

Semantic Webs for Learning: A Vision and Its Realization

Arthur Stutt and Enrico Motta

Knowledge Media Institute,
Open University, UK
{a.stutt,e.motta}@open.ac.uk

Abstract. Augmenting web pages with semantic contents, i.e., building a ‘Semantic Web’, promises a number of benefits for web users in general and learners in particular. Semantic technologies will make it possible to reason about the Web as if it was one extended knowledge base, thus offering increased precision when accessing information and the ability to locate information distributed across different web pages. Moreover, it will become possible to develop 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 material, and so on. In this paper we provide a framework to show how Semantic Browsers which use ontologies to identify important concepts in a document as a means of providing access to associated educational services can be used in conjunction with Knowledge Charts (ontologically permeated representations of a community’s knowledge) in a process we call Knowledge Navigation as an important new resource for learning.

1 Introduction

Augmenting web pages with semantic contents, i.e., building a ‘Semantic Web’, promises a number of benefits for web users in general and learners in particular. Semantic technologies will make it possible to reason about the Web as if it was one extended knowledge base, thus offering increased precision when accessing information and the ability to locate information distributed across different web pages. Moreover, it will become possible to develop a range of additional 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 material, and so on.

In this paper we provide a framework to show how semantic web technology can be harnessed to provide a much richer ‘web experience’ than that currently provided by web browsers and static web pages. In particular the ideas presented in these scenarios will be grounded on some of the work currently being carried out at the Knowledge Media Institute on semantic web browsing and on new forms of scholarly publishing.

Of course, the semantic web, like any other attempt at formalizing knowledge, carries a risk: to simplify what is complex, to impoverish what is rich. For this reason it is important not to lose focus of what the technology should be about: it should support users in making connections, engaging in critical analysis, locating the right knowledge and navigating and making sense of alternative teaching narratives. If used correctly, this technology could provide a quantum leap in the level of support available to students.

1.1 The Semantic Web and Learning

If current research is successful there will be a plethora of e-learning platforms making use of a varied menu of reusable educational material or learning objects. For the learner, the Semantic Web will, in addition, offer rich seams of diverse learning resources over and above the course materials (or learning objects) specified by course designers. The annotation registries which provide access to marked up resources will make it possible for more focused, in some cases ontologically-guided (or semantic) search, to take place. This much is already in development. But we can go much further. Semantic Learning Webs (we believe there will never be a singular Learning Web) depend on four things: annotated educational resources, a means of reasoning about these, a means of retrieving the most suitable and a range of associated services. One important class of tool here is the Semantic Browser which uses ontologies to identify important concepts in a document and to provide access, via services, to relevant material.

1.2 Knowledge Charts, Navigation, Neighbourhoods

Briefly put, Knowledge Neighbourhoods are locations on the Web where communities collaborate to create and use representations of their knowledge - Knowledge Charts. These are browsed - in a process we call Knowledge Navigation - using Semantic Browsers. They are constructed, communally, using Semantic Constructors. Knowledge Charts are thus ontologically permeated representations of a community's knowledge or point of view.

In the rest of this paper we will say more about the notions of Semantic Browsers/Constructors, Knowledge Navigation and Knowledge Neighbourhoods. We will also examine the current state of the technologies needed to produce these, suggest ways in which they can be extended and combined and illustrate how they will act together to provide Semantic Learning Webs.

Before we do so we will say a bit about the pedagogic and practical needs of learners as they seek to interpret, and navigate through, the future Web. We will end with a brief discussion of how, armed with these new semantic tools, learners may be capable of becoming robustly critical thinkers, able not only to move easily through the surfeit of information sources but also to examine and critically assess the varied religious, scientific, economic, ethical and political claims and counter-claims which will find fertile ground on the Web.

2 Learning and Cognition

At a cognitive level students need environments which are congruent with what goes on in learning. We can distil learner needs into three categories: structure, relatedness and interpretation.

