Supporting Ontology Driven Document Enrichment within Communities of Practice

John Domingue, Enrico Motta, Simon Buckingham Shum, Maria Vargas-Vera, Yannis Kalfoglou

Knowledge Media Institute

International Centre for Distance Learning

Nick Farnes

The Open University Walton Hall, Milton Keynes, MK7 6AA, UK {j.b.domingue; e.motta; s.buckingham.shum; m.vargas-vera; y.kalfoglou; n.farnes}@open.ac.uk

ABSTRACT

Formative work by Lave and Wenger has articulated how *practices* emerge through the interplay of informal processes with symbolic codifications and artifacts. In this paper, we describe how ontologies can serve as symbolic tools within a community of practice supporting communication and knowledge sharing. We show that when a community's perspective on an issue is stable, it opens the possibility for introducing knowledge services, based on an ontology co-constructed by knowledge engineers with stakeholders. Using a case study we describe our approach, ontology driven document enrichment, looking at how ontology construction and population can be supported by web based technologies.

Keywords

Ontology, Semantic Web, Communities of Practice, Knowledge Management.

INTRODUCTION

Formative work by Lave and Wenger [13, 22] has articulated the nature of the *practices* from which the term *community of practice* derives its name. Practices emerge through the interplay of informal processes with symbolic codifications and artifacts:

...Such a concept of practice includes both the explicit and the tacit. It includes what is said and what is left unsaid; what is represented and what is assumed. It includes language, tools, documents, images, symbols, well-defined roles, specified criteria, codified procedures, regulations, and contracts that various practices make explicit for a variety of purposes. But it also includes all the implicit relations, tacit conventions, subtle cues, untold rules of thumb, recognizable intuitions,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

K-CAP'01, October 22-23, 2001, Victoria, British Columbia, Canada. Copyright 2001 ACM 1-58113-XXX-X/01/0010...\$5.00.

specific perceptions, well-tuned sensitivities, embodied understandings, underlying assumptions, and shared world views. Most of these may never be articulated, yet they are unmistakable signs of membership in communities of practice and are crucial to the success of their enterprise. ([22], p. 47)

In this paper, we describe how ontologies [9] can serve as symbolic tools within a community of practice. We show that when a community's perspective on an issue is stable (i.e. there is reasonable consensus), it opens the possibility for introducing knowledge services, based on an ontology co-constructed by knowledge engineers with stakeholders. The ontology reflects a "shared world view", codifying "well-defined roles", "specified criteria" and "codified procedures." Throughout, we regard representations such as ontologies as *boundary objects* [2] whose role is to support communication and negotiation over meaning between stakeholders within and across communities of practice.

Once an ontology has been constructed a population phase uses the ontology to describe web documents from a communal viewpoint. Two key questions which arise in this type of enterprise and that we address in this paper are: who develops the ontology? and how is the ontology population phase supported?

We believe that knowledge engineers are crucial in the ontology development phase. The main reason for this choice is that a careful design of the ontology is crucial to ensure the success of any particular document enrichment initiative. The ontology specifies the selected communal viewpoint, circumscribes the range of phenomena we want to deal with and defines the terminology used to acquire domain knowledge. In our experience small errors/inconsistencies in any of these aspects can make the difference between success and failure. Moreover, ontology design requires specialist skills which are normally not possessed by the members of our target user communities.

Our approach is to develop the ontology using a participatory design methodology. The ontology is developed during a series of face-to-face meetings between knowledge engineers, who are concerned with issues such

as representational consistency and completeness, and a representative group of the target community.

In contrast it is essential that ontological enrichment occurs without the aid of knowledge engineers. Unless enriched web resources are a "living archive" the resultant services will soon fall into disuse. In describing the APECKS personal ontology server Tennison and Shadbolt [20] make a case for "living ontologies".

In the rest of this paper we shall illustrate our approach, which we term *ontology driven document enrichment* [16], using a case study. We start by outlining the domain, the architecture of the application and one of the knowledge services that we created. We then describe the design of the ontology and four ways in which we support the ontology population process. Related work is briefly summarized before ending with some conclusions.

CASE STUDY AN OBSERVATORY ON LIFELONG LEARNING INITIATIVES Case Study Background

In its Green Paper, 'The Learning Age', the UK Government set out its vision of 'a learning society in which everyone, from whatever background, routinely expects to learn and upgrade their skills throughout life.' One of the significant steps carried out by the UK Government to fulfil this vision was the creation of the University for Industry (Ufi) in the autumn of 2000. The overall goal for Ufi was to provide flexible learning packages which would improve the quality of life of individuals and to boost business competitiveness.

