Using Knowledge Acquisition and Representation Tools to Support Scientific Communities

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

Widespread access to the Internet has led to the formation of geographically dispersed scientific communities collaborating through the network. The tools supporting such collaboration currently are based primarily on electronic mail through mailing list servers, and access to archives of research reports through ftp, gopher and world wide web. However, electronic communication can support the knowledge processes of scientific communities more directly through overtly represented knowledge structures. This paper describes some experiments in the use of knowledge acquisition (KA) and representation (KR) tools to define and analyze major policy and technical issues in an international research community responsible for one of the test cases in the Intelligent Manufacturing Systems (IMS) research program. It is concluded that distributed knowledge support systems in routine use by world-class scientific communities collaborating through the Internet will provide a major impetus to artificial intelligence research.

Introduction

Sociologists of science have characterized scientific communities as forming *invisible colleges* which monitor and manage the changing structure of knowledge in their domain [Crane 1972]. 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 [Merton 1973; Blume 1974; Cole 1992], and, in recent years, have begun to use artificial intelligence concepts and methodologies in such studies [Collins 1990; Thadgard 1992]. Studies have been described in which knowledge acquisition and representation methodologies and tools have been used to elicit and model conceptual structures in scientific communities [Gaines and Shaw 1989]. In the 1990s major scientific communities are beginning to use communication through the Internet as a major channel for scientific discourse [Landow and Delany 1993], and it has become feasible to make KA and KR methodologies and tools routinely available. The conjecture is that the discourse will be improved by more overt representation of the underlying knowledge structures resulting in the systematic acceleration of the scientific research. Improved communication and information management tools should also make it easier to manage the large-scale international collaborative research projects now being supported through the Internet.

This paper reports some studies of the application of knowledge acquisition and representation tools that are in routine use in expert system applications to the analysis of the knowledge processes of a scientific community. The project studied is IMS TC7 'GNOSIS', one of 6 one-year test cases under the international Intelligent Manufacturing Systems research program which started in the second quarter of 1993. The project [GNOSIS 1994] involves over 100 participants in 31 industry and university organizations in 14 countries, with the objective of developing a *post mass production manufacturing paradigm* involving *reconfigurable artifacts*. The project has made extensive use of electronic mail and electronic document archives to coordinate its activities, and the studies reported are part of an investigation to improve such coordination in the main 10-year study commencing in 1994.

The studies have used a wide range of KA and KR tools ranging from hypermedia to manage heterogeneous knowledge sources, through text-analysis of conceptual associations, concept mapping for knowledge visualization, repertory grid and induction tools to develop structures from the knowledge sources, to semantic network, knowledge representation and inference tools to support formal representation and reasoning using the knowledge structures. The objective of the studies has been to determine whether KA and KR tools can play a useful role in supporting the intellectual management of the research program. In particular, one focus has been to reconstruct the knowledge processes that resulted in the mission statement developed for the GNOSIS funding application, since this has determined the primary research activities, groupings and basis for evaluation. It has also been a major topic for critical analysis as the community prepares its long-term research program for the next decade.

We thank colleagues in GNOSIS, particularly Martti Mäntylä, Doug Norrie, Tetsuo Tomiyama and Moriki Toyama.

Text Analysis of IMS Research Objectives

Hypermedia tools have proved effective in managing the diverse forms of knowledge collected in KA, and in linking this to derived knowledge structures both for system development and for user explanation [Gaines and Shaw 1992] The hypermedia components of the KA tools were used to manage some of the knowledge flows in the early GNOSIS meetings when short intensive workshops were held with 20-30 participants where research capabilities were described, sub-goals developed and tasks allocated. For example, some 27 papers exchanged and presentations given at a workshop in Tokyo in mid-March 1993 were digitized through OCR after the meeting and made available by the end of March as a 300-page conference volume in a uniform style [GNOSIS 1993]. A CD-ROM version also contains embedded QuickTime videos of a laboratory presentation and software demonstrations.

The availability of this material in digital form within the KA system made it possible to use text analysis tools to analyze the conceptual structures of the papers. Figure 1 shows a concept map generated automatically through analysis of the co-occurrence of words in sentences, a technique commonly used in information retrieval systems [Callon, Law and Rip 1986]. The document analyzed is one on *The Technical Concept of IMS* [Tomiyama 1992] that played a major role in the design of the research program. 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 to 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 socio-technical 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 (TW's) 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 1 but are neutral to the major issues of the IMS program on the right.

