Magpie: Supporting Browsing and Navigation on the Semantic Web

John Domingue

Martin Dzbor

Enrico Motta

Knowledge Media Institute,
The Open University
Milton Keynes, United Kingdom

+44 1908 655014
j.b.domingue@open.ac.uk

+44 1908 858524
m.dzbor@open.ac.uk

+44 1908 653506
e.motta@open.ac.uk

ABSTRACT

We describe several advanced functionalities of Magpie – a tool that assists users with interpreting the web resources. Magpie is an extension to the Internet Explorer that automatically creates a semantic layer for web pages using a user-selected ontology. Semantic layers are annotations of a web page, with a set of applicable semantic services attached to the annotated items. We argue that the ability to generate different semantic layers for a web resource is vital to support the interpretation of web pages. Moreover, the assignment of semantic web services to the entities allows users to browse their neighbourhood semantically. At the same time, the Magpie suite offers *trigger functionality* based on the patterns of an automatically updated semantic log. The benefits of such an approach are illustrated by a semantically enriched browsing history management.

Categories and Subject Descriptors

H.5.4 [**Hypertext/Hypermedia**]: Architecture, Navigation, User Issues – *semantic web browsing, semantic services.*

General Terms

Performance, Experimentation, Human Factors

Keywords

Semantic Web, browsing history management, semantic web services, named entity recognition

1 INTRODUCTION

A lot of research has gone into supporting the task of finding web resources – by means of 'standard' information retrieval mechanisms or by means of semantically enhanced search [6, 13]. Less attention has been paid to the task of supporting the *interpretation* of web pages. Annotation technologies [8, 14] allow users to associate meta-data with web resources, which can then be used to facilitate their interpretation. The annotation technologies provide a useful way to support shared interpretation, but they are very limited; mainly because the annotation is carried out manually.

Hence, the quality of meta-data depends on the authors or librarians annotating the web page.

The majority of web pages are not semantically annotated. This is a great obstacle in a move towards the Semantic Web [1]. Magpie is a tool supporting the interpretation of web pages and acting as a complementary knowledge source, which a user can call upon to gain instantaneous access to the background knowledge relevant to a web resource. Magpie follows a different approach from that used by most other annotation techniques: it automatically associates a *semantic layer* to a web resource, rather than relying on a manual annotation.

This ability relies on *ontology* [5] – an explicit, declarative representation of a discourse. Ontologies are the cornerstone of the emerging semantic web: they provide conceptual interoperability, allow agents to 'understand' information on the web and to collaborate with other semantically aware agents. Magpie uses ontologies to associate meaning with the information found on a web page. Based on the identified meanings, relevant services can be invoked, or value-added functionalities offered to the user. The association between an ontology and a web resource provides an *interpretative viewpoint* or *context* for the resource in question. Indeed the overwhelming majority of web pages are created within a specific context.

Some readers of a web page might be very familiar with such a context, while others might not. In the latter case, Magpie is especially beneficial, given that the context is made explicit to the reader and context-specific functionalities are provided. One incentive for this kind of research was summed up by a seminal study of how users browse the web. Tauscher and Greenberg [16] presented the following statistics on the types of actions users carry out:

- 58% of pages visited are revisits,
- 90% of all user actions are related to navigation,
- 30% of navigation actions use the 'Back' button,
- less than 1% of navigation actions use a history mechanism

