

Usability of Annotation Tools: comparison and outlook

Diana Maynard¹, Martin Dzbor², and Siegfried Handschuh³

¹ Department of Computer Science, University of Sheffield
Regent Court, 211 Portobello Street, Sheffield, UK
`diana@dcs.shef.ac.uk`

² Knowledge Media Institute, Open University,
Walton Hall, Milton Keynes, UK
`m.dzbor@open.ac.uk`

³ National University of Galway,
University Road, Galway, Ireland
`Siegfried.Handschuh@deri.org`

Abstract. This paper investigates the state of the art in automatic textual annotation tools, and examines the extent to which they are ready for use in the real world. We define some benchmarking criteria for measuring the usability of annotation tools, and examine those factors which are particularly important for a real user to be able to determine which is the most suitable tool for their use. We discuss factors such as usability, accessibility, interoperability and scalability, and evaluate a set of annotation tools according to these factors. Finally, we draw some conclusions about the current state of research in annotation and make some suggestions for the future.

Key words: annotation, benchmarking, usability

1 Introduction

Ontology-based annotation tools are used to populate ontologies with instances from text, and/or to annotate text with conceptual information from an ontology. This task forms an important part of ontology creation and management, by enabling us to combine and associate existing ontologies, perform more detailed analysis of the text, and to extract deeper and more accurate knowledge. Ontology population is a crucial part of knowledge base construction and maintenance that enables us to relate text to ontologies, providing on the one hand a customised ontology related to the data and domain with which we are concerned, and on the other hand a richer ontology which can be used for a variety of semantic web-related tasks such as knowledge management, information retrieval, question answering, semantic desktop applications, and so on.

The term "annotation" can be used to mean many different things: for example, simply adding comments to a document; adding information about document structure or composition, author, document type etc; adding linguistic

information such as part-of-speech tagging, and so on. Here we refer to textual annotation as the process (or result) generally performed by means of some kind of ontology-based information extraction (OBIE). This consists of identifying the key terms in the text (such as named entities and technical terms) and then relating them to concepts in the ontology. In this paper we are concerned with tools that perform automatic rather than manual annotation.

There has been some previous work aimed at comparing and/or evaluating different annotation tools. However, this largely examines only the performance of the information extraction component and is based on a single set of task-specific data (see for example [1]). Their aim is simply to examine the state of the art in annotation performance, paying little heed to the demands of real users of semantic web technologies. What such previous work ignores is that there are many more aspects to evaluating a tool than just measuring its performance in a particular situation. In our view, a complementary comparative study is required to look into issues going beyond performance as measured by precision and recall in quantitative evaluations. In this paper we analyse usability, accessibility, scalability and interoperability, and investigate the extent to which some current annotation tools are ready for the demands of users. This aims to enable users to draw some conclusions about which (if any) tools are best suited to their needs. We make no claims that any one tool is a priori better than any other, because it depends on many factors such as who will be using it, what they will be using it for, and which features they care about most (e.g. speed vs. usability).

2 Requirements

According to [2], there are 3 groups of annotation tool users:

- Annotators: those who just want to annotate data without worrying about representation, design, or how the tool works.
- Annotation Consumers: those who want to use annotated data for various purposes, and who need to be able to visualise, search and query annotated data and annotations.
- Developers: there are two main types:
 - Corpus developers: corpus designers who may wish to add new annotation schema, map existing data to a new structure, etc.
 - System developers: programmers who may wish to alter or extend the functionality of the tool, and who therefore need to understand how it works, add new components etc.

In reality these groups may not be so well defined, for example, corpus developers are very often the same people as the annotators. However we should be aware of the differing abilities of these groups of users. We would not necessarily expect annotators to have linguistic skills, let alone computer skills, although they will probably have domain knowledge, whereas system developers can be expected to have computer skills but not necessarily domain knowledge or

linguistic skills. Corpus developers will generally have linguistic skills and probably domain knowledge, while annotation consumers will probably have domain knowledge but not necessarily linguistic or computing skills.

