

Information Extraction, Automatic

Hamish Cunningham

*Department of Computer Science, University of Sheffield, Regent Court, 211
Portobello Street, Sheffield S1 4DP, UK. hamish@dcs.shef.ac.uk*

Contents

| | | |
|-----|--|----|
| 1 | Introduction: Extraction and Retrieval | 2 |
| 2 | Language Computation Without Understanding | 3 |
| 2.1 | Complexity vs. Specificity | 4 |
| 3 | Application Scenarios | 5 |
| 3.1 | Financial Analysts | 6 |
| 3.2 | Marketing Strategists | 6 |
| 3.3 | PR Workers | 6 |
| 3.4 | Media Analysts | 6 |
| 4 | Five Types of IE | 7 |
| 4.1 | Named Entity Recognition | 8 |
| 4.2 | Coreference Resolution | 9 |
| 4.3 | Template Element Production | 10 |
| 4.4 | Template Relation Production | 11 |
| 4.5 | Scenario Template Extraction | 12 |
| 4.6 | An Extended Example | 13 |
| 4.7 | Multilingual Extraction | 16 |
| 5 | IE After MUC | 17 |
| 5.1 | Portable IE | 17 |
| 5.2 | ACE: Automatic Content Extraction | 18 |
| 5.3 | Ontology-Based IE | 19 |

Abstract

This article describes Information Extraction (IE), the process of deriving disambiguated quantifiable data from natural language texts in service of some pre-specified precise information need. The article: covers the origins of IE and the factors relevant to its deployment
Preprint submitted to Elsevier Science *18th November 2004*

in applications contexts; presents scenarios in which the technology has been applied; breaks down the task into five subtasks and defines them; looks at recent developments in the field.

Keywords

Information Extraction, Natural Language Analysis, Human Language Technology, Natural Language Engineering.

ELL2 Cross-References

“Natural Language Analysis”. “Text Retrieval Conference and Message Understanding Conference”. “Natural Language Processing, System Evaluation”. “Human Language Technology”. “Language Processing, Statistical Methods”

Acknowledgements

Thanks to Valentin Tablan, Diana Maynard and Kalina Bontcheva for important contributions to this article. The author was partly supported by EPSRC grant GR/N15764/01 (AKT) and by EU grants SEKT, PrestoSpace and KnowledgeWeb.

1 Introduction: Extraction and Retrieval

Information Extraction (IE) is a technology based on analysing natural language in order to extract snippets of information. The process takes texts (and sometimes speech) as input and produces fixed-format, unambiguous data as output. This data may be used directly for display to users, or may be stored in a database or spreadsheet for later analysis, or may be used for indexing purposes in Information Retrieval (IR) applications such as Internet search engines like Google.

IE is quite different from IR:

- an IR system finds relevant texts and presents them to the user;
- an IE application analyses texts and presents only the specific information from them that the user is interested in.

For example, a user of an IR system wanting information on trade group formations in agricultural commodities markets would enter a list of relevant words and receive in return

a set of documents (e.g. newspaper articles) which contain likely matches. The user would then read the documents and extract the requisite information themselves. They might then enter the information in a spreadsheet and produce a chart for a report or presentation. In contrast, an IE system would automatically populate the spreadsheet directly with the names of relevant companies and their groupings.

There are advantages and disadvantages with IE in comparison to IR. IE systems are more difficult and knowledge-intensive to build, and are to varying degrees tied to particular domains and scenarios. IE is more computationally intensive than IR. However, in applications where there are large text volumes IE is potentially much more efficient than IR because of the possibility of dramatically reducing the amount of time people spend reading texts. Also, where results need to be presented in several languages, the fixed-format, unambiguous nature of IE results makes this relatively straightforward in comparison with providing the full translation facilities needed for interpretation of multilingual texts found by IR.

The rest of this article:

- discusses the origins of IE and the parameters that condition its performance in application settings (section 2);
- presents some applications scenarios in which IE is being deployed (section 3);
- gives a canonical definition of five subtasks that collectively characterise the spectrum of IE systems (section 4);
- looks at the development of IE research subsequent to the MUC (Message Understanding Conferences) programme of the late 20th century (section 5).

See also the article on “Natural Language Analysis”. Other overview sources on IE include: Cowie and Lehnert (1996); Appelt (1999); Cunningham (1999); Gaizauskas and Wilks (1998); Pazienza (2003).

2 Language Computation Without Understanding

Information Extraction in the sense discussed here grew out of work in the late 1980s and 1990s¹ around the Message Understanding Conferences (MUCs Grishman and Sundheim (1996); Sundheim (1995); Chinchor (1998) – see also the article “Text Retrieval Conference and Message Understanding Conference). These events were particularly distinctive because they employed a strict quantitative evaluation procedure where different research sites first defined a precise task, then implemented competing systems that were measured against human-annotated data in a controlled environment. Although this type of experimental procedure seems obvious in many other branches of science, its use prior to MUC was uncommon in many language processing contexts (perhaps due to the early association between the work

¹ Previous work in computation with human language had attempted similar things in some cases, and indeed a related term ‘fact extraction’ was in use as far back as the 1960s. A formal and widely accepted definition of IE did not emerge until the work discussed here, however.

and phenomenological fields such as linguistics, as opposed to empirical experimentation or engineering – Boguraev et al. (1995.)). The experimental cycle of the MUC competitions drove progress in IE and led to the canonical definition that is presented below in section 4. The series also contributed much of our understanding of how to measure the performance of IE systems (see also the article “Natural Language Processing, System Evaluation”).

The MUCs were held over the decade from 1987; following this IE work took several directions, including commercialisation of some of the basic techniques, additional work on portability (see section 5.1) and, most recently, connecting with efforts around the general project of the Semantic Web Berners-Lee (1999) (see section 5.3).

An important outcome of the last 15 years work on IE is that the community now understands reasonably precisely the various parameters that influence the performance of the technology in diverse applications settings. (In other words, the scientific work has laid the basis for the engineering of practical systems.) The next section looks at these parameters.

