

Ontology-Driven Document Enrichment: Principles, Tools and Applications

Enrico Motta, Simon Buckingham Shum and John Domingue

Knowledge Media Institute
The Open University
Walton Hall, MK7 6AA
Milton Keynes, UK

{e.motta, s.buckingham.shum, j.b.domingue}@open.ac.uk

Abstract. In this paper we present an approach to *document enrichment*, which consists of developing and integrating formal knowledge models with archives of documents, to provide intelligent knowledge retrieval and (possibly) additional knowledge-intensive services, beyond what is currently available using 'standard' information retrieval and search facilities. Our approach is *ontology-driven*, in the sense that the construction of the knowledge model is carried out in a top-down fashion, by populating a given *ontology*, rather than in a bottom-up fashion, by annotating a particular document. In the paper we give an overview of the approach and we examine the various types of issues (e.g., modelling, organizational and user interface issues) which need to be tackled to effectively deploy our approach in the workplace. In addition we also discuss a number of technologies we have developed to support ontology-driven document enrichment and we illustrate our ideas in the domains of electronic news publishing, scholarly discourse and medical guidelines.

1. INTRODUCTION

An important activity in knowledge management is "to convert text to knowledge" (O'Leary, 1998). This activity is central to knowledge management for two reasons: i) work practices and information flow in organizations tend to be *document-centred* and ii) documents themselves do not normally exhibit the amount of structure required to support semantically-aware search engines or other forms of intelligent services. For these reasons there has been much interest in technology to support the specification of structured information in textual documents, especially web pages. The web standardisation community has focused on the underlying representational infrastructure: XML (XML, 1999) has been proposed as the basic annotation formalism to support the specification of structured information in web pages, while RDF builds on the XML syntax to provide a standard declarative representation, which allows users to express semantic relationships between items on the Web. Approaches such as Ontobroker (Fensel et al., 1998) and Shoe (Heflin et al., 1998) provide formalisms and associated interpreters which make it

possible to embed knowledge representation structures in web pages and use them to perform inferences.

In this paper we look at the wider issues concerning "the conversion of text to knowledge" and discuss a comprehensive approach to *document enrichment* (Sumner et al., 1998), which we are trying out in a number of projects here at the Knowledge Media Institute. The approach is characterized in terms of a set of activities, with associated informal guidelines. In the paper we also describe a number of technologies, which we have developed to support our approach to document-centred knowledge management. These technologies include a *knowledge modelling language*¹, form-based interfaces for adding and retrieving knowledge from a model, and a web-based browser/editor, which supports the collaborative development of knowledge models over the World-Wide-Web. Finally, we discuss the application of our approach to three domains: *electronic news publishing* (Domingue and Motta, 1999), *scholarly discourse* (Buckingham Shum et al., 1999) and *medical guidelines* (PatMan, 1998).

The paper is organized as follows: in the next section we give an overview of our approach, in terms of the underlying methodological assumptions and the associated process model. In section 3 we describe the technology we have developed to support the approach. In sections 4, 5 and 6 we discuss the application of the approach to the three aforementioned domains. Finally, in sections 7 and 8 we discuss related work and reiterate the main contributions of this paper.

2. ONTOLOGY-DRIVEN DOCUMENT ENRICHMENT

Our approach is *ontology-driven*, in the sense that the construction of the knowledge model is carried out in a top-down fashion, by populating a given *ontology* (Gruber, 1993), rather than in a bottom-up fashion, by annotating a particular document. Figure 1 underlines this point graphically, by emphasizing that the construction of a knowledge model is driven by a pre-existing ontology, a set of documents and other sources of knowledge, such as appropriate (human) experts. Following Gruber, we use the term "ontology" to indicate "a specification of a reusable conceptualization". More simply, an ontology can be seen as providing a vocabulary for describing a range of models. For instance, an ontology for medical guidelines provides a generic set of concepts and relations (e.g., medical condition, diagnostic guideline, guideline user type), which can then be instantiated for particular guidelines to build guideline-specific models, in domains such as stroke management or prevention of pressure ulcer.

An ontology-driven approach to model construction affords several advantages. Instantiating an ontology is usually simpler and speedier than developing a model from scratch. In addition,

¹ Here we use the term "knowledge modelling" as a short form for "knowledge-level modelling", an expression introduced by Allen Newell (1982) to describe models of knowledge-intensive behaviour which abstract from the way this behaviour is implemented and focus instead on the knowledge employed by an agent and the goals the agent is trying to achieve.

because an ontology makes explicit the conceptualization underlying a particular model, it becomes easier to maintain, reuse and interoperate the model with other components. Finally, reasoning modules can be associated with an ontology and these are then applicable to all models built by instantiating the ontology in question. For instance, in the case of medical guidelines, one can envisage building ontology-specific guideline verification tools, which can then be used to verify individual guidelines developed by instantiating the same generic guideline ontology.

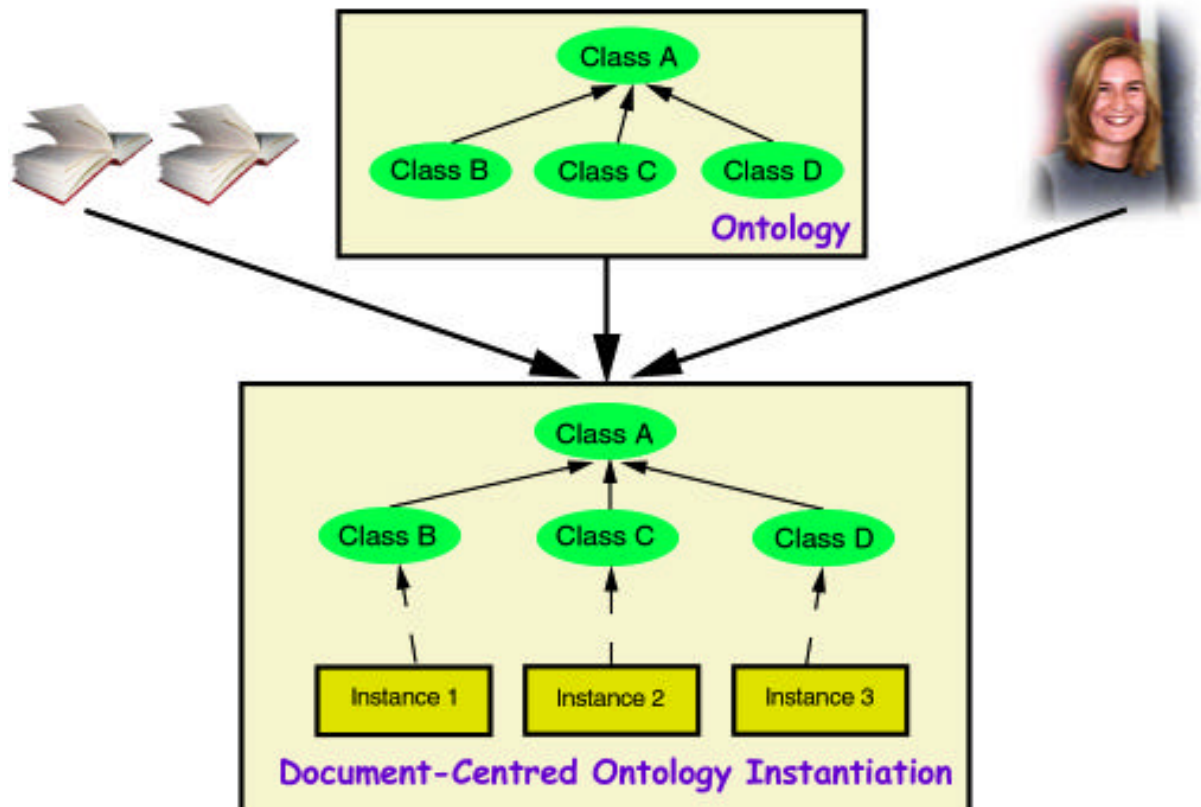


Figure 1. Ontology-driven Document Enrichment

Because our model construction process is ontology-driven, we prefer to use the term "enrichment" (Sumner et al., 1998), rather than "conversion" or "annotation", to refer to the process of associating a formal model to a document (or set of documents). In general, a representation, whether formal, graphical or textual, can be enriched in several different ways - e.g., i) by providing information about the context in which it was created, ii) by linking it to related artefacts of the same nature, or iii) by linking it to related artefacts of a different nature. Although in our document-centred knowledge management work we provide multiple forms of document enrichment, such as associating *discussion spaces* to documents (Sumner and

Buckingham Shum, 1998), in this paper we will primarily concentrate on the association of formal knowledge models to documents².

Thus, an important facet of an ontology-centred approach to document enrichment is that the formalised knowledge is not meant to be a translation of what is informally specified in the associated document. Hence the knowledge model typically plays a different role from the associated text. For instance, in the medical guideline scenario the knowledge model helps to verify that all the kinds of knowledge expected to be found in a document describing a medical guideline are indeed there. In the scholarly discourse scenario the knowledge model is meant to capture the meta-knowledge required to structure academic debates (e.g., theory X contradicts theory Y), which is often expressed only implicitly in publications (i.e., acquiring it typically requires some interpretation effort) and is not modelled at all in traditional libraries. In a nutshell, the emphasis in our approach is in identifying the *added value* (in terms of enabling semantic retrieval and document indexing capabilities, or other reasoning services), which can be provided by a formalised knowledge model. Our methodology comprises the following six steps.

1. Identify use scenario.
2. Characterize viewpoint for ontology.
3. Develop the ontology.
4. Perform ontology-driven model construction.
5. Customise query interface for semantic knowledge retrieval.
6. Develop additional reasoning services on top of knowledge model.

These steps are briefly described in the next sub-sections.

2.1 Identify Use Scenario

At this stage the services to be delivered by the knowledge management system are defined. In particular, issues of feasibility and cost are investigated. Addressing the latter involves answering questions such as: “What is the added value provided by the knowledge model, considering the non-trivial costs associated with the development and instantiation of an ontology?”, “Is there the need for a ‘full-blown’ knowledge model and for going beyond the facilities provided by off-the-shelf search engines?”, “What additional reasoning services will be provided, beyond deductive knowledge retrieval?”. Addressing feasibility issues requires assessing (among other things) whether or not it is feasible to expect the target user community to perform document enrichment or whether specialized human editors will be needed. This latter solution introduces a significant bottleneck in the process and moreover assumes that to introduce a central editor in the model development is actually feasible. This is definitely not the case in some of our application domains. For instance, in the scholarly discourse scenario our

² Having said so, the medical guideline scenario described in section 6 does integrate a formal model with a set of discussion spaces, to provide multiple forms of document enrichment.

aim is to allow members of academic community to express their own views about academic contributions, whether their own or whether proposed by other scholars. In this scenario it is not feasible to envisage a central editor mediating between an author and a shared knowledge base. In addition, we also believe that aiming to enable non-experts to take part in the development of knowledge models is a good research objective per se, not only on the grounds of cost-effectiveness and sustainability, but also because of the perceived benefits deriving from being engaged in the modelling activity. Specifically, there is already some evidence that involving students in constructing knowledge models improves their understanding of a discipline (Wideman and Owston, 1993). Although no hard data are available, there is no obvious reason why these benefits should not occur in scenarios outside the educational domain.

Of course, successfully enabling non-experts to engage in knowledge modelling is far from a trivial proposition. There are subtle trade-offs involved here, which we address in section 2.4, when discussing the ontology instantiation task.

2.2 Characterize viewpoint for ontology

The previous step was concerned with assessing the feasibility of the approach and defining the functionalities of the envisaged system. This step focuses on the ontology: having decided on an ontology-based approach, it is important to characterize the particular viewpoint that the envisaged ontology will impose on the documents. For instance, in the electronic publishing domain we focus on the events characterizing academic life, while in the scholarly discourse domain the focus is on academic debate. Clearly the distinction between this step and the steps immediately before and after can in some cases be fuzzy. For instance, the specification of the viewpoint in the scholarly discourse domain is tightly integrated with the characterization of the use scenario. However, we believe that it is useful to explicitly distinguish a viewpoint specification task for two reasons. The first one is that it is feasible to envisage scenarios in which different viewpoints can be taken as the starting point for the ontology development process. Examples abound in the literature, especially with respect with highly generic ontologies - compare for example the *Cyc* ontology (Lenat and Guha, 1990) with the work by Sowa (1995). Another reason is methodological: it is useful to separate the issues concerning the functionalities, scope and purpose of the envisaged system from the specification of the functionalities, scope and purpose of the ontology. For instance, in the electronic publishing scenario the knowledge management system is concerned with providing semantic retrieval capabilities and services supporting story identification and personalised news feeds - see section 4 for more details. Within this scenario the chosen viewpoint focuses on modelling academic events and the key 'actors' in these events: technologies, people, organizations and projects.

2.3 Develop the Ontology.

Having defined a particular viewpoint over a domain (in this case, a set of documents), ontology development is largely a technical enterprise - i.e., the issues here concern modelling and

formalization, rather than scope and purpose. Several approaches to ontology development have been proposed in the literature, which introduce distinctions along different dimensions - contrast, for instance, the *bottom-up* development style of van der Vet and Mars (1998) with Sowa's (1995) *top-down* approach. Uschold and Gruninger (1996) argue that a *middle-out, purpose-driven* approach is most effective, in which the basic concepts in a domain are identified first (e.g., dog), and later generalised (mammal) and/or specialised (cocker spaniel).

In our scenarios we have followed a task-independent, middle-out approach, and we use the selected viewpoint to help us to identify the *key concepts* in the class of models we want to construct. For instance, in the electronic newsletter scenario, a starting point was the notion of *news item*, which in turn was characterised as relating a number of *events*. Thus, our main modelling effort centred on identifying and modelling the various types of events relevant to our scenario (academic life). In the scholarly discourse domain, the viewpoint is 'academic debate' and the ontology then concentrates on characterising the most common relations which exist between academic theories, methodologies, models, etc.

Another important issue concerns who constructs the ontology? As pointed out earlier, we want to support scenarios in which knowledge models are constructed collaboratively. But what about the ontology itself? Is this constructed collaboratively? Our answer is negative. In all the projects we have carried out so far, we have centralised the ontology development. 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 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 or failure. Moreover, ontology design requires specialist skills which are normally not possessed by the members of our target user communities.

Our ontology design approach is informed by two main modelling guidelines. The rationale for both of these is user-centred: because we want non-expert users to be able to engage in the model construction process, the ontologies we produce need to be easy to understand and to instantiate. These user-centred considerations have priority over any other design criterion.

