# Semantic Technologies to Support Teaching and Learning with Cases: Challenges and Opportunities

Patrick Carmichael and Agustina Martinez Garcia

University of Cambridge, Cambridge, CB2 1SB
{wpc22, am857}@cam.ac.uk

**Abstract.** Visions of a 'semantic web' and the technologies, standards and services associated with it have the potential to support and enhance teaching and learning. As yet, this potential has not been practically demonstrated in ways that are accessible to teachers and students. 'Ensemble: Semantic Technologies to Support the Teaching and Learning of Case Based Learning' is currently exploring the potential of semantic technologies to support and enhance teaching and learning in fields in higher education where knowledge is complex, changing or contested, and where as a result case based learning is the pedagogy of choice. This paper describes how the wide range of case based learning approaches has informed the selection, development and deployment of semantic technologies, and identifies a number of key challenges and areas for development which would enable more widespread adoption of semantic technologies by teachers and students.

**Keywords:** semantic technologies; collaborative design; digital repositories; triplestores; visualization; case based learning; higher education

## 1 Introduction

The Semantic Web is conceptualized as "an extension of the current Web in which information is given well-defined meaning, better enabling computers and people to work in cooperation … data on the Web [is] defined and linked in a way that it can be used for more effective discovery, automation, integration, and reuse across various applications" [1]. Even though progress towards the realization of a wholly semantic 'web' has been slow [2] and in educational settings most applications have been in the area of resource discovery and description [3,4], technologies that have been developed to enable it have the potential to enhance and extend existing web technologies and to offer new opportunities for teaching and learning [5,6].

Semantic web technologies allow the integration of content, both user-generated and from digital repositories, web services and 'non-semantic' data such as 'legacy' databases. They offer users advanced search tools and a range of representations and visualisations of data and also support collaborative functions such as reviewing, rating and annotation. With the development of more accessible semantic authoring and data visualization tools, such as those produced by the SIMILE project at MIT [7], the opportunities for teachers and students to be more closely involved in the design and realization of semantic applications have been greatly increased.

The Ensemble project[1] ("Ensemble: Semantic Technologies for the Enhancement of Case Based Learning") is exploring the potential of semantic technologies to support and enhance teaching and learning in a variety of settings in higher education: specifically advanced undergraduate courses at the University of Cambridge and postgraduate and professional courses at City University, London. The work of the project is focused on settings in which complexity, contestation or rapid change makes some kind of case based learning the pedagogy of choice. As well as substantive research settings in which learning with cases is the focus of attention, a series of pilot projects and technical demonstrators has informed the work of the man project and served to engage potential participants. Members of the project team are also undertaking more wide ranging work on digital repositories, knowledge representation in different fields, visualization of complex data and the role of semantic technologies in student assessment.

This paper reflects on several areas of the project's work, highlighting some distinctive aspects of its view on the development of teaching and learning applications; and identifies six areas of current and future development for the project, each with broader implications for developers of semantic technologies and the web applications they enable.

## 2 Learning with Cases: a range of practice

The project has undertaken extensive and detailed studies of a range of teaching and learning environments in which cases are used in teaching and learning, observing teaching and learning activities (using still photography, audio and video where possible); collecting course documentation, student work and other resources; and interviewing staff and students both about general pedagogical practices and about the specific activities observed in 'sense-making' discussions. Courses being researched are diverse, including plant sciences and archaeology (at Cambridge) and journalism and marine operations and management (at City University London).[2]

While there are some elements in common across the different research settings, such as a commitment to engage students with authentic data and to support the development of disciplinary practices, the nature, role and scope of cases, and the practices and discourse that accompany them, vary widely. In some of the settings being researched, what was observed was an approach in which cases are designed and resources provided by the teacher in order to illustrate particular issues, problems;

---

[1] "*Ensemble: Semantic Technologies for the Enhancement of Case Based Learning*" is a project of the ESRC-EPSRC Technology Enhanced Learning Programme, funded under Grant RES-139-25-0403. Patrick Carmichael at the Centre for Applied Research in Educational Technologies, University of Cambridge is director of the project. Other team members include: Uma Patel, Lawrence Solkin and David Bolton (City University London); Keith Johnstone, Agustina Martinez (University of Cambridge); Rob Walker (University of East Anglia); Louise Corti (University of Essex); Michael Tscholl, Fran Tracy and Katy Jordan (jointly held posts at City University London and the University of Cambridge). Full details are available on the project website at: http://www.ensemble.ac.uk

