

Advanced GATE Embedded

Module 5

Seventh GATE Training Course
June 2014

© 2014 The University of Sheffield

This material is licenced under the Creative Commons

Attribution-NonCommercial-ShareAlike Licence

(<http://creativecommons.org/licenses/by-nc-sa/3.0/>)

Outline

- 1** GATE in Multi-threaded/Web Applications
 - Multi-threading and GATE
 - Servlet Example
 - The Spring Framework
 - Making your own PRs duplication-friendly
- 2** Extending GATE
 - Adding new document formats

Outline

- 1** GATE in Multi-threaded/Web Applications
 - Multi-threading and GATE
 - Servlet Example
 - The Spring Framework
 - Making your own PRs duplication-friendly
- 2** Extending GATE
 - Adding new document formats

Introduction

- Scenario:
 - Implementing a web application that uses GATE Embedded to process requests.
 - Want to support multiple concurrent requests
 - Long running process - need to be careful to avoid memory leaks, etc.
- Example used is a plain HttpServlet
 - Principles apply to other frameworks (struts, Spring MVC, Metro/CXF, Grails...)

Setting up

- GATE libraries in `WEB-INF/lib`
 - `gate.jar` + JARs from lib
- Usual GATE Embedded requirements:
 - A directory to be "gate.home"
 - Site and user config files
 - Plugins directory
- Alternatively use `Gate.runInSandbox`, but certain things can only be configured from the `gate.xml` files.

GATE in a Multi-threaded Environment

- GATE initialization needs to happen once (and only once) before any other GATE APIs are used.
- The `Factory` is synchronized internally, so safe for use in multiple threads.
- Individual PRs/controllers are *not* safe – must not use the same PR instance concurrently in different threads
 - this is due to the design of runtime parameters as Java Beans properties.
- Individual LR (documents, ontologies, etc.) are only thread-safe when accessed read-only by *all* threads.
 - if you need to share an LR between threads, be sure to synchronize (e.g. using `ReentrantReadWriteLock`)

Initializing GATE using a ServletContextListener

`ServletContextListener` called by container at startup and shutdown (only startup method shown).

```
1 public void contextInitialized(ServletContextEvent e) {
2     ServletContext ctx = e.getServletContext();
3     File gateHome = new File(
4         ctx.getRealPath("/WEB-INF"));
5     Gate.setGateHome(gateHome);
6     File userConfig = new File(
7         ctx.getRealPath("/WEB-INF/user.xml"));
8     Gate.setUserConfigFile(userConfig);
9     // default site config is gateHome/gate.xml
10    // default plugins dir is gateHome/plugins
11    Gate.init();
12 }
```

Initializing GATE using a ServletContextListener

You must register the listener in `web.xml`

```
1 <listener>
2   <listener-class>
3     gate.web.example.GateInitListener
4   </listener-class>
5 </listener>
```

Handling Concurrent Requests

Naïve approach – new PRs for every request

```
1 public void doPost(request, response) {
2     ProcessingResource pr = Factory.createResource(...);
3     try {
4         Document doc = Factory.newDocument(
5             getTextFromRequest(request));
6         try {
7             // do some stuff
8         }
9         finally {
10            Factory.deleteResource(doc);
11        }
12    }
13    finally {
14        Factory.deleteResource(pr);
15    }
16 }
```

Many levels of try/finally
– make sure you clean up
even when errors occur

Problems with Naïve Approach

- Guarantees no interference between threads
- But inefficient, particularly with complex PRs (large gazetteers, JAPE grammars, etc.)

Take Two: using ThreadLocal

Store the PR/Controller in a thread-local variable

```
1 private ThreadLocal<CorpusController> controller =
2     new ThreadLocal<CorpusController>() {
3
4     protected CorpusController initialValue() {
5         return loadController();
6     }
7 };
8
9 private CorpusController loadController() { ... }
10
11 public void doPost(request, response) {
12     CorpusController c = controller.get();
13     // do stuff with the controller
14 }
```

An Improvement...

