



## **Magpie: User's Manual and Functionality Overview**

by

**Martin Dzbor**

<http://kmi.open.ac.uk/people/dzbor>

2003-2004 © Knowledge Media Institute

## Executive summary

This document describes a tool that assists users with browsing the Web. Browsing activity involves *finding* the right web page and *making sense* of its content. Most effort so-far was focused on finding web pages using ‘standard’ information retrieval mechanisms. Less attention was given to finding a web page visited in the past (browsing history), and even less to the sense making process. Magpie as a suite of tools provides support for the interpretation of web pages and semantically enriched browsing history management. It automatically associates an ontology-based semantic layer to web resources, allowing relevant services to be invoked within a standard web browser. Relevant research papers are (Domingue, Dzbor et al. 2003; Dzbor, Domingue et al. 2003). More information about Magpie is available from the project’s home page<sup>1</sup>.

## What is Magpie?

Magpie is a tool that plugs into a standard web browser – Microsoft Internet Explorer. Its novelty comes from the fact that through Magpie, the web user can access not only the information provided by web page authors, but also a significant amount of background information and knowledge. Typically, it is this ‘background information’ the users are seeking on the web, and sometimes, it takes long time to find appropriate and *relevant* information (Dzbor, Domingue et al. 2003). While there is a lot of information on the web, it becomes increasingly difficult to:

- Make sense of large amount of information,
- Formulate the search engine queries so that they find reliable and relevant information, and
- Filter out reliable and useful information from the vast amount available.

The issue of sense-making may seem to be fairly straightforward for most people who interacted with the web-based resources. We simply ‘know’ whether a particular link proposed by a search engine makes sense in relation to our intention. With Magpie, we suggest a strategy how such a sense-making capability could be built into a web browser. Before discussing the details of how this is accomplished, let us mention what Magpie is and is not:

- Magpie is an *enhancement* to the standard browser. It is not supposed to replace it.
- Magpie works with the *explicit* information that appears on the web pages and in various knowledge repositories. It only associates explicit information at one end with a potentially relevant one at another place. It does not ‘read in between the lines’ or ‘second guess’ the authors’ intentions.
- Magpie provides *means* for a structured and informed exploration of the web resources. It helps with complex reasoning but is not ‘omniscient’.

Consider the following scenario as a typical example of how Magpie can be used to make sense of a particular organization and its internal relationships. Looking at a list

---

<sup>1</sup> URL: <http://kmi.open.ac.uk/projects/magpie>

of publications one member of the organization published on the web, we can guess a few possible relationships hidden on such a page. For example, it shows the member's community of practice, people s/he interacts with, topics and issues s/he specializes in, and similarly. However, what about information such as what specific projects is that member involved in? What generic interests does that member have? What about other collaborators and colleagues in addition to the paper co-authors?

## System requirement for Magpie

Magpie is a plug-in written for Microsoft Internet Explorer, and this fact determines most of the requirements your system must satisfy to use Magpie and thus semantic browsing effectively:

- Operating system: *MS Windows* (Win32 platform) ... this is due to the fact that Magpie is shipped as DLL, which only runs on Win32 systems. We have tested the plug-in on the following operating systems: *Windows 2000 Pro*, *Windows XP Home* and *Pro*.
- Web browser: *MS Internet Explorer*, releases 5.5 or higher. We strongly recommend using currently most up-to-date *Internet Explorer 6*.
- Hard disk space: At least 20MB free space on your C: drive.
- Connection to Internet: This is not needed if you wish to only highlight the ontological entities on off-line web pages. However, Internet link *is essential* if you want to take advantage of ontology-supported semantic browsing and semantic services.

Magpie has been tested on all systems and applications mentioned above. Please, inform us if you find additional operating systems or IE versions where it works.

## Magpie as a Semantic Web browser

The background information mentioned in the previous section is essential for semantic browsing. In case of Magpie, this information is stored by means of ontologies (Gruber 1993), which are formalized explicit representations of domains and/or problem matters. In order to provide the full functionality, the user has to have access to one or more ontologies s/he intends to use to make sense of browsed web pages.

To achieve as high flexibility as possible, Magpie has been designed modularly. This means that only a relatively small part resides locally in the user's web browser. Majority of the functionality is provided remotely by accessing various semantic services, knowledge-based reasoning tools, and data over network. Two standard parts of Magpie needed for a basic operation include:

- Magpie toolbar loadable into a web browser; this is delivered as a dynamically loaded library (DLL).
- Ontology-based lexicon that can reside on the local computer (e.g. user's 'private' representation of a problem domain)<sup>2</sup>.

---

<sup>2</sup> We are currently developing functionality that would enable users to subscribe to shared, public ontological lexicons stored e.g. in XML/RDF or OWL formats.

## ***Installation and basic setup***

We start with the installation and setup of Magpie extension on the user's computer. The current configuration of beta release of Magpie assumes that working folder is "C:\Magpie". While the user may change the settings, this folder will be used for the storage of web page snapshots, as well as ontology storage. The entire functionality of Magpie is implemented as dynamically loaded library (DLL). This library has to be registered with your operating system before anything can be tested. In the case you are installing Magpie from an executable installer (`magpie-setup.exe`), the registration should happen automatically.