Promoting and supporting lifelong learning is a very difficult activity which requires knowledge of a number of disparate research areas including learning theory, organisation science and sociology. For the Ufi to be successful associated researchers and policy makers would need to discover and disseminate good practice on lifelong learning. It was decided that the main supporting mechanism for this would be a Web portal, termed the (available **Observatory** National at www.lifelonglearning.ac.uk), which was setup in the early part of 2000. By the time the Ufi was launched the observatory contained a number of resources including a bulletin board and a web based newsletter. The main resource was a 'Good Practice' database which held several hundred hand-coded summaries of articles describing lifelong learning initiatives. Although the database entries were highly regarded the text based search mechanisms provided a poor method of accessing relevant items.

Our goal in this project was to provide a semantic query service for lifelong learning researchers and policy makers who wanted to analyse relevant case studies, and for organisations that required help in understanding their learning needs.

Approach and Overall Design

The semantic query service was constructed collaboratively by knowledge engineers at the Knowledge Media Institute (KMi) within the Open University (OU), lifelong learning researchers at the International Centre for Distance Learning (ICDL) also at the OU and a number of external lifelong learning researchers. The lifelong learning researchers specified a number of questions that the observatory should be able to answer. The questions were categorised into three main themes deemed important by the lifelong learning research community. Each theme contained three or four sub-themes. The themes were:

- Widening participation,
- Organisational change, and
- Funding.

The specified questions were relatively broad and high level. For example, one of the questions associated with widening participation is "What techniques are needed to target the needs of socially excluded groups?" and one of the questions associated with organisational changes is "What strategies appear most effective in attracting SMEs to learning?".

The main concepts and relations within the themes and questions were used as the basis for an initial observatory ontology. The ontology was then expanded over a period of four months so that the formulated questions could be answered whilst ensuring that any new concepts and relations conformed to the view of the lifelong learning researchers.

The ICDL researchers then populated the ontology with instances which reflected the knowledge content of the learning initiatives in the Good Practice database. In this paper we describe how we supported these researchers in their population task.

Architecture

The overall architecture of the system is shown in Figure 1. At the centre of the architecture is a knowledge server whose main role is to retrieve appropriate learning initiatives from the database from end-user queries. The main components of the server are as follows:

- *LispWeb* a customised HTTP server [18] which offers a library of high-level Lisp functions to dynamically generate HTML pages.
- *WebOnto Server* WebOnto [3], composed of a central server and a Java based client, enables users to collaboratively browse and edit knowledge models over the web.
- OCML An operational knowledge modelling language [15], which provides the underlying representation for our ontologies and knowledge models.
- Observatory Library a set of knowledge models which includes the observatory ontology used to index the learning initiatives in the good practice database.

Connected to the central server are:

- The Good Practice Database a database containing several hundred summaries of documented examples of lifelong learning.
- Named Entity Recognizer this uses the Marmot and Badger systems from Riloff [17] in combination with

a regular expression matcher to support the automatic creation of OCML entities from text in web pages.

- *WebOnto Client* a Java based client to the WebOnto server.
- Semantic Search Service a service for retrieving learning initiatives from high level queries.



Figure 1. The architecture of the Observatory.

In contrast with other approaches to semantic annotation we decouple the knowledge structures from the web resources. This architecture allows us to provide multiple knowledge services, possibly for different communities of practice, over the same set of web documents. For example, a community of graphic designers may be interested in the typography and layout of a set of web pages whereas experienced website developers may be interested in the structure of the underlying HTML code. Another feature of this architecture is that the interfaces are directly connected to the ontology – there is no intermediate web crawling or compilation phase.

An Semantic Search Service

The semantic search service is designed to be easy-to-use by non-IT specialists and to provide answers to policy level questions. Figure 2 shows a screen snapshot of a web interface, constructed in FlashTM, for finding learning initiatives according to the type of funder or the characteristics of the targeted learning community. In the figure the user is asking for a government funded learning initiative which involved a socially excluded community.

n	Aniaca	er Marchmant Observatory		10.0
	Welcome to th Using the menus below y learning initiatives accord learning	e Marchmont Ober ou can find entries it ing to the type of fur community targeted	vatory. In our database of oder or the type of	1
	Funder - Tree	Gaver	ment-Organization	
Learni	ng Communities - Seciely	-Excluded 📕 Yes		
Qu	ery -			_
	Funder -	Type =	Government-Organization	
	Learning communities -	Socially-Excluded •	Tep	
		OFT		_

Figure 2. A screen snapshot showing the query interface asking for a government funded learning initiative which involved a socially excluded learning community.