2.1 Complexity vs. Specificity

The difference between the general project of language analysis (or understanding) and the more restricted enterprise of extraction is a response to the low performance of analysis in the general case. To illustrate, imagine that we wish to build an application which uses language analysis to support some of its functions. Then consider that items of information to be extracted can vary in complexity (e.g. we may be interested simply in people names, or in complex events that involve multiple participants) and in specificity (e.g. we may be interested in general information reported in any way in any text, or in specific domains reported in particular ways in certain types of text). If we plot the acceptable performance level of analysis components on a graph of complexity vs. specificity of the information to be extracted (figure 1), then there is an obvious trade-off to be made between the two: the more complex the data to be extracted, the more specific must be the domain of discourse; the simpler the data, the more generally the extraction algorithms may be applied. This two-dimensional view of IE performance hides a number of subtle issues. Specificity of an IE task is in itself multi-dimensional, being influenced by both text type and genre. For example, a seminar announcement text (a very specific text type, easy for IE systems to analyse) may be harder or easier to process depending on genre. For example, if it is a seminar series about famous historical figures there may be ambiguity between the speakers’ names and those of their subjects. Conversely, if the series relates to mobile telecommunications, the regularity and specificity of the terminology typical of this genre should make structural analysis easier. Therefore we must also consider the parameters text type and subject (or genre, or domain):

Text type: the kinds of texts we are working with, for example: Wall Street Journal articles, or email messages, or novels, or the output of a speech recogniser.

Domain: the broad subject-matter of those texts, for example: world events, or symptoms of a deficiency in a particular enzyme, or seminar announcements, or financial news, or requests for technical support, or tourist information, and the style in which they are

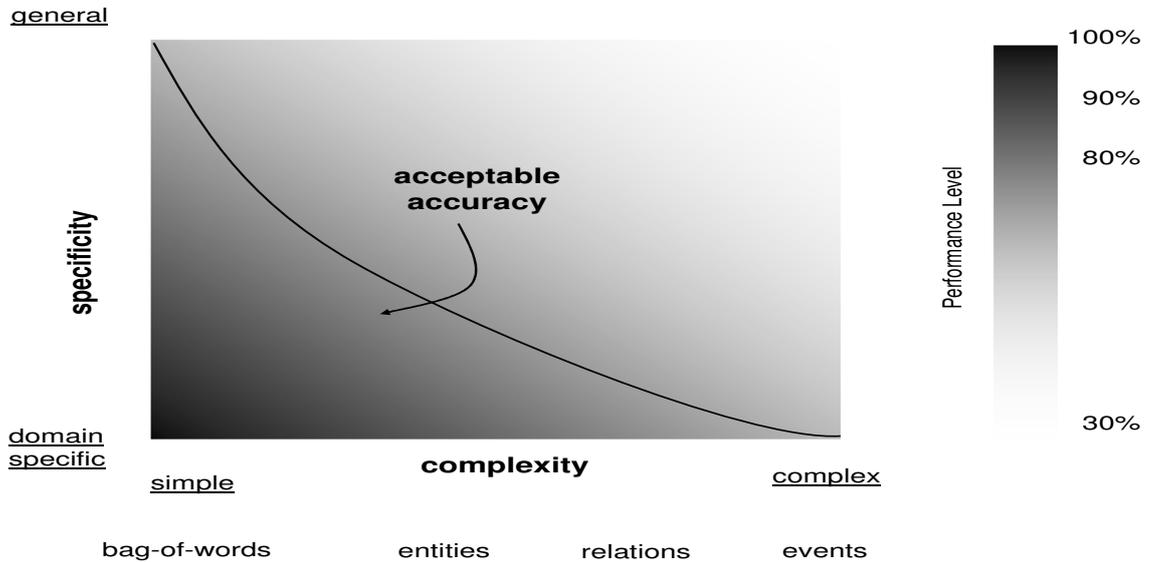


Figure 1. Performance trade-off relative to specificity and complexity.

written, e.g. informal, formal.

Finally there is the question of the particular types of information that the IE user is interested in, for example: personal names, or rocket launches, or mergers between companies, or problems experienced with a particular software package, or descriptions of how to locate parts of a city. We return to this issue in section 4.

3 Application Scenarios

This section contains some scenarios that illustrate how IE applications software can mediate between the text and the structured information needs of various types of user. In each case the user specifies an information need to IE development staff, who then prepare a custom extraction system. The input materials are increasingly defined as subsections of the World Wide Web, either the formal materials of news, company sites etc., or the informal Web of weblogs, mailing lists and wikis. The IE developers must analyse the problem according to the dimensions outlined in section 2.1 above, and determine how to combine human and machine processing so as to attain the required performance profile (see section 5.1 below).

The types of scenario given here are of increasing relevance due in large part to the explosive growth, dynamic nature and multilingual content of the Web.

3.1 Financial Analysts

The Web contains various indications of how a company is viewed, and whether or not it can be expected to perform strongly in the coming period. Some of this data is already highly analysed, e.g. financial news in English. Other data, in other languages, or from less well scrutinised sources, is both voluminous and obscure. IE can enable analysts to answer questions such as

‘How many instances predicting strong performance for a particular company are out there?’

‘Over the past year how has the profile of predictions for this company changed?’

‘How many positive/negative sentiments were expressed for the company?’

3.2 Marketing Strategists

The dynamic adjustment of marketing spend is currently made more difficult by the lack of metrics that indicate the impact of campaign elements. IE can support campaign tuning today based on yesterday’s results, producing outputs such as:

‘In this morning’s IT press 7% of articles discussed your company. The average proportion of the article directly relating to your company was 33%. The figures for the other key players in your sector are summarised in the following table....’

Similarly, we can measure the extent of media coverage relative to spend events:

‘Company Y exhibited at Comdex. In the week following the exhibition 20% of the press that covered Comdex mentioned Y.’

3.3 PR Workers

Public Relations staff are concerned to identify negative reporting events as quickly as possible in order to respond. IE can be configured to report like this:

‘The table below summarises 12 negative reporting events concerning your company in the last 24 hours of IT news....’

3.4 Media Analysts

IE can be used to create a range of media metrics, for example the ‘media distance’, or extent of collocational association, between concepts and products/companies:

‘The media distance between your company and the subject of XML is 0.09; for IBM the value is 0.2.’

4 Five Types of IE

By the time that it ended in 1998² the MUC programme had arrived at a definition of IE split into five tasks:

Named Entity recognition (NE)

Finds and classifies names, places, etc.

Coreference resolution (CO)

Identifies identity relations between entities.

Template Element construction (TE)

Adds descriptive information to NE results (using CO).

Template Relation construction (TR)

Finds relations between TE entities.

Scenario Template production (ST)

Fits TE and TR results into specified event scenarios.

In simpler terms: NE is about finding entities; CO about which entities and references (such as pronouns) refer to the same thing; TE about what attributes entities have; TR about what relationships between entities there are; ST about events that the entities participate in. Consider these sentences:

The shiny red rocket was fired on Tuesday. It is the brainchild of Dr. Big Head. Dr. Head is a staff scientist at We Build Rockets Inc.

