Semantic Learning Webs

Arthur Stutt, Enrico Motta

By 2020, microprocessors will likely be as cheap and plentiful as scrap paper, scattered by the millions into the environment, allowing us to place intelligent systems everywhere. This will change everything around us, including the nature of commerce, the wealth of nations, and the way we communicate, work, play, and live. This will give us smart homes, cars, TVs, jewellery, and money. We will speak to our appliances, and they will speak back. Scientists also expect the Internet will wire up the entire planet and evolve into a membrane consisting of millions of computer networks, creating an "intelligent planet." The Internet will eventually become a "Magic Mirror" that appears in fairy tales, able to speak with the wisdom of the human race.

Michio Kaku, Visions: How Science Will Revolutionize the Twenty-First Century, 1998

If the semantic web needed a symbol, a good one to use would be a Navaho dream-catcher: a small web, lovingly hand-crafted, [easy] to look at, and rumored to catch dreams; but really more of a symbol than a reality.

Pat Hayes, Catching the Dreams, 2002

Commentaries:

All JIME articles are published with links to a commentaries area, which includes part of the article's original review debate. Readers are invited to make use of this resource, and to add their own commentaries. The authors, reviewers, and anyone else who has 'subscribed' to this article via the website will receive e-mail copies of your postings.

1   Peering myopically into the future

Though it is almost impossible to envisage what the Web will be like by the end of the next decade, we can say with some certainty that it will have continued its seemingly unstoppable growth. Given the investment of time and money in the Semantic Web (Berners-Lee et al., 2001), we can also be sure that some form of semanticization will have taken place. This might be superficial - accomplished simply through the addition of loose forms of meta-data mark-up, or more principled – grounded in ontologies and formalised by means of emerging semantic web standards, such as RDF (Lassila and Swick, 1999) or OWL (Mc Guinness and van Harmelen, 2003). Whatever the case, the addition of semantic mark-up will make at least part of the Web more readily accessible to humans and their software agents and will facilitate agent interoperability.

If current research is successful there will also be a plethora of e-learning platforms making use of a varied menu of reusable educational material or learning objects. For the learner, the semanticized Web will, in addition, offer rich seams of diverse learning resources over and above the course materials (or learning objects) specified by course designers. For instance, the annotation registries, which provide access to marked up resources, will enable more focussed, ontologically-guided (or semantic) search. This much is already in development. But we can go much further. Semantic technologies make it possible not only to reason about the Web as if it is one extended knowledge base but also to provide a range of additional educational semantic web services such as summarization, interpretation or sense-making, structure-visualization, and support for argumentation.

1.1   The nature of knowledge on the semantic web

Marshall and Shipman (2003) argue that "The second perspective, the one that finds the Semantic Web in the role of a true Knowledge Navigator, seems out of reach for both theoretic and pragmatic reasons" (p. 65). By a true Knowledge Navigator they mean a "network of knowledge which can be used by personal agents' (p. 59). They object on the grounds that this relies on what they call decontextualized knowledge representation. Hence, they say (p. 65) "at least in the short term — [there may be] many semantic webs rather than The Semantic Web". More precisely there will be a multiplicity of community-based Semantic Learning Webs (SLWs) each with its own, perpetually changing ontologies, knowledge bases, repositories and ways of making sense of the world. That there will be more than one such community need not cause alarm. As writers such as Star and Griesemer (1989) have shown, our intellectual

world is already structured as a set of sometimes overlapping knowledge communities. Semantic technologies can be used to foster and support these communities with ontologies and ontologically informed artefacts providing a means for trans-boundary communication — i.e., using *boundary objects* in Star and Griesemer's sense of the term. We refer to these communities as *Knowledge Neighbourhoods*.

1.2   Semantic Learning Webs and Knowledge Neighbourhoods

Like any other semantic web resources, SLWs will depend on three things: annotated educational resources, a means of reasoning about these, and a range of associated services. Note that since one possible service is the automated identification of suitable material in un-annotated documents – for instance by using *information extraction* techniques customised for the semantic web (Vargas-Vera et al., 2002; Ciravegna et al., 2002), annotated materials (or learning objects) may not be strictly necessary. One important class of tools here are *Semantic Browsers*, such as *Magpie* (Dzbor et al., 2003), which use ontologies to identity important concepts in a document and provide access to relevant material. We use the term *Knowledge Charts* for the ontologically permeated representations of a community's knowledge or point of view produced and used by these tools. They are an important new resource for learning. We refer to the active processing of them as *Knowledge Navigation*.

Our notion of Knowledge Navigation is not intended as a 'true' Knowledge Navigator in Marshall and Shipman's sense. Knowledge Charts are much more context specific representations and we use Knowledge Navigation only for the process of traversing these (localized) representations.

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. 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 SLWs. 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 then say something about the pedagogic

uses of argumentation and how it lends itself to support from semantic technologies. 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 find fertile ground on the Web.

The aim of this paper, then, is not to predict or warn but to suggest how some current developments might be used to produce learning environments which will excite rather than control learners, and which will make them into responsible but critical members of civil society. That's not to say that one should be unduly optimistic. The trend over the past few years has been towards a view of learning as skill acquisition. This is even truer of online learning than traditional learning. Many companies view e-learning as a part of their knowledge management systems. It performs the role of transmitting the knowledge necessary to the company's survival or success from its knowledge repositories to the members of its work force. Sadly, even governments tend to look at learning more and more in terms of information transfer and work-related skill acquisition. As Brown and Duguid (2002) say: "The idea of learning as the steady supply of facts or information, though parodied by Dickens 150 years ago [in Hard Times], still prevails today. Each generation has its own fight against images of learners as wax to be molded, pitchers to be filled, and slates to be written on" (p. 135). Or, we might add, passive recipients of managed knowledge. We can and should do better, as there is a great thirst for learning out there. We can only hope that we don't stifle it with our learning environments, semantic or otherwise.

2   Problems and opportunities for learners

Tempting as it is for technologists, let's start not from what the technology is likely to give us but from what the learner needs. We can do this at many different levels. Pragmatically, what learners need are things like time to study, a quiet place to study, tips on how to pass examinations, peer to peer and learner to tutor communication, someone to talk to when problems arise, and help with reading and interpreting documents (taken as including written, graphical, audio and video sources).