Structure. As Laurillard indicates [6], one central component of learning is coming to see structure. As the Web grows this ability will become even more important. Unless the learner can find a way to successfully navigate through and filter out irrelevancy, it will be more or less impossible to make use of the rich resources available on the web. In our view there are two main structures which can be used here and both of them can be aided by the use of ontologies: argumentation/debate, and narrative. Debate here includes the various scientific controversies which arise about notions such as continental drift, GM technology, global warming. These controversies are in themselves multi-dimensional since they often have ethical and economic/political aspects as well as scientific. Narrative here includes the historical account of how ideas change and evolve as well as the ‘stories’ we tell as a means of making sense of something (e.g., the story of the rise and fall of working class politics in the UK).

Relatedness. Part at least of the importance of structure is that it is a means of seeing something (a theory, concept, equation) as a whole. Equally important are the relations which link these to other ideas and theories. Both relate to Laurillard’s discussion [6] of the need for the integration of parts.

Interpretation. The learner needs to be able to take a piece of learning material and situate it in a multidimensional space which includes at least: the social, economic and political context (an obvious example here is the GM debate); its place in the meta-narrative of advancement a science tells itself (e.g., Newton supersedes Copernicus, Quantum Physics finesses Newton); and, its role in an ongoing debate or conversation among academic stakeholders (e.g., the discovery of Archaeopteryx and the more recent Chinese feathered dinosaurs supports the notion that birds are dinosaur descendants).

In what follows we develop a vision of the use of the Semantic Web where:

- Learning is contextualized to specific locations in the Semantic Web (i.e., it is community related rather than generic);
- The structure of pieces of knowledge is given by Knowledge Charts and their Navigation tools;
- The charts represent structures (such as narratives, and arguments) using ontologies and provide access to them using graphical representations;
- Relatedness is given by these charts and by the links they provide to further learning resources;
- Interpretation is facilitated by the contextual knowledge these charts provide.

3 What Current Research on the Semantic Web and eLearning Offers: Getting Away from the Obsession with Learning Objects

Without going into a great deal of detail we can describe current research on learning and the (semantic) web as being centrally concerned with so-called Learning Objects (i.e., separable units of educational material which can be combined and reused in a variety of contexts). Central to their reusability are the descriptions which their designers provide using a variety of metadata schemes. Currently there are a number of standards in this area (see, for example, [1]) but it is likely that this number will be reduced with some standards combining and others being discarded as commercial and other pressures come into play.

Another development has been the growth in educational repositories and peer-to-peer networks for sharing these. One example here is the Edutella network [13].

At the same time as the means of sharing these objects has developed, work has also proceeded on adding detail to the metadata schemes in order that particular learning goals, object sequences, roles and activities (in short, a pedagogy) can be specified (see, for example, work on the Educational Modelling Language [4], [5]).

Most of this work has been accomplished without the use of explicitly semantic technologies. However, a natural development of the repositories and networks is the notion of ontology-based brokerages which match learners with learning materials [1] and course construction tools [15] which attempt to automatically combine learning objects into “courses” or sequences of objects.

More recently we have seen the development of educational semantic web services. An example here is the Smart Space for Learning approach using the Elena mediation infrastructure [14]. The services here range from assessment, to short lectures, courses and degree programmes.

There are two main points here. Firstly, the Semantic Web technologies depend for their success on the viability of the strategy of depending on reusable learning components, on the possibility of capturing the characteristics of these in formal descriptions using metadata schemes, on the widespread acceptance of these schemes and, finally and most importantly, on the likelihood that there will be enough incentive for learning object providers or others to annotate their products with the accepted metadata.

Secondly, while, Nejdil and his colleagues have advanced the notion of Educational Web Services, these are still based largely on learning objects. To this extent, their project may fail if learning objects fail to live up to their promise. However, it seems that their architecture may be generic enough to make use of any properly annotated World Wide Web content. Unlike Nejdil and colleagues, we envisage a form of educational service which can only be provided by the Semantic Web. For example it is possible to foresee a service which automatically re-creates the chain of reasoning used in discovering, say, the cause of the SARS outbreak, using unannotated Web pages.