- *Minimal ontological definitions.* That is, to try and provide only the minimal set of attributes needed to define a particular class. This approach has the advantage that, when populating the ontology, users don't have to face lots of irrelevant attributes. For instance, in the electronic news publishing domain we initially reused the definition of class event given in the HPKB upper layer (HPKB, 1997), but we then removed about 90% of its slots. The reason for this was that the HPKB definition aims to cover all potential attributes which *can* be relevant to a *generic instance* of class event. However, typically only a few slots will *actually* be relevant for any *specific subclass* of the class. For instance, slot *damages* is only relevant to events which can cause damage. Hence, we

would introduce this slot when characterising a subclass of class event, such as *damaging-event*, rather than associating it with class event itself.

- *User-centred definitions.* This guideline requires that the terminology used by the ontology needs to be easy to understand for a user who is not necessarily a knowledge engineer. There are two aspects here: heavily technical modelling concepts – e.g., sophisticated modelling solutions for representing time – ought to be avoided. Moreover, the terminology should be as context-specific as possible. For instance, while we can talk about “agents performing events”, when describing events in general, we should use the class-specific terminology, “awarding body assigns awards”, when talking about an award-giving type of event. This latter guideline implies that the underlying modelling language should support slot renaming along *isa hierarchies* – i.e., inherited slots should get subclass-specific names. The importance of domain-specific, user-oriented terminology has been recognised in knowledge acquisition for a long time (Musen, 1989) and arguably provides an important difference between the criteria associated with modelling for knowledge acquisition and those associated with modelling for system development.

2.4 Perform ontology-driven model construction

We are acutely aware that many schemes for registering shared resources and providing structured descriptions founder on the crucial 'capture bottleneck' - the envisaged beneficiaries of the system simply do not have the motivation or time to invest in sharing resources to reach a critical mass of useful material. Sobering lessons on this theme have been drawn for group and organizational memory systems (Selvin, 1999), and indeed, for any system that requires users to formalise information (Shipman and Marshall, 1999). Why should we succeed where others have failed? Our working hypothesis is that our domains have unique characteristics lacking in domains in which collaborative development has failed.

- *Co-operation rather than competition.* We are selecting domains where co-operation, meant here as “willingness to share knowledge”, is either a basic work premise or is enforced by external constraints. For instance, academic analysis and publishing require scholars to read, refer to and praise/criticise each other's work. In sum, the dynamics of academic publishing requires co-operation. A similar situation occurs in the medical guidelines scenario. Institutions and individual scientists may compete with each other, but the outcome of consensus conferences (by definition) are shared knowledge resources. In other scenarios, for instance when constructing an organizational memory, other forces (e.g., directives from higher management) may force co-operation, even when competition would be the norm.
- *Benefits outweigh costs.* Motivation is a crucial aspect. Motivation essentially boils down to a cost-benefit analysis. For instance, in the scholarly discourse scenario, we assume

that the basic motivation of an academic is to disseminate his/her work. Hence, having completed a new document, the author will want to maximise its 'digital presence' on the net by carefully encoding its contributions and connections to the existing literature.

- *Compatibility with organizational work practices.* This requires the seamless integration of our document enrichment model with existing work practices. For instance, in the case of electronic publishing, the ontology is used to enrich news items, which are submitted either through email or through a web-based form. Hence, at least for those users who submit through the latter mechanism, instantiating the ontology becomes an additional form-filling activity, carried out using the same medium (i.e., the web browser) and at the same time. Analogously, in the case of scholarly discourse, at least in some academic communities, authors are used to submitting papers to digital repositories and providing metadata. Filling an ontology-derived form should then be perceived as a small, additional step³.

2.5 Customize query interface for semantic knowledge retrieval

At this stage the appropriate query interface is designed, in accordance to the use scenario and the expected functionalities. To support this step we have developed a flexible form-based interface, called *Lois*, which can be customised for each specific application domain. *Lois* is described in section 3.3.

2.6 Develop additional reasoning services on top of the knowledge model

This final step is where the real benefit of the approach lies. Once a knowledge model has been produced, then it becomes possible to provide additional intelligent functionalities and ensure that the benefits outweigh the costs. These reasoning services tend to be application specific. For instance, in the scholarly publishing scenario, we are planning to develop specialised agents, whose goal is to identify emerging scholarly perspectives, using heuristic knowledge and machine learning techniques. For instance, an agent could discover a 'European perspective' on a particular issue, if a structural pattern in the knowledge model - e.g., use of formal methods - also matched the geographic location of the relevant researchers. In the electronic publishing domain we have designed two agents, which reason about the contents of the knowledge model to identify new, potentially 'hot' stories and to provide personalised news feeds - see section 4.2.5.

³ Actually this is not the case in the medical guideline scenario. In this case accommodating the guideline management technology described in this paper will normally require a change in organizational work-practices, as computerised guideline support is still fairly rare in hospital settings.

3. TECHNOLOGIES FOR ONTOLOGY-DRIVEN DOCUMENT ENRICHMENT

In this section we describe the main modelling support technologies we have developed to enable ontology-driven document enrichment. These are as follows:

- *OCML*. An operational knowledge modelling language, which provides the underlying representation for our ontologies and knowledge models.
- *WebOnto*. A tool providing web-based visualization, browsing and editing support for developing and maintaining ontologies and knowledge models specified in OCML.
- *Lois*. A flexible form-based interface for knowledge retrieval.
- *Knote*. A form-based interface for populating an ontology.

In addition to these modelling support technologies, we also have a number of document-centred technologies, such as the *D3E shell* (Sumner and Buckingham Shum, 1998), which supports a tight integration of discourse and debate facilities within documents. Because the focus of this paper is on ontology-driven knowledge management, we circumscribe the technology review in this section to modelling support technologies. Other relevant technologies, such as D3E, will be introduced ‘opportunistically’ in the application-specific sections.

3.1 OCML, an operational knowledge modelling language

The development of the OCML language (Motta, 1999) was primarily motivated by dissatisfaction with existing ontology specification languages. Languages such as *Ontolingua* (Gruber, 1993) benefit from a well-defined semantics and extensive libraries, but provide very limited support for model operationalization. No interpreter exists for *Ontolingua* models and the translation approach proposed for operationalizing *Ontolingua* models into executable languages suffers from two serious limitations. First of all it does not effectively support rapid prototyping and incremental model development - essentially it proposes a ‘batch mode’ style of operationalization. Secondly, it assumes the availability of smart translators - e.g., for translating *Ontolingua* models to knowledge representation languages, such as *Loom* (MacGregor, 1991). Unfortunately these are currently beyond the state of the art of language translation and are “unlikely to be forthcoming in the foreseeable future” (Uschold et al., 1998). As a result *Ontolingua* models are typically operationalized through extensive manual effort, an approach which does not lend itself to rapid prototyping. Similar arguments apply of course to non-executable specification languages in general and it does not come as a surprise that the importance of executable specifications has been argued also in mainstream software engineering (Fuchs, 1992).

OCML was therefore conceived as a kind of ‘operational *Ontolingua*’. Our goal was to develop a language that would be as compatible as possible with *Ontolingua*, thus benefiting from the advantages in terms of well-formed semantics, availability of libraries and large diffusion enjoyed by *Ontolingua*, while at the same time providing support for rapid prototyping and

incremental development. Hence, consistently with the notion of knowledge-level modelling OCML focuses on logical, rather than implementation-level primitives. It provides constructs for specifying relations, functions, classes, instances and axioms, and, in particular, supports the variety of relation specification keywords used by Ontolingua - e.g., `:iff-def`, `:def`, `:sufficient` and `:axiom-def`.

```
(def-relation <> (?x ?y)
  "True if ?x and ?y do not denote the same object"
  :iff-def (not (= ?x ?y)))
```

Figure 2. Sample OCML definition.

In order to allow rapid prototyping, OCML provides a function interpreter, a control interpreter and a proof system, which seamlessly integrates inheritance and function evaluation with a backward chaining inference engine. Hence, an important aspect of OCML is that it tries to fulfil a dual role: on the one hand it aims to provide a rich specification environment, compatible with *de facto* standards, such as Ontolingua. On the other hand it also aims to provide effective operationalization support, by means of interpreters and an inference engine. This dual role is normally supported by OCML in a transparent way. For instance, let's consider the definition given in figure 2. The `:iff-def` option in the relation specification plays a dual role: it helps to characterize the semantics of the relation and at the same time it is used by the OCML proof system to check whether the relation `<>` holds for a pair of arguments.

```
(def-relation SUPERCLASS-OF (?super ?c)
  "The inverse of subclass-of"
  :iff-def (and (class ?super)
                (class ?c)
                (member ?super (all-superclasses ?c)))
  :prove-by (or (and (variable-bound ?c)
                    (member ?super (all-superclasses ?c)))
                (and (variable-bound ?super)
                    (member ?c (all-subclasses ?super)))
                (and
                 (not (variable-bound ?super))
                 (not (variable-bound ?c))
                 (class ?super) (class ?c)
                 (superclass-of ?super ?c))))
```

Figure 3. OCML definition which separates specification and operationalization.

In other cases the user may require more control over the operational use of a definition - for instance, in cases where the given specification will be too inefficient. To this purpose OCML provides extra keywords, which allow a user to associate *procedural attachments* (Weyhrauch, 1980) to a definition, either using OCML itself, or by means of Lisp code.

Figure 3 provides an example of an OCML definition which separates specification and operationalization. The `:iff-def` keyword provides the formal definition of relation `superclass-`

of, while the `:prove-by` keyword defines a way to infer whether or not the relation holds for a given pair of arguments. Because a `:prove-by` option has been specified, the OCML interpreter will not use the `:iff-def` option when trying to reason about this particular relationship. Such an integrated use of specification and operationalization in the same language is similar to the way the *epistemological* and *heuristic* levels are amalgamated in Cyc (Lenat and Guha, 1990). Indeed, we found this feature very useful in our knowledge management work, because it allows us to quickly move from ontology specification to a knowledge base and then incrementally improve the efficiency of the knowledge base opportunistically, in response to eventual performance bottlenecks.

OCML modelling is supported by a library of reusable definitions, which is structured according to the basic categories of our application modelling framework: *task*, *method*, *domain* and *application* (Motta, 1999). The library also relies on a number of *base ontologies*, which provide definitions for basic modelling concepts, such as numbers, sets, relations, tasks, methods, roles, etc. The OCML base ontology also includes the constructs specified in the Ontolingua frame ontology, thus ensuring compatibility between frame-based specifications in the two languages. Constructs for which OCML does not provide operational support - e.g., axioms - are simply added to the model, but they are not used in the reasoning process.

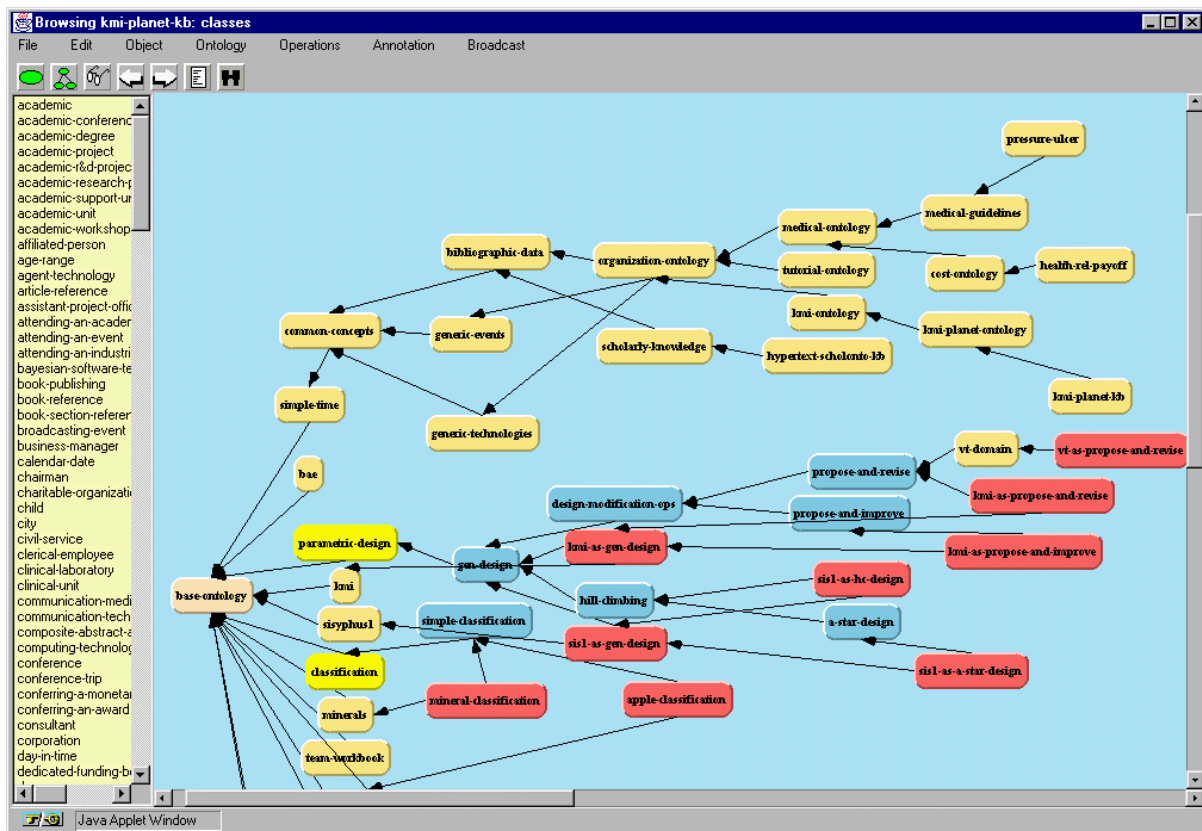


Figure 4. WebOnto visualization of part of the KMi library of knowledge models.

Our library of OCML models can be accessed through the WebOnto browser at URL <http://webonto.open.ac.uk> - see figure 4. WebOnto is described in the next section.

3.2 WebOnto: browsing and editing knowledge models on the web

WebOnto (Domingue, 1998) enables knowledge engineers to collaboratively browse and edit knowledge models over the web, using a standard web browser. The architecture is composed of a central server and a Java applet. The central server is built on top of a customised web server, *LispWeb* (Riva and Ramoni, 1996). In addition to implementing the standard HTTP protocol, the *LispWeb* server offers a library of high-level Lisp functions to dynamically generate HTML pages, a facility for dynamically creating image maps, and a server-to-server communication method. The WebOnto server interacts with the ontology libraries represented in OCML and is responsible for the following tasks.

- Converting OCML structures into ASCII strings which can be sent via HTTP.
- Sending requested OCML source code to WebOnto clients.
- Updating server source code from WebOnto clients.
- Locking an ontology when it is being edited (sub-ontologies remain unlocked). Locking an ontology ensures that conflicting edit operations, such as two users changing the same OCML structure, cannot occur.
- Managing administrative data about the knowledge models in the library and WebOnto users.