[2] These are described in detail in another paper within this symposium: Tscholl, M., Tracy, F. and Carmichael, P. (2009) 'Case Methods, Pedagogical Innovation and Semantic Technologies' SemHE Symposium, EC-TEL: 4th European Conference on Technology Enhanced Learning Nice, France, 30 September 2009

in others, however, students had greater opportunities to interpret problems in different ways, draw on a wider range of paper and online resources, and establish the boundaries of cases themselves. In one setting (Archaeology at the University of Cambridge), teachers drew on their own experience and introduced resources and physical artifacts and cases emerged through a rich classroom discourse.

This variation in turn means that the opportunities for the development and integration of semantic technologies into these teaching and learning environments differ across disciplines and settings, and the need to engage teachers and students in collaborative design activities is greater than if the intention were to support a generic or ideal model of case based learning. In some of the settings, web technologies and online resources were already used, while in others their role was limited to supporting out-of-class research activities. The opportunities for semantic technologies to enhance teaching and learning activities had to be carefully balanced with a sensitivity to existing practices, both those where technologies (such as specialist software applications, databases and visualization tools) were already used, and those where the introduction of such technologies might reduce the richness of the students' experiences and the opportunities for new learning opportunities and outcomes to emerge.

## 3 Design and Development: patterns of participation

The approach taken by the project has been to use a range of cooperative and participatory design approaches, involving rapid prototyping as a means of developing what Anderson and Crocca call 'crucial artifacts' [8] with applications in the teaching and learning environments described above. Critical to this prototyping process has been the availability of a set of software applications and services provided by the SIMILE project at MIT, including the Babel data conversion utility and the Exhibit visualization toolkit [7]. These allow browser-based application based on small, exemplary data sets to be produced and then used as a focus for further discussions and user testing.

Several important issues have emerged from the early work of the project in this area, each of which has implications for design processes and for the visions, specifications and ultimately the teaching and learning applications that might emerge from those processes:

- A recognition of the fact that project participants may play multiple roles in research settings (teacher, researcher, learner, manager) and therefore that they may have multiple, and potentially conflicting, motivations and perceptions as to what is 'of worth' in the applications that might be built. In one pilot study, for example, discussions with the key teacher participant highlighted their desire to support students in overcoming specific conceptual difficulties, while at the same time helping them to understand how elements of the course 'fitted together'. Furthermore, in role as teaching coordinator, they saw the application as having the potential to support curricular management, review and innovation.
- A view of students as 'expert learners': meaning that they might have expert standing in the field or discipline but not be experienced in formal learning settings ("experts, learning"), or that they might be successful as students in the formal settings ("experts *at* learning"). The recognition of these student perspectives

alongside those of teachers and course designers means that the nature, content and implementation of the 'crucial artifact' are problematized.

- A recognition of the need to reconcile the opportunities offered by advanced learning technologies to support emergent, personal responses to the curriculum, context or case, with disciplinary/professional norms and/or the demands of assessment systems. This is an issue for educational innovations more generally [9], with the introduction of new technologies highlighting the tensions caused by a lack of 'constructive alignment' between advanced learning outcomes and assessment systems in particular [10].

Each of these issues would be relevant even if no new technologies were being designed and developed. In the settings studied there are already discourses concerned with the role of learning from and with cases (where this itself represents a relatively new departure) or about the impact of changing student demographics or how the demands of professional or academic practice might be addressed in undergraduate or postgraduate courses. The nature of teaching and learning with cases is necessarily emergent and discursive, and the lack of pre-ordained learning outcomes means that characteristically these discourses accompany not just curriculum planning, but continue throughout and beyond the teaching and learning activities themselves. This means that rather than teachers and users being engaged in one set of activities which then inform a set of 'requirements' which become the focus of a separate discourse amongst programmers, software developers and user interface designers, these latter groups are inducted into a ongoing, generative discourse drawing on subject knowledge, technological development and pedagogical issues. They will of course have discourses and practices of their own, but these are intrinsically linked into dynamic and fluid pedagogical ones.

Out of these discourses, a series of priority areas for the research and development of semantic technologies has been identified. These represent 'boundary objects' [11] within the project, as they are areas of common interest across the different groups involved, although each of which may understand them in different ways and relate them to their own specific pedagogical or technological practices. The majority of the participants amongst the research team and in the research settings as a legitimate area for further project activity would recognize each of the following.