At a more abstract level what problems will the learner face? We note the following: information overload, and information authentication. Given the amount of information available (as well as the possibilities inherent in the technology) learners could also benefit from increased customization or personalization.

Information overload. Most of the problems which come readily to mind are to do with finding a path though the stupendous amount of information currently available not only on the Web but in books, newspapers, television and films – not to mention the input from texting buddies (see Nelson, 1994, for one discussion of this phenomenon). Google currently indexes over three *billion* web pages. Even 20 years ago the student learnt mainly from written material in the form of journal material or textbooks, accessed as photocopies or by visiting a library. Now, much of the journal material is available on the World Wide Web and can be accessed online (providing that the student's institute has a subscription). Even books, and especially expensive monographs, are to be found online. Apart from these digitized versions of previously available sources, the learner is now faced with a bewildering barrage of information intermediaries. It is too easy to overstate this however. Learners have always been faced, at least from the invention of the printing press, with a large array of material. Indeed the ability to sift and evaluate information is one of the skills which a student must acquire. Nonetheless, it is this perceived mass of information which has led to many of today's most visited sites — search engines such as Google — and motivated some of the basic Semantic Web tools — semantic or ontology-based search. As an example of the latter, the OntoBroker system (Decker et al., 1999) provides an ontology-based crawling and answering service.

Information authentication. As well as more conventional indexing facilities provided by librarians and archivists, there now exist web sites where information is collated, indexed and sometimes stored by a variety of actors some of whom are informal digital archivists but with similar standards of objectivity. Others, however, have an axe to grind. Most alarmingly sites exist which present themselves as impartial research conduits when in fact they are funded by commercial and other interests. For example, the number of sites surrounding the issue of genetically modified organisms is immense and not a few may be indirectly or directly funded by companies with large stakes in the future of the technology.

Customization of course/information discovery. Given the amount of information available, the problem of matching learner to material, which is relevant to his or her needs at a particular point in time, becomes more and more acute.

Informal/formal support. A significant factor in any learning context is access to the support networks that learners need in order to be able make sense of the material which they can all too readily access. In the traditional classroom or lecture hall this is usually a human being appointed by the institute for the purpose of fielding questions. In other situations this might take the form of a tutor or mentor (who may

only be slightly more experienced than the learner). This *apprenticeship* mode of learning extends beyond trades such as plumbing and carpentry into seemingly abstract 'trades' such as mathematics, as Brown, Collins and Duguid (1989) have shown. Of greatest importance however are the informal, face to face or online, groups of learners who master a subject by asking questions however trivial, by offering each other differing interpretations of more formal material, by testing these and by negotiating a consensus interpretation – in other words by *constructing* a piece of knowledge. We say more about constructivism in Stutt (1997). These groups, usually of a few students in face-to-face situations (including more formal seminars and tutorials), may become much larger in the online world. Here, informal groups or *virtual communities of learning* (see Stutt et al., 2002) may number hundreds.

3   Learning and cognition

While the above represent the practical problems and needs of learners, it is possible to give an even more abstract account of learner needs which we can use to guide our thoughts about future learning environments. At a more cognitive level students need environments which are congruent with what goes on in learning. From what we have said already we can distinguish between three types of learner needs: for *structure*, *relatedness* and *interpretation*. These correlate more or less with the first two items in Laurillard's (2002) characterization of learning as: *apprehending structure, integrating parts, acting on the world, using feedback, and reflecting on goals.*

3.1   Structure

As Laurillard indicates, 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 three main structures which can be used here and all of them can be aided by the use of ontologies: *argumentation/debate, narrative* and *analogy.*

Debate 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 narrative 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).

Analogies are an important factor in how we make sense of something. Learners can better understand something (e.g., the structure of the atom) if they can relate it to something already well understood (e.g., the planetary system). Researchers, too, often make use of analogies in making a case for something. For example, interpretations of the Anasazi culture in the American South-West often rely on analogies with practices among modern Pueblo peoples.

We recognize that in this paper we are mainly dealing with a particular kind of learning - theoretical knowledge acquisition - in which structure and interpretation are important. In other kinds of learning where skills and competence are more important (e.g., learning a language) other structures are also needed. Note, however, that at least argumentation (especially when the learner participates) has a skill-based component.

While we foreground debate, narrative and analogy, other ontologies, for example, for the set of fundamental theories about a domain are also important - but less easy to capture. For example, cosmology appeals to a limited set of principles in it interpretation of events in the Universe. In other cases we may need causal models or simulations. This is true of disciplines such as ecology where it is important to see the how a model of an ecosystem changes over time.

3.2   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 recognition of the need for the integration of parts.

3.3   Interpretation

The learner needs to be able to take a segment of learning material and situate it in a multidimensional space which includes at least: the scholarly, 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 and articulation of plate tectonics served to refute the arguments of those who contested Wegener's claims about continental drift; the discovery of Archaeopteryx and the more recent Chinese feathered dinosaurs supports the notion that birds are dinosaur offspring).

3.4   Summary: how the semantic web can help addressing learning needs

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

- Learning is contextualized to separable 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, arguments and analogies) using ontologies and provide access to them using graphical representations;

- Relatedness is given by these objects and by the links they provide to further learning resources;

- Interpretation is facilitated by the contextual knowledge these objects provide.

4   The importance of argument

While it is likely that many different kinds of structure (and structuring tool) will be deployed, we will concentrate in this article on those based on argumentation and scientific controversy. This is for a variety of reasons.

a. Our exposition will be more focussed thereby.

b. Argumentation is already the subject of much research as the basis for learning. For instance Kirschner et al. (2002) and Andriessen et al. (2003) both contain extensive collections of papers on the important role of argumentation (and its visualization) in learning. Indeed an attention to critical thinking (both the analysis and construction of arguments) has long been an important component in a well rounded education, as books such as that by Fisher (1988) indicate.

c. Our research in the Knowledge Media Institute has resulted in a variety of argument-based tools not least the D3E discussion tool (Sumner and Buckingham Shum, 1998), which provides the discussion spaces for this and other issues of JIME!