However, the DLL registration may fail in the case you have a running instance of Internet Explorer! If this happens, follow the instructions below to find out whether a running instance of Internet Explorer is a problem. Once there are no applications or processes of Internet Explorer running, go to the directory you installed Magpie to (e.g. C:\Magpie), and run/double-click on the file:

```
register-magpie-dll.bat
```

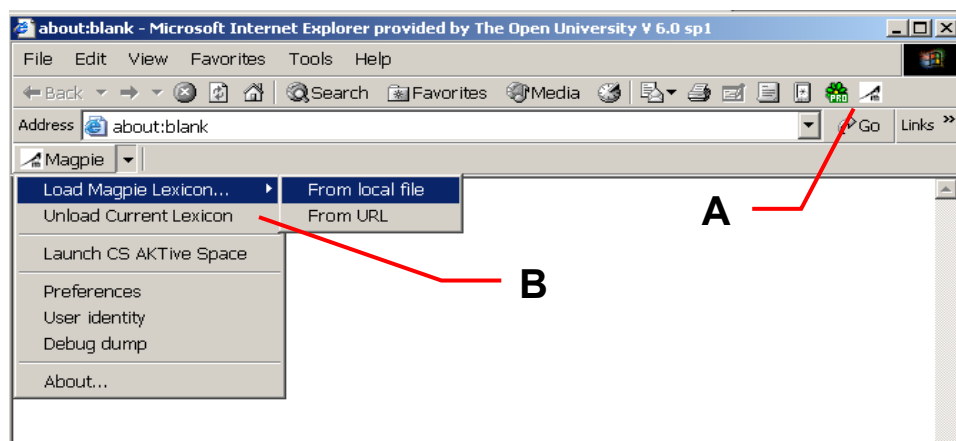
Alternatively, you may do the registration manually. The following command must be executed in the DOS prompt (launch the DOS prompt window using Start → Run... and type "cmd" – a DOS window with a prompt should show):

```
regsvr32 C:\Magpie\Magpie.dll
```

If the DLL registration was accomplished successfully, a window with a respective system message should appear on your screen. In case that Magpie registration with the system was not successful, the most typical reasons are:

- You don't have sufficient privileges to install/register a new application.  
*Resolution:* Please log in as 'Administrator', or ask your system administrator for assistance.
- You have an Internet Explorer application open, and system cannot modify its settings while running.  
*Resolution:* Close all windows of Internet Explorer, and re-run the above command. Since Internet Explorer occasionally stays in the computer memory as a process (with no visible windows showing), it may be necessary to have a look at the running processes, too.
- I see no Internet Explorer windows on my desktop, what about the processes?
- *Resolution:* Right-click on your system bar (the one where 'Start' button is), and in the menu choose option 'Task Manager...'. In the Task Manager window, go to tab labelled 'Processes', find one called 'IEXPLORE.EXE', and highlight it. Next, click on button 'End Process', and confirm it.

These are the most common reasons for failing to operate Magpie extension properly. If you encounter anything else, please, let us know. To test whether Magpie really is in your web browser, launch Internet Explorer. Once it starts up, you should see a little button (usually towards the right-hand end of the Internet Explorer's button bar). The button is depicted as A in Figure 1, and if you click on it an initial Magpie toolbar appears showing a single button labelled "Magpie" and a little arrow (pointer B).



**Figure 1. Magpie toolbar as an integral part of Internet Explorer showing a button for showing/hiding Magpie toolbar (A) and initial appearance of the toolbar with menu (B)**

### ***Preparation – ontology definition/download***


Once you can get as far as described in the previous section, Magpie is likely to be installed and registered correctly. To test its functionality, you will need to load at least one ontological lexicon. A lexicon is a mapping from strings to entities within an ontology (Domingue, Dzbor et al. 2003). Each entity in the ontology is referred by its instance name and class. Instances may have more than one lexical representation. For instance, assume there is an instance known as “BKD-P” in the ontology “AKT-KNOWLEDGE-BASE”. Possible textual phrases that are to be associated with the aforementioned instance include: “Bayesian knowledge discoverer”, “Bayesian knowledge discovery”, and “BKD”. Similarly, instance “JOE-DOE” would translate to “Joe Doe”, “J Doe”, “Doe J”, “Doe Joe”, etc.

Lexicon containing the instances such as those mentioned earlier represents items the Magpie extension would be able to work with and find in the web page. Moreover, thanks to the lexicon and underlying ontology, Magpie is able to offer a suite of semantic services that are dependent on a class of a particular instance. For example, our “BKD-P” is known as a “RESEARCH-PROJECT” specifically and “PROJECT” more generally. Similarly, “JOE-DOE” is classified as a “SENIOR-RESEARCH-FELLOW-IN-ACADEMIA”, or more generally as a “PERSON”. Hence, semantic services related to the “PERSON” class would show when clicked on “Joe Doe” in the web page.

Lexicon comes presently (i.e. in this beta release) in form of two files. The main lexicon has extension “\*.onto”, and contains the actual instances from the ontology. The second file – with the extension “\*.menu” is self-explanatory; it contains the definitions and labelling of the semantic menus associated with the classes from this particular ontological lexicon. Unless you want to create your own ontology, you do not need to worry too much about the syntax of each file. In the future releases, we shall move towards standard web language such as RDF for lexicon definition. Similarly, we will try to offer a web service creating a Magpie-compatible lexicon from ‘any’ arbitrary ontology submitted as an input. Nonetheless, these are future extensions!

