

Representations for semantic learning webs: Semantic Web technology in learning support

M. Dzbor, A. Stutt, E. Motta & T. Collins

Knowledge Media Institute, Centre for Research in Computing, The Open University, Walton Hall, Milton Keynes, UK

Abstract

Recent work on applying semantic technologies to learning has concentrated on providing novel means of accessing and making use of learning objects. However, this is unnecessarily limiting: semantic technologies will make it possible to develop a range of educational Semantic Web services, such as interpretation, structure-visualization, support for argumentation, novel forms of content customization, novel mechanisms for aggregating learning material, citation services and so on. In this paper, we outline an initial framework that extends the use of semantic technologies as a means of providing learning services that are owned and created by learning communities.

Keywords

interpretation, learning services, Semantic Web.

Introduction

Polsani (2003) has recently expressed criticism about learning objects by pointing out that 'the term learning object has become the holy grail of content creation and aggregation in the computer-mediated learning field'. Learning objects promise flexible, individualized learning, where courses and training programmes customized for specific users can be assembled efficiently and economically. With the emergence of Semantic Web technologies (Berners-Lee *et al.* 2001), there has also been interest in combining these technologies with learning objects, for example, to enable 'smart' course production (Stojanovic *et al.* 2001) or to broker and match learning materials to the users (Anido *et al.* 2002). Although there is merit in this research, it is our view that there is much more that Semantic Web technologies can do for learning than merely providing ways to locate and compose learning objects. Indeed, as argued later, we find the very notion of learning objects problematic. In spite of perceiving some epistemic shortcomings, our

approach should not be seen as an alternative to the use of learning objects but rather as a complementary, value-adding development.

In this paper, we propose an approach, not tied to (but also not excluding) learning objects, which extends the use of Semantic Web technologies as a means of providing learning services that are owned and created by a learning community. It draws upon the established values of collaborative learning (Morris & Joiner 2002), conversational learning (Laurillard 2002) situated learning (Lave & Wenger 1991) and community-based learning (Preece 2000). The open and distributed nature of the Web offers not only the infrastructure for collaboration and participation as part of a learning community but also provides a wide range of distributed resources from domain experts that are suitable for learners. In our view, semantic learning services can be seen as the means through which members of specific learning communities may actively engage with the resources on the Web at large and inside the learning repositories, in order to further their own knowledge and the knowledge of the communities to which they belong.

Our approach makes better use of the potential of Semantic Web (or ontology-aware) technologies,

Accepted: 21 August 2006

Correspondence: Martin Dzbor, Knowledge Media Institute, The Open University, Walton Hall, Milton Keynes MK7 6AA, UK. E-mail: M.Dzbor@open.ac.uk

offering a range of educational Semantic Web services, such as interpretation or sense-making, structure-visualization, support for argumentation, novel forms of content customization, novel mechanisms for aggregating learning materials, citation services, domain exploration and so on. This proposition is based on one particular view of the Semantic Web, namely that of *a networked space in which concepts acquire meanings by their commitment to a shared ontological model* or possibly, multiple models. Ontology awareness enables the learning application to intertwine domain-specific structures with problem- or task-specific ones, and also to re-use these structures across a wider range of applications.

We thus emphasize the active and engaging learning experiences that Semantic Web services may provide. While we may be critical of learning objects, we do not argue that they cannot provide such experiences. Nor do we suggest that ours is the only approach that situates learning in a community context. However, we do claim that we are providing a more coherent framework, derived largely from advances in ontological engineering on formal representations of knowledge. This framework can be used to harness the power of semantic technologies in capturing, manipulating, navigating and reasoning about the essential conceptual structures of particular domains in a way that is familiar to experts and that makes sense to learners. For instance, the use of an argumentation ontology in constructing representations for use by learners makes it possible for a system not only to retrieve and assemble but to *reason about the parts* of an argument and how they are related.

Of course, when we use Semantic Web technologies, we are, as with any other attempt at formalizing knowledge, in danger of simplifying what is complex. With this in mind, we must clearly focus on the end results of their use: our goal is to help users to create novel learning narratives while making sense of the vast resources now available to them on the Web at large. For this reason, much of the focus of this paper is on interpretation and on supporting the creation of novel connections between different resources that already exist, albeit they may not exist in a form of formalized learning objects; more often, they would be web pages, reports, publications, etc. The approach discussed in the paper is very different in spirit from

the idea of individualized content aggregation out of reusable components.

This paper is organized as follows: the next section presents an overview of some of the current applications of Semantic Web technologies to learning. An alternative approach focusing on semantic services is introduced in the following section. The section after that further explains how this approach can be used in an example-learning scenario. The penultimate section outlines a framework for future learning services. And, the last section concludes the paper with a summary of the proposed approach.

While the central ideas in this framework (*knowledge charts, knowledge navigation and knowledge neighbourhoods*) are not yet fully implemented, we will be drawing on the concrete technologies developed at the Knowledge Media Institute on Semantic Web browsing, new forms of scholarly publishing and support for Semantic Web services, to illustrate the technical feasibility of the approach. In essence, this combines a number of leading-edge technologies into an original framework, which can provide innovative learner-oriented services. This paper also adds a more tangible technological framework to and extends early thoughts about the conceptual basis of this work, which appeared in (Stutt & Motta 2004). Moreover, these technologies have been applied and evaluated in the context of various learning and nonlearning tasks, and there seems to be early evidence of semantics positively affecting users' performance in a given task; the details of this study are given in (Uren *et al.* 2005).

Current applications of Semantic Web technologies to learning

One view of the Semantic Web is as an attempt to complement the existing web by providing ontology-based formal annotations of web resources. An ontology can be seen as an agreed, canonical set of concepts and relations, which captures the essential structure of a particular universe of discourse – e.g. for the purposes of reuse and interoperability. Thus, in archaeology, for example, we might have a range of concepts such as *artefact, layer, feature* and *period* with relations such as *associated-with (layer, period)*. It is therefore possible to create knowledge bases containing instantiations of these concepts and relations,

and to reason about domains in order to produce new knowledge, such as *must-be-later-than* (*object1, period1*). The annotation of web resources in effect turns the classic web into an extensive distributed knowledge base that can be reasoned about, mined for knowledge and applied in knowledge-rich activities and tasks.

