Knowledge Acquisition and Representation Techniques in Scholarly Communication

by Brian R Gaines and Mildred L G Shaw
Knowledge Science Institute
University of Calgary
Calgary, Alberta, Canada T2N 1N4
{gaines, mildred}@cpsc.ucalgary.ca

Abstract: Paper journals, conferences and workshops have been the major media for scholarly discourse for 300 years. However, in the 1990s access to low-cost personal computing and Internet communications is leading to radical changes in the operation of scholarly communities. Electronic publication and conferencing is becoming common in all disciplines using commonly available Internet facilities such as ftp archives, list servers, gopher and world-wide web. Some scholarly communities that had not previously achieved a critical mass have done so through the net, others have launched major collaborative projects managed through the net, and others are questioning the value of conventional conferences that are limited by being localized in space and time compared with the flexibility of continuous international electronic conferencing through the web. However, the majority of current electronic scholarly discourse emulates paper-based media in relying primarily on text and diagrams for knowledge communication. It is beginning to take advantage of some of the multimedia capabilities of electronic publishing for color diagrams, pictures, movies and sound. Hypertext and hypermedia capabilities are being used to develop webs of linked material. Concept maps and formal knowledge structures are being used to provide a framework for knowledge expression, interchange and collaborative development. This article focuses on the extension of current documentation technologies to provide knowledge-level support for scholarly communities.

Introduction

The Knowledge Science Institute (KSI) was set up in 1985 with a dual mandate: to study and model the knowledge economy [3] with a view to forecasting trends and modeling knowledge processes; and to innovate in selected areas that might lead to insights into the impact of knowledge support systems [14]. Figure 1 shows the major research programs addressing these objectives:

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association of Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

DOC 94-10/94 Banff, Alberta, Canada
© 1994 ACM 0-89791-681-6/94/0010...$3.50

- *Information systems infrastructure* modeled as a tiered structure of learning curves for electronic devices, computer architecture, software engineering, interactive systems, knowledge representation and inference, knowledge acquisition, autonomous systems, and socially organized systems [16].

- *Knowledge dynamics* modeled in social, psychological and logical terms as a process of distinction creation, dissemination, application and evaluation [13, 27].

- *Collective stance model* of the formation of expertise though positive feedback processes in the human species viewed as a single organism recursively partitioned in space and time into sub-organisms that are similar to the whole; these parts include societies, organizations, groups, individuals, roles, and neurological functions.

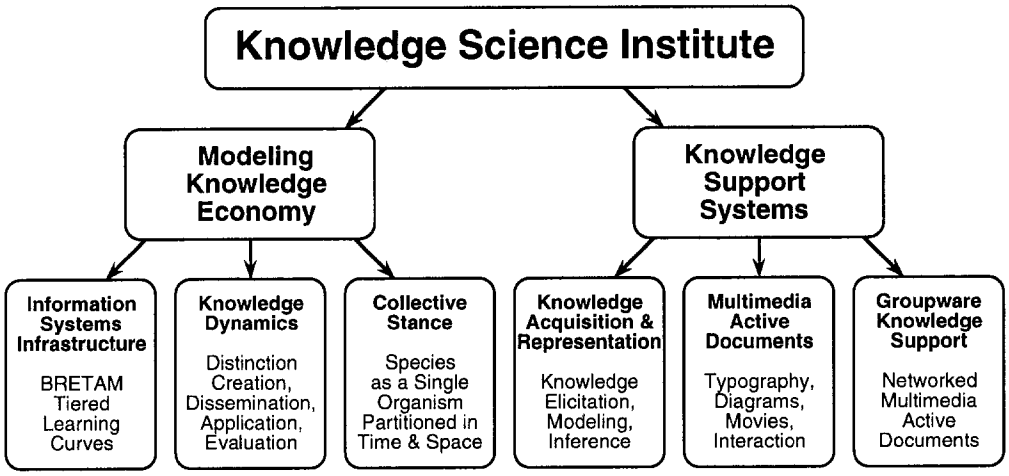

Figure 1 Knowledge Science Institute research program

- *Knowledge acquisition and representation* methodologies and tools for eliciting, modeling and reasoning with human knowledge [27, 28].
- *Multimedia active documents* supporting the dissemination of knowledge through systems integrating text, diagrams, pictures, movies, sound, concept maps, semantic nets, animation, simulation and access to general information systems [29].
- *Groupware knowledge support systems* targeted on the systematic acceleration of scientific research through networked access to knowledge-based systems and multimedia active documents [23, 30, 53].

One conclusion of the KSI studies is that a major anachronism in human knowledge processes has, until recently, been the continued reliance on paper media for knowledge dissemination [18]. While book production technology has gone through several major technological advances, the end-product is basically the same as it was for Gutenberg and Caxton over five hundred years ago. Similarly, scholarly journals have changed little in their presentation and operation since the inception of the *Philosophical Transactions* of the Royal Society in March 1665. When one considers the limitations of the paper medium in communicating knowledge, and contrasts with the capabilities of current multimedia computer and communication technology, the incentives to make major changes in scholarly publishing are very high. Books and journals primarily support typographic text and monochrome diagrams and photographs, and require separate indexing for purposes of contextual search. Computer-readable multimedia documents support the same quality of typographic text, color diagrams and photographs, sound and video, animation, simulation, computation, and are intrinsically contextually searchable. They can be produced and distributed on CD-ROM at a substantially lower costs than paper books or journals, and can be read on personal computers and selectively printed on laser printers. They can also be produced and distributed on the Internet, and this provides an active medium supporting many forms of scholarly discourse.

In developing systems for the support of human knowledge processes one has to ensure that the technology is effective relative to the underlying social dynamics. Sociologists of science have characterized scientific communities as forming *invisible colleges* which monitor and manage the changing structure of knowledge in their domain [11]. Individuals playing major roles in these communities are experts not only in overt knowledge of the domain but also through skills in its management and development. Sociologists and philosophers of science have undertaken empirical studies to elicit and model such expertise [4, 9, 41], and, in recent years, have begun to use artificial intelligence concepts and methodologies in such studies [10, 58]. In the 1990s major scientific communities are beginning to use communication through the Internet as a major channel for scientific discourse [36], and it is this existing usage which provides a foundation on which to build more advanced tools and systems incorporating artificial intelligence techniques.