More importantly, the ScholOnto project – see below for more details, is actively engaged in producing the ontologies and tools needed for the construction of a semantic scholarly web (Buckingham Shum et al., 2000).

d. Finally, if as is our view, learning should be clearly distinguished from training, then one means of doing so is by ensuring that learners are provided with the tools for accessing the claims and counter-claims of a range of voices which make up a discipline's conversation with itself as well as the means for making a contribution to this ongoing conversation. Information can't be understood unless its historical context and the dialectic which produced it is understood. Thus the learning process is as important as the product.

5   What exactly is the Semantic Web?

We might as well ask what the World Wide Web is. For the explorer it is a means of communication with home-base, for the CEO a way to sell more product, for the teacher a new way to teach. Abstract definitions in terms of computer networks and hypertext languages simply do not capture the range of meanings we can apply to the phrase "World Wide Web". Similarly, we can define the Semantic Web abstractly as an extended World Wide Web where content bears it own description so that applications as well as humans can make sense and use of it. We can talk about *ontologies* as systems of concepts, properties and relations which can be used in these descriptions, about the *notations* or *languages* for representing these, about *annotation* as the means of adding descriptions to web pages, about *populated ontologies* or *knowledge bases* which capture the knowledge contained in web pages, of *agents* as computer programs which can reason about these descriptions in order to carry out specific tasks, about *services* which coordinate multiple agents and other services as a means of satisfying human needs. But this only gives us the technical infrastructure. As we can clearly see from Berners-Lee et al. (2001), the Semantic Web is an aspiration, a vision of what might be done, a dream even of a future web which becomes more amenable to human needs by relinquishing much of the hard work of publishing, locating and retrieving web content to automatic processes. Like any other dream, the Semantic Web is contested. There are some who suggest that the current generation of web-based applications are too entrenched for such a novel architecture to succeed (e.g., Ewalt, 2002). Yet others suggest that the representational infrastructure is unnecessarily complicated (Hayes, 2002).

Note that while it often seems from accounts of the Semantic Web that the various ontologies (debate, analogy, narrative, domain) are of the greatest importance, in fact

the Semantic Web will not be fully realized until a range of applications is built on top of these ontologies. Thus it is likely that semantic web services, where agents seek out suitable web services using semantic descriptions, and more particularly educational semantic web services will be most important in the next ten years. The ontologies provide the interoperable data while the services do the interesting things (like supporting argumentation).

6   What current research on the Semantic Web and eLearning offers

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*— or 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 here 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. Anido et al. (2002) provide a good overview here.

Another development has been the growth in educational repositories and peer-to-peer networks for sharing these. One example here is the Edutella network (Nejdl et al. 2002). 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. Work on the Educational Modelling Language is key here (Koper, 2000; 2001)

Most of this 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 (Anido et al., 2002) and course construction tools (Stojanovic et al., 2001) which attempt to automatically combine learning objects into "courses" or sequences of objects. Finally, 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 (Simon, et al., 2003.). The services here range from assessment, to short lectures, courses and degree programmes.

There are three main points here. Firstly, none of these technologies has reached full maturity as yet or been deployed widely so it is hard to gauge their success or failure.

Secondly, the Semantic Web technologies depend for their success on the viability of the strategy of depending on reusable atomic learning components, on the possibility of capturing the characteristics of these in formal descriptions using metadata schemes and the widespread acceptance of these descriptions 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. There have been numerous criticisms of the whole project of e-learning (Dreyfus, 2001), of the suitability of Learning Objects outside an individualistic pedagogic model — one moreover based on the needs of military trainers — (Friesen, 2003) and of the adequacy of the metadata used for descriptions (Friesen, 2002; Nilsson et al, 2002).

Thirdly, while, Nejdl and his colleagues have advanced the notion of Educational Web Services, these are still based largely on Learning Objects: "Educators and (semi-) automated tutoring systems compose learning services out of learning objects and other educational resources" (Simon et al. 2003). To this extent, their project may fail if Learning Objects fail to deliver. However, it seems that their architecture may be generic enough to make use of any properly annotated World Wide Web content. It is also true to say that their educational services do not differ from those provided by educational institutions in face-to-face contexts. It is possible to 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.

7   What's possible?

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 (LOM, SCORM, etc.) 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. In addition metadata schemes can be used to configure or sequence a learning object as part of some overall course.

What is lacking 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 CO_2 emissions (as concluded by the Intergovernmental Panel on Climate Change) there are others who either dispute the cause, the extent or the reality of global warming. For instance Lomborg (2001) has cast doubt on the quality of the IPCC models and suggested that the costs of limiting CO_2 emissions far outweigh the benefits. In turn others have disputed the case Lomborg makes.

Another example (Frankel, 1987; Stutt and Shennan, 1990) is provided by the various moves in the argument originating in Wegener's theory of continental drift. Indeed as Frankel shows, this sort of academic debate can be seen in terms of a variety of strategies. One of Wegener's claims was that there was once (200 million years ago) a continent, called Pangea by scientists, which broke up forming the present continents as the parts moved away from each other. One of the main ways of arguing against this theory was by showing anomalies in Wegener's account – specifically that the Permo-Carboniferous Ice Cap couldn't have formed if the whole of the Earth's surface land was massed around the Pole.

These controversies are characterized by hard-fought negotiations by political, economic and business stakeholders as well as the scientists. They rarely come to a complete conclusion though there may be resolutions, closures, and abandonment (McMullin, 1987).

A semantic service could be used as a means of browsing documents and, where appropriate, accessing not only services such as glossaries, but also more wide-ranging services along possibly many dimensions. For instance a semantic system could indicate to the climate science learner that the concept of global warming is a 'hot topic' unlike say that of the cause of anticyclonic winds. Having thus alerted the learner to the importance of the topic, the system could present material relating to the ongoing debate on the global warming issue using a sophisticated graphical interface to aid navigation through the various competing arguments.

8   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, some 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 — automatically highlights portions of the text which it can assist with, using technologies based on KMi's ClaimSpotter (Sereno, 2003) and Magpie (Dzbor et al., 2003). In this case SWEL can offer a way into the scientific debate about global warming and/or explain the scientific concepts involved.