In order to test your Magpie, you need to load the lexicon first. Assuming you have downloaded/installed the default Magpie application, follow these steps:

1. Click on Magpie button to show the Magpie toolbar in your web browser. You should see a situation similar to that in Figure 1.
2. Next, click on the little arrow next to the label “Magpie” on the Magpie toolbar (see pointer B in Figure 1). A menu similar to that depicted in the figure should appear.
3. Choose the first option labelled “Load Ontology”; and then “From local file”. A file selection dialog window appears, letting you choose the lexicon (with \*.onto extension).
4. Navigate to the location (folder) where you stored the lexicon; usually it would be found in “C:\Magpie” or “Ontology Store” sub-folder of the main installation directory (mostly “C:\Magpie”).
5. Choose “Org-Intranet-Knowl-Management.onto” for the purposes of interacting with a knowledge management application of Magpie.
6. Clicking on “Open” button Magpie starts loading lexicon. Note that this may take a while (around 2-5 seconds for “Org-Intranet-Knowl-Management.onto” lexicon) depending on the size of the lexicon, number of instances, and of course computer power.
7. If the loading and lexicon initialization is successful, you should see Magpie toolbar changing its appearance and show the top-level classes/categories from a particular ontology. Each class is a button with a different colour label. For details of the screen, see Figure 2.



**Figure 2. Magpie toolbar after loading a lexicon and showing the top-level classes.**

Having come so far means that the Magpie extension is properly set up and ready to be used for browsing. Now you may start browsing the web by typing a URL (web address) in your web browser, or by searching your favourite web engine. We shall look at the actual behaviour during browsing in the next section. However, let us mention two simple configuration changes you can perform at this stage first.

First, you can change a button label for each of the top-level classes/categories by using the first option (“Change Button Text”) in the pull-down menu shown by pointer A in Figure 2. To access the pull-down menu, you click on the little arrow next to the button you want configure, and then choose the option you want.

Similarly, you may want to change the colours used by Magpie to highlight entities of a particular class found in the web page. Currently, all research areas that are known within the loaded ontological lexicon would be highlighted in sky blue. While this is a good colour for a page with white background, it may be inappropriate for say light blue background. In such a case, you may easily change the colour of highlighters. To perform the change, access the pull-down menu for a particular button/category, and choose second option (“Change Highlight Colour”) as shown in Figure 2. A colour picker dialog would show allowing you to choose your preferred shade.

## Using Magpie while browsing

In order to get in grips with Magpie functionality, we suggest starting with the organizational memory lexicon (Org-Intranet-Knowl-Management.onto file). This file provides a lexicon for the KMi reference ontology<sup>3</sup>, which describes the life of a research institute in academia, and would help you identify various people, projects, research areas, and their relationships as they are known in this particular institute. To see Magpie with KMi lexicon in its 'native' environment we suggest the following: <http://kmi.open.ac.uk/people>. Type this URL into your Magpie-enabled browser, and wait until it appears on screen. Once the site is visible, you may load the KMi lexicon if you haven't done so yet (see previous section for how to go about it).

There are essentially two basic ways for a user to interact with Magpie and any web page enrichments facilitated by Magpie extension:

- First, the user may highlight concepts of any particular class/category shown in the Magpie toolbar that happen to appear on the web page.
- Second, for each highlighted item or phrase, the user may invoke a semantic or contextual menu for browsing the value-added information about the instance.