Solution 9	hackney-stamford-learning-community					
earning-community2 hackney-stamford-learning-community	hackney-stamford-learning-community	olution 9				
		learning-community2 <u>hac</u>	ey-stamford-lear	aming-comm	nunity	
learning-initiative1 hackney-learning-initiative (query rationale)Take me to t	ckney-learning-initiative (query rationale)Take me to the	learning-initiative1 hackn	learning-initiativ	ve (querv ra	tionale)Take	me to the
relevant initiative. (related initiative rationale)	ated initiative rationale)	elevant initiative. (related	tiative rationale))	monde / r dite	mo to mo

Figure 3. A screen snapshot showing the results of the query in figure 2.

赉Lo	ois Qu	Jery ()	utpu	t - Netscape									×
Eile	<u>E</u> dit	⊻iew	<u>G</u> o	<u>Communicator</u>	<u>H</u> elp								
The hac	<u>hacl</u> kney	kney- -learn	stam ing-i	Expl	anati community ne learning g	on of has a low	f Re educations hackne	esult onal <u>leve</u> ey-stamf	S	initiat arning	ive -comr	nunity.	
6	D =			Document	t: Done				0 .₽,	d P	EA)	L	11.

Figure 4. The explanation generated for the query formulated in figure 2.

The query is run in OCML on the knowledge server. A set of rules link OCML knowledge items to relevant learning initiatives within the good practice database. Figure 3 shows the 9^{th} (of 11) solutions. Each solution contains links to a knowledge item, a related learning initiative and links to an explanation of why they were returned. The explanation, shown in figure 4, describes why the target learning community, the members of the Stamford housing estate, were considered to be socially excluded.

Slot Name	Documentation	Value Type
Has-title	The title of the initiative.	A string.
Has-location	The location of the initiative. This includes information on the social geography of the area.	A learning related location.
Has-initiative-date	The starting date for the initiative.	An integer representing the year (used within the existing database).
Has-rationale The underlying rationale for the initiative.		A rationale for learning.
Has-funder The funding organisation or person.		Either an organisation or person.
Other-involved-parties	Organisations, individual people and communities which take part in the initiative.	Either an organisation, generic- organisation, person or community.
Has-learner	The target audience for the initiative.	A learning-community.
Has-deliverable	The tangible results of the initiative.	A document, technology or organisation (a project may create a new organisation).

Table 1. The definition of the learning-initiative class.

Ontology Design

As we outlined earlier in this paper, there were several constraints which had to be satisfied when creating the observatory ontology. The ontology had to characterise the domain such that a) the types of questions posed by the lifelong learning policy makers and researchers could be answered, b) there was a mapping to the existing database of learning initiatives, and c) the characterisation conformed to the viewpoint of the researchers.

We should emphasise the importance of the last constraint. It was important that all of the 'observatory team' understood and had ownership of the ontology. Also as outlined in [5] in their analysis of the KA^2 initiative, and in [11] in their description of a SHOE case study, ontology development and representing specific resources are intertwined activities.

The conceptual design of the ontology was developed in a series of weekly meetings involving the whole observatory team. A number of the meetings included external policy makers and lifelong learning researchers the end users of the observatory. Once an initial version of the ontology had been implemented in WebOnto a sample population phase followed. In the early part of this phase the knowledge engineers and populators collaboratively coded 10 practices in the database. Coding difficulties would either result in immediate changes to the ontology or be logged and changed later. The populators then coded a further 20 practices on their own reporting problems by phone or email. Additionally, the team continued to meet face-toface weekly to discuss problems and changes to the ontology. These discussions would invariably result in changes to the ontology and occasionally in the addition of new tools. WebOnto's architecture meant that any changes to the ontology (or to WebOnto itself) were immediately available to the populators.