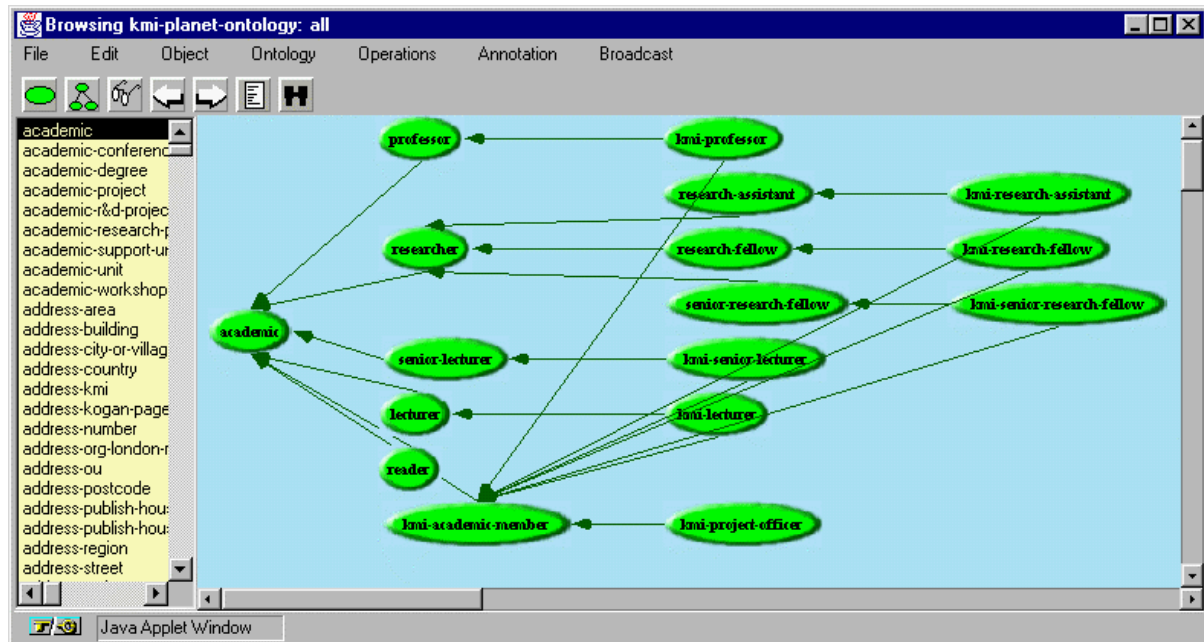


Figure 5. The main WebOnto window showing a class hierarchy.

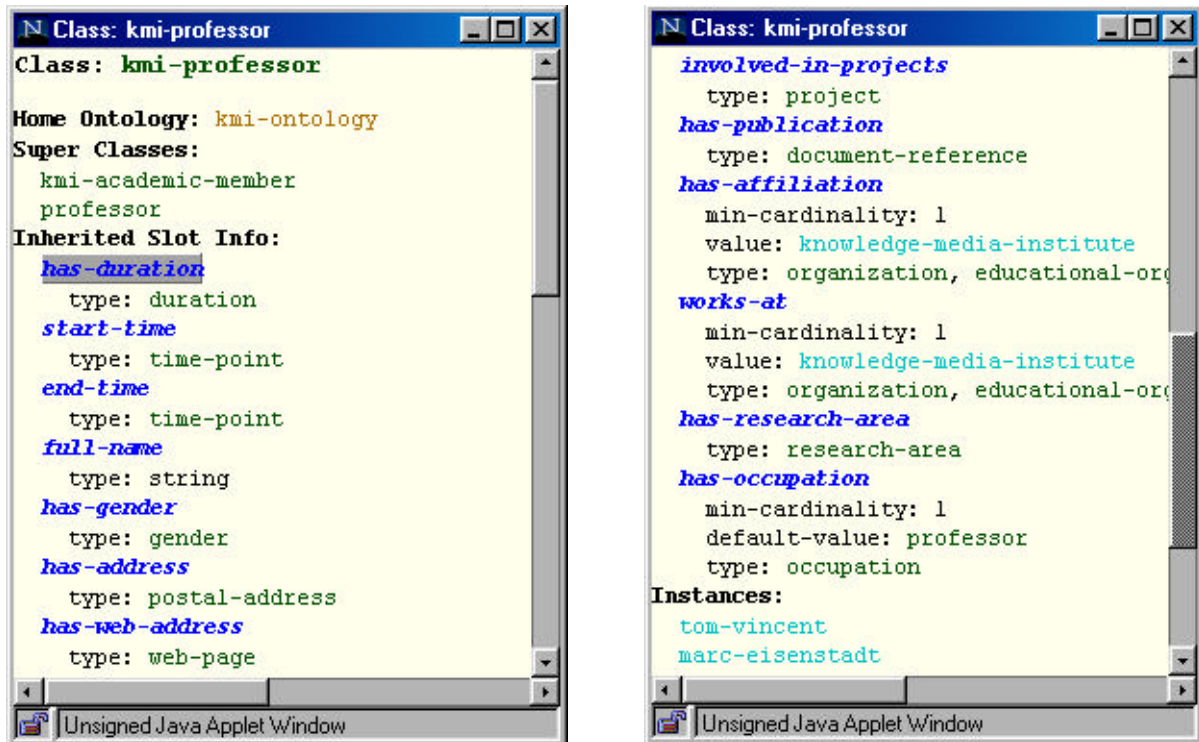




Figure 6. The top and bottom halves of a fine-grained view window giving the details of a class.

The Java applet provides multiple visualizations of OCML knowledge models, a direct manipulation and forms interface for creating new knowledge structures, and a groupware facility which supports both synchronous and asynchronous model building by teams of knowledge engineers. The direct manipulation interface integrates both graphical and textual views of a knowledge model and allows the user to work in either modality, while ensuring that changes are propagated immediately. More specifically, WebOnto gives automatic semantic feedback, in visual form, whenever the user performs an incremental edit of the knowledge model. Tanimoto (1990) describes this as *level 3 liveness*.

Another important aspect of the WebOnto environment is the tight integration of both *coarse-grained* and *fine-grained* visualizations, a feature which is essential to support the development of large models (Eisenstadt et al., 1990).

Figure 5 shows a screen snapshot of the main WebOnto Java Applet window. The class hierarchy shown in the figure was created by selecting the class `academic` in the list shown in the left panel and then clicking on the “draw hierarchy” toolbar button (). Figure 6 shows a detailed view of one of the classes included in the hierarchy in figure 5, `kmi-professor`. This view was obtained by selecting the `kmi-professor` node in figure 5 and then clicking on the “inspect” toolbar button (). Any text displayed in the fine-grained view window is automatically parsed and colour-coded, in accordance with the relevant OCML type (e.g., relation, function, instance, etc.). It can

then be inspected simply by selecting and clicking on it. As an example, we have highlighted slot `has-duration` in figure 6.

Since WebOnto was made publicly available on the web in October 1999, we have already had several hundred visitors to the site and the tool has attracted significant interest from both academic and industrial organizations. The university of Amsterdam has carried out a comparative evaluation of several knowledge modelling tools (Duneiveld et al., 1999), including WebOnto. This was evaluated very favourably, in particular being judged as the most user-friendly and as the one requiring the shortest learning curve.

Because WebOnto is primarily aimed at expert model builders, its relationship to knowledge management is indirect: its role is to support the developers of the underlying ontology, rather than the users of the target knowledge management application. However, i) given that WebOnto has been designed to be as easy to use as possible and ii) in many cases end users want to inspect an ontology directly (for instance, to gain a better understanding of the underlying organization), we usually include pointers to WebOnto in our application interfaces. Hence, in practice we can consider WebOnto as part of our suite of tools for supporting knowledge management applications.

3.3 Lois, a flexible form-based interface for knowledge retrieval.

The aim of Lois is to provide a web-based interface for posing queries to a knowledge model at a level which abstracts from the underlying modelling language. This goal has been accomplished by developing a form-based interface which allows users to select 'key concepts' in the ontology and then construct a query by navigating the structure of the ontology (i.e., by following relations between concepts). This navigation leads to the creation of a query as a *list of rows*, which are linked by logical connectives. For instance, figure 7, which is taken from our electronic publishing domain, shows a query which asks for a member of the KMi academic staff, who is involved in the development of software visualization technology. This query was constructed by selecting class `kmi-member` first, specializing it to `kmi-academic-member`, selecting the relation `has-research-area` and then circumscribing the range of this relation to `res-area-software-visualisation`. More specifically, the first row of the query was created by a) selecting the "Member of KMi" button, b) selecting "concept-sub-type" in the "Attributes" window and c) selecting "kmi-academic-member" in the "Attribute Types" window. The second row was created by selecting "More about kmi-member1" in the pop-up menu underneath "Organisation", then selecting relation "has-research-area" in the "Attributes" window, then "kmi-research-area" in the "Attribute Types" window, and then "res-area-software-visualization" in the "Values" window. The items in each window pane dynamically show the choices relevant to the currently selected item in the pane immediately on the left. For example, the window "Attributes" in figure 7 shows all the attributes appropriate to the currently selected item under the "Concepts" window, which is "More about kmi-member1". Analogously, once we select item "has-research-

area” in the “Attributes” window, then window “Attribute Types” is automatically updated, to show the types of research areas known to the underlying knowledge base.




Figure 7. Finding a KMi researcher who works on software visualization.

The Lois interface is created automatically once the key classes for a knowledge model have been specified. In the example, these are: “Story Event Type”, “KMi Technology”, “KMi Project”, “KMi Member” and “Organization”. These act as the entry points to the space of possible Lois queries. In other words, Lois can cover any query, as long as this can be specified as a path through the knowledge model, starting from one of the key classes. The usability of the Lois interface also depends on the terminology used by the underlying ontology. The knowledge engineer is therefore required to employ class and slot names which can be understood by the Lois user. If this requirement is satisfied, then the user only needs to learn to construct queries through the accumulation of rows.

3.4 Ontology instantiation using Knot

Our goal is to enable as wide an audience as possible to take part in the ontology-driven construction of a shared knowledge model. Knot was therefore designed to be ‘low entry’, so that users would not necessarily require a background in knowledge modelling. At the same time Knot should allow experienced ontology engineers the freedom to create arbitrarily complex

OCML expressions. Knoté provides instance forms which are modelled on the *dynamic forms* of Girgensohn et al. (1995). The key difference between these two types of forms is that instance forms in Knoté are generated directly from the ontology and not from a user description.

The Knoté forms are created on the fly from a class definition and are thus domain independent. This approach has the benefit that it is possible to edit classes and instances iteratively, a feature which is often required during model development. This capability is not supported by systems such as Protégé-II (Eriksson et al., 1994), where the instance forms are created by a distinct meta-tool. Recent work on Protégé-2000 (Eriksson et al., 1999) seeks to alleviate this problem by allowing classes to be viewed as instances of meta-classes.

Slot Name	Value	Type	Options
Name	prevention-of-pressure-ulcer		
has-duration		duration	None
start-time		time-point	None
end-time		time-point	None
has-author	ahcpr	(or person organization)	ahcpr
has-subcomponent	pressure-ulcer-risk-	generic-planning-entity	pressure-ulcer-risk-assessment
has-goals	"identifying at-risk i	string	patient-repositioning provision-of-mechanical-loading-a pressure-ulcer-skin-care-and-earl ensure-adequate-dietary-intake friction-minimization skin-cleansing skin-inspection pressure-ulcer-risk-assessment
additional-decision-support-model		decision-support-model	
has-plan-specification		plan-specification	
outcome-measure	pressure-ulcer-inci	medical-variable	
has-main-goal	"predicting and pr	string	
target-population	people-at-risk-from-	population-specification	
full-name	"pressure ulcer in	string	None
associated-medical-condition-class	pressure-ulcer	medical-condition-type	None
temporal-constraints		string	None
location-constraints		guideline-application-location	None
associated-documents	ahcprpub92-0047	document-reference	ahcprpub92-0047
has-guideline-user-type	generic-care-giver	guideline-user-type	None

Figure 8. Class instantiation with Knoté.

An example of a Knoté form is presented in figure 8. The figure shows a form partially instantiating a medical guideline for the prevention of pressure ulcers - see section 6 for more details on this domain. The structure of the form was derived from the definition of class `preventive-guideline` in the ontology `pressure-ulcer`. Knoté provides the user with quite a lot of help in filling the form. When a slot is typed, Knoté allows the user to navigate the subclasses and the instances of the given type, to select an existing instance or to create a new one. In the latter case, the appropriate Knoté form is then generated. The user can also click on a slot name

(listed in the leftmost column of the form) to get examples of the use of the slot. Our experience suggests that this example-centred support tends to be more useful than the generic documentation associated with a slot. Figure 9 shows the help provided when the user clicks on slot associated-medical-condition-class in the leftmost column of figure 8.

With the description of the Knot form-filling support we have completed our brief overview of our technology for knowledge modelling. In the next sections we will illustrate examples of the approach in three application domains.

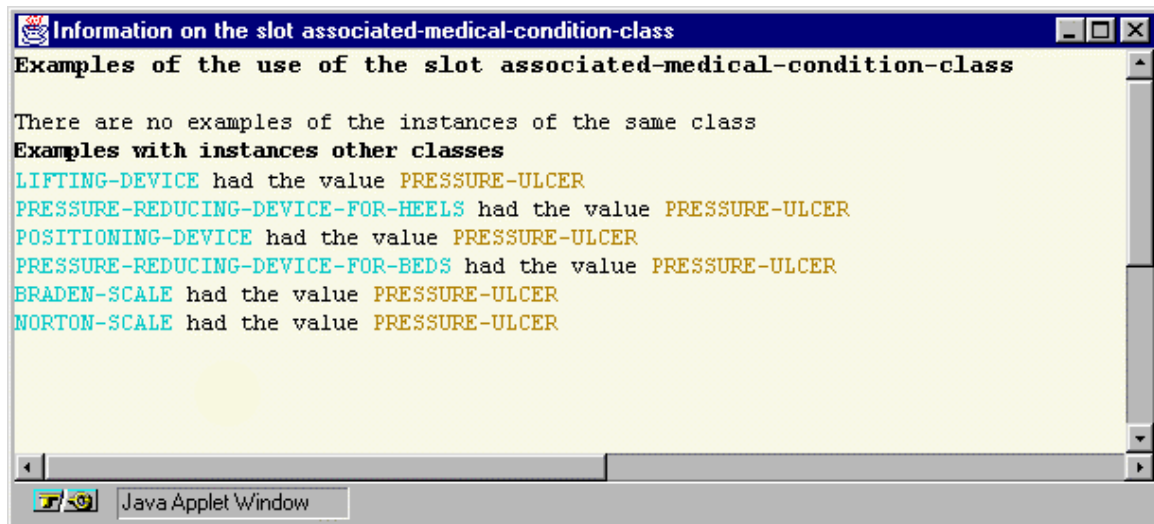


Figure 9. Example-based slot documentation.