However, our main problem with the learning object approach is that it seems to entail what we might call a Fordist model of education where the design process is seen as an assembly line with learning objects as the standardized components allowing mass production. This is a view of the Semantic web which limits the ways in which knowledge technologies can be deployed. We should be going beyond the static

web with its Learning Objects, Portals, Pages, and Databases to a flexible, multi-layered, multiple-viewpoint, reasoning-oriented approach.

While much effort has been expended on learning objects and learning object repositories, from the Semantic Web perspective, the ontological commitments of the various contending metadata schemes are limited. In essence the metadata is intended to describe a learning object in sufficient detail for a human or other agent to be able to select it as appropriate in some learning context.

What we lack are any tags which can be used to indicate to a learner how the learning object may be contextualized. That is to say that there are no ontological relations in the learning object description which can indicate how an object should be interpreted, or how it fits into the central debates in the field. Faced with the current state of affairs a learner can successfully navigate the space of possible learning objects but cannot navigate the space composed of the far more important structures of relations which knit topics, concepts, examples and so on into the fabric of the disciplinary field.

For example in Earth/Climate Science there is much controversy about the notion of global warming. While most scientists accept that global warming is a reality and that it is caused by increased anthropogenic CO₂ emissions (e.g., the Intergovernmental Panel on Climate Change) there are others who either dispute the cause, the extent or reality of global warming. For instance Lomborg [8] has cast doubt on the quality of the IPCC models and has suggested that the costs of limiting CO₂ emissions far outweigh the benefits. In turn others have disputed the case Lomborg makes.

4 A Scenario

In the following scenario we explore the possible affordances of a Semantic Learning Web. The scenario is intended to give a more concrete form to our vision of a future learning environment based on Knowledge Charts, Knowledge Navigation, Knowledge Neighbourhoods and Semantic Browsing. It is important to emphasize that, while this scenario is visionary, many of the details are derived from ongoing work on the modelling of argumentation and semantic browsing. See the section on Realizing the Vision below for more details of this work.

We can imagine that our learner is reading a web page/document/learning object on climate change as part of some course on environmental studies. While some mention is made of alternative and competing viewpoints, this is not dealt with fully in the text. As she reads, our semantic system – let's call it SWEL for Semantic Web E-Learning – indicates portions of the text which it can assist with. In this case it can offer a way into the scientific debate/controversy about global warming and/or explain the scientific concepts involved.

The learner opts to access the scientific controversy. SWEL provides a graphical interface to the principal components of this debate. Figure 1 shows two levels of Knowledge Chart. Level 1 shows the structure of an argument linking CO₂ rise to climate change. Level 2 shows part of the ongoing scientific controversy about this linkage. Figures 2 and 3 show Lomborg's argument in slightly more detail.