NE discovers that the entities present are the *rocket*, *Tuesday*, *Dr. Head* and *We Build Rockets Inc.* CO discovers that *it* refers to the rocket. TE discovers that the rocket is *shiny red* and that it is Head’s *brainchild*. TR discovers that *Dr. Head* works for *We Build Rockets Inc.* ST discovers that there was a rocket launching event in which the various entities were involved.

The scenarios presented above in section 3 deploy these five types of information in various combinations (and with various typical performance profiles).

These various types of IE provide progressively higher-level information about texts. They are described in more detail below, with examples (see also the extended example in section 4.6).

² The final conference was MUC-7 SAIC (1998).

4.1 Named Entity Recognition

The simplest and most reliable IE technology is *Named Entity recognition* (NE – see also the article “Named Entity Extraction”). NE systems identify all the names of people, places, organisations, dates, amounts of money, etc. So, for example, if we run the text in figure 2³ through an NE recogniser, the result is as in figure 3. (The results shown here and below are from IE software distributed with the GATE system Cunningham (2002) – see also the article on Computational Language Systems, Architectures.)



Figure 2. An example text

All things being equal, NE recognition can be performed at up to around 95% accuracy (see also section 2.1). Given that human annotators do not perform to the 100% level (measured in MUC by inter-annotator comparisons), NE recognition can now be said to function at human performance levels, and applications of the technology are increasing rapidly as a result.

The process is weakly domain dependent, i.e. changing the subject-matter of the texts being processed from financial news to other types of news would involve some changes to the system, and changing from news to scientific papers would involve quite large changes.

³ The Independent, London, 4th August 2004.

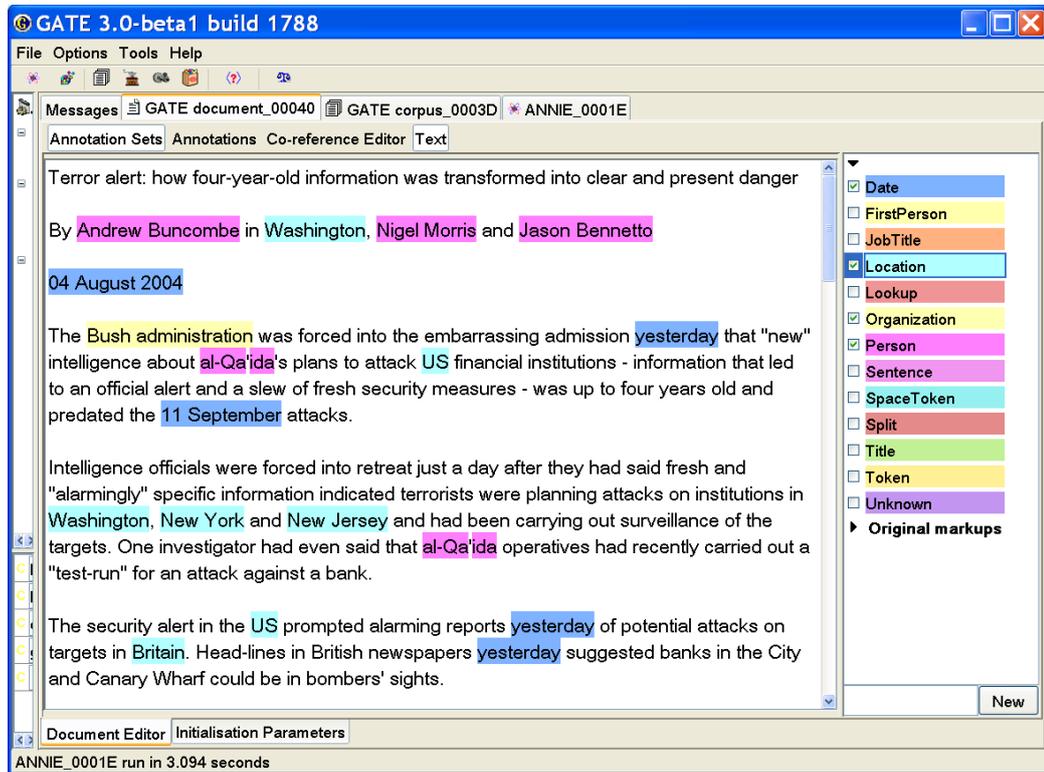


Figure 3. Named Entity recognition

4.2 Coreference Resolution

Coreference resolution (CO) involves identifying identity relations between entities in texts. These entities are both those identified by NE recognition and anaphoric references to those entities. For example, in

Alas, poor Yorick, I knew him Horatio.

coreference resolution would tie ‘Yorick’ with ‘him’ (and ‘I’ with Hamlet, if sufficient information was present in the surrounding text).

This process is less relevant to users than other IE tasks (i.e. whereas the other tasks produce output that is of obvious utility for the application user, this task is more relevant to the needs of the application developer). For text browsing purposes we might use CO to highlight all occurrences of the same object or provide hypertext links between them. CO technology might also be used to make links between documents. The main significance of this task, however, is as a building block for TE and ST (see below). CO enables the association of descriptive information scattered across texts with the entities to which it refers. To continue the hackneyed Shakespeare example, coreference resolution might allow TE or TR analysis to situate Yorick in Denmark. Figure 4 shows results for the example text from figures 2 and 3.