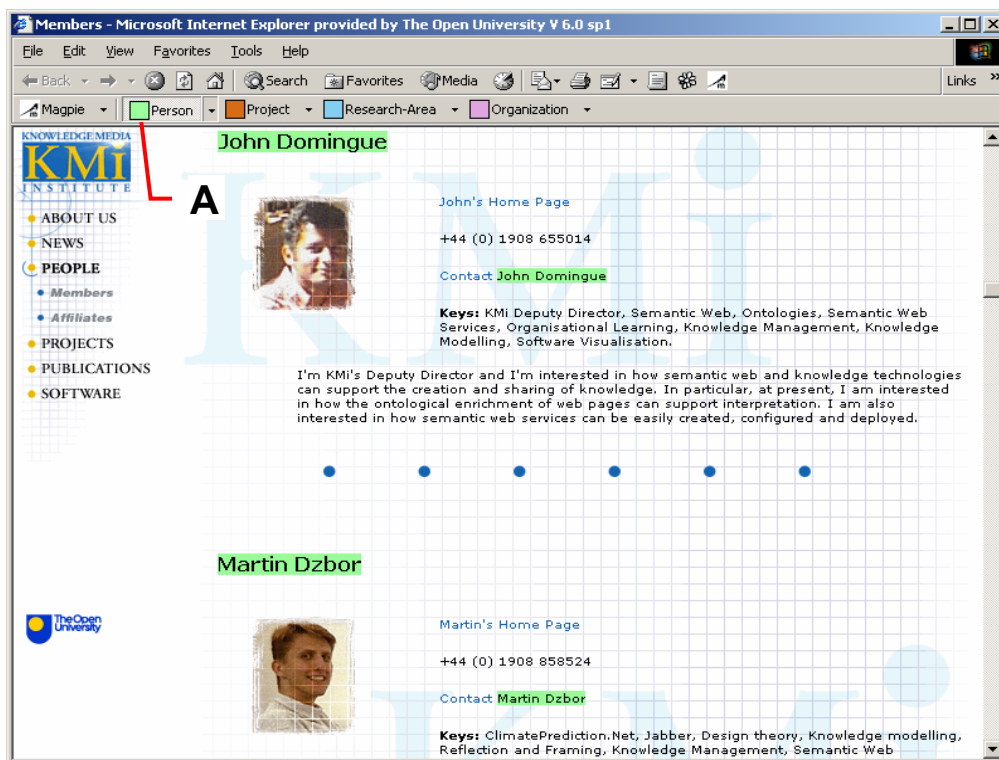


Figure 3. Web page after clicking on button “Person”; i.e. highlighting people in the page who are known in the loaded ontological lexicon (Org-Intranet-Knowl-Management.onto)

## Concept highlighting

As we mentioned above, we explain the highlighting functionality on KMi People's web page. Assuming you already navigated to the appropriate URL<sup>4</sup>, and loaded Org-Intranet-Knowl-Management.onto lexicon, start with clicking on button labelled “Person” on the Magpie toolbar (see pointer A in Figure 3). This button instructs

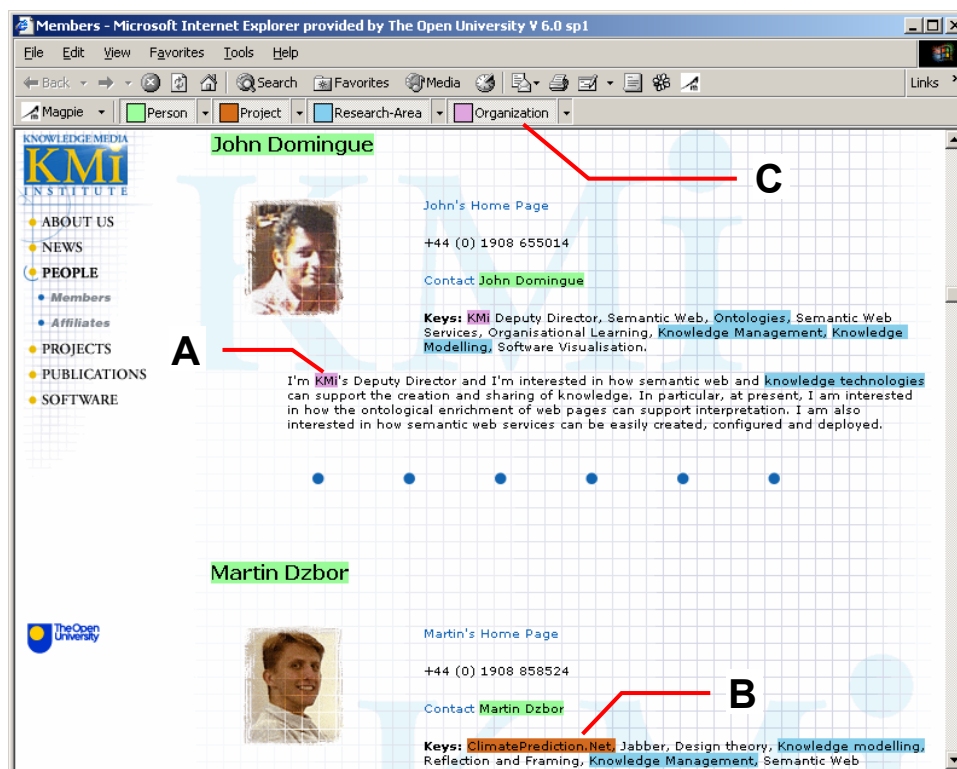
<sup>3</sup> Available in several formats from <http://www.aktors.org/publications/ontology>

<sup>4</sup> <http://kmi.open.ac.uk/people>

Magpie to find any phrase on the page that might be considered as a KMi member. By default, the button highlights concepts with light green colour (see snapshot in Figure 3). As could be expected, Magpie found its authors, and highlighted them in green.

Before we start exploring the semantic menus and associated semantic services, we can highlight other – perhaps less obvious concepts and phrases on the web page. For example, to annotate known research areas, click on the appropriate button. In this case it would one with a sky-blue icon and label “Research-Area” (see pointer B in Figure 3). Similarly, you can repeat the same clicking action on each of Magpie buttons. Each button would highlight instance from the respective category in the web page. The fully annotated web page is shown in Figure 4, and shows KMi projects, research topics and areas of KMi, and key organizational bodies.

More specifically, Magpie correctly recognized phrase “KMi” as an organization (see pointer A in Figure 4), and “ClimatePrediction.Net” as Martin’s project (see pointer B in Figure 4). In addition, several research topics/areas are highlighted – “Ontologies”, “Knowledge Management”, “Knowledge Modelling”, and “Knowledge Technologies” (all with light blue background in Figure 4). The fact that a particular category is actually highlighted is shown on Magpie toolbar by a ‘relief button’ (see compressed buttons pointed by C in Figure 4).



