

E-LEARNING IN THE SEMANTIC AGE

Matthias Palmér
Ambjörn Naeve
Mikael Nilsson
The KMR (Knowledge Management Research) group
CID (Centre for user oriented IT Design)
NADA (Numerical Analysis and Computing Science)
KTH (Royal Institute of Technology)
100 44 Stockholm, Sweden
[amb|mini|matthias@nada.kth.se]

ABSTRACT

Today educational technologies are reaching a state that allows interoperability and reuse of learning resources. The underlying techniques rely heavily on the standards movement for metadata representation. On top of this, monolithic reference platforms are being developed with the aim to ease application development. However, we do not think this approach is flexible enough to embrace future learning techniques. In contrast, we suggest a learning framework where services can be developed and exchanged between as well as within systems. A fundamental part of this framework is the semantic layer, which builds on the structure of the Semantic Web. Hence we do not regard metadata as something 'objective' that you have to download from some central server. On the contrary, metadata should be allowed to consist of subjective views of resources that are distributed and shared in contexts that can evolve dynamically. In support of such requirements, our learning framework consists of a combination of semantic web techniques and peer-to-peer services for search, retrieval, publication, replication and mapping of metadata.

KEYWORDS

Learning Framework, Semantic Web, Conceptual Web, interoperability, peer-to-peer, concept browser, Conzilla, Edutella.

1. INTRODUCTION

Computer Based Instruction (CBI) and Intelligent Tutoring systems (ITS) are two important areas within the field of IT-supported learning, which in recent years have drifted somewhat apart. While CBI has taken a more pragmatic engineering approach, ITS has developed more into a field of advanced research. With the success of the web, their respective agendas have come to focus on issues of interoperability and reuse [35]. This has been acknowledged by several key international players including IEEE [33], IMS [34], ARIADNE [37] and AICC [38], resulting in a multitude of standards that can be used for building interoperable learning systems. Fortunately, many of these standards complement each other or seem to be converging towards a common form, which allows applications to use them together. In order to ease this integration, several standards have been collected and harmonized and often a reference platform has been given as proof of concept. One example is SCORM [35] by ADL [36], which provides a reference model that includes collected/developed standards for content aggregation as well as a run-time environment. A more recent example is OKI [46], which is a project initiated by MIT¹ in order to provide all of their courses online. They plan to deliver an architecture and some APIs² for basic services. Furthermore, IMS provides a set of basic standards but they also have an enterprise standard (message oriented approach) which specifies how an LMS³ should exchange messages regarding management of courses and students, such as tracking, assessment etc. Blackboard [48] is an example of a commercial LMS that uses IMS enterprise.

These attempts at building learning platforms for interoperability are mainly targeted to ease the need of LMSes for adaptation to standards, but as a consequence, learners can be expected to gain more freedom. For example, the goal of SCORM is to provide a reference model for content that is durable (survives system changes), interoperable (between systems), accessible (indexed and searchable) and reusable (able to be modified in different contexts and by different tools). This will hopefully allow students to move more freely between LMSs and even to combine several services from different LMSs.

1.1. Some problems with current approaches

We see at least three problems with current approaches. First, most providers of content have large monolithic systems where adaptation to e.g. SCORM will not significantly change the underlying teacher-learner model. Students will be presented with material in a (maybe personalized) context often leading up to some form of (standardized) test. New and more interesting methods for learning - such as techniques for collaboration, annotation, conceptual modeling etc. - will not profit from such adaptation⁴.

Second, even though monolithic, closed or proprietary systems will be able to exchange learning resources, course-like structures and keep track of students with the help of those standards, they will need to go through yet another process of adaptation to the next big batch of agreements on learning technologies, such as e.g. profiling and tracking of student performance.

¹ Massachusetts Institute of Technology.

² Application Programming Interfaces.

³ Learning Management System.

⁴ Our learner-centric educational architecture called a *Knowledge Manifold*, which supports inquiry-based and customizable forms of learning is discussed in [6], [7] as well as in a separate article [9] in these proceedings.

