Natural Language Engineering **10** (3/4): 349–373. © 2004 Cambridge University Press
DOI: 10.1017/S1351324904003468 Printed in the United Kingdom

Evolving GATE to meet new challenges in language engineering†

KALINA BONTCHEVA, VALENTIN TABLAN, DIANA MAYNARD and HAMISH CUNNINGHAM
Department of Computer Science, University of Sheffield, Sheffield S1 4DP, UK
e-mail: {kalina,valyt,diana,hamish}@dcs.shef.ac.uk

(*Received 1 July 2003; revised 9 February 2004*)

Abstract

In this paper we present recent work on GATE, a widely-used framework and graphical development environment for creating and deploying Language Engineering components and resources in a robust fashion. The GATE architecture has facilitated the development of a number of successful applications for various language processing tasks (such as Information Extraction, dialogue and summarisation), the building and annotation of corpora and the quantitative evaluations of LE applications. The focus of this paper is on recent developments in response to new challenges in Language Engineering: Semantic Web, integration with Information Retrieval and data mining, and the need for machine learning support.

1 Introduction

Natural Language Engineering (LE) is a rapidly changing field, driven by the push of internal technological advances on the one hand, and by the pull of new application contexts on the other. These drivers include:

- the project of the Semantic Web, adding a machine-tractable semantic layer to the existing World Wide Web (Fensel, Wahlster and Lieberman 2002; Davies, Fensel and van Harmelen 2002);
- the growing need for easily portable multilingual LE applications in face of increasing integration in an expanding Europe;[1]
- the continuing growth in methods combining information theory and machine learning with symbolic and linguistic systems;

† Work on GATE has been supported by the Engineering and Physical Sciences Research Council (EPSRC) under grants GR/K25267 (Large-Scale Information Extraction), GR/M31699 (GATE 2), GR/N15764/01 (AKT), EMILLE, and by several smaller grants. We would like to thank the numerous people who have contributed to this work. Particularly Marin Dimitrov, Cristian Ursu, Oana Hamza, Borislav Popov, Atanas Kiryakov, Robert Gaizauskas, and Yorick Wilks for their contributions to GATE v2; and finally to Donia Scott and the three anonymous reviewers of this paper.

[1] And perhaps in view of decreasing integration and increasing strife internationally.

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

- the emergence of applications combining Information Extraction (IE) and Information Retrieval (IR) methods in very large scale contexts.[2]

These changes all create new challenges to Language Engineering architectures and infrastructures, a field which we may call Software Architecture for Language Engineering, or SALE. SALE systems aim to provide engineering support for research and development of language processing software, in a way similar to the support offered to programmers by software development environments like Visual C++.

GATE (Cunningham, Humphreys, Gaizauskas and Wilks 1997; Cunningham 2002; Cunningham, Maynard, Bontcheva and Tablan 2002), a General Architecture for Text Engineering, was first released in 1996, then completely re-designed, re-written, and re-released in early 2002. The system is now widely-used (by thousands of people at hundreds of sites) and is a relatively comprehensive infrastructure for language processing software development. Key features of GATE are:

- Component-based development reduces the systems integration overhead in collaborative research.
- Open source, documented and supported software enhances the repeatability of experiments.
- Automatic performance measurement of the LE components promotes quantitative comparative evaluation.
- Distinction between low-level tasks such as data storage, data visualisation, discovery and loading of components and the high-level language processing tasks.
- Clean separation between between data structures and algorithms that process human language.
- Consistent use of standard mechanisms for components to communicate data about language, and use of open standards such as Unicode and XML.
- Insulation from indiosyncratic data formats – GATE performs automatic format conversion and enables uniform access to linguistic data.
- Provision of a baseline set of LE components that can be extended and/or replaced by users as required.

The focus of this paper is on the recent novel aspects of the infrastructure, that have been introduced in GATE version 2 (GATE v2) – the majority in response to the recent challenges discussed above. Some of these new elements were added to GATE because of problems identified with its first version (Cunningham 2002):

- Non-extensibility of the tool model (extensibility was only available for processing modules, not for viewers, or other tools).
- Problematic multilingual support.
- Need for a pattern-matching engine to allow rapid prototyping and to simplify the development of LE components.

[2] See, for example, IBM's WebFountain product.

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

- Need for better support for lexical resources, e.g. lexicons, thesauri, and their use within the processing components.
- Need to extend towards new data formats, e.g. XML, DAML+OIL, that did not exist at the time when GATE version 1 (GATE v1) was implemented.
- Inappropriate processing model for language generation systems.
- Insufficient portability with respect to operating systems.

The rest of the paper is structured as follows. Section 2 introduces the GATE framework with a focus on important new elements of the architecture, such as support for distributed resources, application execution strategies, and the JAPE finite-state processing engine. Section 3 introduces the reusable LE components in GATE v2, including machine learning and ontology-based modules. Section 4 discusses the new types of language resources: ontologies and lexicons, their representation, editing, and use. Section 5 presents the new extendable model for reuse and customisation of viewers, editors, and other interface tools. A number of recent LE applications is then discussed in section 6, followed by a review of related work (section 7). Finally, we conclude by revisiting some recent challenges for language engineering and outline ongoing and future work on GATE in these directions.

2 GATE: A framework for robust tools and applications

GATE[3] is an architecture, a framework and a development environment for LE (Language Engineering) applications. As an *architecture*, it defines the organisation of an LE system and the assignment of responsibilities to different components. As a *framework*, it provides reusable implementations for LE components and a set of prefabricated software building blocks that language engineers can use, extend and customise for their specific needs. As a *development environment*, it helps its users minimise the time they spend building new LE systems or modifying existing ones, by aiding overall development and providing a debugging mechanism for new modules. Because GATE has a component-based model, this allows for easy coupling and decoupling of the processors, thereby facilitating comparison of alternative configurations of the system or different implementations of the same module (e.g. different parsers). The availability of tools for easy visualisation of data at each point during the development process aids immediate interpretation of the results.

Applications developed within GATE can be deployed outside its Graphical User Interface (GUI), using programmatic access via the GATE Application Programming Interface (API).[4] In addition, the reusable modules, the document and annotation model, and the visualisation components can all be used independently of the development environment.

[3] GATE v2 is freely available for download from `http://gate.ac.uk`. It is implemented in Java and runs on a wide range of platforms.

[4] `http://www.gate.ac.uk/releases/gate-2.2-build1350--ALL/doc/javadoc/index.html`

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

GATE is engineered to a high standard and supports efficient and robust text processing. It is tested extensively, including regression testing, and frequent performance optimization. GATE v2 has proved capable of processing gigabytes of text and thousands of documents (most recently tested on the British National Corpus (Burnard 1995)). The system is supported by extensive documentation, online tutorials and an established mailing list.

In the rest of this section, we provide an overview of the GATE architecture, discuss the new annotation graph based data representation, and introduce important new elements, such as multilingual support, distributed resources, execution strategies, and the JAPE finite-state processing engine.

2.1 Overview of the architecture

The GATE architecture distinguishes between both data and algorithms, and their visualisation.[5] Following the terminology established in version 1, GATE components are one of three types:

- *Language Resources* (LRs) represent entities such as lexicons (e.g. WordNet), corpora or ontologies;[6]
- *Processing Resources* (PRs) represent entities that are primarily algorithmic, such as parsers, generators or n-gram modellers;
- *Visual Resources* (VRs) – added in version 2 – represent visualisation and editing components that participate in GUIs.

These resources can be local to the user's machine or remote (available over the (Inter)net), and all can be extended by users without modification to GATE itself.

One of the main advantages of separating the algorithms from the data they require is that the two can be developed independently by language engineers with different types of expertise, e.g. programming and linguistics. Similarly, separating data from its visualisation allows users to develop alternative visual resources, while still using a language resource provided by GATE.

Collectively, all resources are known as CREOLE (a Collection of REusable Objects for Language Engineering), and are described in XML configuration files, which declare their name, implementing class, parameters, icons, etc. This component metadata is used by the framework to discover and load available resources.