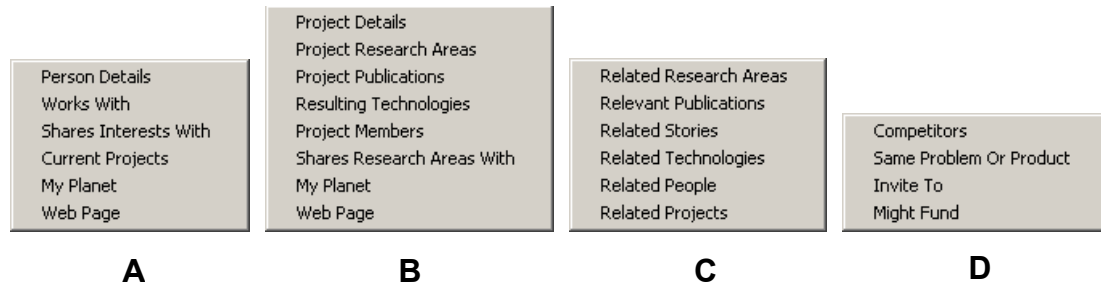
**Figure 4. Fully annotated web page showing all Magpie buttons (categories) turned on and highlighted in appropriate colours: people – light green, projects – chocolate brown, research areas – sky blue, and finally, organizations – light lilac.**

## Contextual, semantic menus

The next after successfully highlighting concepts and phrases in the text, we want to invoke some of the semantic services. The adjective ‘semantic’ means that the actual suite of services is chosen on the class membership of a clicked on phrase. Please,

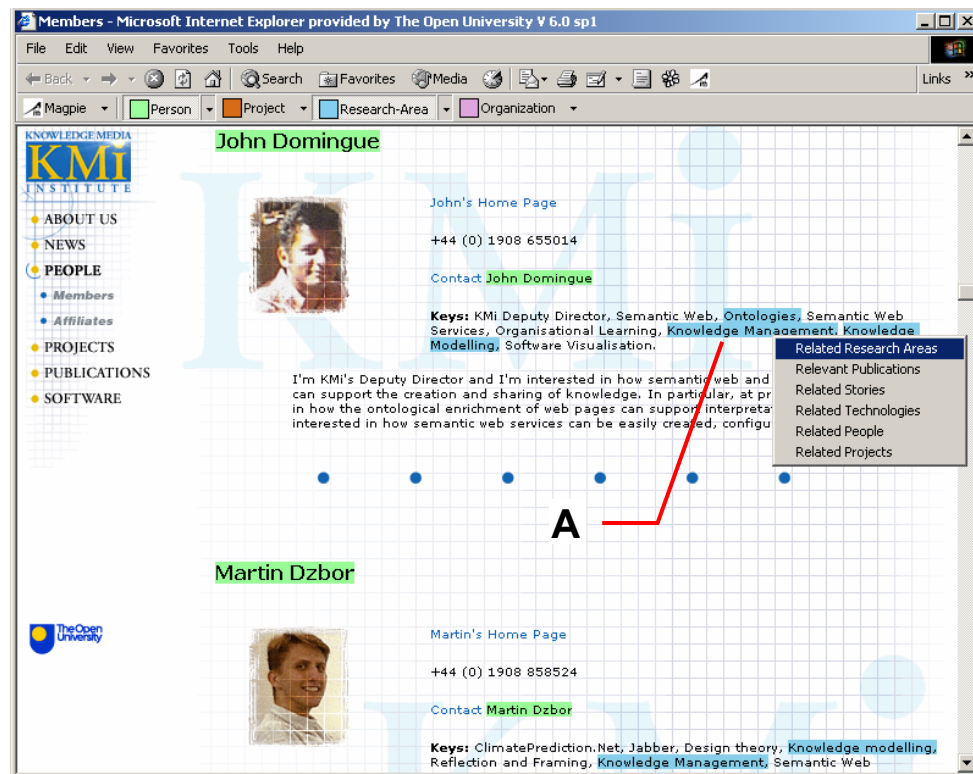


note that only phrases that were highlighted by Magpie would respond with a menu of semantic services! To invoke a semantic menu, you need to move with the mouse pointer over a particular entity, and click with the *right button*. For AKT.onto, four distinct semantic menus could appear depending on the category membership of a particular entity. They are shown in Figure 5.



**Figure 5. Four types of semantic menus defined in AKT.onto lexicon: (A) for the category “Person”, (B) for the category “Project”, (C) for the category “Research-Area”, and (D) for the category “Organization”.**

Now, click on one entity/phrase in the text denoting a person. For example, let us assume we right-click on highlighted phrase “Knowledge Management”. Since this phrase became visible after pressing button “Research-Area”, we expect menu with semantic services specific for this category (option C in Figure 5). The actual appearance of how your browser may look like is shown in Figure 6. As you see, the menu appears exactly where the user clicked with the mouse. It has a usual appearance of contextual menus under given operating system<sup>5</sup>.



**Figure 6. Web page with entities of two categories annotated, and semantic menu invoked over phrase “Knowledge Management”.**

<sup>5</sup> Note that the style may differ if you choose a different scheme or fonts in your Windows settings.

