August 2004

The Role of Ontologies in Teaching and Learning

Ruth Wilson
rw_consult@hotmail.com

Table of Contents

List of Figures

Ruth Wilson
rw_consult@hotmail.com

1. Executive Summary

Ontologies are currently a buzzword in many communities, hailed as a mechanism for making better use of the Web. They offer a shared definition of a domain that can be understood by computers, enabling them to complete more meaningful tasks. Although ontologies of different descriptions have been in development and use for some time, it is their potential as a key technology in the Semantic Web which is responsible for the current wave of interest. Communities have different expectations of the Semantic Web and how it will be realised, but it is generally believed that ontologies will play a major role.

In light of their potential in this new context, much current effort is focusing on developing languages and tools. OWL (Web Ontology Language) has recently become a standard, and builds on top of existing Web languages such as XML and RDF to offer a high degree of expressiveness. A variety of tools are emerging for creating, editing and managing ontologies in OWL.

Ontologies have a range of potential benefits and applications in further and higher education, including the sharing of information across educational systems, providing frameworks for learning object reuse, and enabling intelligent and personalised student support. The difficulties inherent in creating a model of a domain are being tackled, and the communities involved in ontology development are working together to achieve their vision of the Semantic Web.

This Technology and Standards Watch report discusses ontologies and their role in the Semantic Web, with a special focus on their implications for teaching and learning.

2. Keywords

Ontologies, Semantic Web

3. The Technology

Ontologies promise “a shared and common understanding of a domain that can be communicated between people and application systems” [1]. They attempt to formulate a thorough and rigorous representation of a domain by specifying all of its concepts, the relationships between them and the conditions and regulations of the domain. Ontologies can express hierarchical links between entities as well as other semantic relations. An example of part of an ontology is provided in Figure 1, in which it is specified not only that an author *is a* person and that a book *is a* publication, but also that an author *writes* a book and that a book *has* chapters.

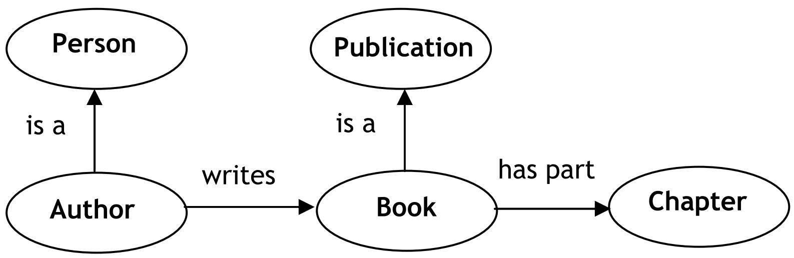

Figure 1. Example of a small ontology

Many different specifications have been termed ontologies, from glossaries to taxonomies to logical theories [2]; examples include CYC (a representation of human knowledge and common sense), the Dewey Decimal Classification (used in libraries) and the DMOZ Open Directory (http://www.dmoz.org/). What they all seek to achieve is to represent some area of knowledge or an aspect of the world.

Recently, ontologies have been proposed as an enabling technology for the Semantic Web. Most information on the Web is written in natural language, intended for humans to read. The Semantic Web proposes to make this information also understandable by computers. To illustrate the difference, finding information on the Web usually involves conducting a keyword search and then filtering and analysing a list of retrieved resources. On the Semantic Web, software agents would be able to deal with more complex queries, filtering and extracting meaning on behalf of the user, and returning a direct response. Some believe that such reasoning services will become so integrated into everyday life that they will be “as necessary to us as access to electric power” [1].

Making the Semantic Web a reality involves adding metadata to documents and creating structures to represent knowledge. By modelling domains using ontologies, all the different kinds of resources on the Web can be given a common semantics, enabling users to complete more complicated tasks.

4. The Technology and Standards Watch Issue

Ontologies, and the increased functionality they imply, will bring new opportunities to e-learning and e-teaching. They offer possibilities for organising and visualising didactic knowledge, and for this knowledge to be shared and reused by different educational applications [3]. Further, learners will be able to interact with Web-based courses and other educational systems in new, more intricate and personalised ways [4]. Ontologies will enable complex and dynamic learning requirements to be met automatically, and will assist learners in comprehending the domain and building their own concept map. In short, they promise to make the Web a more effective resource for teachers and learners.

This report will introduce ontologies to the further and higher education community, explaining why they are being developed, what they hope to achieve, and their potential benefits to the community. Current ontology tools and standards will be described, and the emphasis will be on introducing the technology to a new audience and exploring its risks and potential applications in teaching and learning. At a time when educational programmes based on ontologies are starting to be developed, the author hopes to increase understanding of the key issues in the wider community.