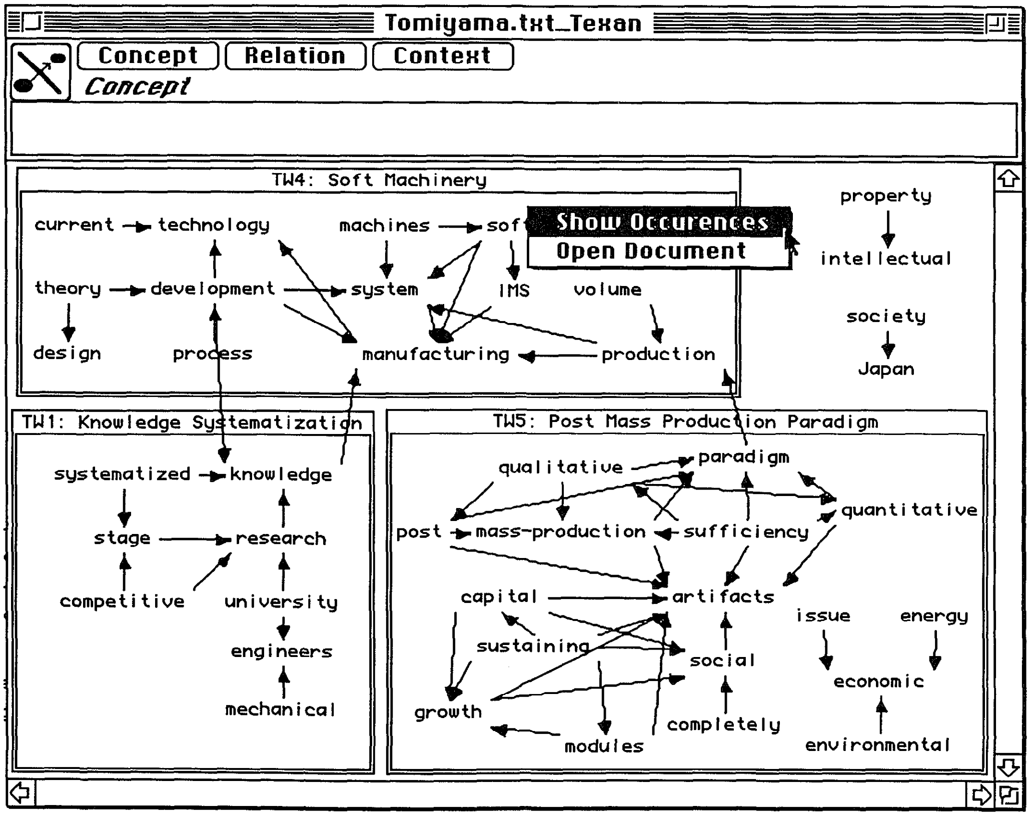

Figure 1 Analysis of paper describing objectives of IMS program

Concept Map of Mission Statement

Analyses like that of Figure 1 was used to develop concept maps clarifying relations between work packages. Concept mapping is another KA 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 [Novak and Gowin 1984; Lambiotte, Dansereau, Cross and Reynolds 1989].

Concept maps are sorted directed graphs in which nodes have types and labels. Many concept mapping systems resemble early semantic networks [Quillian 1968] 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 [1958] analysis of scientific arguments can be given precise semantics [Cavalli-Sforza, Gabrys, Lesgold and Weiner 1992], and *Coreview* [Wan and Johnson 1992] prescribes a well-defined set of ontological primitives.

Figure 2 shows the *mission statement* of the GNOSIS project as a concept map. This concise statement of the project objectives was taken from the introduction to the legal agreement signed by all participants, and much effort went into its formulation. The upper right part of the map is concerned with the *post mass production paradigm* studies (TW5) that show up on the lower right of Figure 1.

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 Figures 1 and 2 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 popup menu on the right shows the user opening a document on environment issues through the access provided by the concept map environment.