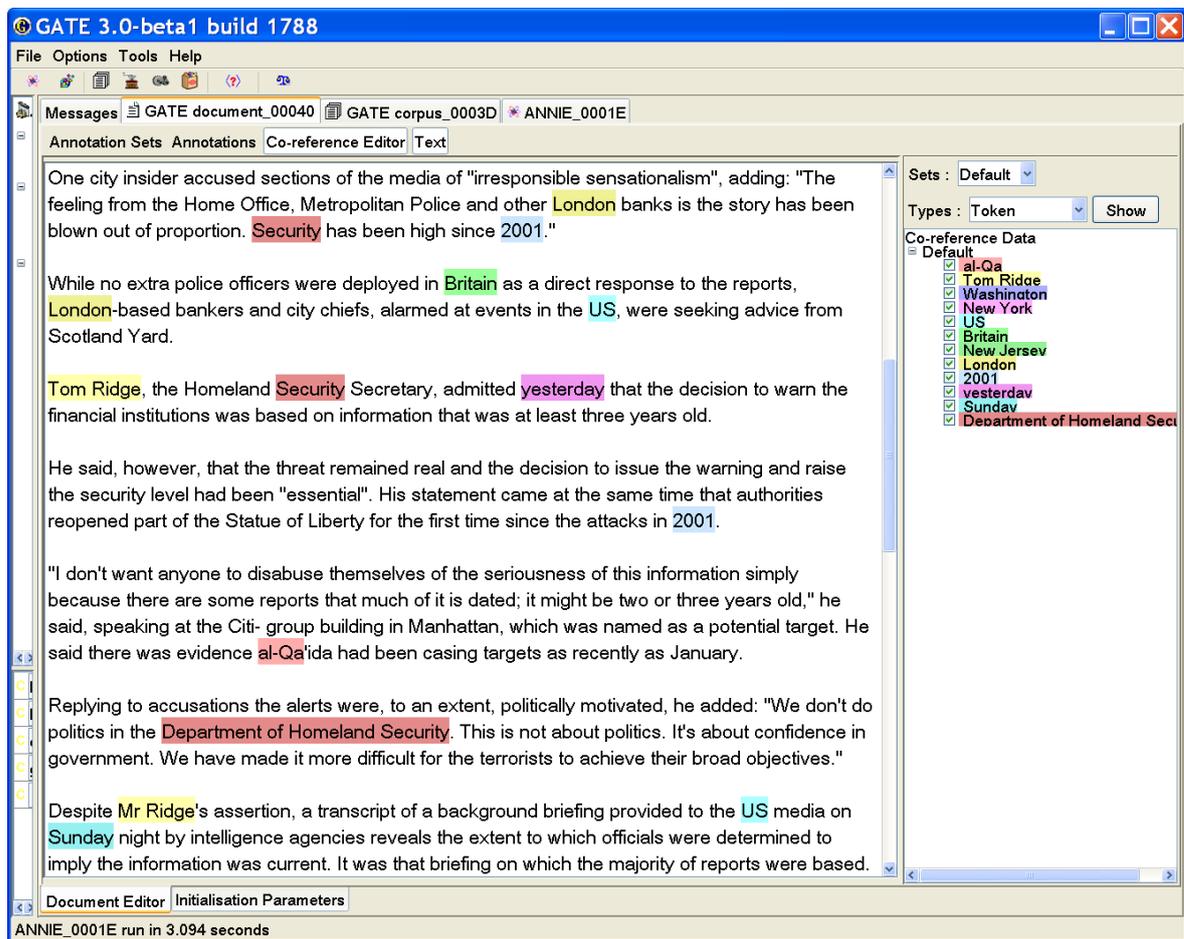


Figure 4. Coreference resolution

CO breaks down into two sub-problems: anaphoric resolution (e.g. 'I' with Hamlet); proper-noun resolution. Proper-noun coreference identification finds occurrences of same object represented with different spelling or compounding, e.g. 'IBM', 'IBM Europe', 'International Business Machines Ltd.', ...). CO resolution is an imprecise process, particularly when applied to the solution of anaphoric reference. CO results vary widely; depending on domain perhaps only 50-60% may be relied upon. CO systems are domain dependent.

4.3 Template Element Production

The TE task builds on NE recognition and coreference resolution, associating descriptive information with the entities. For example, from the figure 2 text the system finds out that the 'Bush administration' is also referred to as 'government officials', and adds this as an alias.

Template elements for the example text are given in figure 5. The format is an arbitrary one; the point to note is that it is essentially a database record, and could just as well be

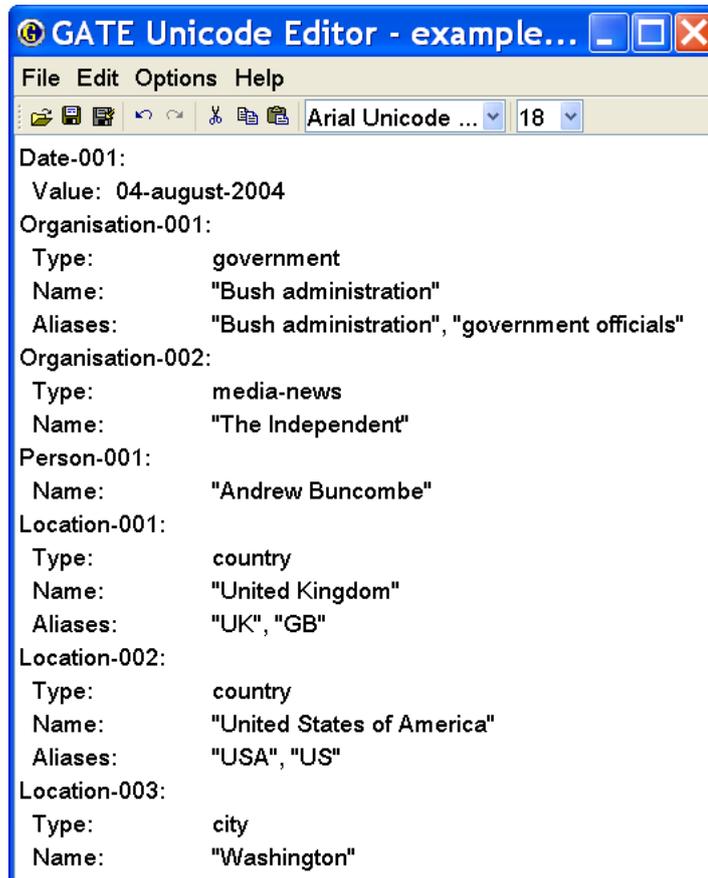


Figure 5. Template Elements

formatted for SQL⁴ store operations, or reading into a spreadsheet, or (with some extra processing) for multilingual presentation. See below for an extended example in a simplified format.

Good scores for TE systems are around 80% (on similar tasks humans can achieve results in the mid-90s, so there is some way to go). As in NE recognition, the production of TEs is weakly domain dependent, i.e. changing the subject-matter of the texts being processed from financial news to other types of news would involve some changes to the system, and changing from news to scientific papers would involve quite large changes.

4.4 *Template Relation Production*

Before MUC-7, relations between entities were part of the scenario-specific template outputs of IE evaluations. In order to capture more widely useful relations, MUC-7 introduced the TR task – see figure 6. As described in Appelt (1999), “The template relation task requires the identification of a small number of possible relations between the template elements identified

⁴ Structured Query Language – a common way for manipulating relational databases.