5. What are Ontologies?

5.1 Some Definitions

The term "ontology" is inherited from philosophy, in which it is a branch of metaphysics concerned with the nature of being. It began being used in artificial intelligence in the 1980s, and is now frequently used by computing and information science audiences.

Understanding what ontologies are can be problematic, since they are often spoken about in high-level terms. Gruber's commonly cited definition, for example, states that "an ontology is an explicit specification of a conceptualisation" [5]. (A conceptualisation is "an abstract, simplified view of the world that we wish to represent for some purpose" and comprises the concepts that exist in an area of knowledge and the relationships between them [6]). Put simply, an ontology is a common vocabulary for describing a domain that can be used by humans as well as computer applications.

When it comes to discussing the practical issues surrounding ontologies, it is typical for the individual communities involved in research and development to discuss them using their own terms and according to their own familiar frameworks. Artificial Intelligence experts, for example, may talk about them in terms of logical theories and frame languages. Information Scientists, on the other hand, may equate them to taxonomies and thesauri. A good summary of artefacts which have been described as ontologies, and their relative levels of expressiveness, is provided by Deborah McGuinness in her "ontology spectrum" [2], which includes:

- Controlled vocabularies
- Thesauri
- Classification schemes
- Taxonomies
- Topic maps
- Frame languages
- Logical theories

At the "simplest" end of the spectrum are lists of the things in a domain, such as controlled vocabularies and glossaries. More expressive are thesauri and hierarchies, which describe the relations between things. And at the most expressive end of the spectrum are ontology languages which allow the expression of logical constraints and detailed relationships such as disjoint, inverse, part-of, and so on.

5.2 Ontologies and the Semantic Web

So why is there so much talk of ontologies at the moment? The Library of Congress Subject Headings, WordNet and Yahoo! would all have their place in McGuinness's spectrum, but are certainly not new. In fact, the knowledge management, database design and information retrieval fields have been considering ontologies for some time [7], but it is their potential in the context of the Semantic Web that accounts for their recent popularity as a research topic in many communities.

The idea behind the Semantic Web is, as the name suggests, to add a level of meaning to the Web so that it can be more easily manipulated by computer programs, and thereby used more effectively by humans.

Berners-Lee et al, in their seminal paper “The Semantic Web” [8], put forth the notion of a global knowledge base, envisaging interacting with a software agent to complete sophisticated tasks. The example running through their paper involves booking a series of physical therapy sessions based on a number of variables, including:

- The doctor’s prescribed treatment
- The location of therapists
- Available appointments
- The quality of therapy provided
- The patient’s schedule
- Availability of relatives to chauffeur the patient

All of these variables are discovered and coordinated automatically by instructing an agent through a Web browser, and the user is presented with a simple plan.

This notion of the Semantic Web as a global knowledge base comprising a framework of agents is just one view, and is considered by many to be more harmful than useful. Another view stems from the library perspective, in which order and structure are brought to the highly heterogeneous set of documents available on the Web via metadata and classification, resulting in the “ultimate digital library” [9]. The functionality of a well-organised information source is brought to the Web, on which anyone, anywhere can publish anything.

A third perspective treats the Semantic Web as an infrastructure for communicating and sharing data with independently developed components, using common terminologies through ontologies for integration. It is this vision, of weaving together multiple, domain-specific ontologies, which drives much of today’s development. A detailed analysis of the three perspectives is provided in “Which Semantic Web?” by Marshall and Shipman [9].

XML goes some way towards providing the kind of functionality required to realise this vision, by enabling the structure, rather than just the content and layout, of documents to be defined. RDF goes further still, by providing a language for describing the content of resources. Meaning is encoded in “triples” of subject, predicate and object, and this simple structure can be used to describe much of what we might want to express. For instance, we can specify that John (the subject) studies at (the predicate) Glasgow University (the object), or Sarah (subject) lectures in (predicate) English Literature (object). Each resource, property and value is linked to a Universal Resource Identifier (URI) which provides a unique machine-interpretable definition. Content which is marked up in RDF can then be manipulated by software which is able to access these structures. RDF Schema (RDFS), moreover, allows classes and properties to be associated with resources so that basic models of knowledge can be built for specific applications. For example, undergraduate and postgraduate can be defined as subclasses of student, and lecturer as a subclass of university staff.