A fairly obvious conclusion from these statistics is that web users need support in capturing what they have seen previously. Current browsing history and bookmark tools are not effective. Magpie can automatically track concepts found during a browsing session using a *semantic log*. The log allows *trigger services* to be activated when a specific pattern of concepts has been found. The same log can be used as a conceptual representation of the user's browsing history. Since all Magpie abilities are underpinned by ontological reasoning, this enables the users to use the history semantically rather than as a purely linear and temporal record of their activities.
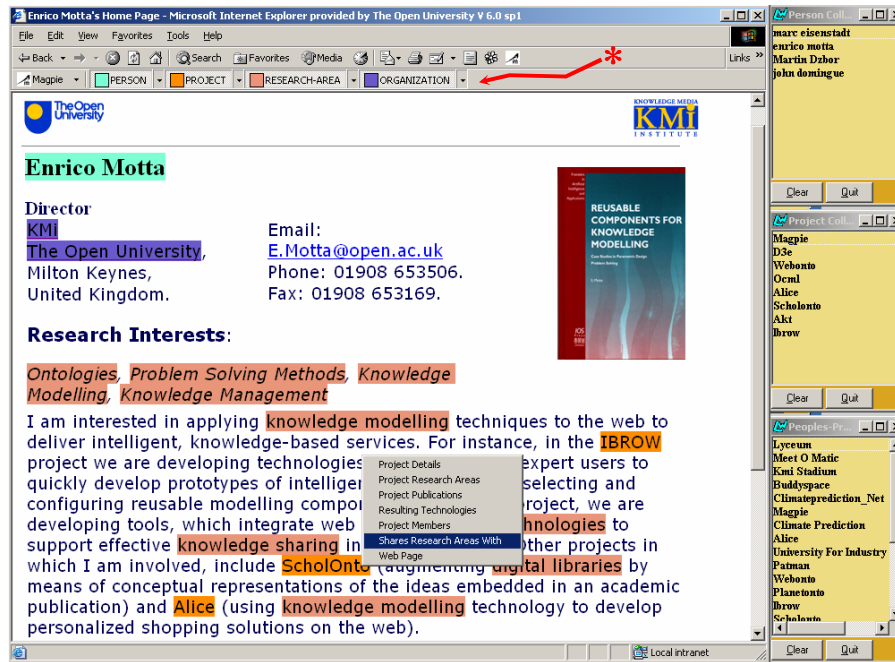
Fig. 1. Enrico Motta's home page viewed through Magpie. Known *people*, *organizations*, *projects* and *research areas* are highlighted using the Magpie toolbar (marked by '∗'). On the right-hand side are three Magpie collectors – the top two log the people and projects found in the browsing session. The bottom one shows the (not explicitly mentioned) projects associated with the people found.

2 MAGPIE USAGE SCENARIO

Assume a journalist is writing an article on the Knowledge Media Institute (KMi) for a magazine. She needs to gather information about the key projects led by senior KMi staff. Using a web browser with a Magpie extension, she visits the home page of the lab's director Enrico Motta. After loading it, she wants to quickly recognize interesting concepts denoting researchers, collaborating organizations, projects, and research areas in the page. These concepts draw on an existing ontology of academic organizations, which was populated by mining databases and web resources, and is available to the external users[1].

Fig. 1 shows the journalist's browser with the concepts of interest highlighted using the Magpie toolbar, which extends the functionality provided by Internet Explorer. As can be seen, Magpie preserves structure of the page, and highlights the concepts upon user's request. This approach reduces the confusion, which may occur when the content and/or appearance of a web page are altered. The Magpie toolbar (see marker '*' in Fig. 1) allows users to toggle highlighting of the specific class of entities, which were annotated using an ontology-derived lexicon. The classes are ontology dependent – changing the ontology generates new toolbar buttons. As ontology represents an *interpretative viewpoint* we leave the choice of ontology to the user.

On the right-hand side of Fig. 1 are three Magpie *collectors*. These are automatically filled by Magpie *trigger services* as the user browses. During a browsing session, the entities found on accessed web pages are asserted into a *semantic log* knowledge base (KB). Collectors show a semantically filtered view of the semantic log. For instance, the top two collectors in Fig. 1 show the people and projects that were recognized on any page visited during the current browsing session. The bottom collector shows the projects associated with any people recognized during the browsing session, which were not mentioned explicitly in any page but are known in the domain ontology. Fig. 1 shows a number of projects the four researchers from the top-right collector are associated with.