The next step is choosing one of the options in the displayed semantic menu. Let us start with the simple one first: *Which other research areas are related to [Knowledge Management]*? Translated and abbreviated into semantic services labels, this question corresponds to the first option, which is shown by pointer A in Figure 6. In response to requesting this particular semantic information about phrase “Knowledge Management”, a small window pops up resembling the one in Figure 7. According to AKT ontology, areas such as *knowledge acquisition*, *knowledge maintenance*, and so on are relevant in the context of clicking on “Knowledge Management”.

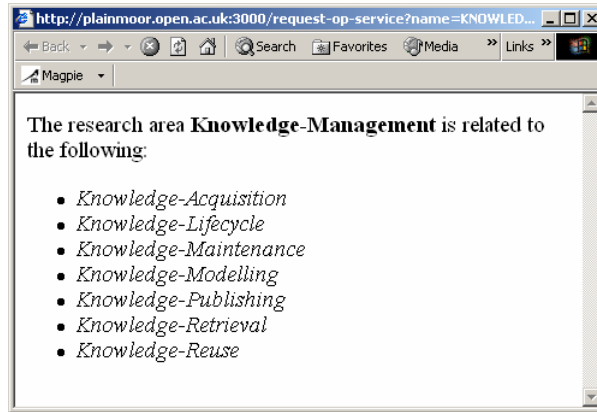


Figure 7. Response from an ontological reasoning engine to the query about other research areas related to “Knowledge Management”.

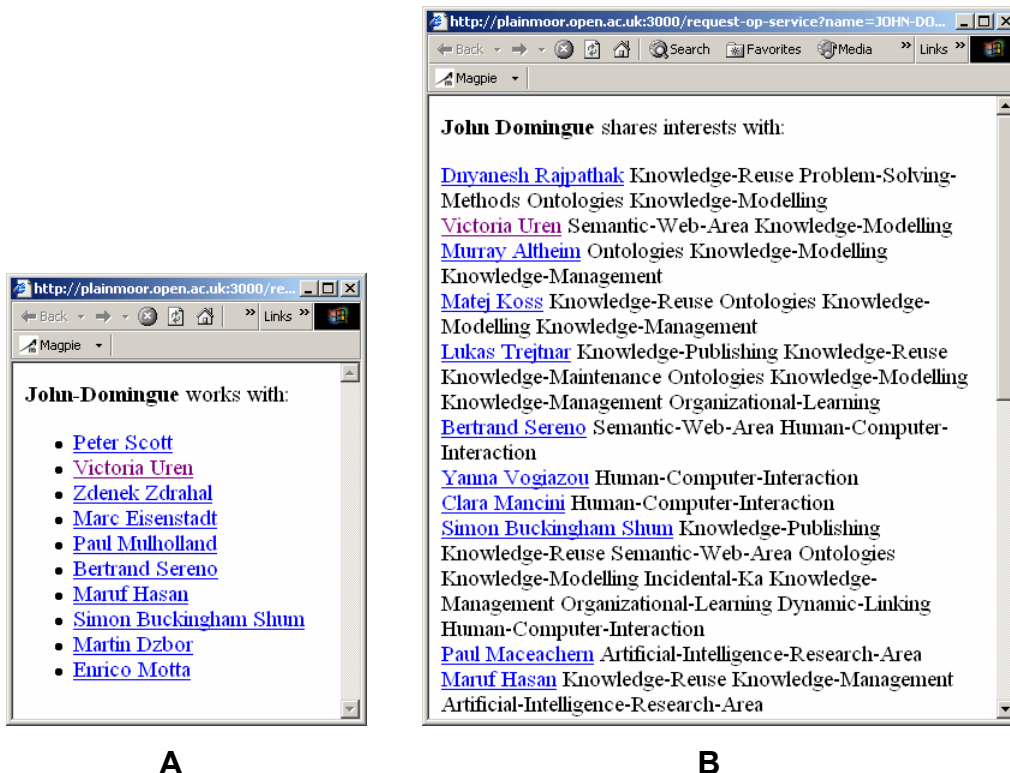


Figure 8. Two ontological responses related to the instance “John Domingue”: (A) Who does John work with? and (B) With whom does John share research interests?

For slightly more sophisticated reasoning, we may now invoke a suite of semantic services contextually relevant to phrase “John Domingue” – from category “Person”. Right-click on any highlighted occurrence of “John Domingue”, and menu shown in Figure 5A should appear. Choose menu option “Works With”. Translated into human-friendly language, this semantic service requests a list of all people with whom John collaborates. The answer is given in Figure 8A.

Finally, let us choose option “Shares Interests With” for the same instance – “John Domingue. Effectively, we are asking who are the other people [in KMi] with interests similar to those of John. This is one of services requiring more complex reasoning. Hence, it may take a few seconds for the answer to show in the pop-up window of your web browser. Nonetheless, the expected answer shall resemble that in Figure 8B.

As we can see from the two figures, some of the responses of the ontological reasoning engine come back as a plain text (see Figure 7). Other response (especially those related to the people or other entities that have their own web sites) come back as URLs. You may therefore click on any of the people listed in Figure 8. Even better, click and drag one of the URLs to your browser window. This trick will leave Magpie toolbar loaded and available to annotate personal pages of KMi members.