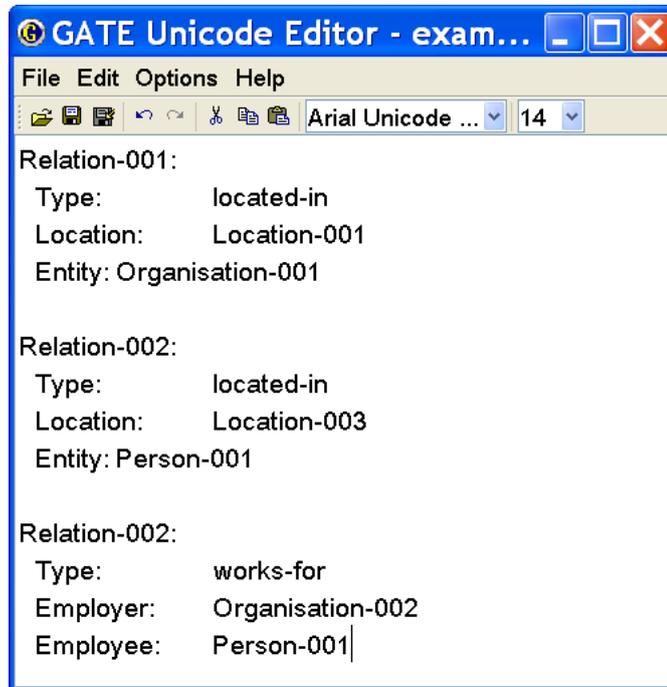


Figure 6. Template Relations

in the template element task. This might be, for example, an employee relationship between a person and a company, a family relationship between two persons, or a subsidiary relationship between two companies. Extraction of relations among entities is a central feature of almost any information extraction task, although the possibilities in real-world extraction tasks are endless.” In general good TR scores reach around 75%. TR is a weakly domain dependent task.

4.5 Scenario Template Extraction

Scenario templates (STs) are the prototypical outputs of IE systems, being the original task for which the term was coined. They tie together TE entities and TR relations into event descriptions. For example, TE may have identified Isabelle, Dominique and Françoise as people entities present in the Robert edition of Napoleon’s love letters. ST might then identify facts such as that Isabelle moved to Paris in August 1802 from Lyon to be nearer to the little chap, that Dominique then burnt down Isabelle’s apartment block and that Françoise ran off with one of Gerard Depardieu’s ancestors. A slightly more pertinent example is given in figure 7, adapted from MUC-6 ARPA (1995).

ST is a difficult IE task; the best MUC systems score around 60%. The human score can be as low as around 80+%, which illustrates the complexity involved. These figures should be taken into account when considering appropriate applications of ST technology. Note, however, that it is possible to increase precision at the expense of recall: we can develop ST systems that do not make many mistakes, but that miss quite a lot of occurrences of relevant

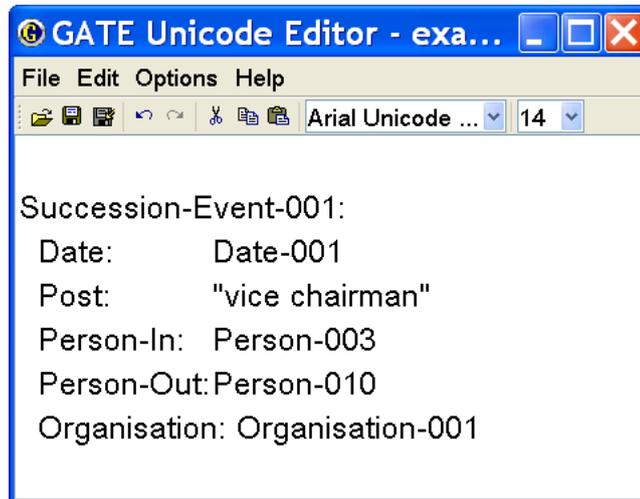


Figure 7. Scenario Template

scenarios. Alternatively we can push up recall and miss less, but at the expense of making more mistakes.

The ST task is both domain dependent and, by definition, tied to the scenarios of interest to the users. Note however that the results of NE, TR and TE feed into ST. Note also that in MUC-6 and MUC-7 the developers were given the specifications for the ST task only 1 month before the systems were scored. This was because it was noted that an IE system that required very lengthy revision to cope with new scenarios was of less worth than one that could meet new specifications relatively rapidly (see section 5.1). As a result of this, the scores for ST in MUC-6/7 were probably slightly lower than they might have been with a longer development period. Experience from previous MUCs and from subsequent work suggests, however, that current technology has difficulty attaining scores much above 60% accuracy for this task.

4.6 *An Extended Example*

So far we have discussed IE from a general perspective. In this section we look at the capabilities that might be delivered as part of an application designed to support analysts tracking international drug dealing.

When the system is specified, our imaginary analyst states that “the operational domains that user interests are centred around are... drug enforcement, money laundering, organised crime, terrorism, legislation”. The entities of interest within these domains are cited as “person, company, bank, financial entity, transportation means, locality, place, organisation, time, telephone, narcotics, legislation, activity”. A number of relations (or “links”) are also specified, for example between people, between people and companies, etc. These relations are not typed, i.e. the kind of relation involved is not specified. Some relations take the form of properties of entities – e.g. the location of a company – whilst others denote events – e.g.

a person visiting a ship.

Working from this starting point an IE system is designed that:

- is tailored to texts dealing with drug enforcement, money laundering, organised crime, terrorism, and legislation;
- recognises entities in those texts and assigns them to one of a number of categories drawn from the set of entities of interest (person, company, ...);
- associates certain types of descriptive information with these entities, e.g. the location of companies;
- identifies a set (relatively small to begin with) of events of interest by tying entities together into event relations.

For example, consider the following text:

Reuter – New York, Wednesday 12 July 2005.

New York police announced today the arrest of Frederick J. Thompson, head of Jay Street Imports Inc., on charges of drug smuggling. Thompson was taken from his Manhattan apartment in the early hours yesterday. His attorney, Robert Guliani, issued a statement denying any involvement with narcotics on the part of his client. “No way did Fred ever have dealings with dope”, Guliani said.

A Jay Street spokesperson said the company had ceased trading as of today. The company, a medium-sized import-export concern established in 1989, had been the main contractor in several collaborative transport ventures involving Latin-American produce. Several associates of the firm moved yesterday to distance themselves from the scandal, including the mid-western transportation company Downing-Jones.

Thompson is understood to be accused of importing heroin into the United States.

From this IE might produce information such as the following (in some format to be determined according to user requirements, e.g. SQL statements addressing some database schema).

