Open Research Online

The Open University's repository of research publications and other research outputs

Magpie: customizing users' experiences when browsing on the semantic web

Conference or Workshop Item

How to cite:

Dzbor, Martin; Domingue, John and Motta, Enrico (2004). Magpie: customizing users' experiences when browsing on the semantic web. In: International Conference on Intelligent User Interfaces, 13-16 Jan 2004, Island of Madeira, Portugal.

For guidance on citations see [FAQs](#).

© 2004 ACM

Version: Version of Record

Link(s) to article on publisher's website:
<http://dx.doi.org/doi:10.1145/964442.964479>

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's data [policy](#) on reuse of materials please consult the policies page.

oro.open.ac.uk

Magpie: Customizing Users' Experiences when Browsing on the Semantic Web

Martin Dzbor

John Domingue

Enrico Motta

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

Knowledge Media Institute,
The Open University
Milton Keynes, United Kingdom
+44 1908 655014
j.b.domingue@open.ac.uk

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

ABSTRACT

We describe several advanced functionalities of Magpie – a tool that assists users in the interpretation of web resources. Magpie is an extension to Internet Explorer that automatically creates a *semantic layer* for browsed web pages from a user-selected ontology. Semantic layers are annotations of web pages, 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 services* based on patterns of an automatically updated semantic browsing log. The enriched log also supports semantic 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

One behavior that is increasingly becoming more common and widespread in the emerging “wired world” involves the activity of *browsing the Web*. There have been significant advances in research into supporting the task of finding web resources by means of ‘standard’ information retrieval mechanisms or by means of semantically enhanced search [6, 9]. In this paper, we look at this activity from the perspective of supporting the *interpretation* of web pages. Annotation technologies [7, 10] allow users to associate meta-data with web resources, which can then be used to facilitate their interpretation. Annotation is a useful way to support

shared interpretation, but manual annotation poses a serious bottleneck – especially in terms of re-usability and relevance.

The majority of web pages are not semantically annotated, which hinders the vision of the Semantic Web [1]. The vision links web resources and enables browsing using the shared meanings and machine-interpretable data. Magpie¹ is a tool supporting the interpretation of web pages by 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 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 domain.

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*, in which user’s web browsing behavior takes place. Web pages are created within a specific context, and some users 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.

2 BEHAVIOR PATTERNS IN FOCUS

One incentive for this kind of research was summed up by a seminal study of how users browse the web. Tauscher and Greenberg [12] presented the following statistics on the types of behavior users exhibit while engaging in browsing the web:

- 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 see in a current web page and what have they seen previously. Magpie addresses this need by automatically tracking concepts found during a browsing session and storing them in 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.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IUI’04, January 13–16, 2004, Madeira, Funchal, Portugal.

Copyright 2004 ACM 1-58113-815-6/04/0001...\$5.00.

¹ Demo downloadable from <http://kmi.open.ac.uk/projects/magpie>

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 and bottom ones log the people and projects found in the browsing session. The middle one lists (not explicitly mentioned) projects associated with the people found.

A trigger service represents a semantically filtered view of a semantic log customized for a particular user. The purpose of trigger services is to deliver additional information that is relevant to a particular web page from the perspective of a selected ontology. In other words, through trigger services Magpie allows different users to browse the same web resource and yet focus on and interpret it using different terminology.

Another behavioral aspect we take into account is based on the assumption that users select a specific ontology to facilitate a specific viewpoint. First, the ontology enables semantic annotation and visual highlighting of concepts relevant to that viewpoint. Second, using ontological commitments Magpie can associate with each annotated concept a set of *on-demand* semantic services that offer users higher degree of customization of their browsing behavior than any traditional technologies (e.g. such as open hypermedia based on syntactic anchors).

The combination of ontology-specific on-demand and trigger services in the context of web browsing facilitates a new behavior. To a certain extent, it is possible to perceive this emergent activity as *semantic web browsing*. Magpie can be consequently seen as a semantic web browser, and as we argue in [4], adaptability and customization are key premises of *semantic-izing* the existing Web with its vast legacy HTML resources.

3 REASONING ABOUT BEHAVIORS

Let us highlight the key principles of how Magpie reasons about user behaviors through a short scenario. Consider a journalist or market analyst browsing web site of the Knowledge Media Institute (KMi) to gather information about key research projects and technologies developed at KMi. The starting point is the KMi director’s summary of activities he leads, and the analyst draws draw on an existing ontology of academic organizations, which

was populated by mining databases and web resources, and is available to the external users².