The objective of the KSI research program reported in this article and its sister articles is to provide a new generation of open architecture knowledge support systems [21] supporting the collaboration of scientific communities through with a view to achieving systematic acceleration of scientific research. Specifically, a new layer of knowledge level tools has been developed that complements existing list servers and electronic document archives with research coordination facilities. The tools are designed to support the knowledge processes of scientific communities through active document, hypermedia, concept mapping, knowledge representation and active agent technologies. They use existing Internet protocols for mail, file transfer and hypermedia to integrate with related systems. They provide import and export facilities for common document formats to integrate with existing document production. The data formats are designed for parallel use on CD-ROM to provide alternative means of distribution to the Internet. The pragmatic objective is to achieve widespread use of greatly improved facilities for scientific discourse and collaborative project management. A deeper objective is to study the impact of the proposed technologies on actual research practice with a view to designing a next generation of tools targeted on the changes in patterns of research behavior already apparent as new information technologies have come to play a significant role in scholarship.

The other articles report on developments in visual language tools for knowledge representation [34], and multimedia active document tools for knowledge dissemination [54]. This article is primarily concerned with supporting scholarly communities through the application of techniques for knowledge acquisition and representation developed in artificial intelligence research. The next section gives an overview of relevant techniques and tools, followed by some illustrative applications.

Knowledge Acquisition and Representation

Knowledge engineering is a branch of information system engineering that places particular emphasis on the epistemological status of information that is classified as *knowledge*, that is as "justified, true belief" [47] and "more than opinion, less than truth" [48]. These definitions provide a context in which the *credibility* and *derivation* of information are significant, and need to be taken into account as overt data in their own right. That is, the knowledge-based components of an information system will tend to be those in which *meta-information* and *meta-information-processing* is required [20]. Such systems will have components that do not only process information but also process meta-information *about* that information and its processing. A knowledge-based system may be characterized succinctly as an essentially *reflective* information system [39].

It is these characteristics of knowledge engineering that make it of interest in the support of scientific communities. Conventional information technology already plays a major role in the support of such communities through data collection, storage and processing, and is playing an increasing role in knowledge dissemination through digital media. However, in current applications the computer only supports the flow of information, not the reflective modeling of the nature of that information—knowledge *about* the information, its sources, credibility, completeness, operational interpretation, and so on. What is now feasible is to add a layer of reflective information processing to existing computer-based systems supporting scholarly communities, that will provide meta-information of various kinds.

Figure 2 is taken from a survey of knowledge acquisition methodologies and tools [26], and shows how current systems come together to provide an integrated methodology for knowledge-based systems development:

- Stage consists of developing an informal knowledge base from interviews, protocols and media. This stage involves managing a heterogeneous data base of knowledge sources represented in a variety of media, and is often supported through hypermedia storage, indexing and retrieval tools.
- Stage ① consists of using the informal knowledge to elicit the major sub-domains that encompass the significant phenomena in the overall domain. This stage involves giving structure to the material gathered in the previous stage, and is often supported through concept mapping tools [30].

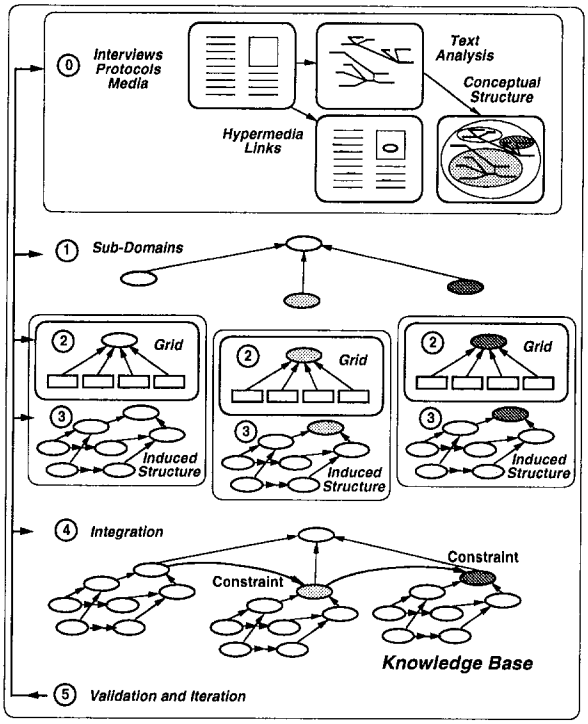

Figure 2 Knowledge acquisition and representation methodologies

- Stage ② consists of using conceptual modeling methodologies in each sub-domain to elicit relevant attributes and critical cases. This stage involves defining the terminology and distinctions that provide a conceptual framework for the problem domain, and is often supported through repertory grid elicitation and analysis [27].
- Stage ③ consists of deriving the relations constraining distinctions, the *rules* of the domain. This stage involves modeling the dynamics of the domain, and is often supported through conceptual induction [17].
- Stage √ consists of linking the sub-domains together, generally by specifying as a constraint in one domain that some role in that domain has as value an individual in another. This stage is the most difficult to automate, and is often supported through formal visual languages allowing people to refine and test the knowledge base through graphic interaction [15].
- Stage ⑤ consists of testing the overall knowledge base and iterating back through any of the earlier stages in order to develop and refine it further. This process is essentially cyclic since knowledge changes with experience—the technology requirement is to support iterative development of the knowledge base at each of the previous stages.

One of the major lessons learnt in the past decade of knowledge-based systems development has been that the products of all the stages shown in Figure 2 constitute the *knowledge base*, not just the final formal representation. The derivational linkages between the knowledge base and its sources are vital to the provision of reflective *explanation* systems [2, 57], and to the maintenance and improvement of the knowledge base [22, 24].

Concept Map Representing Mission Statement

The initial stage of knowledge acquisition and representation shown in Figure 2 makes use of hypermedia tools to manage a large body of heterogeneous material. What is interesting in the context of this article is the tools used to derive knowledge structures from that material. Concept mapping is a knowledge acquisition technique used in the initial stages of system development to structure the domain and task. It has been widely used in educational studies to elicit the changing

conceptual structures of students as they interact with the educational system [35, 44]. Concept maps are sorted directed graphs in which nodes have types and labels. Many concept mapping systems resemble early semantic networks [46] in having imprecise semantics, and there is a continuum between concept maps and the precise visual languages used to define knowledge bases. For example, Toulmin's [60] analysis of scientific arguments can be given precise semantics[8], and *Coreview* [61] prescribes a well-defined set of ontological primitives.

