Linking with Meaning: Ontological Hypertext for Scholars

Simon Kampa, Tim Miles-Board, Leslie Carr and Wendy Hall

IAM Research Group, University of Southampton, UK

E-mail: {srk98r, tmb99r, lac, wh}@ecs.soton.ac.uk

ABSTRACT

The links in ontological hypermedia are defined according to the relationships between real-world objects. An ontology that models the significant objects in a scholar's world can be used toward producing a consistently interlinked research literature. Currently the papers that are available online are mainly divided between subject- and publisher-specific archives, with little or no interoperability. This paper addresses the issue of ontological interlinking, presenting two experimental systems whose hypertext links embody ontologies based on the activities of researchers and scholars.

KEYWORDS: Ontological Hypermedia, Navigation, Ontologies, Link Authoring, Semantic Web

INTRODUCTION

Ontology is the study of "things that exist" that began as a branch of philosophy and is now popular in the field of knowledge management [25]. Ontologies are formal models that allow reasoning over concepts and objects that appear in the real world.

Hypermedia is the study of "what can be said" using computer media, databases and links. Hypermedia provides computer-mediated extensions to familiar textual communication. This is important because real-world objects, or the "things that exist", have a complex relationship to each other, and so complex structures are required for expressing and exploring those relationships when we make hypermedia statements about them.

Ontological hypermedia is then the kind of hypermedia whose structure and links are derived from the relationships of objects in the real world. It is closely related to schematic hypertexts that define a linking schema as a design mechanism for structuring document collections [21]; ontological hypertexts extend the schema into the domain of world knowledge (allowing inferencing mechanisms to deduce meaningful document links).

This paper describes two hypertext system prototypes that use versions of a simple ontology for scholars. The ontology models the significant features of the real world with an appropriate degree of detail for that particular community. The distinguishing attribute of scholars is their interest in research: the people who perform it, the projects in which it occurs and the literature in which it is reported. This knowledge allows links to be inferred that may be irrelevant for other communities, *e.g.* a link between two articles by the same scientist may be intrinsically useful where a link to another story by the same journalist may not be useful to the reader of a newspaper.

Since the mid-1930s commentators have assumed that it has been technologically possible to make the whole research literature available through a form of machine storage, with the difficult problems being recognized as knowledge maintenance [47] and knowledge interconnection [6]. We have come some way towards this goal with the World Wide Web (WWW), but experience shows that the research literature is still not available as a 'whole' and is neither comprehensively indexed nor interlinked. Consequently, the process of intellectual research using the Web involves detective work (What did the author of this paper go on to write? What other papers describe this project? Did this work influence any standards or other software?) because the literature is disconnected from itself and from the records and reports of the research activities which produced it.

In this paper we investigate the use of ontologies to improve the linking of research literature together with the Web sites and home pages of projects, institutions and individual researchers, by providing a principled way of describing both the topic under discussion, and the process by which it was produced in order to allow other researchers to better understand the work. We describe two systems, the first, OntoPortal, is a portal structured on a simple ontology, which describes and provides links to external research pages on the Web. The second, E-Scholar Knowledge Inference MOdel (ESKIMO), uses a proxy service to show a linked structure as a researcher navigates the external "research web", and introduces the concept of inferring new scholarly knowledge.

BACKGROUND

The WWW has become a popular publishing medium for scholars in many fields [28, 29]. We are becoming e-Scholars, either through the actions of primary and secondary publishers placing their archives online and adapting to e-commerce opportunities or else because of the actions of researchers themselves in using the Web to extend free access to their own work [33]. E-scholars not only publish their research results in articles on the Web, they also perform their research on the Web, a process of systematic investigation to collect information on a topic by reading and searching.

The potential advantages that the WWW can offer the scholar are easy to enumerate: instantaneous access to the entire research literature, completely linked through citation and reference and improved scholarly collaboration including commentary and review. In point of fact the current benefits fall somewhat short of these ideals, partly because of the apparently conflicting rights and responsibilities of authors and publishers of the research literature.

In some disciplines (*e.g.* physics and biomedicine) research papers can be easily accessed by the research community through centralised repositories and can be reviewed by peers months in advance of traditional paper journal or conference proceedings publications [26, 42]. Even in the absence of controlled peer review, the 'invisible hand' may still enforce the continuing quality of WWW-published research [27].

As well as access and speed, the Web is increasingly providing interconnectivity, allowing cited research to be linked so that the e-Scholar is able to easily navigate across the related research literature. The resulting hyper-web enables e-Scholars to quickly become familiar with a research field in terms of its literature, activities and authors, in order to subsequently make pertinent and appropriate contributions.

The emergence of recent standards such as XML [53] and RDF [48] heralds the change of the WWW into a Semantic Web [3]. The Semantic Web can enhance the scholarly Web by annotating resources with semantic knowledge, improving indexing and introducing inference capabilities. Introducing supplementary information to the Web in the form of metadata (knowledge that is relevant to scholars but not the central focus of their reading) allows the e-Scholar to become immersed in the sum of the activities of the research community. Modelling and formalizing this network of resources and relationships enables inference methods to be applied allowing non-trivial questions about a research field to be resolved.

In practical terms, the initial effort in the e-Scholar movement has been publishing the scholarly literature on the Web. E-Journals and Digital Libraries have provided the main platforms on which to accomplish this. Examining some of the leading on-line e-journals and digital libraries [1, 44, 45, 46], it is immediately evident that although a large number of articles have been placed on the Web, comprehensive support to position these with respect to the rest of the literature and the community is lacking.