However, RDF and RDFS do not allow classes and properties to be defined precisely enough for some of the envisaged Semantic Web scenarios. For example, they cannot describe classes’ disjointedness or cardinality, or the characteristics of properties. And this is where ontology languages come in. Building on RDF, they offer greater expressiveness and machine interpretability and enable concepts in a domain and the relations between them to be defined explicitly. Web pages can be embedded with metadata that refers to an ontology, and software that can interpret the ontological relations and rules is then able to act on the knowledge it has about the semantics of the pages. For example, a set of educational resources on the Web can be marked up so that they link to definitions of the topics they explain; software can then use this data to locate resources on a particular topic, or to make inferences based on the relations between topics, as defined in the ontology. Multiple resources may draw upon and perform reasoning tasks with these definitions, effectively adding a layer of meaning to the Web.

6. Languages and Tools

6.1 Ontology Languages

OWL (Web Ontology Language) builds on RDF and enables more sophisticated semantics to be defined; it is currently the focus of much work and attention. Its development arose from the need for machine-readable descriptions of the meaning of terminology used in Web pages, to improve the "poorly mapped geography" of the Web so that computational agents can use it more effectively [10]. It has been driven by the following goals:

- To make ontologies publicly available so that their semantics can be shared.
- To provide a consistent approach to ontology revision.
- To enable ontologies which model concepts in different ways to be mapped.
- To detect inconsistencies between ontologies.
- To express different kinds of knowledge, and be able to reason with it.
- To be easy to learn, encouraging adoption.
- To be compatible with other standards.
- To support multilingual ontologies and cross-cultural views.

OWL has recently become a World Wide Web Consortium Recommendation [11], and these goals are reflective of the current issues and concerns facing ontology developers as a whole. It provides three sub-languages, intended to meet different users' needs:

- OWL Lite can provide a class hierarchy together with some simple constraints.
- OWL DL is based on Description Logic formalisms which allow for powerful expression. The ontology can be reasoned but a "reasoner" is required.
- OWL Full enables maximum expressiveness, without regard for computational limits.

OWL is the product of much prior work, and is based on another ontology language called **DAML+OIL**, developed by a joint EU-USA, DARPA-sponsored ad hoc committee. Built on RDF(S), DAML+OIL is a markup language which allows semantics to be expressed in XML. It, in turn, supersedes **DAML+ONT**, and both are largely based on OIL.

OIL (Ontology Inference Layer) provides semantics for describing term meanings. It comprises four levels, each one more expressive than the last. Programs capable of only processing the core layer are still compatible with ontologies expressed in higher layers.

Other Web ontology languages to note include:

- **SHOE** (Simple HTML Ontology Extensions), an extension of HTML with tags for incorporating semantic knowledge into documents. Based on an ISA hierarchy, it is very simple but less expressive than some other specifications.
- **OML** (Ontology Markup Language) and **CKML** (Conceptual Knowledge Markup Language) adapt SHOE to XML and are based on Conceptual Graphs. OML was later redesigned to be RDF compatible.

6.2 Ontology Tools

At the same time as these ontology languages have been developed, tools have emerged for creating, editing and managing ontologies written in the various languages. Protégé 2000 is briefly described here to indicate the kind of functionality on offer, and some other popular tools are outlined below.

Protégé 2000 [12] was developed at Stanford University as an Open Source editing environment where ontologies can be constructed through a graphical user interface. It has been developed using a plug-in architecture, where new services can be added easily to the environment, and is perhaps the most widely used ontology tool. It can handle ontologies in XML, RDF(S), XML Schema, DAML+OIL and OWL.

Figure 2 shows a small ontology being created using Protégé's interface. In this example, several instances of the class "colour" have been defined, including colour encoding, colour systems, true colour, bicolour and greyscale. Any "slots" (attributes) can be defined - in this case, each instance has a name and three potential relations to other instances (Has Part, Requires and Suggested Order). "Forms" allow the user interface for entering ontology information to be manipulated for ease of use, and "queries" allow the ontology to be searched based on the attributes of instances. The Ontoviz plugin (Figure 3) allows ontologies to be viewed graphically; visualisation tools such as this can be helpful in the development process.