This list, which is divided into two parts (the first concerned with access to semantically rich data resources and the second with the support of situated teaching and learning activities) is not offered as a general 'development agenda' for semantic web activity, but rather as a model of how semantic technologies will be developed and deployed in the particular settings of the Ensemble project and its partner settings. At the same time we are aware that many of the issues we raise here will have relevance or resonance with other developers of learning technologies and their integration into teaching and learning environments.

## 4 Areas for Development: Teacher and Student Access to Semantic Data, Metadata and Contextual Information

The first three areas for research and development are concerned with identifying and gaining access to online resources, converting these into appropriate formats, and then aggregating data and metadata: so that they can then be incorporated into teaching and learning environments. In all of the research settings in which we have worked, these have been explicitly identified as prerequisites to further discourses, design activities and potential implementation.

### 4.1 Access to Data, Metadata, Ontologies and Service Descriptions

Much of the online data that might be integrated into teaching and learning environments through the use of semantic technologies exists in legacy databases, data-driven websites and in commercial environments, such as publishers' web environments. While some data providers offer resources in formats oriented towards their use in semantic web applications, these remain in the minority. One may be fortunate enough to work in a domain where projects like DBpedia [12] or Freebase [13] have already begun to aggregate data sources into appropriate format, but in most cases there is work to do in converting data. This has been the case as researchers and participants in project research settings develop ideas of how semantic technologies might enhance teaching and learning: for example, in plant sciences there are some existing ontologies and data-driven websites which also allow downloads of their datasets as XML, but for the most part data either has to be extracted via existing web interfaces or even 'scraped' from websites where this is permitted by the authors or originators.

What has also emerged from work across research settings is the need for a wide range of contextual information around both quantitative and qualitative data sets. As well as the consistent application of metadata and terms drawn from established taxonomies and ontologies, there is also a need for a further level of metadata associated with resources related to the context of its collection, sampling and selection protocols and information about other associated methodologies [14]. Initiatives such as the Data Description Initiative (DDI)[15] have provided useful frameworks for qualitative data. Even apparently unproblematic quantitative data is made more useful in semantic applications when accompanied by contextual and methodological information. For example: in a technical demonstrator built to support the exploration of patterns of plant distribution [16], datasets from the Global Biodiversity Information Facility (GBIF) [17] were used as a source of observational data. While these data are all presented in standard formats, the basis on which they were collected and selected varies (as is to be expected in a community-contributed collection); some data reflect intensive data collection with fine grained GPS locations in restricted geographical areas; while others have greater coverage but are limited to reporting the appearance of plants in 1-kilometre or 10-kilometre 'quadrats'. Attempts to aggregate or reason across such varied data could potentially lead to misleading representations being generated and fallacious conclusions being reached.

Ultimately, the success of attempts to integrate existing data sources into semantic portals and applications may be dependent upon the development of machine-readable 'service descriptions' for data, analogous with the existing web service description (WSDL) system [18], allowing data sources to be 'self describing' to semantic agents and aggregators.

### 4.2 The Role of 'RDFizers'

At the same time as there is a need for rich metadata to describe resources and the context of their generation, the requirement to be able to link heterogeneous data and potentially to support machine reasoning across these means that there has been a significant effort, across the semantic technology developer and user communities, to develop conversion utilities. This means that emerging semantic web standards such as RDF and N3 have come to be more than simply means of expressing metadata or relationships between existing resources. While RDF was originally conceived of as a language for describing information resources so that these descriptions can be exchanged without loss of meaning [19], it can also be used as a standardized way of describing or even *expressing* heterogeneous data in differing formats that can be used across applications.

A number of projects, community efforts and commercial organizations provide 'RDFizers' and this extended role for RDF becomes apparent when looking at the range of these tools now available. Some are designed to convert metadata from established resource description systems (such as MARC or BibTex records); others represent responses to the availability of ready but non-semantic metadata (such as the tags associated with images in public 'Flickr' libraries), but others such as SIMILE's 'Babel' Tool [20] and the University of Maryland's 'RDF123' [21] represent means of converting *any* structured data in spreadsheet, CSV or TSV format into RDF (and other formats) regardless of whether these are associated with a schema or ontology.