The expectation of a well-linked hyper-web of knowledge is also lacking. Inside each digital library (as well as outside on the Web) the ability to locate information requires an efficient search engine. The literature is also split between competing libraries run by separate publishers and professional societies. Citation linking across these artificial boundaries improves the e-Scholar's access by linking research papers to cited material [7, 30, 33, 34]. Careful analysis of these citations can reveal research fronts, trends and perspectives [11, 19, 20]. However, the Web itself has been seldom observed to exhibit such associatively linked hypermedia to its full potential [9, 10] - documents that are closely related are rarely linked together, and the e-Scholar must derive these links implicitly through other means.

Even accepting the limited nature of the research Web, hypertext systems may pose problems of disorientation, as well as cognitive and information overload [12, 55]. A number of solutions have been suggested, including the use of navigational metaphors [24, 35], sophisticated overview maps [22, 36, 40], link reduction algorithms [18] and collaborative systems [43], although none have been implemented on a large scale. Although others have argued that these concerns for the hypertext reader are exaggerated [4, 13], the scholar as a hypertext writer incurs a significant effort in the authoring and maintenance of hypertext links [39]. Can an e-Scholar reasonably be expected to link her research metadata appropriately and completely into the wider community and then maintain these links as the community grows and changes?

The vision of a semantic web of scholarly resources is also yet to be fully realised. Pioneering initiatives such as SHOE [31, 38] and (KA)2 [2] have demonstrated how scholars can annotate their WWW documents with additional semantic information to improve indexing and unlock the possibility to inference new knowledge. However, this annotation process adds considerable authoring overhead.

The absence of a comprehensive scholarly web leaves e-Scholars with no alternative than to use traditional and less effective methods of information harvesting. In this paper, we present our contribution to the continuing improvement of the e-Scholar's environment: the design and implementation of a hypertext system using metadata, ontological reasoning and navigation via structured linking.

PRINCIPLES

The design goal of an ontological hypertext is to produce a methodology for a building a hypertext system to improve the navigation facilities available to e-Scholars. We wanted to take advantage of the benefits of both hypermedia and the Semantic Web that are not in widespread evidence on the WWW (such as large-scale associative linking, and the ability to annotate resources with metadata), and produce a semantic hyper-web of scholarly information that encapsulates the knowledge required to become thoroughly immersed in a research field. We describe the fundamental design principles here, before visiting actual implementations of this methodology in later sections.

Navigation

In order to embody the vast knowledge repository that makes up an entire research field, we employed an ontological approach. Ontologies allow us to model real-world domains through explicitly specifying concepts, instances, relations, functions and axioms. By modelling the domain of a research field using an ontology, a hypertext system has the power to intelligently communicate, analyse and reason over knowledge.

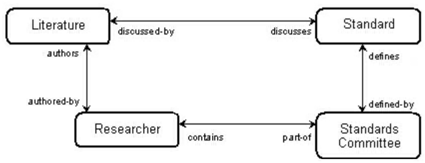

Figure 1 : Ontological modeling of the scholarly activities surrounding a standards committee.

The ontology makes explicit the different types of resources available to a researcher in any field, and formally defines the relationships between them. The ontology is not domain-oriented (it can be applied to any research field) nor simply bibliography-oriented (hence transcending the features available through independently published Digital Libraries), but community- and scholarly activity- oriented. For example, Figure 1 shows how the activities surrounding a Standards Committee in a research field (*e.g.* W3C or ISO) could be modelled.

The design further specifies that rather than keeping this underlying ontological data model hidden from the user, we promote it to the forefront of the interface for exploring the research field that it encapsulates. Using the ontology in the interface layer is only possible if the ontology is relatively focused and intuitive. Using an ontology as a navigation tool has many advantages, as both typed links and a fixed linking structure are enforced. Named relationships between ontological concepts naturally provide a link taxonomy from which the interface can derive presentation techniques for displaying different link types. These relationships also restrict the range of permitted links between different concepts (for example, in Figure 1, 'Literature' can only be linked to 'Researcher' and 'Standard'). Using the underlying ontological model as a navigational metaphor also enforces bi-directional linking between related concepts, and *n*-ary links in the case that the ontological model contains a one-to-many relationship.

Specializing the ontology allows different relationships to be reified for different research communities. For example, computer scientists, physicists and psychologists have very different ideas as to what constitutes 'appropriate' research literature in terms of the acceptability of conference proceedings, technical reports, journal articles or books. A pragmatic approach to modelling "what exists" in the real world and allows the links inferred in the virtual world to reflect those distinctions.

Linking

Typed links [12, 14] address the potential problem of disorientation and cognitive confusion arising from large associatively-linked hyper-sets, since the user is able to predict the effect of traversing a link before the act of traversal has actually taken place.

As the facilities of XLink have yet to make an impact on Digital Library and WWW site design, bi-directional and n-ary linking is seldom exhibited in these environments, and hence users are unaccustomed to the navigational improvements that this linking model offers in the exploration of their research field (Table 1).

Typical Digital Library link	**Navigational possibilities**
Link from research paper to a paper that it cites	Link between all cited research papers Link between multiple versions of that paper, published and e-print, and its reviews and discussion Link between all papers that describe the project/system discussed in that paper
Link from a researcher's homepage to information about their research group	Link between all researchers' homepages in the group Link between that researcher and contact information, personal homepage, and details of the researcher's role in current and previous projects worked on by that research group.

Table 1: Contrasting typical Digital Library links with the navigational possibilities of complex hyperlinking.

Using an ontological navigation tool also allows users to effectively answer queries using a "query-by-linking" approach (rather than "query-by-searching"), using facts that they are able to assert in order to discover new facts through exploration. For example, a researcher having just read a research paper describing a particular standard wishes to find out whether any other papers describe the standard, perhaps with contrasting viewpoints. Using WWW and Digital Library technology as it stands, resolving this type of query usually involves resorting to searching for similar papers with a search engine, using appropriate keywords as the terms ("query by searching"). However, using the ontological navigation, the researcher can realize this query quickly and effectively by exploring a link between all research papers that describe the standard ("query by linking").