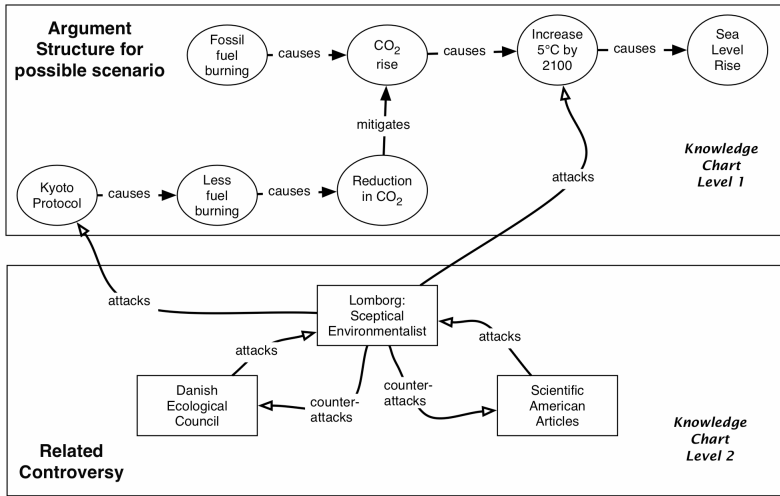


Fig. 1. A web of argumentation related to Global Warming




Fig. 2. Lomborg's arguments

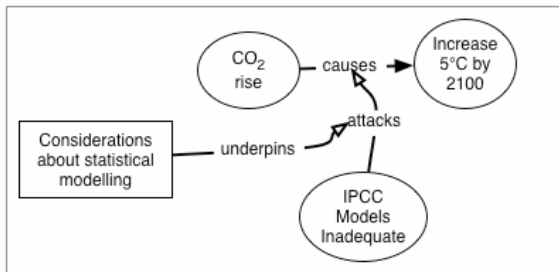


Fig. 3. Lomborg's arguments in more detail

The learner clicks on the “Lomborg Sceptical Environmentalist” node. This opens up to provide a more detailed version of Lomborg’s argument. Basically Lomborg makes two points: (a) that the models used in the IPCC’s calculations about the effects

of CO₂ emissions are inadequate and (b) the cost of reducing the world's emissions by the negligible amount Kyoto would attain far outweighs the benefits.

Since Lomborg's argument about models is based on a view of what statistical models can do, the learner can now opt to follow a link to either a description of statistical models or a deeper view of Lomborg's argument here.

And so on. At each point in the debate model, the learner can access the original documents of which the model is a summary. In order for these to be of maximum use they can be annotated with the argument steps illustrated above using, for example, highlighting in the same colour as the graphical model for claims and grounds and so on. It is up to the learner how much annotation is displayed at any point.

Of course, any new document or Chart could have further Knowledge Charts associated with it which the learner can pursue in turn. For instance, Lomborg might appeal to various economic models in his reasoning. The learner could now decide to follow up links to pages or other Knowledge Charts on classical and ecology-based economic models.

5 Navigation of Knowledge Charts Using Semantic Browsers

As we have said, we envisage a new type or set of types of learning object (we call them Knowledge Charts) whose purpose is to provide pathways through controversies and narratives, and other structures such as analogies and expositions of scientific principles. Indeed, given the prevalence of documents filled with domain content, the system we envisage stands or falls on the existence of these meta-models. A whole new discipline concerned with the production of Knowledge Charts may spring up though it is more likely that they will be crafted by members of a particular learning community. Knowledge Charts (such as the several levels of argumentation and scientific controversy in Figures 1-3) differ from standard learning objects in that: they are built using ontologies, they include both content, annotation and associated graphical representations, they have a taxonomy, and, they are used both for navigation (viewed hypertextually) and interpretation (viewed conceptually). Knowledge Charts will probably be pre-constructed in the first instance. However, given the number of possible meta-learning objects for any course component, it is likely that we will have to find a means of automating their construction.

In order for the scenario to become a reality we also need a system which can perform Knowledge Navigation (i.e., a Semantic Browser). We can rely on our Semantic Browser to identify important concepts in learning objects using domain ontologies without demanding explicit annotation (although this will be used if available). The various traversals (from text to Chart, from Chart to Chart and from Chart to new material) require either explicitly expressed relations among the Charts (or Chart components) or depend on inferencing made possible by the domain and structure ontologies. For example, the system could have a set of rules which allow linkages from argument nodes (where theories are used to warrant particular claims) to Charts which represent the theory.

Note that while the illustrations given above in Figure 1 may indicate that Knowledge Charts are fixed, this is not so. Knowledge Charts reflect the points of view of an

individual, a group or a community and as this knowledge may change it will be necessary for the individual or community to update their Knowledge Charts.

6 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 knowledge about some topic. Knowledge Neighbourhoods are locations for possibly many virtual learning communities. Communities tend to be organized around their interest in particular topics which may range from something as wide-ranging as particle physics to something as specific as the use of mythology in Ovid. Communities are composed of a variety of members who fulfil different roles and enter into a variety of relations with each other. Members can belong to more than one community or group. For example a member may act as a leader in a particular community and have relations such as *sets-the-agenda-for* to other members. These roles and relations may change over time.