4. PLANET-ONTO: ENRICHING NEWS STORIES

4.1 The scenario

KMi Planet (Domingue and Scott, 1998) is a web-based news server which facilitates communication within our laboratory (the Knowledge Media Institute) and allows the wider community to access lab-related items of interest. Unlike most news-of-the-future projects, which focus on personalised news feeds, KMi Planet replicates some of the functionalities provided by newsrooms. 'Journalists' send stories using the lowest common-denominator medium - an email message - and the Planet news server acts as an intelligent layout editor, formatting the news item and related figures, integrating it into the news archive and presenting it in several different modes. The most obvious presentation style is the newspaper-like multiple story page, where the submitted story is presented in the context of other stories, as shown in figure 10. Planet is used as the 'front door' to our laboratory, both metaphorically and physically (a dedicated machine running Planet is stationed at the entrance of the laboratory, so that visitors to KMi tend to be shown or 'play with' Planet first). Our archive is growing steadily and now contains about a hundred stories, submitted by 13 journalists. We currently have 480 registered 'readers', i.e., users who subscribe to the Planet alert services. Planet has been a 'success story'

and numerous versions of the newsletter have now been produced for other organizations, both within and outside the Open University.

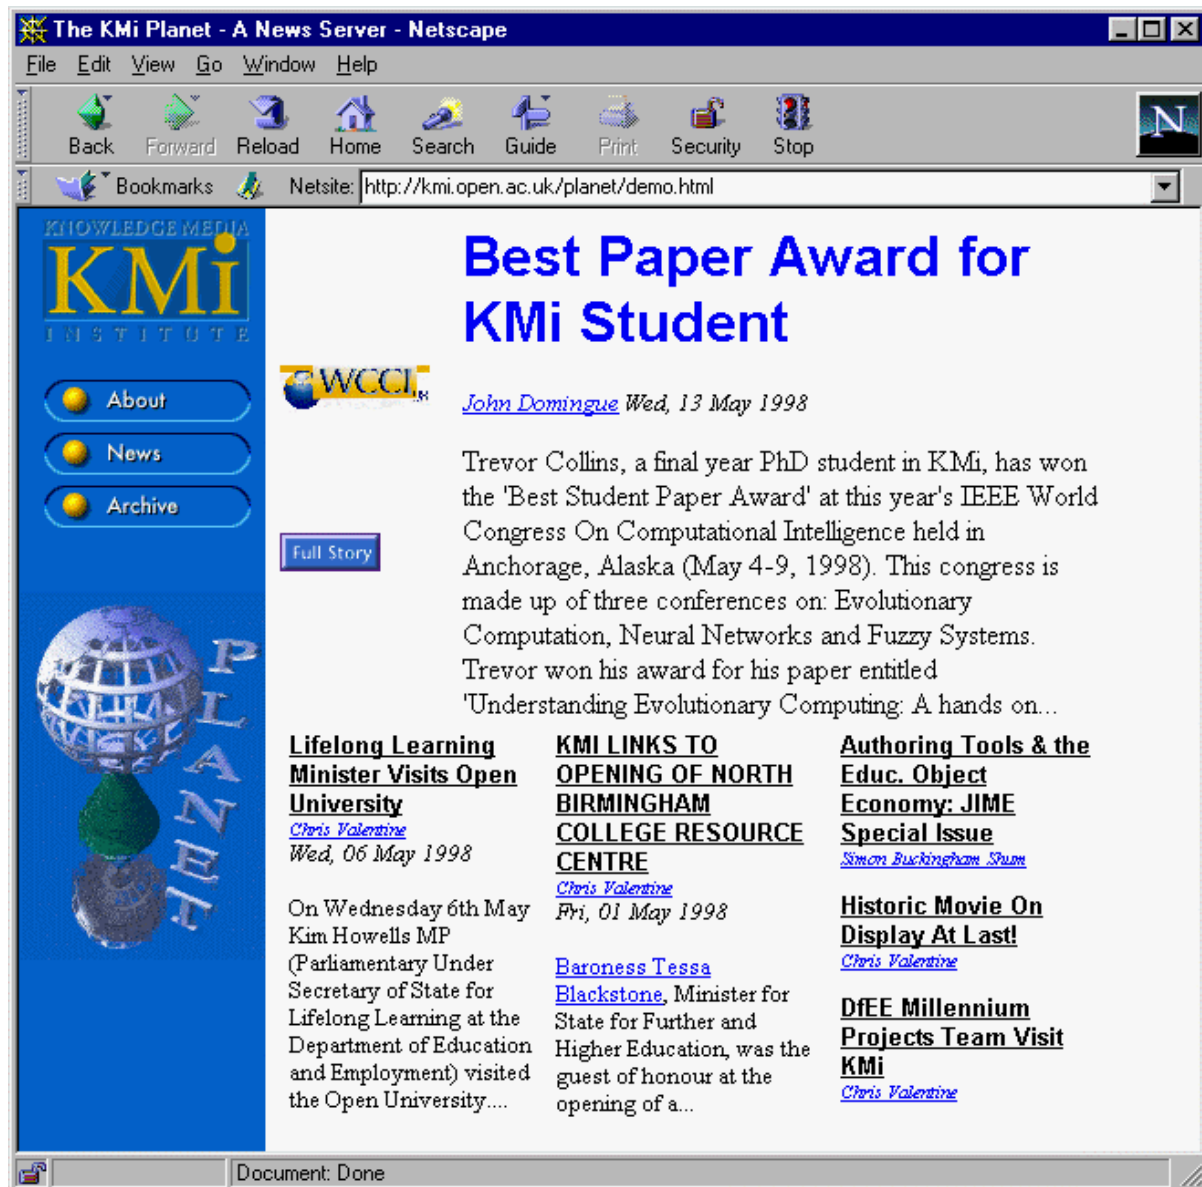


Figure 10. The front page of KMi Planet.

For all this success, it is apparent that, as the Planet archive and readership grow, more sophisticated mechanisms supporting semantic searches and individualised presentations and alerts are needed. Users of Planet often come across interesting news items but they cannot easily follow-on with obvious queries. For instance, having read a story about an award for a paper about visualizing genetic algorithms, a user may want to find out who else works on software visualization, what other projects are going on in this area, etc. Of course, many answers can be found by browsing our web pages and through standard search mechanisms.

However, even in well-organized sites, many important relations between people, technologies, projects and organizations are often missing, leading to the 'standard' knowledge management problem of finding out who does what and who knows what.




Figure 11. Taxonomy of events in the Planet knowledge base.

In addition to the need for better search and retrieval facilities, the experience of a day-to-day use of Planet over more than two years has highlighted a number of other issues. An obvious one is the need for individualised news feeds and presentations – currently, registered readers periodically get a standard news update message. A less obvious issue is the need for Planet to move away from simply being a passive news archive and become a 'real newspaper', where news is not just passively received but proactively identified and requested, in accordance with events in the department and the observed interests of the readership. To address these issues we have developed an integrated suite of tools, *PlanetOnto*, which extends the original Planet news server by providing support for ontology-driven document enrichment, integrated browsing and deductive knowledge retrieval, personalised news feeds and alerts, and proactive identification of potentially interesting news items.

4.2 Instantiating the approach

4.2.1 Characterize viewpoint for ontology.

The viewpoint is centred on the range of academic events which characterise the life of an academic department. Hence, we have developed a rich taxonomy of events, which is shown in figure 11. This taxonomy includes both very abstract events, such as *event-involving-movement*, which are defined in a highly generic ontology (*event-ontology*), as well as specific events, related to academic life, such as *giving-a-lecture*. In total, we have about 40 different types of events, of which 25 have been defined specifically for the KMi Planet application. In addition, we have identified five key classes, which have driven both the development of the ontology and the design of the Lois interface shown in figure 7. These are *people*, *organisations*, *stories*, *projects* and *technologies*.

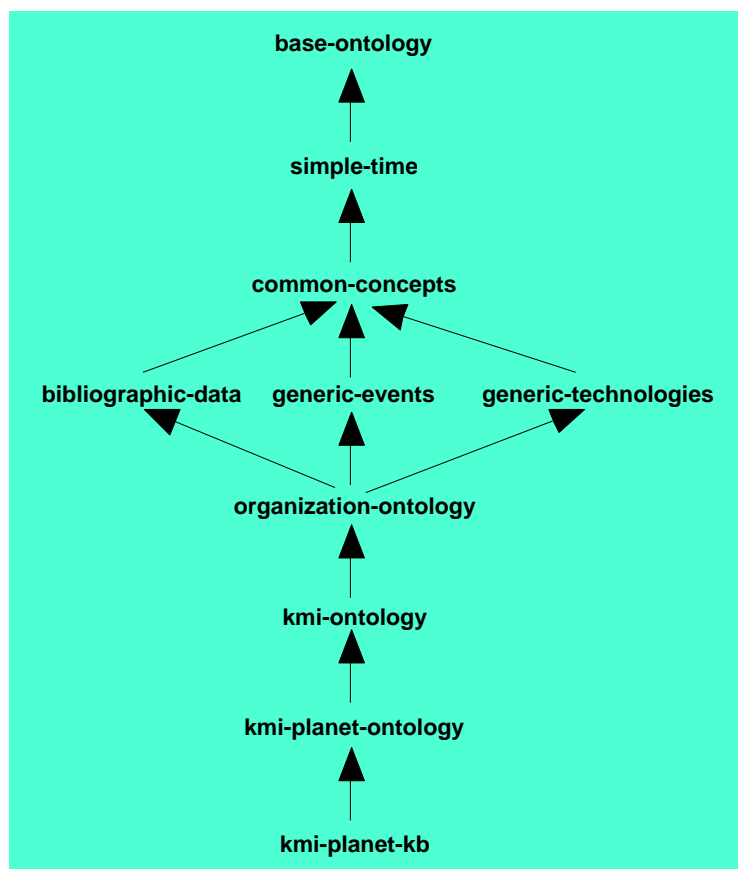


Figure 12. Ontology inclusion in the Planet Ontology.

4.2.2 Develop the Ontology.

The Planet ontology builds on a number of other ontologies included in the library, as shown in figure 12. Some of these ontologies were reused from the Ontolingua repository at Stanford - e.g., *simple-time* and *bibliographic-data*, others were developed from scratch. It is difficult to

quantify the overall development effort, as these ontologies are held in the WebOnto central repository and grow steadily over time. As an indication, we can point out that the first version of the Planet ontology took about two weeks to complete and proved adequate to support the range of queries envisaged in the use scenario.

4.2.3 Perform ontology-driven model construction.

A prototype model was developed using a subset of the stories. This initial prototype was used to gather initial feedback about the technology and the approach, and used as a showcase demo to publicise the approach to KMi visitors. The experience from the first model construction exercise led to the implementation of slot-renaming mechanisms in OCML. A new, complete version of the Planet knowledge base was developed during the summer of 1999. Full-scale trials of the Planet-Onto system will begin in January 2000.

Solutions Found To Your Query

There were 4 solutions to the query.

- [Solution 1](#)
- [Solution 2](#)
- [Solution 3](#)
- [Solution 4](#)

Solution 1

?kmi-member1 [marc-eisenstadt](#) [Find Associated Stories](#)

Solution 2

?kmi-member1 [paul-mulholland](#) [Find Associated Stories](#)

Solution 3

?kmi-member1 [john-domingue](#) [Find Associated Stories](#)

Solution 4

?kmi-member1 [stuart-watt](#) [Find Associated Stories](#)

Figure 13. Solutions to query about academic KMi members active in software visualization.

4.2.4 Customize query interface for semantic knowledge retrieval.

This step consisted of developing the Lois form shown in figure 7. As pointed out earlier, the design of the form was driven by the identification of the key concepts in the ontology. Experience so far shows that the design covers the type of queries which are asked to Planet-

Onto - i.e., queries tend to be centred around key concepts. Figure 13 shows part of the web page produced by the Lois server, when answering the query asked in figure 7. In this case four solutions were found - i.e., four KMi academic members working in the area of software visualization. The user can click on each item to inspect it - i.e., to find what is known about the item in the PlanetOnto knowledge base - and can also try to find stories relevant to the item. Our PlanetOnto knowledge base includes dozens of heuristics to associate items in the knowledge base to stories. For instance, a story is related to a KMi member, say M, if it describes a technology developed by M, or if it talks about a project run by M. Analogous heuristics have been defined to associate stories to technologies, organizations, projects and events. Thus, in this application we use the PlanetOnto knowledge base ‘to enrich’ the Planet newsletter in two main ways: we allow users to find answers to queries which cannot be answered simply by browsing stories and we use the knowledge model as a semantic indexing device for Planet stories.

4.2.5 Develop additional reasoning services on top of knowledge model.

We have designed (but not yet implemented) two intelligent agents, whose purpose is to provide additional services by reasoning about the formalized knowledge model. These are described in the next sub-sections.

4.2.5.1 Story chasing with NewsHound

NewsHound will periodically gather data about ‘popular’ news items and will use these data to solicit potentially popular stories from the appropriate journalists. This will be accomplished by identifying ‘gaps’ in the knowledge base - e.g., projects about which there is no information. The design of NewsHound is consistent with one of the main goals in the design of KMi Planet: the system should try and emulate a news room team. One of the tasks that a news editor carries out is to identify potentially popular stories and assign them to one of the journalists in the staff. NewsHound is meant to emulate this behaviour. In order to identify potentially interesting stories, NewsHound will use two main types of data: statistics on access to individual stories in KMi Planet and records of the queries posed through Lois. Each story within Planet keeps a record of its popularity by counting the number of times the full text is requested from the KMi Planet server. Once NewsHound identifies a story as ‘popular’, then it tries to identify related stories which have the potential to be popular. To perform this task NewsHound analyses the knowledge base trying to find items of interest that have not yet been covered by Planet stories. Typically, these are projects and technologies which i) are known to NewsHound, ii) are ‘related’ to ‘popular’ projects and technologies, but iii) have not yet been covered by a story. The term ‘related’ is the key here. NewsHound will use various heuristics to define ‘relatedness’. For instance, direct subclasses of the same class are considered related; technologies are related if they build on the same underlying technology; projects are related if they tackle the same areas. These heuristics are of course completely ‘soft’ and modular and therefore any new one can be added without affecting the existing ones.

An interesting feature of NewsHound has to do with the unique scenario in which it examines a knowledge base for gaps. Typically, completeness in a knowledge base is defined with respect to logical or task-related properties (van Heijst, 1995). In our scenario incompleteness is defined in pragmatics terms: publications need popular stories.