Upon finding a particularly interesting paper, the researcher may then decide to navigate to other papers written by the authors. This query may be answered by exploring a link between the paper and all its authors, and then exploring a link between each author and each of the papers they have been involved in writing.

Architecture

The architectural design of the first of our ontological hypertexts (OntoPortal) takes the form of a portal that projects an exploratory meta-layer over the resources applicable to a research field (Figure 2). Normally the scholar (1) sees the WWW as a weakly linked set of resources (2). A process of detective work using search engines must derive relations between resources explicitly. The enhanced e-Scholar (3) views resources on the WWW (5) through the meta-layer provided by the portal (4). The portal collates metadata about these resources, structures it using an ontological model, and presents this metadata as a layer through which the resources and their relationships can be explored.

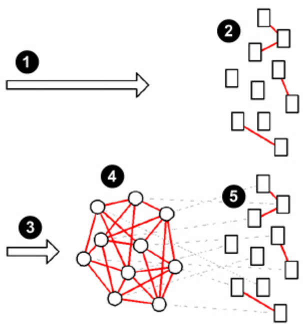

Figure 2 : Instead of seeing WWW resources as weakly linked, the e-Scholar views them through the meta-layer provided by the portal.

The design of the second system (ESKIMO) includes a reasoning engine to infer important facts about the research community from the ontological knowledge base. For example, given the corpus of literature appropriate to a particular research field, the reasoner infers the experts in the field from the community at large, based on a set of heuristics such as the number of research papers published and the number of citations in other research papers. The particular novelty of this approach is in the use of the explicit metadata to generate new forms of knowledge and hence new kinds of links.

ESKIMO is an open environment for applying the semantic meta-layer to the WWW. By embedding an inference engine into the linking and navigation layer, it ensures that the user does not lose the benefits of the semantic meta-layer when an external resource is visited, and the user hence leaves the confines of the meta-layer. That is, ESKIMO is able to derive and work within the context of the WWW resource itself, using this as an entry point into the ontological knowledge base in order to provide knowledge and inferences inline with the resource.

General Related Work

The practice of adding semantics to web resources (using standards such as XML [53], RDF [48], Dublin Core [16], OML [49], and OIL [17]) in order to build a knowledge base has been the focus of previous research.

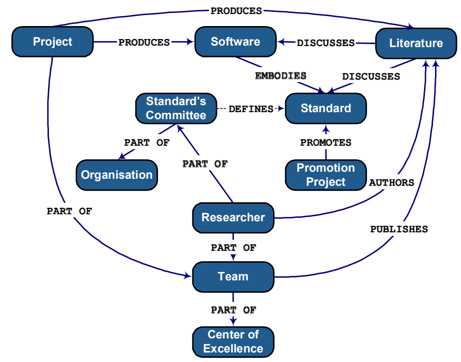

Figure 3: OntoPortal Ontology

SHOE [31, 38] allows researchers to annotate their WWW resources with metadata, in order to build a distributed knowledge base. Ontologies are used to declare the desired characteristics and relationships of a Web resource and SHOE-specific markup is used to annotate the resource and make its properties explicit. The real potential behind SHOE is the ability to draw on the ontology to infer supplementary knowledge not directly stated within the facts describing the web resource. (KA)2 [2] applies a similar approach although places a greater emphasis on the ontological engineering process as well as supporting inference. [51] describes how the (KA)2 initiative has now evolved to provide a coherent set of tools with which to design community Web portals. Finally, topic maps [41] are used to classify concepts (topics) and their relationships in a web site. This semantic layer is then used as a concept browser where users are provided with consistent and accurate access to the information.

COHSE [8, 23] is an open hypermedia system whose links are defined by an independent Ontology Service. ScholOnto [5] uses an ontology of argumentation to model relationships in the literature, focusing on the claims scholars make in their articles. The ontology supports patterns of argumentation between literature, such as refutation, support, extension or modification. These relationships are recorded in a central knowledge base, from which higher-level relationships between the literatures can be inferred, such as *Has anyone built on the ideas in this paper, and in what way?*

ONTOPORTAL

Ontoportal demonstrates how a semantic meta-layer can be projected over existing WWW resources, meaningfully describing the resources themselves, and the relationships between them. The underlying ontological model is discussed, before examining the architecture of the system, knowledge gathering techniques, and some example interactions.

The Ontology

The ontology was constructed from knowledge and experience of the research community as well as recommendations from peers. The type of a resource in a research field includes obvious concepts such as literature and researcher, but also project, software, standard, organisation, promotional project, team and centre of excellence, and these were used as building blocks for the ontology (Figure 3). Originally, the ontology was formalised using a DTD that OntoPortal used to structure the data (Figure 4).

```
<!ELEMENT project (title, description?,
    url?, url_desc?,misc? internal_url?,
    project_relations?, editorial?)>
<!ELEMENT title (#CDATA)>
<!ELEMENT description (#CDATA)>
<!ELEMENT platform (#CDATA)>
<!ELEMENT url (#CDATA)>
<!ELEMENT url_desc (#CDATA)>
<!ELEMENT misc (#CDATA)>
```

Figure 4 - Fragment of the DTD defining a 'project'

Due to the difficulty of precisely constraining elements using a DTD, an XML Schema was developed instead (figure 5). By constraining the elements more rigorously, such as by specifying the number of times an element may occur, the possibility of conflict or error was diminished.

```
<schema
   xmlns="http://www.w3.org/2000/10/XMLSchema">
<element name="project"
   type="t_project_content"/>
<complexType name="t_project_content">
<element name="title" type="string"
   minOccurs="1" maxOccurs="1"/>
<element name="description" type="string"
   minOccurs="0" maxOccurs="1"/>
<element name="url" type="string" minOccurs="0"
   maxOccurs="1"/>
<element name="url_desc" type="string"
   minOccurs="0" maxOccurs="1"/>
<element name="misc" type="string"
   minOccurs="0" maxOccurs="1"/>
<element ref="editorial" minOccurs="0"
   maxOccurs="1"/>
<element ref="relations" minOccurs="0"
   maxOccurs="1"/>
</complexType>
   </schema>
```

Figure 5 - Fragment of the XML Schema

Architecture

Figure 6 illustrates an overview of the architecture for the OntoPortal system. The knowledge, corresponding to facts in the ontology, is stored in an XML-based knowledge base (3). (OntoPortal will use any ontological structure rather than just the scholarly application described here.) Requests from client browsers (1) to view information about a particular resource initiate a query mechanism which retrieves the appropriate metadata from the knowledge base (2), combining where necessary (*e.g.* a request for all literature metadata).