Third, the current perspective on metadata is too limited. Anyone who has something to say about a learning resource should be able to do so. This includes learners, teachers and content contributors such as authors and providers. Communicating this metadata is equally important as it can help, direct or encourage others to actively participate and learn. Proposed solutions, such as [35], will result in learning resources (and their metadata) that will reappear in different versions and formats rather than dynamically evolve and improve. Important synergy effects will be missing.

1.2. Proposed solutions

We will begin by discussing the third problem, and then, in response to the first two problems, we will introduce the notion of a learning framework.

To share information such as SCORM metadata and CSFs⁵ is simple, or at least straightforward, when the standards are well designed and documented. However, the techniques mentioned above will need resources to be identified and talked about in a distributed way that is similar to what is done on the web⁶. In order to enable this, we need an extendable (machine-understandable) "semantic language" for expressing the semantics of anything that is identifiable. To be extendable means allowing new schemas to be dynamically added as well as allowing tools and learners to interact in new ways not yet conceived within a standards body. RDF⁷ [39] and the vision of the Semantic Web [43] provide a good starting point for such a semantic language. RDF allows meaning, i.e. semantics, to be expressed in a distributed manner similar to how the web works. What the semantic web can do for e-learning is discussed in section 3.2.

We also need mechanisms that enable searching for, locating and exchanging semantic information expressed in this language. These mechanisms, together with the semantic language, constitute the basic ingredients of a learning framework. SCORMs 'reference model' is a first step in this direction, since it describes how content should be described and exchanged in a general way. We think that a learning framework should consist of layers that include a semantic (web) layer as well as a service layer on top of this. Certainly more layers will be needed, e.g. a communication layer and an application layer. This will be briefly discussed in section 2.3.

A learning framework should support applications in working more flexibly with learning technologies and should provide a platform similar to the Java platform enterprise edition. Hence, we think that a future LMS should consist of a set of independent but cooperating non-monolithic services/applications that are available in public pools. Applications/services can then be built on top of existing services - providing a better interface - or by defining new independent services. Some examples of such applications/services developed by the KMR group at CID will be presented in chapter 4.

⁵ Content Structure Formats.

⁶ This is beyond the current scope of SCORM.

⁷ Resource Description Framework.

2. INTRODUCING OUR LEARNING FRAMEWORK

There is a growing recognition for the need of a learning framework⁸, which captures relevant standards in a layered fashion in order to allow for interoperability and simplified application development. We will now present an alternative, which allows a more capable and flexible design based on techniques from the semantic web community⁹.

2.1. Supporting standards

The strength of standards lies in their capacity to set the stage for interoperability between LMS:es as well as between applications within a system. Such interoperability can be achieved in several ways. Let us examine three different ways to interoperate¹⁰:

- Integration at the software level, with components talking directly to one another using APIs.
- Message-based integration, with components exchanging messages in plain text using a protocol such as SOAP¹¹.
- Data-based integration using XML files or records in relational databases to share information.

The first alternative presents several problems. It is rather inflexible and risky, since it forces you to have a tight coupling between your system and services. The second alternative, messaging, is better, since it allows more flexibility. However, the semantics of the messages needs to be understood, at least in part, by all involved parties. In some sense this common understanding relies on the third alternative, data integration, which pre-supposes a common understanding of the schemas to use. The basic standards, e.g. IMS metadata [14] and content packaging, are mainly concerned with defining data models (schemas) and operate on the level of data integration.

In the Learning Framework presented here, we focus on this view and think of standards as either schemas in isolation or schemas together with the services that make use of them in order to accomplish different tasks. Messages between or inside systems should be seen as transmitting facts about resources (content, courses, students etc.) conforming to certain schemas that specific services have declared an interest in. So in this case the distinction between the message integration and the data integration approach is only a matter of context, i.e. it depends on how a service treats the information, e.g. like facts about resources, descriptions of processes or a piece of information in an ongoing dialogue between two services.

One example could be a student performing a search based on a query. If this search is successful and she wants to save it for future use, what should be stored - the end results or the initiating query? The results are certainly a set of facts and can therefore without problems be stored as something similar to bookmarks. However, if the query is stored as just another result in such a bookmark system, whenever it is retrieved, it should be treated as a process.