As learning can be viewed as the acquisition and continuous development of knowledge in a particular domain, Semantic Web technologies would seem, *prima facie*, to be of interest to the educational community. At the very least, semantic technologies will provide a more focused means of searching for items on the Web. Thus, we might ask a *semantic* (i.e. ontology-aware) search engine to provide us with a solution to the ambiguous query: ‘Get me the web page of the author of the book *The Sceptical Environmentalist?*’ and receive a more focused reply than that provided by conventional, keyword-based search engines.

Of course, semantic technologies can do much more than search, as we shall see in this paper. However, as far as the educational sphere goes, the use of semantic technology for learning has been primarily confined to the annotation of web resources, such as learning objects. Not surprisingly, much recent work has focused on adding detail to metadata schemes so that particular learning goals, object sequences, roles and activities (in short, a pedagogy) can be specified. For example, the work on the Educational Modelling Language (Koper 2000, 2001) and its more recent development as part of the IMS Learning Design Specification¹ follows this trend. Anido *et al.* (2002) provide an overview of the various metadata schemes currently in use. There have also been developments in educational repositories and on peer-to-peer networks for sharing these, such as the Edutella network (Nejdl *et al.* 2002).

While most of this work has been accomplished without the use of explicitly semantic technologies, a natural development of the repositories and networks is the notion of ontology-based brokerages. These match learners with learning materials (Anido *et al.* 2002), and course construction tools (Stojanovic *et al.* 2001), which, in turn, attempt to combine learning objects automatically into ‘courses’ or sequences of

learning objects. One of the few other examples of learning services is the Smart Space for Learning approach using the Elena mediation infrastructure (Simon *et al.* 2003). Work also proceeds on producing customized displays of learning materials that take into account the needs of users (Simon *et al.* 2003).

All this research to some extent depends on learning objects (LOs). However, the use of learning objects is associated with a number of difficulties:

- The field of LOs and, in particular, their annotation is relatively immature, with many competing metadata schemes. There is no single widely accepted standard for the metadata scheme for LOs annotation or mark-up;
- Metadata schemas for learning objects are not powerful enough to capture important scholarly aspects of learning materials (e.g. argument structures);
- There are costs as well as benefits of annotation. There are substantial costs involved in annotating resources to such an extent that it is unlikely that authors will spend much time in providing more than the basic tags (e.g. Dublin Core tags such as author, title and date);
- Automatic aggregation of LOs is still a holy grail.
- Learning objects are mainly applied to support individual learning rather than community-centred learning.

One approach to e-learning, which is not as dependent on learning objects, has been developed at the Royal Institute for Technology (KTH) in Sweden by Naeve and colleagues. Their *Gardens of Knowledge* (Naeve 1997) are learning environments that can be used to explore networks of ideas. They are also developing the idea of the *Conceptual Web* (Naeve *et al.* 2001), as a layer above the Semantic Web intended to make the Semantic Web more accessible to humans using graphical context maps, which include concepts and relations among concepts, as well as access to associated content.

Our approach: learning via knowledge charts, knowledge navigation and knowledge neighbourhoods

Our approach to the provision of learning services has three main components: knowledge charts, knowledge

¹<http://www.imsglobal.org/learningdesign/>

navigation and knowledge neighbourhoods. In the example below, we illustrate the approach using a learning service that provides contextualization via structures representing a scientific controversy. A semantic browser – a web browser that is also capable of using ontologies and ontological relations for navigation – is used to (a) provide access to this service and (b) navigate through the controversy structures as an aid to sense-making. Both these ways of using ontologies and their inherent, shared structures, represent forms of knowledge navigation. The learning service and the controversy structure (we call it a knowledge chart) are both owned and constructed by learning communities focused on particular disciplines, topics or interests, inhabiting bits of the Semantic Web that we refer to as knowledge neighbourhoods.

While learning objects are usually annotated with metadata, these are principally intended to facilitate discovery. For us, effective learning requires the ability to situate a thought in its context within or across disciplines as part of a narrative, a scientific controversy or an analogical argument. These functions are best provided by a range of learning services, such as sense-making, structure-visualization, support for argumentation, novel forms of content customization or novel mechanisms for aggregating learning

materials. As argumentation is a central skill in any academic discipline (Kirschner *et al.* 2003; Andriessen *et al.* 2003), our example service provides sense-making support coupled with a means for visualizing an argument's structure. This structure can be extended by the learners themselves as active participants in the learning community.

Example scenario: the SciControversy learning service

Figure 1 shows a coarse-grained view of how the main components of our approach relate to each other to form the basis for delivering learning services. The following scenario is intended to give a more concrete form to our view of a learning environment based on knowledge charts, knowledge navigation, knowledge neighbourhoods and semantic browsing that we estimate could be realized within months rather than years. It is important to emphasize that, while this scenario is intentionally forward-looking, many of the details are derived from ongoing work on the modelling of argumentation, Semantic Web services and semantic browsing, and are based on existing technologies developed at the Knowledge Media Institute.

In our scenario, we imagine that our learner is reading a document on climate change as part of some

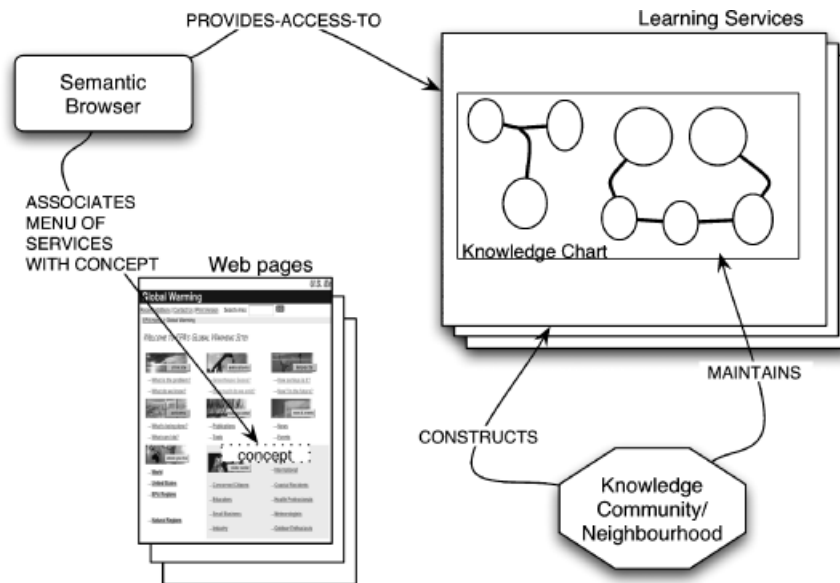


Fig 1 Elements of the framework act together to deliver a learning service: the semantic browser highlights concepts from a domain ontology in a text the learner is reading. Clicking on a highlighted concept provides access to a range of learning services, which include knowledge chart navigators. Knowledge charts are created and maintained by knowledge neighbourhoods.

course on environmental studies. While some mention is made of alternative and competing viewpoints, these are not dealt with fully in the text. As she reads, our semantic plug-in of a standard web browser, called Magpie (see Dzbor *et al.* 2003), automatically identifies portions of the text relevant to a specific course. These concepts are not only identified and visually highlighted but also associated with a number of available semantic educational services. In this case, it can offer a service that displays an interactive view of the scientific controversy about global warming using the SciControversy learning service. Note that Magpie is already fully operational and could support this part of our scenario; some snapshots of the actual tool in use are shown later in ‘Availability of the necessary technologies’.