One of the highlighted concepts that have semantic meaning in a given ontology is 'ScholOnto'. A right-click on the 'ScholOnto' term invokes a *semantic services menu* as shown in Fig. 1. The menu options depend on the class of the selected entity within a particular ontology. In our case, 'ScholOnto' is a Project, so project-related options are displayed. The user selected the 'Shares Research Areas With' service, results of which form a list of related items and are shown in Fig. 2.

Some of the concepts listed in Fig. 2 are not explicitly present in Fig. 1. In other words, the journalist now takes advantage of using the context in which a particular page was written. This allows her to browse orthogonally to the syntactic, author-defined links using implicit relationships among different concepts known in a particular ontology.

3 MAGPIE ARCHITECTURE

The overall goal of this project is to support interpretation of web documents with no a-priori mark-up by means of adding an ontology-derived semantic layer. The main design principles to emphasize are: the ability to extend a standard web browser, preservation of web page appearance, and separation of content and semantic annotation. The full list of principles underlying the design of Magpie is discussed in [4].

[1] This example uses the AKT reference ontology, which is available at http://www.aktors.org/publications/ontology/.

Fig. 2. Results of the '*Shares Research Areas With*' semantic query invoked for the '*ScholOnto*' project by the semantic menu action depicted in Fig. 1. Each bullet shows a project followed by a list of overlapping research areas.

Magpie is essentially a bridge, a *mediator* between formal descriptions used by the ontology-based service providers and semantically unstructured web documents. The Magpie architecture comprises a *Service Provider* and a *Service Recipient* component. The web services terminology emphasizes multiple roles played by a web browser and an ontology-based server. In line with web services paradigm there may be many providers of the same service and many different services [15]. Currently, the Magpie central service provider is built around a suite of tools accessing a library of knowledge models containing domain ontologies, populated KBs, semantic services and a semantic log KB.

The ontological representations are shared by different services, so that they interpret the same concept in the same way. Services are semantically annotated to the ontological classes, which enables us to abstract from their implementation. Some services use a semantic log KB, a feature discussed in detail further in the paper. The actual services are also discussed further in this paper.

Details both conceptual and technical of the different components of Magpie infrastructure are beyond the scope of this particular paper. They are discussed in detail in earlier publication [4]. For the purposes of this paper, it is sufficient to highlight the key principles of how entities are annotated in the web page.

3.1 Magpie browser extension – IE plug-in

Magpie Browser Extension (further plug-in) is embedded in the user's web browser, and is responsible for managing the interaction between the user and the semantically enriched browser. The toolbar is a graphical user interface (GUI) to the underlying functionality. The web browser plug-in is built around a fast named entity recognition (NER) engine that recognizes and highlights ontological entities. Our NER engine uses an ontology-derived lexicon enhanced by a few simple heuristics, which work extremely fast.

The lexicon entries are generated from the instances within the ontological knowledge base, which is populated from various sources (e.g. databases). The heuristics include e.g. recognition of abbreviations or people's initials. The specific rules applicable to our scenario from the previous section use the AKT reference ontology. Instead of adding more complex NER techniques to the plug-in, we are experimenting with implementing the advanced NER algorithms as (semantic web) services available upon a user's request. This leaves the actual plug-in thin and fast.

When an entity of interest is recognized in the web page, the plug-in annotates it with customized `<SPAN…>` tags, and links it with a relevant ontological instance/class within the chosen ontology. This process creates *a semantic layer* over the original document, original content of which remains untouched. The interesting concepts and the corresponding text on the page are highlighted in response to user pressing a particular button on the Magpie toolbar. Simultaneously with the annotation, the recognized entities are passed on to the semantic log KB, where they are recorded and used by appropriate trigger services.

4   SEMANTIC SERVICES IN MAGPIE

In the previous section we briefly described how a semantic layer is created, displayed and activated. The main benefits of using Magpie however are generated from the ability to deploy semantic services on top of the semantic layer. These services are provided to the user as a physically independent layer over a particular HTML document. Magpie distinguishes between two types of semantic services, each having a specific user interaction model: *on-demand* and *trigger* services.