In order to complicate the example, the student might want to have another *representation* of the successful query in her personal bookmarks environment. This

⁸ Sometimes called a 'platform' or a 'model'.

⁹ See [1], [11] and [43].

¹⁰ As described by Mark Norton in his speak at the IMS symposium in Ottawa in August 2001, reported by Scott Wilson in [20].

¹¹ Simple Object Application Protocol.

could be achieved by introducing a rule where the pre-condition consists of the query itself and the post-condition consists of the desired *representation* of the query. Here we see something quite interesting, the query schema is reused in the rule schema, which corresponds to piping the search services into the rule matching service. This illustrates a situation where services - capable of treating certain schemas - can reuse other services in order to avoid duplication of work¹². A reference implementation of the Learning Framework should provide a minimal set of services supporting the most common schemas.

2.2. Layers and services

Computer science is constantly changing. Individual techniques and even complete paradigms can become obsolete during a very short period of time. Naturally, systems that have not been "designed for change" will not survive very long. A common approach to designing for change is to divide the system into layers, which communicate only via well-defined interfaces. In a closed layered architecture each layer can interface with only the immediate bordering layers, while in an open layered architecture any layer can interface with all layers beneath it. This is very similar to the design approach of computer networks¹³. We think that a learning framework should be layered but with well-defined interfaces only for the lowest (physical and transport) layers. The semantic layer, which comes next, should serve as the "middle ground" for the different sub-systems in the service and applications layers above.

In [5] there is a discussion of how layered information models should achieve interoperability between applications, and a special reference model, Information Model Interoperability (IMI), is defined. This model contains a syntax layer (serialization, e.g. XML/RDF), an object layer (facts about resources) and a semantic layer (the schemas used). In a knowledge representation system there is no well-defined border between facts and schemas¹⁴ [5]. This is also noted in [20]. Hence our semantic layer will actually embrace both the object layer and the semantic layer defined in IMI. However, the distinction is still there, and the schemas will probably be the preferable way for services to specify the domains within which they work.

2.3. Suggested layers

Figure 1 depicts some of the suggested layers in our Learning Framework¹⁵. We will now perform a bottom up walk-through and discuss some of the techniques occurring in each layer.

¹² Reuse of schemas saves a lot of work as well, which is beneficial for tool builders as well as for standards bodies.

¹³ For an overview, see standard textbooks such as [17].

¹⁴ In fact, the ULM (Unified Language Modeling) methodology that we have developed takes advantage of this fact. See [7] and [10] for further details.

¹⁵ This is by no means a complete picture.

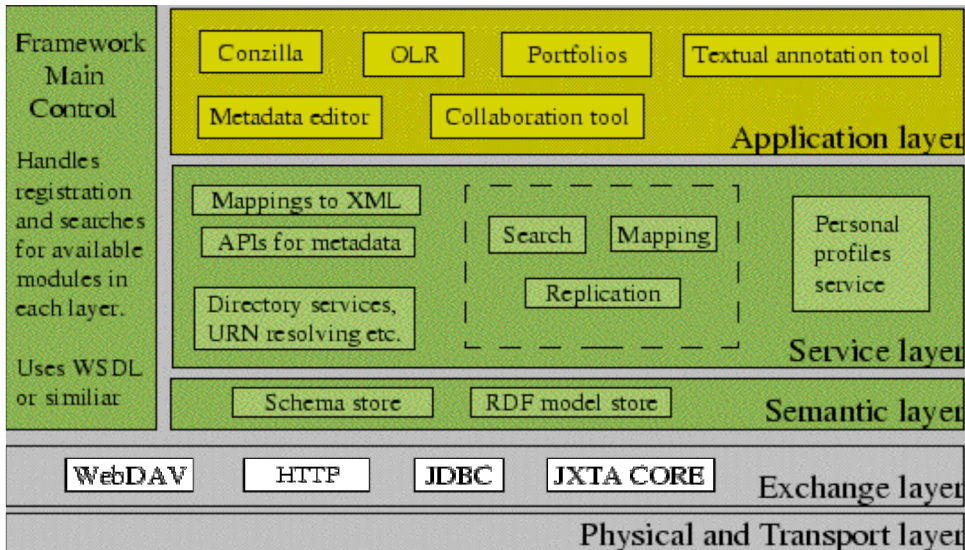


Figure 1. Some suggested layers of our Learning Framework.