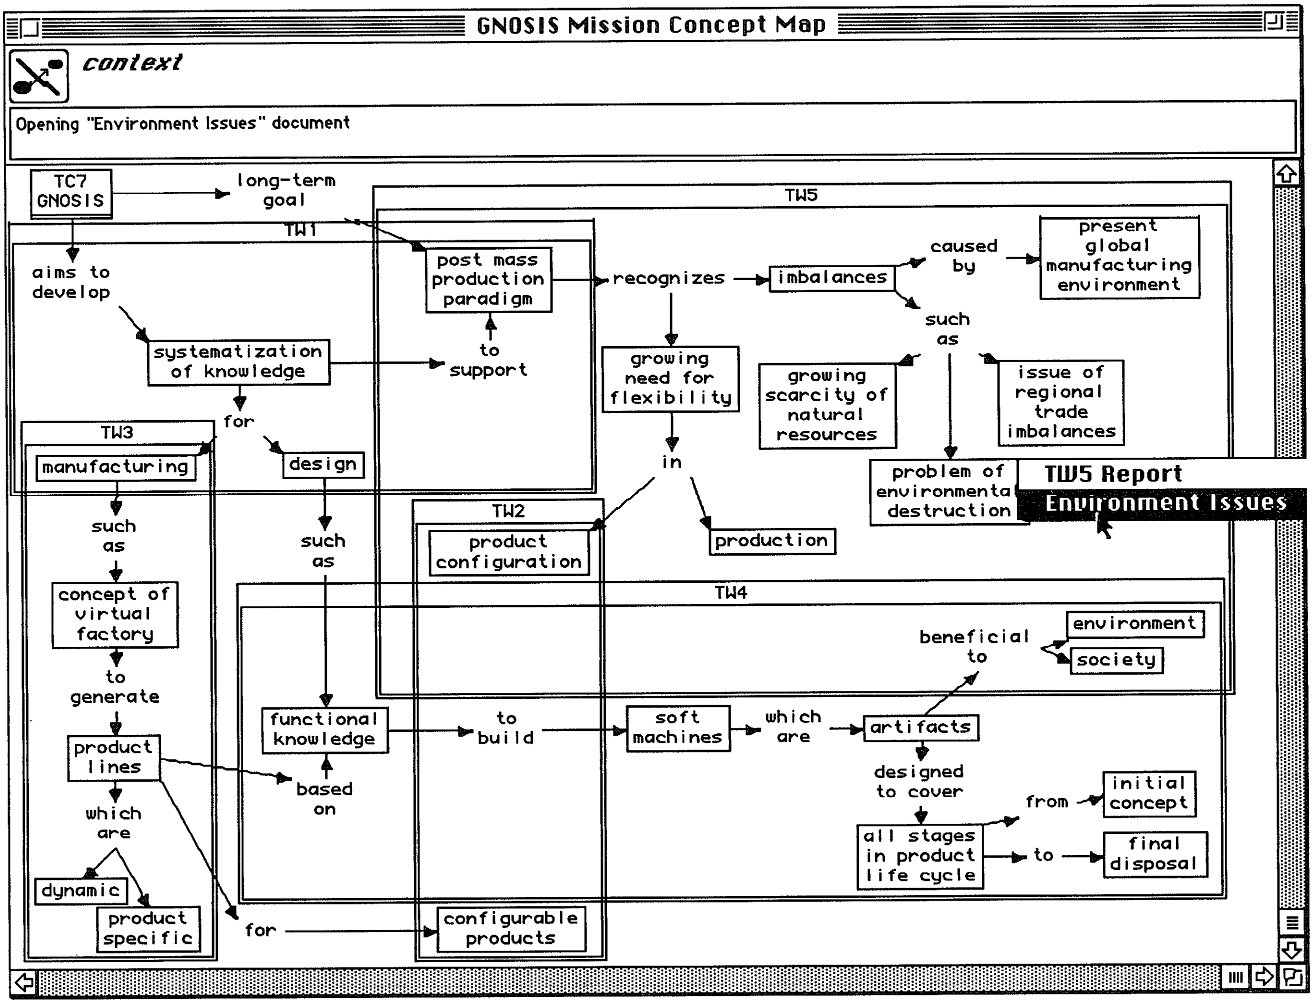

Opening "Environment Issues" document

Figure 2 Concept map of GNOSIS mission statement

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 derived from personal construct psychology [Kelly 1955], that of *repertory grid elicitation* [Gaines and Shaw 1980; Boose 1984; Gaines and Shaw 1993a].

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 [Gaines and Shaw 1993b].

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 3 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.