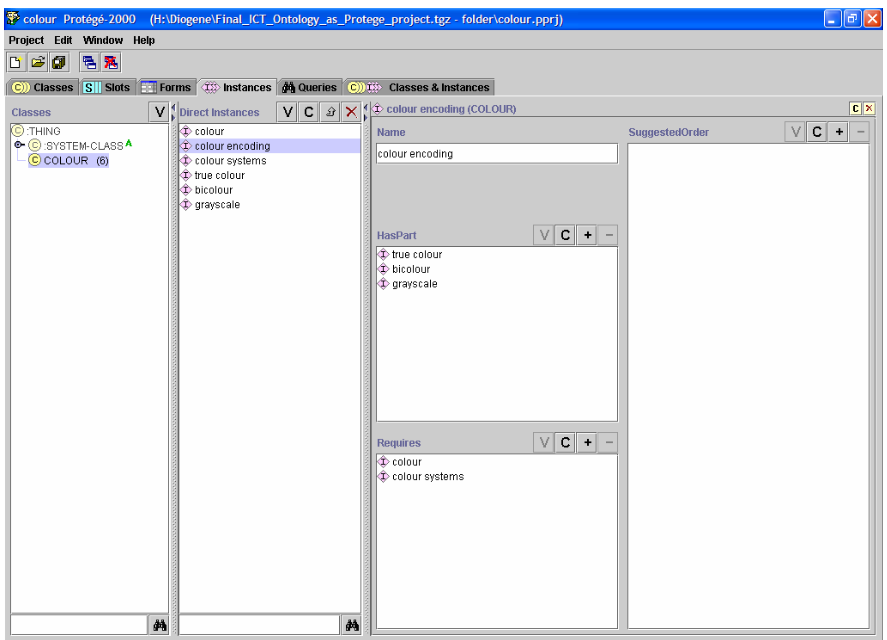

Figure 2. Small ontology in Protégé 2000

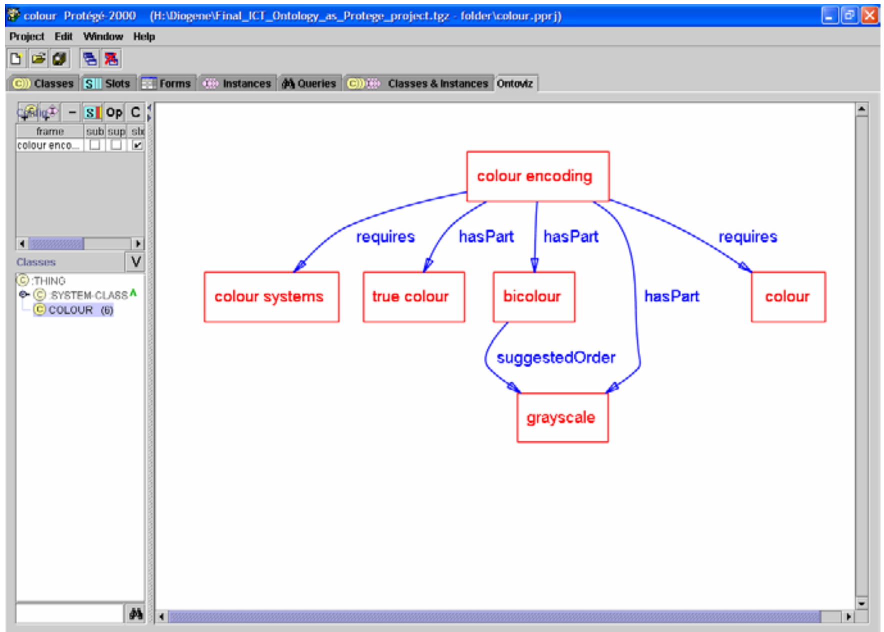

Figure 3. Ontoviz plugin for Protégé 2000

JISC has funded the CO-ODE project which unites the most popular native OWL editor (OilEd) with Protégé 2000 [13]. It has been designed to be used directly by experts in a particular domain rather than knowledge engineering experts. This has generated a lot of interest in the ontology development community.

Other well-used ontology development tools include:

- **OntoEdit** [14] was developed by the University of Karlsruhe and consists of a repository of ontologies, an inference and query engine and various translators. It supports the development and maintenance of ontologies by using graphical means. It is available in free and professional versions.
- **OilEd** [15] was developed at the University of Manchester for easy development of OIL languages (including DAML+OIL). It is a simple tool enabling users to create ontologies and check them for consistency, and is available as an Open Source project under the GPL licence.
- **WebODE** [16] supports most of the activities in the ontology lifecycle, and its ontologies can be integrated with other systems. It is available free on registration.
- **Ontolingua** [17] is an ontology editing environment which enables ontologies to be constructed collaboratively by distributed groups. It includes a library of reusable ontologies, which means that new ontologies can be created quickly from existing ones. Available free on registration.
- **Ontosaurus** [18] comprises an ontology server and an ontology browser server that dynamically creates HTML pages that display the ontology hierarchy. A live demonstration version is available which allows browsing but no editing.
- **LinkFactory** [19] is a formal ontology management system, designed to build and manage very large and complex language-independent formal ontologies. It consists of two components: the LinkFactory Server, which stores data in a relational database, and the LinkFactory Workbench, which allows the user to browse and model several ontologies and align them. Both components are developed in Java. Pricing and other information is available from Language and Computing Inc.