Figure 2 represents the conceptual model of this controversy in an illustrative form: a real knowledge chart for this controversy could be (admittedly) more complex. The figure shows two levels of knowledge chart. Level 1 (top of Fig 2) shows the causal scientific structure of the argument linking a rise in carbon dioxide (CO₂) to a chain of several climate changes. Level 2 (bottom of Fig 2) shows part of the ongoing scientific controversy about this causal linkage. Imagine now that the learner clicks on the ‘Lomborg Sceptical Environmentalist’ node. This will open up to provide a more detailed version of Lomborg’s argument (Lomborg 2001); e.g. focusing on the fun-

damental assumption that the temperature is indeed going to rise by 2100.

As Lomborg’s (2001) argument about models is based on a view of what statistical models can do, the learner can further explore the feasibility of the argument. For example, they can now opt to follow a link to either a description of statistical models or a deeper view of Lomborg’s argument. And so on. At each point in the debate model, the learner can access the original web resources of which the model is a summary. Of course, any new document or chart could have further knowledge charts associated with it that the learner can pursue in turn.

Moreover, because the navigation is mediated by the semantically enriched browsing framework, the opportunistic association of services with concepts comes at a fairly low cost. In principle, the Magpie framework assumes that the service developer or provider publishes the service to the framework by annotating it using the same ontology that is later used for highlighted and annotating terms in the web pages. When the lexicon is created for the Magpie plug-in, these published services are collated depending on which categories are included in the lexicon. The process of collation draws on a shared ontology and the Magpie framework is thus able to associate existing services to items of interest automatically. New services will opportunistically become available either as (i) the context changes (i.e. as the learner selects

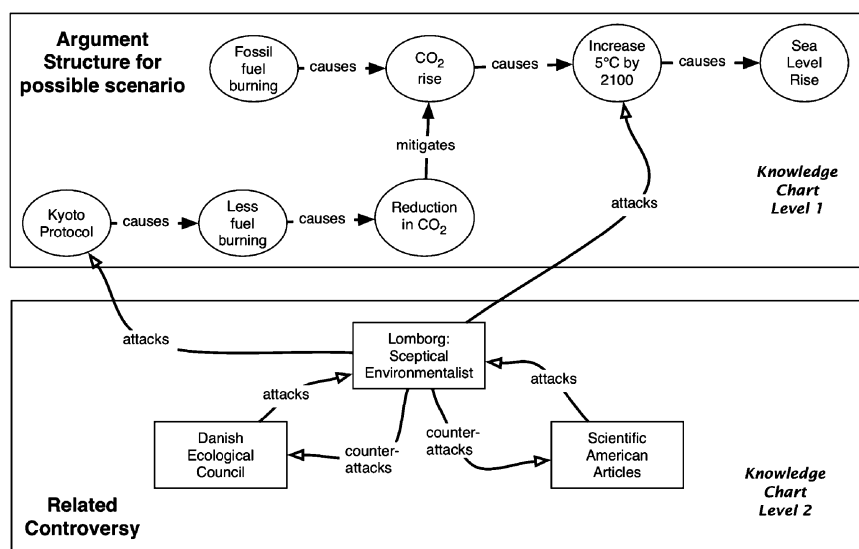


Fig 2 A pair of sample knowledge charts with their argumentation network (based on the argument from Lomborg 2001).

different ontologies) or (ii) the provider decides to implement new functionalities. For instance, as Lomborg becomes the focus of attention, Lomborg-centred services become available (e.g. bibliographic information). More technical details on how this can be realized appear in (Dzbor *et al.* 2004).

What has happened in the above-described scenario was not a mere composition of fixed learning objects. On the contrary, the services have drawn upon the underlying structure of the climate science domain, of the scholarly discussion domain, on the pedagogic domain, and offered the user a set of exploratory paths through a complex issue. They used a real scientific argument to guide the learner through the domain, instead of often simplified relation ‘precedes-follows’ between two learning objects. Furthermore, the services were situated in a particular web page the user decided to interpret; there was no need to switch from web browsing as the primary user interaction modality to interacting with repositories, search engines, etc.

A framework for implementing learning services

The scenario presented in the previous section uses a service that contextualizes a particular chunk of information found in text via representations of controversies. This service provides sense-making support and, because it represents controversies, a means of visualizing an argument structure, which could, in principle, be changed by the learner as part of a learning community. The scenario briefly illustrates the key components of our approach: knowledge charts, their navigation and knowledge neighbourhoods. In what follows, we further discuss the components of our approach and present an overview of the ontologies and tools required to realize its application.

Community knowledge: knowledge charts

The purpose of knowledge charts is to provide pathways through controversies and narratives, as well as other structures such as analogies and expositions of scientific principles. While we have introduced controversies in the above scenario, we consider three types of knowledge structure to be crucial: *debates/controversies*, *narratives* and *analogies/alternatives*. Currently, these three types are developed to different extents in our working prototypes – mostly focusing on narratives and analogies/alternatives.

Debates/controversies are structured exchanges of positions, factual statements, rebuttals, attacks and so on. Scientific controversies may be seen as a special sort of debate in which the exchanges are aimed at testing the validity of particular theoretical positions that are not widely accepted. For instance, Wegener’s theory of continental drift was the topic of a scientific controversy in the last century. It is typical of controversies that they reach some sort of closure. No one *now* doubts that tectonic plates exist; yet, being aware of the controversy and its subsequent resolution further informs our appreciation of the theory. Realizing the learning process as re-visiting specific controversies that led to establishing accepted theories is more engaging and more realistic. The learner does not only learn *about* the theory but also learns *to use* the theory, *to pick* its core and *to apply knowledge*.