Unlike version 1, GATE v2 offers comprehensive multilingual support using Unicode as its default text encoding. It also provides a means of entering text in various languages, using virtual keyboards where the language is not supported by the underlying operating platform (McEnery, Baker, Gaizauskas and Cunningham 2000; Tablan, Ursu, Bontcheva, Maynard, Hamza and Wilks 2002).

[5] Analoguous to the model/controller/view architecture common in GUI toolkits.

[6] Ontologies are often not considered as describing language and, therefore, their classification as LRs could be somewhat un-intuitive. Probably a more appropriate term will be data resources, but we adhere to the original terminology adopted in GATE v1. Fundamentally, GATE treats all declarative data resources like lexicons, ontologies, and documents in the same fashion, i.e. they can be provided as inputs or parameters to PRs and are visualised and edited in VRs.

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

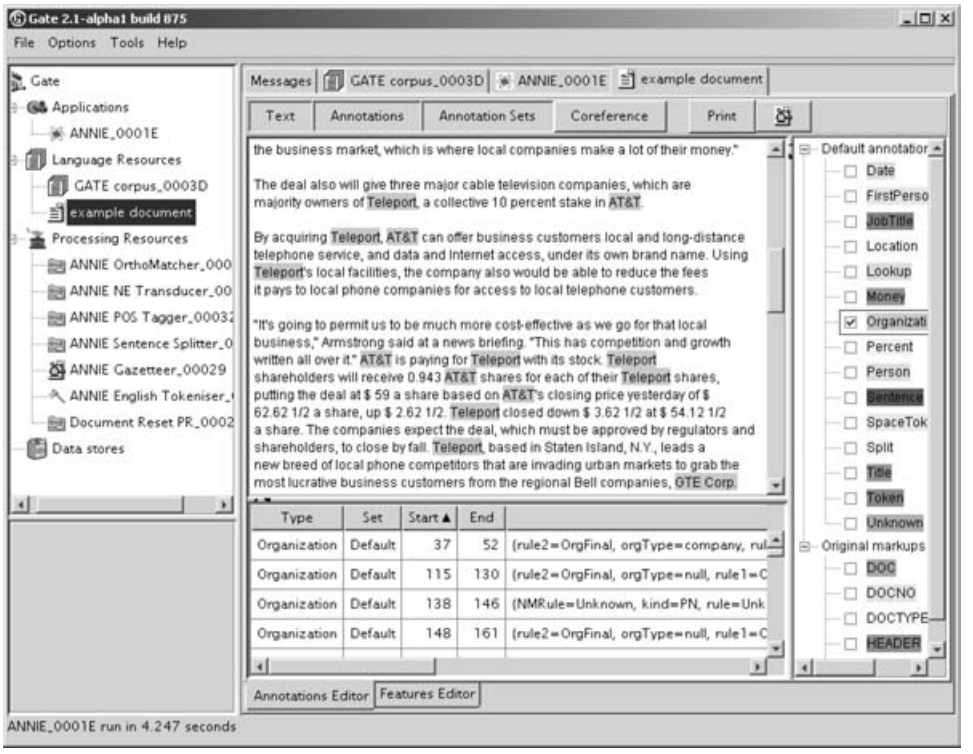


Fig. 1. GATE's document viewer/editor.

When an application is developed within GATE's graphical environment, the user chooses which processing resources go into it (e.g. tokeniser, POS tagger), in what order they will be executed, and on which data (e.g. document or corpus). The application results can be viewed in the document viewer/editor (see Figure 1). Applications can be saved, reloaded, and embedded in other systems.

2.2 Data representation and handling

GATE supports a variety of formats including plain text, HTML, XML, RTF, and SGML to enable the processing of a wide range of documents and improve interoperability with other LE infratructures. In all cases, when a document is opened in GATE, 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 1997) (originally used in GATE v1), which has been made largely compatible with the Atlas format (Bird, Day, Garofolo, Henderson, Laprun and Liberman 2000), and is isomorphic with "stand-off markup" as widely adopted by the SGML/XML community (Ide, Bonhomme and Romary 2000). The annotations associated with each document are a structure central to GATE, because they encode the data read and produced by all processing modules, as well as the formatting information originally present in the documents.

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

Annotations are organised in *annotation graphs*, where the vertices are anchored in the document content. Annotations are the arcs in the graph; they have a start node and an end node, an identifier, a type and a set of features. Nodes have pointers into the content, e.g. character offsets for text, milliseconds for audio-visual content. There can be more than one annotation graph associated with each document. The support for separate annotation graphs (called families in IBM's TEXTRACT architecture (Neff, Byrd and Boguraev 2004)) is needed to allow separation of information into different categories.

GATE has a single model for encoding information that describes documents, collections of documents (corpora), and annotations on documents, based on attribute name/attribute value pairs, called *features*. Attribute names are strings; values can be any object. Figure 1 shows annotations of type `Organization` and their features: `orgType` for the type of organisation, `rule` specifying which rule produced this annotation, etc. At present GATE uses features only as a representation formalism and does not provide operations like unification.

In order to facilitate the interpretation and validation of features associated with each annotation type, GATE v2 provides annotation schemas, which are encoded in the XML Schema language supported by W3C, and define all attributes and the type of their value (e.g. `orgType` for `Organization` annotations). If the value needs to be one of a predefined set of values (e.g. a part-of-speech tagset), then this set of values is also given in the annotation schema. GATE's annotation schemas were inspired by TIPSTER's annotation type declarations (Grishman 1997), however GATE also uses them to facilitate and validate users' input during manual annotation (see section 5.1).

2.3 JAPE: Finite state processing over annotations

GATE v2 has new facilities for finite state processing over annotations based on regular expressions. JAPE – a Java Annotation Patterns Engine – is similar to the finite state pattern matching transducer part of the TEXTRACT architecture (Neff *et al.* 2004). 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. This cascade is similar to the finite state cascade used for parsing (Abney 1996) in that rules in the current phase can refer to annotations created in earlier phases and there is no recursion.

The left-hand-side (LHS) of the rules consist of an annotation pattern (or a macro describing such a pattern), that may contain Kleene regular expression operators (e.g. `*`, `?`, `+`). The right-hand-side (RHS) consists of annotation manipulation statements (in JAPE or Java). 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.

At the beginning of each grammar, several options can be set:

- Control – this defines the method of rule matching: `brill`, `appelt`, or `first`. The `brill` style means that when more than one rule matches the same region of the document, they are all fired. With the `appelt` style, only one rule can be fired for the same region of text (the choice depends on the length

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

of the matching region and rule priority). With the `first` style, a rule fires for the first match that is found and no other matching attempts are made for that region of the document.

- Input annotation types – restrict the input to the grammars in order to improve execution speed and control the range for matching in the document.