Because the domain, the intersection of learning and social policy, was relatively broad we created and reused a number of higher level ontologies. Figure 5 shows the structure of

the relevant portion of our library. The arrows indicate that an ontology *uses* its parent ontology (i.e. inherits all of the OCML entities). The observatory knowledge base currently indexes several hundred good practice case studies.



Figure 5. Each node represents an ontology or knowledge base. The shadowed nodes indicate knowledge models which were created during the project.

The core of the ontology is based on a learning initiative class which represents a single documented case in the Good Practice database. As we can see from table 1 the main attributes of learning initiatives are the title, location, date, learning rationale, funders, organisations involved, target learners and the tangible results. Often the descriptions of learning initiatives describe generic rather than specific entities. For example, involved parties are sometimes described using phrases such as "a local college" or "a few mechanical engineering SMEs". These types of statements are captured using the generic-organisation class – the instances of this class are classes of type organisation.

The other key definition within the ontology is the learning-community class. We do not have space here to include this definition but the key attributes include the affiliation, ethnic group, occupation, gender, age, skill level and dependents. This broad range of slots reflects the diverse attributes that learning and social policy researchers argue can affect access to learning within a community.

Ontology Population

Although WebOnto is primarily aimed at expert model builders we have recently provided a number of tools to allow non-experts to populate ontologies. Integrating support for ontology creation and population within WebOnto contrasts with the approach taken in tools such as Protégé [8] where ontology construction and population are separated.

Help in WebOnto is provided in four main ways:

- *Multiple visualizations* aid in reviewing what has been created.
- *Automatically generated instance forms* support the addition of instances.
- *Knowledge items from web pages* information extraction techniques have been coupled with direct manipulation techniques to enable OCML entities to be created from web pages.
- *Automatic type checking* automatically checking for undefined values and constraint violations.

Multiple Visualizations

The use of visualizations has long been acknowledged to be important in the creation of knowledge models [4]. The key is to provide support for high level or coarse grained views which are tightly coupled to multiple fine grained views. WebOnto provides high level graphical views of class hierarchies tied to fine grained views which use font and colour to differentiate between types of OCML entities.

A significant task where visualizations can aid populators is in validation. Populators need easy-to-read detailed descriptions of the entered knowledge structures. Often the ontological enrichment of a web resource is based on a single class or on a set of related classes - typically class A constrains the type of a slot in class B. Specific resources are represented by a set of connected instances. This heuristic provides the basis for the design of a connected instances visualization. This view displays all the instances connected to a selected instance. Figure 6 shows a connected instances view of the hackney-learninginitiative. Within this view instance names are shown in black, classes in green and slot names in a light blue. Knowledge items which were entered by the user are shown in bold. Any slot values which are instances are expanded. Each instance is picked out using background shading.



Figure 6. A screen snapshot of a connected instance based visualization. Items in bold were defined by the ontology populators. Colour coding distinguishes between instances, classes and relations. Individual instances are picked out with background shading (enhanced for this paper).

Within figure 6 we can see that the hackney-learninginitiative is an instance of learning-initiative. The has-location slot has the value hackney-li-location which is an instance of learning-related-location. The has-premises-type slot of hackney-li-location has two values - the classes community-centre-premises The and library-premises. department-foreducation-and-employment instance was created by the user but the values of its slots were not. The depth of the inline expansion is defined by the user. Selecting any instance in the view creates a new connected instances view. We elected to provide these visualizations in HTML format so that they could easily be printed and viewed in hardcopy format – a requirement from the lifelong learning researchers populating the ontology.

Automatically Generated Instance Forms

Many errors in semantic annotation occur because of errors in naming existing entities and in selecting the class of new instances [5]. The forms in WebOnto seek to alleviate this by prompting users with the names of relevant knowledge items.

An example of an automatically generated form for editing an instance of a learning community is shown in figure 7. Each slot is displayed as a row. The slot name is a button which displays examples of the values that have been given to the slot within other instances. Figure 8 shows the result of selecting the 'other-involved-parties' button.

N Instance of learning-init	iative				_ 🗆 ×
Name:	hackney-learning-initiativ	/e			
Click on a slot name to see e	examples of its use				
has-title	demonstration outreacl	string	•	None	•
has-project-name	"the hackney project"	string	•	None	•
has-location	hackney-li-location	learning-related-location	•	None	•
has-initiative-date	1997	integer	•	None	Ŧ
has-rationale	hackney-improve-acces	rationale-for-learning	•	None	×
has-funder	department-for-educatic	organization	•	None	·
other-involved-parties	london-borough-of-hack	organization	•	None	•
has-learner	hackney-pembury-leam	learning-community	•	None	•
has-deliverable	hackney-learning-initiati	document	•	None	•
	OK	Cancel			

Figure 7. A screen snapshot showing an automatically generated learning-community instance edit form.