- Only initialise resources once per thread
- Interacts nicely with typical web server thread pooling
- But if a thread dies (e.g. with an exception), no way to clean up its controller

One Solution: Object Pooling

- Manage your own pool of Controller instances
- Take a controller from the pool at the start of a request, return it (in a finally!) at the end
- Number of instances in the pool determines maximum concurrency level

Simple Example of Pooling

Setting up and cleaning up:

```
1 private BlockingQueue<CorpusController> pool;
2
3 public void init() {
4     pool = new LinkedBlockingQueue<CorpusController>();
5     for(int i = 0; i < POOL_SIZE; i++) {
6         pool.add(loadController());
7     }
8 }
9
10 public void destroy() {
11     for(CorpusController c : pool) {
12         Factory.deleteResource(c);
13     }
14 }
```

Simple Example of Pooling

Processing requests:

```
15 public void doPost(request, response) {
16     CorpusController c = pool.take();
17     try {
18         // do stuff
19     }
20     finally {
21         pool.add(c);
22     }
23 }
```

This blocks when the pool is empty. Use poll for non-blocking check.

Creating the pool

- Typically to create the pool you would use `PersistenceManager` to load a saved application several times.
- But this is not always optimal, e.g. large gazetteers consume lots of memory.
- GATE provides API to *duplicate* an existing instance of a resource: `Factory.duplicate(existingResource)`.
- By default, this simply calls `Factory.createResource` with the same class name, parameters, features and name.
- But individual Resource classes can override this by implementing the `CustomDuplication` interface (more later).
 - e.g. `DefaultGazetteer` uses a `SharedDefaultGazetteer` — same behaviour, but shares the in-memory representation of the lists.

Other Caveats

- With most PRs it is safe to create lots of identical instances
- But *not all*
 - e.g. training a machine learning model with the batch learning PR (in the *Learning* plugin)
 - but it is safe to have several instances *applying* an existing model.
- When using `Factory.duplicate`, be careful not to duplicate a PR that is being used by another thread
 - i.e. either create all your duplicates up-front or else keep the original prototype “pristine”.

Exporting the Grunt Work: Spring

- <http://www.springsource.org/>
- “Inversion of Control”
- Configure your business objects and connections between them using XML, Groovy or Java annotations.
- Handles application startup and shutdown
- GATE provides helpers to initialise GATE, load saved applications, etc.
- Built-in support for object pooling
- Web application framework (Spring MVC)
- Used by other frameworks (Grails, CXF, ...)

Initializing GATE via Spring XML

```
1 <beans
2   xmlns="http://www.springframework.org/schema/beans"
3   xmlns:gate="http://gate.ac.uk/ns/spring">
4   <gate:init gate-home="gate-files"
5       plugins-home="gate-files/plugins"
6       site-config-file="gate-files/gate.xml"
7       user-config-file="gate-files/user-gate.xml">
8     <gate:preload-plugins>
9       <value>gate-files/plugins/ANNIE</value>
10    </gate:preload-plugins>
11  </gate:init>
12 </beans>
```

- Paths can be full URLs (`file:/...`) or *resource* paths that are resolved appropriately by Spring

Loading a Saved Application

To load an application state saved from GATE Developer:

```
1 <gate:saved-application
2   id="myApp"
3   location="gate-files/application.xgapp"
4   scope="prototype" />
```

- `scope="prototype"` means create a new instance each time we ask for it
- Default scope is “singleton” — one instance is created at startup and shared.

Duplicating an Application

- Alternatively, load the application once and then duplicate it

```
1 <gate:duplicate id="myApp" return-template="true">
2   <gate:saved-application location="..." />
3 </gate:duplicate>
```

- `<gate:duplicate>` creates a new duplicate each time we ask for the bean.
- `return-template` means the original controller (from the saved-application) will be returned the first time, then duplicates thereafter.
- Without this the original is kept pristine and only used as a source for duplicates.

Worked example – a Spring Boot webapp

- *Spring Boot* is a framework to get a Spring-based application up and running with a few lines of code.
- “Convention over configuration” approach, providing sensible defaults that you can override.
- Builds using Maven 3 (or Gradle)
- Example is a Spring MVC web application
 - see hands-on materials for the code.

Setting up a Spring Boot app

- One “parent” and one “plugin” to add to Maven POM
- Add dependencies on the relevant “starter” modules, which themselves depend on the required libraries
 - `spring-boot-starter-web` for a basic Spring MVC application
 - your preferred view technology, in our case `spring-boot-starter-thymeleaf`
- And whatever other libraries you require, e.g. GATE

Setting up a Spring Boot app

- Entry point is a simple boilerplate class with annotations.
- Automatically scans sub-packages for other annotated classes (controllers, etc.).

```
1 package gatetutorial;
2 import org.springframework.boot.SpringApplication;
3 import org.springframework.boot.autoconfigure.
   SpringApplication;
4
5 @SpringBootApplication
6 public class TutorialApp {
7
8     public static void main(String... args) {
9         SpringApplication.run(TutorialApp.class, args);
10    }
11 }
```

A simple example controller

- Our example is a single controller that presents an HTML form to enter text.
- When form is submitted, process the document with a GATE application and show the document features as a table.