Tools are also available specifically for merging, annotating, storing and querying ontologies. Some, such as KAON (the KArlsruhe ONtology and Semantic Web tool suite), provide a framework for building ontology-based applications [20]. A detailed comparison is provided by the OntoWeb Consortium [21].

These tools support ontology development and management to different extents. Some support collaborative development, in which multiple members of a community of interest can have input to the process. However, no matter how advanced the tool, this can be a complicated process. This is discussed further in section 8.2.

7. Developments

The Artificial Intelligence community began developing ontologies in the 1980s, in a bid to enable natural language processing. These generally took the form of very large and ambitious knowledge bases, attempting to encompass all human knowledge and common sense. Examples include:

- CYC, "a formalised representation of a vast quantity of fundamental human knowledge: facts, rules of thumb, and heuristics for reasoning about the objects and events of everyday life" [22].
- SENSUS, an extension of WordNet (a combination dictionary and thesaurus for the English language), creating a large taxonomy for use in machine translation systems [23].
- Generalised Upper Model, another linguistically motivated ontology for expressing information in natural language, independent of task and domain [24].

More recently, the potential for ontologies in furthering the goals of the Semantic Web has become clear [25] and they have captured the imagination of wider computational audiences. Development of general ontologies has continued (CYC is now in its 21st year), but a trend has also emerged towards developing domain-specific ontologies. The DAML Library lists 282 ontologies [26]; collections are also available in the Protégé and OWL Ontology Libraries [27, 28].

The distributed nature of development has resulted in a degree of overlap, and it is envisaged that this will only increase with time; recent research has focused on methods for mapping ontologies to achieve interoperability.

In the 1990s, and running parallel to this formal ontology development, the library and information science (LIS) community's approach to bringing order to an increasingly unwieldy Web was to develop subject gateways. These organise internet resources (manually selected by subject experts or LIS professionals) under subject headings, providing easier access to quality resources for users. Gateways are often based on library classification schemes and existing thesauri, usually with some degree of modification to sufficiently represent the subject area on the Web. BUBL LINK [29], Humbul [30] and ADAM [31], for example, all draw upon the Dewey Decimal Classification (DDC) system to some extent; DDC was first produced in 1876, is currently in its 22nd edition and is widely used in libraries around the world.

A proliferation of gateways and other digital library initiatives has prompted the LIS community to focus efforts on interoperability and cross-searching activities. Approaches include the HILT Project [32], which has set up a pilot terminologies "route map", "to investigate issues relating between various subject schemes, namely DDC, LCSH, the UNESCO thesaurus, and the MESH thesaurus, in order to cater for cross-browsing and cross-searching across distributed digital collections". Renardus provides a single subject browsing service across a range of European subject gateways, via manual mapping of a range of schemes to DDC [33]. Attention has recently turned to more expressive ontologies, and how they might be used to enhance digital library developments [34].

And so the two communities – LIS and computer science - find themselves in similar places, working on ontologies for sharing data between a variety of knowledge organisation systems to make better use of the Web. One has expertise in creating and using ontologies, the other in formalising their expression, and there is much that they can learn from one another.

As ontology research and development grows, so demands on representation languages and editing tools increase. For example, currently under development are probabilistic extensions for OWL, facilities for associating ontologies with problem-solving methods (PSMs) and ontology rule languages like the Semantic Web Rule Language (SWRL). SWRL proposes to extend OWL's semantics by combining it with Rule Markup Language (RuleML) so that rules can be made. Specifically, this will enable implications to be made based on the relations between concepts or objects; for example, we can specify that if John has a mother and his mother has a sister then John has an aunt. Effectively, it will extend the reasoning capabilities of OWL.

Editing tools will need to evolve quickly in order to meet these developments. Moreover, Denny's recent survey of ontology software providers revealed growing demands for higher-level abstraction of ontology language constructs to enable more powerful knowledge modelling, enhanced visual/intuitive navigation mechanisms, reasoning facilities, and methods for aligning domain and core ontologies [35].

8. Assessment

8.1 Applications and Benefits in HE/FE

The application of ontologies in learning environments is still at an early stage, but recent research hints at future directions this might take.