Figure 3 shows a concept map representing the mission statement of a major research project, IMS Test Case 7 'GNOSIS', part of the international Intelligent Manufacturing Systems research program. The project [31] involves over 100 participants in 31 industry and university organizations in 14 countries. The mission is stated to be:

> The long-term goal of this proposal is to develop an approach to a post mass-production paradigm. This is a paradigm which recognizes the imbalances caused by the present global manufacturing environment—the growing scarcity of natural resources, the problem of environmental destruction, and the issue of regional trade imbalances. It also recognizes the growing need for flexibility in production and product configuration leading towards integrated networks of producers and suppliers.
>
> To support this long-term goal, this test case aims to develop systematization of knowledge for design and manufacturing and the concept of the virtual factory.
>
> The systematization of functional knowledge is aiming to build soft machines—artifacts considered beneficial to society and to the environment and designed to cover all stages in the product life from initial concept to final disposal.
>
> The concept of the virtual factory to generate dynamic product-specific product-lines for configurable products with functional knowledge will be developed and demonstrated.

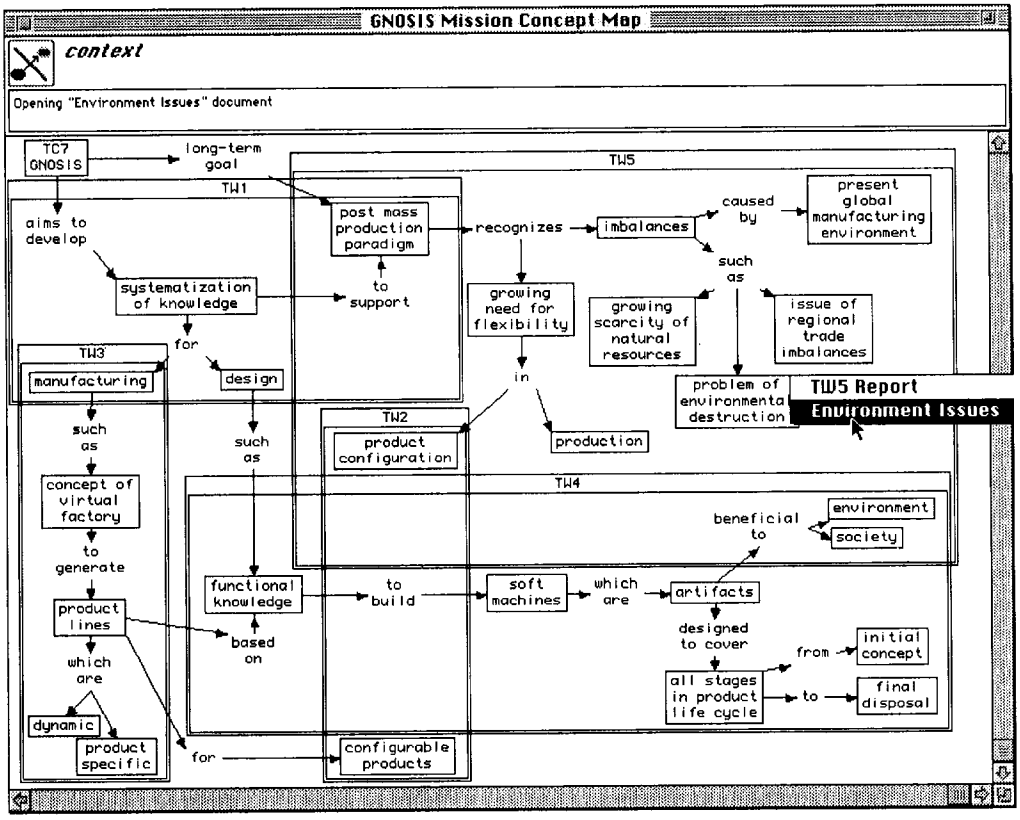

Figure 3 Concept map of GNOSIS mission statement

Five technical work packages were established to carry out the mission , and the concept map in Figure 3 has been partitioned to show their roles and interactions. The upper right part of the map is concerned with the *post mass production paradigm* studies (TW5). The pivotal role of *knowledge systematization* studies (TW1) shows as linking TW5 to *configurable production systems* (TW3) and to *soft machines* studies (TW4). Studies of *configuration management systems* (TW2) are visible only in the nodes *product configuration* and *configurable products*. In practice, TW2 took off more slowly than the other work packages, and the analysis of the project documents tends to indicate that the objectives for this work package were not as clearly formulated or integrated into the overall project as for the other work packages. Examination of Figure 3 also suggests that there is a gap that needs filling between the very high level socio-economic goals of the IMS program and the very specific technical objectives of the work packages other than TW5.

The concept map and the text form of the mission statement are alternative ways of representing the same knowledge. The map being two-dimensional gives more scope to show links between concepts. Its diagrammatic form also provides a natural user interface supporting access to related material. For example, the concept mapping tool allows arbitrary actions to be provided through context-sensitive popup menus associated with each node. The popup menu on the right shows the user opening a document on environment issues through the concept map.

Automatic Text Analysis of Research Domain

The concept map in Figure 3 was developed manually from the mission statement. Automatic text analysis tools can be used to generate conceptual maps directly from documents, extracting the significant vocabulary and showing the associations between terms. The objective is not to 'understand' the text but to use it to avoid wasting an expert's time through the elicitation of the basic vocabulary and conceptual structure. The methodology used is the analysis of the co-occurrence of words in sentences, a technique commonly used in information retrieval systems [7].

Figure 4 shows a concept map derived automatically by analysis of an article, *The Technical Concept of IMS* [59], that was significant in the GNOSIS project proposal. The document is treated as a set of entities which are sentences whose features are the words they contain. Rules are derived using empirical induction in which the premise is that if one word occurs in a sentence then the conclusion is that another will occur. The graph shows the links from premises to conclusions derived in this way. The tool is interactive and, as shown at the top center, provides access through a popup menu associated with each word to a list of occurrences of that word in context, and the original document.