Table 1 shows the extent to which each kind of requirement is important for the different users. A requirement which is not applicable (e.g. correctness of performance is not applicable when manually annotating) is denoted by "-". A single "+" indicates that a requirement may be important or is slightly important; "++" indicates a greater importance; "+++" would indicate that a requirement is absolutely essential.

User	Manual Annotator	Annotation Consumer	Corpus Developer	System Developer
Usability	++	+	++	+
Flexibility	+	+	++	++
Performance	-	++	+	+
Scalability	-	++	+	++
Interoperability	+	+	+	++

Table 1. User Requirements for annotation tools

One of the main problems with designing automatic annotation tools (common to many other software tools) is the tradeoff between generality and specificity. Even though annotation may appear to be a quite straightforward task with a clear objective, both the uses and users of annotation tools may differ widely, and unless the tool is designed for a specific clear purpose with a particular kind of user in mind, there will almost always be dissatisfaction somewhere. One of the most important criteria is therefore that the annotation tool should be flexible and easily adaptable and/or extendable to the user's needs. This could range from something as simple as being able to change the colour of the annotation, through to enabling a whole new kind of visual representation or even modality.

3 Annotation Tools

The textual annotation tools we have chosen to investigate are GATE, KIM, OntoMat, MnM and Magpie. These have been chosen for a number of reasons, but mainly because they all perform in some way automatic annotation of textual data with respect to an ontology (although they are quite diverse); they are all XML-based, open source, readily available and do not require extensive training to use. These tools have all been previously evaluated to some extent regarding performance, but have not been directly compared with respect to other criteria such as those we present.

3.1 GATE

GATE [3] is an architecture for human language technology which contains, amongst other things, tools and resources to perform textual annotation both manually and automatically (using information extraction techniques). GATE comes with a default information extraction system called ANNIE [4], which is designed to be a starting point for adaptation to a specific domain, language, ontology or application. The ANNIE system identifies generic concepts such as persons, locations, organisation, dates, etc., so most tasks require the development of new rules or adaptation of existing rules in order to find instances of different kinds of concepts. GATE also contains components for ontology-based information extraction [5].

3.2 KIM

KIM [6] offers an end-to-end, extendable system which addresses the complete cycle of metadata creation, storage, and semantic-based search and includes a set of front-ends for online use, that offer semantically enhanced browsing. The essence of the KIM IE engine is the recognition of named entities with respect to the KIM ontology, which is achieved using a modified version of ANNIE [7]. The entity instances all bear unique identifiers that allow annotations to be linked both to the entity type and to the exact individual in the instance base. For new (previously unknown) entities, new identifiers are allocated and assigned; then minimal descriptions are added to the semantic repository. The annotations are kept separately from the content, and an API for their management is provided.

3.3 Magpie

Magpie [8] is a suite of tools which supports the interpretation of webpages and "collaborative sense-making", by annotating a text with instances from a known ontology. These instances can be used as a confidence measure for carrying out some services. It uses an ontology to provide a very specific and personalised viewpoint of the webpages the user wishes to browse. In terms of its named entity recognition capabilities, Magpie relies mainly on recognising existing instances in the ontology, which are stored in a simple lexicon to speed up search. The main method used is string matching between text and ontology instance, although some lexicon-based pattern matching is also carried out, using components from GATE (for example to combine a known first name of a person with an unknown surname).

3.4 MnM

MnM [9] is an annotation tool which provides both automated and semi-automated support for annotating web pages with semantic content. MnM integrates a web browser with an ontology editor and provides open APIs to link to ontology servers and for integrating information extraction tools. This enables a user to

mark up web resources by instantiating generic concepts from a standardised terminology (an ontology) for a particular domain. MnM learns how to recognise new objects that require annotation, by learning from a collection of previously annotated documents (markup derived from pre-existing ontologies).

MnM, along with OntoMat, was one of the earliest supervised IE systems for annotation; it was designed to mark up training data (corpus) for IE tools rather than as an annotation tool per se. This means that it stores marked up documents as tagged versions of the original (storing annotations directly in the document), rather than the RDF formats used by the Semantic Web community. In the Semantic Web, documents and their annotations are usually stored separately, so that documents and annotations can be owned by different people or organisations and stored in different places.