Just as XML is a meta-language rather than a language, RDF is a very general model (in fact, even more general than XML) and 'converting to RDF' can, as outlined above, mean everything from applying consistent schemas, ontologies and namespaces, to a one-time conversion of a spreadsheet with no semantic information other than that expressed in the column headings. The analogy with XML still holds: in XML-based environments where standard DTD's are used and enforced (again, the DDI is a good example), documents can be created, shared and interpreted in a consistent way, preserving the initial intended semantics [22], but this is the exception rather than the rule, and most XML data and metadata are not tightly bound to a DTD or schema.

There is a temptation, particularly when building rapid prototypes, to develop one-off RDF-izers oriented towards a single data source and a single web application: an RDF-izer capable of converting the GBIF datasets mentioned above would be an example. In the same way, the development of networks of data in RDF, referring to stable URIs and associated with consistent ontologies represents a significantly greater effort. What this means for the Ensemble project is that, in each of the research settings, the available data sources need to be carefully assessed and the process by which they should be integrated will need to involve more robust, scalable and generalised approaches to their integration, rather than simple conversion processes determined by current applications.

### 4.3 Close Coupling of Repositories and Triplestores

Following on from the issues of data acquisition and conversion, the project has also had to address issues of data storage, aggregation and presentation to users. The key features and functionalities of several digital repositories were analyzed, as a result of which the Fedora digital repository [23] was selected as the basis for storage of data and metadata. The main difference between Fedora and other digital repositories is that Fedora provides a semantic web-ready solution in terms of storing not only digital objects in different formats, directly inside the repository or by managing references to external resources stored in external servers or databases, but also because it is closely coupled with the Mulgara Semantic Triplestore.

The flexibility of the Fedora object model allows multiple streams of metadata to be associated with an object, so that, for example, a digital asset can be associated with descriptive metadata in different formats; contextual information of the type described above in section 4.1; and user-generated annotations and commentary. As well as experimenting with the use of Fedora not just as a provider of resources (images, texts, datasets) and rich metadata, we have also used it as the as a data source in which for 'preconversion' formats like plain text, or CSV but also, and most significantly for converted 'native' XML/RDF data which can be ingested into a semantic triplestore (in this case Mulgara) and can be streamed directly to other applications or queried using SPARQL from a dedicated SPARQL endpoint.

This has involved the development of custom applications which first query object metadata to discover what data streams are available, and, if suitable data are found, then incorporate these into the triplestore to be aggregated with other data sources, reasoned across or exposed via endpoints or other web applications. In our example of plant distribution data, this means that what exists in the Mulgara Triplestore is not only metadata *about* available datasets and relevant ontologies, but RDF representation *of the data themselves*. With this step-change in provision, the applications that can be constructed can be transformed from being primarily concerned with semantically enhanced *resource identification and retrieval*, to *exposing data for manipulation* directly by teachers and students. In applications concerned with plant distribution, students are therefore able to interact with data obtained from GBIF (as set out in section 4.1), converted to RDF (section 4.2) and stored, along with metadata and ontologies in the Triplestore, exposed via a SPARQL endpoint and then presented using the SIMILE 'Exhibit' toolkit as faceted post-search applications, maps and timelines [24]. This is a critically important affordance for teachers and students involved in case based learning, where access to 'authentic' data is often a distinctive and important aspect of learning activities and a focus of learning discourses.

## 5. Query Endpoints and the concept of the 'Data Portal'

What has been described in the foregoing section is a set of prerequisites for the implementation of semantic technologies in teaching and learning environments, involving:

- Access to 'self-describing data' via a prototype service description mechanism
- 'Future-proof' rather than 'one-time-only' conversion of these data and metadata into appropriate formats, with consistent metadata and sufficient ontology development to permit data linking
- Availability of metadata and, where appropriate and available, qualitative and quantitative data, for use by teachers and students in learning from and with cases

With these general requirements in mind, the priority for the project has been to develop the idea of a SPARQL endpoint into the basis of a richer teaching and learning resource: a configurable 'data portal' appropriate to the teaching and learning environment in which it is deployed. Such a portal would make available pre-converted data sources, together with metadata and contextual information; ontologies where these exist (also accompanied by information about source and scope); and access to other providers of supplementary data and metadata through API's and web services.