The learner opts to access the scientific debate. SWEL provides a graphical interface to the principal components of this debate. In Figure 1 the environmental studies course component is part of the Climateprediction.Net experiment's web site - see http://www.climateprediction.net. There are two levels of Knowledge Chart. Level 1 shows the structure of an argument linking CO_2 rise to climate change. Level 2 shows part of the ongoing scientific controversy about this linkage.

Atmospheric concentrations of carbon dioxide have been increasing in the past 200 years or so since the Industrial Revolution began. The source is mainly the burning of fossil fuels. The rest is due to land use change, such as deforestation. **Figure 13** shows the atmospheric concentration of carbon dioxide in the past 1000 years (data have come from ice cores, direct measurements in recent years etc., if you're interested in this, read 'The two-mile time machine' by Richard B. Alley) and various estimates of how carbon dioxide concentrations will behave in the next 100 years, depending on how we react to legislation on carbon emissions. The concentrations used in the standard and doubled CO_2 experiments of climate*prediction*.net are marked.

Environmental Studies Course Component
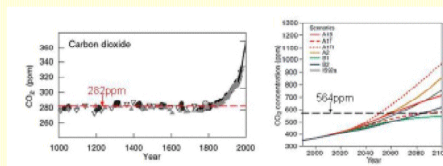
Figure 13. *The global atmospheric concentration of CO_2 in parts per million (ppm) (left) measured over the past 1000 years and estimated for the next 100 years (right). Source: IPCC Third Assessment Report. The CO_2 concentrations used in climateprediction.net experiments; 282ppm and 564ppm, are marked.*

Scientists are still uncertain exactly how the Earth-climate system will respond to such changes in carbon dioxide and other changes to the composition of the atmosphere.
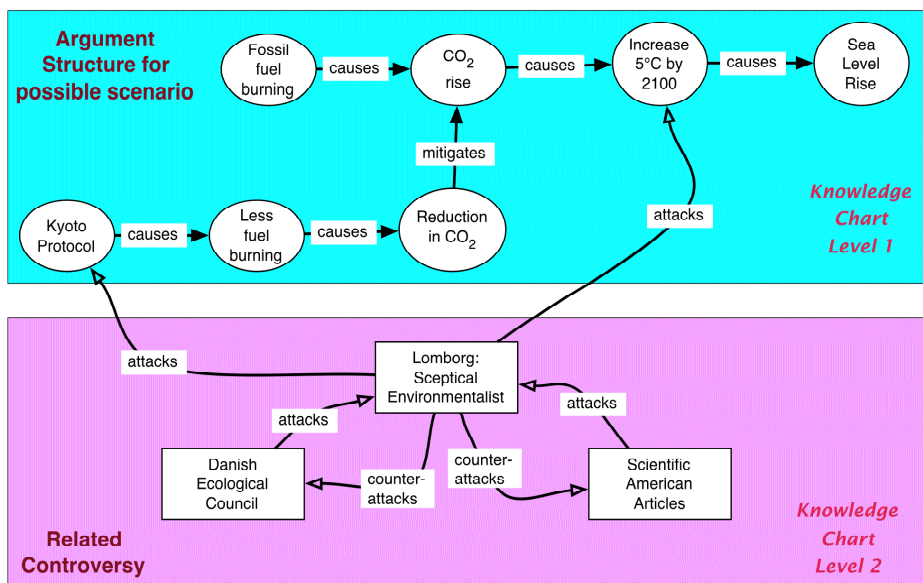

Figure 1: A course fragment (from the Climateprediction.net site) and its associated web of argumentation.

The learner clicks on the "Lomborg Sceptical Environmentalist" node. This opens up the node 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.
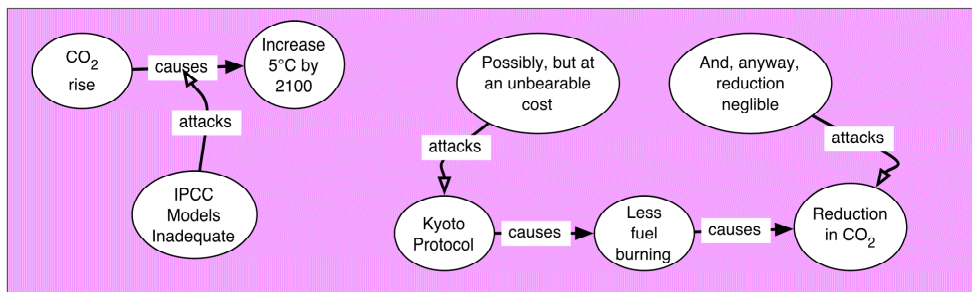

Figure 2: Lomborg's arguments

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.
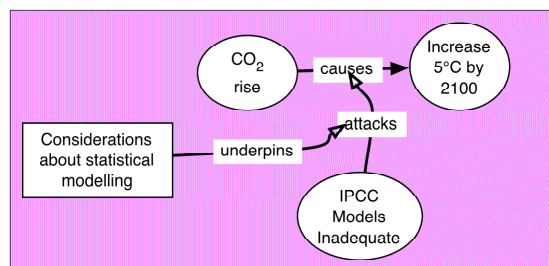
Figure 3: More detailed version of Lomborg's first argument

And so on. At each point in the debate model, the learner can access the original documents of which the model is a summary, using semantic search. In order for these documents to be of maximum use to the learner, in the context of this knowledge chart, the parts in the text which are relevant to these argument steps have been already semantically annotated, e.g., by a SWEL crawler using ClaimSpotter-

like technology or by the author himself, using scholarly semantic annotation tools, such as ScholOnto (Buckingham Shum et al., 2000).

Figure 4: One web document among many where Lomborg puts forward his views. In this case from the spiked science web site.

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 meta-models on classical and ecology-based economic models.

8.1   Our vision part I – Navigation of Knowledge Charts using Semantic Browsers

The key role of Knowledge Charts is to provide pathways through controversies, analogies and narratives 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 for analogical, narrative and argumentational models 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 content (summaries), 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. As already mentioned, work is already proceeding on the use of human language technologies to support the extraction of argument structures from texts (Sereno, 2003; Vargas-Vera and Moreale, 2003) and it is possible that this can be extended to search the Web components of the sorts of debates or controversies common in science (see section on Realizing the Vision below for more on this).