3.5 OntoMat

OntoMat-Annotizer [10] is an interactive annotation tool for web pages, supporting the user in the task of creating and maintaining ontology-based OWL markups, i.e. creating instances, attributes and relationships. It includes an ontology browser for the exploration of the ontology and instances, and an HTML browser that displays the annotated text. It is Java-based and provides a plugin interface for extensions. The intended user is the individual annotator, i.e. somebody who wants to enrich their web pages with OWL metadata. Instead of manually annotating the page with a text editor, OntoMat allows the annotator to highlight relevant parts of the web page and create new instances via drag'n'drop interactions.

4 Evaluation Criteria

In this section, we describe and discuss the evaluation of the annotation tools previously described, in terms of criteria such as usability, interoperability and accessibility issues. We describe our own experiences with the tools, augmented by comments and insights via a small survey of actual users of the tools. Table 2 shows a summary of the results (the codes are explained in the description below).

4.1 Interoperability issues

Interoperability is concerned with how well the tool interacts with other tools and systems. Annotation is a task that is often combined with other applications, such as browsing, search and retrieval, indexing, etc., so it is important that annotation tools can easily interact with other systems. This is best achieved by conformance to existing standards.

Interoperability evaluation not only covers annotation format, but also issues such as:

Tool	MAGPIE	MnM	KIM	OntoMat	GATE
Version tested	v1.0	v1	v1.05	0.8a	v3.0
Interoperability					
Platform	W,L,M	W,L	W	W,L,M	W,L,M
Browser	IE,MF	own	IE	own	own
Browser variation	yes	n/a	n/a	n/a	n/a
Data format	HTML,TXT	HTML,TXT	All	HTML,TXT	All
Ontology format	own	DAML+OIL, RDF	OWL Lite	OWL Lite	OWL, Lite
Source available	no	no	no	yes	yes
Usability					
Installation ease	+2	+1(Win)	+2	+2	+2
Installation doc	-1	0	+1	0	+1
Doc quality	+2	-1	+2	-1	+1
Doc format	SI,I,M	TI	TI,I,M	TI,I	TI,I,M
Linked help	no	no	yes	no	yes
Configuration	+2	-1	+1	+2	-1
Aesthetics	+2	-1	+2	-1	+2
Range of tasks	+1	-1	+1	+1	+2
Accessibility					
Graphics	+2	+1	+1	0	0
Tooltips	yes	no	some	some	yes
Mouse alternative	no	no	no	some	no
Changeable colours	yes	yes	n/a	yes	yes
Changeable menu fonts	no	no	no	no	yes
Changeable text fonts	n/a	no	annots only	no	yes

Table 2. Assessment of functionality of textual annotation tools

- data format: what kinds of text format can be processed, e.g. xml, html, sgml, txt, etc.;
- annotation schemes: whether annotation schemes can be imported/exported from other tools;
- plugins: if it is possible to plug in other tools and applications;
- API: whether the tool provides some API to programmatically access it.

Under the topic of interoperability we investigate factors such as which platforms and browsers the tool runs on, ontology formats possible and how easily the tool can be modified.

Platform This factor considers which platforms the tool can be run on, according to the documentation and/or discussion with the providers. We consider only Windows (W), Linux (L) and Mac (M), as these are the 3 main platforms

most likely to be used. All tools worked with both Windows and Linux apart from KIM which only worked with Windows. GATE, OntoMat and MAGPIE also worked with Macs.

Browser Here we consider which browsers the tool works with: Internet Explorer (IE), Mozilla or Mozilla Firefox (MF) or its own proprietary browser (own). Here MAGPIE was the only tool which worked with both IE and Firefox. KIM worked only with IE, while MnM and OntoMat have their own proprietary browser and GATE does not use a browser as such but imports the documents into its own interface (which is not browser-like). GATE can, however, be run via its API as a web service on any browser.