In the context of the plant distribution example mentioned previously, such a portal would allow teachers and students to gain access to the GBIF data sets and other open-access data, pre-converted and accompanied by information about the generation, reliability, granularity and currency of observations; taxonomic frameworks; definitions and rules for the interpretation of these; access to 'live' data feeds and news sources; supplementary information such as that provided through the Geonames web services [25]; and federated search of relevant publications. These elements, together with user generated content (such as students' own field observations uploaded and converted – and this is where the lightweight, general conversion utilities have an important role to play), represent some of the resources which teachers and students would then be able to mobilize as they explore the scope and boundaries of cases, develop narratives, formulate problems and test hypotheses.

## 6. Further Issues: Supporting Situated Knowledge and Practice

The provision of data sources (whether directly, through a digital repository or a more sophisticated and evolving 'data portal') is only one aspect of how semantic technologies could enhance teaching and learning around cases. While this would equip teachers and students with existing cases on which to draw, or with data and literature which they could used in the construction and reconstruction of cases, the work of the project with teacher and student participants across research settings suggests that further opportunities exist.

When students draw on, discuss or construct cases, the interpretations are frequently tentative and emergent, as they establish the boundaries and scope of the case, what data and reasoning is appropriate and even 'of *what* is this a case'? Data may be incorporated into multiple cases; in archaeology, for example, a particular artifact being mobilized in relation to cases of social life in the past, the technological development, trade or religious belief. Alternatively, it may form part of a case illustrating past or current archaeological practices (such as excavation or preservation); or be introduced as a case of particular trends in curation, interpretation or public engagement with the discipline. One important aspect of learning with cases, both in established "CBL" case based learning and in the broader range of practice, outlined in section 2 above, is the importance of personal and reflective narrative.

All of these practices are highly situated and suggest that the introduction of any semantic technologies into teaching and learning environment will involve negotiation of new discourses and practices. Initial analysis together with a series of pilot projects and technical demonstrators has allowed us to identify three further areas for semantic technology development which would contribute both to these emergent discourses and practices, and to the development of software applications with the potential to enhance teaching and learning. As in section 4, these are issues with relevance and applications that extend beyond the research settings involved in the Ensemble project, and we would suggest that these are relevant to other initiatives and activities that are seeking to establish what the semantic web for education might become.

### 6.1 Support for multiple, local and 'soft' ontologies

In addition to established taxonomies and ontologies which form part of what Polanyi calls the 'mutual authority' of 'self-coordinating' disciplines [26], additional ways of organizing knowledge are evident in teaching and learning environments. The most obvious example of this are 'curricula', organizational mediation of which may play a far more significant role in shaping student understanding than any ontological commitments of disciplinary experts. Even when an area of study is well theorized and there appear to be well-established formal taxonomies and ontologies there are, as Eraut reminds us, 'many meanings of theory and practice' [27], which may need to be represented in teaching and learning environments, particularly when 'experts, learning' with extensive 'craft knowledge' are engaging with cases which may be theorized in different ways, or when teachers are mediating student engagement with complex content.

One example with which was identified in an early pilot provides a good example: items in a directory of online learning resources, designed to support the learning of mathematics by engineering students, were described using a local 'mathematics for engineers' ontology derived the collective experience of teachers and students involved in the course concerned. While this included mathematical concepts and terms, this subset of mathematical vocabulary was organized in a way that reflected the conceptual and curricular perspectives of the course in question, rather than a complete and formal mathematical framework [28]. So the 'data portal', and possibly the applications which draw on its contents may need to be able to represent not only formal taxonomies and ontologies, but also those which represent pedagogical perspectives on a knowledge domain (as with 'maths for engineers') and 'soft' or 'working' ontologies which reflect a variety of student perspectives, and which may change as their learning informs ontological elaborations and transformations.

### 6.2 Visualization tools appropriate to specialist teaching and learning

A further area for research and development is that of appropriate modes of representation and visualization of data, metadata and other information. Research in different settings has revealed the extent to which specific visual representations are seen as part of the distinctive practice of a discipline or sub-discipline. So, for example, in archaeology, there are established ways of representing vessels, or stratigraphical cross-sections, and while photographic images are increasingly used, these are seen as supplements rather than replacements for line drawings which allow the highly nuanced representation of expert 'ways of seeing'.