4.1   On-demand semantic services

Once the semantic entities on a web page are annotated, the contextual (right-click) menu of a web browser is overridden by an *on-demand services menu* whenever the mouse hovers over a recognized entity. The 'on-demand services' menu is also context-dependent as could be expected; however, in this case, it is a semantic context defined by the membership of a particular entity to a particular ontological class. The information on class membership is contained in the ontology or a lexicon generated from ontology.

One specific ontology on our ontology server formally defines what services can be attached to particular classes, and the semantics of their operations. The semantic services are defined and published in line with standards of the emerging web services technology [15]. In our scenario of Magpie serving as a semantic portal for organizational research, the services were defined for the individual ontological classes on the ontology server without any brokering. An example of a service for class *Project* is shown as a semantic menu displayed in the center of Fig. 1. Similarly to parsing and annotation done by the plug-in automatically, the 'on-demand services' menu is also generated on the fly.

Selecting an option in semantic services menu generates a request to the Magpie dispatcher to contact the appropriate service provider and perform the requested reasoning. The knowledge-level reasoning facilitated by the service provider provides context for a particular entity. This is delivered back to the web browser to be annotated and displayed. An example of a response is visible as a new browser window in the foreground of Fig. 2.

Hence, the Magpie plug-in in co-operation with the standard browser functionality facilitates two complementary methods of web browsing. First, syntactic browsing using the `<A HREF=…>` anchors inserted into a document by its author. The second browsing method uses the customized semantic anchors created during

Fig. 3. "Where am I" – a tool for semantic management, browsing, and search in the user's browsing history

the automatic annotation, and the dynamically generated semantic services. The former method accesses a <u>physically linked content</u>, whereas the latter method makes available the <u>semantic context</u>. Our interface differentiates the two methods to minimize confusion, and to emphasize the complementary nature of the two access mechanisms.

4.2 Trigger semantic services

User-requested (on-demand) semantic services are one method for interacting with the relevant background knowledge. A different type of service gaining popularity are various agents-recommenders or advisers. These are active or *push* services, and they differ from the on-demand ones by their tendency to "look over the user's shoulder", gather facts, and present conclusions. In other words, they tend to be data-driven.

In our case, a pre-condition for having active services is to keep *history logs* of browsing, particularly a log of the recognized entities. The label 'browsing history' is appropriate because a log accumulates findings not only from the current web page, but also from previously visited pages in the same browsing session. While an annotated web page is displayed in a browser, the recognized entities are asserted as facts into the Magpie semantic log KB.

Several *watchers* monitor the patterns in the asserted facts. When the relevant assertions have been made for a particular watcher, a semantic service response is *triggered*, and applicable knowledge delivered to the Magpie plug-in that in turn displays it in a dedicated window next to the user's web browser. In principle, this interaction is *asynchronous* – the service provider starts the communication, contacts the user's dispatcher, and pushes potentially relevant information to the user.

The information that can be pushed in this way may range from simple collections of relevant items to sophisticated guidance on browsing or browsing history visualization. Since the service provider (watcher) taps into a knowledge base constructed potentially from the logs of community members, the guidance or history visualization may draw on community knowledge and behaviors. This type of setup may seem surprising in the scenario presented earlier because a journalist is clearly *not a member* of KMi community. Does it make sense to send her community-relevant information?

We believe that this approach corresponds to a journalist adopting the viewpoint of a specific community to *interpret* and *make sense* of a given web resource *from the perspective* of that community. Thus, a formal membership of a particular community and the utilization of their ontological viewpoints are two different roles that can be distinguished when using Magpie. Since a trigger service can be (in principle) subscribed to, it is useful to tap into the knowledge of a community of which the user is *not a formal member*. On the contrary, this enables him or her to see the document in its 'native' context.