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 (perhaps in combination with information extraction techniques, see Vargas-Vera et al., 2002) without demanding explicit annotation. The various traversals (from text to Chart, from Chart to Chart and from Chart to new material) rely either on explicitly expressed relations among the Charts (or Chart components) or 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. We might also envisage a situation where agents such as SWEL could construct new Knowledge Charts. This should be possible by reasoning from the system's knowledge base of available resources, as well as its internal model of the sorts of components needed in a particular kind of Chart (as given by its ontology).

Note that while the illustrations given above in Figures 1–3 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.

In summary, SWEL needs: an ontology of types of Knowledge Chart (debate, story, analogy, causal model); ontologies for each type (claim, ground, support, refute for argument; event, actor, relationship for narrative); knowledge bases representing these; pre-designed meta-learning objects, such as an argument with all the possible links or, given the number of available resources, an automatic means of population perhaps using human language technologies; and, finally, a means of expressing the pedagogic purpose of these Charts (i.e., an extension to EML).

The real conceptual and technical difficulties will arise when we try to work out how to keep track of the complex multidimensional sets of relations for debate, community structure, community roles, practices and so on. It is likely that the populated ontologies for domain contents (i.e., the contents of learning objects) will be immense. If we combine this with the populated ontologies for all the meta-models it is possible that this will grow exponentially since for each web page or learning objects, there are many possible Knowledge Charts each of which can have in turn its own possible links to other Charts.

8.2   Graphical representation for Knowledge Charts

Taking the example of a debate or scientific controversy-based Knowledge Chart, it is important here that it can be viewed at many different levels. We need a view which encapsulates all of the moves but which also shows some of the complexities. There should also be views which select sections of the argument or, indeed, draw together themes into a single view (even if this isn't initially presented by SWEL). For example, the learner could select all the arguments made by particular participants. At the most detailed level, the nodes should represent all the significant points made in an argument. Similar points could be made about other types of Chart.

A sophisticated graphical representation is thus needed as a means of navigation by the learner. The sorts of multi-dimensional structures we envisage have much in common with hypertexts. They differ in that we do not consider the path through the various types of material as a single text. The course is the central 'text' with the various alternative resources, debate structures, narratives and analogies merely diversions along the way.

In order to provide a graphical interface to the scientific debate outlined above, SWEL must be provided with or derive at least a set of *debate* moves made by debaters linked by a set of *rhetorical relations* (i.e. a *debate ontology*). In this case the debaters are Lomborg, the Danish Ecological Council, and contributors to Scientific

American. In essence there are two main debate moves: Lomborg's attack on the notion of global warming and the subsequent rubbishing of his book by a variety of debaters. In addition there is Lomborg's detailed riposte to the Danish Ecological council. The relations are all attack or refute here but we can imagine a whole set of possible relations including support, restate, and confirm. In the case of a narrative, SWEL would have a different model made up of applicable ontology instances (e.g., a narrative has events, temporality, protagonists, plots, flash-backs etc.).
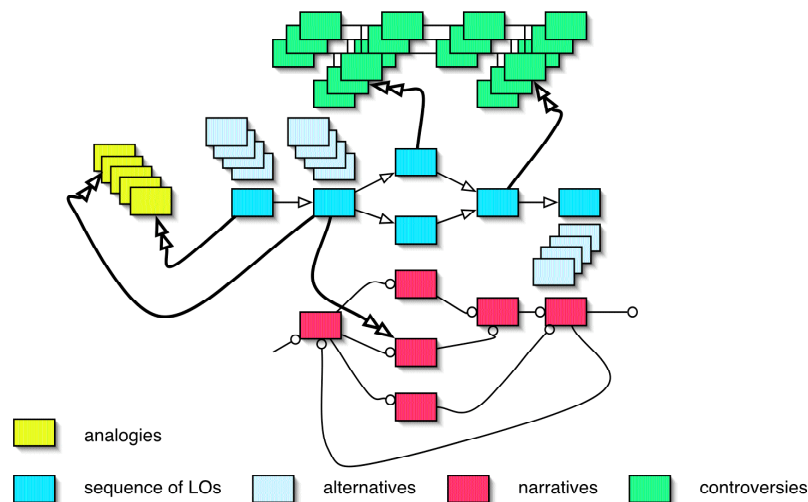
Figure 5: A schematic space with a course (or sequence of Learning Objects) which may be linked by a Semantic Browser to various types of Knowledge Charts

In Figure 5 we see part of a notional course which has associated with each of the learning objects sets of alternative expository material (such as web sites or individual web pages retrieved by semantic search), and Knowledge Charts for narratives, debates/controversies and analogies.

8.3   Our vision part II – Knowledge Neighbourhoods

While we can already provide all the necessary ontologies for the climate change or continental drift examples, we could go further and situate the learner in what we call *Knowledge Neighbourhoods*. These *semantic neighbourhoods for learning* combine community support with educational semantic web services and are operationalized by providing portals though which their members interact. Underlying this are

mappings from what is known about communities to services they provide. The notion originates in work by Domingue at al (2003) on the Parish framework for providing web services to the community of users of a local council's facilities.

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. A Knowledge Neighbourhood is thus a *Virtual Settlement* in Jones' (1997) sense of the term – "the cyber-place within which a virtual community operates". Knowledge Neighbourhoods are thus 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 also organized in terms of the practices they are concerned with or the tasks they perform. They are thus *communities of practice* (Wenger, 2000). 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 perform as a *leader* in a particular community with relations such as *sets-the-agenda-for* with other members. These roles and relations may change over time. These ideas were developed during the Alice project (Domingue et al., 2003) where they were applied to a customization service for ecommerce. They are even better suited to a community oriented approach to eLearning. Knowledge Neighbourhoods are composed of a variety of spaces: both public (for storing important documents, for debate; for publishing; for visualizing the community) and private (personal notes, calendars).