Browser variation This question looked at if and how there are any differences when the tool is run with different browsers. This only therefore applied to MAGPIE as the other tools only work with a single browser. It was found that there are quite a few differences with different browsers: when running on Firefox the performance was slightly substandard and there were some unexpected happenings; there were also conflicts between browser and plugin commands on Firefox. It is expected that later versions of MAGPIE will not have this problem, however.

Ontology format This looks at which ontology formats are compatible with the tool. All the tools are compatible with OWL except for MnM (which is an older tool and no longer supported) which uses DAML+OIL and RDF, GATE is also compatible with RDF. At the time of testing, MAGPIE used its own proprietary ontology format but its current version supports OWL.

We also investigated specifically the OWL interoperability of the tools, by testing the ability of the tools to import and export OWL files without data loss. In fact, only GATE and KIM (which use the same underlying ontology repository, OWLIM [11]) could be benchmarked in this way, out of the annotation tools under investigation, as the others either did not support OWL or could not be tested for various reasons. Compared with other semantic web tools, GATE (and KIM) actually performed very well in this respect, only suffering from some minor data loss. In particular, it interacted very well with Protege, KAON, JENA and SWI-Prolog, which were also among the best performing tools in the experiment. More details of this experiment can be found in [12]. We can conclude from this that there seems to be a benefit for a semantic annotation tool to make use of a semantic repository such as OWLIM, which assists interoperability with ontology editors, support for ontology formats, etc.

Data format Here we look at what format the text needs to be in for the tool to be able to process it. GATE scores highly here by being able to process many kinds of textual format, including plain text, SGML, XML, HTML, RTF, Word, and PDF. KIM also can process the same formats as GATE as it uses the

functionalities from GATE for document format analysis. The other tools only really process texts that can be displayed in a browser, i.e. HTML, plain text, XML etc.

Source available This question looks at whether the source code is freely available so that developers can extend or modify the tool as required, for example adding new annotation sets, new visualisation capabilities, new processing resources, etc. This is a very important part of the flexibility and extensibility of the tool, since as discussed in Section 2, if the type of user and tasks for which the tool is to be used remain unknown or unpredictable at the time of design, then the tool needs to be able to cater for such flexibility if it is to fulfil interoperability requirements and be widespread in its use. The only tools of those tested that, to our knowledge, have this capability are OntoMat and GATE. However, there may be extensibility advantages with tools such as KIM that have open APIs.

4.2 Usability

Under the topic of usability, we categorise the quite broad issues such as documentation, ease of setup and installation, aesthetics of design, and range of tasks possible. Some more specific issues concerning accessibility are categorised separately in the following section. Responses in this section are mostly allocated a score ranging from -2 to +2. The lowest score (-2) generally reflects a "very poor" answer, -1 reflects "poor", 0 reflects "OK", +1 reflects "good" and +2 reflects "very good". In the case of questions such as "setup" the scores could be interpreted as ranging from "very difficult" through to "very easy".

Installation ease First we look at how easy the tool is to install. While this is clearly related to the following question about the installation documentation, the two may be orthogonal as the installation documentation could be very poor but installation may still be very easy. Indeed we can see this from the results, for example MAGPIE was deemed very easy to install (+2) but the installation documentation was deemed poor (-1). In fact, all tools were deemed very easy to install except for MnM, which was deemed as easy to install for Windows but not for Linux.

Installation documentation As mentioned above, the quality of the installation documentation did not necessarily correlate with the ease of installation of the tool. The installation documentation was evaluated separately from the general tool documentation or user guide. MAGPIE was the only tool to have a negative score for installation documentation, which correlates with the difficulty of installation. In the case of KIM, GATE and OntoMat, the quality of the installation documentation turned out not to be that important since the tools were easy to install anyway.

Documentation quality We then turned our attention to the quality of the main documentation. Here we look first of all at the quality in general, and then we look at some aspects more specifically (the format of the document and how easy it is to find help). MAGPIE and KIM had very good documentation, which was clear, comprehensive and easy to follow. MnM's documentation was clear but very basic, while KIM's documentation was not very clear. GATE's documentation was very comprehensive but quite confusing and the users found it hard to locate the relevant information because it was so detailed.