This mechanism also determines the relationships between the requested resource and the other resources in the meta-layer, and adds this information to the XML records as hyperlinks that invoke further queries.

The resulting XML document (4), which conforms to the schema (Figure 5), is either sent directly to the client browser (5) if the browser supports XSLT [54] transformations, or transformed into a presentation format (usually HTML) on the server-side (7). In the former case, the user can use the XSLT stylesheet from the server (6), or provide an alternative stylesheet to meet individual presentation needs.

Knowledge Gathering

The knowledge gathering process is facilitated through simple HTML forms used by editors to enter metadata describing a resource. To integrate the resource with its proper context, valid ontological relationships with other resources already in the meta-layer are presented to the

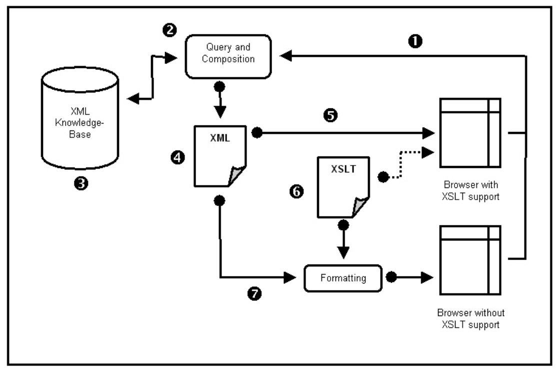

Figure 6 - OntoPortal Architecture

editor, from which the most appropriate are selected.

The metadata research field can be subdivided into several specialist areas - we found it useful in the implementation of OntoPortal to model these areas as "filters" or "themes" through which only appropriate parts of the meta-layer are exposed. This also suited the data population model, with each editor taking responsibility for a "theme" which reflected their research expertise.

Discourse Scholarly discourse is a vital part of the research process [26]. OntoPortal provides support for discourse in the form of capturing the editorial opinion and analysis of the editors as part of the metadata describing a resource as it is added to the meta-layer. Users without editorial privileges are currently restricted to simply making form-based suggestions for possible modifications or additions to the meta-layer, which are delivered to the editors. Suggestions can then be moderated and integrated into the meta-layer.

OntoPortal provides facilities for the research community to engage in threaded discussions seeded on any resource described in the semantic meta-layer. This enables members of the research community to receive peer-level feedback on their current work. Careful attention is paid to the potential abuse of these public access forums, and mechanisms for providing editors with censorship control are in place. The D3E system [52] has demonstrated the success of this approach in promoting good quality feedback and understanding in the community.

Sample Interaction

Figure 7 contrasts the OntoPortal system with a standard Digital Library in a common research situation. The user,

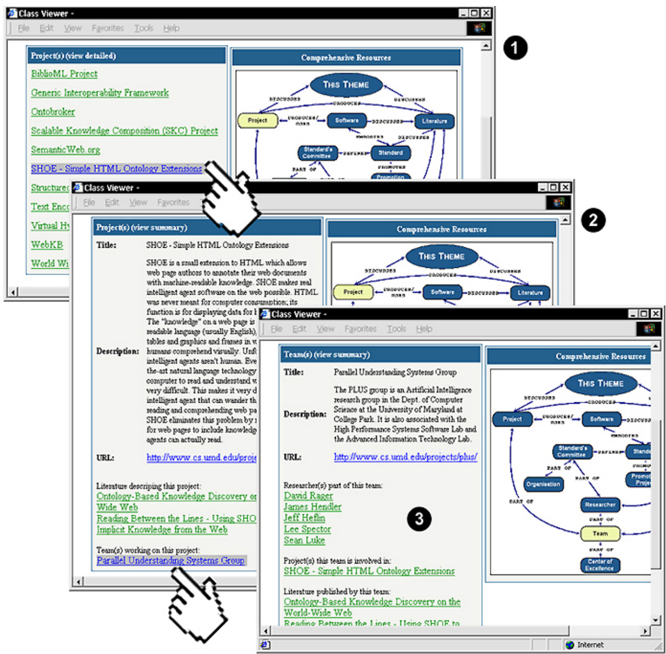

Figure 7: Exploring author/literature relations

having located a research paper relevant to a task, wishes to discover further papers by the same authors in order to increase understanding of the concepts the research paper discusses. In a Digital Library, realizing this intent involves resorting to a search engine and searching for papers written by these authors. However, the search results are limited by the publishing rights of the Digital Library, and will not return links to relevant research papers in other Digital Libraries. Furthermore, the user has no understanding of how the papers returned by the search engine are related to the original paper.