First, a list of entities and associated descriptive information. Relations of property type are made explicit. Each entity has an id, e.g. ENTITY-2, which can be used for cross-referencing between entities and for describing events involving entities. Each also has a type, or category, e.g. company, person. Additionally various type-specific information is available, e.g., for dates, a normalisation giving the date in standard format.

```
Reuter
  id:          ENTITY-1
  type:        company
  business:    news
New York
```

id: ENTITY-2
 type: location
 subtype: city
 is_in: US
 Wednesday 12 July 2005
 id: ENTITY-3
 type: date
 normalisation: 12/07/2005
 New York police
 id: ENTITY-4
 type: organisation
 location: ENTITY-2
 Frederick J. Thompson
 id: ENTITY-5
 type: person
 aliases: Thompson; Fred
 domicile: ENTITY-7
 profession: managing director
 employer: ENTITY-6
 Jay Street Imports Inc.
 id: ENTITY-6
 type: organisation
 aliases: Jay Street
 business: import-export
 Manhattan
 id: ENTITY-7
 type: location
 subtype: city
 is_in: ENTITY-2
 Robert Guliani
 id: ENTITY-8
 type: person
 aliases: Guliani
 1989
 id: ENTITY-9
 type: date
 normalisation: ??/??/1989
 Latin-America
 id: ENTITY-10
 type: location
 subtype: country
 Downing-Jones
 id: ENTITY-11
 type: organisation

```

    business:      transportation
heroin
    id:           ENTITY-12
    type:         drug
    class:        A
United States
    id:           ENTITY-13
    type:         location
    subtype:      country

```

(These results correspond to the combination of NE and TE tasks; if we removed all but the type slots we would be left with the NE data.)

Second, relations of event type, or scenarios:

```

narcotics-smuggling
    id:           EVENT-1
    destination:  ENTITY-13
    source:       unknown
    perpetrators: ENTITY-5, ENTITY-6
    status:       on-trial
joint-venture
    id:           EVENT-2
    type:         transport
    companies:    ENTITY-6, ENTITY-11
    status:       past

```

(These results correspond to the ST task.)

4.7 Multilingual Extraction

The results described above may be translated for presentation to the user or for storage in existing databases. In general this task is much easier than translation of ordinary text, and is close to software localisation, the process of making a program's messages and labels on menus and buttons multilingual. Localisation involves storing lists of direct translations for known items. In our case these lists would store translations for words such as "entity", "location", "date", "heroin". We also need ways to display dates and numbers in local formats, but code libraries are available for this type of problem.

Problems can arise where arbitrary pieces of text are used in the entity description structures, for example the descriptor slot in MUC-6 TE objects. Here a noun phrase from the text is extracted, with whatever qualifiers, relative clauses etc. happen to be there, so the language is completely unrestricted and would need a full translation mechanism.

5 IE After MUC

In this section we bring the story up-to-date by looking at developments in IE system portability, the ACE programme which succeeded MUC, and the project of automated annotation for Semantic Web applications.

5.1 Portable IE

A particular IE application might be configured to process financial news articles from a particular news provider (written in the house style) and find information about mergers between companies and various other scenarios. The performance of the application would be predictable for only this conjunction of factors. If it was later required to extract facts from the love letters of Napoleon Bonaparte as published on wall posters in the 1871 Paris Commune, performance levels would no longer be predictable. Tailoring an IE system to new requirements is a task that varies in scale dependent on the degree of variation in the parameters discussed in section 2.1.

Therefore a central track of IE research addresses the issue of *portability*, and this has been one of the most fruitful areas of development since the end of the MUC programme in the late 1990s.

We can distinguish three broad currents of work in this area:

- (1) Learning extraction rules of models from annotated examples.
- (2) Embedding learning systems within end-user systems.
- (3) Supporting the development of rules/models by skilled specialist staff.

The first approach is to learn part or all of the extraction system from annotated training data. The advantage is a reduction in the need for skilled staff to perform system porting. The disadvantages are:

- only simple data can be extracted, or complex data from simple texts, such as seminar announcements (in fact many of the algorithms currently common in this areas were developed for screen scraping⁵, which is a simpler task than most language analysis);
- large volumes of training data may be required.

Surveys of work in this area are to be found in Cardie (1997); Daelemans and Osborne (2003).

Secondly, researchers have attempted to embed learning within end-user tools where the

⁵ Screen scraping is the process of deleting the presentational elements of data or text, e.g. as displayed on a Web page, in order to allow further processing of the data. Sites that provide comparative shopping lists are often based on scraping the pages of competing suppliers and re-presenting the data for comparison purposes.

users correct IE suggestions. This approach addresses the problem of the costs of producing training data associated with the learning approach by speeding up the annotation process. A cyclical method known as *mixed-initiative learning* is used, where the user begins by doing all the annotation work manually while the system learns in the background. When the quality of the learnt models are high enough, the system can then propose annotations to the user; correction of mistakes is fed back into the learning algorithm. See Day et al. (1997); Grishman (2001). Mixed initiative learning was reinvented as ‘active learning’ in later work.

Lastly, IE system portability can be maximised by providing a development environment for skilled staff to adapt a core system. The advantages are:

- the core system can be designed for robustness and portability;
- extraction data complexity is not limited by a learning algorithm;
- all the engineering aspects of the process can be taken care of by the infrastructure (from data visualisation to performance evaluation).

The disadvantage is that the adaptation process is labour intensive, and it is difficult for end-users to acquire the necessary skills. A survey of work is available in Cunningham and Scott (2004); see also the article “Computational Language Systems, Architectures”.

5.2 ACE: Automatic Content Extraction

The Automatic Content Extraction programme (ACE, ACE (2004); Maynard et al. (2003)) is a successor to MUC that has been running since a pilot study in 1999 and which has continued the competitive quantitative evaluation cycles of its predecessor. ACE differs from MUC in three significant ways:

- (1) Several of the five MUC tasks defined in section 4 are conflated in ACE. The NE and CO tasks become a single ‘Entity Detection and Tracking’ (EDT) task in ACE, and the TE and TR tasks a single ‘Relation Detection and Tracking’ task. The ST task is renamed ‘Event Detection and Characterisation’.
- (2) The ACE tasks are more complex than their MUC counterparts in several ways. In the EDT task, for example, there is a more fine-grained taxonomy of entities, and it is necessary for systems to interpret metonymic entity mentions, requiring semantic analysis of the texts under consideration. Also, multiple domains and sources are used, including materials output from automatic speech transcription and optical character recognition programs.
- (3) The evaluation results from the ACE programme are not public. Whereas in MUC all the competition results were made public, ACE results are restricted to the participants. The utility of the programme for non-participants is therefore much lower than for MUC.