The Scholarly Ontologies Project from the Knowledge Media Institute, Open University, has built a prototype system called ClaiMaker which provides an infrastructure for making claims about the significance of research publications. Claims and challenges in a scholarly research field can be modelled as an "argument network", facilitating debate and interrogation:

> 'Claims' are made by making connections between ideas. Any claim is of course open to counterarguments. The connections are grounded in a discourse/counterargument ontology, which makes possible innovative services for navigating, visualising and analysing the network as it grows [36].

It is intended that such an environment would be used by students, lecturers, researchers and librarians, either individually or collaboratively, engaging in discourse about new ideas.

The Edutella Project, on the other hand, aims to provide an RDF-based metadata infrastructure for JXTA applications [37]. JXTA is "a set of open protocols that allow any connected device on the network ranging from cell phones and wireless PDAs to PCs and servers to communicate and collaborate in a P2P [Peer-to-Peer] manner" [38]. The first application will involve a metadata-enhanced P2P network for the exchange of educational resources between universities in Germany, Sweden and the US. This will allow institutions to maintain control over their own content while enabling other institutions to access it. Despite the use of different schemas by different institutions, each peer can describe its metadata information in RDF, and Edutella's querying, replication and mapping services will enable transparent access to distributed resources.

The "courseware watchdog" tool also uses an ontology to find and organise remote learning resources on the Web [39]. The ontology fulfils several roles. It:

- Supports a browsable interface to distributed resources
- Provides search terms for focused crawling, based on user interests
- Offers a way of querying annotated collections of learning resources
- Enables different “views” of discovered material (such as content, format, educational level, and so on), using clustering techniques

The watchdog is intended as a flexible way of supporting the individual needs of students and teachers in a lifelong learning environment where requirements are constantly changing.

Another example of how ontologies might be put to use in a learning context is provided by the Diogene Project [40], which is developing a Web training environment for ICT (Information and Communications Technology) professionals. Its ontology is based on the ACM Computing Classification (CCS), used for classifying books, journal articles and conference proceedings in the field of computing in a four-level subject hierarchy; however, rather than a simple hierarchy, the following relations link its concepts:

- Has Part: $HP(x, y_1 \dots y_n)$ means that concept x is composed of the concepts y_1 *to* y_n; that is to say, to learn x it is necessary to learn y_1 to y_n.
- Requires: $R(x, y)$ means that, to learn x, it is first necessary to learn y.
- Suggested Order: $SO(x, y)$ means that it is preferable to learn x and y in this order.

The relations form the basis of a dynamic course-generation process, suggesting a sequence of topics based on learners’ objectives. In the examples given in Figures 2 and 3 (above), they specify that a user wishing to study colour encoding will first be *required* to know about colour and colour systems, then he will learn about true colour, bicolour and greyscale (which are *part of* colour encoding), and it will be *suggested* that he study greyscale before bicolour. Training material from registered providers and from the Web is categorised within the ontology, and delivered to learners to explain each concept in the learning path.

The structure of Diogene’s ontology is based on the course management system used in the m-Learning project [41], which aims to deliver learning materials to young adults across Europe via mobile devices as well as via Web and television.

Other work focuses on developing ontologies describing the learning process. One example includes creating an ontology describing collaborative learning tasks, goals and actions, which can then be used to guide designers of learning environments in the process of specifying new scenarios, and in directing group discussion [42].

The potential benefits of such ontologies to further and higher education are clear:

- Students are provided with advanced browsing and searching support in their quest for relevant material on the Web. Especially where their understanding of a topic is low, they can be directed intelligently towards resources of relevance.
- All the work involved in creating an ontology can directly benefit learners by helping them to visualise and comprehend the relationships between concepts in their domain, as understood by more experienced practitioners. This can trigger “associative ways of processing, reflecting and analysing information” [3].
- Syntactically different but semantically similar resources can more easily be located.
- Information can be shared across educational applications, enabling reuse not only of learning objects but also of domain knowledge and pedagogical programmes.
- Distance learners can be provided with the intelligent and personalised support that they would otherwise miss out on. As demonstrated by the Diogene project, a well-constructed ontology can be used as part of a wider system to intelligently guide the student from one topic to another, suggesting related subjects along the way.
- Personalised courses can be generated on demand.

- Finally, Packer takes a different view of pedagogical benefits [43], promoting an ontological account of the learning process as a way of understanding and articulating the developmental changes that occur in individuals as they learn.