The initial output was a digraph consisting of one major connected component and some minor ones which correspond to significant topics such as intellectual property rights that did not directly relate to the main issues. The user noted that the major component itself consisted of 3 loosely connected sub-components, and added the context boxes shown to distinguish and name these parts. The significance of these parts is that they correspond to 3 of the 5 technical work packages of the research program. What is particularly significant is the missing work packages, TW2 concerned with product configuration management systems, and TW3 concerned with configurable production systems. These were added to the GNOSIS research program during its formation through amalgamation of interests with other potential proponents of IMS test cases. These packages link technically to the knowledge systematization activities on the left of Figure 4 but are neutral to the major issues of the IMS program on the right.

Thus, automated text analysis can provide similar information to that developed manually. The synopsis of what the document is about, with links back to the original, supports project members in assimilating and coordinating complex goals and activities. However, it does not in itself result in *artificial intelligence* in the sense of an inferential computational system. The knowledge acquired is being fed back to facilitate the intelligence of people rather than being used by a computer to provide direct support of knowledge processes through computational intelligence. One objective of the studies reported here has been to investigate the potential for knowledge-based systems to

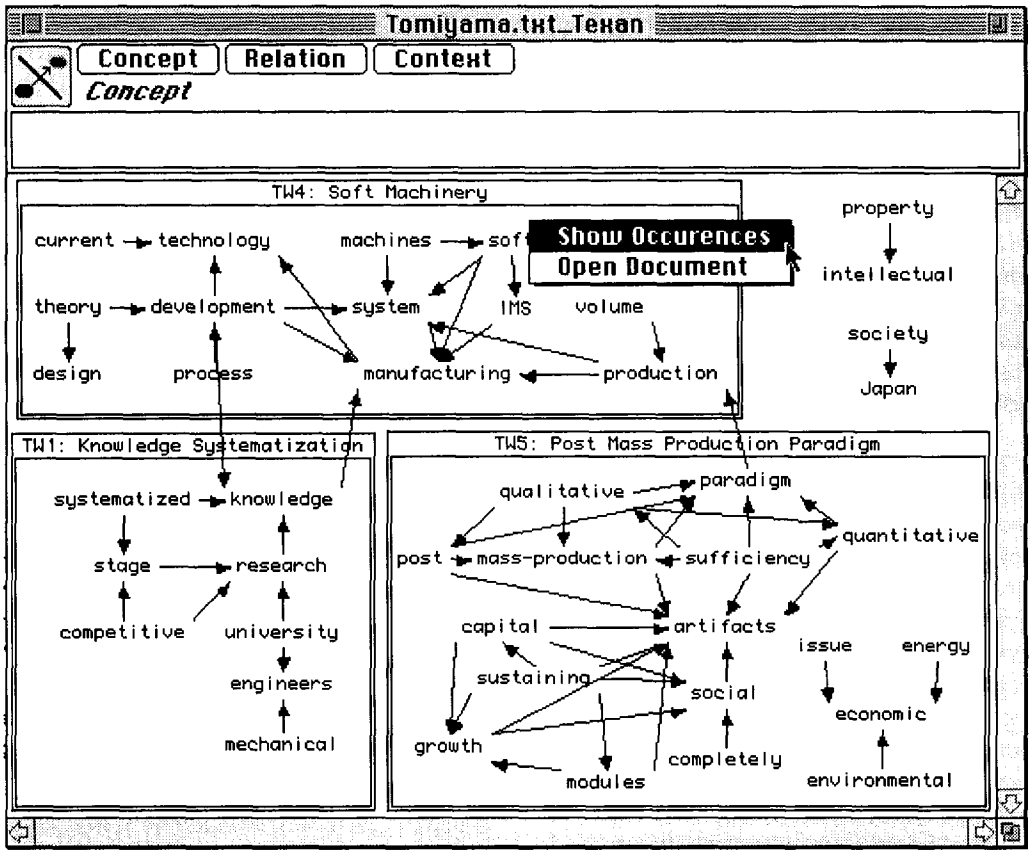

Figure 4 Analysis of an article describing the objectives of the IMS program

play a role in clarifying: project goals and activities, the relations between them, and, in particular, the inconsistencies between goals and activities that seem to arise in all major projects. The next section exemplifies how this may be done.

Objectives as Formal Knowledge Structures

Formal knowledge structures were developed for the top level goals of the IMS program as described in the document, *The Technical Concept of IMS*[59], by detailed analysis of the sub-section on *ideal artifacts*, extracts from which are shown below.

> They should be congenial to mankind.
>
> Safety: This is surely incontrovertible.
>
> Ease of use: This means more than user friendliness or outstanding man-machine communications capabilities. It means they should be prepackaged with the necessary innate intelligence so that they are the least possible trouble to use.
>
> Trouble free: Inseparable from ease of use.
>
> They should be congenial to the environment.
>
> Possible to recycle: By this I mean that manufacturing industry should actively seek to develop its own functional equivalents of veins and kidneys. Only thus will it be possible to eliminate waste and achieve complete reclamation.
>
> Not dirty: We don't even need to look as far as the automobile industry and the problem of exhaust gases to see the need for minimizing the load imposed on the environment.
>
> Not a consumer of energy: This again is self-evident.
>
> They should be congenial to Society.
>
> Compatibility that is welcome to society: This compatibility with society is an extremely important factor for the future. The advent of new types of man-made product often has revolutionary effects on daily life and society. However, we will need to be extremely careful introducing any man-made

> objects that promise to change society. By this I mean that manufacturing industry should be most concerned with the ripple effect of the introduction of new man-made objects on economic problems, labor problems, environmental problems, etc.
>
> New market creation, not generation of trade friction: We need to consider not only the influence on society but also, and indeed primarily, that upon exports. We also need to consider how we are going to retrench when our manufacturing responsibilities develop to the point when this becomes necessary.

The text describes the requirements for an *ideal artifact* that would provide a suitable target for the IMS research program, and then discusses in detail the factors underlying these requirements. The knowledge acquisition methodology appropriate to the elicitation of such a conceptual structure is another technique derived from personal construct psychology, that of *laddering* [27, 32]. Laddering tools take a concept such as *congenial to mankind* and ask two types of question: laddering up, *why* should an artifact be congenial to mankind?; laddering down, *how* can an artifact be *congenial to mankind*? In interactive elicitation the expert is taken up and down the conceptual structure by sequences of such how and why questions. For example, laddering up from "why should an artifact be *congenial to mankind*?" to "to be an *ideal artifact*" would lead to laddering down through the question "are there other ways of being an *ideal artifact*" and the elicitation of further requirements such as *congenial to the environment*. Laddering up from *ideal artifact* would lead to the concept in this case *to be a suitable target for IMS research*.