The example below shows a JAPE rule with pattern and corresponding action. (On the LHS, each such 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"},
```

A large proportion of the processing resources bundled with GATE use JAPE.

2.4 Execution strategies

The execution strategies in GATE determine how the algorithms (i.e. processing resources) are combined to form a complete application and how that application is executed on the provided input (i.e. language resources such as corpora). These execution strategies are implemented in GATE v2 as controllers and the user chooses the one most suitable for their application.

The two main execution strategies commonly found in LE applications are *pipeline* (i.e. sequential) and *parallel*. For example, GATE's pipeline controller executes the PRs in the specified order and with the given parameters (e.g. a document). A variation of this controller is the corpus pipeline, which executes the PRs sequentially over a given corpus, taking care of loading the documents from the disk (if necessary), processing each of them, and saving the results.

At present GATE does not provide a controller which supports parallel execution of PRs, although it is planned for future work. However, it does support execution of applications on different machines over data shared on the server (see Section 2.5), which enables simultaneous processing of documents.

Other types of execution strategies are *conditional* (equivalent to an `if` statement in programming languages) and *iterative* (equivalent to loop statements). The conditional controller (see Figure 2) was introduced in GATE v2 in order to allow flexible processing of heterogenous data (e.g. a corpus consisting of documents from several domains, styles, and genres). The controller allows the user to specify for each

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

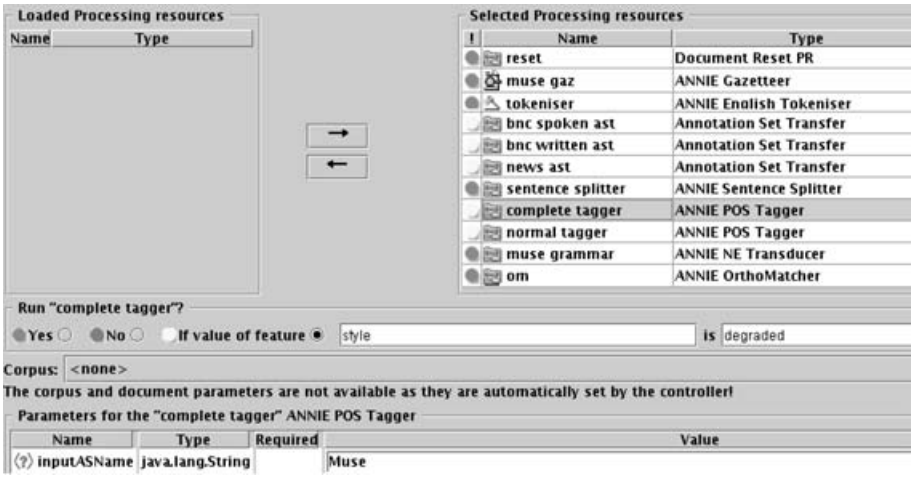


Fig. 2. GATE's conditional controller.

PR the conditions when it can be executed: `always`; `never`; or `conditionally` – only if a particular feature with a given value is present on the document.

For example, a categorisation module that always fires can assign different genres to each document as features (e.g. sport, political, business). Then the application can contain PRs trained specifically for the given type of documents and at execution time, based on the document type, the controller will determine which PRs are executed. For example, a text about sport would use a different named entity recogniser from the one about politics (England is a team in the first case and a location in the second) or a text in all uppercase would use a different POS tagger (amongst other things) from a text in mixed case (see Figure 2). However, the application will list all alternative PRs (e.g. two POS taggers and three entity recognisers), but only those appropriate for the current document will be executed on it, whereas a different set of PRs might be executed on the next one.

2.5 Distributed language resources

With the availability of wider bandwidth and the increasingly distributed nature of research comes a new need for distributed language resources (LRs) which are best stored on server machines with large amounts of disk space, so they are installed once and shared among many researchers. An important step towards the creation of such distributed support is the provision of an efficient and user-friendly client-server architecture. This problem can be decomposed into two major tasks: (i) providing support for distributed resources stored on a server; and (ii) providing client tools that are easy to use.

GATE v2 was extended to offer support for server-based language resources by storing them in relational databases (similar to the way in which ATLAS corpora and annotations are stored and shared (Ma, Lee, Bird and Maeda 2002)). The shared LRs are accessed by the GATE visual environment that acts as a client. Currently Oracle and PostgreSQL are supported, of which the latter is freely available.

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

Uniform models for storing and accessing resources in files and databases is an actively researched area in Information Retrieval (e.g. Cutting, Pedersen and Halvorsen 1991) for storing indexes and document collections, and for more complex annotation models as required by language processing applications, e.g. the TIPSTER architecture (Grishman 1997). The difference with GATE's and ATLAS-based uniform models comes from the annotation graphs which are more generic and support the association of more than one graph per document. However, this generality comes at the cost of more complex database schemas, which leads to slower performance due to the higher number of join and select operations required.

In order to avoid performance bottlenecks, GATE assumes that each resource is stored entirely in one place, e.g. file or a given database. It is not possible for parts of a resource (e.g. document attributes) to reside on different servers or storage media. Dynamically evaluated attributes and alignment of corpora are also not supported, which makes the GATE LR model a sub-set of more powerful models like the CQP (Christ 1994). However, unlike in CQP, recursive attribute structures are supported (e.g. syntax trees).

GATE users are isolated from the ways in which LRs are stored. When choosing to store an LR in a distributed database, users just have to choose which of the available servers they want to use and provide their user name and password. Then GATE transparently converts all LRs in the corresponding database objects and stores them via JDBC. In this way, all GATE applications are separated from the technicalities of using a database server for distributed LR storage – distributed LRs are accessed in the same way as LRs stored on the local file system.

An important aspect of supporting distributed LRs is security and access control. GATE has role-based access control and every language resource is associated with security properties specifying the actions that certain users and groups may perform with this resource. When users create distributed LRs, they specify access rights for them, their user group and other users. For example, LRs created by one user/group can be made read-only to others, so they can use the data, but not modify it. The access rights available in GATE are similar to those of a Unix file system.

3 New GATE processing resources

Provided as part of the GATE infrastructure is a set of reusable processing resources for common LE tasks. (None of them are definitive, and the user can replace and/or extend them as necessary.) In version 2, these are packaged together to form ANNIE (A Nearly-New IE system), which is similar to the Vanilla Information Extraction (VIE) system in GATE v1. The PRs can also be used individually or coupled together with new modules in order to create new applications. For example, many other NLP tasks might require a sentence splitter and POS tagger, but would not necessarily require resources more specific to IE tasks such as a named entity recogniser.

ANNIE consists of the following main processing resources: tokeniser, sentence splitter, POS tagger, gazetteer, name entity tagger (based on JAPE), orthomatcher and coreference resolver (Cunningham *et al.* 2002). The resources communicate via GATE's annotation model discussed in Section 2.2. Since these PRs have analogues

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

in the VIE system, supplied with GATE v1 (Cunningham 2002), they will not be discussed here due to space liminations.

The Information Retrieval (IR) module is a new optional PR that allows the combination of IE and IR in a single application. First, the IR module can be used to find relevant documents in a given corpus and then IE modules can be run to extract relevant information, e.g. in a question-answering application. Also, Lucene[7], the underlying retrieval engine, allows the inclusion of arbitrary features in the document index, facilitating experiments in NLP-based IR, for example.

The architectural implications of integrating IR were that it lead to the definition of a generic IR API in GATE which shares commonalities with the TIPSTER IR model (Grishman 1997). It consists of three types of resources: the IR index itself can be viewed as a GATE LR, the indexing and retrieval is a PR, and the user interface showing the ranked IR results is a VR. These distinctions are similar to those made in the object-oriented IR architecture discussed in Cutting *et al.* (1991). The difference is mainly in GATE's corpus abstraction model which is more generic and ehnahced to meet the requirements of language processing modules, in addition to IR. The integration of Lucene follows this three-partite GATE IR model and is done via wrapper code which translates between the two APIs.

3.1 Integration of machine learning

In order to be able to use Machine Learning (ML) in GATE, a new model was designed for interfacing to ML-based classifiers. The aim is by providing support for a wide range of ML algorithms to facilitate comparative experiments on the same dataset.