These semantic neighbourhoods for learning combine community support with educational semantic web services and are operationalized by providing portals through which their members interact. Underlying this are mappings from what is known about communities to services they provide.

The semantic web can support these communities in a variety of ways. Firstly, by providing ontologies for communities, community structures, roles, relations, spaces, topics, tasks, practices and so on, they can provide an accepted lingua franca for the community as a whole. And, since 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 which ensure that the community is built, maintained and flourishes. Specific services can assist with community tasks, such as intelligent search for topic-related information.

Knowledge important to the community will be annotated with ontologies relevant to the community. As we said in our introduction, it is unlikely that these ontologies will be generic - that in fact there will be a single Semantic Web for Learning. It is more likely that the Semantic Web for Learning, like Ancient Greece or Medieval Italy, will be composed of loosely related Knowledge Neighbourhoods.

7 Realizing the Vision: Technologies Available Now and Needed for SWEL

While the complete framework for Knowledge Navigation does not yet exist, enough of the necessary components are available for us to be confident that it could be available within the next few years.

Domain Ontologies. Both the ontologies and the means of representing them are already available. For instance, work is ongoing to devise a detailed ontology for the climate change domain.

Semantic Web Services for Learning. We need to do empirical research on the kinds of services learners require as well as on the tools we should provide for community construction of these. Currently we are considering services such as the following: Discovery Services, Summarization, Citation-makers, Interpretation, Structure visualization, Argumentation support, Argumentation discovery.

Service Infrastructure. The Knowledge Media Institute is one of the few research bodies involved in work on semantic web services (where an agent or service can reason about appropriate services and retrieve and configure these). The IRS-II infrastructure [9] provides service designers with the means of registering their services both in conventional and semantic registries. It also provides an environment for creating applications configured out of a number of these more primitive services. Currently, the demonstration system can schedule operations abroad for patients with arthritis and who need urgent hip replacement surgery, using primitive services such as hospital availability checkers and currency converters.

Ontologies for Argumentation-Based Knowledge Charts. There is already extensive support for scholarly argumentation in the ScholOnto project [2]. The ScholOnto project is developing an ontology-based digital library server to support scholarly interpretation and discourse with a rich set of relations (such as agrees with, is evidence for and so on) among papers.

Other Types of Knowledge Charts. There is embryonic support for representing and capturing narratives in the CIPHER project [10] which aims to support the exploration of national and regional heritage. This is accomplished by supporting online Cultural Heritage Forums (CHFs) where a community focussed on a specific theme or interest can browse or construct narratives relating to the theme or interest. For example, a CHF supports a community interested in communicating/recording narrative accounts of relevant experiences at Bletchley Park in Milton Keynes, UK where the Enigma encryption machine was deciphered during the Second World War.

Repositories for Structures/Charts. The ScholOnto tool provides the basis for a Knowledge Chart repository.

A Semantic Browser. This is perhaps the most important element in the idea of successful Knowledge Navigation. Our thinking here has been extensively influenced by the ongoing experiments with the Magpie semantic browser [3] in the Knowledge Media Institute. While Magpie can be used as a generic semantic web browser, it originated, in part, as a means of assisting in sense-making for participants in the Climateprediction.net experiment. This experiment, like the Seti@home project, makes use of the distributed computing resources of thousands of home computers, in this case, to run different versions of a climate model. Magpie provides access (via a contextual menu) to complementary sources of knowledge which can be used in contextualizing and interpreting the knowledge in a Web page. This is done by automatically associating a semantic layer to a Web page. This layer depends on one of a number of ontologies which the user can select. When an ontology is selected, the user can also decide which classes are to be highlighted on the Web page. Clicking on a highlighted item (i.e., an instance of a class from the selected ontology) gives access to a number of semantic services. For instance an ontology might contain the class 'Project'. Click-

ing on an instance of this class would provide access to Project details, Research Areas, Publications, Resulting Technologies, Members, Shared Research Areas and project Web Page. In the Climateprediction.net project access is to material which will help to make sense of statistical analyses of complex climate models as well as to the rich literature on climate modelling and climate change.