Communication across the boundaries of the various communities can be achieved in two main ways: (a) communities with members who also belong to other communities can rely on these members' synthesizing and translation abilities; (b) *boundary objects* can be used to effect communication. Leigh Star defines these as "objects which are both plastic enough to adapt to local needs and the constraints of the several parties employing them, yet robust enough to maintain a common identity across sites. They are weakly structured in common use, and become strongly structured in individual use. These objects may be abstract or concrete. They have different meanings in different social worlds but their structure is common enough to more than one world to make them recognizable, a means of translation. The creation and management of boundary objects is a key process in developing and maintaining coherence across intersecting social worlds." More succinctly Arias and Fischer (2000) view them as "objects that serve to communicate and coordinate the perspectives of various constituencies". In our view, Knowledge Charts are boundary

objects: they are both abstract representations and concrete entities; they are important for knowledge acquisition in particular communities but will be recognizable to at least closely related domains; and, the meanings they embody (through the ontologies they employ) and give access to, will vary across domains.

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. It is more likely that the Semantic Web, like Ancient Athens or Medieval Italy, will be composed of loosely related Knowledge Neighbourhoods.

8.4   Connecting the visions

While understanding may come about through awareness of ongoing scientific debate or controversy (vision I) it often requires the learner to *participate* in the chains of argumentation. Since only qualified peers will be able to enter fully into the controversy, practical experience in arguing these points comes about by means of discussions or debates among learners.

Knowledge Neighbourhoods are an important means of contextualizing Knowledge Charts. It is unlikely that learning object providers will be able to provide the sophisticated representations needed for these. It is probably unlikely that a new class of developers will arise (though this is possible) dedicated to their production. Knowledge Neighbourhoods will be provided with the necessary tools for designing and storing Knowledge Charts and, at the same time, the motivation for expending the necessary energy. Building Charts may become a community enterprise. We envisage that each learning community will have its own set of specialized Charts although it is possible that some of these may be shared (as boundary objects) with other learning communities.

Tools such as Semantic Browsers can provide the mappings from current learning object to Knowledge Chart to alternative learning materials. Other tools will be needed which deploy ontologies in constructing Knowledge Charts and in mediating community formation and intra-community discussion.

9   Realizing the visions: technologies needed for SWEL

9.1   A technology-centred analysis

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 (Naeve, 1997) are learning environments which can be used to explore networks of ideas. They are also developing (Naeve et al., 2001) the idea of the *Conceptual Web* —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 need another tool — a Semantic Writer/Constructor — which can construct or assist in the construction of Knowledge Charts. Underlying these we need various domain and structure-related ontologies (e.g., for argumentation) and an infrastructure for creating and publishing semantic web services.

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.

Service Infrastructure. Research on semantic web services is concerned with developing infrastructure, modelling and reasoning support to allow automatic discovery, composition and execution of web services. The IRS-II infrastructure (Motta et al., 2003) 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. Recently we have developed a demonstration system which can schedule operations abroad for patients with arthritis and needing urgent hip replacement surgery. The system relies on the availability of distributed web services, such as yellow pages for hospital treatments, hospital scheduling systems, ambulance providers, and currency converters.

Ontologies for argumentation-based Knowledge Charts. As we have indicated above, there is already extensive support for scholarly argumentation in the ScholOnto project (Buckingham Shum et al., 2000). In order to support disagreement and conflicting perspectives in academic research fields, we need tools which support the user in making sense of the relations among documents. The ScholOnto project is developing an ontology-based digital library server to support scholarly interpretation and discourse. Researchers can articulate their view of where a document fits in the ongoing academic conversation thus creating a semantic network of scholarly discourse. A tool — ClaiMaker — has been developed to model the rhetorical relations (proves, refutes, is consistent with, is analogous to, and so on) among claims in research papers, publish these on a server and make queries about the relations and the documents containing them (see Li et al., 2002). While not intended primarily as a learning tool, both the access to a web of inter-related scholarly papers and the opportunity to add further annotations (i.e., to extend the semantic web of research papers) have educational and well as research applications.

Other Knowledge Charts. There is embryonic support for representing and capturing narratives in the CIPHER project (Mulholland et al., 2002) 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. There is less support currently for analogy-based and causal Knowledge Charts although there is a great deal of work in the AI literature on both of these (e.g., Gentner, 1983, on analogy; Forbus, 1990, on qualitative physics).

A means of visualizing Knowledge Charts. The ScholOnto project is already exploring the use of the Compendium tool (www.CompendiumInstitute.org) for representing the network of scholarly discourse.

Repositories for structures/charts. Again, the ScholOnto tool provides the basis for a Knowledge Chart repository.

Automated Construction of Knowledge Charts. Knowledge Chart construction is the other side of the coin from Semantic Browsing, and an important activity for learning communities. Some 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 (Vargas-Vera et al. 2002; Ciravegna et al., 2002) more is needed for the identification of concepts from argumentation, narrative and other structural ontologies. While work here is in its early stages there are promising results from work on automated argumentation extraction in the Knowledge Media Institute. Vargas-Vera and Moreale (2003) report on experiments in extracting arguments from student essays, while there is currently ongoing work on a tool — ClaimSpotter — for automatic discovery of scholarly claims as part of the ScholOnto project. Work on story extraction from news articles is reported in Vargas-Vera and Celjuska (2003).

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 (Dzbor et al. 2003) 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 mentioned above. 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 the ontology could contain the class 'Project'. Clicking 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 much of the necessary infrastructure is already available or could be readily adapted, there remains a great deal of work to be done on quite central components of our envisaged system.

9.2   Envisaged extensions to existing semantic technologies

Extensions to Magpie.   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 text to Knowledge Charts, Knowledge Charts to Knowledge Charts and Knowledge Charts to other web resources. 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, 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. Finally, as Laurillard points out (pp. 43-4), an academic argument may extend over an entire article with its components scattered throughout the text.   Hence, a means of identifying and collecting the parts of complex concept instances scattered through a text is needed. This consolidation is especially important for instances of concepts such as *argument, claim* and *premises*. The work mentioned above on ClaimSpotter could form the basis for new tools here.

Narrative, causal and analogical ontologies and their representations. As we have already said much more needs to be done on producing pedagogically relevant ontologies for narrative, causation and analogy. More importantly, we need some means of visualizing these in a manner which is succinct enough to be used for

navigation but detailed enough for the learner to make sense of the debate, narrative, causal model or analogy.