Narratives as we view them are (a) the high-level stories or meta-narratives that a discipline tells itself, and (b) lower-level stories that a community adopts as a means of educating new members. For example, archaeology currently sees itself as a specialized, discipline concerned with access to the ‘archaeological record’, rather than the discovery of buried treasures. It tells this meta-narrative as a result of the pioneering endeavours of individuals such as Worsaae in Denmark, who moved the profession away from poorly thought out and experimentally inadequate excavations.

Similarly, the aforementioned climate science was once the domain of a few highly scientific discussions in a handful of specialized journals. Nowadays, there is barely a newspaper or popular magazine that has not made some point about climate change. Thus, climate change becomes less a problem of physical and chemical reactions, and more a political point. What was once available to a few scientists now finds its way to the general public – consider, for example, the response of scientific climate models such as the climate *prediction.net* to the release of the popular ‘Day After Tomorrow’ movie. Instead of rebuffing it as an urban myth, the project offered people an opportunity to verify the probability of such events ever taking place. Thus, a rigorous scientific experiment with Monte Carlo statistical simulations of climate was spun off from a more captivating narrative of a sci-fi movie.

Analogies and alternatives may be taken as a form of argument in which a discipline proceeds by map-

ping its current state of knowledge onto the results or theories in other disciplines and using these to derive new results for itself. For example, social sciences such as psychology often work by using analogies from other scientific disciplines. Cognitive psychology is one such discipline that has derived models of cognition based on an analogy between the working of the brain and the processing mechanisms of computers.

An important characteristic of this particular form of knowledge charts is the inherent divergence they bring into a particular learning interaction. In the aforementioned climate science chart, the learner first explored the causal chain of climate phenomena contributing to climate change. An alternative view on the roles of some of these phenomena (e.g. the assumption of temperature actually rising by a given margin) may completely re-define the chart. In our case, the alternative view could have been accessed through first presenting a controversy, and then allowing the learner to delve deeper into this alternative composition of the facts. These alternatives are particularly hard to realize with traditional learning objects – typically, LO has a certain role, which is too tightly defined to change the principal meaning of the LO.

While knowledge charts may be considered as a new, complex form of learning object, they differ in that they are built using ontologies (i.e. ontologies form the basis for the selection of components and how they are related, while learning objects use ontologies mainly for mark-up). Unlike ontologies themselves, which are mostly schemas, the charts include all the content (e.g. summaries), the annotation and the associated graphical representations. They may draw upon taxonomy, but more often there would be multiple taxonomies the elements of a knowledge chart may fall into. And finally, the charts are used both for navigation (if we view them as hypertext structures) and for interpretation (if we view them as conceptual maps).

Currently, knowledge charts will have to be hand-built in the first instance. However, given the increased sophistication of information extraction technology (Ciravegna *et al.* 2003), we expect that it will become possible to construct these structures semi-automatically or even automatically. In particular, our colleagues Sereno *et al.* (2004) have prototyped a tool, ClaimSpotter, which supports the semi-automatic identification of the key scholarly

claims in a text, and another group developed a tool supporting authors in expressing and recording scholarly claims using explicit language of argumentation ontologies (Li *et al.* 2002). While this work is still at the prototype stage, it shows a glimpse of a future in which semantic agents will be able to index automatically published information according to generic argumentation ontologies. Indeed, the automatic construction of these charts may well be necessary as they are not static, and will need to be updated to reflect the changes in the points of view of an individual, a group or a community.

Another strand of research that some authors of this paper are pursuing is based on re-using shallow associations among entities that can be discovered in a large corpus of text documents (e.g. web pages, text books, reports, etc.) With sufficient corpora, it is possible, first, to base the associations on the expected relations as defined by a specific ontology, and, second, to calculate the statistical significance of the associations actually fitting the expectations of a particular ontology. For further technical details, see e.g. Zhu *et al.* (2005).

It is already possible to build approximations and simple forms of these charts with the tools we have. Arguments or debates can be modelled using nodes to represent the parts of an argument with links between, for example, grounds and claims. These nodes could be clicked on to access web pages containing fuller versions of the grounds and claims. More work needs to be carried out on implementing the navigation of large knowledge charts; i.e. charts containing several hundred nodes.

Knowledge neighbourhoods

A knowledge neighbourhood can be viewed as a location in cyberspace where learners can congregate into groups or larger communities with the goal of acquiring and sharing knowledge about some topic. As with the representations used in knowledge charts, we do not claim that there is anything novel in our use of a community orientation, apart from the name and the tight inter-relations between communities and ontology-based representations of their knowledge.

Virtual communities tend to be organized around their interest in particular topics, which may range from something as broad as particle physics to

something as specific as the use of mythology in Ovid. Members can belong to multiple communities that are composed of a variety of members who take on different roles and enter into a variety of relations with each other. For example, a member may take on the role of a leader in a particular community and have relations with other members, such as setting the agenda for meetings. Roles and relations typically change over time.

These semantic neighbourhoods for learning combine community support with educational semantic web services and are operationalized by providing web portals through which their members interact. Underlying these are mappings from what is known about the communities to the services provided in their portals. The cultural heritage forums developed in the CIPHER project to support the exploration of national and regional heritage illustrate how semantic web services can be deployed in this way (Mulholland *et al.* 2004). Their authors propose, among other topics, useful metaphors for community-customized interaction with domain resources, which are based on semantic relationships. One such metaphor is the *spotlight browsing* (shown in Fig 3), whereby a dynamic neighbourhood of the central term may depend on the focus, interest or past knowledge of the community members.

The Semantic Web can support these special-interest communities through a variety of means, firstly,

by providing ontologies for communities, community structures, roles, relations, spaces, topics, tasks, practices and so on. And, because these neighbourhoods are relatively circumscribed, there will be fewer problems in formulating, negotiating and accepting these ontologies than if we attempted to provide global ontologies. Secondly, the Semantic Web community can provide a range of Semantic Web services that facilitate community building, maintenance and communication as well as the creation of knowledge charts. For example, specific services can assist with community tasks, such as intelligent search for information related to the community's topic of interest or expert discovery (Zhu *et al.* 2005).

Knowledge important to the community will be annotated with ontologies relevant to the community. In our view, it is unlikely that these ontologies will be generic – that, in fact, there will be a single Semantic Web. It is more likely that the Semantic Web, like Ancient Greece or Medieval Italy, will be composed of loosely related knowledge neighbourhoods and specialized, networked ontologies.