From an architectural perspective, the ML interfaces were designed to be generic and to make it easier to integrate different classifiers from different toolkits. In addition, the design provides linguistic information as input to the machine learning algorithms directly from GATE's model of annotations. In practice, this means that a generic mechanism for collecting instances from GATE's data model was implemented. It is based on a declarative XML-based specification of mappings between GATE's annotation types and their attributes and the input features for the classifier. Once collected, the data is exported in the format required by the ML algorithm, which is often a table where each row is an instance and each column is a feature.

When collecting training data, all the annotations of the type specified as instances are found in the given corpus (using GATE's annotation API), and for each of them the set of attribute values is determined. All attribute values, provided as features to the learning algorithm, refer either to the current annotation or to one situated at a specified relative position (e.g. +1 is the next annotation).

Boolean attributes refer to the presence (or absence) of a particular type of annotation overlapping at least partially with the given instances annotation. *Nominal* and *numeric* attributes refer to features on a particular type of annotation that

[7] `http://jakarta.apache.org/lucene/docs/index.html`

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

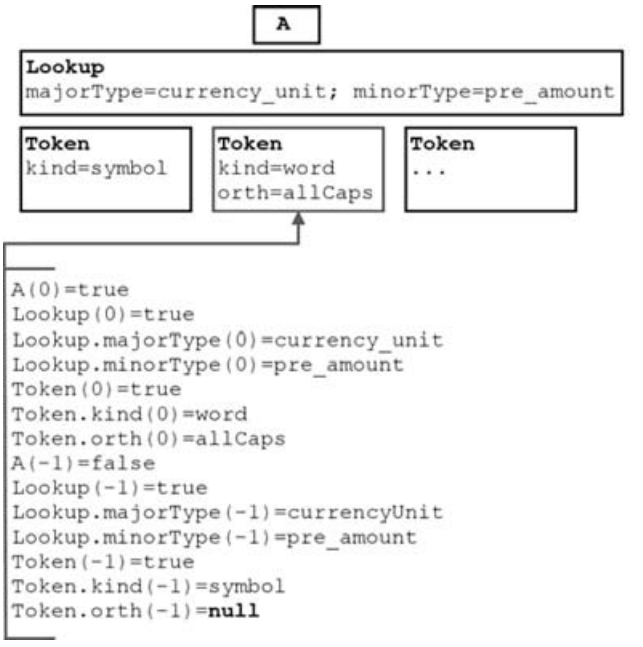


Fig. 3. Example of attributes and their values.

(partially) overlaps the given annotation. One of the boolean or nominal attributes is marked as the *class attribute*, and the values which that attribute can take are the labels for the classes to be learned by the algorithm.

Since linguistic information is often used as attributes for ML in LE applications, GATE provides support for supplying a wide range of linguistic information – part-of-speech, sentence boundaries, gazetteer lists, and named entity class – as attributes for the machine learning algorithms. This information, together with tokenisation information (kind, orthography, and token length) is obtained by using the ANNIE PRs supplied with GATE (Cunningham *et al.* 2002). The user chooses which of this information should be provided as features to the learning algorithm.

For example, Figure 3 shows that the attributes supplied to the learning algorithms are derived from the gazetter annotation overlapping the current instance (0), the token annotation overlapping the current instance (0), the token annotation of the previous instance (−1), etc. In this case, instances for learning algorithm are defined to be every occurrence of the Token annotation and what is being learned is the presence or absence of the boolean attribute A.

An ML implementation has two modes of functioning: training – when the model is being built, and application – when the built model is used to classify new instances. Our implementation consists of a GATE processing resource that handles both the training and application phases. It is responsible for detecting all the instances in a corpus and collecting the attribute values for them. The output writer is format specific, so new ones will need to be implemented if new ML libraries are integrated, which use format different from Weka's.

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

Depending on the type of the attribute, different actions will be performed when classification occurs. For boolean attributes, a new annotation of the type specified in the attribute definition will be created (e.g. annotation of type A). Nominal attributes trigger the addition of the feature on the annotation of the required type (Token in the example above), situated at the position of the classified instance. If no such annotation is present, it will be created.

Currently GATE supports the ML algorithms implemented in the open-source WEKA library (Witten and Frank 1999)[8]; other implementations are being added in the SEKT project (see below). WEKA was chosen because it also provides tools for performance evaluation, testing, and attribute selection. In addition, most of the ML algorithms in WEKA can provide a probability distribution rather than a simple classification. This can be used by GATE applications for setting a confidence threshold, thus enabling a choice of balance between precision and recall.

3.2 Ontology-based PRs

The addition of support for ontologies in GATE (section 4.2) enabled the creation of ontology-based PRs, i.e. PRs that use the ontology as one of their data resources. One such new component is the ontology-based gazetteer, which treats phrases as instances of concepts from a given ontology. Based on this information, subsequent PRs can access the same ontology and perform reasoning with the information produced by the ontology-based gazetteer, e.g. to obtain the semantic distance between two concepts. The new gazetteer is not tied to any specific ontology format, because it uses GATE's generic ontology model (see section 4.2).

One of the reasons behind connecting ontologies to the gazetteer lists is to make the process of editing gazetteer entries part of ontology maintenance. Another advantage is that it enables GATE PRs to annotate automatically in the texts the instances already available in the user's ontology.

The ontology-based gazetteer PR comes with a corresponding VR that supports editing of the gazetteer lists. The user can also specify how they are connected to ontology classes. Figure 4 shows the ontology with all instance names displayed on the right as a gazetteer list. There can be more than one list of instance names per concept, as shown in this example. In this way, lists provided with the GATE system can be used simultaneously with lists derived from instances in the user's ontology. In cases where text is annotated as matching a gazetteer list entry coming from the ontology, the annotation provides information about its ontological class and the identifier of the ontology itself. Using this information, subsequent PRs can find and query the ontology for reasoning purposes.

The JAPE transducer was also extended to provide access to ontologies on the right-hand side of JAPE rules. This allows JAPE rules to add new information to ontologies (e.g. instances) or to consult them (e.g. to obtain semantic distance between concepts). We are also planning to extend the JAPE transducer to take

[8] WEKA homepage: `http://www.cs.waikato.ac.nz/ml/weka/`

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

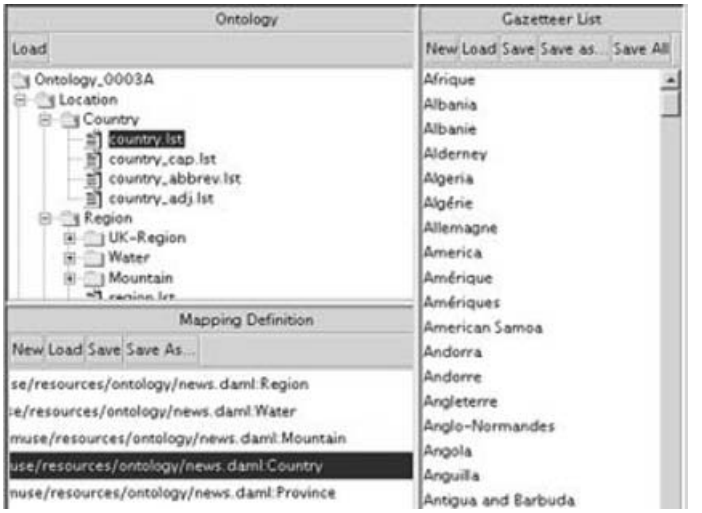


Fig. 4. Connecting the user's ontology to the GATE gazetteers.

into account subsumption relations in the ontology when matching on the left-hand side. So for example, a rule might look for an organization followed by a location, in order to create the `locatedAt` relationship between them. If JAPE takes into account subsumption, then the rule will automatically match all sub-classes of `Organisation` in the ontology, e.g. `Company`, `GovernmentOrg`.

4 New types of language resources

The two most frequently used language resources in GATE are documents and corpora (collections of documents). Documents have content, features (e.g. encoding, language), and sets of annotations organised in annotation graphs (see section 2).

However, experience from GATE v1 showed the need for other types of LRs, such as lexicons, ontologies, and thesauri, because they are used frequently in many LE applications and common APIs would greatly facilitate their reuse.

As already discussed, existing GATE PRs like gazetteer lookup and JAPE transducer were extended to use the new ontology LR. Similarly, other existing PRs would need to be extended (e.g. by adding a new parameter) if they need to access the new LRs, e.g. Wordnet. One way around this problem is to pre-compile certain information from lexicon LRs as gazetteer lists, which are used instead at run-time by the original gazetteer PR. For example, lists of jobtitles were compiled semi-automatically from WordNet and then used for named entity recognition and anaphora resolution. Naturally, this approach is feasible only when a small, already known number of queries need to be made to the lexicon.

4.1 Lexicons

Lexicons are a required resource in many language analysis modules. More specifically, in the context of applying NLP technology for the Semantic Web, semantic

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

resources like thesauri and semantic lexicons like WordNet play an important role (Buitelaar and Declerck 2003). Consequently, GATE's LR model was extended to cover certain new types of lexical resources.

A difficulty in lexicon reuse is that each resource has its own representation syntax and covers a particular subset of linguistic phenomena. GATE models lexicons in an object-oriented way, yielding an initial set of lexicon-specific classes. The commonalities between resources are exploited by subsequent, incremental integration of these classes into one common object model. The representation of objects within this model is geared towards generally accepted linguistic concepts.