Using the OntoPortal system, in investigating the metadata associated with the original research paper, (1), the user is able to find out directly which other research papers the authors of the paper have been involved in writing. The user moves to the “Author of Literature” link information, and from here is able to take advantage of familiar bibliographic author/literature relationships in order to explore the research papers produced by each author. For example, following the link to metadata about a particular researcher (2), the user is able to view a list of publications by that researcher, visiting metadata about those which seem relevant (3). This transcends the artificial boundaries of Digital Libraries, and by exploring other available links in the meta-layer (as well as “Author of Literature” link information, the metadata in describing the original research paper (1) also meaningfully links the paper with its

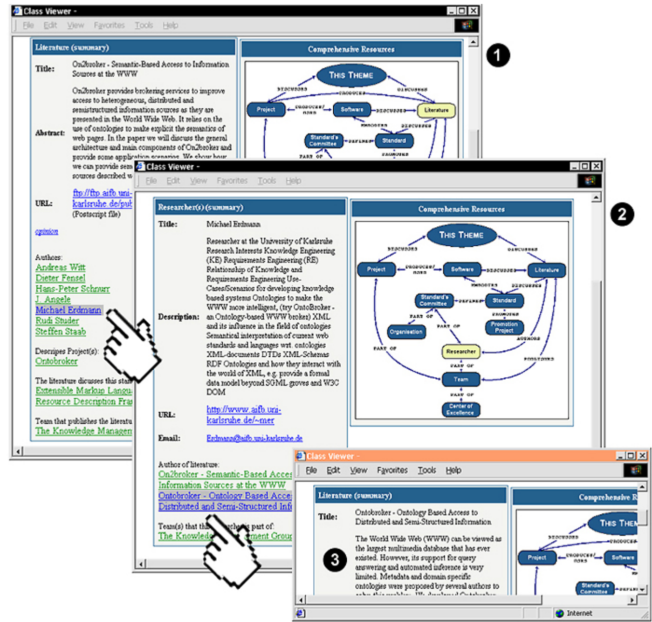

Figure 8: Exploring project/team relationships

associated projects, the standards that it discusses, and the research team that produced it), the user is able to become fully immersed in the research community and context surrounding the research paper.

A pictorial representation of the ontology is displayed as a navigation tool in each of OntoPortal’s views, allowing convenient exploration of other key areas in the research field.

In Figure 8, the user follows through the project metadata for a particular theme, and decides to find out who is involved in a specific project (1). The metadata about the project includes a relationship with the team working on the project (2), and from the metadata describing this team, the user is able to discover the researchers involved (3).

The portal currently contains five themes, each with a dedicated portal editor working on ongoing metadata population and maintenance. Over 400 resources in the metadata research field have been recognised and classified according to the ontological concepts. The metadata describing the resources is interconnected by some 500 complex links.

ESKIMO

The E-Scholar Knowledge Inference MOdel (ESKIMO) system demonstrates how ontologically modelled data can be used to infer new facts and resolve links based on analytical queries, such as “Who are the experts in hypertext?” or “What are the seminal papers in metadata research?” As with the OntoPortal system, users are still provided with direct access to the meta-layer. The system also works in conjunction with a scholar’s browsing environment to ensure the context of ESKIMO matches the

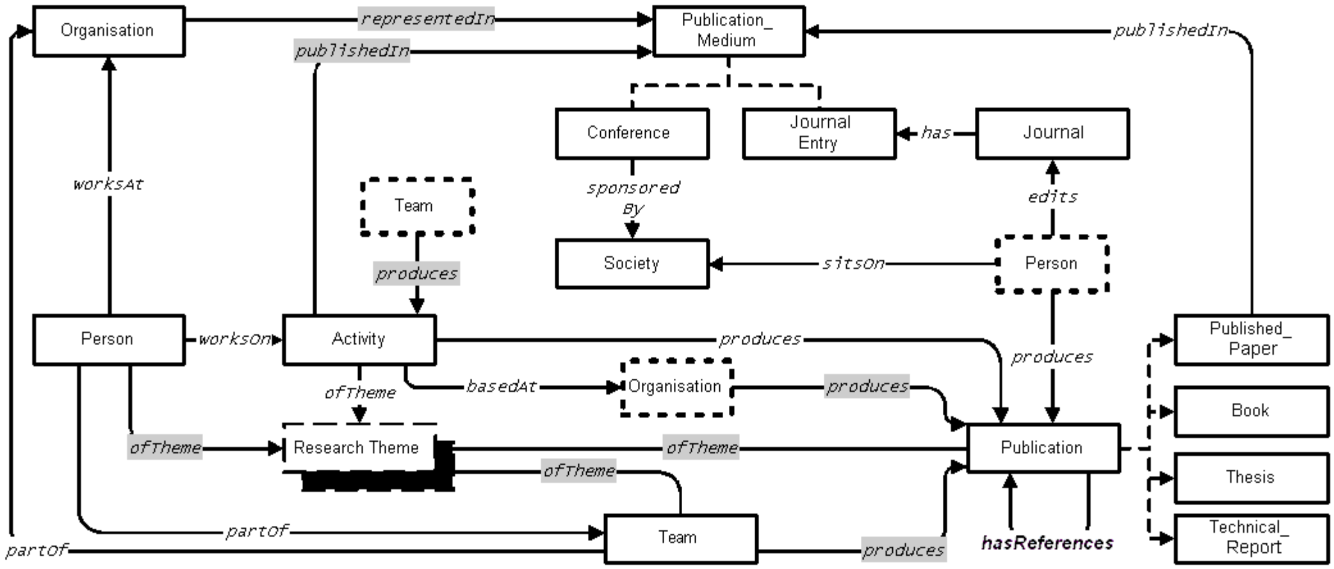

Figure 9 - ESKIMO Ontology

current focus. This ensures that, unlike OntoPortal, the scholar does not leave the realms of the support environment and is always able to refer to it.

ESKIMO's ontology (Figure 9) models the general academic community making it possible to model disparate disciplines in engineering and literary fields. It represents an evolution of the OntoPortal ontology and is more literature oriented with a greater emphasis on producing research (such as detailed bibliographic information for correctly citing research papers). Importantly, citations have been modelled by the ontology to enable the discovery of further work, research fronts, trends and perspectives [11, 19, 20]. The inclusion of the *Research Theme* class (figure 10) not only enables more focused queries, but allows inference based on subject domain.