The lowest level concepts developed through laddering down are expected to be *operational* in that it is possible to determine whether they apply to an entity through reference to assertions about it. In personal construct psychology terms this means that they will each be one pole of construct that can be completed by definition of its opposite, for example, what would one term a system that is not *safe*. Hence, the final stage of a laddering process is to elicit the opposite poles of the lowest level concepts.

A formal knowledge structure was developed for the research objectives using a laddering tool with answers derived from the text above. The tool represents the structure in a visual language for terminological knowledge representation systems[15]. The language is a formal one with semantics that provide a direct translation into the constructs of CLASSIC [6] and similar knowledge representation systems. Concepts are represented by ellipses, primitive concepts by ellipses with short horizontal markers, inheritance by an arrow from a concept to one subsuming it, disjointness by a line between disjoint primitives, individuals by a rectangle, role by unboxed text, rules by a rectangle with a double line at each end, and so on. Contexts are represented by a container rectangle.

At the top of Figure 5 is the top level decision logic expressed in the boxed text above, that an *IMS ideal artifact* should be *congenial to mankind*, *congenial to the environment* and *congenial to society*. This definition is typical of that found in the analysis of codified legal material—apparently simple requirements but expressed in terms of high-level concepts that require other knowledge structures for their interpretation [49]. The text provides the interpretations, for example, that an artifact *congenial to mankind* is *safe*, *easy to use* and *trouble-free*.

The third level knowledge structure at the bottom of Figure 5 is elicited at the final stage of laddering when the operational concepts at the lowest level are made into dimensions by the elicitation of opposites. The importance of this can be seen by considering that some of the second-level interpretations such as *safe* appear capable of operational definitions in terms of factual attributes of entities. However, the positive attributes required are not themselves operational but are characterized in terms of their opposites—rules that characterize a person as *dishonest* or an artifact as *unsafe* are more readily operationalized than those for *honest* or *safe*. These definitions are not given explicitly in the text because they are part of the 'commonsense knowledge' that the reader brings to its interpretation, knowledge that one would expect to find overtly expressed in a commonsense knowledge base such as CYC [38].

The knowledge structure of Figure 5 is a semantic network that can be loaded directly into the into the KRS inference engine[19] and used to reason about artifacts. The pattern of inference is reasoning to a contradiction. If an artifact is asserted to be an *ideal*

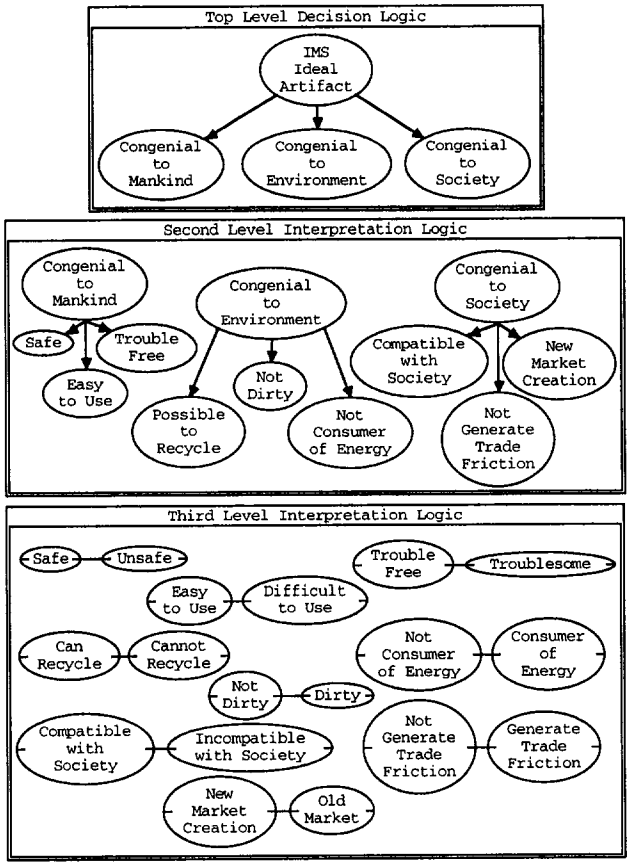

Figure 5 Formal knowledge structures for concept of IMS

artifact it inherits all the properties of being *safe*, *easy to use* and so on. If an inference rule concerned with its being *unsafe* or being *difficult to use* fires, the disjoint primitive relations cause the knowledge base to be incoherent, the original assertion is automatically retracted and the knowledge base can be interrogated to determine the reason for the error.

This is the inference pattern of CLASSIC used to detect constraint violations in the management of large-scale software projects [12]. The pattern also characterizes scientific research in which conjectures are made that may be refuted [45]. The conjectures are never verified and remain forever fallible, and new knowledge as it is added may lead to the refutation of past conjectures.

The dimensions at the lowest level in Figure 5 also define a repertory grid relating artifacts and research objectives. The top-level concept, *ideal artifact,* can be entered in the grid as an abstract object, an 'ideal element' in grid terminology, that is rated as being on every 'preferred pole.' The second level concepts. such as *congenial to mankind*, can also be entered as abstract objects rated as being on the relevant preferred poles, and as being 'any' on the irrelevant ones. More concrete objects such as actual artifacts, or designs, can be entered and rated on the dimensions developed from the text. The clustering tools then show the relations between artifacts, and between them and the idealized abstract artifacts.