This approach has several advantages: the linguistic knowledge is modelled in a conceptual and maximally uniform way; the linguistic resources are represented by maximally uniform data structures; overlaps and differences between resources can be identified by means of the common object model; the model provides a level playing field for the evaluation of lexical resources. However, in describing the more specific parts of the model GATE sticks closely to the terminology and structures of the existing resources.

In other words, the lexicon model starts at a very abstract level defining a lexicon as a collection of lemmas which have part-of-speech and several other features, which tend to be shared among most lexicons (e.g. definitions). If a concrete lexicon integrated with this model does not provide some of the information (e.g. definitions) then the lexicon wrapper will return an empty value. At lower level of the lexicon class hierarchy are more specific models of lexicons, e.g. the WordNet model which also contains synonym sets and support for its semantic relations like hyponymy, hypernymy, and meronymy.

Each lexicon is integrated with GATE via a wrapper which is a class towards the bottom of GATE's lexicon class hierarchy. Typically the wrapper follows the structure of the given lexicon and thus defines the lexicon-specific model. However, due to the existence of the lexicon class hierarchy, a PR which only needs to perform lexical lookup in order to obtain part of speech information can use the more general lexicon model, and access the newly integrated lexicon via that more general model, without need for any adaptation.

This idea is similar to the way GATE supports different document formats, but the common data model of annotations insulates the PRs from having to handle each of the formats. At the same time, however, the formatting information is preserved, so the PR can use or analyse it if required.

The most extensive support is currently provided for WordNet, because it is the most widely used lexical database, is freely available, and also has semantic relations, which make it suitable for Semantic Web applications. The GATE object-oriented model of WordNet follows closely that of the original resource and is based on synonym sets and the semantic relations defined between them (Miller 1995). This makes it accessible programmatically by any GATE PR via this WordNet API or the more generic lexicon API. In addition, a WordNet VR (see Figure 5) is provided to enable people to browse the resource using a graphical interface similar to that of the original WordNet browser. The advantages of implementing a new Java-based

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press



Fig. 5. The GATE WordNet browser.

browser as part of GATE are:

- reusability: it is easy to reuse and integrate with other Java applications. It is also easily extendable, because it has a well-defined, public API and open source code.
- generality: it uses the GATE lexicon object-oriented model, which means that it can be used to display other lexicons with a similar structure, not just WordNet.

Recently the lexicon class hierarchy in GATE was extended to support editable lexicons, as required by language generation applications, which often contain domain-specific terms not contained in existing NLP lexicons. A corresponding lexicon editor was also implemented. This is ongoing work carried out as part of the MIAKT project (Bontcheva 2004).

4.2 Ontologies

Recent Semantic Web (Fensel *et al.* 2002) developments have increased the need for LE applications that use ontologies (Fensel 2001) and knowledge bases in order to obtain a formal representation of and possibly reason about their domain (e.g. Saggion, Cunningham, Bontcheva, Maynard, Hamza and Wilks 2003a; Handschuh, Staab and Ciravegna 2002). In addition, LE methods can be used to populate an ontology with new instances discovered in texts or even to construct an ontology automatically.

Consequently, GATE was extended to provide support for importing, accessing and visualising ontologies as a new type of Language Resource. Strictly speaking, ontologies (and knowledge bases) should not be called LRs, because they are language independent. However, for backwards compatibility, we decided to keep the original GATE terminology of LRs. Under this definition Language Resources

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

cover all types of declarative resources used for language processing, e.g. grammars, lexicons, ontologies, as well as corpora and documents.

Some of the ontology functionality is provided through the integration of the Protégé editor (Noy, Sintek, Decker, Crubzy, Ferguson and Musen 2001). We also developed GATE-specific ontology visualisation, because it facilitates the process of developing and testing PRs that use ontologies (e.g. gazetteers and JAPE-based modules). This viewer was also incorporated recently in a new GATE editor for connecting NLG lexicons to ontologies (Bontcheva 2004).

The proliferation of ontology formats for the Semantic Web, e.g. DAML+OIL (Horrocks and van Harmelen 2001), OWL (Bechhofer, van Harmelen, Hendler, Horrocks, McGuinness, Patel-Schneider and Stein 2003), RDF(S) (Lassila and Swick 1999), means that LE applications need to deal with these different formats. In order to avoid the cost of having to parse and represent ontologies in each of these formats in each LE application, a common object-oriented model of ontologies with a unified API (Application Programming Interface) was created.

This approach has well-proven benefits, because it enables each PR to use this format-independent ontology model, thus making the PRs immune to changes in the underlining formats. If a new format needs to be supported, the PRs will automatically start using ontologies in this format, due to the seamless format conversion provided by GATE. From a language engineer's perspective the advantage is that they only need to learn one API and model, rather than having to learn many different and rather idiosyncratic ontology formats.

Since OWL, DAML+OIL and RDF(S) have different expressive powers, GATE's ontology model consists of a class hierarchy with a growing level of expressivity. At the top is a taxonomy class which is capable of representing taxonomies of concepts, instances, and inheritance between them. Multiple inheritance is not supported.

At the next level is an ontology class which can represent also properties, i.e. relate concepts to other concepts or instances. Properties can have cardinality restrictions and be symmetric and/or transitive. There are also methods providing access to their sub- and super-properties and inverse properties. The property model distinguishes between object (relating two concepts) and datatype properties (relating a concept and a datatype such as string or number).

The expressivity of this ontology model is aimed at being broadly equivalent to OWL Lite. In the case of a DAML+OIL ontology, GATE uses a sub-set of Jena's API to read in the model and instantiate the GATE ontology classes. Any information outside the GATE model is ignored. When reading RDFS, which is less expressive than OWL Lite, GATE only instantiates the information provided by that model. If the API is used to access one of these unsupported features, it returns empty values.

The ontology support is currently being used in several projects that develop Information Extraction (IE) to populate ontologies with instances derived from texts (Multiflora (Wood, Lydon, Tablan, Maynard and Cunningham 2003), hTechSight (see Section 6)) and also in the MIAKT project (Bontcheva 2004) for language generation from ontologies. The IE systems produce annotations referring to the concepts and instances in the ontology, based on their URIs (Uniform Resource Identifiers).

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

In order to enable quantitative evaluation of such ontology-based IE systems, as part of the EU-funded SEKT project,[9] GATE will be extended with document annotation facilities that support semi-automatic annotation given an ontology, so appropriately annotated corpora can be created. This editor will be based on GATE's annotation schemas which will be extended with optional `ontology`, `class`, and `instance` features with string values.

5 An extendable model of visual resources

Visual resources are providing a graphical user interface for viewing and editing of LRs (e.g. documents) and PRs (e.g. gazetteer lists). The component model of GATE was extended in version 2 to include VRs, in order to enable extensibility of GATE's development environment and reuse of VRs in other applications, thus overcoming an important limitation of the earlier GATE release. In other words, the need for reusable and flexible GUI components resulted in one of the biggest changes in the GATE architecture from version 1.

VRs are discovered by GATE using the CREOLE XML meta-data configuration, used for PRs and LRs as well. The meta-data for VRs is as follows:

```
<RESOURCE>
  <NAME>Annotations Editor</NAME>
  <CLASS>gate.gui.DocumentEditor</CLASS>
  <!-- type values can be "large" or "small"-->
  <GUI TYPE="large">
    <MAIN_VIEWER/>
    <RESOURCE_DISPLAYED>gate.Document</RESOURCE_DISPLAYED>
  </GUI>
</RESOURCE>
```

The name and class attributes are the same as for any GATE resource. The GUI information specifies the size of the window in which it is to be displayed, whether it is the main viewer for this resource (resources can have more than one VR associated with them), and which resource type is displayed/editied by this VR.

The main GATE GUI uses this information to configure itself automatically. This means that if a new type of resource is introduced (e.g. ontology), a corresponding VR can be implemented and added to the CREOLE repository (e.g. an ontology editor) and without further programming this VR will become available in GATE's visual development environment, associated with the resources it can handle. Inheritance is used when determining which VRs can show which resources, e.g. a VR associated with documents will show the three subclasses of database, file-based, and transient documents. If a new (more specialised) type of LR/PR is introduced, existing less specialised VRs can be used until a new VR, supporting the full functionality of the new resource, is developed.

[9] http://sekt.semanticweb.org. The SEKT partners are: British Telecommunications Plc.; Empolis GmbH; University of Sheffield; University of Karlsruhe; Jozef Stefan Institute; Institut fur Informatik der Universitat Innsbruck; Intelligent Software Components S. A.; Kea-pro GmbH; Ontoprise GmbH; Sirma AI Ltd; Vrije Universiteit Amsterdam; Autonomous University of Barcelona.

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

Currently GATE has 11 VRs in its default configuration and 6 optional ones. Example of available VRs are:

- Document VRs: display document content and the associated annotations (see Figure 1), including coreference chains.
- Annotation VRs: display/editing of annotations (see section 5.1), including specialised editors for tree structures (e.g. sentence trees).
- Corpus VR: allows the creation and editing of corpora.
- Other VRs associated with LRs – ontology VRs, e.g. the onto-gazetteer VR (Figure 4); lexicon VRs: handle lexicons, such as WordNet (see Figure 5); Information Retrieval VR: displays the IR results as a ranked list of documents.
- VRs associated with applications: allow creation and editing of NLP applications (see Figure 2).
- VRs associated with PRs: enable editing of the resources used by PRs, such as the gazetteer lists used by the gazetteer lookup PR; also show the values of the PR parameters that were specified at creation time.

5.1 Corpus annotation facilities

Since many NLP algorithms require annotated corpora for training, GATE's development environment now provides easy-to-use and extendable facilities for text annotation. These facilities were tested extensively by building corpora of name entity annotated texts for the MUSE, ACE, and MUMIS applications (Maynard, Tablan, Ursu, Cunningham and Wilks 2001; Saggion, Cunningham, Bontcheva, Maynard, Hamza and Wilks 2003b). The annotation tools follow the new VR model and are easy to embed and reuse in other applications.

The annotation can be done manually by the user or semi-automatically by running some processing resources over the corpus and then correcting/adding new annotations manually (similar to the pre-tagging functionality in the Alembic workbench (Day, Aberdeen, Hirschman, Kozierok, Robinson and Vilain 1997)). Depending on the information that needs to be annotated, some ANNIE modules can be used or adapted to bootstrap the corpus annotation task. For example, users from the humanities created a gazetteer list with 18th century place names in London which, when supplied to the ANNIE gazetteer, enabled the automatic annotation of location information in a large collection of 18th century court reports from the Old Bailey in London (Bontcheva, Maynard, Cunningham and Saggion 2002b).

Since manual annotation is a difficult and error-prone task, GATE v2 tries to make it simple to use and yet keep it flexible. To add a new annotation, one selects the text with the mouse (e.g. "Mr. Clever") and then clicks on the desired annotation type (e.g. Person) in the list of types on the right-hand-side of the document viewer (see Figure 1). If however the desired annotation type does not appear there or the user wants to provide more detailed information (not just the type), then an annotation editing dialogue can be used. The XML-based annotation schemas (see Section 2.2) are used to facilitate and validate the user input. This approach is

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

similar to NITE's use of XML schemas for data validation and XML stylesheets for dynamic creation of annotation interfaces (Carletta, Evert, Heid, Kilgour, Robertson and Voormann 2003).

6 Applications

In this section we describe briefly some of the LE applications we have developed using the new features introduced in GATE v2. The MUSE system (Maynard *et al.* 2001) is a multi-purpose Named Entity recognition system which is capable of processing texts from widely different domains and genres, thereby aiming to reduce the need for costly and time-consuming adaptation of existing resources to new applications and domains. MUSE used extensively many of the new features like the regression testing tools (Cunningham *et al.* 2002) and the conditional execution strategy (section 2.4). The MUSE system also exploited the machine learning facilities in GATE in order to perform NE recognition in Arabic and Chinese.

The hTechSight project[10] performs ontology-based IE for semantic tagging of job adverts, news and reports in the chemical engineering domain. The system is built using the ontology support in GATE v2. More specifically, the domain ontology (in DAML+OIL) is accessed via GATE's ontology model, which makes the IE components independent from the concrete ontology format, i.e. the same IE system could be run with an OWL ontology without any change to the PRs. The terminological gazetteer lists were linked to classes in the ontology model, so gazetteer lookup provided ontological information. In addition, the GATE ontology model was used on the right-hand side of the hTechSight JAPE grammars, so newly discovered instances (e.g. new job position) were added to the ontology. The GATE ontology visualisation tools were used during development and debugging.

Another project that benefitted from the new ontology and lexicon support in GATE v2 is MIAKT (Bontcheva 2004). It is developing a set of GATE-compatible tools for editing, storage, and maintenance of NLG lexicons. The model of NLG lexicons is an extension of the GATE lexicon model with methods for adding new lexical entries. The ontology model was used as a basis for a mapping tool which connects lexical entries to concepts and instances in the ontology. This was required by the MIAKT application in order to generate reports from ontologies in diverse formats, e.g. OWL, DAML+OIL.

GATE v2 has also been used for corpus annotation (e.g. in the American National Corpus project (Macleod, Ide and Grishman 2002)), multimedia indexing (Saggion *et al.* 2003b), automatic summarisation (Saggion *et al.* 2003a), e-science (Wood *et al.* 2003), and digital libraries (Bontcheva, Brewster, Cuiravegna, Cunningham, Guthrie, Gaizauskas and Wilks 2002a) (e.g. the Perseus digital library (Crane 2002)). Further details about some of the projects using GATE are available at `http://gate.ac.uk/projects.html`.

[10] `http://http://prise-serv.cpe.surrey.ac.uk/techsight/`

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

7 Related work

GATE draws from a large pool of previous work on infrastructures, architectures and development environments for representing and processing language resources, corpora, and annotations. Due to space limitations here we will discuss only a small subset. For a detailed review and its use for deriving the desiderata for this architecture see Cunningham (2000).

The system most similar in terms of functionality and spirit to GATE is IBM's TEXTRACT architecture (Neff *et al.* 2004). Similar to GATE, TEXTRACT has a graphical development environment, a number of reusable components (called plugins), a finite state transducer that allows rapid application prototyping, and support for processing of large corpora. The main differences from GATE are in the data representation model, the implementation language (C++ vs Java), and the restriction to language analyser applications. TEXTRACT has led to a new system called UIMA (Ferrucci and Lally 2004).