4.3 Semantic 'bookmarking'

The opportunity to adopt a particular viewpoint for interpreting web pages from a perspective of a given community has other benefits. One of them is the possibility to offer a dedicated tool for managing browsing history – using a particular viewpoint. Going back to the study of Tauscher and Greenberg [16] mentioned in the introduction, as many as 58% downloads of web pages are re-visits. One design recommendation from their study is that bookmarks should have a *meaningful* representation.

History management based on the semantics of the visited pages, and implemented by a *triggered* semantic layer may help to alleviate issues with the syntactic and linear (access time ordered) methods. Instead of search browsing history records in an unnatural (for humans) space of URLs and access dates, we give users the opportunity to search a conceptual space containing entities with specific semantic meanings. Instead of plain URLs, our semantic logging works with URLs that are annotated and associated with concepts such as "Projects" or "Research Areas".

A snapshot of Magpie "Where am I" interface is shown in Fig. 3. It consists of an iconic bar (A) displaying visited pages graphically; list of neighbouring pages (B), ontological footprint of a particular page (C), and a semantic filter interface (D). The visited pages can be either browsed in a standard linear fashion, or filtered and searched using concepts from a particular ontology. This is the same ontology that was originally used to annotate the pages. An ontological footprint can be defined as a set of annotations that were automatically extracted from a particular web page. A footprint is thus a summary of a particular page from the perspective of a given ontology. The combination of seeing a web page in its 'syntactic'[2] neighbourhood superimposed by the ontological footprint is the primary novelty introduced by our Magpie framework.

Semantic filters (see pointer D in Fig. 3) can be used to reduce the number of visible web pages from the user's browsing history. An example of such a conceptual query trying to find a particular subset of web pages is shown in Fig. 4. The underlying ontology supports inference over a user-formulated query. For example, in our case, no member of KMi is working directly in the area of visualization but in related areas (such as software debugging and telepresence). Similarly, the research interests are associated with 'research staff', which is recognized as a sub-class of 'academic'.

```
(and
    (academic ?X)
    (has-research-interest ?X visualization))
```

Fig. 4. An example of a semantic filter using concepts from the ontology, which finds and shows only the web pages containing academics who work on visualization (see results in Fig. 5)

5   OVERVIEW OF SIMILAR WORK

One of the inspirations for Magpie was the COHSE system [2]. COHSE combines an Open Hypermedia System with an ontology server into a framework for ontological linking – an ontology-derived lexicon is used to add links to arbitrary web pages. The links are added either by proxy server or by an augmented Mozilla™ browser. The distinctions between Magpie and COHSE are in their differing design goals. The goals for COHSE were (i) to separate web links and web pages, and (ii) to make these links conceptual (i.e. ontology-based). The goal for Magpie is to support interpretation and information gathering. Magpie's interface enables entities to be recognized and annotated on the client side. Instead of embedding new relevant links into a document, Magpie offers class-dependent semantic services for each entity found. Magpie also offers *trigger* services via semantic logs. Neither type

of Magpie service replaces traditional links; they are an auxiliary and extendible knowledge source available at the user's fingertips.

A number of tools support annotation of web pages. A classic example is the Amaya HTML editor that implements the Annotea infrastructure [8]. Annotea users may add various meta-statements to a document, which are separate from the document itself and are accessible to collaborating teams via a central annotation server. The annotation in this sense attaches additional information to a web page. This makes Annotea a powerful tool for joint authoring with a small group of collaborating agents sharing a common goal. However, this approach makes it difficult to facilitate annotation sharing in 'open' user communities. In these cases, there is no guarantee that a freely articulated annotation would convey the same meaning to the different users.

Unlike Magpie, Annotea assumes that someone (author?) invests additional effort into making a page semantically rich. Magpie is more liberal and assumes a reader only subscribes to a domain ontology, which is then used to channel background knowledge. It may be argued that ontology creation takes more effort than manual document mark-up. This is true; however, ontology as a domain model can be re-used for different purposes, not only the annotation of one document. The effort spent on designing a shared ontology is greater in the short term but in the longer term, it is a more cost-effective way of recording a shared point of view. Moreover, ontologies are increasingly available for download, so often, no development is actually required.