3.1 Semantic annotation and highlighting

Fig. 1 shows the journalist’s browser with the concepts of interest highlighted using the Magpie toolbar. Magpie preserves the structure of the page, and highlights the *relevant* concepts upon user’s request. Users may toggle the highlighting of a specific class of entities, which were annotated using an ontology-derived lexicon. The classes are ontology dependent. If the user selects a different ontology, this will generate new set of toolbar buttons (see marker ‘*’ Fig. 1). Thus, the user has a different interface (UI) available to him or her, and this UI can be modified on-the-fly.

All modifications of UI belonging to this category are driven by a single action of the user – the selection of an *interpretative viewpoint*. Since the lexicon is derived from an ontology, the information about which particular top-level classes are displayable in Magpie toolbar can be incorporated into a particular domain model/ontology. This seemingly simple constraint is an important prerequisite of the successful deployment of Magpie (or indeed, of any other semantic tools). In addition to having access to a particular web resource, the user will also need access or subscription to appropriate domain ontology. Web resources provide the content, whilst the ontology serves as a basis for the contextual semantic layers (for details see [4]).

3.2 On-demand semantic services

Conceptually, the purpose of on-demand semantic services is to deliver contextually relevant knowledge. As can be seen in Fig. 1,

² This example uses the AKT reference ontology, which is available at <http://www.aktors.org/publications/ontology/>.

the on-demand services are delivered through a right-click, *contextual menu* of the web browser. The ‘on-demand services’ menu is also context-dependent; 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 lexicon generated from ontology.

The semantic services are defined and published in line with standards of the emerging web services technology [11]. Hence, the appearance of the contextual pop-up menu is dynamic, and will look differently for different classes and different ontologies. The menu shown in the center of Fig. 1 was customized for concepts classified in the selected ontology as *Projects*, and the options displayed were generated by Magpie Services Agent taking into account the availability of published web services associated with a particular category.

Selecting an option in semantic services menu generates a request to the appropriate service provider to perform the reasoning (if applicable). The knowledge-level inference of the service provider is delivered back to the user’s web browser to be annotated and displayed. An example of a response to service ‘*Shares Interests With*’ invoked for project ‘*ScholOnto*’ is visible in the foreground of Fig. 3. It consists of other known projects whose research themes overlap with those of ScholOnto.

This particular form of UI customization is again based on the ontology selected by the user. However, in addition to the user’s choice, an equally important part in the reasoning about the final effect of the user’s right-click is retained by the agent managing semantic web services. The agent checks for service availability but may also ensure the user is authorized to view a particular option, and/or restrict the menu based on preferences/subscriptions.

3.3 Trigger semantic services

Trigger services are a Magpie variant of a popular form of delivering information used by various recommenders or advisers [2, 8]. 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. The pre-condition for *trigger* services is a *history log* of browsing with the recognized entities. The label ‘browsing history’ reflects the fact that a log aggregates findings throughout a 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.

The semantic log is monitored by watchers – each corresponding to a unique trigger service. The appearance of a pattern activates a respective watcher, which in turn sends applicable knowledge to the Magpie hub. This is transmitted as an XML-encoded message that can be displayed in a 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 (hub), and pushes the relevant information to the user.

The user’s only action is to select which triggers s/he is interested in; these are to be activated whenever appropriate messages arrive from a service provider. So far we have implemented simple trigger services in form of *collectors*, example of which is in Fig. 3. This collector lists all technologies authored by people whose names are recognized in a web page; note that the technologies might not be mentioned explicitly. In our scenario, Enrico Motta (co-)authored ‘*OCML*’ and ‘*IRS-2*’. Each of these can be browsed using a similar contextual menu as discussed in section 3.2; only this time it is customized for ontological class ‘*Technology*’.

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.

4 CUSTOMIZABLE ASPECTS

As we mentioned earlier, 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. 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.

Details both conceptual and technical of the different components of the Magpie infrastructure are beyond the scope of this particular paper. They are discussed in detail in earlier publications [4, Domingue, 2004 #54]. For the purposes of this paper, we consider the key issues related to the users’ behavior and appropriate changes to the Magpie UI.

Fig. 3. Magnified view of a user’s interaction with a “*People’s Technologies*” trigger service. The contextual menu is again class-specific; with an added option “*Related URI-s*” serving as an automated semantic bookmark engine.

There are three basic components that can be modified either by the users' deliberate decisions or their (contextual) behavior. The former category comprises the selection of ontology-derived lexicon and subsequent amendment of the Magpie toolbar GUI. On-demand and trigger services are two examples of how a contextually bound action may provide an opportunity for customization. Magpie delivers the contextually bound customization through a dynamically generated set of semantic services (and respective menu options). These menus/services not only respond to the users' actions, but also facilitate subsequent browsing behavior.