While Magpie is already capable of semantic browsing (linking from salient concepts to auxiliary material via services) it is currently unable to perform the sorts of complex linkages that would be required for Knowledge Navigation. We need extensions to enable it to link from text to Knowledge Charts, from Knowledge Chart to Knowledge Chart and from Knowledge Chart to alternative web documents. We also need extensions to its concept recognition abilities. Currently it is able to recognize only instances of concepts which are contained in its knowledge base. For example that ScholOnto is a project. There are already plans to make use of Human Language Technology, and in particular, information extraction tools, as a means of automatically identifying instances. This would ensure that textual items such as phrases referring to a project as well as texts in languages other than English would be identified as concept instances.

Note that Magpie does not require annotation of the Web page it operates on since it uses the ontology to pick out concepts to highlight. There is no reason why a full-scale Semantic Browser for Knowledge Charts should not make use of metadata when it is available.

The Semantic Writer/Constructor. This component of our system will assist in the creation of Knowledge Charts. Its main role is in providing the environment for learners and/or their teachers (and communities more generally) to construct, collaboratively, new Knowledge Charts which can represent the point of view of the community, the group of learners, the tutor or the individual learner. As well as assisting in the construction of Knowledge Charts the Semantic Constructor should be able to do all that is necessary to publish these (e.g., notify the registry, or reformat in RDFS). New work just started on Magpie will result in a system which can use information extraction to identify instances. This can also be used to create mark-up within documents thus providing a basis for the Semantic Construction of Knowledge Charts. This means that the extended Magpie will act as both Semantic Browser and Semantic Constructor.

Automated Construction of Knowledge Charts. As we said above, Knowledge Chart construction is the other side of the coin from Semantic Browsing, and an important activity for learning communities. However, some automated assistance in this will be needed. While there is ongoing work on the use of Human Language Technologies such as information extraction to identify concept instances in a text [16] more is needed for the identification of concepts from argumentation, narrative and other structural ontologies. Web resources could, but need not be, annotated with concepts from the various structural ontologies.

8 Related Work

One approach to e-learning, which has much in common with our notion of Knowledge Charts and their Navigation, has been developed at the Royal Institute for Technology (KTH) in Sweden by Naeve and colleagues. Their Gardens of Knowledge [11]

are learning environments which can be used to explore networks of ideas. They are also developing [12] the idea of the Conceptual Web as a layer above the Semantic Web intended to make it more accessible to humans using graphical context maps which include concepts and relations among concepts. Conzilla is a concept browser which allows the user to navigate through a space of context maps to access associated content.

While the idea of graphical representations of domain concepts and their navigation are similar, our approach concentrates more on the elaboration of a typology of high level representations (of arguments, stories and so on) which can be used for navigation and sense-making. In addition, our Knowledge Charts are embedded in the social context of Knowledge Neighbourhoods and our Knowledge Navigation is performed by a tool - the Semantic Browser - which can make the necessary mappings from learning objects to learning objects as well as provide a range of ontologically-directed, community-oriented services such as automated argument construction. We also envisage another tool - a Semantic Writer/Constructor - which can construct or assist in the construction of Knowledge Charts.

9 Conclusion - Learning Webs and Critical Thinkers

The Web presents both challenges (e.g., the danger of information overload) and opportunities (e.g., the wealth of competing viewpoints). The Semantic Web (or Webs) will provide yet more opportunities for learning in the form of greater access to a multiplicity of diverse learning objects. It can also, as we have suggested, provide the means for learners to navigate through the plethora of sources, find help in their interpretation of material by contextualizing it to debates and narratives, and actively enter into these debates or construct these stories as members of living online communities of learners.

If we are to avoid the reductionism inherent in learning objects and related approaches and to support users in making connections, in engaging in critical analysis, in locating the right knowledge and in making sense of pedagogic narratives, we must provide a means of traversing the various links possible from web document to web document by means of Knowledge Charts which express the associated debates, narratives, and so on. By providing this in combination with a means of learner participation in these debates, using Knowledge Neighbourhoods, the learner becomes, not a passive recipient of knowledge, but the sort of critical thinker able to deal with the complexity of the material available in any future knowledge based society.