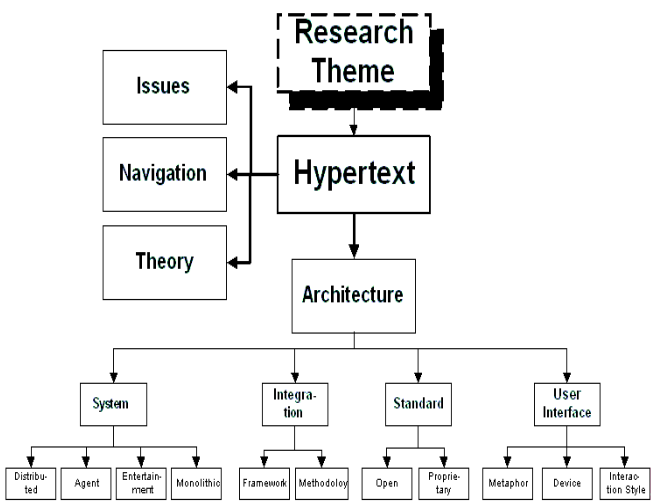

Figure 10 - ESKIMO Ontology

The navigational benefits are inherited from ESKIMO, so queries such as those listed below are possible.

On what projects has Mark Bernstein worked?
Who has cited the paper, 'Linking considered harmful'?
Which projects are based at Xerox PARC?
Which journals does George Landow edit?

Significantly, we can also draw on that ontology to infer implicit information, in addition to the usual reflexive inferences completed on ontologies. The highlighted axioms in the ontology (figure 9) represent those that the system infers, rather than those where the user has supplied the data. For example, to determine which papers a team has published, one could follow:

```
∀x,p,t
PersonProducesPub(x, p) ∧ PersonPartOf(x, t)
⇒ TeamHasPublished(t, p)
```

Such rules can be combined to produce more complex and analytical queries, such as:

Who are the experts in user interface research?
What are the seminal papers in adaptive hypermedia?
Which institutes/teams collaborate?
Which papers provide a broader view on this topic?

A series of heuristic rules are used to solve the analytical queries. For example, to determine the experts in user interface research, the following rules are applied:

```
∀x,y,c,p
ArticleHasTheme(x, 'User Interface') ∧
AuthorOf(x,y) ∧
NumberofTimesCited(x, c) ∧ greaterThan(c, 20) ∧
NumberofPapersPub(y, p) ∧ greaterThan(p, 10) ∧
WorksOnActivity(y, 'User Interface')
⇒ Expert(p, x)
```

The numbers *20* and *10* represent user definable thresholds.

Furthermore, the scholarly data captured in the knowledge base can be used to augment traditional bibliographic analysis tools, such as co-citation and impact factors. For example, drawing citation, researcher and institute facts from the ontology allows us to identify possible collaborations between institutes.

The ACM Hypertext conference proceedings from 1989 to 2000 (represented as XML metadata) were used as the dataset on which to experiment. Naturally, the entire ontology could not be populated using just the proceedings, and therefore concepts, such as activity and society, were filled manually. The Southampton Framework for Agent Research (SoFAR) [50], which includes support for ontology implementation, was utilised to embody the knowledge base.

Sample Interaction

Figure 11 displays an example interaction with the ESKIMO system. The user views a Digital Library-based resource, (1). ESKIMO derives the context of the current focus and displays any known information in a separate browser window, (2). As with OntoPortal, users can navigate the knowledge base to discover further information, (3).

Figure 12 demonstrates the system using the assertions in its ontology to infer new facts. A user is viewing a paper (1) that the ESKIMO system displays information on (2). As the user is unfamiliar with the topic a list of experts is retrieved (3), that can be used a starting point for further reading.

Both interactions demonstrate how, as with the OntoPortal system, e-Scholars are immersed in their community and are quickly able to gain access to a wealth of knowledge and understand how it relates to the corpus of material. Users instruct the system to draw new facts from the ontology and thereby analyse the knowledge in an entirely new context.

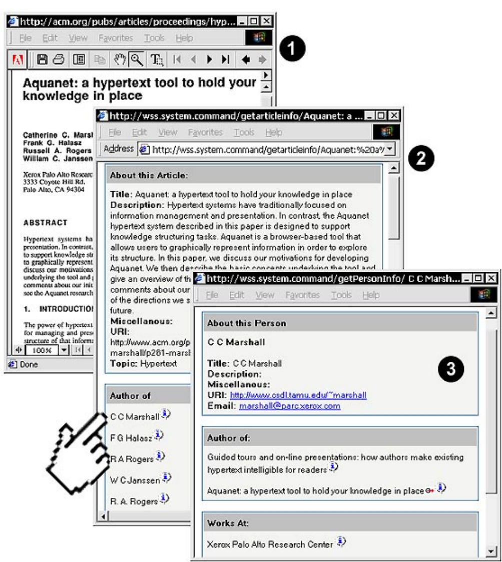

Figure 11 - Sample interaction with ESKIMO - exploring the meta-layer in parallel with browsing

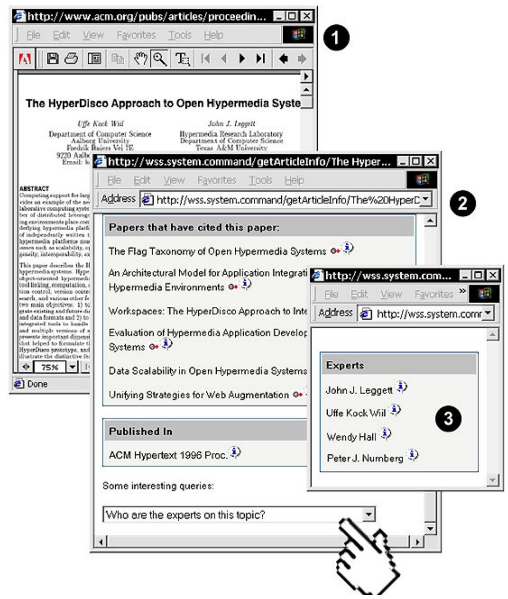

Figure 12: Using ESKIMO's inference engine to uncover new facts about a research field

CONCLUDING REMARKS

The hypertext systems described here base their construction (and the users' navigation) upon an ontological model of the objects that the user encounters in the 'real world'. The effect of this is to provide a principled approach to setting up a coherently linked resource. In particular, by defining an ontology that closely reflects the tasks undertaken in a literature survey, useful support of the online knowledge worker (or 'e-scholar') is obtained. The links derived from the ontology help to transform a weakly linked collection of articles, project home pages and institutional sites into a richly interconnected hypertext resource.