The CREAM-based Ont-O-Mat/Annotizer [7] is another tool integrating ontologies and information extraction tools. Annotations in this framework are close to those advocated in this paper. Any ontological instance, attribute or relation may be an annotation hook. A key feature of this tool is its use of *discourse representations* to structure the relatively flat output of information extraction tools according to the chosen ontology. CREAM's annotation inferences resemble our trigger services produced by a data-driven reasoning. On the other hand, our 'on-demand' services seamlessly address the awareness of the existing relationships and the actual context of ontological instances.

Sticky Notes [9] is a tool moving from the annotated web pages to using the annotations for *finding* items of interest. Their model focuses on the manual annotation, a kind of 'notes on the margin' of a particular web page. Being user-driven, a sticky note can reflect the user's intention or internal conceptualization (e.g. '*this idea shall improve turnover*' in addition to the explicit meta-data (e.g. '*authored-by*'). The authors propose a high-level language for querying such annotations. Magpie creates conceptual annotations of web pages automatically. Therefore, it cannot express users' intentions. However, the concepts used in annotation are sufficiently high-level to allow users to query their browsing history in these conceptual terms. While the 'Sticky Notes' approach helps with managing personal experiences, it does not lend itself to an automated support for sense-making. Our Magpie assists with *finding* the web pages, as well as their contextual *interpretation*.

Another strand of research relevant to Magpie framework involves Letizia [11], and the idea of a reconnaissance agent. Such an agent "looks out for the information relevant to the user". In terms of web browsing, pre-filtering the links from a web page may improve the relevance and usefulness of browsing. The functionality similar to that of Letizia ("local reconnaissance") is implemented in Magpie by semantic logging and ontological reasoning with the semantic log. Unlike Letizia, Magpie does not offer any recom-

---

[2] As before, by 'syntactic' we mean pages linked through `<A HREF=.../>` anchors as defined by web page author.
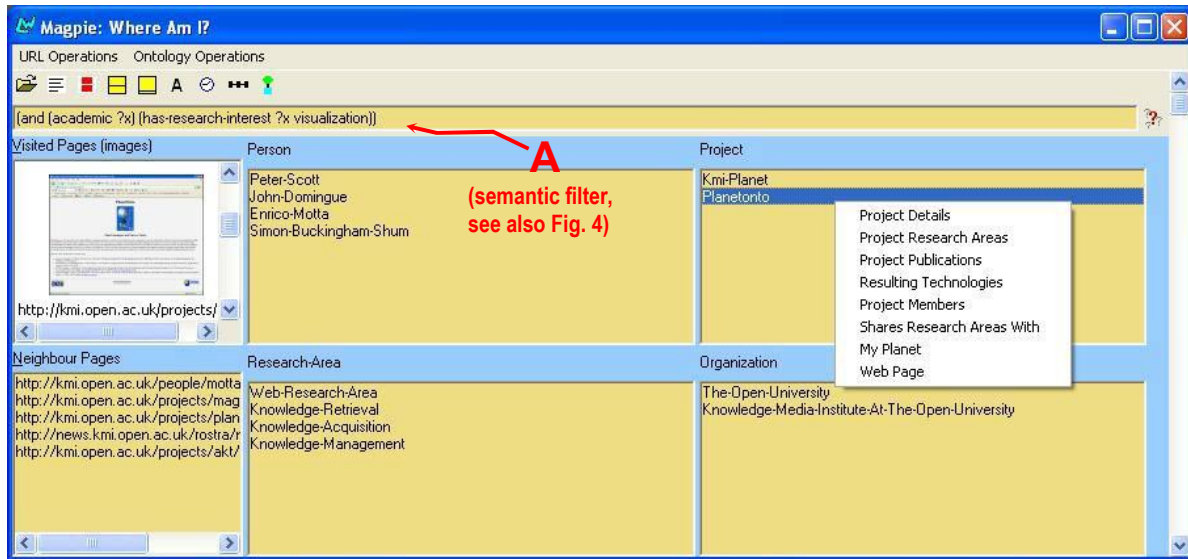
Fig. 5. "Where am I" – a semantic filter applied on the user's browsing history reduces the number of records shown