## **Related and future functionalities**

In addition to the basic functionality of highlighting ontological entities and providing additional information about them, we are currently working on several other tools. These additional tools shall extend the foundation of current Magpie toolbar, and provide more value added services to the user. We are mentioning them here in order to see Magpie in a particular context rather than an isolated product. The additional tools include:

- Mechanism for trigger or push services, and
- Mechanism for semantic management of browsing history and bookmarks.

### ***Trigger services in Magpie***

All services mentioned so-far in the previous sections are essential user-driven. The user has to do some action, choose a particular concept or phrase, and request what additional information s/he would like to obtain. This request then goes to the ontology-based reasoning engine, which answers the requested query. In addition to the user-requested services we also develop so-called *trigger* or *push* services. These behave differently from the requested ones.

The user subscribes to a particular collector or adviser service, and Magpie sends whatever was found on each web page to a central data store. Different agents monitor this central store, and if a set of data occurs that complies with a particular pattern, respective service is triggered. In other words, using data-driven reasoning, new knowledge is inferred. This chunk of knowledge is then pushed to the user’s desktop, where it appears in a dedicated window.

Consider example in Figure 9. Using standard, user-requested services, we may easily find out more information about a person (e.g. “John Domingue”). However, using trigger services, these projects can be easily listed next to a browsing window (see

pointer A in Figure 9). Moreover, each project pushed by a trigger services is also associated with a URL of the page, which helped trigger a particular response. This enables the user to make semantic bookmarks rather than syntactic. In other words, remembering page that is relevant to a particular concept, phrase or topic, rather than recording a URL without the context it was accessed. From Figure 9 we can see that other projects that are related to John (though not mentioned explicitly on the currently viewed web page) include:

- Lyceum, Meet-O-Matic and KMi Stadium (these are Marc Eisenstadt’s contributions to potentially relevant projects);
- BuddySpace, *climateprediction.net* and Magpie (Martin Dzbor’s contribution);
- Patman, WebOnto, PlanetOnto, etc. (John Domingue’s contribution)

Obviously, there are some overlaps in this list – for instance, *climateprediction.net* or Magpie involve all four individuals listed in box B; thus each of them would contribute with these particular projects as potentially relevant.

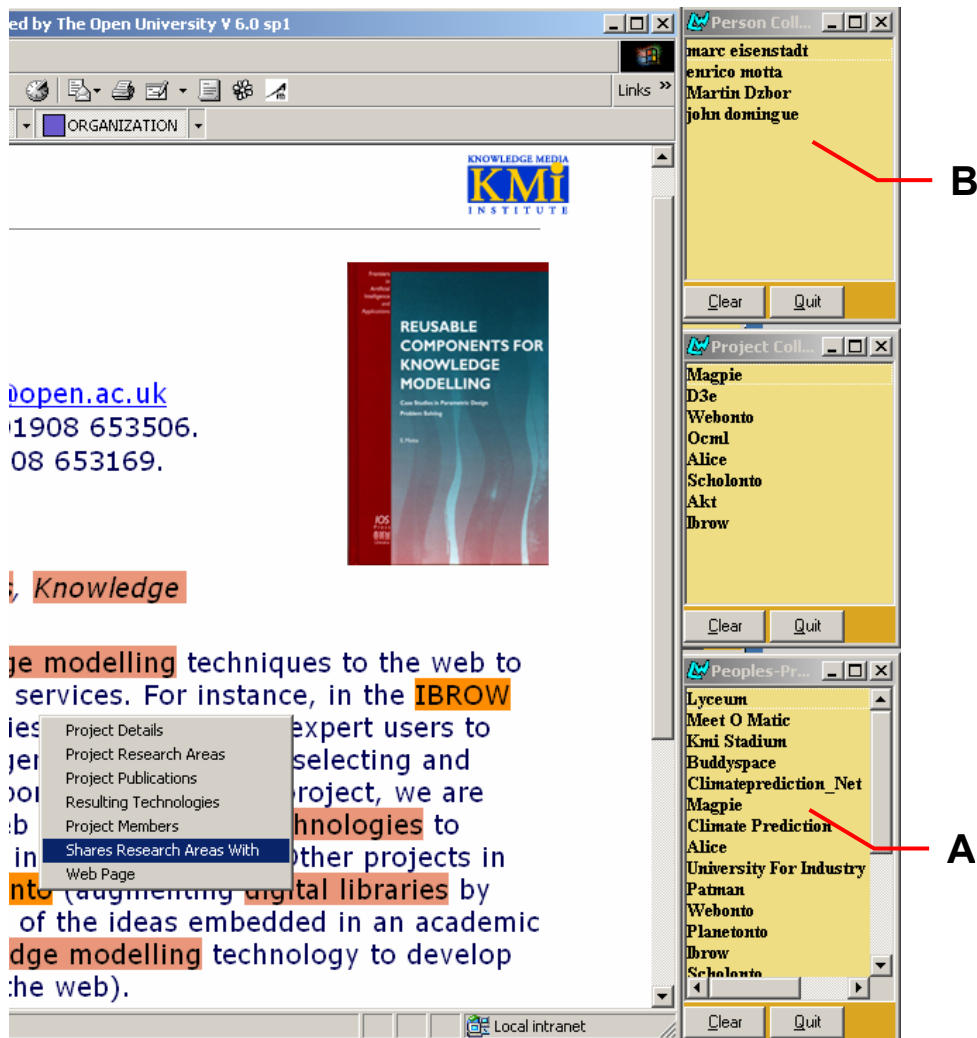


Figure 9. Magpie collectors as an example of trigger services – other projects relevant to the people mentioned on a page are listed.