N Information on the slot other-involved-parties	×
Examples of the use of the slot other-involved-parties	-
Examples with instances of the same class	
FOOTBALL-CLUBS-LEARNING-INITIATIVE had the value LOCAL-EDUCATION-AUTHORITY	
CHINATOWN-LEARNING-INITIATIVE had the value LIA-NETWORK-ORGANIZATION	
VIDEO-LINKS-LEARNING-INITIATIVE had the value PLYMOUTH-COLLEGE	
BOLTON-TELEMATICS-LEARNING-INITIATIVE had the value BOLTON-BUSINESS-SCHOOL	-
4	•
Ger Unsigned Java Applet Window	

Figure 8. A screen snapshot of the help given when selecting the other-involved-parties button of the form shown in figure 7.

The second column is a simple text field into which the name of a value can be entered. Within our underlying knowledge modelling language OCML [15] slots can be typed using a class or a combination of classes (e.g. (or organization person)). These classes and all of their descendants appear in alphabetical order the third column of the form. Figure 9 shows a user selecting the training-organization class for the other-involved-parties slot. When a class is selected the instances of the class appear in the menu in the fourth column. Figure 10 shows a user selecting the focus-central-london instance.

other-involved-parties london-borough-of-hackne	organization	•	None	•
	senior-manager	•		
has-learner hackney-pembury-learning	senior-research-fellow		None	*
	student			
has-deliverable hackney-learning-initiative	system-manager		None	•
	trade-association			_
OK	trade-union			
	trainer			
Unicod Jaus Applet Window	training-organization			

Figure 9. A screen snapshot showing a user selecting the training-organization class for the other-involvedparties slot of a learning-initiative instance.

other-involved-parties	london-borough-of-hack	training-organization	v None	
has-learner	hackney-pembury-learn	learning-community	None	
has-deliverable	hackney-learning-initiati	document	afp-training-organization amtec-uk-ltd	
	OK	Cancel	aztec birmingham-and-solihull-tec bradford and district tor	
🖥 JUnsigned Java Applet Wind	low		business-training-and-advisory-serv	
			calderdale-tec celtec	
			cewtec	
			dorset-tec	
			essex-tec fas-irish-training-and-employment-;	
			focus control london	2

Figure 10. A screen snapshot showing a user selecting the focus-central-london instance for the otherinvolved-parties slot of a learning-initiative instance.

The forms here are in some respects similar to the forms provided in Protégé-II [8]. The key difference is that instance forms in WebOnto are generated directly from the ontology whereas the forms in Protégé-II use an extra set of form specific definitions. The extra information means that the generated forms can use non-trivial layouts but require an extra compilation cycle. Within WebOnto any changes to the ontology are immediately reflected within the forms.

Knowledge Items from Web Pages

As with the majority of our application domains a proportion of the elements referred to in the observatory knowledge base appear within web documents, specifically, within the entries within the Good Practice database. To aid in the generation of knowledge items from web documents WebOnto contains an interface to a *named entity recognizer*. Named entity recognizers are used to extract items of a pre-specified type from grammatical text. We currently use Marmot [17] to tokenize the text (identifying the nouns) and Badger [17] extract the named entities. We also use a regular expression matcher (written in Perl) because Badger relies on the input text being composed of grammatical sentences (nouns, verbs and prepositions) and this is not always the case for the learning initiatives.

The interface between OCML and the entity recognizer is implemented with two types of constructs: pattern definers and templates. A pattern definition consists of the name of an OCML class or instance and a set of strings which represents patterns using the using the standard notation for regular expressions. The pattern for a college is:

```
(def-pattern college
"(capital_word)* College"
"(capital_word)* College of (capital_word)*")
```

Within the observatory case we have created patterns to identify organizations, ethnic groups, peoples' names and dates.

Templates are used to create new OCML structures from the results of the entity recognizer. Currently three types of template are used:

- *New class instance* this specifies how text can be used to create a new instance of a class.
- *New class subclass* this specifies how subclasses of a class can be created.
- *Fill instance* specifies how an existing instance is filled.