In the course of another of the project pilot studies, an interactive timeline of plant evolution [29] was designed and developed to present undergraduate biology students with an interactive and extensible version of a well-known graphical image in a textbook by Niklas [30] which shows the relative number of species of plants in major groups over geological time. What was clear from discussions with teachers and students was the importance of preserving the core visual representation of the Niklas diagram: even though it proved possible to integrate a wide range of heterogeneous qualitative and quantitative data into the web application, the distinctive stacked line graph remains a central feature.

This has important implications for the representation of the highly complex data that might be made available to teachers and learners through interaction with semantic technologies. Currently many of the visual representations that mediate access to semantic resources draw on established information science metaphors (browsable lists, timelines, thumbnail galleries or object viewers) or use the semantic web's own 'directed graph' underpinnings to present visualizations of social or conceptual networks. Close collaborative design and development activities, of the kind that took place in the development of the plant evolution timeline, will be necessary in order to develop visualization tools that not only offer teachers and learners the affordances of semantic technologies that support and enhance learning with cases, but that also align with and extend existing practice in the area of representation and visualization.

### 6.3 User generated content, annotations and social semantic web

A final area of research and development is the integration of semantic technologies with the concepts, practices and technologies of 'Web 2.0' or the 'social web': including blogs, wikis, recommender systems and person-to-person social networks [31, 32]. This is an emerging area of research and development across all aspects of semantic technology development, but in the context of higher education in particular, where 'Web 2.0' have transformed not only teaching and learning practices but also student expectations, the potential overlap between emerging semantic technologies and 'social' software is particularly apparent.

For the Ensemble project, some aspects of the 'social semantic web' clearly have relevance in those teaching and learning settings in which students work together or participate in extended knowledge-building activities. Students of journalism, for example, work collaboratively to build up networks of contacts and information sources on which they subsequently draw when taking part in group or individual assignments; and small groups of plant science students already use a wiki to record progress in a case based learning activity.

At the same time, the processes of narrative development and 'case building' that have been highlighted across project research settings seem to point to the value of 'blog-like' or 'wiki-like' environments with semantic technologies either embedded or closely coupled. The idea of the 'data portal' outlined in section 5 would then be extended by the addition of the 'read-write' functions familiar from 'Web 2.0' applications.

## 7. Emergent Technologies, Emergent Understandings

What we have described in this paper is a continuing process in which teaching and learning practices and the discourses of case construction and reconstruction that accompany them have been explored through participatory activities and close engagement of the research team and research participants in different settings.

Rather than there being distinct and time-constrained processes of requirements gathering, design and development, technological discourses have emerged or been introduced and are better seen as being 'interwoven' with pedagogical ones upon which they draw and which they, in turn, inform and advance.

These braided discourses around the nature, scope and representation of cases has informed the concept of the data portal as a general model with which participants in different settings can engage, elaborating and configuring it to reflect their existing and emergent practice. These more situated practices will be the focus of further discourse and co-design activity which focus on the way learners understand, mobilize, construct and reconstruct cases in specific teaching and learning environments and settings.

In the same way, there is no assumption that once new semantic applications are developed, they can simply be deployed or 'nursed' into the established practice of teachers and learners. Discourses between teachers and learners, researchers and developers will necessarily continue as the meanings, impacts and, potentially, creative and unexpected uses of the technologies emerge and themselves change future teaching and learning practice and discourse. We would contend that this is necessary for any group of technologies as complex and expansive as those associated with the semantic web. Indeed, recognition of the need for this continuing and open discourse may be an important prerequisite for teacher and student engagement, both with specific applications that are enabled by semantic technologies, and the broader vision of the semantic web.

## References


1 Semantic Web Activity Statement (2001) *W3C Technology and Society domain Semantic Web Home* Online at: http://www.w3.org/2001/sw/Activity.

2 Shadbolt, N., Berners-Lee, T. and Hall, W. (2006) The Semantic Web Revisited *IEEE Intelligent Systems* 21(3) pp. 96-101.

3 Devedzic, V. (2006). *Semantic Web and Education*. Springer: Integrated Series in Information Systems.

4 Fensel, D., Hendler, J. and Liederman, H. (2003) *Spinning the Semantic Web: Bringing the World Wide Web to Its Full Potential* Cambridge, MA: The MIT Press.

5 Allert, H. (2004) Coherent Social Systems for Learning: an approach for contextualised and community-centred metadata *Journal of Interactive Media in Education* 2004(2) Online at: http://www-jime.open.ac.uk/2004/2.