Physical and Transport layer: The bits have to be transported over some medium, e.g. cable, infrared beam etc. On top of this there is a stack of protocols such as e.g. TCP/IP.

Exchange layer: For retrieving resources http is often enough. WebDAV¹⁶ is an extension, which allows intelligible write/access/versioning of the resources as well as of their metadata. Peer-to-peer systems, such as JXTA¹⁷ [42], allow decentralized repositories, which lowers the threshold for publication of content and metadata. Also database connectivity such as JDBC¹⁸ fits nicely into this layer.

Semantic layer: It is somewhat unclear to what extent this layer exists besides from serving as a temporary store for models and as a middle-ground between services as well as between services and the exchange layer. What is clear though is that the semantic layer should be compatible with the Semantic Web [43] and hence use RDF/RDFS¹⁹ [39], [40], and all possible extensions in terms of schemas, ontologies etc.

Service layer: Hopefully a multitude of services will soon be crafted in order to solve problems such as profiling and resolving of resources. Several more advanced services will need active cooperation from other peers in order to work properly, e.g. search for metadata, definition of mappings between schemas and publication of content/metadata. Such cooperation can take place via a peer-to-peer system such as JXTA²⁰. Since not all services and applications will be able to use RDF and schemas to express themselves, there will always be a need for export to XML or special APIs. Some services will deal with this explicitly, probably as part of a 'pipe' of

¹⁶ Web Distributed Authoring and Versioning.

¹⁷ Juxtapose.

¹⁸ Java Data Base Connection.

¹⁹ Resource Description Framework / Resource Description Framework Schema.

²⁰ The KMR group at CID is actively participating in an international collaborative JXTA-based project called *Edutella* [3], [31], where several such services will be developed, see section 4.2. In Figure 1, services that will be part of *Edutella* are drawn within the dashed box.

services. A very important example is provided by 'unintelligent' light-weight clients that need to work with simpler representations of metadata.

Application layer: This is the topmost layer where man and machine meet. Hence it will require a multitude of different capabilities with user-friendly interfaces. See chapter 4 for an introduction to some of our existing and future applications. Probably there should be a sub-layer where highly reusable applications belong, e.g. so-called plugins.

Framework main control: This layer constitutes the main control structure where all modules register their capabilities via some language such as WSDL²¹ [47] or maybe OIL²² [41]. Typically an application registers which schemas, protocols, APIs and/or mime-types that it is capable of exposing to the user via some UI²³. If it has no UI it is probably a service and not an application.

3. SUPPORTING SEMANTICS IN A LEARNING FRAMEWORK

The Learning Framework suggested above is open to a multitude of new services. In order to be effective, it needs a powerful language for expressing facts about resources and schemas that will allow machines as well as humans to understand how these facts are related without relying on heuristics. Moreover, there is a need for expressing facts about remote (identifiable) objects without accessing remote data stores.

If the semantic web did not already exist, we would probably have to invent something similar in order to be able to realize our Learning Framework. But the semantic web is more than just a necessary prerequisite for this framework. It allows the definition of a large variety of innovative services that are not supported by the structure of the present web. In the next sections we will give a brief introduction to the Semantic Web and its capabilities within a learning context.

3.1. The Semantic Web

The present web is a highly intertwined graph, constructed mainly for human consumption, where heuristics are needed in order to digest the meaning (semantics) of a resource. Hence the construction of agents that could help to automatically sort out information is hard. This does not mean that the existing web lacks semantics. However, it is presently not accessible in a standardized machine-readable way.

The *Semantic Web* is a W3C initiative [43] that will increase the power of the existing web because it will adhere to the basic principles of the latter: "anyone can say anything about anything" [1]. The emerging Semantic Web will not replace the existing web. Instead it will become a complement that will allow the creation of more powerful tools for searching and making inferences²⁴.

