Wilks. Access to Multimedia Information through Multisource and Multilanguage

- Information Extraction. In *Proceedings of the 7th Workshop on Applications of Natural Language to Information Systems (NLDB 2002)*, Stockholm, Sweden, 2002.
- [Saggion *et al.* 02b]
H. Saggion, H. Cunningham, D. Maynard, K. Bontcheva, O. Hamza, C. Ursu, and Y. Wilks. Extracting Information for Information Indexing of Multimedia Material. In *Proceedings of 3rd Language Resources and Evaluation Conference (LREC'2002)*, 2002. http://gate.ac.uk/sale/lrec2002/mumis_lrec2002.ps.
- [Saggion *et al.* 03a]
H. Saggion, K. Bontcheva, and H. Cunningham. Robust Generic and Query-based Summarisation. In *Proceedings of the European Chapter of Computational Linguistics (EACL), Research Notes and Demos*, 2003.
- [Saggion *et al.* 03b]
H. Saggion, H. Cunningham, K. Bontcheva, D. Maynard, O. Hamza, and Y. Wilks. Multimedia Indexing through Multisource and Multilingual Information Extraction; the MUMIS project. *Data and Knowledge Engineering*, 48:247–264, 2003.
- [Saggion *et al.* 03c]
H. Saggion, J. Kuper, H. Cunningham, T. Declerck, P. Wittenburg, M. Puts, F. De-Jong, and Y. Wilks. Event-coreference across Multiple, Multi-lingual Sources in the Mumis Project. In *Proceedings of the European Chapter of Computational Linguistics (EACL), Research Notes and Demos*, 2003.
- [Saggion *et al.* 07]
H. Saggion, A. Funk, D. Maynard, and K. Bontcheva. Ontology-based information extraction for business applications. In *Proceedings of the 6th International Semantic Web Conference (ISWC 2007)*, Busan, Korea, November 2007.
- [Scott & Gaizauskas. 00]
S. Scott and R. Gaizauskas. The University of Sheffield TREC-9 Q&A System. In *In Proceedings of the 9th Text REtrieval Conference*, 2000.
- [Shaw & Garlan 96]
M. Shaw and D. Garlan. *Software Architecture*. Prentice Hall, New York, 1996.
- [Stevenson *et al.* 98]
M. Stevenson, H. Cunningham, and Y. Wilks. Sense tagging and language engineering. In *Proceedings of the 13th European Conference on Artificial Intelligence (ECAI-98)*, pages 185–189, Brighton, U.K., 1998.
- [Tablan *et al.* 02]
V. Tablan, C. Ursu, K. Bontcheva, H. Cunningham, D. Maynard, O. Hamza, T. McEnery, P. Baker, and M. Leisher. A Unicode-based Environment for Creation and Use of Language Resources. In *3rd Language Resources and Evaluation Conference*, Las Palmas, Canary Islands – Spain, 2002. ELRA. <http://gate.ac.uk/sale/ies103/ies103.pdf>.

[Tablan *et al.* 03]

V. Tablan, K. Bontcheva, D. Maynard, and H. Cunningham. Ollie: on-line learning for information extraction. In *SEALTS '03: Proceedings of the HLT-NAACL 2003 workshop on Software engineering and architecture of language technology systems*, volume 8, pages 17–24, Morristown, NJ, USA, 2003. Association for Computational Linguistics. <http://gate.ac.uk/sale/hlt03/ollie-sealts.pdf>.

[Tablan *et al.* 06a]

V. Tablan, T. Polajnar, H. Cunningham, and K. Bontcheva. User-friendly ontology authoring using a controlled language. In *5th Language Resources and Evaluation Conference (LREC)*, Genoa, Italy, May 2006. ELRA.

[Tablan *et al.* 06b]

V. Tablan, W. Peters, D. Maynard, H. Cunningham, and K. Bontcheva. Creating tools for morphological analysis of sumerian. In *5th Language Resources and Evaluation Conference (LREC)*, Genoa, Italy, May 2006. ELRA.

[Tablan *et al.* 08]

V. Tablan, D. Damjanovic, and K. Bontcheva. A natural language query interface to structured information. In *Proceedings of the 5th European Semantic Web Conference (ESWC 2008)*, volume 5021 of *Lecture Notes in Computer Science*, pages 361–375, Tenerife, Spain, June 2008. Springer-Verlag New York Inc.

[Ursu *et al.* 05]

C. Ursu, T. Tablan, H. Cunningham, and B. Popav. Digital media preservation and access through semantically enhanced web-annotation. In *Proceedings of the 2nd European Workshop on the Integration of Knowledge, Semantic and Digital Media Technologies (EWIMT 2005)*, London, UK, December 01 2005.

[van Rijsbergen 79]

C. van Rijsbergen. *Information Retrieval*. Butterworths, London, 1979.

[Wang *et al.* 05]

T. Wang, D. Maynard, W. Peters, K. Bontcheva, and H. Cunningham. Extracting a domain ontology from linguistic resource based on relatedness measurements. In *Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence (WI 2005)*, pages 345–351, Compiègne, France, September 2005.

[Wang *et al.* 06]

T. Wang, Y. Li, K. Bontcheva, H. Cunningham, and J. Wang. Automatic Extraction of Hierarchical Relations from Text. In *Proceedings of the Third European Semantic Web Conference (ESWC 2006)*, Budva, Montenegro, 2006.

[Witten & Frank 99]

I. H. Witten and E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*. Morgan Kaufmann, 1999.

[Wood *et al.* 03]

M. M. Wood, S. J. Lydon, V. Tablan, D. Maynard, and H. Cunningham. Using parallel texts to improve recall in IE. In *Recent Advances in Natural Language Processing*, Bulgaria, 2003.

[Wood *et al.* 04]

M. Wood, S. Lydon, V. Tablan, D. Maynard, and H. Cunningham. Populating a Database from Parallel Texts using Ontology-based Information Extraction. In *Proceedings of NLDB 2004*, 2004. <http://gate.ac.uk/sale/nldb2004/NLDB.pdf>.

[Yourdon 89]

E. Yourdon. *Modern Structured Analysis*. Prentice Hall, New York, 1989.

[Yourdon 96]

E. Yourdon. *The Rise and Resurrection of the American Programmer*. Prentice Hall, New York, 1996.

Colophon

Formal semantics (henceforth FS), at least as it relates to computational language understanding, is in one way rather like connectionism, though without the crucial prop Sejnowski's work (1986) is widely believed to give to the latter: both are old doctrines returned, like the Bourbons, having learned nothing and forgotten nothing. But FS has nothing to show as a showpiece of success after all the intellectual groaning and effort.

On Keeping Logic in its Place (in *Theoretical Issues in Natural Language Processing*, ed. Wilks), Yorick Wilks, 1989 (p.130).

We used \LaTeX to produce this document, along with TeX4HT for the HTML production. Thank you Don Knuth, Leslie Lamport and Eitan Gurari.