Documentation format Here we look at the format of the documentation. This links very closely to the quality of the documentation, showing very clearly that the more modalities in the documentation, the higher the quality. We investigated 4 aspects of the document format: whether it had step-by-step instructions (SI), images/screenshots (I), movies/demos (M) or just simple textual instructions (TI). MAGPIE was the only tool which had clear step-by-step instructions, and also had both images and movies. OntoMat and GATE both had textual instructions, images and movies, while KIM had textual instructions and images, and MnM just had textual instructions. Clearly the combination of all text, movies and images was the clearest for users, although step-by-step instructions were not necessarily an improvement on simple basic instructions (both MAGPIE and Ontomat scored +2 on documentation quality although only MAGPIE had the step-by-step instructions). As with the installation instructions, however, a tool might be easy to use even though the instructions are poor.

Linked help A final part of the instruction examination looked at whether there was a help facility available directly from the tool. This could be either just a link to the documentation (as in the case of GATE) or a specific help function similar to most proprietary programs. It is very useful to have direct access to help without having to revert back to the website or downloaded instructions somewhere in the user's file system. MAGPIE had a button linking back to the website, from which the user guide could be found, but no direct link to a help facility. KIM had no help facility, while OntoMat and GATE both have linked help available: GATE links directly to the user guide via a button on the menu.

Configuration The configuration criterion looked at how easy the tool was to set up in the way that the user wanted. This does not include installation, but rather things like changing the appearance of the GUI to suit the user's needs, changing options such as whether the tool should save session on exit, layout of menus, different skins or "look and feel", altering the fonts, etc. It does not consider the actual presence or absence of such options, but is based on how easy it is to set up the tool according to the options given, i.e. how easy these options are to use. MAGPIE and OntoMat were deemed very easy to set up (possibly because not so many options were available), while KIM was deemed fairly easy.

MnM was considered quite hard to set up because of unnecessary and confusing dialogues, while GATE was deemed quite hard simply because many options were available and it was not always clear what they did, for example "Add space on markup unpack if needed". It is also not obvious in GATE that some options (such as changing the colour of annotations) are even possible, without reading the user guide.

Aesthetics This question looked at how pleasing the tool was generally for the user, in terms of appearance, attractiveness etc. MAGPIE, KIM and GATE all scored highly here, and were considered to be colourful and interesting. MnM and OntoMat were rated quite poorly, and considered to be dull and unappealing. The aesthetics of a tool seems to be quite highly linked with the overall goals of the tool: MnM was designed as a very simple tool for a fairly limited range of tasks, and therefore it seems that not much consideration was given to its attractiveness. GATE on the other hand is designed for a very broad range of users and applications, and since one of its objectives is to be used as widely as possible, much consideration has gone into the look and feel of the tool.

Range of tasks This question looked at how complex the tool was and at how wide a range of tasks could be achieved. All tools were classified as having a fairly wide range of tasks except for MnM, while GATE was considered to have a very wide range.

4.3 Accessibility

Software accessibility is essentially about making tools that are usable, easy to use and attractive to use for everyone (not just for people with disabilities) [13]. Some of the most important examples of accessibility problems stem from *inflexibility*. A well designed tool will have options to change the user's preferences regarding colours, layout, font sizes and styles, and so on, and the ability to save and restore latest sessions, etc.

Even though a user should be able to choose such options, the default options should also be well designed. For example, text should be in a mixture of upper and lower case, and colour schemes should incorporate contrasting colours in the same range. Icons should be clearly understandable, not just with alternative text on mouseover, but should also use clear symbols and be large enough to click on easily. Mouse alternatives should also be widely available, although in practice it is hard to annotate regions of text manually without using a mouse.

In general, GATE performed the best on the accessibility tasks, with KIM a close second. Again, this is largely because GATE is much more flexible than the other tools, allowing the user to set it up how they want and to change the look and feel according to their preference.