References

1. Anido, L. E., Fernández, M.J., Caeiro, M., Santos, J.M., Rodríguez, J.S., Llamas, M.: Educational metadata and brokerage for learning resources. *Computers & Education* Volume 38, Issue 4, (2002) 351-374
2. Buckingham Shum, S., Motta, E., Domingue, J.: ScholOnto: An Ontology-Based Digital Library Server for Research Documents and Discourse. *International Journal on Digital Libraries* 3 (3) (2000) 237-248

- 3 Dzbor, M., Domingue, J. B., Motta, E.: Magpie - towards a semantic web browser, In: Proc. Of the 2nd Intl. Semantic Web Conference, October 2003, Florida US (2003)
- 4 Koper, R.: From Change to Renewal: Educational Technology Foundations of Electronic Environments. Educational Technology Expertise Center, Open University of the Netherlands. (2000) Available from: eml.ou.nl/introduction/docs/koper-inaugural-address.pdf
- 5 Koper, R.: Modeling Units of Study from a Pedagogical Perspective: The Pedagogical Meta-Model Behind EML. Educational Technology Expertise Center, Open University of the Netherlands. (2001) Available from: eml.ou.nl/introduction/docs/ped-metamodel.pdf
- 6 Laurillard, D.: Rethinking University Teaching: A Conversational Framework for the Effective Use of Learning Technologies. 2nd Edition. RoutledgeFarmer, London (2002)
- 7 Li, G., Uren, V., Motta, E., Buckingham Shum, S., Domingue, J.: ClaiMaker: Weaving a Semantic Web of Research Papers. 1st Int. Semantic web Conference (ISWC 2002), Sardinia (2002)
- 8 Lomborg, B.: The Skeptical Environmentalist. Cambridge: Cambridge UP, Cambridge (2001)
- 9 Motta, E., Domingue, J., Cabral, L., Gaspari, M.: IRS-II: A Framework and Infrastructure for Semantic Web Services. 2nd International Semantic Web Conference (ISWC2003) 20-23 October 2003, Sundial Resort, Sanibel Island, Florida, USA (2003)
- 10 Mulholland, P., Zdrahal, Z., Collins, T.: CIPHER: Enabling Communities of Interest to Promote Heritage of European Regions. Cultivate Interactive, issue 8, 15 November 2002. (2002) Available from: <http://www.cultivate-int.org/issue8/cipher/>
- 11 Naeve, A.: 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 (2002)
- 12 Naeve, A., Nilsson, M., Palmér, M.: E-Learning in the Semantic Age. Proc. 2nd European Web-based Learning Environments Conference (WBLE 2001), Lund, Sweden (2001)
- 13 Nejdil, W., Wolf, B., Qu, C., Decker, S., Sintek, M., Naeve, A., Nilsson, M., Palmér, M., Risch, T.: EDUTELLA: A P2P Networking Infrastructure Based on RDF. Proceedings Eleventh International World Wide Web Conference WWW2002, May 2002 (2002)
- 14 Simon, B., Miklós, Z., Nejdil, W., Sintek, M., Salvachua, J.: Smart Space for Learning: A Mediation Infrastructure for Learning Services. Proceedings of the 12th World Wide Web Conference, May, 2003 (2003)
- 15 Stojanovic, L., Staab, S and Studer, R.: ELearning Based on the Semantic Web. Proceedings WebNet2001, Orlando, Florida, USA, 2001 (2001)
- 16 Vargas-Vera, M., Motta, E. Domingue, J., Lanzoni, M., Stutt, A., Ciravegna, F.: MnM: Ontology Driven Semi-Automatic and Automatic Support for Semantic Markup. In: ed. Gomez-Perez, A. Proc. 13th International Conference on Knowledge Engineering and Management (EKAW 2002), Springer Verlag, Berlin (2002)