Repertory Grid Analysis of Soft Machines

One problem for GNOSIS has been the presentation of the project objectives and activities to funding and reviews agencies concerned with the international and national programs. There was a common theme of *reconfigurable systems* in the three major technical work packages, but it became clear that this was inadequately projected in the project documents. In preparation for a major review meeting in June 1993 an analysis was made of the *soft machine* concept using a KA technique

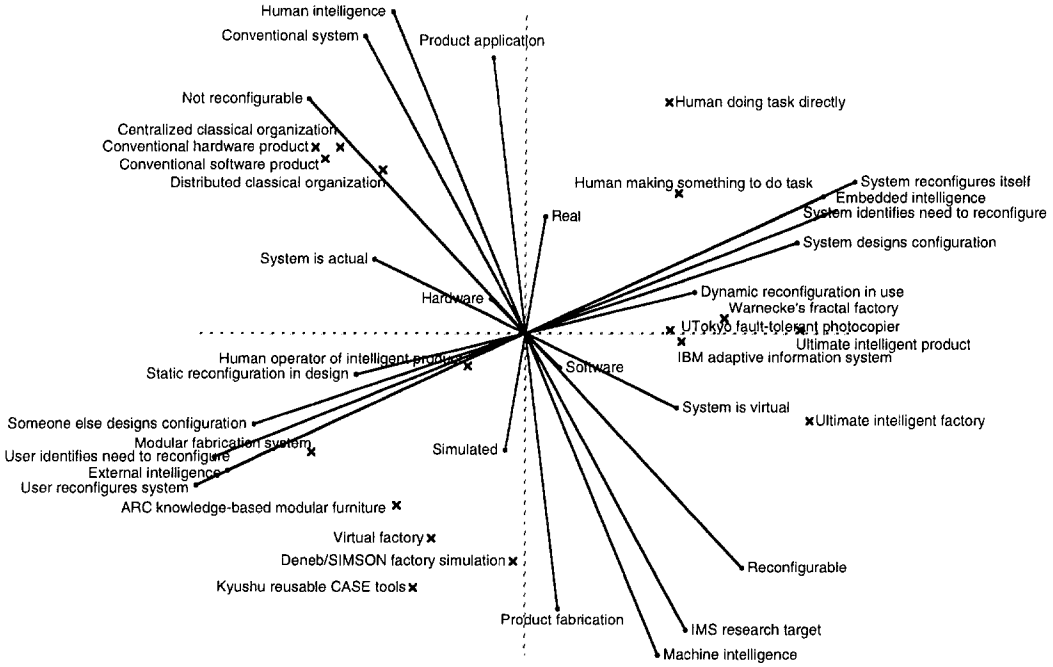

Figure 6 Principal components analysis of soft machine repertory grid

derived from personal construct psychology [33], that of *repertory grid elicitation* [5, 25, 27].

The repertory grid is used in knowledge elicitation when an expert finds it easier to provide exemplary cases rather than develop a knowledge structure directly. The technique elicits the significant distinctions between cases all the time feeding back matches between cases to elicit new distinctions, and matches between distinctions to elicit new cases. The resultant grid is clustered to feed back to the expert the overall conceptual structure for validation, and rules may be induced from the comparatively small dataset which are usually meaningful because the feedback has eliminated spurious correlations. The grid and rules can be exported as a knowledge base covering the specific domain characterized by the cases [28].

Six major GNOSIS sub-projects were used as initial elements, and the ensuing repertory grid elicitation process resulted in the addition of another 10 elements, including human operators and organizational structures that provided contrasts to some aspects of the technological projects. Eleven distinctions were elicited that provided detailed insights into the complexity of the notion of reconfigurability, and these were presented on viewfoils to the review body to explain the roles of the GNOSIS projects and the relations between them relevant to issues of soft machinery.

Figure 6 shows the grid clustered through a principal components analysis to bring together similar distinctions and similar elements, and to show the relationships between them. It shows, for example, that some past distinctions in manufacturing are no longer as critical as they used to be—the *hardware—software* distinction did not characterize other distinctions—GNOSIS treats software manufacturing and hardware manufacturing alike. The main dimensions apparent are those typified by *system reconfigures itself—user reconfigures system* and that of *human intelligence—machine intelligence*. The IMS projects are on the machine intelligence side of the second dimension and cluster into two groups on the first, those concerned with self-reconfiguring systems and those concerned with user-reconfiguring systems.

Repertory grids have been used extensively in education to model students' changing conceptual models [50], and to support creative processes [52]. They can also be used in group situations to elicit and compare the conceptual structures of different people, and to analyze these structures in terms of differences in terminology and its usage. For example, RepGrid-Net [51, 53] is a com-

puter-based message system that integrates conventional electronic mail with repertory grid elicitation and analysis. The coordinators of such a network provide a basic focus for it through statements of intent, topics and issues which are handled on a bulletin board basis. They also provide one or more kernel grids listing specific topics and the concepts which they apply to them. These kernel grids can be developed by others interested in the groups, using the stated topics and concepts, and adding to them. General similarities between grids are analyzed to provide a socionet of people with common viewpoints, and this may be used to access the mail system to communicate with them. Detailed comparisons of similarities and differences between viewpoints may be made, and individual concept structures can be analyzed.

Related Research

The 1992 AAAI Workshop on *Communicating Scientific and Technical Knowledge* illustrated the potential for the use of artificial intelligence research in supporting the scientific community [56]. Computer-based concept mapping tools have been used to provide hypertext support for scientific knowledge processes [8, 55].

Research on shared ontologies and knowledge interchange formats is relevant to the development of knowledge support systems as envisioned in this article [43]. It is not reasonable to expect individual projects to develop the complete ontological framework to support the lower level inference rules necessary to completion of the knowledge structures in Figure 5. It would be practical for a particular project to make its top-level logic overt, and then develop interpretations down to a level where they could be made operational by interfacing to shared ontologies available in interchange formats.

Significant areas of related research are in communities focusing on issues of computer-supported cooperative work and collaboratories [37], and on the empirical study of the social and knowledge processes underlying scientific activity [1, 40, 42]. Most of the current tools for scientific communication on the Internet such as email, list servers, ftp archives, gopher, archie, WAIS and world wide web, have been built on a pragmatic basis to solve problems of textual communication and information retrieval. These communication systems are neutral to the content of the data communicated and can be extended to carry knowledge structures facilitating structured communication at the knowledge level. The development of effective knowledge support systems requires the joint efforts of the largely disjoint communities concerned with these diverse research and development areas.

Conclusions

Paper journals, conferences and workshops have been the major media for scholarly discourse for 300 years. However, in the 1990s access to low-cost personal computing and Internet communications is leading to radical changes in the operation of scholarly communities. Electronic publication and conferencing is becoming common in all disciplines using commonly available Internet facilities such as ftp archives, list servers, gopher and world-wide web.