5 Scalability

One aspect of scaling up to the real world is to combat the problems of accuracy when systems are deployed in industrial applications, such as the enabling of public metadata-on-demand services. However, even in a closed domain with a tightly focused application, there are more mundane issues concerning the sheer volume of data, with respect to storage, processing speed and power. Scalability has been identified as a critical issue for the processing of large volumes of data, so that statistical IE algorithms can be designed and trained (since these require large amounts of annotated training data in order to be accurate). In this section, we discuss some research designed to investigate the feasibility of scaling annotation tools up to the demands of the real world, i.e. performing annotation on an industrial level with massive volumes of data and/or huge ontologies.

SWAN⁴ was a recent experiment in scaling up automated metadata extraction for industrial strength Semantic Web application development. The SWAN project leverages language technology to advance a number of existing elements covering both basic research and applications demonstrators. The SWAN experiment enabled its developers to determine the issues and problems in large-scale deployment of the technology in applications such as HLT for the Semantic Web, EAI and e-commerce applications, and to gain experience in scaling up their systems and in their application in next-generation knowledge access and management. It involved 3 systems: GATE, KIM and a focused crawler called SECO.

Following on from this, the KIM Cluster Architecture has been developed, partly within the context of the SEKT project⁵, with the specific aim of extensive scaling of the existing KIM model [14]. The system has been demonstrated on a corpus of over 1.2 million news articles, resulting in the recognition of over 1 million named entities. These were all stored in the semantic repository based on OWLIM, along with their semantic descriptions (attributes and properties). The average speed of annotation in this experiment was 10KB per annotator node/component.

The scalability of GATE has been demonstrated through a number of experiments running various language processing tasks over large textual collections [15]. The machine learning API is capable of running on huge amounts of data, e.g. very large ontologies (dependent on the hard disk size) and thousands of documents for training (depending on memory size and – for some learning algorithms – hard disk size). This is mainly due to the datastore mechanism in GATE, which stores immediate data and results on the hard disk rather than in memory, loading and unloading documents as they are processed.

The other 3 annotation tools examined were not really designed for large-scale applications, and this is their biggest drawback when using such tools for industrial rather than testbed research applications. The main aim of MnM was to provide proof-of-concept rather than a serious industrial-strength tool. MnM

⁴ <http://old.deri.ie/projects/swan/>

⁵ <http://www.sekt-project.com>

has serious problems when used with even a medium-sized ontology, and since it relies on training, a user first has to manually annotate documents, which is infeasible on a large scale and with an ontology of any significant size.

One aspect of scalability that showed particularly strongly in the case of Magpie was the relationship between the desire to scale up as opposed to the real time responsiveness of the application to the user's interaction. Thus pragmatically, scale is only an added value of a particular annotation tool as far as the user is willing to wait for the outputs: for GATE, this is easily in the extent of several hours or even days, since the output is a marked-up collection of documents. For Magpie, the time limits are in the range of seconds – otherwise, the user is less likely to achieve the main objective of this tool: navigation on the Web using semantic relationships and links.

In summary, GATE and KIM are designed to be scalable and to be applicable in real world scenarios. Extensive experiments - both on the tools themselves and with their use in industrial settings - have shown that they can cope reasonably well with large volumes of data and real-life ontologies without sacrificing performance. There are, however, some limitations: GATE applications may be quite slow depending on both the approach used and the size and content of the texts. MnM was never designed to be used in a large-scale environment and therefore suffers from a number of problems (some of which are indeed insurmountable without modifying the actual tool) when applied to large ontologies. Similarly, OntoMat was never designed for scalability and can suffer problems when used in such an environment. In general, any system that relies on learning will have problems when applied to a new ontology, as it will require a large amount of manually annotated training data which is time-consuming to produce.

6 Recommendations and Conclusions

In this paper, we have introduced a set of different annotation tools and investigated various aspects such as their aims, usability and scalability. Our aim is not to try to identify which tool is the best, but rather to examine the features a user should look for when deciding on a tool to use, and to highlight some issues which should be factors in that user's choice.