mendation. Furthermore, Magpie uses an available (and shared) ontological lexicon for NER and annotation rather than specific examples provided to an agent to generalize and learn patterns [10].

Our framework is focused on using the neighbourhood information for filtering and making sense of web pages visited in the *past*. The history access and awareness is important because it enables users to trace particular concepts back to the pages, where they were first encountered. Concept tracing throughout a browsing session was piloted in e-commerce applications such as the ZStep-based Woodstein tool [12] that used a (browsing) history to debug user-computer interaction during an e-commerce session. Another stream of research using high-level conceptual annotations of web pages to reason about them comprises the vision of a pervasive 'Memex' environment [3] for recording both an individual's as well as community browsing experience.

Another similarity between the concepts of reconnaissance agents and user interaction debuggers is the idea of "zero input" or "one-click" interaction. Similarly to Letizia and Woodstein, our Magpie uses the information already present in the web page without asking the user to provide any queries or search keywords. Our approach performs the annotation on the client's browser and provides a host of relevant service 'one click away' from the annotated web page. The services might be distributed around the Web. The only time when we break this rule in Magpie is our browsing history manager. In that case, it may be useful to allow the user formulate a simple query in a high-level conceptual language rather than logic (as most search engines do).

6   CONCLUSIONS

Reducing the information overload from the expanding web is often cited as the premise for work on supporting the retrieval of relevant documents. But finding relevant documents is only half of the story. Their interpretation involves a reader in understanding the context, in which the document was created. To gain the full insight, a reader requires knowledge of the specific terms mentioned and the implicit relationships contained both within the document and between the document and external knowledge sources. Magpie addresses this issue by capturing context within

an ontology, which then is used to enrich web documents with a semantic layer. Semantic services expose relevant segments of the ontology according to the user's needs. The choice of ontological viewpoint for interpreting a particular web page drives the interpretation bottom-up – by the user rather than domain expert or knowledge engineer.

Magpie users browse the web in a standard way with negligible differences in their user experience. Magpie achieves this by extending standard web browsers with standard mark-up languages, without altering the layout of the web page and imposing any significant time overhead. The key principle is that the user controls to what extent semantic browsing comes to the fore. The Magpie toolbar enables concepts highlighting according to their ontological class, and the Magpie infrastructure enables arbitrary semantic actions to be triggered by patterns of items found within a semantic log. Trigger services also allow certain tasks to be delegated. In the scenario we showed how discovered entities could be used for a later inspection. However, Magpie allows more complex trigger services to be implemented.

One example of such a complex service has been described in this paper as a tool for the semantic management of browsing history. This tool draws on the enriched semantic annotation of the web pages within a user's browsing history. In addition to the access times and URLs it offers an ontological footprint, which was created by an automatic annotation at the time user browsed the page. The combination of semantic footprints with web pages automatically annotated using the concepts constituting the footprints facilitates a conceptual search of previously visited pages. We believe that the abilities to *find* the right page and *make sense* of it easily are two fundamental activities contributing to the wider adoption of Semantic Web technologies.

Attention as opposed to information is now widely acknowledged to be the scarce resource in the Internet age. Consequently, tools that can leverage semantic resources to take some of the burden of the interpretation task from the human reader are going to be of enormous use. We believe that Magpie is a step towards achieving this goal.