ACKNOWLEDGEMENTS

We wish to thank our editors Sean Bechofer and Nick Gibbins, for imparting their knowledge in helping to populate the OntoPortal knowledge base and providing valuable feedback regarding the editorial process.

The development of OntoPortal was funded in part by the Defence Evaluation and Research Agency (DERA), UK.

BIBLIOGRAPHY

[1] ACM Digital Library, http://www.acm.org/dl/.

[2] Benjamins, R. V., Fensel, D., and Perez, A. G., Knowledge Management through Ontologies, in Proc. PAKM'98 Practical Aspects of Knowledge Management (Basel, Switzerland, 29-30 October, 1998).

[3] Berners-Lee, T., Semantic Web Road Map, http://www.w3.org/DesignIssues/Semantic.html, 1998.

[4] Bernstein, M., Deeply Intertwingled Hypertext: The Navigation Problem Reconsidered, Technical

Communication, 38, 1 (1991), 41-47.

[5] Buckingham Shum, S., Motta, E., and Domingue, J., Representing Scholarly Claims in Internet Digital Libraries: A Knowledge Modeling Approach, in Proc. 3rd European Conference on Research and Advanced Technology for Digital Libraries (Paris, France, September 22-24, 1999), Springer Verlag.

[6] Bush, V., As We May Think. Atlantic Monthly, 1945. http://www.isg.sfu.ca/~duchier/misc/vbush/.

[7] Cameron, R., A Universal Citation Database as a Catalyst for Reform in Scholarly Communication, First Monday 2,4 (April 7th, 1997).

[8] Carr, L., Bechhofer, S., Goble, C., Hall, W., and De Roure, D., Conceptual Linking: Ontology-based Open Hypermedia, in Proc. 10th International WWW Conference, (Hong Kong, May 2001).

[9] Carr, L., Hall, W., and Miles-Board, T., Is the WWW Killing Hypermedia?, in Poster Proc. 9th International WWW Conference, (Amsterdam, May 2000), 42-43.

[10] Carr, L., Hall, W., Miles-Board, T., Writing and Reading Web Hypermedia, Technical Report ECSTR-IAM00-1, Electronics and Computer Science Department, University of Southampton, 2000.

[11] Chen, C., and Carr, L., Trailblazing the Literature of Hypertext: Author Co-Citation Analysis (1989-1998). in Proc. 10th ACM Hypertext Conference (Darmstadt, Germany, February 21-25, 1999), 51-60.

[12] Conklin, J., Hypertext: An Introduction and Survey, IEEE Computer, 2, 9 (1987), 17-41.

[13] Davies, C. and Houghton, D., The Web: Hyperspace, Hypermedia, or Just Hyped?, in Proc. British HCI Group Symposium, The Missing Link: Hypermedia Usability Research and the Web (May 1, 1996).

[14] DeRose, S.J., Expanding the Notion of Links, in Proc. 2nd ACM Hypertext Conference, (Pittsburgh, November 5-8, 1989), 249 - 257.

[15] De Young, L., Linking considered harmful, in Proc. European Hypertext Conference, (INRIA, France, November 1990), 238-249.

[16] Dublin Core Initiative. http://www.purl.org/dc/.

[17] Fensel, D., Horrocks, I., Van Harmleden, F., Decker, S., Erdmann, M., and Klein, M., OIL in a Nutshell, http://www.cs.vu.nl/~ontoknow/oil/downl/oilnutshell.pdf, October 2000.

[18] Furuta, R., Shipman, F. M., Marshall, C., Brenner, D., and Hsieh, H., Hypertext Paths and the World Wide Web: Experiences with Walden's Paths, in Proc 3rd Hypertext Conference (Southampton, UK, April 6-11, 1997), 167-176

[19] Garfield, E., Citation Indexes for Science : A New Dimension in Documentation Through Association of Ideas, Science, 122 (1955), 108-111.

[20] Garfield, E., Citation Indexes for Retrieval and Research Evaluation, Presentation (Unpublished) at the G7 Consensus Conference on the Theory and Practice of Research Assessment, (Capri, Italy, October 7, 1996), 1-30, http://www.garfield.library.upenn.edu/papers/ciretreseval-capri.html.

[21] Garzotto F., Schwabe, D., and Paolini, P., HDM - A Model-Based Approach to Hypertext Application Design, ACM Transactions on Information Systems, 11, 1 (January 1993), 1-26.

[22] Gloor, A., CYBERMAP Yet Another Way of Navigating in Hyperspace, in Proc. 3rd ACM Hypertext Conference, (San Antonio, Texas, 1991).

[23] Goble, C., and Carr, L., COHSE: Informed WWW Link Navigation Using Ontologies, in Proc. IEEE Colloquium "Lost in the Web", (London, UK, November 11th, 1999).

[24] Golovchinsky, G., and Chignell, M., The Newspaper as an Information Exploration Metaphor, Information Processing and Management, 33, 5 (September 1997), 663-683.

[25] Guarino, N., Formal Ontology and Information Systems. in Proc. FOIS'98 (Trento, Italy, 6-8 June, 1998).