The current web consists of resources (often documents) that are interconnected via their identifiers (URIs²⁵). The Semantic Web will add semantics to these resources by adorning them with *properties* expressed in RDF/RDFS [39], [40]. A property is a *triple* that consists of *subject*, *object* and *predicate*. The subject is the

²¹ Web Services Description Language.

²² Ontology Inference Layer.

²³ User Interface.

²⁴ See [1] for an excellent article on the promises of the Semantic Web.

²⁵ Uniform (or Universal) Resource Identifier.

resource that the property applies to, the object is another resource or an opaque string saying something about the subject and the predicate is the type of the property. Adding properties to a resource can be done by anyone anywhere, i.e. the properties do not need to be collected into one document or conceived of at the same time as the document is created. Instead, the properties are spread out and there is no easy way of collecting them. This will present a challenge for the search engines of tomorrow.

3.2. What we can do with the Semantic Web in a learning context

Traditionally, a learning resource can be made accessible on the web and then (possibly) updated by the provider. Using a learning framework equipped with semantic web techniques we may also do the following to a resource:

Describe it: Since a resource can have uses outside the domains foreseen by the provider, any given description (metadata instance) is bound to be incomplete. Because of the distributed structure of RDF, a description can be expanded or new descriptions following new formats (schemas) can be added. This allows for creative uses of content in new and unforeseen ways. Hence, one of the most important features of the current web - the fact that anyone can link anything to anything - has been carried over into RDF.

Certify it: There is no reason why only big organizations should certify learning resources. Individuals, such as e.g. teachers, may want to certify certain content as a quality learning resource that is well suited for specific learning tasks. Handling this kind of certification will be an important part of the Semantic Web as pointed out in [1].

Annotate it: Everything that has an identifier (URI) can be annotated. There are already attempts under way in this direction. Annotea [44] is a project where annotations in RDF format are created locally or on a server. The annotations apply to HTML or XML documents and are automatically fetched and incorporated into web pages via a special feature in the experimental browser Amaya [45].

Extend it: Structured content (typically in XML format) is becoming common. Successive editing (diffs) can be done via special RDF-schemas allowing private, group consensus or author-specific versions of a common base document. The versioning history will be a tree with known and unknown branches, which can be traversed with the help of the next generation of versioning tools.

Use it everywhere: RDF is application independent. Since the metadata is expressed in a standard format, which is independent of the underlying schemas, even simplistic applications can understand parts of complex RDF graphs. If your favorite tool does not support the corresponding schemas, it can at least present them in a rough graph, table or whatever standard form it has for describing resources and their properties. If more advanced processing software is available (such as logic engines), more advanced treatment of the RDF descriptions is possible.

And more: Apart from these uses, you can invent new schemas describing structures, personalization, results from monitoring and tracking, processes and interactions that can enrich the learning environment in various ways.

4. APPLICATIONS AND SERVICES - OUR CONTRIBUTIONS

The Semantic Web is a huge, interconnected, distributed graph that is not appropriate for human consumption. We have proposed a structure called the *Conceptual Web* [11], which will live on top of the Semantic Web and which will contribute to making the latter accessible to humans in a more appealing and understandable way. The Conceptual Web will be structured as a *Knowledge Manifold* [6], [9] and will be explored through a *concept browser* [7], [10]. This is a new type of knowledge management tool that lets you navigate different contexts by traversing the contextual neighborhoods of each concept and view the content of a concept and/or concept relation under a dynamically configurable set of aspects [15].

4.1. Conzilla

Conzilla [12], [13] is a first prototype of a concept browser, which we have developed within the KMR group at CID during the last 3 years²⁶. Using a concept browser, there are mainly two modes of exploration:

Browsing is performed on prefabricated structures such as context maps, which typically have a layout and a presentation that conforms to some visual language such as UML²⁷. Such maps typically represents a selection of what is relevant in a certain context. The linking between different contexts for the same concept is performed similarly to how it is done in html. More advanced structures are also available such as neighborhoods of maps and inspection of metadata. A concept browser maintains a clear separation of content from context and allows you to link each concept to a set of content components that can be viewed in an external viewer such as an ordinary web browser. The reader is referred to [10] for details.