8.2 Risks

However, there are some difficulties involved in creating ontologies. The ontology development process can be difficult and costly: arriving at a representation of a domain requires thorough knowledge of that domain as well as an ability to think abstractly about it, choosing its boundaries, selecting which concepts to define and at what level of detail, and deciding how these concepts should be described. The more expressive the ontology, the more complex and time-consuming this task, and achieving an objective representation is next to impossible. Further, concepts will require different definitions depending on the context in which the ontology is being used. For ontologies to interoperate and data to be shared, this contextual information must be understood [9]. Efforts are now underway into developing systems that are better able to cope with uncertainty [44, 45].

Moreover, ontologies which go beyond a hierarchical representation (such as those employed in the Diogene Project) and use more complex relations to describe their domain may indeed have greater expressive power, but can be awkward to construct and for end-users to comprehend, especially when large. Research is underway into making such ontologies more easily browsable [46]. Complicated ontologies are not, however, a prerequisite of the Semantic Web.

Ontology development brings together communities from diverse backgrounds, including library science, knowledge engineering and business. This has, in the past, meant fractured development, leading to a considerable degree of overlap and reinvention, and the creation of ontologies which define the same things differently [47]. Soergel, for example, cited CYC and WordNet as examples of where “the huge intellectual capital embodied in many classification schemes and thesauri is largely ignored” [48]. Recently, however, participating communities have been making real efforts to overcome the communication barriers and to adopt a more unified approach. By coordinating their efforts, communities will be able to draw on each others’ expertise, resulting in more streamlined and better-informed ontology development.

8.3 Conclusion

Clearly, ontologies have a role to play in learning and teaching on the Web. Gateways already provide enhanced subject access via classification schemes; applications in which relationships and dependencies are more richly defined could offer more intelligent, personalised access to networks of resources. OWL already enables the specification of such ontologies, and examples are emerging of ontologies being employed to improve access to heterogeneous collections of educational content.

Moreover, the difficulties involved in ontology development, such as the problems of building representations of a domain, and the involvement of parties from diverse backgrounds, are being overcome through collaboration and focused effort.

The Semantic Web is still in its early days and applications are only just beginning to emerge. Judging from current research directions, the near future will hold greater interoperability, shareability and reusability among existing Web applications via a modular approach involving domain-specific classifications and ontologies, rather than an all-encompassing knowledge model.

9. Glossary

DDC	Dewey Decimal Classification
LCSH	Library of Congress Subject Headings
LIS	Library and Information Science
Ontology	A semantically rich vocabulary for describing a domain that can be understood by computers
OWL	Web Ontology Language
RDF	Resource Description Framework
Semantic Web	An extension to the Web in which documents are described using well-defined meanings and therefore accessed more effectively
XML	eXtensible Mark-up Language

10. References

[1] Davies, J., Fensel, D. and van Harmelen, F. (eds). *Towards the Semantic Web: Ontology-driven Knowledge Management*. Chichester: John Wiley and Sons, 2003.

[2] McGuinness, D.L. Ontologies Come of Age, in, Dieter Fensel, Jim Hendler, Henry Lieberman and Wolfgang Wahlster (eds), *Spinning the Semantic Web: Bringing the World Wide Web to its Full Potential*. MIT Press, 2002.

[3] Aroyo, L. and Dicheva, D. 'Summary of Workshop on Concepts and Ontologies in Web-based Educational Systems'. *International Conference on Computers in Education*. 3-6 December 2002, Auckland, New Zealand, http://wwwis.win.tue.nl/~laroyo/ICCE2002_Workshop/CFP.htm

[4] Devedzic, V. 'The Semantic Web – Implications for Teaching and Learning'. *International Conference on Computers in Education*. 11-14 November 2001, Seoul, Korea.

[5] Gruber, T. 'A Translation Approach to Portable Ontology Specifications'. *Knowledge Acquisition*, 5 (2), 1993.

[6] Genesereth, M. R. and Nilsson, N.J. *Logical Foundations of Artificial Intelligence*. CA: Morgan Kaufmann Publishers, 1987.

[7] Guarino, N. 'Formal Ontology and Information Systems'. In *Proceedings of Formal Ontology in Information Systems*, June 1998. Netherlands: IOS Press.

[8] Berners-Lee, T., Hendler, J. and Lassila, O. 'The Semantic Web'. *Scientific American*, 17 May 2001.

[9] Marshall, C.C. and M. Shipman, F.M. 'Which Semantic Web?' In, *Proceedings of the fourteenth ACM Conference on Hypertext and Hypermedia*. Nottingham, UK, 2003. New York: ACM Press.