A template consists of the name of a class or instance, a list of variables and the template body. Within the template body variables are denoted by the prefix '\$', and, \$class-name and \$instance-name are special variables which represent the name of the class and instance respectively. The template used to create the hackney-community-college instance was:

(def-new-instance-template organization (name)
 (def-instance \$name \$class-name))

Other examples of how we have combined our knowledge modelling infrastructure with information extraction technologies can be found in [21].

Automatic type checking

The late 80s and 90s saw a considerable effort into creating tools for validating and verifying knowledge bases [14]. We have found that even relatively simple tools can aid ontology populators. OCML contains a general purpose real-time constraint checker. The output of checking the observatory knowledge base is shown in figure 11. Any of the instances or relations shown in figure 11 can be inspected by simply clicking on them.

N Results of Checking the observatory-kb2 ontology	
The result of checking Observatory-Kb2 was:	-
In the instance hackney-learning-initiative-course	1323
the slot has-content has the unknown value food-hygiene-cou	urse 🔄
In the instance swt-learning-community-delivery-method	
the slot medium-used has the unknown value networked-comput	ers
In the instance hattersley-cc-deliverer-delivery-method	
the slot has-producer has the unknown value wea	-
	•
Generation Contractions Contraction Contra	

Figure 11. The result of carrying out consistency checking on the observatory knowledge base. Items within the knowledge base are highlighted using colour and can be selected and inspected. Colour is used to distinguish between instances and relations.

RELATED WORK

The KA^2 initiative [1] shares a number of commonalities with our work. As with the case described here the aim of KA^2 is to allow a community to build a knowledge base collectively, by populating a shared ontology. The knowledge base is constructed by annotating web pages with special tags, which can be read by a specialised search engine cum interpreter, Ontobroker [6]. In this paper we have described and approach which learns from the early problems reported in that initiative [5].

A number of tools such as the CEDAR toolkit [10] and OntoAnnotate [19] provide support based on a web browser integrated with a view of an ontology. The CEDAR annotation tool allows segments of text from web pages to be associated with OCML structures stored on a WebOnto server. Within OntoAnnotate text can be selected from a web page and dragged to fill in the value of an instance. OntoAnnotate also contains mechanisms for managing annotations after an ontology is altered, a text pattern matcher similar to the one described here and links to an ontology based information extraction system. Both the CEDAR annotation tool and OntoAnnotate are designed to use ontologies to annotate web pages whereas goal of the technologies described here are to facilitate the population of ontologies. Hence, rather than creating a separate tool we elected to extend WebOnto thus tightly coupling the ontology development and resource description activities.

In terms of the underlying architecture, as we stated earlier the main difference between our approach and the above approaches to adding semantic information to web pages is that we decouple the web pages from the knowledge model. We should state however that the WebOnto server is now able to export knowledge models in OIL RDF syntax [7]. This facility was used to incorporate parts of our library into an OIL based ontology server as part of a dynamic link service (see [12] for more details).

CONCLUSIONS

In this paper we have described how ontologies can support knowledge sharing within communities of practice. To be successful it is important that all stakeholders are able to participate in the ontology development process and that this process is ongoing and integrated with ontology population. Moreover, ontology population requires support from a mixture of technologies and as far as possible should be integrated into existing working practices.

We have now been using this approach over a number of years in a variety of projects, in domains ranging from managing best practice in the aerospace industry, to supporting the application of medical guidelines. Our experience to date suggests that our approach appears to provide both the technology and the methodological framework required to minimize risk and ensure the participating community's acceptance.

ACKNOWLEDGEMENTS

This work was funded by the Marchmont Project under the Adapt Programme and the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01. The AKT IRC comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University.

The authors would like to thank Maureen Nichols and Sophie Farnes for their coding effort and tolerance.