Spring pooling support

- Spring's built-in AOP features offer support for object pooling.
- Given a bean definition, we can expose a *proxy* object with the same behaviour backed by a pool of instances
- *Each method call* on the proxy is dispatched to one of the objects in the pool.
- Each target bean is guaranteed to be accessed by no more than one thread at a time.
- When the pool is empty, can configure further requests to block or fail.

The `gate:pooled-proxy` helper

- The machinery is all Spring but complex to configure.
- GATE provides a helper to automate this in Spring XML configuration.
- Don't pool GATE applications directly, instead pool a helper class that calls GATE.

```
1 <bean id="gateService"
2     class="gatetutorial.service.GateService">
3     <gate:pooled-proxy max-size="3"
4         initial-size="3" />
5 </bean>
```

More advanced pooling options

- Many more options to control the pool, e.g. for a pool that grows as required and shuts down instances that have been idle for too long, and where excess requests fail rather than blocking:

```
1 <gate:pooled-proxy
2     max-size="10"
3     max-idle="3"
4     time-between-eviction-runs-millis="180000"
5     min-evictable-idle-time-millis="90000"
6     when-exhausted-action-name="WHEN_EXHAUSTED_FAIL"
7 />
```

- Under the covers, `<gate:pooled-proxy>` creates a Spring `CommonsPoolTargetSource`, attributes correspond to properties of this class.
- See the Spring documentation for full details.

The GateService

- The GateService is written assuming single-threaded access.

```
1 public class GateService {
2     // will be injected automatically
3     @Autowired private CorpusController application;
4
5     private Corpus corpus;
6
7     @PostConstruct
8     public void init() throws GateException {
9         corpus = Factory.newCorpus("GateService");
10        application.setCorpus(corpus);
11    }
```

The GateService

```
13 @PreDestroy
14 public void destroy() {
15     Factory.deleteResource(corpus);
16     // not strictly necessary as gate:duplicate will handle this
17     Factory.deleteResource(application);
18 }
```

The GateService

```
20 public FeatureMap processWithGate(Document doc)
21     throws GateException {
22     try {
23         corpus.add(doc);
24         application.execute();
25         return doc.getFeatures();
26     } finally {
27         corpus.clear();
28     }
29 }
30 }
```

Tying it together

- Initialize GATE and configure service pool in XML (in src/main/resources)

```
1 <gate:init ... />
2 <gate:duplicate id="gateApplication"
3     return-template="true">
4     <gate:saved-application
5         location="gate-files/application.xgapp" />
6 </gate:duplicate>
7
8 <bean id="gateService"
9     class="gatetutorial.service.GateService">
10     <gate:pooled-proxy max-size="3"
11         initial-size="3" />
12 </bean>
```


Tying it together

- Add an annotation to TutorialApp to load the XML.

```
1 import org.springframework.context.annotation.  
   ImportResource;  
2 // ...  
3  
4 @SpringBootApplication  
5 @ImportResource("/gate-beans.xml")  
6 public class TutorialApp {
```

Tying it together

- And finally, autowire the GateService into controller, and call its methods without having to worry about threading.

```
1 // imports as required  
2  
3 @Controller  
4 public class GateController {  
5  
6     @Autowired  
7     private GateService gateService;
```

Tying it together

```
9 @RequestMapping(value="/",  
10                method = RequestMethod.POST)  
11 public String process(  
12     @ModelAttribute("params")  
13     AnnotationRequest params,  
14     Map<String, Object> model) throws GateException {  
15     Document doc = // extract text from request  
16     try {  
17         FeatureMap features =  
18             gateService.processWithGate(doc);  
19         model.put("features", features);  
20         return "index";  
21     } finally { Factory.deleteResource(doc); }  
22 }  
23 }
```

Exercise 1: The Spring Boot example

- In hands-on/webapps you have the source code for the Spring Boot example we've been discussing.
- You will need Maven 3 installed to run it.
- What's provided:
 - the pom.xml with the necessary dependencies,
 - source code for the controller and GateService (in src/main/java),
 - the Thymeleaf view with the text entry form and results table,
 - configuration in src/main/resources, including GATE config files and the bean definition XML.