Further work on repositories of Knowledge Charts. Work remains to be done on the exact nature of the repositories needed for the storage and retrieval of Knowledge Charts. For instance, deciding whether a central registry with distributed storage nodes or a peer-to-peer network would be more appropriate.

The Semantic Writer/Constructor. This component of our system will assist in the creation of Knowledge Charts. It will have some ability to recognize instances of structural ontologies as discussed above. However, 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 we have argued elsewhere, the construction of complex representations for knowledge is in itself a useful exercise (Stutt, 1997). 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, reformat in RDF). New work just started on Magpie will result in a system which can use information extraction 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.

Community-based tools. Tools are needed for identifying, creating, supporting, fostering, and tracking communities and integrating them with knowledge representations into Knowledge Neighbourhoods. While we have begun work on community-oriented customization in the e-shopping arena (Domingue et al., 2003) much more work needs to be done on identifying the characteristics and structures of communities as well as their dynamics.

10   Conclusion — Learning Webs and Critical Thinkers

According to Fisher (1988, p. 1):

> "it is possible to rely too heavily on experts and this approach to learning and knowledge tends to encourage passivity and receptiveness rather than inventiveness and imagination… One object of this book is… to impress on the reader what a long way one can get in understanding any subject by thinking it through for oneself… We shall do this by concentrating on the arguments experts have produced for believing a wide range of things and

showing how it requires only a relatively slight knowledge of the subject to evaluate the arguments oneself."

This is true whether we learn via the Web or from books. The Web, however, 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.

Instead of the oracular pronouncements of Kaku's Magic Mirror, a means of traversing the various links possible from web document to web document by *means of a meta-model or models— a Knowledge Chart — expressing the associated debate, narrative, analogy and so on* will be most valuable to learners. 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 a knowledge based society.

References

Andriessen, J., Baker, M. and Suthers, D. eds. (2003) *Arguing to Learn: Confronting Cognitions in Computer-Supported Collaborative Learning Environments*. Dordrecht: Kluwer Academic Publishers.

Anido, L. E., Fernández, M.J., Caeiro, M., Santos, J.M., Rodríguez, J.S. and Llamas, M. (2002) Educational metadata and brokerage for learning resources. *Computers & Education Volume 38*, Issue 4, 351-374.

Arias, E. G. and Fischer, G. (2000) Boundary Objects: Their Role in Articulating the Task at Hand and Making Information Relevant to It. *International ICSC Symposium on Interactive and Collaborative Computing (ICC'2000)*, December 2000. Available from: http://www.cs.colorado.edu/~gerhard/papers/icsc2000.pdf

Berners-Lee, T., J. Hendler, and 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 284(5): 10.*

Bishop, A. (2002) Come into my Parlour Said the Spider to the Fly: Critical Reflections on Web-Based Education from a Student's Perspective. *Distance Education*, 23(2), 231-236.

Brown, J.S. and Duguid, P. (2002) *The Social Life of Information*. Boston, MA: Harvard Business School P.

Brown, J.S., Collins, A. and Duguid, P. (1989). Situated Cognition and the Culture of Learning. *Educational Researcher, 1989*. pp. 32-42. Available from: http://www.ilt.columbia.edu/Publications/papers/JohnBrown.html

Buckingham Shum, S., Motta, E. and Domingue, J. (2000) ScholOnto: An Ontology-Based Digital Library Server for Research Documents and Discourse. *International Journal on Digital Libraries (2000), 3 (3)*, pp. 237-248.

Ciravegna, F., Dingli, A., Petrelli, D. and Wilks, Y. (2002). User-System Cooperation in Document Annotation based on Information Extraction. Proceedings of the 13th International Conference on Knowledge Engineering and Management (EKAW 2002), *Lecture Notes in Artificial Intelligence 2473*, Springer Verlag, 2002.

Decker, S., Erdmann, M., Fensel, D. and Studer, R (1999) Ontobroker: Ontology Based Access to Distributed and Semi-Structured Information. Meersman, R et al. (Eds), *Semantic Issues in Multimedia Systems*. Proceedings of DS-8. Boston: Kluwer Academic Publisher. pp. 351-369.

Domingue, J. (2003) Parish: A community centred framework for Semantic Web Services. *Unpublished Memo. Knowledge Media Institute*, The Open University, UK.

Domingue, J., Stutt, A., Martins, M., Tan, J., Petursson, H., and Motta, E., (2003) Supporting Online Shopping through a Combination of Ontologies and Interface Metaphors, *International Journal of Human Computer Studies, Volume 59, Issue 5,* November 2003, Pages 699-723.

Dreyfus, H (2001) *On the Internet*. London: Routledge.

Dzbor, M., Domingue, J.B., Motta, E. (2003). Magpie - Towards a Semantic Web Browser. *Proceedings of the 2nd International Semantic Web Conference (ISWC 2003)*, October 2003, Florida, US.

Ewalt, D. E. (2002) The Next Web: The Semantic Web, the next step in the Web's evolution, promises even more dramatic changes. *InformationWeek*, 14 Oct. 2002. Available from: http://www.informationweek.com/story/showArticle.jhtml?articleID=6503831

Fisher, A. (1988) *The Logic of Real Arguments*. Cambridge: Cambridge UP

Forbus, K.D. (1990) Qualitative Physics: Past, Present, and Future. In D. S. Weld and J. de Kleer (eds), *Qualitative Reasoning about Physical Systems*. Kaufmann. pp. 11–39.

Frankel, H. (1987) The Continental Drift Debate. In: eds. Engelhardt, H.T., Jr. and Caplan, A. L. *Scientific Controversies: Case Studies in the Resolution and Closure of Disputes in Science and technology*. Cambridge: Cambridge UP. pp. 203-248.

Friesen, N. (2002) Semantic Interoperability and Communities of Practice. In Mason, Jon. (Ed.) *Global Summit of Online Learning Networks: Papers.* Adelaide: Educationau, 2002 Available from:
http://www.educationau.edu.au/globalsummit/papers/nfriesen.htm

Friesen, N. (2003) Three Objections to Learning Objects. In McGreal, R. (ed.), *Online Education using Learning Objects.* London: Taylor & Francis Books Ltd. Available from: http://phenom.educ.ualberta.ca/~nfriesen

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy, *Cognitive Science, 23*. pp.155–170.