Some scholarly communities that had not previously achieved a critical mass have done so through the net, others have launched major collaborative projects managed through the net, and others are questioning the value of conventional conferences that are limited by being localized in space and time compared with the flexibility of continuous international electronic conferencing through the web.

However, the majority of current electronic scholarly discourse emulates paper-based media in relying primarily on text and diagrams for knowledge communication. It is beginning to take advantage of some of the multimedia capabilities of electronic publishing for color diagrams, pictures, movies and sound. Hypertext and hypermedia capabilities are being used to develop webs of linked material. Concept maps and formal knowledge structures are being used to provide a framework for knowledge expression, interchange and collaborative development. This article has focused on the extension of current documentation technologies to provide knowledge-level support for scholarly communities.

Acknowledgments

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada. We are grateful to colleagues in GNOSIS, particularly Tetsuo Tomiyama, for their cooperation and support.

References

[1] P. Abelson, "Mechanisms for evaluating scientific information and the role of peer review," *Journal American Society Information Science*, vol. 41, no. 3 pp. 216-222, 1990.

[2] S. Abu-Hakima and F. Oppacher, "Improving explanations in knowledge-based systems: RATIONALE," *Knowledge Acquisition*, vol. 2, no. 4 pp. 301-343, 1990.

[3] D. Bell, *The Coming of Post-Industrial Society*, New York: Basic Books. 1973.

[4] S.S. Blume, *Toward a Political Sociology of Science*, New York: Free Press. 1974.

[5] J.H. Boose, "Personal construct theory and the transfer of human expertise," in *Proceedings AAAI-84*. American Association for Artificial Intelligence: California. p. 27-33, 1984.

[6] A. Borgida, R.J. Brachman, D.L. McGuiness, and L.A. Resnick, "CLASSIC: a structural data model for objects," in *Proceedings of 1989 SIGMOD Conference on the Management of Data*. ACM Press: New York. p. 58-67, 1989.

[7] M. Callon, J. Law, and A. Rip, ed. *Mapping the Dynamics of Science and Technology*. MacMillan: Basingstoke, UK. 1986.

[8] V. Cavalli-Sforza, G. Gabrys, A.M. Lesgold, and A.W. Weiner, "Engaging students in scientific activity and scientific controversy," in *AAAI-92 Workshop on Communicating Scientific and Technical Knowledge*, K. Swaminathan, Editor. AAAI: Menlo Park, California. p. 99-114, 1992.

[9] S. Cole, *Making Science: Between Nature and Society*, Cambridge, Massachusetts: Harvard University Press. 1992.

[10] H.M. Collins, *Artificial Experts: Social Knowledge and Intelligent Machines*, Cambridge, Massachusetts: MIT Press. 1990.

[11] D. Crane, *Invisible Colleges: Diffusion of Knowledge in Scientific Communities*, Chicago: University of Chicago Press. 1972.

[12] P. Devanbu, P.G. Selfridge, B.W. Ballard, and R.J. Brachman, "A knowledge-based software information system," in *IJCAI'89: Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*. Morgan Kaufmann: San Mateo. p. 110-115, 1989.

[13] B.R. Gaines, "Social and cognitive processes in knowledge acquisition," *Knowledge Acquisition*, vol. 1, no. 1 pp. 251-280, 1989.

[14] B.R. Gaines, "Knowledge support systems," *Knowledge-Based Systems*, vol. 3, no. 3 pp. 192-203, 1990.

[15] B.R. Gaines, "An interactive visual language for term subsumption visual languages," in *IJCAI'91: Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*. Morgan Kaufmann: San Mateo, California. p. 817-823, 1991.

[16] B.R. Gaines, "Modeling and forecasting the information sciences," *Information Sciences*, vol. 57-58, pp. 3-22, 1991.

[17] B.R. Gaines, "The tradeoff between knowledge and data in data acquisition," in *Knowledge Discovery in Databases*, G. Piatetsky-Shapiro and W. Frawley, Editor. MIT Press: Cambridge, Massachusetts. p. 491-505, 1991.

[18] B.R. Gaines, "An agenda for digital journals: the socio-technical infrastructure of knowledge dissemination," *Journal of Organizational Computing*, vol. 3, no. 2 pp. 135-193, 1993.

[19] B.R. Gaines, "A class library implementation of a principled open architecture knowledge representation server with plug-in data types," in *IJCAI'93: Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*. Morgan Kaufmann: San Mateo, California. p. 504-509, 1993.

[20] B.R. Gaines, "Modeling and extending expertise," in *EKAW'93: Knowledge Acquisition for Knowledge-Based Systems*, N. Aussenac, G. Boy, B. Gaines, M. Linster, J.G. Ganascia, andY. Kodratoff, Editor. Springer: Berlin. p. 1-22, 1993.

[21] B.R. Gaines, "Class library implementation of an open architecture knowledge support system," *International Journal Human-Computer Studies*, vol. 40, pp. to appear, 1994.

[22] B.R. Gaines, M. Linster, and M.L.G. Shaw, "An integrated knowledge support system," in *Proceedings of FGCS'92: International Conference on Fifth Generation Computer Systems*. ICOT: Tokyo. p. 1157-1164, 1992.

[23] B.R. Gaines and D.H. Norrie, "Mediator: information and knowledge management for the virtual factory," in *SIGMAN AAAI-94 Workshop: Reasoning about the Shop Floor*. AAAI: Menlo Park, California. 1994.

[24] B.R. Gaines, A. Rappaport, and M.L.G. Shaw,

"Combining paradigms in knowledge engineering," *Data and Knowledge Engineering*, vol. 9, pp. 1-18, 1992.

[25] B.R. Gaines and M.L.G. Shaw, "New directions in the analysis and interactive elicitation of personal construct systems," *International Journal Man-Machine Studies*, vol. 13, pp. 81-116, 1980.

[26] B.R. Gaines and M.L.G. Shaw, "Integrated knowledge acquisition architectures," *Journal for Intelligent Information Systems*, vol. 1, no. 1 pp. 9-34, 1992.

[27] B.R. Gaines and M.L.G. Shaw, "Basing knowledge acquisition tools in personal construct psychology," *Knowledge Engineering Review*, vol. 8, no. 1 pp. 49-85, 1993.