Exercise 1: The Spring Boot example

- What's missing:
 - the GATE application itself. . .
- Use the document statistics PR from earlier.
- In GATE Developer, create a “corpus pipeline” application containing a tokeniser and your statistics PR.
- Right-click on the application and “Export for GATECloud.net”.
 - This will save the application state along with all the plugins it depends on in a single zip file.
- Unpack the zip file under
`src/main/resources/gate-files`
 - don't create any extra directories – you need `application.xgapp` to end up in `gate-files`.

Exercise 1: The Spring Boot example

- Now you can run the application – in hands-on/webapps run `mvn spring-boot:run`
- Browse to `http://localhost:8080/`, enter some text and submit
- Watch the log messages. . .
- Notice the result page includes a feature “handledBy” – each service instance in the pool has a unique ID.
- Multiple submissions go to different instances in the pool.
- Try editing `src/main/resources/gate-beans.xml` and change the pooling configuration.
- Test concurrent requests – the service has a built-in delay to simulate a slow application.

Not Just for Webapps

- Spring Boot (and Spring in general) isn't just for web applications
- You can use the same tricks in other embedded apps
- GATE provides a `DocumentProcessor` interface suitable for use with Spring pooling, which exposes one `void` method `processDocument`

```
1 <bean id="processor"
2   class="gate.util.LanguageAnalyserDocumentProcessor">
3   <property name="analyser" ref="gateApplication"/>
4   <gate:pooled-proxy max-size="3" initial-size="3" />
5 </bean>
```

A simple command-line app

```
1 @Component
2 public class GateCommand implements CommandLineRunner {
3     @Autowired private DocumentProcessor proc;
4
5     public void run(String... args) throws Exception {
6         Document doc = Factory.newDocument(args[0]);
7         try {
8             proc.processDocument(doc);
9             // ...
10        } finally { Factory.deleteResource(doc); }
11    }
12 }
```

The main entry point `TutorialApp` is unchanged from the web example.

A JMS message consumer

```
1 @Component
2 public class Receiver {
3     @Autowired private DocumentProcessor proc;
4
5     @JmsListener(destination = "someQueue",
6                 concurrency = "3")
7     public void receive(String stringMessage) {
8         Document doc = Factory.newDocument(stringMessage);
9         try {
10             proc.processDocument(doc);
11             doStuffWithResults(doc);
12         } finally { Factory.deleteResource(doc); }
13     }
14 }
```

In this case we need to add `@EnableJms` to the entry point class, and relevant dependencies to the POM.

Conclusions

Two golden rules:

- Only use a GATE Resource in one thread at a time
- Always clean up after yourself, even if things go wrong (deleteResource in a finally block).

Duplication and Custom PRs

- Recap: by default, `Factory.duplicate` calls `createResource` passing the same type, parameters, features and name
- This can be sub-optimal for resources that rely on large read-only data structures that could be shared
- If this applies to your custom PR you can take steps to make it handle duplication more intelligently
- For simple cases: *sharable properties*, for complex cases: *custom duplication*.

Sharable properties

- A way to share object references between a PR and its duplicates
- A JavaBean setter/getter pair with the setter annotated (same as for `@CreoleParameter`)

```
1 private Map dataTable;
2
3 public Map getDataTable() { return dataTable; }
4
5 @Sharable
6 public void setDataTable(Map m) {
7     dataTable = m;
8 }
```

Sharable properties

- Default duplication algorithm will get property value from original and set it on the duplicate before calling `init()`
- `init()` must detect when sharable properties have been set and react appropriately.

```
1 public Resource init() throws /* ... */ {
2     if(dataTable == null) {
3         // only need to build the data table if we weren't given a shared one
4         buildDataTable();
5     }
6 }
7
8 public void reInit() throws /* ... */ {
9     // clear sharables on reinit
10    dataTable = null;
11    super.reInit();
12 }
```

Sharable properties – Caveats

- Anything shared between PRs *must* be thread-safe
 - use appropriate synchronization if any of the threads modifies the shared object (e.g. a `ReentrantReadWriteLock` which is itself `@Sharable`).
 - or (for the `dataTable` example), use an inherently safe class such as `ConcurrentHashMap`
 - for shared counter, use `AtomicInteger`
- If you use sharable properties, take care not to break `reInit`

Exercise 2: Multi-threaded cumulative statistics

- `hands-on/shared-stats` contains a variation on the `DocStats` PR that keeps a running total of the number of Tokens it has seen.
- Build this (using the Ant build file), load the plugin, create an application containing a tokeniser and a “Shared document statistics” PR, export for GATECloud.net and unzip into your webapp as before.
- Try posting some requests to the webapp.
- You will see a `running_total` feature, but this is per handler, not global across handlers.