## Semantic browsing history

One of the purposes of the aforementioned data store and semantic log keeping is to manage users' browsing history better and easier. Since Magpie associates concepts and phrases found in the text with the actual URLs, it enables us to browse the history using semantic means. In contrast to the linear bookmark and history management provided by the standards web browsers, semantic history is more powerful. In addition to seeing the browsing history in a temporal manner, our "Where am I" management tool is able to provide several complementary views:

- Organization by ontological concepts and phrases,
- Organization using iconic representation of web pages rather than URLs,
- Search of large history/bookmark repositories using high-level query language that focuses on web page content rather than label (URL) or access time.

Figure 10 shows an example of the "Where am I" tool in practice. This particular snapshot works with the same ontological lexicon as we mentioned earlier (AKT.onto). Therefore, we are able to make sense of the entities found in the web pages we browsed in the previous sections. We organize history browser around the same categories of entities ("Person", "Project", etc.). Similarly, the tool accepts the entities from the given ontology to form a part of the user's query.

Thus, instead of going blindly through an extensive browsing history, our tool enables the user to ask directly what s/he wants/needs. Assume the user wants to see a page, *which s/he accessed last week, and which was presenting some visualization work by some academic research institution; however, s/he can't remember the names of the people, institution, or project.* It would be futile to ask all the possible combinations of the keywords occurring in the previous sentence.

However, using our ontology-driven history management tool, the same can be done fairly easily. Find me a page that satisfies the following conditions in a conjunction:

- (academic ?X)
- (has-research-interest ?X 'visualization)
- and ?X is a concept known in the ontological lexicon AKT.onto

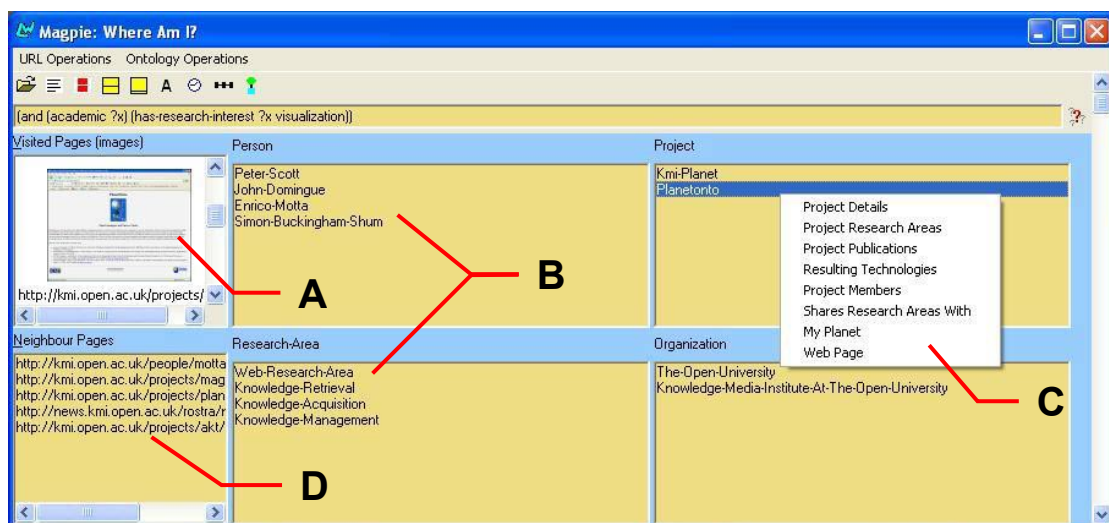


Figure 10. "Where am I" – bookmark and browsing history management tool

As we see in Figure 10, in response to the ambiguous query from above, our tool produced a small list of web pages (pointer A shows their iconic representation). Each of the web pages is annotated using the concepts from the used ontological lexicon (AKT.onto). These concepts constitute what we call semantic footprint of a web page in the space of a given ontology, and in Figure 10 they are in four rectangles (B). In addition, each of the concepts behaves similarly as it would behave on a web page. In other words, it is possible to invoke semantic services for the concepts (C), and investigate the details. Thus, we can arrive at the information we were originally seeking without knowing exactly what it was and where it was. At the same time, the semantic awareness of our tool reduces the total number of *relevant* pages and their physically linked neighbourhood (D).

## References

Domingue, J., M. Dzbor, et al. (2003). Semantic Layering with Magpie. Handbook on Ontologies in Information Systems. S. Staab and R. Studer, Springer Verlag.

Dzbor, M., J. Domingue, et al. (2003). Magpie: Towards a Semantic Web Browser. Proc. of the 2nd Intl. Semantic Web Conf., Florida, USA.

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