Knowledge navigation as semantic browsing

In order for the scenario to become a reality, we also need a system that can perform knowledge navigation. This can be viewed as navigation into and out of knowledge charts, from one region to another within



Fig 3 'Spotlight' area showing categories directly and indirectly related to the chosen concepts (left). Directly related categories are shown in the focus of the torch beam; indirectly related categories are to the side of the torch beam (right). Source: Collins *et al.* (2005).

such charts and from knowledge charts to appropriate external web resources. Because this relies on an understanding of the various ontologies involved, we refer to this as *semantic browsing*. Our semantic browser is already capable of identifying important concepts in the texts or learning objects that a learner is interacting with, by using domain ontologies, without demanding explicit annotation (although this should be used if available and should form part of the knowledge charts that a community produces). Traversals will require either explicitly expressed relations among the charts (or chart components) or depend on inferences made possible by the domain and structure ontologies. For example, the system might have a set of rules that allow linkages from argument nodes (where theories are used to warrant particular claims) to charts that represent the theory.

One example of such navigation through semantic relationships has been used in our semantically enriched application developed to educate learners about the Semantic Web. Among other facilities, this application enables traversals between different subspaces of a large domain about Semantic Web studies. For example, it is desirable that a student acquires the skill to perform lateral or horizontal (Brusilovsky & Rizzo 2002) cognitive transformations. Our application supports the association of a specific concept (say 'hypertext') (i) with other topics conceptually close to the central one (e.g. 'XML' or 'publishing'), or (ii) with researchers publishing on this or a close topic. These associations are navigationally enabled, which means that if one obtained (say) publications about 'hypertext', then navigating through 'XML' node could correspond to narrowing the results so that they satisfy both – the central theme ('hypertext') in the context of the marginal sub-theme ('XML') (Fig 4).

Further examples on how semantic browsing may assist in learning are shown in 'Availability of the necessary technologies', part dedicated to semantic learning tools. Figures 5 and 6 show that thanks to the capability of our semantic browser Magpie (Dzbor *et al.* 2003), one can extend static knowledge content, and easily complement it with dynamically generated learning narratives. This narrative is built from simple chunks (akin to learning objects) that can be linked to form different structures, and thus emphasize different aspects. As shown in Fig 5, one can take a prescriptive narrative 'definition → fact', or in Fig 6 one may

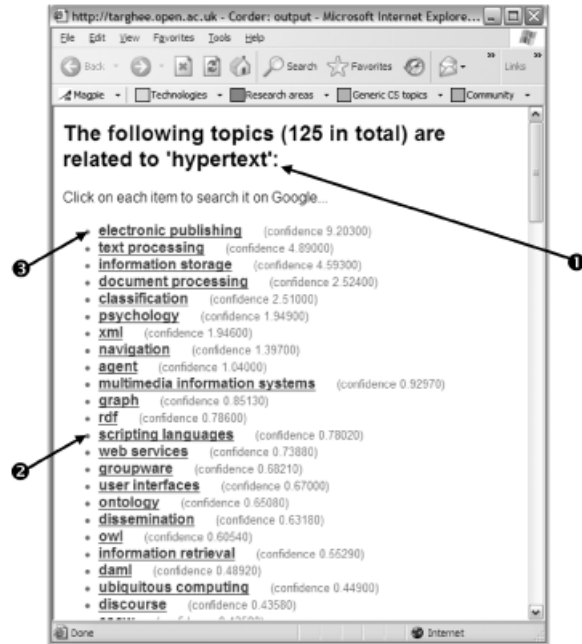


Fig 4 Example of how dynamically harvested lateral linkages among concepts using the Magpie semantic browsing framework may lead to several possible re-formulation of the information retrieval query. Initially, the user requested the term 'hypertext' (see marker 1); marker 2 and 3 show two distinct paths to amend the original query: 2... "interpret concept 'hypertext' in the context of 'scripting languages' and retrieve appropriate publications" or 3... "interpret concept 'hypertext' in the context of 'electronic publishing' and retrieve appropriate publications"

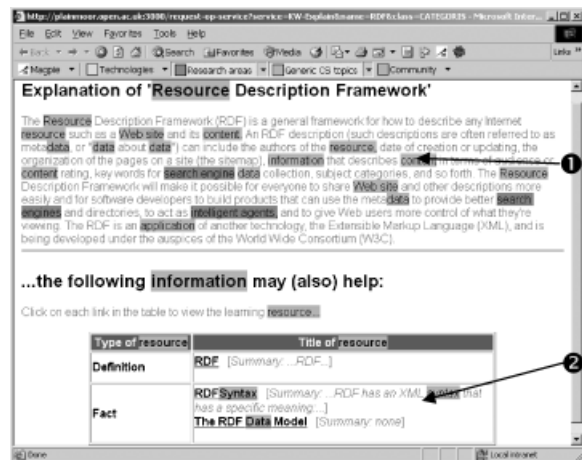


Fig 5 Example of how the Magpie semantic browser allows students to follow alternative learning paths: an existing (static) summary of the explanation of Resource Description Framework (1) combined with an automatically composed sequence of similar learning objects (2).

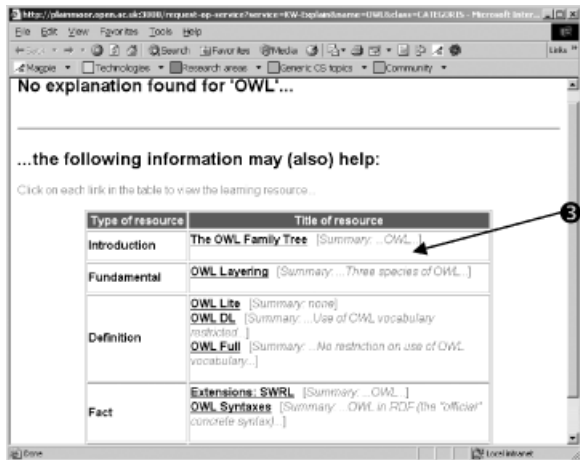


Fig 6 Example of how the Magpie semantic browser allows students to follow alternative learning paths: an automatically composed sequence serving as a fallback for a missing explanation (⊖).

emphasize one specific feature ('fundamental'), before defining anything. The links shown in the two figures are different forms of the aforementioned traversals through a particular knowledge space.

Availability of the necessary technologies

Next, the technologies that are considered critical for realizing the aforementioned conceptual functionalities are briefly reviewed and commented upon – from the viewpoint of their current availability and status quo.