Querying is performed when you do not find what you want in the prefabricated structures and therefore need to actively search for knowledge. Both the query and the results may be highly structured graphs that need user interfaces that are appealing and easy to use. With the help of our Learning Framework we plan to extend the querying functionality of *Conzilla* beyond the basic filtering capabilities that are available today [15]. This will be achieved through various forms of peer-to-peer services that will be developed within the *Edutella* project described in section 4.2.

4.2. Edutella

RDF is a descriptive language and although it is distributed it needs application/services in order to perform useful tasks. Actually the tasks of finding and distributing information require separate solutions in order to avoid well-known access points such as search engines. These have a serious drawback, since the information contained in the Semantic Web may change rapidly leaving the search engines of today far behind with their three months update-cycles²⁸. The true power of the Semantic Web will not be unleashed until there are ways to discover facts via

²⁶ *Conzilla* is presently being developed as an open source project at SourceForge [22] and can be downloaded from that site.

²⁷ Unified Modeling Language [16] is an industry standard for object-oriented modeling. In fact, when drawing context maps, we use our own modification of UML called ULM (Unified Language Modeling) [10], which is more concept-oriented and has been especially designed in order to depict how we speak about things.

²⁸ See [19] for a discussion on wide respectively deep search.

queries that not only access remote repositories but also combine and logically process the results from several sources.

The KMR group at CID is participating in an international collaboration project called PADLR²⁹, whose driving vision is a learning web infrastructure which will make it possible to exchange/annotate/organize and personalize/navigate/use/reuse modular learning resources, supporting a variety of courses, disciplines and universities. Within this project, we are collaborating with research groups at the universities of Uppsala [23], Stanford [26], Hannover [27] and Karlsruhe [28] in order to develop Edutella [3], [31], an infrastructure and a search service for a peer-to-peer network that will facilitate the exchange of educational resources.

Edutella, which will be a set of services implemented within the JXTA system, is aiming (among other things) to solve the search problems described above. The envisioned services will include searching, mapping and replication. Searches will be routed to anyone who has registered a matching answering capability. Mapping will enable translation between schemas. This will allow very flexible reuse of information, since an application will not need to adapt to competing or more capable schemas because these schemas can be mapped to something that the application already understands. There will be no closed formats. Replication will allow metadata about learning resources to be spread across the web, which will simplify the discovery of the corresponding resources.

4.3. Virtual Workspace Environment

VWE is a distributed Learning Management System, which is designed to support the construction of customizable learning environments by enabling the composition of learning resources³⁰. It has many features in common with the "framework main control" sketched in section 2.3. In fact, VWE is a small configurable operating system that can run in a web browser, which allows you to access your own learning environment from everywhere.

5. CONCLUSIONS AND FUTURE WORK

Learning, as well as other human activities, cannot be confined within well defined boundaries such as course systems. In a broader perspective, we want to enhance existing and future environments and make them more interesting from a learning perspective. As stated in [9] we believe that "nobody can teach you anything, a good teacher can inspire you to learn". Hence we acknowledge the crucial importance of learner motivation. Moreover, the learning environment has to support trust building and rich forms of communication between teacher and learner as well as between learners³¹. In order to be powerful, the environment must be inspiring and trigger curiosity for the learning task³².

²⁹ Personalized Access to Distributed Learning Resources [30]. This project is funded by the Wallenberg Global Learning Network [29].

³⁰ VWE has been developed under the coordination of Fredrik Paulsson of the KMR group at CID and the Swedish National Agency for Education. He is also chair of the Swedish subcommittee of the ISO/IEC standardization body for e-learning technology, ISO/IEC JTC1 SC36 [32].

³¹ An interesting experiment in this direction is described and analyzed in [4].

³² A virtual reality based shared distributed 3D-environment called *CyberMath* [18], which aspires to fulfill these requirements has been developed at CID. See www.nada.kth.se/~gustavt/cybermath for quicktime movie demos of some learning experiences in the CyberMath system.