4.2.5.2 Providing personalised alerts with NewsBoy

Lois has been designed (among other things) to help users to track down Planet stories with very specific characteristics. However, a significant number of users prefer to work with *push technology*, that is they prefer to be automatically notified about potentially interesting stories, rather than having to query Lois about them. We have therefore designed an agent, *NewsBoy*, to provide a mode of use that is complementary to the one supported by Lois. NewsBoy enables users to create a personalised front-page to which interesting stories are 'pushed'.

When a new story is formalised and added to the Planet-Onto repository, NewsBoy matches the story against the specified interests of registered readers. Readers whose interests match that of the story are notified by email that a new story has been added to their personal Planet page. To make an explicit declaration a reader simply specifies a number of queries using the Lois interface. The reader is then updated when a new story matches at least one of the logged queries. Alternatively, a reader can state that she would like NewsBoy i) to log all the queries she makes using Lois and ii) to create a user profile from the log. The resulting user profile is simply the logical disjunction of the queries contained within the log.

It is interesting to compare NewsBoy to other approaches which attempt to infer user profiles from analysing patterns of access to documents – e.g., see (Liebermann, 1995; Krulwich and Burkey, 1997). These approaches try to induce user interests using empirical methods. Our approach is semantic-centred: the user herself specifies the range of documents of potential interest through unambiguous declarative specifications.

4.3 Summing up

The PlanetOnto application historically defined our approach to ontology-driven document-enrichment and provides a showcase for both the technology and the approach. It shows that it is possible to use ontologies to circumscribe a range of phenomena in the workplace, that it is feasible for a motivated community of non-expert users to perform ontology-driven collaborative model development, and it embeds a number of additional facilities which enrich a textual archive with semantic indexing, deductive knowledge retrieval and intelligent agents. In particular, it shows the advantages derived from adopting a modular approach, which starts from highly generic ontologies to arrive to an application-specific knowledge base, whose main role is to support heuristic knowledge retrieval and semantic indexing of documents.

5. SCHOLONTO: SUPPORTING SCHOLARLY DEBATE

5.1 The scenario

Contextualising ideas in relation to the literature is a fundamental task for authors and readers—are they new, significant, and trustworthy? Scholars accomplish this task firstly by bringing to bear their own knowledge of the field. This process leads to *commentary and discourse* of various kinds, which reflect the extent to which peers regard an author’s work as authoritative. These can take the form of private exchanges, formal peer review of conference/journal submissions, or published reviews of literatures and books. We can think of conventional scholarly publication and debate as a document-centred, text-based process. Text is a rich medium in which to publish and discuss ideas in detail and with subtle nuances, but the corresponding disadvantage is that it takes a long time to read and is hard to analyse computationally.

A complementary approach focuses on the *conceptual models* implicit in textual documents and discourse. The goal is to provide a *summary* representation of ideas and their interconnections, in order to assist literature-wide analysis. We believe that this approach has advantages over textual media for tracing the intellectual lineage of a document’s ideas, and for assessing the subsequent impact of those ideas—that is, how they have been challenged, supported and appropriated by others. In addition, the availability of explicit conceptual models opens possibilities for automatic analysis of a community’s collective knowledge.

We begin with the idea that an author’s goal is to persuade the reader to accept his/her *perspective*, which constitutes a set of claims about the world. Usually, the author has some new ideas that s/he is contributing, and asserts particular relationships between these and existing ideas already published in order to demonstrate both the reliability of the conceptual foundation on which s/he is building, and the innovation and significance of the new ideas. The scholarly reader’s task is to understand which ideas are being claimed as new, and assess their significance and reliability. Let us switch from a reading scenario to the scenario of literature search and analysis. In this case, the scholar has some ideas and relationships in mind that s/he is trying to locate in the literature—has anyone written about them, or perhaps these ideas exist but not yet in a single document? The interpretative task includes i) formulating the ideas of interest in ways that may uncover relevant documents, ii) reading the documents (as just described), and then iii) interpreting them to characterize any patterns that appear to emerge. This is a similar scenario to that of a newcomer to a scholarly community - e.g., a student, librarian, lecturer or researcher from another discipline, who wants to know, for instance, what the seminal papers are, or if there are distinctive perspectives on problems or techniques that define that community.

We contend that scholars are very poorly supported in these tasks by conventional library and technological environments, but that digital libraries open up new possibilities which have yet to be exploited. Consider the document interpretation scenario. In the non-digital world, there is

currently no way beyond following citations (only those provided by the author), or using citation indices (to find others citing him for some reason), to ask questions such as: “Has anyone built on the ideas in this paper, and in what way?”, “Has anyone challenged this paper?”, “Has anyone proposed a similar solution but from a different theoretical perspective?”. These are arguably the kinds of phenomena of most interest to scholars when they read or write papers, engage in debate or search the literature. These are also the kinds of questions asked by researchers unfamiliar with a literature. Our goal is to support these kinds of queries by means of our enriched-document approach. Specifically, we have developed an ontology to support scholarly debate and we plan to use this ontology to characterise scholarly relations between documents.

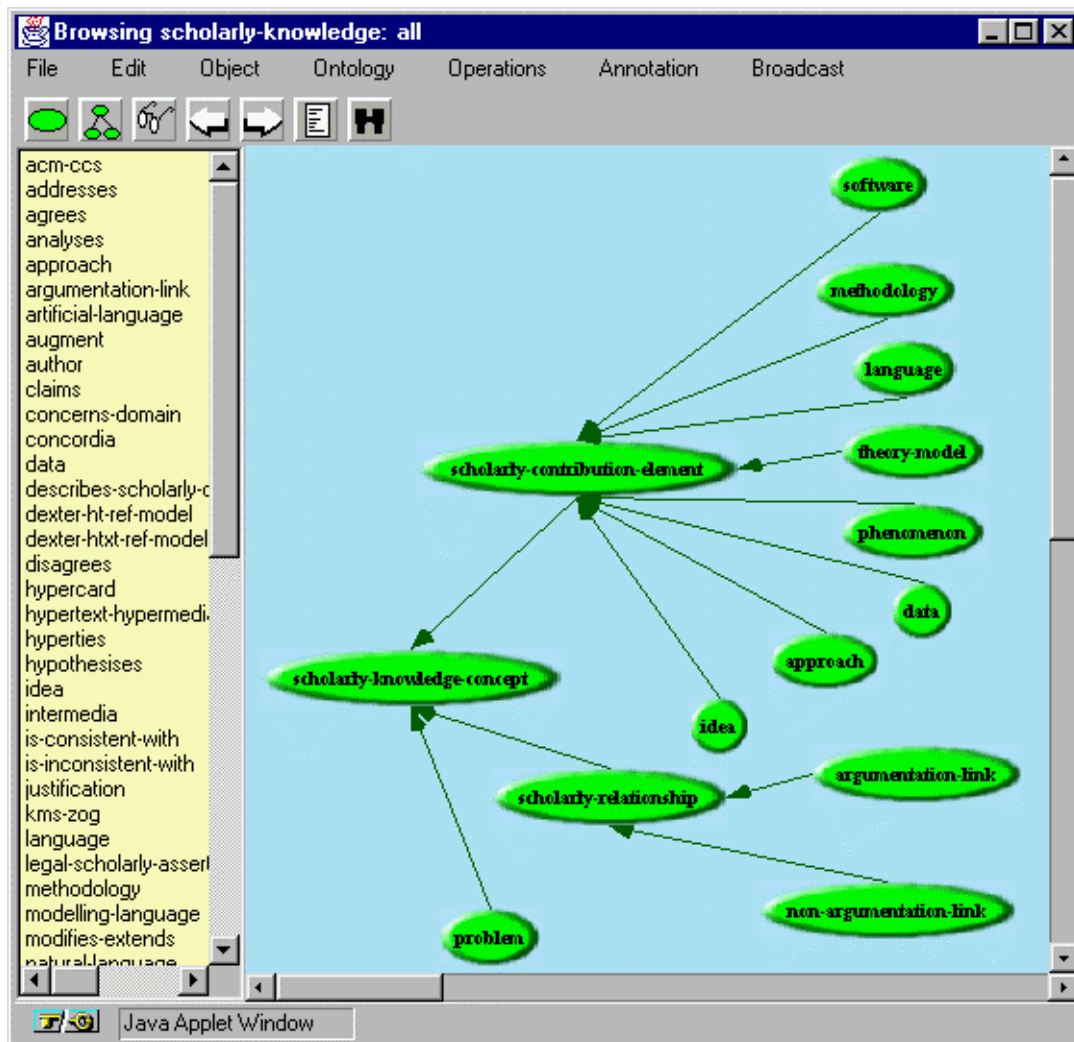


Figure 14. Main taxonomy of the ScholOnto ontology.

This document concerns:

▼ Design Domain Select from ACM CCS keywords (optional)

Hypertext/hypermedia

New...

Citation:
Halasz, F. and Schwartz, M. (1994). Comm. of the ACM, 37 (2), 30-39

URL: www.acm.org/pubs/citations/journals/cacm/1994-37-2/p30-halasz/

▼ ACM CCS

▼ H: Information Systems

H.0 GENERAL

H.1 MODELS AND PRINCIPLES

H.2 DATABASE MANAGEMENT (E.5)

H.3 INFORMATION STORAGE

H.4 INFORMATION SYSTEMS APPLICATIONS

H.5 INFORMATION

This article describes a new:

▼ Theory/Model

Dexter Hypertext Reference Model

Submit...

▼ <none>

Submit...

New...

▼ Theory/Model

▼ Design Domain

Hypertext/hypermedia

Dexter Hypertext Referenc

Relationships to other articles/concepts:

Dexter Hypertext Reference Model

▼ Analyses

▼ Software

Augment
Concordia
Hypercard

▼ Predicts

▼ Software

Theoretically possible Dexter compliant syst

▼ Uses/Applies

▼ Specification Language

Z

Analyses	Software
Addresses	Methodology
Uses/Applies	Language
Modifies/Extends	Theory/Model
Predicts	Phenomenon
Supports	Idea
Raises Issues With	Problem
Refutes	

Sub-Problem of...
Variation on...

New...

▼ Language

▼ Design Domain

State Transition Networks
Task Action Grammer
User Action Notation
VDM

Z

▼ Problem

▼ Design Domain

Hypertext/hypermedia

Commercial hypertexts
...absence of
Educational hypertext
...weak evidence
Navigation
...disorientation
Standards
...absence of

Figure 15. Data entry interface to Scholonto.

5.2 Instantiating the approach

5.2.1 Characterise viewpoint for ontology.

It might appear paradoxical to propose the use of ontologies to support scholarly communities in managing their knowledge, since conflicting worldviews, evidence and frames of reference lie at the heart of research and debate. Of course, the key issue is what is represented. It is hard to envisage when scholarly communities will no longer need to make claims about, or contest, the nature of a document's *contributions* - e.g., "this is a new theory, model, notation, software, evidence", or its *relationships* to other documents - e.g., "it applies, modifies, predicts, refutes...". Our approach builds on this relatively stable dimension of what are otherwise constantly evolving research fields, by representing scholars' *claims about the significance* of ideas and concepts—a focus on *discourse* and *argumentation* (how scholars support and contest claims), and on *context* (the conceptual network in which an idea is embedded). Representing concepts separately from claims about them is critical to supporting multiple perspectives.

5.2.2 Develop Ontology.

Figure 14 shows the top level structure of the ontology. Both nodes and links in the semantic network created by scholars' submissions are *scholarly knowledge concepts*. Nodes are *scholarly contribution elements*, and links *scholarly relationships*. These are subdivided into *argumentation* and *non-argumentation links*. The ontology is designed to support scholars in making *claims* by asserting relationships between concepts. Other scholars may *support, raise-issues-with, or refute* these claims. A claim is formally defined as a relation between a set of authors, who make a legal - scholarly - assertion, with some justification.

The design of the ontology was based on the analysis of scholarly articles from a range of different fields, and took about two person weeks' effort. Once the top level structure stabilised, it required only two days to develop the first version of the ontology in OCML.

5.2.3 Perform ontology-driven model construction.

We are currently seeding a knowledge base with document descriptions to test the ontology's ability to support the scenarios described in this paper. We have also designed a domain-specific version of Knote, which guides users through the modelling schema using dynamic menus, and enables them to browse and search for existing concepts to assist their reuse. This interface is shown in figure 15.

5.2.4 Customize query interface for semantic knowledge retrieval.

A Lois form has been designed, which is shown in figure 16. In the example the user is trying to assess the motivation behind, and impact of, a theoretical model called the *Dexter Hypertext Model*. The queries specify, respectively, i) what problems does it analyse? ii) are there any theory-models which modify-extend it? and iii) is there any software which uses or applies it? Analogously to the use of heuristic search patterns in PlanetOnto to support semantic retrieval of

news items associated with entities in the knowledge base, we plan to provide links to online bibliographic databases, to automatically retrieve papers which exemplify a certain approach/technology/idea/etc. This will provide a much more powerful bibliographic search tool than the conventional, keyword-centred ones, which are currently available in digital libraries.