Ontologies

Three main types of ontologies are required to create the learning services we foresee: specifically, ontologies for describing topics or domains, ontologies for creating community-oriented knowledge structures (referred here to as knowledge charts) and ontologies for describing learning communities (here, knowledge neighbourhoods).

Ontologies for various domains abound. However, we are particularly interested in the concepts that are regarded by community members as the most salient and foundational for their domain. In physics, this would be notions such as relativity. In Semantic Web studies, this would be notions such as ontology.

Ontologies for knowledge charts are less plentiful. Knowledge charts include a range of high-level re-

presentations of the most important, or highly controversial, knowledge in a domain. For example, we could have representations of the processes that underlie fossilization in palaeontology or the current controversy about global warming in climate science. Thus, each of the three main types of knowledge charts discussed earlier (debates/controversies, narratives and analogies) requires both a structural and a domain ontology. Debates require structural concepts such as *claim*, *ground*, *evidence* and *theoretical backing*, while stories need *characters*, *events* and *motives*. Our research in the Knowledge Media Institute includes the ScholOnto project, which is actively engaged in producing the ontologies needed for the construction of a semantic, scholarly web of debate (Buckingham Shum *et al.* 2000).

Examples of the support available for representing and capturing narratives can be found in the CIPHER project (Mulholland *et al.* 2002) – for instance, the *spotlight* metaphor in Fig 3 was used in constructing narratives in the domain of browsing cultural digital collections. We thus have already available the basis for the ontologies needed in two main types of knowledge chart.

Finally, we need detailed ontologies for different types of communities and for the different roles that individuals play in them. For example, communities of interest come together to explore a shared interest, such as stamp collectors or real-ale enthusiasts. On the other hand, communities of practice are formed with the purpose of sharing a practice, such as medical professionals learning about the latest techniques for hernia surgeries. Within communities, as we mentioned above, members have different statuses and roles. Some members are concerned with the body of knowledge that the community creates and preserves, while others have a more casual interest. Some members act as teachers as well as contributors to the body of knowledge, while others are solely consumers of knowledge.

The use of ontologies to support communities is a relatively new area of research that promises to be extremely valuable in this educational context.

Semantically enriched navigation tools for learning

A whole range of ontologically informed tools will be needed for the creation of learning services and

knowledge structures, as well as for the nurturing and support of learning communities. Central to these are the semantic browser and the knowledge chart constructor. The semantic browser provides a means for navigating knowledge structures, while the chart constructor provides a means for creating and updating the knowledge structures. The former already exists, and the latter remains at the design stage.

Note that while the semantic browser does not require that the web pages that the user starts from have any explicit mark-up (although it could use this if available), the knowledge charts that it provides access to, as a part of a learning service, will be annotated with the various domain and structural ontologies.

This notion of a semantic browser is based on the work carried out on Magpie (Dzbor *et al.* 2003). Essentially, Magpie works as follows: the community or resource designer provides an ontology for a domain. At the same time, a set of learning services is created by members of the community who have been given the appropriate publishing rights and made available to the Magpie tool. The service providers also specify mappings between concepts and services. These services can be as simple as glossary-lookups or as complex as operational simulations.

For example, when navigating the web using Magpie, a user first selects the lexicalized version of an ontology about climate science. This ontology provides the 'semantic spectacles', allowing the user to look at the Web through the eyes of a climate science expert. During browsing, Magpie automatically associates textual elements in the current web document to concepts in the ontology and highlights these. When the user selects a textual element and right clicks, a menu provides access to the range of associated services. For instance, in Fig 7, the user has selected the highlighted term 'Coriolis effect' and has now access to three services, providing explanations of the concept, facilitating access to additional relevant material from the course and to an online encyclopaedia.

An alternative scenario has been realized for the domain of supporting learners in the domain of Semantic Web studies – as a test case for a dedicated learning application. One service where the opened framework of semantic services and distributed nature of resources became clearly advantageous concerned the discovery of a learning resource explaining a particular concept; e.g. Resource De-

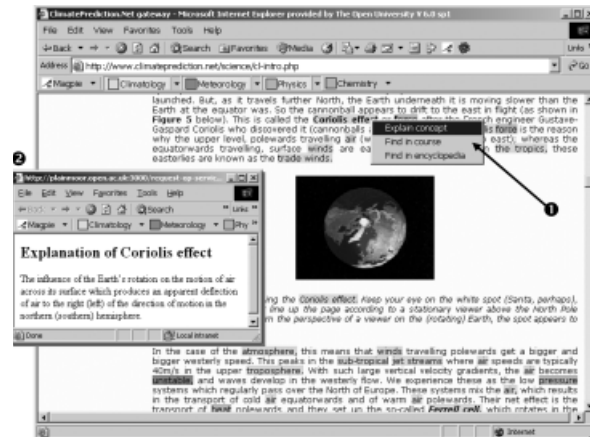


Fig 7 Example of how the Magpie semantic browser can be used to support student learning in climate science. A menu of semantic services (see marker ①) provides a quick and user-friendlier way to composing queries for retrieving appropriate resources and information - e.g. an explanation of a scientific concept as shown by marker ②.

scription Framework (RDF; Lassila & Swick 1999) or Web Ontology Language (OWL; Dean *et al.* 2002). The experimental repository contained an annotated resource about RDF but there was no static resource about OWL. Thus, the user wanting to learn about OWL received a standard 'Error 404' message telling her that a resource has not been found. This is not appropriate in any learning scenario, as part of the learning is the appreciation of alternatives, similarities and the like.

In order to rectify this behaviour, we implemented a fallback method that extracted individual slides from several tutorial and presentations that existed in the repository, and used these slides to compose a learning sequence on the fly. This has been achieved with no a priori annotation of the presentations and without treating the slides as basic learning objects. The composition has been performed based on the statistical relevance of classifying a given slide into one of standard pedagogic roles (e.g. Definition, Example, Conclusion, etc.) Once the core slide has been identified, its content served to construct a chain of appropriate examples, facts and conclusions.

An example of this functionality is shown in Figs 5 and 6. First, Fig 5 shows a combination of a pre-annotated explanation (marker ①) with an automatically composed sequence of extracted slides that may deepen the student's understanding of the concept (marker ②). The deepening is carried out by means of

offering alternative facts and building up one possible narrative about the 'RDF' concept.