In summary, we want to use the power and the flexibility of the Semantic Web within our Learning Framework in order to develop tools, standards and environments that support the following areas:

- **Content management**, allowing dynamic creation of courses as well as distributed annotations by both teachers and learners. The supporting technologies include: content creation with distributed versioning capabilities, personal portfolios³³ and personal profiles.
- **Knowledge navigation**, focusing on the organization of ideas/concepts and their relations in a conceptually clear context³⁴. Separation of content and context allows you to maintain an overview of the learning landscape. Navigation is carried out between contexts³⁵ that are connected through contextual neighborhoods. The Conceptual Web [11] could form the basis for a global knowledge project, intended to evolve and capture more and more of the accumulated human knowledge.
- **Experience-orientated environments**³⁶, where objects can be annotated by anyone. If you want, you can share your annotations with other people - maybe fellow learners, teachers or just people that are present for the experience. You can communicate with them directly, launch a discussion or why not perform or engage in some activity together.

The possibilities are limited only by your imagination.

REFERENCES

Papers and books

- [1] Berners-Lee, T., *Semantic Web Roadmap*, www.w3.org/DesignIssues/Semantic.html.
- [2] Berners-Lee, T. & Hendler, J. & Lassila, O., *The Semantic Web*, Scientific American, May 2001, www.scientificamerican.com/2001/0501issue/0501berners-lee.html.
- [3] Decker, S. & Manning, C. & Naeve, A. & Nejdil, W. & Risch, T. & Studer, R., *Edurella - An Infrastructure for the exchange of Educational Media*, Part of the PADLR [30] proposal to WGLN [29].
- [4] Knudsen, C. & Naeve, A., *Presence Production in a Distributed Shared Virtual Environment for Exploring Mathematics*, Proceedings of the 8th International Conference on Advanced Computer Systems (ACS 2001), Szcecin, Poland, Oct. 17-19, 2001.
- [5] Melnik, S. & Decker, S., *A Layered Approach to Information Modeling and Interoperability on the Web*, Database group, Stanford University, Sept. 2000, <http://www-db.stanford.edu/~melnik/pub/sw00>.
- [6] Naeve, A., *The Garden of Knowledge as a Knowledge Manifold - A Conceptual Framework for Computer Supported Subjective Education*, CID-17, TRITA-

³³ The KMR group at CID, in collaboration with Uppsala Learning Lab [24], have developed a special portfolio system for this task.

³⁴ This is (part of) the aspiration of a concept browser [10], as exemplified by the Conzilla program [22].

³⁵ e.g. visualized as context maps expressed in ULM [10].

³⁶ e.g. shared distributed 3D-environments like CyberMath.

- NA-D9708, Department of Numerical Analysis and Computing Science, KTH, Stockholm, 1997, http://cid.nada.kth.se/sv/pdf/cid_17.pdf.
- [7] Naeve, A., *Conceptual Navigation and Multiple Scale Narration in a Knowledge Manifold*, CID-52, TRITA-NA-D9910, Department of Numerical Analysis and Computing Science, KTH, 1999. http://cid.nada.kth.se/sv/pdf/cid_52.pdf.
- [8] Naeve, A., *The Work of Ambjörn Naeve within the Field of Mathematics Educational Reform*, CID-110, TRITA-NA-D0104, KTH, Stockholm, 2001, www.amt.kth.se/projekt/matemagi/mathematics_educational_reform.doc.
- [9] Naeve, A., *The Knowledge Manifold - an Educational Architecture that supports Inquiry-Based Customizable Forms of E-Learning*, Proceedings of the 2nd European Web-Based Learning Environment Conference (WBLE 2001), Lund, Sweden, Oct. 24-26, 2001.
- [10] Naeve, A., *The Concept Browser - a New Form of Knowledge Management Tool*, Proceedings of the 2nd European Web-Based Learning Environment Conference (WBLE 2001), Lund, Sweden, Oct. 24-26, 2001.
- [11] Naeve, A. & Nilsson, M. & Palmér, M., *The Conceptual Web - our research vision*, Proceedings of the first Semantic Web Working Symposium, Stanford, July 2001, www.semanticweb.org/SWWS/program/position/soi-nilsson.pdf.
- [12] Nilsson, M. & Palmér, M., *Conzilla - Towards a Concept Browser*, CID-53, TRITA-NA-D9911, KTH, 1999, http://cid.nada.kth.se/sv/pdf/cid_53.pdf.
- [13] Nilsson, M., *The Conzilla Design - The definitive reference*, CID/NADA/KTH, 2000, <http://conzilla.sourceforge.net/doc/conzilla-design/conzilla-design.html>.
- [14] Nilsson, M. (ed.), *IMS/LOM-RDF binding*, www.imsproject.org/rdf/index.html.
- [15] Pettersson, D., *Aspect Filtering as a Tool to Support Conceptual Exploration and Presentation*, TRITA-NA-E0079, CID/NADA/KTH, Dec. 2000.
- [16] Rumbaugh, J. & Jacobsson, I. & Booch, G., *The Unified Modeling Language Reference Manual*, Addison Wesley Longman Inc., 1999.
- [17] Tanenbaum, A. S., *Computer Networks*, Prentice-Hall, 3rd ed., 1997.
- [18] Taxén G. & Naeve, A., *CyberMath - Exploring Open Issues in VR-based Learning*, SIGGRAPH 2001 Educators Program, In SIGGRAPH 2001 Conference Abstracts and Applications, pp. 49-51, www.nada.kth.se/~gustavt/cybermath.
- [19] Waterhouse, S., *JXTA Search: Distributed Search for Distributed Networks*, <http://search.jxta.org/JXTAsearch.pdf>.
- [20] Wilson, S., *The next big thing? - Three architectural frameworks for learning technologies*, CETIS, www.cetis.ac.uk/content/20010828163808/viewArticle.