Work on standard ways to deal with XML data is also relevant here, such as the LT XML work at Edinburgh (Thompson and McKelvie 1997), WHAT (Schäfer 2003), and NITE (Carletta *et al.* 2003), as is work on managing collections of documents and their formats, e.g. Brugman, Russel and Wittenburg (1998), Grishman (1997) and Zajac (1998). We have also drawn from work on representing information about text and speech, e.g. Brugman *et al.* (1998), Mikheev and Finch (1997), Zajac (1998) and Young, Kershaw, Odell, Ollason, Valtchev and Woodland (1999), as well as recent annotation standards, such as the ATLAS project (an architecture for linguistic annotation) at LDC (Bird *et al.* 2000b). Our approach is also related to work on user interfaces to architectural facilities such as development environments, e.g. Brugman, Russel and Wittenburg (1998), Grishman (1997) and Zajac (1998), and to work on comparing different versions of information, e.g. Paggio (1998).

Another relevant area is work on indexing complex structures in annotations. The Corpus Query System (Christ 1994) is a frequently cited source in this area, providing indexing and search of corpora and later of WordNet. Similar ideas have been implemented in CUE (Mason 1998) for indexing and search of annotated corpora. Some work on indexing in the LT XML system is reported in McKelvie and Mikheev (1998). Bird *et al.* (2000a) propose a query language for the LDC annotation graph model, called AGQL. Cassidy (2002) discusses the use of XQuery as an annotation query language and concludes that it is good for dealing with hierarchical data models like XML, but needs extending with better support for network data models like annotation graphs. In comparison, GATE v2 indexes and retrieves annotations by storing them in a relational database, indexed by type, attributes and their values, allowing more than one annotation graph per document. In this way, it is possible to retrieve all documents that contain a given attribute and/or value or retrieve all annotations of a given type in a corpus, without having to traverse each document separately. The query language used is SQL.

Flexible and extendable visualisation of the complex linguistic structures frequently encoded in corpora was the focus of the NITE project (Carletta *et al.* 2003). The goal is to provide XML-based tools to help users with building their own specialised graphical interfaces to manipulate complex data in heavily-annotated

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

multimodal corpora. Some of the ideas are similar to GATE's VR model and both GATE and NITE allow the use of XML schemas to validate user input.

What makes this work particularly novel is that it addresses a very wide range of issues in LE application development in a flexible and extensible infrastructure. In addition, it promotes robustness, re-usability, experimental repeatability and scalability as important principles that help with the construction of practical LE systems. This paper showed how the original architecture was extended in response to new challenges by implementing models of lexical and ontological resources, machine learning, IR, and reusable VRs. Where possible, existing and widely used implementations of such components were integrated with GATE.

8 Discussion

The GATE architecture has been created with generality in mind. The rationale behind that design decision is to make GATE maximally extendable and independent of any particular NLP application (e.g. IE, IR). However, generality often entails more complex data structures and non-optimised implementations. For example, the new annotation graph data model is more generic than the TIPSTER one and can accomodate NLG components. On the other hand, we found it to be less efficient when storing LRs in databases, due to the larger number of database tables and existence of more than one annotation graph per document. In other words, there is a trade-off between generality and performance and while GATE's emphasis is generality, industrially-motivated LE architectures such as TEXTRACT (Neff *et al.* 2004) make performance their main priority.

The new types of LRs: lexicons and ontologies – did not require any changes to the architecture. Instead, the LR hierarchy was enhanced with new classes. The enhancement of the LR model to include ontologies created a slight terminological problem because ontologies are not typically considered as modelling language. However, the essential distinction between data resources (LRs) and algorithms that use them (PRs) remains valid and maybe, in future releases of GATE, LRs need to renamed to DRs (data resources) to reflect their wider range and function.

The main change to the architecture, as already discussed in section 5, came from the addition of Visual Resources (VRs) which enabled reuse and extensibility of the GATE GUI components.

9 Conclusion and future work

In this paper we have described recent extensions to a widely-used open-source infrastructure for language engineering software – GATE – which aims to assist the development of robust applications and resources for NLP through a generic set of tools and models. We believe that some of the main reasons for its success are its extensible and highly customisable nature, which allows the system to evolve easily in response to new challenges, and its open-source distribution and extensive documentation and online user support. The system now has thousands of users at hundreds of sites, and we hope to build on this user base to continue to meet the challenges posed by the rapid evolution of the LE field.

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

One future direction is extending the system to handle language generation modules, to enable the construction of applications which require language production in addition to analysis, e.g. intelligent report generation from IE data. In order to facilitate reuse and interoperability we plan on building on the RAGS architecture for language generation (Mellish, Reape, Cahill, Evans, Paiva and Scott 2004). We are currently developing generation tools for the Semantic Web in the context of the MIAKT project. These requirements for NLG support come from the need for automatic documentation of ontologies and also for intelligent knowledge publishing (Bontcheva *et al.* 2001).

Another future direction is towards providing better facilities for writing and debugging GATE's finite-state grammars, similar to those provided by the TEXTRACT architecture (Neff *et al.* 2004). The main concern is to enable non-expert users to modify existing and write their own language processing tools, based on GATE's JAPE pattern-action language. The requirements arose in the context of our digital library projects (Bontcheva *et al.* 2002).

The biggest challenge comes from grid-based e-science, the Semantic Web, and Web services, where LE applications can and should play a vital role. A number of existing projects, e.g. MIAKT, hTechSight, use Web services to provide interoperability between the system's components, which frequently run on different hardware and software platforms. Therefore, the existing GATE facilities need to be extended further in order to enable:

- easy deployment of grid-based LE applications, capable of utilising the vast computing facilities of the Grid to perform high-speed language processing on big corpora;
- providing a better synergy between Semantic Web and LE technology by extending the support for ontologies and knowledge bases and integration of reasoning and ontology support infrastructures such as Sesame[11] or KIM (Popov, Kiryakov, Kirilov, Manov, Ognyanoff and Goranov 2004).
- use of Web Services for building distributed LE applications.

A closer synergy between Semantic Web and LE infrastructures and standards is starting to emerge from the perspective of improving resource reusability and interoperability through development of standards for encoding annotation data. Recent efforts are aimed at the creation of a framework for linguistic annotations based on XML, RDF, and DAML+OIL/OWL standards for defining the semantics of the annotations (Ide and Romary 2004). GATE will benefit directly from such standardisation efforts, because they will improve the compatibility between NLP components by specifying common definitions of language resources such as lexicons.

Among these, Grid-based computing and Web services support are most likely to have architectural implications for GATE, amounting to more than integration of existing tools or implementation of format converters. This work will be carried out as part of the SEKT project.

[11] http://sesame.aidministrator.nl/

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

References

Abney, S. (1996) Partial parsing via finite state cascades. *Natural Language Engineering*, **2**(4): 337–344.

Bechhofer, S., van Harmelen, F., Hendler, J., Horrocks, I., McGuinness, D. L., Patel-Schneider, P. F. and Stein, L. A. (2003) OWL Web Ontology Language Reference. Technical report, W3C Proposed Recommendation 15 December 2003. `http://www.w3.org/TR/2003/PR-owl-ref-20031215/`, 2003.

Bird, S., Buneman, P. and Tan, W. (2000a) Towards a query language for annotation graphs. *Proceedings Second International Conference on Language Resources and Evaluation*, Athens.

Bird, S., Day, D., Garofolo, J., Henderson, J., Laprun, C. and Liberman, M. (2000b) ATLAS: A flexible and extensible architecture for linguistic annotation. *Proceedings Second International Conference on Language Resources and Evaluation*, Athens.