It is clear from our research that the most important thing to consider is the task to which the tool will be put and the situation in which it will be used. For this, we must first examine two main factors: the type of user, and the scope of tasks on which the tool will be put to use. Once these have been identified, we can turn our attention to more details: for example, if the tool is to be used solely by annotators, then a high degree of usability straight from the box may be important; if the tool is to be used only for a single, standalone kind of task, interoperability may be less of an issue than if the tool is to be used within a larger suite of tools, and so on. In Section 2 we established a set of requirements and investigated the importance of each with respect to the kind of user. In Table 3 we summarise the extent to which the annotation tools we have analysed fulfil these requirements. Here ”-” indicates the absence or near

absence of that requirement, while the number of + signs indicate the level. While it appears that GATE is in general the best performing tool overall, it really depends what a user’s requirements are. For example, if a simple or highly scalable tool is required, it may not be the best choice.

Tool	GATE	KIM	MnM	Magpie	OntoMat
Usability	++	++	-	++	+
Flexibility	++	+	-	+	+
Performance	+++	++	+	++	++
Scalability	++	+++	-	++	+
Interoperability	++	++	-	+	++

Table 3. Requirements fulfilled by annotation tools

In this paper, we have not approached the subject of performance of the tools. This is for two reasons: (1) performance evaluations of annotation tools have been documented elsewhere, e.g. [15, 1, 16], and in the Dot.Kom and SEKT projects; (2) previous evaluations have shown that it is difficult to compare different tools on the same documents and with the same ontology, since they are all designed for different purposes and work best on the kind of data they were designed for, so any comparison may lead to unfair prejudice. For example, in the case of Magpie, the precision on annotating text that is within the domain of the loaded lexicon is far superior to most other tools; however, the recall tends to decline rapidly if the text moves beyond the domain of the lexicon. This is a well-known feature of ontology-driven annotation, often presented as ontology brittleness [17]. We refer the reader, however, to [15] for a look at the performance of the tools in question in a number of experiments.

The problem of domain dependence is also related to the issue of algorithmic reuse. The IE systems underlying the annotation tools are largely tailor-made for specific domains and applications, with the result that not only are they hard to adapt for new tasks, but that it is difficult to extricate potentially reusable components or sub-components which are buried deep in the architecture. For example, there may not be a distinction between foreground information, which is dependent on the domain and application, and background information, which can be reused as it stands; and consequently, between the tools needed to access and manipulate these two types of information.

Annotation tools have already proved their success in a number of real world applications, such as Del.ici.ous, Flickr, digital libraries such as Perseus⁶, Garlik⁷ (which mines data about consumers present in various sources including the web), Fizzback⁸ (which provides real-time customer feedback from SMS and

⁶ <http://www.perseus.tufts.edu/>

⁷ <http://www.garlik.com>

⁸ <http://www.fizzback.com>

email feeds) and so on. However, there are several reasons why semantic annotation is not more widespread. First, it is time-consuming and complex to produce annotations in open domains. Second, there are not enough DIY cases: Flickr and Delicious have tapped into obvious needs, and Innovantage hopes to do the same, but in many cases we do not have enough folk for a folksonomy. For example, broadcast archives are a Catch-22 situation: people are not aware that they need them until they use them, and they can't use them until widespread need generates either human effort or funding. Automatic methods for annotation, on the other hand, have made huge advances in recent years to a usable level, but the output still tends to be incomplete or inaccurate, especially in open domains. One solution is therefore to combine automatic methods with human annotation at the lowest possible cost, and preferably done by non-experts. While we have plenty of algorithms, data structures and evaluation protocols emerging, we still currently lack a clear statement of how to specify and implement new annotation tasks, especially those oriented towards non-HLT experts. What is therefore needed is a methodology covering how to:

- decide if annotation is applicable to a problem;
- define the problem with reference to a set of examples;
- identify similarities with other problems and thus to estimate likely performance levels;
- design the annotation workflow, including automatic assistance;
- measure success.

These tasks are beyond the scope of the current work; however, they are currently being pursued in projects such as NeOn⁹ and MUSING¹⁰, by researchers at Harvard Medical School, and by some commercial users.

Acknowledgements The research for this paper was conducted as part of the EU Network of Excellence Knowledge Web.