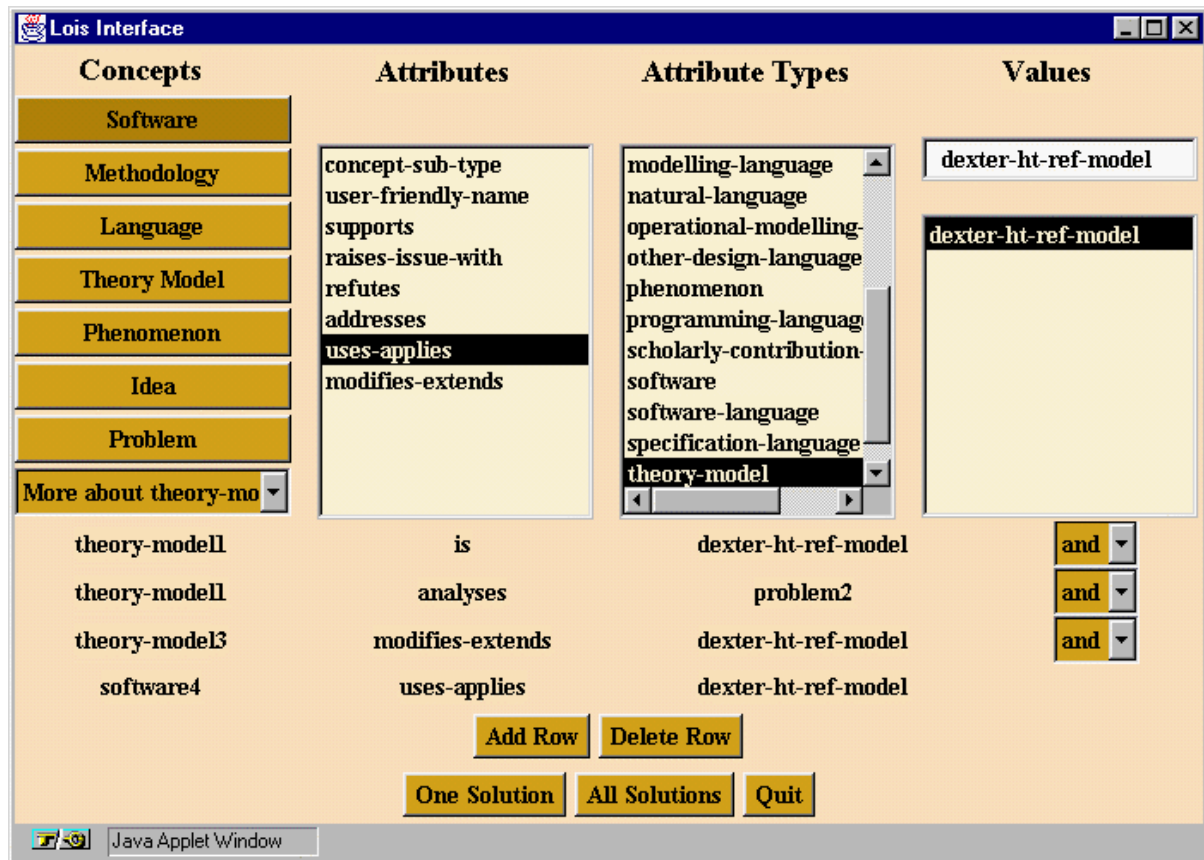


Figure 16. Lois interface to Scholonto.

5.2.5 Develop additional reasoning services on top of knowledge model.

A knowledge model enables inference-based searching and alerting. It will be possible to ask the system questions such as “What impact did Theory T have?”, since “impact” can be defined, for example, in terms of the number of subsequent documents using or modifying the theory, the number of different domains in which it has been applied, the number of problems addressed which drew on the theory, and so forth. Our knowledge modelling environment makes it simple for us (as system maintainers) to write heuristics that could assist in finding relevant documents - e.g., “if Method Y extends Method X, and Method X is challenged, then Method Y may be challenged”. Moreover, as already mentioned, it will be possible to develop specialized agents whose goal is to identify emerging perspectives, using heuristic knowledge and machine learning techniques. As these machine-discovered assertions are added to the knowledge model, software agents effectively become actors in the scholarly debate. This scenario provides an example of a

hybrid agent community and therefore raises a whole host of interesting issues, from the dynamics of social interaction to the design of epistemic agents - see (Masterston, 1998) for an another example and for a detailed discussion of the relevant issues. More in general, the scenario is an example of the general trend towards reducing the boundaries between humans and machines (Stutt and Motta, 1998).

6. KNOWLEDGE MANAGEMENT OF MEDICAL GUIDELINES

The area of medical guidelines provides a very good example of document-centred knowledge work. Guidelines are produced as the result of extensive discussions involving large numbers of experts and stakeholders, discussions which are based on both practical experience and existing literature and which result in ever refined and updated versions of collaboratively produced documents. The aim of the EC-funded PatMan project (PatMan, 1998) is to develop a number of technologies to support guideline-centred patient management. These technologies include an editor/interpreter for medical guidelines, which has been produced by the Medical Informatics group at the University of Pavia, as well as a configuration of our knowledge management technology for this domain.

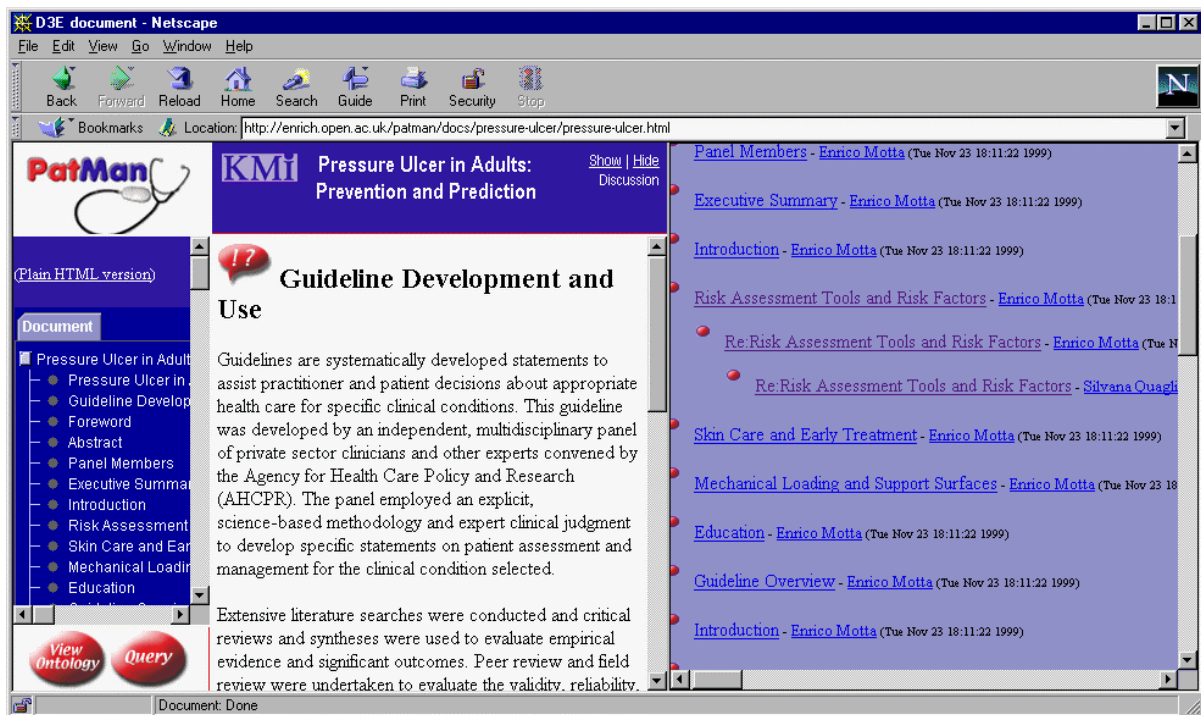


Figure 17. The PatMan repository for enriched guidelines.

Figure 17 shows the *repository of enriched guidelines*, which we have developed in the course of the PatMan project. This repository integrates the D3E support for discourse and debate (Sumner and Buckingham Shum, 1998) with the ontology-driven approach to document enrichment illustrated in this paper. Thus, we aim to support various knowledge-intensive tasks, both during

the design and the application of a guideline. For the sake of explanation, we will start our discussion by illustrating the use of D3E technology in this domain.

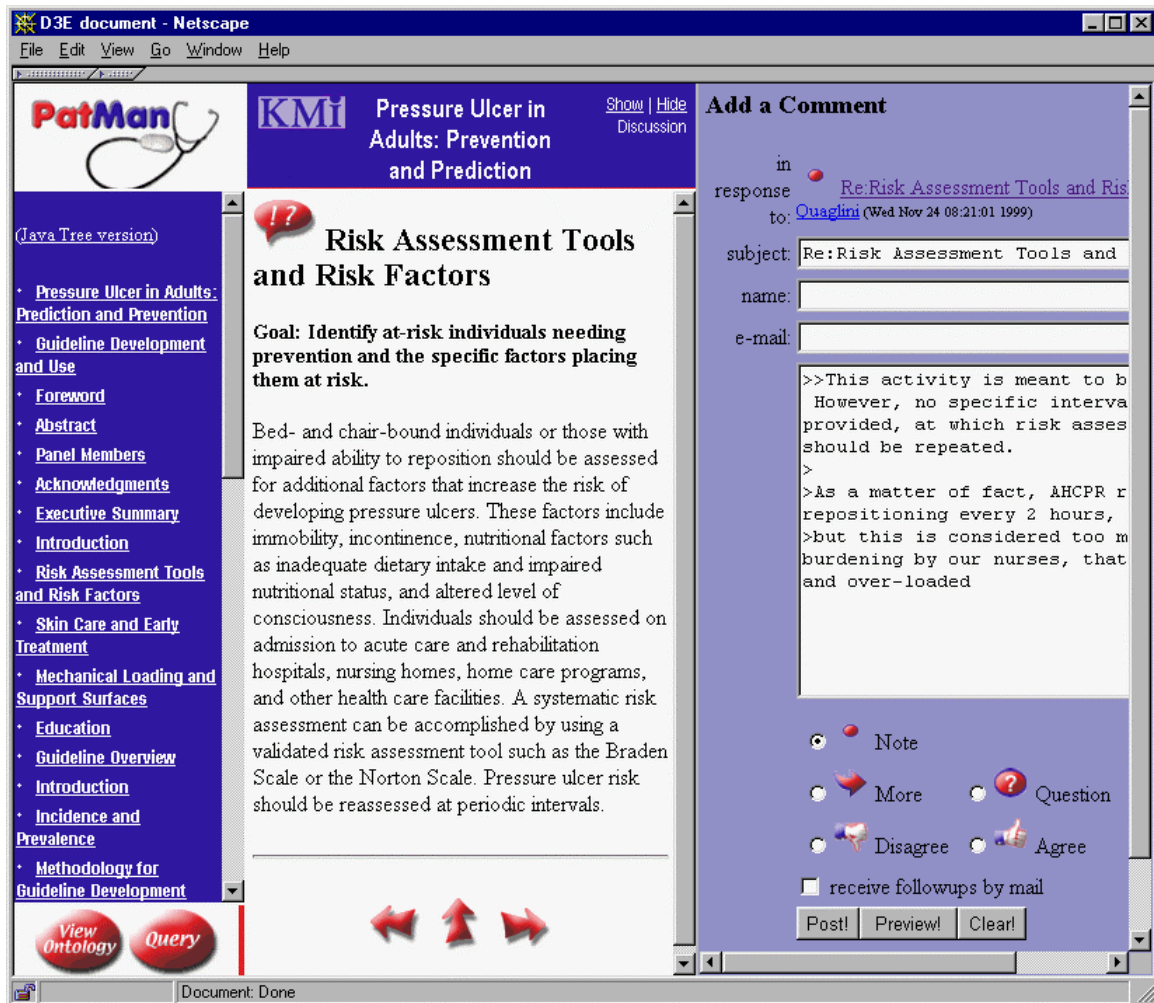


Figure 18. Adding a comment to the discussion space associated with the risk assessment task.

6.1 Supporting discourse and debate in the area of medical guidelines

The D3E technology for discourse and debate can be used to facilitate both the design and the application of a guideline. During the design phase we can use multiple discussion spaces to structure and record the debate associated with the drafting of a document; in the healthcare practice, it can be used (for instance in association with an executable model of the guideline in question) to highlight user problems, modifications, feedback and discussions about those parts of the guideline which are under-specified, either because they are deliberately left open for customization in the local setting, or because of a bug in the specification.

The frame shown in figure 17 was produced using the CEDAR toolkit, which allows publishers to specify the layout of the target D3E environment using a high-level interface toolkit. The

window pane in the middle shows part of the document describing a guideline for the prediction and prevention of pressure ulcer, which has been produced by the Agency for Health Care Policy and Research (AHCPR) in the United States. The window pane on the left provides an outline of the main document, to be used as a navigation device. On the right it is possible to see the various discussion spaces associated with the guideline document. By clicking on a particular discussion space, the user can see the thread of the discussion associated with that particular topic (or document section). Once a discussion space has been selected, the user can add a comment, using a simple set of discussion-structuring primitives - see figure 18.

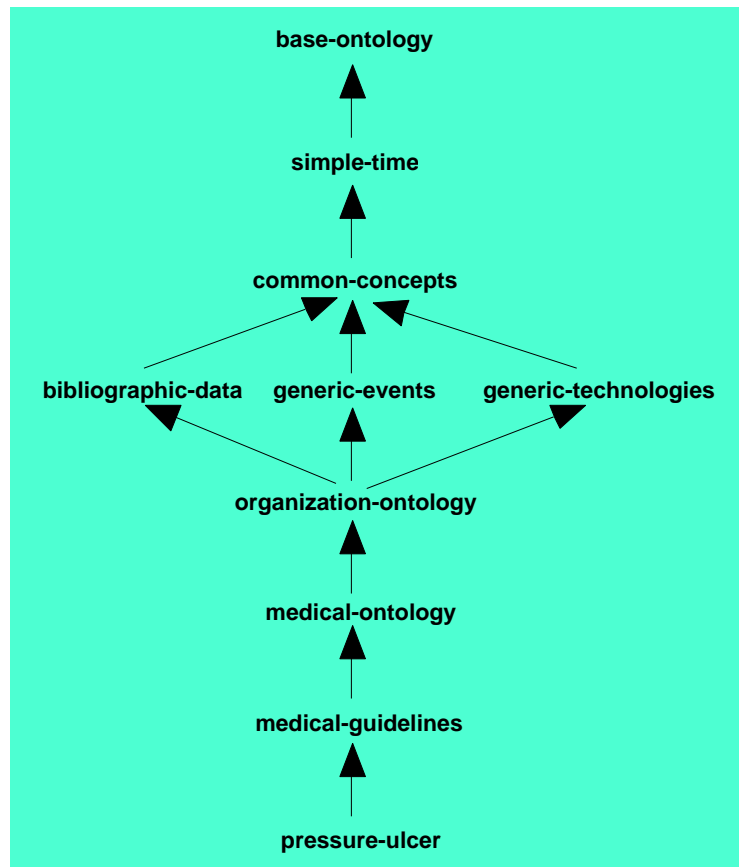


Figure 19. Ontology inclusion in the medical-guidelines ontology.

6.2 Ontology-driven document enrichment in the area of medical guidelines

6.2.1 Characterize use scenario.

The use scenario in the medical guideline domain is very similar to the PlanetOnto one. We want to associate a knowledge model to a set of documents, to provide intelligent knowledge retrieval and semantic indexing facilities. In addition, the availability of a formalized knowledge model opens to the door to the development of additional tools, such as validation and verification tools customized for medical guidelines. Finally, we want to integrate the document-centred support provided by the D3E technology with the knowledge modelling technology. This last task has

been tackled simply by augmenting the D3E-style window frame shown in figures 17 and 18 with two buttons, which allow the user to browse the knowledge model associated with the current guideline and to query it using the Lois interface.

6.2.2 Characterize viewpoint for ontology.

Our viewpoint here is simply ‘medical knowledge’. In other words, when modelling guidelines we focus mainly on the various knowledge types associated with a guideline (such as the associated medical condition, the target population, the classes of users, etc), rather than the algorithmic aspects. The latter have been addressed in the PatMan project by our colleagues at the University of Pavia, who have developed an editor/interpreter for medical guidelines.