The Magpie plug-in together with the standard browser functionality facilitates two different methods of web browsing. The first is syntactic browsing using the `<A HREF=...>` anchors as defined by the document author. The second method uses customized, automatically annotated semantic tags and dynamically generated semantic services. The former accesses a physically linked content, whereas the latter makes available the semantic context. Our interface differentiates between the two methods to emphasize the complementary nature of the two browsing behaviors.

Knowledge that can be delivered through Magpie framework may range from the simple rendering of a database record (e.g. 'Person Details' menu option) to a result of a multi-step inference process in a particular ontology. For example, the 'Shares Research Interests With' menu option invoked for people can lead to a construction of a community of practice for a given individual using data about other people's or institutions' projects or publications.

Similarly, trigger services may be visualized as simple collections of items, but also as semantically enriched browsing history manager. Since watchers tap into a KB 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 an outsider adopting the viewpoint of a specific community to *interpret* and *make sense* of web resources *from the perspective* of that community. Formal membership of a particular community and the utilization of their ontological viewpoints are two different roles Magpie brings together. A trigger service can be subscribed to, thus enabling outsiders to tap into the knowledge of a community of which they are *not formal members*, and see documents in their 'native' context.

5 PRACTICAL DEPLOYMENT

One immediately relevant practical application with a potentially large user base is the aforementioned semantically enriched bookmarking facility. This application of Magpie framework can be formally grounded in Tauscher & Greenberg's recommendation [12] arguing that "*bookmarks should have meaningful representations*". Since Magpie works with domain concepts (e.g. people's names, technologies, etc.) instead of access times and URIs, it is easier to search for a relevant page using rich semantic descriptions. This capability creates ontological footprints for all the pages visited by the user, which essentially summarize a particular web page from the perspective of a given ontology. Further details on this tool are available in [3].

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 sig-

nificant time overheads. The key principle is that the user controls to what extent semantic browsing comes to the fore. The Magpie toolbar enables concept 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. All these features can be seen as UI customizations based on users' behavior.

Our first application of Magpie will take place in a distance education course at The Open University. Students of the course will run a distributed *climateprediction.net* model (similar to the SETI@Home project). Magpie will be used to support the students in interacting with and making sense of highly complex analyses of climate data 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 climate prediction. Thus, Magpie as a semantic browsing paradigm will be an *enabling technology* to increase the public understanding of science.

6 REFERENCES

- [1] Berners-Lee, T., Hendler, J., and Lassila, O., *The Semantic Web*. Scientific American, 2001. **279**(5): p.34-43.
- [2] Chakrabarti, S., Srivastava, S., Subramanyam, M., *et al.* *Using Memex to archive and mine community Web browsing experience*. 9th Intl. WWW Conf. 2000. The Netherlands.
- [3] Domingue, J., Dzbor, M., and Motta, E. *Magpie: Supporting Browsing and Navigation on the Semantic Web*. Conf. on Intelligent User Interfaces (IUI). 2004. Portugal.
- [4] Dzbor, M., Domingue, J., and Motta, E. *Magpie: Towards a Semantic Web Browser*. 2nd Intl. Semantic Web Conf. 2003. Florida, USA.
- [5] Gruber, T.R., *A Translation approach to Portable Ontology Specifications*. Knowledge Acquisition, 1993. **5**(2):p199-221
- [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] Kahan, J., Koivunen, M.-R., Prud'Hommeaux, E., *et al.* *Annotea: An Open RDF Infrastructure for Shared Web Annotations*. 10th Intl. WWW Conf. 2001. Hong-Kong.
- [8] Lieberman, H. and Wagner, E. *End-user Debugging for E-Commerce*. Conf. on Intelligent User Interfaces (IUI). 2003. Florida, US.
- [9] McGuinness, D.L. *Ontological Issues for Knowledge-Enhanced Search*. Proceedings of Formal Ontology in Information Systems. 1998.
- [10] Ovsianikov, I.A., Arbib, M.A., and McNeill, T.H., *Annotation Technology*. International Journal of Human-Computer Studies, 1999. **50**(4): p. 329-362.
- [11] Sadiq, W. and Kumar, S., *Web Service Description (Usage Scenarios)*. URL: <http://www.w3.org/TR/2002/WD-ws-desc-usecases-20020604>. 2002, World Wide Web Consortium.
- [12] 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.