Bontcheva, K. (2004) Open-source tools for creation, maintenance, and storage of lexical resources for language generation from ontologies. *Proceedings of 4th Language Resources and Evaluation Conference (LREC'04)*.

Bontcheva, K., Brewster, C., Ciravegna, F., Cunningham, H., Guthrie, L., Gaizauskas, R. and Wilks, Y. (2001) Using HLT for acquiring, retrieving and publishing knowledge in AKT: position paper. *Workshop on Human Language Technology and Knowledge Management*, Toulouse, France. `http://www.elsnet.org/acl2001-hlt+km.html`.

Bontcheva, K., Maynard, D., Cunningham, H. and Saggion, H. (2002) Using human language technology for automatic annotation and indexing of digital library content. *Proceedings 6th European Conference on Research and Advanced Technology for Digital Libraries (ECDL'2002)*, Rome, Italy.

Brugman, H., Russel, H. and Wittenburg, P. (1998) An infrastructure for collaboratively building and using multimedia corpora in the humaniora. *Proceedings ED-MEDIA/ED-TELECOM Conference*, Freiburg.

Buitelaar, P. and Declerck, T. (2003) Linguistic annotation for the semantic web. In: S. Handschuh and S. Staab, editors, *Annotation for the Semantic Web*. IOS Press.

Burnard, L. (1995) Users Reference Guide for the British National Corpus. `http://info.ox.ac.uk/bnc/`.

Carletta, J., Evert, S., Heid, U., Kilgour, J., Robertson, J. and Voormann, H. (2003) The NITE XML Toolkit: flexible annotation for multi-modal language data. Special issue on Measuring Behavior. *Behavior Research Methods, Instruments, and Computers*, **35**(3).

Cassidy, S. (2002) Xquery as an annotation query language: a use case analysis. *Proceedings 3rd Language Resources and Evaluation Conference (LREC'2002)*, Gran Canaria, Spain.

Christ, O. (1994) A modular and flexible architecture for an integrated corpus query system. *Proceedings 3rd Conference on Computational Lexicography and Text Research (COMPLEX '94)*, Budapest. `http://xxx.lanl.gov/abs/cs.CL/9408005`.

Crane, G. (2002) Cultural heritage digital libraries: needs and components. *Proceedings 6th European Conference on Digital Libraries*, pp. 626–637. Rome, Italy. Springer-Verlag.

Cunningham, H. (2000) *Software Architecture for Language Engineering*. Unpublished PhD thesis, University of Sheffield. `http://gate.ac.uk/sale/thesis/`.

Cunningham, H. (2002) GATE, a general architecture for text engineering. *Computers and the Humanities*, **36**: 223–254.

Cunningham, H., Humphreys, K., Gaizauskas, R. and Wilks, Y. (1997) Software infrastructure for natural language processing. *Proceedings 5th Conference on Applied Natural Language Processing (ANLP-97)*. `http://xxx.lanl.gov/abs/cs.CL/9702005`.

Cunningham, H., Maynard, D., Bontcheva, K. and Tablan, V. (2002) GATE: A framework and graphical development environment for robust NLP tools and applications. *Proceedings 40th Anniversary Meeting of the Association for Computational Linguistics (ACL'02)*.

Cutting, D., Pedersen, J. and Halvorsen, P.-K. (1991) An object-oriented architecture for text retrieval. *Proceedings RIAO '91*, pp. 285–298. Barcelona, Spain.

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

Davies, J., Fensel, D. and van Harmelen, F. (editors) (2002) *Towards the Semantic Web: Ontology-driven Knowledge Management*. Wiley.

Day, D., Aberdeen, J., Hirschman, L., Kozierok, R., Robinson, P. and Vilain, M. (1997) Mixed-initiative development of language processing systems. *Proceedings 5th Conference on Applied Natural Language Processing (ANLP-97)*.

Fensel, D. (2001) *Ontologies: Silver Bullet for Knowledge Management and Electronic Commerce*. Springer-Verlag.

Fensel, D., Wahlster, W. and Lieberman, H. (editors) (2002) *Spinning the Semantic Web: Bringing the World Wide Web to Its Full Potential*. MIT Press.

Ferrucci, D. and Lally, A. (2004) UIMA: An architectural approach to unstructured information processing in the corporate research environment. *Natural Language Engineering*. This issue.

Grishman, R. (1997) TIPSTER Architecture Design Document Version 2.3. Technical report, DARPA. `http://www.itl.nist.gov/div894/894.02/related_projects/tipster/`.

Handschuh, S., Staab, S. and Ciravegna, F. (2002) S-CREAM – Semi-automatic CREAtion of Metadata. *13th International Conference on Knowledge Engineering and Knowledge Management (EKAW02)*, pp. 358–372. Siguenza, Spain.

Horrocks, I. and van Harmelen, F. (2001) Reference Description of the DAML+OIL (March 2001) Ontology Markup Language. Technical report. `http://www.daml.org/2001/03/reference.html`.

Ide, N., Bonhomme, P. and Romary, L. (2000) XCES: An XML-based standard for linguistic corpora. *Proceedings Second International Language Resources and Evaluation Conference (LREC)*, pp. 825–830. Athens, Greece.

Ide, N. and Romary, L. (2004) Standards for language resources. *Natural Language Engineering*. This issue.

Lassila, O. and Swick, R. (1999) Resource Description Framework (RDF) Model and Syntax Specification. Technical Report 19990222, W3C Consortium, `http://www.w3.org/TR/REC-rdf-syntax/`.

Ma, X., Lee, H., Bird, S. and Maeda, K. (2002) Models and tools for collaborative annotation. *Proceedings 3rd Language Resources and Evaluation Conference (LREC'2002)*, Gran Canaria, Spain.

Macleod, C., Ide, N. and Grishman, R. (2002) The American National Corpus: Standardized resources for American English. *Proceedings 2nd Language Resources and Evaluation Conference (LREC)*, pp. 831–836. Athens, Greece.

Mason, O. (1998) The CUE Corpus Access Tool. *Workshop on Distributing and Accessing Linguistic Resources*, pp. 20–27. Granada, Spain. `http://www.dcs.shef.ac.uk/~hamish/dalr/`.

Maynard, D., Tablan, V., Ursu, C., Cunningham, H. and Wilks, Y. (2001) Named entity recognition from diverse text types. *Recent Advances in Natural Language Processing 2001 Conference*, pp. 257–274. Tzigov Chark, Bulgaria.

McEnery, A., Baker, P., Gaizauskas, R. and Cunningham, H. (2000) EMILLE: Building a corpus of South Asian languages. *Vivek, A Quarterly in Artificial Intelligence*, **13**(3): 23–32.

McKelvie, D. and Mikheev, A. (1998) Indexing SGML files using LT NSL. LT Index documentation, from `http://www.ltg.ed.ac.uk/`.

Mellish, C., Scott, D., Cahill, L., Evans, R., Paiva, D. and Reape, M. (2004) A reference architecture for generation systems. *Natural Language Engineering*. This issue.

Mikheev, A. and Finch, S. (1997) A workbench for finding structure in text. In *Fifth Conference on Applied NLP (ANLP-97)*, Washington, DC.

Miller, G. (1995) WordNet: a lexical database for English. *Communications of the ACM*, **38**(11).

Neff, M. S., Byrd, R. J. and Boguraev, B. K. (2004) The Talent System: TEXTRACT architecture and data model. *Natural Language Engineering*. This issue.

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press

Noy, N., Sintek, M., Decker, S., Crubzy, M., Fergerson, R. and Musen, M. (2001) Creating semantic Web contents with Protégé-2000. *IEEE Intelligent Systems*, **16**(2): 60–71.

Paggio, P. (1998) Validating the TEMAA LE evalutation methodology: a case study on Danish spelling checkers. *Journal of Natural Language Engineering*, **4**(3): 211–228.

Popov, B., Kiryakov, A., Kirilov, A., Manov, D., Ognyanoff, D. and Goranov, M. (2004) KIM – Semantic annotation platform. *Natural Language Engineering*. This issue.

Saggion, H., Bontcheva, K. and Cunningham, H. (2003a) Robust generic and query-based summarisation. *Proceedings European Chapter of Computational Linguistics (EACL), Research Notes and Demos*.

Saggion, H., Cunningham, H., Bontcheva, K., Maynard, D., Hamza, O. and Wilks, Y. (2003b) Multimedia Indexing through Multisource and Multilingual Information Extraction; the MUMIS project. *Data and Knowledge Engineering*. To appear.

Schäfer, U. (2003) WHAT: An XSLT-based infrastructure for the integration of natural language processing components. *HLT-NAACL 2003 Workshop: Software Engineering and Architecture of Language Technology Systems (SEALTS)*, pp. 9–16.

Tablan, V., Ursu, C., Bontcheva, K., Cunningham, H., Maynard, D., Hamza, O., McEnery, T., Baker, P. and Leisher, M. (2002) A Unicode-based environment for creation and use of language resources. *3rd Language Resources and Evaluation Conference*.

Thompson, H. and McKelvie, D. (1997) Hyperlink semantics for standoff markup of read-only documents. *Proceedings SGML Europe'97*, Barcelona.

Witten, I. H. and Frank, E. (1999) *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*. Morgan Kaufmann.

Wood, M. M., Lydon, S. J., Tablan, V., Maynard, D. and Cunningham, H. (2003) Using parallel texts to improve recall in IE. *Recent Advances in Natural Language Processing*, Bulgaria.

Young, S., Kershaw, D., Odell, J., Ollason, D., Valtchev, V. and Woodland, P. (1999) *The HTK Book (Version 2.2)*. Entropic Ltd., Cambridge. `ftp://ftp.entropic.com/pub/htk/`.

Zajac, R. (1998) Reuse and integration of NLP components in the Calypso architecture. *Workshop on Distributing and Accessing Linguistic Resources*, pp. 34–40. Granada, Spain. `http://www.dcs.shef.ac.uk/~hamish/dalr/`.

https://doi.org/10.1017/S1351324904003468 Published online by Cambridge University Press