6.2.3 Develop Ontology.

As shown in figure 19, the medical-guidelines ontology builds on eight generic ontologies, seven of which are also used by the PlanetOnto model described in section 4. The odd one out here is the generic medical ontology developed in an earlier medical informatics project (HC-ReMA, 1997), of which the medical-guidelines ontology is a direct specialization.

When developing the medical-guidelines ontology we took the notion of medical guideline as our starting point and most of the effort was spent in characterizing this class and the associated classes and relations - e.g., *outcome-measure*, *guideline-user-type*, etc. Because the users are expected to be health professionals we do not envisage the need to provide ‘deep’ medical models, beyond the formalization associated with a medical guideline.




Figure 20. Subclass-of hierarchy for class generic-planning-entity.

As shown in figure 20, a medical guideline is modelled as a subclass of class `plan`. Hence, the slots associated with a plan specification (i.e., the slots typically used by guideline interpreters) are kept separate from additional information about the guideline. We distinguish between three types of guidelines: preventive, therapeutic and diagnostic.

A medical guideline is characterized in terms of seventeen slots. Here we describe a subset of them, to illustrate the type of information associated with a medical guideline in our model. The complete ontology can of course be browsed on the web, through the WebOnto editor/browser.

- **Has-Guideline-User-Type.** This specifies the class of users which carry out the procedures specified in the guideline. Guidelines can be specified for specific categories of users, e.g., nurses or physicians, or can include procedures targeted at different categories - e.g., some to be carried out by nurses and some to be carried out by physicians. A guideline user type can be patient, software-agent, or a subclass of class `generic-care-giver`, whose subclass-of hierarchy is shown in figure 21. It is important to note that class `guideline-user-type` is in effect a *metaclass*, i.e., a class whose instances are other classes.
- **Associated-Documents.** The documents describing the guideline. These are modelled as instances of class `document-reference`, which is defined in the ontology `bibliographic-data`.
- **Location-Constraints.** Where to carry out the procedures carried out in the guideline - e.g., whether at home, at the local surgery, at a day hospital, etc.
- **Temporal-Constraints.** Whether the guideline as to be carried out at particular times, or before or after other procedures.




Figure 21. Subclass-of hierarchy for class `generic-care-giver`.

- **Associated-Medical-Condition-Class.** The type of medical condition associated with the guideline. For instance, the guideline “Pressure Ulcer: Prevention and Prediction” is associated with the condition “Pressure Ulcer”. In general, a guideline can be associated with more than one type of medical condition.

- **Target-Population.** The population targeted by the guidelines. For instance, the pressure ulcer guideline is targeted at people at risk from pressure ulcers - e.g., people with limited mobility. Instances of this slot should be members of class *Population-Specification*.
- **Has-Plan-Specification.** An executable specification of the guideline. This is the part which is needed by a guideline interpreter to implement computerised workflow support. The class *plan-specification* provides a simple control language to model the algorithmic aspects of a guideline, by means of operators such as *one-of*, *one-or-more-of*, *sequential-all-of*, *non-ordered-all-of*, etc.
- **Has-Subcomponent.** A guideline normally includes a number of procedures, sub-plans and possibly other guidelines. We use the term “sub-component” to refer to a generic component of the guideline, which is meaningful from a planning point of view. The range of this slot is restricted to instances of class *Generic-Planning-Entity*, whose subclass-of hierarchy is shown in figure 20.

The *medical-guideline* ontology contains about 120 definitions, including classes, relations and instances. If we count also the definitions inherited from super-ontologies, then the resulting model comprises over 1000 definitions.

```
(def-instance prevention-of-pressure-ulcer preventive-guideline
  ((has-author ahcpr)
   (has-subcomponent
    pressure-ulcer-risk-assessment
    provision-of-mechanical-loading-and-support-surfaces
    pressure-ulcer-skin-care-and-early-treatment)
   (outcome-measure pressure-ulcer-incidence)
   (has-main-goal "predicting and preventing pressure ulcers")
   (target-population people-at-risk-from-pressure-ulcer)
   (full-name "pressure ulcer in adults: prediction and prevention")
   (associated-medical-condition pressure-ulcer)
   (associated-documents ahcprpub92-0047)
   (has-guideline-user-type generic-care-giver)))
```

Figure 22. Definition of the pressure ulcer guideline.

6.2.4 Perform ontology-driven model construction.

As already mentioned we have tested the *medical-guidelines* ontology in the domain of pressure ulcer. The purpose of the pressure ulcer guideline is to help identify adults at risk of pressure ulcers, to define early interventions for prevention, and to manage Stage I pressure ulcers, the weakest form of pressure ulcers.

```
(def- i n s t a n c e  p a t i e n t - r e p o s i t i o n i n g  h e a l t h - c a r e - p r o c e d u r e
  ((is-performed-at-periodical-intervals? true)
   (has-min-periodicity 2-hours)
   (activity-of provision-of-mechanical-loading-and-support-surfaces)
   (responsible-agent-type generic-care-giver)
   (uses-healthcare-technique si-de-lying-position)
   (requires-healthcare-instrument lifting-device
    pressure-reducing-device-for-beds
    pressure-reducing-device-for-heels
    positioning-device)))
```

Figure 23. Definition of instance patient-repositioning.

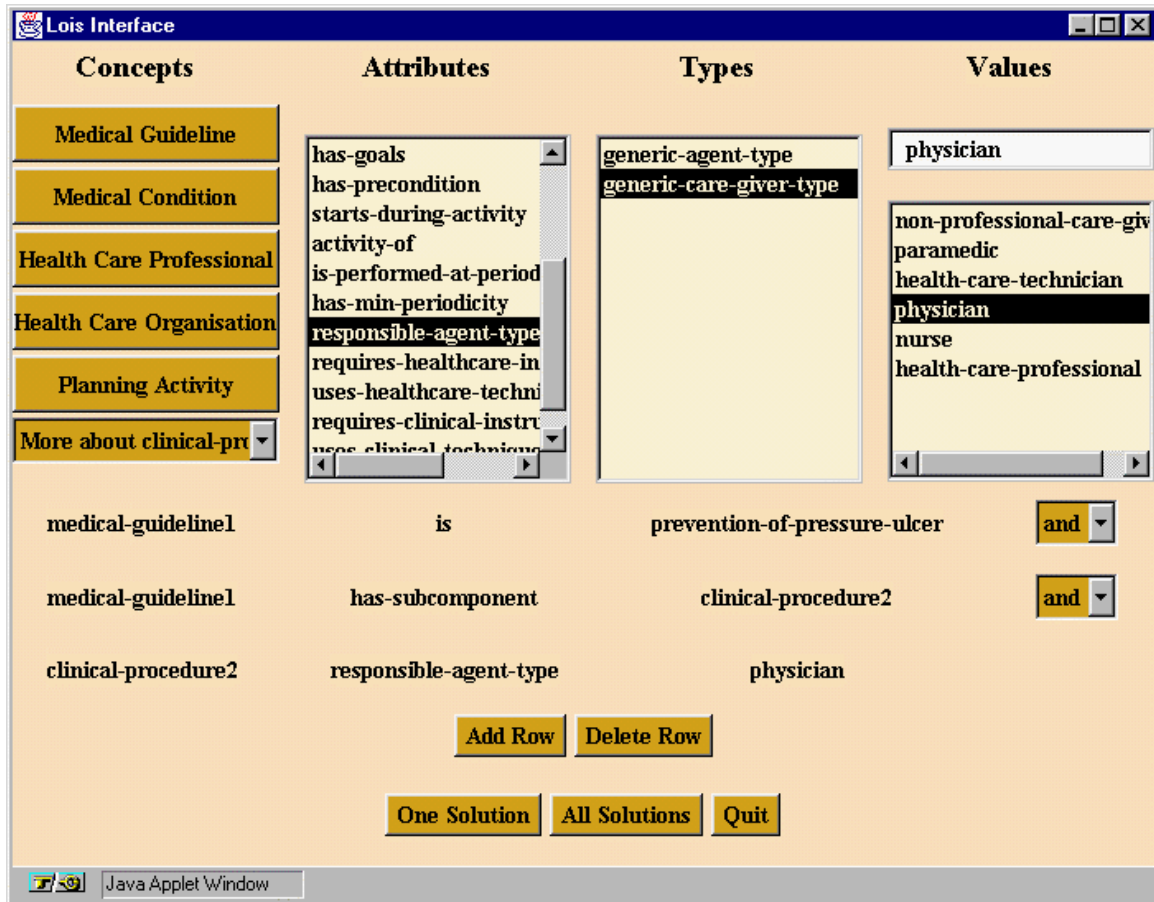


Figure 24. Finding activities in the pressure ulcer guideline needing medical intervention.

A new ontology, called pressure-ulcer, was created on the WebOnto repository and a model of the guideline was created, by instantiating the generic definition of class medical-guideline and by defining the additional entities specifically associated with the guideline. These included the guideline components (e.g., “skin inspection” and “risk assessment”), the notion of ‘pressure ulcer’, its various stages, the notion of ‘people at risk from pressure ulcer’, and a number of healthcare techniques and instruments associated with pressure ulcer. In total, about forty new

definitions were created. The OCML definition of instance `prevention-of-pressure-ulcer` is shown in figure 22⁴. The definition states (among other things) that the guideline has been produced by the AHCPR agency, that it has three direct sub-components and that it is targeted at generic care givers - i.e., it is not specific for a class of healthcare professionals.

Finally, we show in figure 23 the definition of one of the sub-activities specified by the pressure ulcer guideline. In total, about a dozen sub-activities are included in the guideline specification.

6.2.5 *Customize query interface for semantic knowledge retrieval.*

This step consisted of producing the Lois form shown in figure 24. As usual, the form was generated automatically once the ‘key classes’ in this domain were identified. The assumption here is that typical questions to the knowledge model will centre on some medical condition, guideline, or healthcare professional or organization. The “Planning Activity” button makes it possible to ask question directly centred on a particular medical procedure, independently of a medical condition or guideline. Specifically, the query shown in figure 24 asks for activities in the pressure ulcer guideline, which require the intervention of a physician.

As in the PlanetOnto scenario, heuristic rules associate classes of items in the knowledge model (e.g., medical conditions) to relevant sections of associated guideline documents. For instance, clicking on the sentence “Take me to the relevant guideline” next to the string “pressure-ulcer-risk-assessment” in figure 25, takes us back to the PatMan enriched repository, configured so that the risk assessment section and associated discussion site are selected.

6.2.6 *Develop additional reasoning services on top of knowledge model.*

Our main priority for the next few months is to integrate the PatMan configuration of our knowledge management technology with a suitable guideline editor/interpreter. Once this goal is accomplished we plan to develop intelligent validation and verification tools, able to identify errors and omissions in the guideline specification, prior to its transformation to a format suitable for the interpreter.

7. RELATED WORK

From a representational and technical perspective, our approach differs from metadata-centred approaches - e.g., RDF (1999), in that ontology specification languages provide more powerful modelling capabilities than the formalisms currently available in the metadata community - e.g., for specifying relations. Moreover, languages such as OCML and Ontolingua also provide metalevel modelling support which makes it possible to reason about the ontology itself.

⁴ For the sake of brevity not all slots are shown in figure 22. Again, the eager reader can check the full specification of this instance on the WebOnto server.

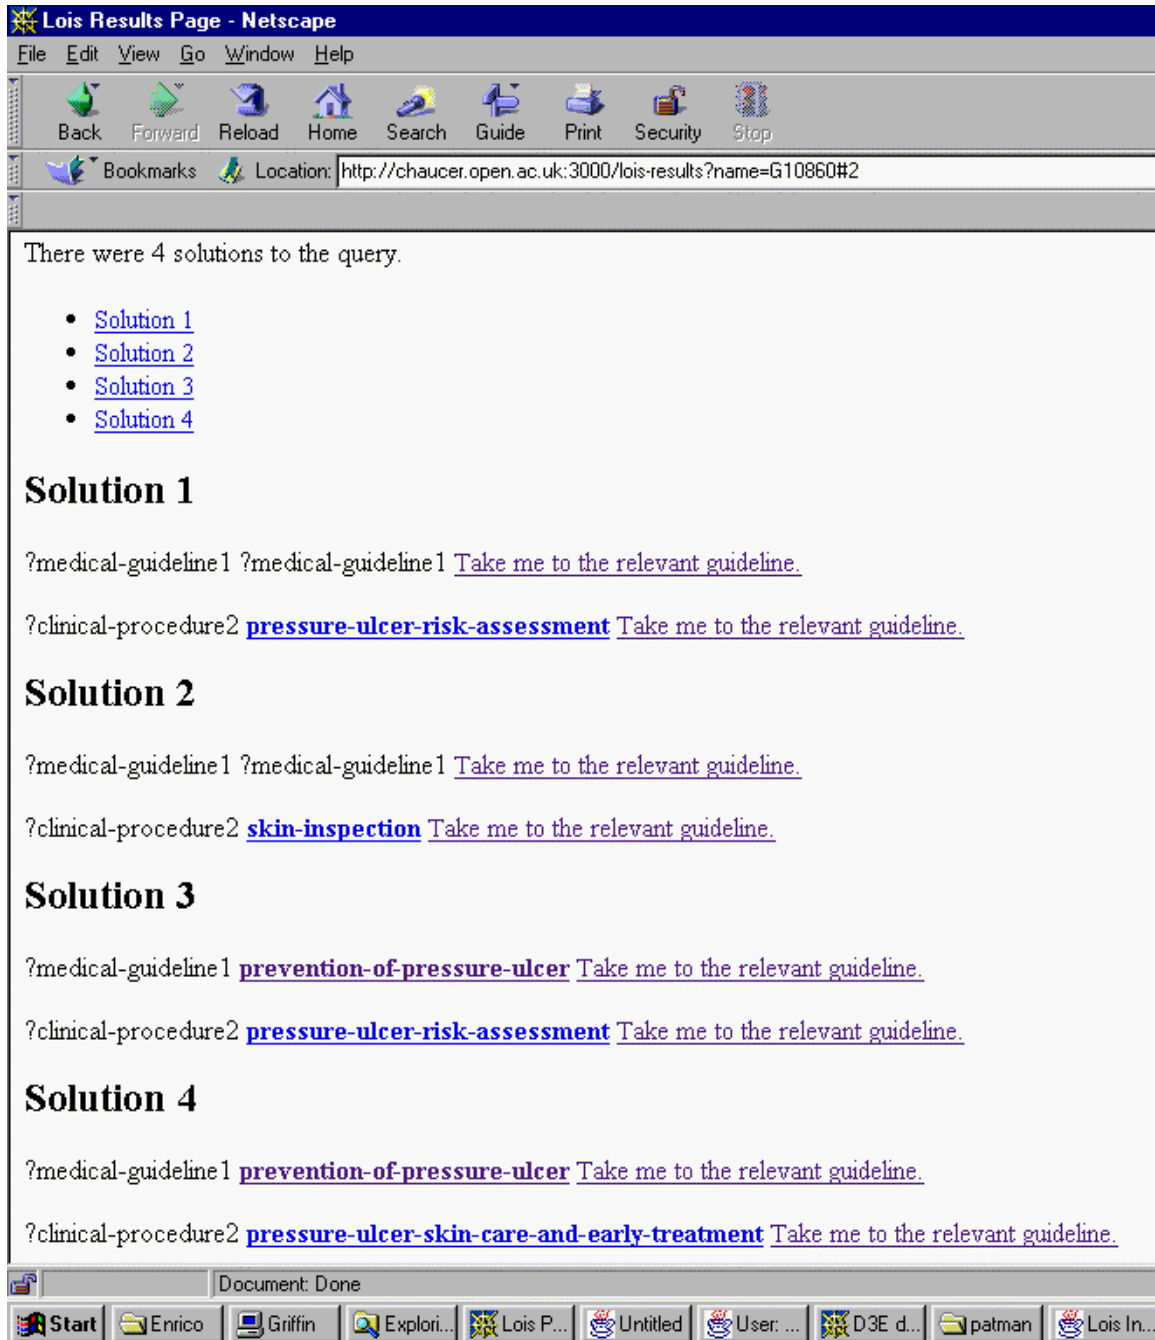


Figure 25. Solutions to the query shown in figure 24.