[28] B.R. Gaines and M.L.G. Shaw, "Eliciting knowledge and transferring it effectively to a knowledge-based systems," *IEEE Transactions on Knowledge and Data Engineering*, vol. 5, no. 1 pp. 4-14, 1993.

[29] B.R. Gaines and M.L.G. Shaw, "Open architecture multimedia documents," in *Proceedings of ACM Multimedia 93*. p. 137-146, 1993.

[30] B.R. Gaines and M.L.G. Shaw, "Concept maps indexing multimedia knowledge bases," in *AAAI-94 Workshop: Indexing and Reuse in Multimedia Systems*. AAAI: Menlo Park, California. 1994.

[31] GNOSIS, ed. *Knowledge Systematization: Configuration Systems for Design and Manufacturing: Final Report of the Test Case*. Knowledge Science Institute, University of Calgary: Canada. 1994.

[32] D.N. Hinkle, The change of personal constructs from the viewpoint of a theory of implications. Ohio State University. 1965.

[33] G.A. Kelly, *The Psychology of Personal Constructs*, New York: Norton. 1955.

[34] R. Kremer and B.R. Gaines, "Groupware concept mapping techniques," in *Proceedings SIGDOC'94*. ACM: New York. 1994.

[35] J.G. Lambiotte, D.F. Dansereau, D.R. Cross, and S.B. Reynolds, "Multirelational semantic maps," *Educational Psychology Review*, vol. 1, no. 4 pp. 331-367, 1989.

[36] G.P. Landow and P. Delany, ed. *The Digital Word: Text-based Computing in the Humanities*. MIT Press: Cambridge, Massachusetts. 1993.

[37] J. Lederberg and K. Uncapher, "Towards a National Collaboratory," in *Report of an Invitational Workshop at The Rockefeller University*. 1989.

[38] D.B. Lenat and R.V. Guha, *Building Large Knowledge-Based Systems*, Reading, Massachusetts: Addison-Wesley. 1990.

[39] P. Maes and D. Nardi, ed. *Meta-Level Architectures and Reflection*. North-Holland: Amsterdam. 1988.

[40] K.W. McCain, "Mapping authors in intellectual space: a technical overview," *Journal American Society Information Science*, vol. 41, no. 6 pp. 433-443, 1990.

[41] R.K. Merton, *The Sociology of Science: Theoretical and Empirical Investigations*, Chicago: University of Chicago Press. 1973.

[42] I.I. Mitroff, *The Subjective Side of Science*, New York: Elsevier. 1974.

[43] R. Neches, R. Fikes, T. Finin, T. Gruber, R. Patil, T. Senator, and W.R. Swartout, "Enabling technology for knowledge sharing," *AI Magazine*, vol. 12, no. 3 pp. 36-56, 1991.

[44] J.D. Novak and D.B. Gowin, *Learning How To Learn*, New York: Cambridge University Press. 1984.

[45] K.R. Popper, *Conjectures and Refutations: The Growth of Scientific Knowledge*, London: Routledge & Kegan Paul. 1963.

[46] M.R. Quillian, "Semantic memory," in *Semantic Information Processing*, M. Minsky, Editor. MIT Press: Cambridge, Massachusetts. p. 216-270, 1968.

[47] A. Quinton, "Knowledge and belief," in *The Encyclopedia of Philosophy: Volume 4*, P. Edwards, Editor. Macmillan: New York. p. 345-352, 1967.

[48] D.D. Runes, *Dictionary of Philosophy*, London: Peter Owen. 1972.

[49] M. Sergot, "Representing legislation as logic programs," in *Machine Intelligence 11*, J.E. Hayes, D. Michie, andJ. Richards, Editor. Clarendon Press: Oxford. p. 209-260, 1988.

[50] M.L.G. Shaw and B.R. Gaines, "Tracking the creativity cycle with a microcomputer," *International Journal of Man-Machine Studies*, vol. 17, no. 1 pp. 75-85, 1982.

[51] M.L.G. Shaw and B.R. Gaines, "Supporting personal networking through computer networking," in *Proceedings of CHI'91: Human Factors in Computing Systems*. ACM Publications: New York. p. 437-438, 1991.

[52] M.L.G. Shaw and B.R. Gaines, "Mapping creativity with knowledge support tools," in *AAAI-91 Workshop on Creativity: Models, Methods and Tools*. AAAI: Menlo Park, California. p. 32-45, 1992.

[53] M.L.G. Shaw and B.R. Gaines, "Group knowl-

edge elicitation over networks," in *Research and Development in Expert Systems X. Proceedings of British Computer Society Expert Systems Conference,* M.A. Bramer and A.L. Macintosh, Editor. Cambridge University Press: Cambridge, UK. p. 43-62, 1993.

[54] M.L.G. Shaw and B.R. Gaines, "Active documents combining multimedia and expert systems," in *Proceedings SIGDOC'94.* ACM: New York. 1994.

[55] P. Smolensky, "Information processing in dynamic systems: foundations of harmony theory," in *Parallel Distributed Processing Volume 1,* D.E. Rumelhart, J.L. McClelland, andT.P.R. Group, Editor. Cambridge, Massachusetts. p. 195-281, 1986.

[56] K. Swaminathan, ed. *AAAI-92 Workshop on Communicating Scientific and Technical Knowledge.* AAAI: Menlo Park, California. 1992.

[57] W.R. Swartout, "Knowledge needed for expert system explanation," *Future Computing Systems,* vol. 1, no. 2 pp. 99-114, 1986.

[58] P. Thadgard, *Conceptual Revolutions,* Princeton, New Jersey: Princeton University Press. 1992.

[59] T. Tomiyama, The technical concept of IMS. RACE Discussion Paper, No. RA-DP2, Research into Artifacts, Center for Engineering, The University of Tokyo. 1992.

[60] S. Toulmin, *The Uses of Argument,* Cambridge, UK: Cambridge University Press. 1958.

[61] D. Wan and P. Johnson, "Supporting scientific learning and research review using COREVIEW," in *AAAI-92 Workshop on Communicating Scientific and Technical Knowledge,* K. Swaminathan, Editor. AAAI: Menlo Park, California. p. 107-114, 1992.*

You can now find SIGDOC in the 'web: http://web.mit.edu/tps/www/NL/SIGDOC_WWW/

Check these pages for conference information.