Exercise 2: Multi-threaded cumulative statistics

- Your task: make the running total global.
- Make the `totalCount` field into a sharable property
 - it's already a thread-safe `AtomicInteger`
 - add a getter and setter, with the right annotation
 - `init()` logic to handle the shared/non-shared cases
 - implement a sensible `reInit()`
- You will need to re-build your PR and re-export (or just copy the compiled plugin to the right place in your webapp).

Custom Duplication

- For more complex cases, a resource can take complete control of its own duplication by implementing `CustomDuplication`
- This tells `Factory.duplicate` to call the resource's own `duplicate` method instead of the default algorithm.

```
1 public Resource duplicate(DuplicationContext ctx)
   throws ResourceInstantiationException;
```

- `duplicate` should create and return a duplicate, which need not be the same concrete class but must “behave the same”
 - Defined in terms of implemented interfaces.
 - Exact specification can be found in the `Factory.duplicate` JavaDoc.

Custom Duplication

- If you need to duplicate other resources, use the two-argument `Factory.duplicate`, passing the `ctx` as the second parameter, to preserve object graph
 - two calls to `Factory.duplicate(r, ctx)` for the same resource `r` in the same context `ctx` will return the same duplicate.
 - calls to the single argument `Factory.duplicate(r)` or to the two-argument version with different contexts will return different duplicates.
- Can call the default duplicate algorithm (bypassing the `CustomDuplication` check) via `Factory.defaultDuplicate`
 - it is safe to call `defaultDuplicate(this, ctx)`, but calling `duplicate(this, ctx)` from within its own custom `duplicate` will cause infinite recursion!

Custom Duplication Example (SerialController)

```
1 public Resource duplicate(DuplicationContext ctx)
2     throws ResourceInstantiationException {
3     // duplicate this controller in the default way - this handles subclasses nicely
4     Controller c = (Controller)Factory.defaultDuplicate(
5         this, ctx);
6
7     // duplicate each of our PRs
8     List<ProcessingResource> newPRs =
9         new ArrayList<ProcessingResource>();
10    for(ProcessingResource pr : prList) {
11        newPRs.add((ProcessingResource)Factory.duplicate(
12            pr, ctx));
13    }
14    // and set this duplicated list as the PRs of the copy
15    c.setPRs(newPRs);
16
17    return c;
18 }
```

Outline

- 1 GATE in Multi-threaded/Web Applications
 - Multi-threading and GATE
 - Servlet Example
 - The Spring Framework
 - Making your own PRs duplication-friendly
- 2 Extending GATE
 - Adding new document formats

Adding new document formats

- GATE provides default support for reading many source document formats, including plain text, HTML, XML, PDF, DOC, ...
- The mechanism is extensible – the format parsers are themselves resources, which can be provided via CREOLE plugins.
- GATE chooses the format to use for a document based on *MIME type*, deduced from
 - explicit `mimeType` parameter
 - file extension (for documents loaded from a URL)
 - web server supplied Content-Type (for documents loaded from an `http: URL`)
 - “magic numbers”, i.e. signature content at or near the beginning of the document

The DocumentFormat resource type

- A GATE document format parser is a resource that extends the `DocumentFormat` abstract class or one of its subclasses.
- Override `unpackMarkup` method to do the actual format parsing, creating annotations in the `Original` markups annotation set and optionally modifying the document content.
- Override `init` to register with the format detection mechanism.
- In theory, can take parameters like any other resource ...
- ... but in practice most formats are singletons, created as *autoinstances* when their defining plugin is loaded.