The Shoe project (Heflin et al., 1998) has proposed an extension to HTML to allow the specification of ontological information. The project team has also developed an editor to support the page annotation process. This work is mainly at the infrastructure level. That is, they suggest a mechanism to allow the representation of information and provide tools to edit and retrieve it. We take a holistic approach to document enrichment and we look at the wider issues concerning usability and sustainability. Thus, we are not just concerned with providing a

mechanism for associating knowledge structures to text but we wish to develop a comprehensive architecture addressing all the relevant issues, from the ‘right’ approach to ontology development to the required visualization and interface tools needed to facilitate the development of knowledge models. Having said so, the technical solutions provided by Shoe could be easily integrated within our framework. For instance OCML structures could be represented in terms of the relevant Shoe tags.

The (KA)2 initiative (Benjamins et al., 1998) shares a number of commonalities with our work. As in most of our scenarios the aim of (KA)2 is to allow a community to build a knowledge base collectively, by populating a shared ontology. In the case of (KA)2 the knowledge base is meant to document the activities of the knowledge acquisition community. Similarly to the approach used in Shoe the knowledge base is constructed by annotating web pages with special tags, which can be read by a specialised search engine cum interpreter, Ontobroker (Fensel et al., 1998). In this paper we have emphasised that the feasibility of the idea of a collective construction of a knowledge base crucially depends on a number of features, including: i) a carefully defined ontology; ii) an underlying modelling language providing user-oriented facilities, such as context-dependent renaming; iii) a user-friendly ontology instantiation environment; and iv) the right motivational stimuli for the participants. In their paper on the (KA)2 initiative, Benjamins et al. (1998) mainly focus on the latter issue. However it seems to us that a careful analysis of all the issues associated with collaborative knowledge modelling is required, in order to manage the risks associated with such enterprises. In particular we believe that a careful design of the underlying ontology is particularly important. For this reason, in contrast with the case of (KA)2, the design of the ontology (but not the ontology population process) is centralised in our approach.

In terms of the underlying architecture the main difference between our approach and other approaches to adding semantic information to web pages is that we decouple the web pages from the knowledge model. This means that we do not directly annotate web pages but the collaboratively constructed knowledge base is held centrally in a server. As a result, changes made to either the ontology or the knowledge base are immediately available to the users. Moreover, decoupling knowledge models from documents means that the relation does not need to be one-to-one. For instance, we have seen how the PlanetOnto knowledge base extends the underlying Planet ontology by providing heuristics for matching entities in the knowledge base to news items. Our approach allows other users to develop alternative knowledge bases (by adopting different heuristics or by subscribing to a different ontological viewpoint) and therefore providing a different set of semantic indexing mechanisms for the same archive.

Finally, decoupling knowledge models from documents also emphasizes that the latter are not the exclusive source of knowledge for the former. We see formalized knowledge servers as playing a similar role to audio, video and ‘vanilla web’ servers. Each type of server plays a distinct role and provides a distinct set of services. A key to successful knowledge management is to

integrate these different media to provide the appropriate services in the relevant scenarios. We refer to the collection of these services as *knowledge media*.

8. CONCLUSIONS

In this paper we have described an ontology-centred approach to knowledge management. The approach and the underlying technology have attracted considerable interest and at the moment we have seven ongoing projects which are testing out this approach in domains ranging from students' help desks to managing best practice in the aerospace industry, to engineering design. In this paper we have tried to highlight the key issues and we have emphasised that the success of these enterprises crucially depends on the successful management of a number of user-centred, task-centred and model-centred issues. 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 success of knowledge management projects.

ACKNOWLEDGEMENTS

The KAW-99 reviewers provided many insightful and stimulating comments on an earlier version of this paper. Silvana Quaglini contributed to the specification of the ontology for medical guidelines; Oscar Corcho Garcia produced the current versions of the ontologies and knowledge bases for the PlanetOnto application; Marek Hatala configured the PatMan repository shown in figures 17 and 18. This research was partially supported by the CEC-funded Enrich (P29015) and PatMan (P4017) projects. Enrich is part of the ESPRIT programme on IT for Learning and Training in Industry; PatMan is part of the Healthcare Sector of the Telematics Application Programme.

REFERENCES

- Benjamins, R., Fensel, D. and Gomez Perez A. (1998). Knowledge Management through Ontologies. In U. Reimer (editor), *Proceedings of the Second International Conference on Practical Aspects of Knowledge Management*. 29-30 October, 1998, Basel, Switzerland.
- Buckingham Shum, S., Motta, E., & Domingue, J. (1999). Representing Scholarly Claims in Internet Digital Libraries: A Knowledge Modelling Approach. In S. Abiteboul and A.-M. Vercoustre (Ed.), *Proc. of EC DL'99: Third European Conference on Research and Advanced Technology for Digital Libraries*. Paris, France, September 22-24, 1999: Springer-Verlag (Lecture Notes in Computer Science). Available at: <http://kmi.open.ac.uk/projects/scholonto/>.
- Domingue, J. (1998). Tadzebao and WebOnto: Discussing, Browsing, and Editing Ontologies on the Web. In B. Gaines and M. Musen (editors), *Proceedings of the 11th Knowledge Acquisition for Knowledge-Based Systems Workshop*, April 18th-23th, Banff, Canada. Available online at <http://kmi.open.ac.uk/people/domingue/banff98-paper/domingue.html>.
- Domingue, J. B. and Motta, E. (1999). A Knowledge-Based News Server Supporting Ontology-Driven Story Enrichment and Knowledge Retrieval. In D. Fensel and R. Studer (editors), *Proceedings of the 11th European Workshop on Knowledge Acquisition, Modelling, and Management (EKAW '99)*, LNAI 1621, Springer-Verlag..

- Domingue, J. and Scott, P. (1998). KMi Planet: Putting the Knowledge Back into Media. In M. Eisenstadt, and T. Vincent, (editors), *The Knowledge Web: Learning and Collaborating on the Net*, Kogan Press, pp. 173-184.
- Duineveld, A. J., Stoter, R., Weiden, M. R., Kenepa, B. and Benjamins, V.R. (1999). Wondertools? A comparative study of ontological engineering tools. Twelfth Workshop on Knowledge Acquisition, Modeling and Management. Banff, Alberta, Canada. 16-21 October, 1999.
- Eisenstadt, M., Domingue, J., Rajan, T., and Motta, E. Visual Knowledge Engineering. *IEEE Transactions on Software Engineering*, 16(10), pp. 1164-1177, October 1990.
- Eriksson, H., Puerta, A. R. and Musen, M. A. (1994). Generation of knowledge-acquisition tools from domain ontologies. *International Journal of Human Computer Studies*, 41, 425-453.
- Eriksson, H., Ferguson, R. Shahar, Y. Musen, M. A. (1999). Automatic Generation of Ontology Editors. Proceedings of the 12th Banff Knowledge Acquisition Workshop, Banff, Alberta, Canada, October 16-22, 1999.
- Fensel, D. Decker, S., Erdmann, M. and Studer, R. (1998). Ontobroker: The very high idea. *Proceedings of the 11th Annual Florida Artificial Intelligence Research Symposium (FLAIRS-98)*.
- Fuchs, N. E. (1992). Specifications are (preferably) executable. *Software Engineering Journal*. pp. 323-334. September 1992.
- Girgenshohn, A., Zimmermann, B., Lee, A., Burns, B. and Atwood, M. E. (1995). Dynamic Forms: An Enhanced Interaction Abstraction Based on Forms. *Proceedings of Interact '95*, pp. 362-367.
- Gruber, T. R. (1993). A Translation Approach to Portable Ontology Specifications. *Knowledge Acquisition*, 5(2).
- Heflin, J., Hendler, J. and Luke, S. (1998). Reading Between the Lines: Using SHOE to Discover Implicit Knowledge from the Web. *AAAI-98 Workshop on AI and Information Integration*. Available online at <http://www.cs.umd.edu/projects/plus/SHOE/shoe-aaai98.ps>.
- van Heijst, G. (1995). *The Role of Ontologies in Knowledge Engineering*. PhD thesis, University of Amsterdam.
- HC-ReMa (1997). Health-Care Resource Management. *Telematics Applications Project HC 3103*. Project Description available at <http://aim.unipv.it/projects/hcrema/>.
- HPKB (1997). High Performance Knowledge Bases. Darpa Project. Project Description available from <http://www.teknowledge.com:80/HPKB/>.
- Krulwich, B. and Burkey, C. (1997). The InfoFinder Agent: Learning User Interests through Heuristic Phrase Extraction. *IEEE Expert Intelligent Systems and their Applications*, 12(5), pp. 22-27.
- Lenat, D.B. and Guha, R.V. (1990). *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*. Addison-Wesley, Reading, MA.
- Lieberman, H. (1995). Letizia: An Agent That Assists Web Browsing. *International Joint Conference on Artificial Intelligence, IJCAI '95*. Montreal, August 1995.
- MacGregor, R. (1991). Using a Description Classifier to Enhance Deductive Inference. *Proceedings of the 7th IEEE Conference on AI Applications*. Miami, Florida, February 1991.
- Masterton, S. (1998). The Virtual Participant: A Tutor's Assistant for Electronic Conferencing. In Eisenstadt, M. and Vincent, T. (editors), *The Knowledge Web*, Kogan Press.
- Motta E. (1999). *Reusable Components for Knowledge Models*. IOS Press, Amsterdam.
- Musen, M. A. (1989). *Automated Generation of Model-Based Knowledge Acquisition Tools*. Research Notes in Artificial Intelligence, Pitman, London.
- O'Leary, D. E. (1998). Knowledge Management Systems: Converting and Connecting. *IEEE Intelligent Systems*, 13(3), pp. 30-33. May/June 1998.
- Patman (1998). Patient Management Workflow Systems. *Telematics Applications Project HC 4017*. Project Description available at <http://aim.unipv.it/projects/patman/>.
- RDF (1999). <http://www.w3.org/RDF/>.

- Riva, A. and Ramoni, M. (1996). LispWeb: a Specialized HTTP Server for Distributed AI Applications, *Computer Networks and ISDN Systems*, 28 (7-11), 953-961. (also available at <http://kmi.open.ac.uk/~marco/papers/www96/www96.html>).
- Selvin, A. (1999). Supporting Collaborative Analysis and Design with Hypertext Functionality. *Journal of Digital Information*, 1, (4). Available at: <http://jodi.ecs.soton.ac.uk/Articles/v01/i04/Selvin/>.
- Shipman, F. M., & Marshall, C. C. (1999). Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in Interactive Systems. *Computer Supported Cooperative Work*, In Press.
- Sowa J. F. (1995). Top-Level Ontological Categories. *International Journal on Human-Computer Studies*, 43(5/6), pp. 669-685.
- Stutt, A. and Motta, E. (1998). Knowledge Modelling: An Organic Technology for the Knowledge Age. In Eisenstadt, M. and Vincent, T. (editors), *The Knowledge Web*, Kogan Press.
- Sumner, T., and Buckingham Shum, S. (1998). From Documents to Discourse: Shifting Conceptions of Scholarly Publishing. *Proc. CHI 98: Human Factors in Computing Systems*, (Los Angeles, CA), 95-102. ACM Press: NY. Available at: <http://kmi.open.ac.uk/techreports/papers/kmi-tr-50.pdf>.
- Sumner, T., Domingue, J. Zdrahal, Z., Hatala, M. Millican, A., Murray, J., Hinkelmann, K., Bernardi, A., Weiss, S., Traphoner, R. (1998). Enriching Representations of Work to Support Organizational Learning. In *Proceedings of the Interdisciplinary Workshop on Building, Maintaining, and Using Organizational Memories (OM-98)*. 13th European Conference on Artificial Intelligence (ECAI-98), 23-28 August, Brighton, UK. Available at: <http://kmi.open.ac.uk/projects/enrich/enrich-oms98-paper.html>.
- Tanimoto, S. (1990). VIVA: A Visual Language for Image Processing. *Journal of Visual Languages and Computing* 2(2), pp. 127-139, June 1990.
- Uschold M., Clark, P., Healy, M., Williamson, W. and Woods, S. (1996). An Experiment in Ontology Reuse. In Gaines, B. and Musen, M. (editors), *Proceedings of the 10th Banff Knowledge Acquisition for Knowledge-Based System Workshop (KAW'96)*, Banff, Canada, November 9th-14th, 1996.
- Uschold M. and Gruninger M. (1996). Ontologies: Principles, Methods and Applications. *Knowledge Engineering Review*, 11(2), pp.93-136.
- van der Vet, P. E. and Mars, N. J. I. (1998). Bottom-up Construction of Ontologies. *IEEE Transactions on Knowledge and Data Engineering*, 10(4), pp. 513-526.
- Weyhrauch, R. W. (1980). Prolegomena to a theory of mechanized formal reasoning, *Artificial Intelligence (13)1-2*, pp. 133-170.
- Wideman, H.H. and Owston, R.D. (1993). Knowledge Base Construction as a Pedagogical Activity. *J. Educational Computing Research*, 9(2), 165-196.
- XML (1999). <http://www.w3.org/XML/>.