[10] 'OWL Web Ontology Language Guide'. *W3C Recommendation*, 10 February 2004, http://www.w3.org/TR/owl-guide/

[11] W3C. *OWL Web Ontology Language Overview*, 10 February 2004. http://www.w3.org/TR/owl-features/

[12] Protégé 2000, http://protege.stanford.edu/

[13] CO-ODE, http://www.co-ode.org/

[14] OntoEdit, http://www.ontoknowledge.org/tools/ontoedit.shtml

[15] OilEd, http://oiled.man.ac.uk/

[16] WebODE, http://delicias.dia.fi.upm.es/webODE/

[17] Ontolingua, http://www.ksl.stanford.edu/software/ontolingua/

[18] Ontosaurus, http://www.isi.edu/isd/ontosaurus.html

[19] LinkFactory, http://www.landcglobal.com/pages/linkfactory.php

[20] [20] KAON, http://kaon.semanticweb.org/

[21] Fensel, D. and Perez, A.G. *OntoWeb deliverable 1.3: a survey on ontology tools,* 2002. http://www.ontoweb.org/download/deliverables/D13_v1-0.zip

[22] CYC, http://www.cyc.com/

[23] SENSUS, http://mozart.isi.edu:8003/sensus2/

[24] Generalised Upper Model, http://www.fb10.uni-bremen.de/anglistik/langpro/Webspace/jb/gum/index.htm

[25] Ding, Y. 'Ontology: the Enabler for the Semantic Web'. *Journal of Information Science*. 27 (6), 2001.

[26] DAML Ontology Library, http://www.daml.org/ontologies/

[27] Protégé Ontology Library, http://protege.stanford.edu/ontologies/ontologies.html

[28] OWL Ontology Library, http://protege.stanford.edu/plugins/owl/ontologies.html

[29] BUBL LINK, http://bubl.ac.uk/link/

[30] Humbul, http://www.humbul.ac.uk/

[31] ADAM, http://adam.ac.uk/

[32] HILT, http://hilt.cdlr.strath.ac.uk/

[33] Renardus, http://www.renardus.org/

[34] Doerr, M., Hunter, J. and Lagoze, C. 'Towards a Core Ontology for Information Integration'. *Journal of Digital Information*. 4 (1), 2003.

[35] Denny, M. 'Ontology Tools Survey, Revisited'. *XML.com*, 14 July 2004, http://www.xml.com/lpt/a/2004/07/14/onto.html

[36] Scholarly Ontologies Project, http://kmi.open.ac.uk/projects/scholonto/

[37] Edutella, http://edutella.jxta.org/

[38] JXTA, http://www.jxta.org/

[39] Tane, J., Schmitz, C., Stumme, G., Staab, S. and Studer, R. 'The Courseware Watchdog: an Ontology-based tool for finding and organizing learning material'. In, *Fachtagung Mobiles Lernen und Forschen*. Universität Kassel, 2003.

[40] Diogene Project, http://www.diogene.org/

[41] m-Learning Project, http://www.m-learning.org/

[42] Barros, B., Verdejo, M.F. Read, T. and Mizoguchi, R. 'Applications of a Collaborative Learning Ontology'. *Mexican International Conference on Artificial Intelligence (MICAI 2002)*, 2002.

[43] Packer, M. 'The Ontology of Learning'. *Annual Meetings of the American Educational Research Association,* Montreal, 1999, http://www.mathcs.duq.edu/~packer/Pubs/AERA99.dir/AERA99A.html

[44] Straccia, U. 'Uncertainty and description logic programs: a proposal for expressing rules and uncertainty on top of ontologies'. *Technical Report*. Institute of Information Science and Technologies (ISTI), 15 April 2004, http://faure.iei.pi.cnr.it/~straccia/projects/ Material/Papers/TR-UDLLPSemWeb/TR-UDLLPSemWeb.pdf

[45] Holi, M. and Hyvonen, E. 'A method for modeling uncertainty in semantic web taxonomies'. *Thirteenth International World Wide Web Conference (WWW2004)*. New York City, 17-22 May 2004, http://www.cs.helsinki.fi/u/eahyvone/publications/www2004/ontoBayes_poster.pdf

[46] Aussenac-Gilles, N. and Mothe, J. 'Ontologies as Background Knowledge to Explore Document Collections'. *Recherché d'Information Assistee par Ordinateur (RIAO 2004)*. Avignon, France, 26-28 April 2004.

[47] Noy, N.F. and Hafner, C.D. 'The State of the Art in Ontology Design: a Survey and Comparative Review'. *AI Magazine*, Fall 1997.

[48] Soergel, D. 'The Rise of Ontologies or the Reinvention of Classification'. *Journal of the American Society for Information Science*. 50 (2) 1999.