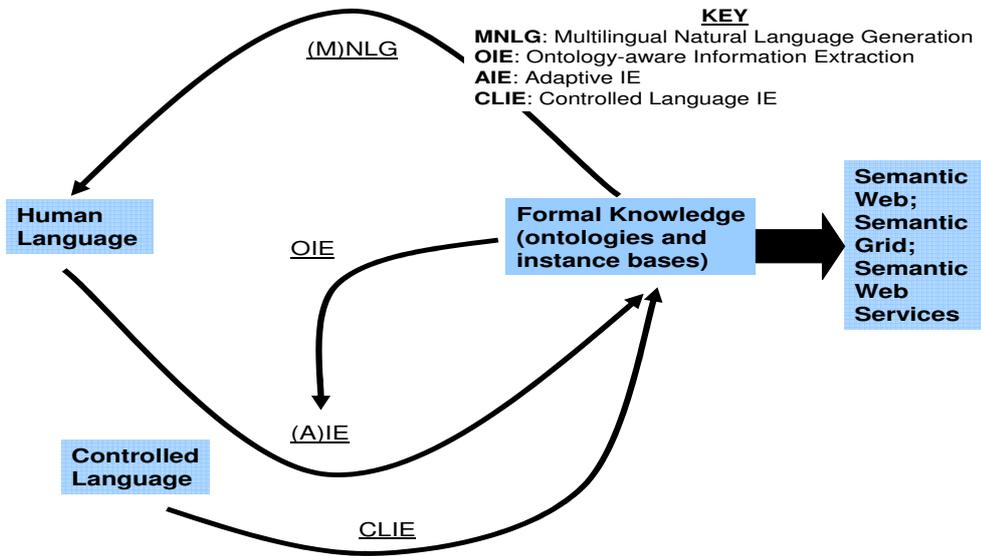


Figure 8. Closing the language loop

5.3 Ontology-Based IE

The Semantic Web aims to add a machine tractable, re-purposeable layer of annotation relative to ontologies to complement the existing web of natural language hypertext Fensel et al. (2002); Davies et al. (2002); Bechhofer et al. (2003). In order to realise this vision, the creation of semantic annotation, the linking of web pages to ontologies, and the creation, evolution and interrelation of ontologies must become automatic or semi-automatic processes, and a significant body of recent work has looked at the application of Ontology-Based IE (OBIE – Bontcheva (2004)) in this context. Figure 8 illustrates the way in which IE and other language technologies can be used to bring together the natural language upon which the current web is mainly based and the formal knowledge needed for a Semantic Web. Figure 9 illustrates OBIE in action: the ontology hierarchy appears on the right, with an annotated document in the centre of the pane. OBIE poses two main challenges:

- the identification of concept instances from the ontology in the text;
- the automatic population of ontologies with instances in the text.

5.3.1 Identification of instances from the ontology

If the ontology in question is already populated with instances, the task of an OBIE system may be simply to identify instances from the ontology in the text. Similar methodologies can be used for this as for traditional IE systems, using an ontology rather than a flat gazetteer. For rule-based systems, this is relatively straightforward. For learning-based systems, however, this is more problematic because training data is required. Collecting such training data is, however, likely to be a bottleneck. Unlike traditional IE systems for which training

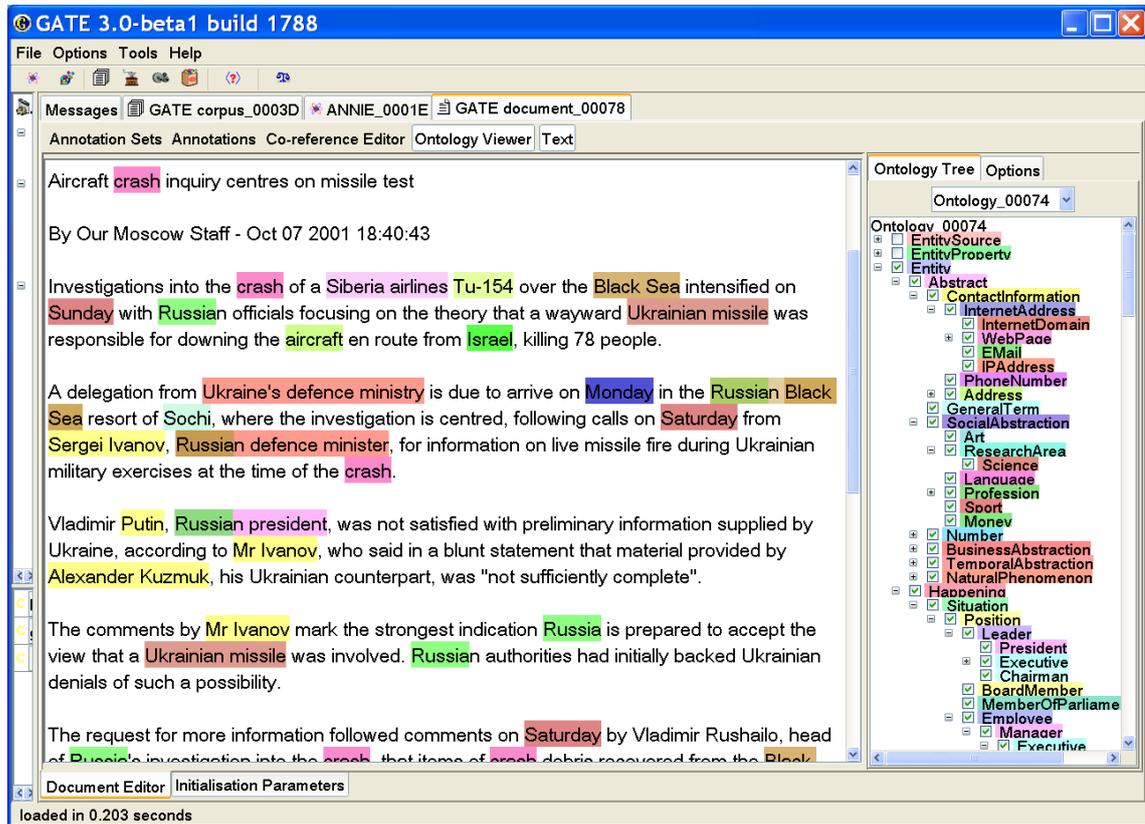


Figure 9. Ontology-based IE

data exists in domains like news texts in plentiful form, thanks to efforts from MUC, ACE and other programmes, there is a dearth of material currently available for Semantic Web applications. New training data needs to be created manually or semi-automatically, which is a time-consuming and onerous task, although systems to aid such metadata creation are currently being developed.