[26] Harnad, S., Implementing Peer Review on the Net: Scientific Quality Control in Scholarly Electronic Journals, in Peek, R., and Newby, G., (eds.) Scholarly Publishing: The Electronic Frontier, Cambridge MA: MIT Press., 1996, 103-118.

[27] Harnad, S. The Invisible Hand of Peer Review, Exploit Interactive, 5 (April 2000).

[28] Harnad, S., Electronic Scholarly Publication: Quo Vadis? Serials Review 21, 1 (1995), 70-72, Reprinted in Managing Information 2, 3 (1995).

[29] Harnad, S., Varian, H., and Parks, R., Academic publishing in the online era: What Will Be For-Fee And What Will Be For-Free? Culture Machine 2, 2000, http://www.cogsci.soton.ac.uk/~harnad/Temp/Varian/new1.htm, 2000.

[30] Harter, S., The Impact of Electronic Journals on Scholarly Communication: A Citation Analysis, The

Public-Access Computer Systems Review, 7, 5 (1996).

[31] Heflin, J., Hendler, J. and Luke, S., Reading Between the Lines: Using SHOE to Discover Implicit Knowledge from the Web, in Proc. AAAI'98 Workshop on AI and Information Integration, (Madison, Wisconsin, July 26-30, 1998).

[32] Hitchcock, S., Web Publishing: Speed Changes Everything, IEEE Computer 29, 8 (August 1996), 91-93.

[33] Hitchcock, S., Carr, L., Harris, S., Hey, J., and Hall, W., Citation Linking: Improving Access to Online Journals. in Proc 2nd ACM Digital Libraries Conference (Philadelphia, 1997), 115-122.

[34] Hitchcock, S., Carr, L., Quek, F., Witbrock, A., Tarr, I., and Hall, W., Linking Everything to Everything: Journal Publishing Myth or Reality? in Proc. ICCC/IFIP Conference on Electronic Publishing - New Models and Opportunities, (Canterbury, UK, 14-16 April 1997).

[35] Ichimura, S. and Matsushita, Y., Another Dimension to Hypermedia Access, in Proc. 5th ACM Hypertext Conference, (Seattle, November 14-18, 1993), 63-72.

[36] Kreutz, R., Euler, B. and Spitzer, K., No longer lost in WWW-based hyperspaces, in Proc. 10th ACM Hypertext Conference, (Darmstadt, Germany, February 21-25, 1999), 133-134.

[37] Langston, L., Scholarly Communication and Electronic Publication: Implications for Research, Advancement, and Promotion, in Proc. Untangling the Web: Proceedings of the Conference Sponsored by the Librarians Association of the University of California, Santa Barbara and Friends of the UCSB Library, (Santa Barbara, 1996).

[38] Luke, S., Spector, L., and Rager, D., Ontology-Based Knowledge Discovery on the World Wide Web, in Proc. AAAI'96 Workshop on Internet-based Information Systems (Portland, Oregon, August 4 - 8, 1996).

[39] Mendes, E., and Hall, W., Towards the Prediction of Development Effort for Web Applications, in Proc 11th ACM Hypertext Conference (San Antonio, Texas, 2000), 242-243.

[40] Mukherjea, S., and Hara, Y., Focus+Context Views of World Wide Web Nodes. in Proc. 8th ACM Hypertext Conference (Southampton, UK, April 6-11,1997), 187-196.

[41] Pepper, S., Euler, Topic Maps, and Revolution, in Proc. XML Europe (Granada, April 1999).

[42] Peters, J., The Hundred Years War Started Today: An exploration of electronic peer review, http://www.press.umich.edu/jep/works/PeterHundr.html.

[43] Pikrakis, A., Bitsikas, T., Sfakianakis, S., Hatzopoulos, M., De Roure, D., Hall, W., Reich, S., Hill, G., and Stairmand, M., MEMOIR - Software Agents for Finding Similar Users by Trails, in Proc. PAAM'98 Practical Application of Intelligent Agents and Multi-Agents, (London March 1998), 453-466.

[44] Postmodern Culture, http://jefferson.village.virginia.edu/.

[45] Psycoloquy, http://www.cogsci.soton.ac.uk/psycoloquy/.

[46] Public-Access Computer Systems Review, http://info.lib.uh.edu/pacsrev.html.

[47] Rayward, W.B., H.G. Wells's Idea of a World Brain: a Critical Re Assessment, JASIS 50, (May 15 1999), 557-573.

[48] RDF, W3C Recommendation, http://www.w3.org/RDF/, 22 February 1999.

[49] Snyder, J., and Flemming, U., The Object Modeling Language (OML) Specification, http://seed.edrc.cmu.edu/ACL/ObjModelAPI/objapi.ps

[50] SoFAR (Southampton Framework for Agent Research), http://www.sofar.ecs.soton.ac.uk..

[51] Staab, S., Angele, J., Decker, S., Erdmann, M., Hotho, A., Maedche, A., Schnurr, H. P., Studer, R., and Sure, Y., Semantic Community Web Portals, in Proc. 9th International WWW Conference (Amsterdam, May 15-19, 2000), 474-491.

[52] Sumner, T., and Buckingham Shum, S., From Documents to Discourse: Shifting Conceptions of Scholarly Publishing, in Proc. Human Factors in Computing Systems, (Los Angeles April 18-23, 1998).

[53] XML 1.0 (Second Edition), W3C Recommendation, http://www.w3.org/TR/REC-xml, 6 October 2000.

[54] XSLT 1.0, W3C Recommendation, http://www.w3.org/TR/xslt, 16 November 1999.

[55] Zellweger, P., Structure, Navigation and Hypertext: The Status of the Navigation Problem, in Proc. 3rd ACM Hypertext Conference (San Antonio, Texas, Dec, 1991).