References

1. Sazedj, P., Pinto, H.S.: Time to evaluate: Targeting annotation tools. In: Proceedings of ISWC 2005 SemAnnot Workshop, Galway, Ireland (2005)
2. Reidsma, D., Jovanovic, N., Hofs, D.: Designing annotation tools based on properties of annotation problems. In: Measuring Behavior 2005, 5th International Conference on Methods and Techniques in Behavioral Research, Wageningen, The Netherlands (2005)
3. Cunningham, H., Maynard, D., Bontcheva, K., Tablan, V.: GATE: A Framework and Graphical Development Environment for Robust NLP Tools and Applications. In: Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistics (ACL'02). (2002)

⁹ <http://www.neon-project.org>

¹⁰ <http://www.musing-project.eu>

4. Maynard, D., Tablan, V., Ursu, C., Cunningham, H., Wilks, Y.: Named Entity Recognition from Diverse Text Types. In: Recent Advances in Natural Language Processing 2001 Conference, Tzigrav Chark, Bulgaria (2001) 257–274
5. Li, Y., Bontcheva, K., Cunningham, H.: Hierarchical, Perceptron-like Learning for Ontology Based Information Extraction. In: 16th International World Wide Web Conference (WWW2007). (May 2007) 777–786
6. Popov, B., Kiryakov, A., Kirilov, A., Manov, D., Ognyanoff, D., Goranov, M.: KIM – Semantic Annotation Platform. *Natural Language Engineering* (2004)
7. Maynard, D., Tablan, V., Cunningham, H., Ursu, C., Saggion, H., Bontcheva, K., Wilks, Y.: Architectural Elements of Language Engineering Robustness. *Journal of Natural Language Engineering – Special Issue on Robust Methods in Analysis of Natural Language Data* **8**(2/3) (2002) 257–274
8. Domingue, J., Dzbor, M., Motta, E.: Magpie: Supporting Browsing and Navigation on the Semantic Web. In Nunes, N., Rich, C., eds.: *Proceedings ACM Conference on Intelligent User Interfaces (IUI)*. (2004) 191–197
9. Motta, E., Vargas-Vera, M., Domingue, J., Lanzoni, M., Stutt, A., Ciravegna, F.: MnM: Ontology Driven Semi-Automatic and Automatic Support for Semantic Markup. In: 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW02), Siguenza, Spain (2002) 379–391
10. Handschuh, S., Staab, S., Ciravegna, F.: S-CREAM — Semi-automatic CREation of Metadata. In: 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW02), Siguenza, Spain (2002) 358–372
11. Kiryakov, A.: OWLIM: balancing between scalable repository and light-weight reasoner. In: *Proc. of WWW2006*, Edinburgh, Scotland (2006)
12. García-Castro, R., David, S., Prieto-González, J.: D1.2.2.1.2 benchmarking the interoperability of ontology development tools using owl as interchange language. Technical report, Knowledge Web (September 2007)
13. Maynard, D.: Benchmarking ontology-based annotation tools for the semantic web. In: UK e-Science Programme All Hands Meeting (AHM2005) Workshop ”Text Mining, e-Research and Grid-enabled Language Technology”, Nottingham, UK (2005)
14. Manov, D., Popov, B.: D2.6.1 Massive Automatic Annotation. Technical report, SEKT EU Project Deliverable (2005)
15. Maynard, D., Dasiopolou, S., Costache, S., Eckert, K., Stuckenschmidt, H., Dzbor, M., Handschuh, S.: Benchmarking of annotation tools. Technical Report D1.2.2.1.3, KnowledgeWeb Deliverable (2007)
16. Maynard, D., Li, Y., Peters, W.: Nlp techniques for term extraction and ontology population. In Buitelaar, P., Cimiano, P., eds.: *Bridging the Gap between Text and Knowledge - Selected Contributions to Ontology Learning and Population from Text*. IOS Press (2007)
17. Dzbor, M., Motta, E.: Study on Integrating Semantic Applications with Magpie. In: 15th Intl. Conference on Artificial Intelligence: Methods, Systems & Applications (AIMSA), Bulgaria (2006)