5.3.2 Automatic ontology population

In this case an OBIE application identifies instances in the text belonging to concepts in the ontology, and adds these instances to the ontology in the correct location. It is important to note that instances may appear in more than one location in the ontology, because of the multidimensional nature of many ontologies and/or ambiguity which cannot or should not be resolved at this level.

5.3.3 Example Systems

The Knowledge and Information Management system (KIM – Popov et al. (2004)) is an extendable platform for knowledge management which offers IE-based facilities for metadata creation, storage, and semantic-based search. It also includes a set of front-ends for online use, that offer semantically enhanced browsing.

Magpie Domingue et al. (2004) is a browser add-in that uses IE to facilitate the interpretation of webpages and ‘collaborative sense-making’. It annotates webpages with metadata in a fully automatic fashion and needs no manual intervention. It automatically populates an ontology from relevant web sources, and can be used with different ontologies. The principle behind it is that it uses an ontology to provide a very specific and personalised viewpoint of the webpages the user wishes to browse.

The SemTag system Dill et al. (2003) uses IE to perform large-scale semantic annotation with respect to the TAP ontology. It first performs a lookup phase annotating all possible mentions of instances from the TAP ontology. In the second, disambiguation phase, SemTag uses a vector-space model to assign the correct ontological class or determine that this mention does not correspond to a class in TAP.

References

- ACE, Feb 2004. Annotation Guidelines for Entity Detection and Tracking (EDT). Available at <http://www ldc.upenn.edu/Projects/ACE/>.
- Appelt, D., 1999. An Introduction to Information Extraction. *Artificial Intelligence Communications* 12 (3), 161–172.
- ARPA, 1995. Proceedings of the Sixth Message Understanding Conference (MUC-6). Defense Advanced Research Projects Agency, Morgan Kaufmann, California.
- Bechhofer, S., van Harmelen, F., Hendler, J., Horrocks, I., McGuinness, D. L., Patel-Schneider, P. F., Stein, L. A., 2003. OWL Web Ontology Language Reference. Tech. rep., W3C Proposed Recommendation 15 December 2003, <http://www.w3.org/TR/2003/PR-owl-ref-20031215/>.
- Berners-Lee, T., 1999. *Weaving the Web*. Orion Business Books.
- Boguraev, B., Garigliano, R., Tait, J., 1995. Editorial. *Natural Language Engineering*. 1, Part 1.
- Bontcheva, K., 2004. 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).
- Cardie, C., 1997. Empirical Methods in Information Extraction. *AI Magazine* 18 (4).
- Chinchor, N. A., April 1998. Overview of Proceedings of the Seventh Message Understanding Conference (MUC-7)/MET-2. In: Proceedings of the Seventh Message Understanding Conference (MUC-7). Fairfax, VA, p. 5 pages, http://www.itl.nist.gov/iaui/894.02/related_projects/muc/.
- Cowie, J., Lehnert, W., 1996. Information Extraction. *Communications of the ACM* 39 (1), 80–91.
- Cunningham, H., May 1999. *Information Extraction: a User Guide* (revised version). Research Memorandum CS-99-07, Department of Computer Science, University of Sheffield.
- Cunningham, H., 2002. GATE, a General Architecture for Text Engineering. *Computers and the Humanities* 36, 223–254.
- Cunningham, H., Scott, D. (Eds.), 2004. Special Issue of *Natural Language Engineering* on

- Software Architecture for Language Engineering. Cambridge University Press.
- Daelemans, W., Osborne, M. (Eds.), May 2003. CoNLL-2003, 7th Conference on Computational Natural Language Learning. Edmonton, Canada.
- Davies, J., Fensel, D., van Harmelen, F. (Eds.), 2002. Towards the Semantic Web: Ontology-driven Knowledge Management. Wiley.
- Day, D., Aberdeen, J., Hirschman, L., Koziarok, R., Robinson, P., Vilain, M., 1997. Mixed-Initiative Development of Language Processing Systems. In: Proceedings of the 5th Conference on Applied Natural Language Processing (ANLP-97).
- Dill, S., Eiron, N., Gibson, D., Gruhl, D., Guha, R., Jhingran, A., Kanungo, T., Rajagopalan, S., Tomkins, A., Tomlin, J. A., Zien, J. Y., 2003. SemTag and Seeker: Bootstrapping the semantic web via automated semantic annotation. In: Proceedings of WWW'03.
- Domingue, J., Dzbor, M., Motta, E., 2004. Magpie: Supporting Browsing and Navigation on the Semantic Web. In: Nunes, N., Rich, C. (Eds.), Proceedings ACM Conference on Intelligent User Interfaces (IUI). pp. 191–197.
- Fensel, D., Hendler, J., Wahlster, W., Lieberman, H. (Eds.), 2002. Spinning the Semantic Web: Bringing the World Wide Web to Its Full Potential. MIT Press.
- Gaizauskas, R., Wilks, Y., 1998. Information Extraction: Beyond Document Retrieval. *Journal of Documentation* 54 (1), 70–105.
- Grishman, R., 2001. Adaptive Information Extraction and Sublanguage Analysis. In: Proceedings of Workshop on Adaptive Text Extraction and Mining at Seventeenth International Joint Conference on Artificial Intelligence. Seattle, USA.
- Grishman, R., Sundheim, B., Jun. 1996. Message understanding conference - 6: A brief history. In: Proceedings of the 16th International Conference on Computational Linguistics. Copenhagen.
- Maynard, D., Bontcheva, K., Cunningham, H., 2003. Towards a semantic extraction of Named Entities. In: Recent Advances in Natural Language Processing. Bulgaria.
- Pazienza, M. T. (Ed.), 2003. Information Extraction in the Web Era. Springer-Verlag.
- Popov, B., Kiryakov, A., Kirilov, A., Manov, D., Ognyanoff, D., Goranov, M., 2004. KIM – Semantic Annotation Platform. Natural Language Engineering.
- SAIC, 1998. Proceedings of the Seventh Message Understanding Conference (MUC-7), http://www.itl.nist.gov/iaui/894.02/related_projects/muc/index.html.
- Sundheim, B. (Ed.), 1995. Proceedings of the Sixth Message Understanding Conference (MUC-6). ARPA, Morgan Kaufmann, Columbia, MD.