6 Koper, R. (2004). Use of the Semantic Web to Solve Some Basic Problems in Education: Increase Flexible, Distributed Lifelong Learning, Decrease Teacher's Workload. *Journal of Interactive Media in Education*, 2004 (6) Online at: http://www-jime.open.ac.uk/2004/6

7 Simile Project. MIT. Main site: http://simile.mit.edu

8 Anderson, W.L. and Crocca, W.T. (1993) Engineering practice and co-development of product prototypes Communications of the ACM 39 pp. 39-46.

9 Trigwell, K. and P. Ashwin (2006). "An exploratory study of situated conceptions of learning and learning environments." *Higher Education* 51 pp. 243-258.

10 Biggs, J. (1996) Enhancing teaching through constructive alignment *Higher Education* 32 pp. 347-364

11 Bowker, G. and Star, S.L. (2000). Sorting Things Out: Classification and its Consequences. Cambridge, MA: MIT Press.

12 Auer, S et al. (2007) DBpedia: A Nucleus for a Web of Open Data. In Aberer et al. (Eds.): The Semantic Web, 6th International Semantic Web Conference, 2nd Asian Semantic Web

Conference, Busan, Korea, November, 2007 (Lecture Notes in Computer Science 4825 Springer 2007, ISBN 978–3-540–76297–3)
13 http://www.freebase.com
14 Carmichael, P. (2008) 'Secondary Qualitative Analysis with Internet Resources' in : Blank, G., Fielding, N. and Lee, R. (eds.) *The Handbook of Online Research Methods* (London: Sage)
15 http://www.ddialliance.org
16 Martinez. A, Carmichael, P. and Corti, L. (2009) 'The Ensemble Project: Using Fedora to Support the Development of the Semantic Web for Education' Paper presented at the Open Repositories Conference (Fedora User Group Meeting), Atlanta, GA, 20 June 2009.
17 http://www.gbif.org
18 http://www.w3.org/TR/wsdl
19 Manola, F and Miller, E. (eds.) (2004) RDF Primer W3C Recommendation 10 Feb 2004 Online at: http://www.w3.org/TR/REC-rdf-syntax/
20 http://simile.mit.edu/babel/
21 Han, L. et al., (2008) RDF123: from Spreadsheets to RDF Proceedings of Seventh International Semantic Web Conference, Karlsruhe, October 2008
22 Erdmann, M. and Studer, R. (2000) 'How to Structure and Access XML Documents With Ontologies' Online at: http://www.aifb.uni-karlsruhe.de/~mer/Pubs/xml-dke-final.ps
23 Lagoze, C., Payette, S., Shin, E. and Wilper, C. (2006) Fedora: an architecture for complex objects and their relationships *International Journal on Digital Libraries*, 6 (2) pp.124-138.
24 Martinez. A, Carmichael, P. and Corti, L. (2009) 'The Ensemble Project: Using Fedora to Support the Development of the Semantic Web for Education' Paper presented at the Open Repositories Conference (Fedora User Group Meeting), Atlanta, GA, 20 May 2009.
25 http://www.geonames.org/
26. Polanyi, M. (1969) 'Knowing and Being' in: *Knowing and Being: Essays by Michael Polanyi* (ed. Grene, M.) University of Chicago Press: Chicago
27. Eraut, M. (2003) 'The Many Meanings of Theory and Practice' Learning in Health and Social Care. 2(2) pp.61-65
28. Carmichael, P. and Watson, J. (2009) 'Maths for Engineers: using Semantic Technologies to Support Student Transitions into Higher Education' Poster, Higher Education Association Conference, Manchester, July 2009
29. Peart, N. and Roberts, B. (2009) 'A Timeline of Plant Evolution - supporting learning about complex relationships' Paper presented at Computer Assisted Learning Conference, Brighton, April 2009. The timeline is available online at http://www.ensemble.ac.uk/projects/plantsci/
30. Niklas, K.J. (1997) The Evolutionary Biology of Plants Chicago, IL: University of Chicago Press.
31. Gruber, T. (2006). 'Where the Social Web Meets the Semantic Web' Presentation at the 5th International Semantic Web Conference, November 7, 2006
32. Mika, P. (2007) *Social networks and the Semantic Web* Springer: Semantic Web and Beyond: Computing for Human Experience Series