Repositioning info

- Some formats are able to record *repositioning info*
- Associates the offsets in the extracted text with their corresponding offsets in the original content.
- Allows you to save annotations as markup inserted into the original content.
- Of the default formats, only HTML can do this reliably.
 - If you're interested, see the `NekoHtmlDocumentFormat`

Implementing a DocumentFormat

- Define a class that extends `DocumentFormat`, with CREOLE metadata

```
1 import gate.*;
2 import gate.creole.metadata.*;
3 import gate.corpora.*;
4
5 @CreoleResource(name = "Example DocumentFormat",
6     autoinstances = {@AutoInstance})
7 public class MyDocumentFormat
8     extends TextualDocumentFormat {
9     // ...
10 }
```

- `autoinstances` causes GATE to create an instance of this resource automatically when the plugin is loaded.

DocumentFormat methods

- Most formats need to override three or four methods.
- `supportsRepositioning` to specify whether or not the format is capable of collecting repositioning info – most aren't

```
1 public Boolean supportsRepositioning() {  
2     return false;  
3 }
```

DocumentFormat methods

- Two variants of `unpackMarkup`
- If you don't support repositioning then best to extend `TextualDocumentFormat` and just override the simple one:

```
1 public void unpackMarkup(Document doc)  
2     throws DocumentFormatException {  
3     AnnotationSet om = doc.getAnnotations(  
4         GateConstants.ORIGINAL_MARKUPS_ANNOT_SET_NAME);  
5     // Make changes to the document content, add annotations to om  
6 }
```

- Other variant (for repositioning formats) is implemented in terms of this one by `TextualDocumentFormat`

DocumentFormat methods

- Finally, `init` to register the format with GATE
- Mostly boilerplate, using protected `Map` fields defined in `DocumentFormat`

```
1 public Resource init() throws  
2     ResourceInstantiationException {  
3     MimeType mime = new MimeType("text", "x-special");  
4     mimeType2ClassHandlerMap.put(  
5         mime.getType() + "/" + mime.getSubtype(), this);  
6     mimeType2mimeTypeMap.put(  
7         mime.getType() + "/" + mime.getSubtype(), mime);  
8     suffixes2mimeTypeMap.put("spec", mime);  
9     magic2mimeTypeMap.put("==special==", mime);  
10  
11     setMimeType(mime);  
12     return this;  
13 }
```

Registering a document format

```
2 MimeType mime = new MimeType("text", "x-special");  
3 mimeType2ClassHandlerMap.put(  
4     mime.getType() + "/" + mime.getSubtype(), this);
```

- Create a `MimeType` object representing the “primary” MIME type for this format.
- Register this object as the handler for this MIME type.

```
5 mimeType2mimeTypeMap.put(  
6     mime.getType() + "/" + mime.getSubtype(), mime);
```

- Establish a mapping between the MIME string “text/x-special” and the primary `MimeType` object.
- To register a format against several different MIME types (e.g. `text/json` and `application/json`), add them to the `mimeType2mimeTypeMap`

Registering a document format

```
7 suffixes2mimeTypeMap.put("spec", mime);
```

- Register the file suffixes (not including the leading dot) that the format will handle, by mapping them to the primary `MimeType`
- Can add several different suffixes for the same type (txt, text, etc.)

```
8 magic2mimeTypeMap.put("==special==", mime);
```

- Add “magic numbers” – strings whose presence within the first 2kB of content will select the format
- E.g. “<?xml” is a strong predictor of XML documents.

Registering a document format

```
10 setMimeType(mime);  
11 return this;
```

- Boilerplate.
- Suffixes and magic numbers are optional – don’t use them if they don’t make sense for your particular format.
- ... but if neither are specified then only documents created with an explicit `mimeType` parameter will use the format.

Exercise: Document format registration

- `hands-on/yam-format` contains a simple document format implementation.
- Processes text files in the “YAM” format (the Wiki markup syntax used on <http://gate.ac.uk>).
- `unpackMarkup` has been written for you.
- Annotates `*bold*`, `_italic_` and `^teletype^ text`, and section headings (lines starting %1, %2, etc.).
- For simplicity, does not modify the text or do repositioning, only adds Original markups annotations.

Exercise: Document format registration

- Your task – write the `init` method registration code
 - Primary MIME type “text/x-yam”
 - File suffixes “.yam” and “.gate”
 - No magic numbers
- To test, `ant jar` to build the JAR file, then load the `yam-format` directory as a plugin in GATE Developer.
 - Note the auto-instance created when the plugin loads
- Create a document from the `overview.yam` file and inspect the Original markups.