Our current effort is focused on deploying the Magpie tools in the climate*prediction*.net project. Using the scheme that was successfully deployed in the SETI@home project, the climate*prediction*.net exploits the idle time on PCs to run multiple versions of the UK Met Office climate model. Running large numbers of perturbed climate models (the project aims to collect 2M users) will overcome uncertainties present in the modeling (and hence prediction) process. During their participation in the project, the users would run climate models on their computers for several months. Magpie will be used for the purposes of interacting with and making sense of highly complex analyses of climate data that will be produced from running a statistical ensemble of perturbed climate models. Magpie will also enable lay members of the public to explore the rich scientific resources that exist in the domain of climatology and climate prediction. Thus, it is hoped that the semantic browsing capabilities of Magpie will serve as an *enabling technology* for the increased public understanding of science.

7 ACKNOWLEDGMENTS

8 REFERENCES

[1] Berners-Lee, T., Hendler, J., and Lassila, O., *The Semantic Web*. Scientific American, 2001. 279(5): p.34-43.

[2] Carr, L., Bechhofer, S., Goble, C.*, et al.* Conceptual Linking: Ontology-based Open Hypermedia. In *Proc. of the 10th Intl. WWW Conf.* 2001. Hong-Kong.

[3] Chakrabarti, S., Srivastava, S., Subramanyam, M.*, et al.* Using Memex to archive and mine community Web browsing experience. In *Proc. of the 9th Intl. WWW Conf.* 2000. The Netherlands.

[4] Dzbor, M., Domingue, J., and Motta, E., Magpie: Towards a Semantic Web Browser. In *Proc. of the 2nd Intl. Semantic Web Conf.* 2003. Florida, USA.

[5] Gruber, T.R., *A Translation approach to Portable Ontology Specifications*. Knowledge Acquisition, 1993. 5(2): p.199-221.

[6] Guarino, N., Masolo, C., and Vetere, G., *OntoSeek: Content-Based Access to the Web*. IEEE Intelligent Systems, 1999. 14(3): p.70-80.

[7] Handschuh, S., Staab, S., and Maedche, A., CREAM - Creating relational metadata with a component-based, ontology driven annotation framework. In *Proc. of the Intl. Semantic Web Working Symposium (SWWS)*. 2001. California, USA.

[8] Kahan, J., Koivunen, M.-R., Prud'Hommeaux, E.*, et al.* Annotea: An Open RDF Infrastructure for Shared Web Annotations. In *Proc. of the 10th Intl. WWW Conf.* 2001. Hong-Kong.

[9] Karger, D.R., Katz, B., Lin, J.*, et al.* Sticky Notes for the Semantic Web. In *Proc. of the Conf. on Intelligent User Interfaces (IUI)*. 2003. Florida, US.

[10] Lieberman, H., Nardi, B.A., and Wright, D., Training Agents to Recognize Text by Example. In *Proc. of the 3rd Intl. Conf. on Autonomous Agents*. 1999. Washington, USA.

[11] Lieberman, H., Fry, C., and Weitzman, L., *Exploring the web with reconnaissance Agents*. Comm. of the ACM, 2001. 44(8): p.69-75.

[12] Lieberman, H. and Wagner, E., End-user Debugging for E-Commerce. In *Proc. of the Conf. on Intelligent User Interfaces (IUI)*. 2003. Florida, US.

[13] McGuinness, D.L. Ontological Issues for Knowledge-Enhanced Search. In *Proc. of Formal Ontology in Information Systems*. 1998.

[14] Ovsiannikov, I.A., Arbib, M.A., and Mcneill, T.H., *Annotation Technology*. International Journal of Human-Computer Studies, 1999. 50(4): p.329-362.

[15] Sadiq, W. and Kumar, S., *Web Service Description (Usage Scenarios)*. 2002, WWW Consortium. (URL: http://www.w3.org/TR/2002/WD-ws-desc-usecases-20020604).

[16] Tauscher, L. and Greenberg, S., *How People Revisit Web Pages: Empirical Findings and Implications for the Design of History Systems*. International Journal of Human Computer Studies, 2001. 47(1): p.97-138.