Hayes, P. (2002). Catching the Dreams [Online] A*vailable from: http://www.aifb.uni karlsruhe.de/~sst/is/WebOntologyLanguage/hayes.htm*

Jones, Q. (1997). Virtual Communities, Virtual Settlements and Cyber-Archaeology: A Theoretical Outline. *Journal of Computer Mediated Communication. 3(3).* [Online] Available from: http://www.ascusc.org/jcmc/vol3/issue3/jones.html

Kaku, M. (1998). *Visions: How Science Will Revolutionize the 21st Century*. Bantam Books.

Kirschner, P.A., Buckingham Shum, S.J. and Carr , C. S. eds. (2003) *Visualizing Argumentation: Software Tools for Collaborative and Educational Sense-Making*. London: Springer-Verlag. Information available from:
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. A*vailable from: 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 from: eml.ou.nl/introduction/docs/ped-metamodel.pdf*

Lassila, O. and Swick, R. R. (1999). Resource Description Framework (RDF) Model and Syntax Specification. *W3C Technical Report*, 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. London: RoutledgeFarmer.

Li, G., Uren, V., Motta, E., Buckingham Shum, S., and Domingue, J. (2002) ClaiMaker: Weaving a Semantic Web of Research Papers. *1st Int. Semantic web Conference (ISWC 2002)*, Sardinia.

Lomborg, B. (2001) *The Skeptical Environmentalist*. Cambridge: Cambridge UP.

Marshall, C. and Shipman, F. (2003) Which Semantic Web? *Proceedings of ACM Hypertext 2003*, pp. 57-66.

McGuinness, D., and van Harmelen, F. (2003). OWL Web Ontology Language Overview. *W3C Technical Report*, http://www.w3.org/TR/owl-features/

McMullin, E. (1987) Scientific Controversy and its Termination. In: eds. Engelhardt, H.T., Jr. and Caplan, A. L. *Scientific Controversies: Case Studies in the Resolution and Closure of Disputes in Science and technology*. Cambridge: Cambridge UP. pp. 49-91.

Motta, E., Domingue, J., Cabral, L. and Gaspari, M. (2003) 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.

Mulholland, P., Zdrahal, Z. and Collins, T. (2002). CIPHER: Enabling Communities of Interest to Promote Heritage of European Regions. *Cultivate Interactive, issue 8*, 15 November 2002. Available from: http://www.cultivate-int.org/issue8/cipher/

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. And Palmér, M. (2001) E-Learning in the Semantic Age. *Proc. 2nd European Web-based Learning Environments Conference (WBLE 2001)*, Lund, Sweden.

Nejdl, W., Wolf, B., Qu, C., Decker, S. Sintek, M., Naeve, A., Nilsson, M., Palmér, M. and Risch, T. (2002) EDUTELLA: A P2P Networking Infrastructure Based on RDF. *Proceedings Eleventh International World Wide Web Conference WWW2002*, May 2002.

Nelson, M.R. (1994) We have the information you want, but getting it will cost you!: Held hostage by information overload. *ACM Crossroads 1(1)*. Available from: http://www.acm.org/crossroads/xrds1-1/mnelson.html

Nilsson, M., Palmér, M. and Naeve, A. (2002) Semantic Web Metadata for e-Learning - Some Architectural Guidelines. *Proceedings Eleventh International World Wide Web Conference WWW2002*, May 2002.

Sereno, B. (2003). From scholarly documents to interpretive claims: an approach to bridge the formalisation gulf. *Unpublished Memo, Knowledge Media Institute*, The Open University, UK.

Simon, B., Miklós, Z., Nejdl, W., Sintek, M. and Salvachua, J. (2003) Smart Space for Learning: A Mediation Infrastructure for Learning Services. *Proceedings of the 12th World Wide Web Conference*, May, 2003. Available from: http://nm.wu-wien.ac.at/research/publications/b164.pdf

Star, S. L. and James Griesemer, J. "Institutional Ecology, 'Translations,' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-1939," *Social Studies of Science, 19*: 387-420 (1989).Reprinted in Mario Biagioli, ed. The Science Studies Reader. Pp. 505-524.London:Routledge.

Stojanovic, L., Staab, S and Studer, R. (2001) ELearning Based on the Semantic Web. *Proceedings WebNet2001 - World Conference on the WWW and Internet*, Orlando, Florida, USA, 2001.

Stutt, A. (1997) Knowledge engineering ontologies, constructivist epistemology, computer rhetoric: A trivium for the knowledge age. *Proceedings Ed-Media'97*, Calgary, Canada, June 1997.

Stutt, A. and Shennan, S. (1990) The Nature of Archaeological Arguments. *Antiquity, 64, 245*, December 1990. pp. 766-767.

Stutt, A., Martins, M.C. and Domingue, J. (2002) Tribalization, E-learning and the Semantic Web. In eds. Kähkönen, E. and Sutinen, E. *Proceedings First Intl. Conference on Educational Technology in Cultural Context*, Finland, Septmber, 2002.

Sumner, T. and Buckingham Shum, S. (1998). From Documents to Discourse: Shifting Conceptions of Scholarly Publishing. *Proceedings of CHI 98: Human Factors in Computing Systems*, 18-23 April, 1998, Los Angeles, CA. ACM Press: New York.

Vargas-Vera, M. and Celjuska, D. (2003) Ontology-driven Event Recognition on Stories. *KMi Tech. Report KMI-TR-135.*

Vargas-Vera, M. and Moreale, E. (2003) A Question-Answering System Using Argumentation. *KMi Tech. Report KMI-TR-132.*

Vargas-Vera, M., Motta, E. Domingue, J., Lanzoni, M., Stutt, A. and Fabio Ciravegna, F. (2002) MnM: Ontology Driven Semi-Automatic and Automatic Support for Semantic Markup. The 13th International Conference on Knowledge Engineering and Management (EKAW 2002), *Lecture Notes in Artificial Intelligence 2473*, Springer-Verlag.

Wenger, E. (2000) *Communities of Practice: Learning, Meaning, and Identity*. Cambridge: Cambridge University Press.