On the other hand, Fig 6 shows how the same service may fulfil the role of a fallback in case no pre-annotated resource or LO is actually available. Although the sequence shown by marker ③ in Fig 6 is not a fully fledged explanation, it is able to approximate some elements of a typical explanatory narrative. Thus, a service-based approach enables our application to facilitate the exploration even in those cases where no applicable objects were annotated or known before. This, we believe, is a strong case supporting the more open-ended service-based framework, which compares favourably with many standard LO repositories. While the Magpie framework and several applications built using this framework are already available, it will need to be modified to some extent to enable it to navigate through knowledge charts – mainly because these include graphical as well as textual elements (and do involve annotation). However, the main points about the current system are that:

- Magpie already uses ontologies to construct linkages and relations about objects dynamically;
- It does not require pre-annotation of web pages or other resources;
- It can be extended with arbitrary learning services in an open and distributed fashion, without the learner actually needing to re-install the core plug-in.

The knowledge chart constructor is not currently available. However, it is our proposition that it could be modelled on the principles of Wiki – i.e. interactive web page creation, so that communities can browse and create knowledge charts using a single combined browser and constructor.

Support for knowledge neighbourhoods is also an area for future research. While we already have a range of tools for supporting discussion (Buckingham Shum *et al.* 2000), in our view, this work needs to be extended. Specifically, a range of environments that 'understand' community dynamics (i.e. which are underpinned by community ontologies) need to be developed that incorporate a range of tools supporting the collaborative construction of knowledge structures such as knowledge charts as well as processes such as learning services.

Work on the Alice project (Domingue *et al.* 2003) provides an example of how community ontologies can be developed, and work on the IRS-II/III semantic web service infrastructure (Motta *et al.* 2003) can be used to support the construction and execution of semantic learning web services.

Conclusion: what Semantic Webs might do for learning

While we think that the framework presented in this paper may lead to the development of exciting and innovative learning services, we make no claims that this system currently exists. We have tried, though, to indicate which components are already available and could be readily adapted to feature the functions identified in our conceptual framework of knowledge charts, neighbourhoods and narratives.

There is still some work to be carried out before we can deploy a fully fledged service-oriented semantic learning system. In particular, we need to investigate which types of services are possible and required, which kinds of knowledge chart are useful, how these should be structured and how we can provide navigation through complex charts. These issues are not only technological; equally important is the consideration of cognitive and user interaction aspects to ensure the users do not get lost in this open-ended system of *exploratory navigation*.

We are currently deploying and testing a simpler version of the envisioned system. This is using Magpie to link from parts of a course or required reading to extra material using a variety of services such as glossaries and searches for additional material by the authors of the piece that the learner is reading. Indeed, Figs 5 and 6, with associated discussion, come from the recently implemented prototype.

Similarly, as learning objects are merely building blocks of a potentially successful teaching strategy, semantic mark-up shall not be seen as a goal in its own right. As Brusilovsky and Rizzo (2002) observe, mark-up is a resource-hungry process, and experience with several annotation tools supports them. Mark-up fulfils the role of a partial solution to problems with interoperability, flexibility and integration, but not directly to problems of pedagogy or argumentation. In this paper, semantic mark-up was seen as a means that helps the

users to address a particular set of problems and needs related to learning and to do it faster, more efficiently and more effectively. This paper developed the idea that one of the critical aspects of supporting tasks with a strong element of associative thinking and knowledge creation (which includes learning) is to complement the raw chunks of data with different kinds of relations, associations and their semantic interpretations.

The key aspect of our approach is that it avoids many problems inherent in much of the existing work on reusable learning objects, by focusing on services rather than content, on communities rather than individualized learning and on navigation and sense-making rather than training. In particular, the decentralized and open nature of tools, based on distributed semantic web services, opens the way to the possibility of creating learning services in a collaborative way by members of a community. These technologies may thus play an important role in the future delivery of learning services as learning becomes more ubiquitous, opportunistic and informal.

This view of the Semantic Web and its role in education presents a set of interesting challenges. Instead of asking how we can annotate raw data or learning objects using formal ontologies, this paper explores how we can draw upon a variety of raw data sources to appreciate a potential, semantically meaningful association and use this to enhance user experience. The prototypes of learning support applications that we showed in Figs 5 and 6 do not assume that someone or something annotates the material for a semantic application. They seek to make the bootstraps an integral part of the application and of the user interaction. And in our view, this approach fits better the open and heterogeneous nature of learning strategies as well as learning resources.

Acknowledgements

This work has been supported by the Advanced Knowledge Technologies (AKT), KnowledgeWeb and NeOn projects. AKT is an Interdisciplinary Research Collaboration sponsored by the UK Engineering and Physical Sciences Research Council by grant no. GR/N15764/01. Knowledge Web and NeOn are supported by the Framework VI of the European Commission's IST programme by grants No. FP6-507482 and FP6-027595, respectively.

References

- Andriessen J., Baker M. & Suthers D., eds (2003) *Arguing to Learn: Confronting Cognitions in Computer-Supported Collaborative Learning Environments*. Kluwer Academic Publishers, Dordrecht.
- Anido L.E., Fernández M.J., Caeiro M., Santos J.M., Rodríguez J.S. & Llamas M. (2002) Educational metadata and brokerage for learning resources. *Computers and Education* **38**, 351–374.
- Berners-Lee T., Hendler J. & Lassila O. (2001) The Semantic Web: computers navigating tomorrow's Web will understand more of what's going on – making it more likely that you'll get what you really want. *Scientific American* **279**, 34–43.
- Brusilovsky P. & Rizzo R. (2002) Map-based horizontal navigation in educational hypertext. *Journal of Digital Information* **3**, 156.
- Buckingham Shum S., Motta E. & Domingue J. (2000) ScholOnto: an ontology-based digital library server for research documents and discourse. *International Journal on Digital Libraries* **3**, 237–248.
- Ciravegna F., Dingli A., Wilks Y. & Petrelli D. (2003) Using adaptive information extraction for effective human-centred document annotation. In *Text Mining, Theoretical Aspects and Applications. Lecture Notes in Computer Science* (eds J. Franke, G. Nakhaeizadeh & I. Renz), pp. 9–14. Springer Verlag, Berlin.
- Collins T., Mulholand P. & Zdrahal Z. (2005) Semantic browsing of digital collections. In *Proceedings of the Fourth International Semantic Web Conference* (Galway, Ireland), (eds Y. Gil, E. Motta & R. Benjamins), pp. 127–141. Springer Verlag, Berlin.
- Dean M., Schreiber G., Bechhofer S., van Harmelen F., Hendler J., Horrocks I., McGuinness D.L., Patel-Schneider P.F. & Stein L.A. (2002) *OWL web ontology language reference*. Available at: <http://www.w3.org/TR/owl-ref/>
- Domingue J., Stutt A., Martins M., Tan J., Petursson H. & Motta E. (2003) Supporting online shopping through a combination of ontologies and interface metaphors. *International Journal of Human Computer Studies* **59**, 699–723.
- Dzbor M., Domingue J. & Motta E. (2003) Magpie – towards a semantic web browser. In *Proceedings of the Second International Semantic Web Conference*, Florida, USA, October 2003, (eds D. Fensel, K. Sycara & J. Mylopoulos), pp. 690–705. Springer Verlag, Berlin.
- Dzbor M., Motta E. & Domingue J. (2004) Opening up Magpie via semantic services. In *Proceedings of the Third International Semantic Web Conference*, Japan, (eds S. McIlraith, D. Plexousakis & F. van Harmelen), pp. 635–649. Springer Verlag, Berlin.