Relevant web sites

- [21] CID (Centre for user-oriented IT Design): <http://cid.nada.kth.se/il>.
- [22] Conzilla development: <http://conzilla.sourceforge.net>.
- [23] UDBL (Uppsala Data Base Laboratory): www.dis.uu.se/~udbl.
- [24] ULL (Uppsala Learning Lab): www.ull.uu.se.
- [25] SweLL (Swedish Learning Lab): www.swedishlearninglab.org.
- [26] SNLPG (Stanford Natural Language Processing Group): <http://nlp.stanford.edu>.
- [27] KBS-Hannover: www.kbs.uni-hannover.de.
- [28] AIFB-Karlsruhe: www.aifb.uni-karlsruhe.de/Forschungsgruppen/index_en.html.
- [29] WGLN (Wallenberg Global Learning Network): www.wgln.org.
- [30] PADLR (Personalized Access to Distributed Learning Repositories) proposal to WGLN, Granted March 2001: www.learninglab.de/pdf/L3S_padlr_17.pdf.
- [31] Edutella: <http://edutella.jxta.org>.

- [32] ISO/IEC JTC1 SC36: <http://jtc1sc36.org>.
- [33] IEEE (Institute of Electrical and Electronics Engineers):
<http://standards.ieee.org>.
- [34] IMS (Instructional Management Systems): www.imsproject.org.
- [35] SCORM (Sharable Content Object Model) Reference Model 1.1, ADL,
Department of Defense, USA, Jan. 2001, www.adlnet.org/scorm/scorm.cfm.
- [36] ADL (Advanced Distributed Learning): www.adlnet.org.
- [37] ARIADNE (Alliance of Remote Instructional Authoring and Distribution
Networks for Europe): <http://ariadne.unil.ch>.
- [38] AICC (Aviations Industry CBT Committee): www.aicc.org.
- [39] RDF (Resource Description Framework): www.w3.org/RDF.
- [40] RDF Schema: www.w3.org/TR/2000/CR-rdf-schema-20000327.
- [41] OIL (Ontology Inference Layer): www.ontoknowledge.org/oil.
- [42] JXTA (The Juxtapose project): www.jxta.org.
- [43] Semantic Web initiative: www.SemanticWeb.org.
- [44] Annotea: www.w3.org/2001/Annotea.
- [45] Amaya: www.w3.org/Amaya.
- [46] OKI (Open Knowledge Initiative): <http://web.mit.edu/oki>.
- [47] WSDL (Web Services Description Language): www.w3.org/TR/wsdl.
- [48] Blackboard: www.blackboard.com.