REFERENCES

- Benjamins, R., Fensel, D. and Gomez Perez A. Knowledge Management through Ontologies. In U. Reimer (editor), *Proceedings of the Second International Conference on Practical Aspects of Knowledge Management*. Basel, Switzerland, 29-30 October, 1998.
- Bowker, G. C., and Star, S. L. Sorting Things Out: Classification and its Consequences. MIT Press: Cambridge, Mass., 1999.
- 3. Domingue, J. Tadzebao and WebOnto: Discussing, Browsing, and Editing Ontologies on the Web. In B. Gaines and M. Musen (editors), *Proc 11th Knowledge* Acquisition *for Knowledge-Based Systems Workshop*, Banff, Canada, April, 1998.
- 4. Eisenstadt, M., Domingue, J., Rajan, T. and Motta, E. Visual Knowledge Engineering. *IEEE Transactions on Software Engineering* Special Issue on Visual Programming, 16(10), 1164-1177, October, 1990.
- Erdmann, M., Maedche, A., Schnurr, H.P., Staab, S. From Manual to Semi-Automatic Semantic Annotation: About Ontology-based Text Annotation Tools. *COLING-2000 Workshop on Semantic Annotation and Intelligent Content,* Centre Universitaire, Luxembourg, 5-6 August, 2000
- 6. Fensel, D. Decker, S., Erdmann, M. and Studer, R. Ontobroker: The very high idea. *Proc* 11th Annual Florida Artificial Intelligence Research Symposium (FLAIRS-98).
- Fensel, D., Horrocks, I., van Harmelen, F., Decker, S., Erdmann, M. and Klein, M. OIL in a Nutshell, *Proc.* 12th Int'l Conf. Knowledge Engineering and Knowledge Management, Lecture Notes in Computer Science, vol. 1937, 2000, Springer Verlag, New York, 1-16
- Grosso, W. E., Eriksson, H., Fergerson, R. W, Gennari, J. H., Tu, S. W., and Musen, M. A. Knowledge Modeling at the Millennium (The Design and Evolution of Protege-2000). *Proc 12th Knowledge Acquisition Workshop*, Banff, Alberta, Canada, October, 1999.
- 9. Gruber, T. R. A Translation Approach to Portable Ontology Specifications. *Knowledge Acquisition*, 5(2). 1993.
- 10. Hatala, M., and Hreno, J. Annotation of documents with knowledge model concepts. *Enrich Report ID44-O-K*, 2000. (Available at http://kmi.open.ac.uk/projects/enrich/id44.pdf).

- 11. Heflin, J. Hendler, J., and S. Luke. Applying Ontology to the Web: A Case Study. In *Proc of the International Work-Conference on Artificial and Natural Neural Networks, IWANN'99*.
- 12. Kalfoglou, Y. Domingue, J., Carr, L., Motta, E., Vargas-Vera, M. and Buckingham Shum, S. On the integration of technologies for capturing and navigating knowledge with ontology-driven services. *KMi Technical Report No. 106*, April, 2001.
- 13. Lave, J. and Wegner, E. *Situated Learning: Legitimate Peripheral Participation*, Cambridge University Press, Cambridge, UK, 1991.
- Meseguer, P. and Preece, A. Verification and validation of knowledge-based systems with formal specifications. *The Knowledge Engineering Review*, 10(4), 331-343, 1995.
- 15. Motta E. *Reusable Components for Knowledge Models*. IOS Press, Amsterdam, 1999.
- 16. Motta, E., Buckingham Shum, S. and Domingue, J. Ontology-Driven Document Enrichment: Principles, Tools and Applications. *International Journal of Human Computer Studies*. 52(5), 1071-1109, 2000.
- 17. Riloff, E. An Empirical Study of Automated Dictionary Construction for Information Extraction in Three Domains. *The AI Journal*, 85, 101-134, 1996.
- 18. Riva, A. and Ramoni, M. LispWeb: a Specialized HTTP Server for Distributed AI Applications, *Computer Networks and ISDN Systems*, 28 (7-11), 953-961, 1996.
- 19. Staab, S., A. Mädche, S. Handschuh. An Annotation Framework for the Semantic Web. In: S. Ishizaki (ed.), *Proc. of The First International Workshop on MultiMedia Annotation*. January, 30 - 31, 2001. Tokyo, Japan.
- Tennison, J. and Shadbolt, N. R. APECKS: a Tool to Support Living Ontologies. Proc. of the 11th Banff Knowledge Acquisition Workshop, Banff, Alberta, Canada, April 18-23, 1998.
- 21. Vargas-Vera, M., Domingue, J., Kalfoglou, Y., Motta, E. and Buckingham-Shum, S. Template-driven information extraction for populating ontologies. *Proc* of the IJCAI'01 Workshop on Ontology, Learning, Seattle, WA, USA 2001.
- 22. Wenger, E. Communities of Practice: Learning, Meaning, and Identity. Cambridge University Press: Cambridge, 1998