- Kirschner P.A., Buckingham Shum S. & Carr C.S., eds (2003) *Visualizing Argumentation: Software Tools for Collaborative and Educational Sense-Making*. Springer-Verlag, London. Information Available at: <http://www.visualizingargumentation.info/>
- Koper R. (2000) *From change to renewal: educational technology foundations of electronic environments*. Educational Technology Expertise Center, Open University of the Netherlands. Available at: <http://eml.ou.nl/introduction/docs/koper-inaugural-address.pdf>
- Koper R. (2001) *Modeling units of study from a pedagogical perspective: the pedagogical meta-model behind EML*. Educational Technology Expertise Center, Open University of the Netherlands. Available at: <http://eml.ou.nl/introduction/docs/ped-metamodel.pdf>
- Lassila O. & Swick R. (1999) *Resource description framework (RDF): model and syntax specification*. World Wide Web Consortium. Available at: <http://www.w3.org/TR/REC-rdf-syntax/>
- Laurillard D. (2002) *Rethinking University Teaching: A Conversational Framework for the Effective Use of Learning Technologies*, 2nd edition. RoutledgeFarmer, London.
- Lave J. & Wenger E. (1991) *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, Cambridge.
- Li G., Uren V., Motta E., Buckingham Shum S. & Domingue J. (2002) ClaiMaker: weaving a semantic web of research papers. In *Proceedings of the First International Semantic Web Conference*, Sardinia.
- Lomborg B. (2001) *The Skeptical Environmentalist*. Cambridge University Press, Cambridge.
- Morris E.J. & Joiner J., eds (2002) JCAL special issue: context, collaboration, computers and learning. *Journal of Computer Assisted Learning* **18**, 52.
- Motta E., Domingue J., Cabral L. & Gaspari M. (2003) IRS-II: a framework and infrastructure for semantic web services. In *Proceedings of the Second International Semantic Web Conference*, Florida, USA (eds D. Fensel, K. Sycara & J. Mylopoulos). Springer Verlag, Berlin.
- Mulholland P., Zdrahal Z. & Collins T. (2002) CIPHER: enabling communities of interest to promote heritage of European regions. *Cultivate Interactive* 8, 15 November 2002. Available at: <http://www.cultivate-int.org/issue8/cipher/>
- Mulholland P., Collins T. & Zdrahal Z. (2004) Story fountain: intelligent support for story research and exploration. In *Proceedings of the ACM Conference on Intelligent User Interfaces*, Madeira, Portugal (eds J. Vanderdonck N.J. Nunes & Ch. Rich). ACM Press.
- Naeve A. (1997) *The garden of knowledge as a knowledge manifold – a conceptual framework for computer supported subjective education*, CID-17, TRITA-NA-D9708. Department of Numerical Analysis and Computer Science, KTH, Stockholm.
- Naeve A., Nilsson M. & Palmér M. (2001) The conceptual web – our research vision. In *Proceedings of the First Semantic Web Working Symposium*. Stanford University, Stanford. (eds I.F. Cruz, S. Decker, J.J. Euzenan & D.L. McGuinness). IOS Press. Available at: <http://kmr.nada.kth.se/papers/>
- Nejdl W., Wolf B., Qu C., Decker S., Sintek M., Naeve A., Nilsson M., Palmér M. & Risch T. (2002) EDUTELLA: A P2P networking infrastructure based on RDF. In *Proceedings of the Eleventh International World Wide Web Conference WWW2002* (eds D. De Roure, A. Iyengar & D. Lassner), Article No. 59. ACM Press.
- Polsani P.R. (2003) Use and abuse of reusable learning objects. *Journal of Digital Information* **3**. Article No. 164. Available at: <http://jodi.ecs.soton.ac.uk/Articles/v03/i04/Polsani/>
- Preece J. (2000) *Online Communities: Supporting Sociability, Designing Usability*. John Wiley & Sons, New York. ISBN: 0471805998.
- Sereno B., Uren V., Buckingham Shum S. & Motta E. (2004) Semantic annotation support in the absence of consensus. In *Proceedings of the First European Semantic Web Symposium*, Crete, Greece (eds Ch. Bussler, J. Davies, D. Fensel & R. Studer), pp. 357–371. Springer Verlag, Berlin.
- Simon B., Miklós Z., Nejdl W., Sintek M. & Salvachua J. (2003) Smart space for learning: a mediation infrastructure for learning services. In *Proceedings of the Twelfth World Wide Web Conference* (eds Y.F. Chen, L. Kovacs & S. Lawrence), Article No. 616. ACM Press. Available at: <http://nm.wu-wien.ac.at/research/publications/b164.pdf>
- Stojanovic L., Staab S. & Studer R. (2001) ELearning based on the semantic web. In *Proceedings of the WebNet2001 – World Conference on the WWW and Internet*, Florida, USA.
- Stutt A. & Motta E. (2004) Semantic learning webs. *Journal of Interactive Media in Education*, Article 10. Special Issue on the Educational Semantic Web (eds T. Anderson & D. Whitelock).
- Uren V., Motta E., Dzbor M. & Cimiano P. (2005) Browsing for information by highlighting automatically generated annotations: a user study and evaluation. In *Proceedings of the Third International Conference on Knowledge Capture (KCAP)*, Banff, Canada (eds P. Clark & G. Schreiber), pp. 75–82. ACM Press.
- Zhu J., Goncalves A.L., Uren V., Motta E. & Pacheco R. (2005) Mining web data for competency management. In *Proceedings of IEEE/WIC/ACM International Conference on Web Intelligence*, Compiègne, France, pp. 94–100. IEEE Press.