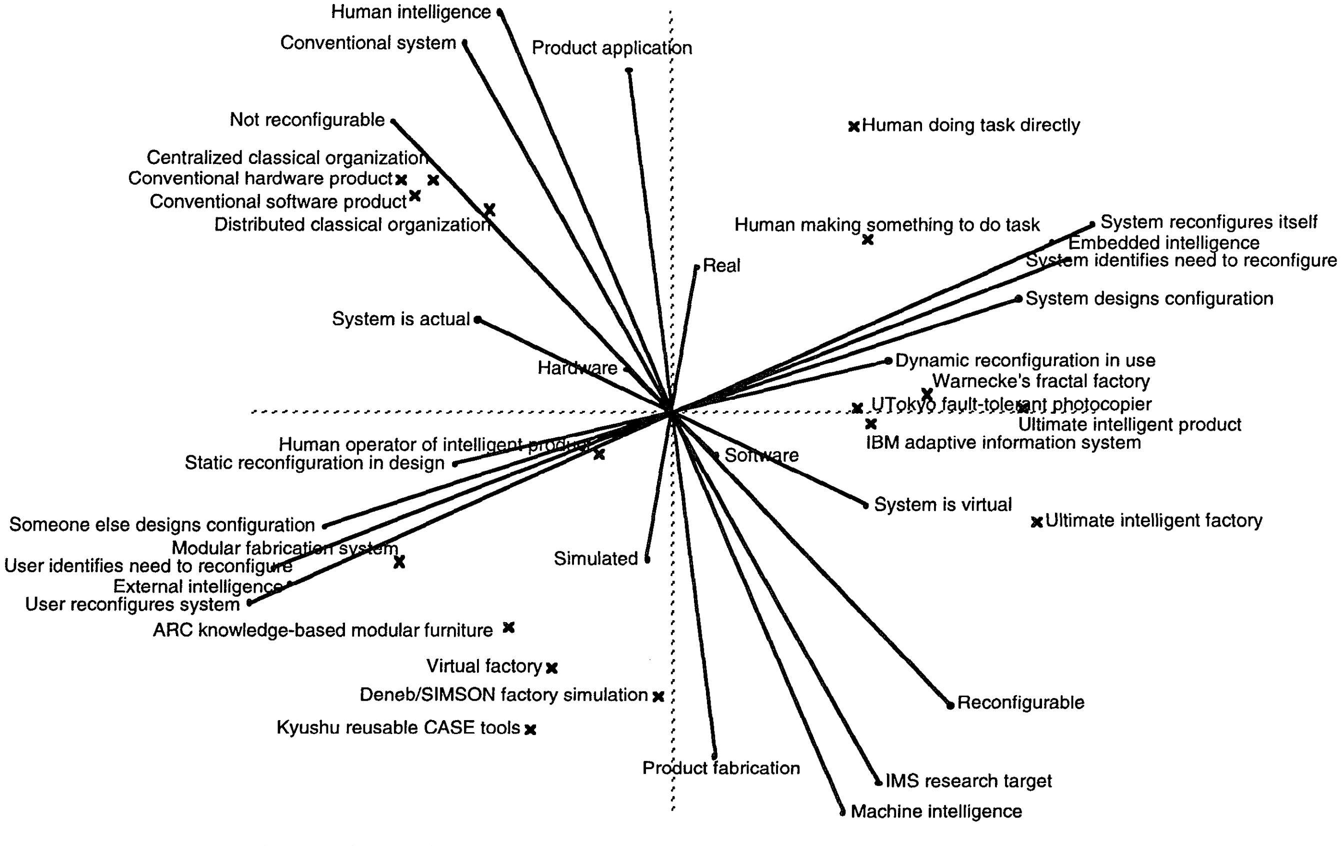

Figure 3 Principal components analysis of soft machine repertory grid

Laddering the Research Objectives

The application of KA techniques described in the previous section, whilst relevant to the understanding, management and presentation of the project 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 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.

To this end, formal knowledge structures were developed for the top level goals of the IMS program as described in the document, *The Technical Concept of IMS* [Tomiyama 1992], 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.

Tomiyama 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 KA methodology appropriate to the elicitation of such a conceptual structure is another technique derived from personal construct psychology, that of *laddering* [Hinkle 1965; Gaines and Shaw 1993a]. 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*.

In moving from the psychological structure elicited by laddering to a formal knowledge structure, some additional meta-characteristics of the structure have to be elicited:

- are the answers to 'how' questions *necessary* characteristics of the concept above, in which case they are represented as a definition but otherwise become the premise of a rule
- if necessary, are the answers to 'how' questions also *sufficient* in which case the concept is fully defined, but otherwise is itself a primitive concept
- if not necessary, are the other answers to 'how' questions *alternatives* in which case they correspond to premises of different rules having a common conclusion.

The lowest level concepts developed through laddering down are expected to be *operational* in that it is possible to determine whether or not 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.

Objectives as Formal Knowledge Structures

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 [Gaines 1991]. The language is a formal one with semantics that provide a direct translation into the constructs of CLASSIC [Borgida, Brachman, McGuiness and Resnick 1989] and similar KR 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 that corresponds to an individual with a role filled by the knowledge base within the container.

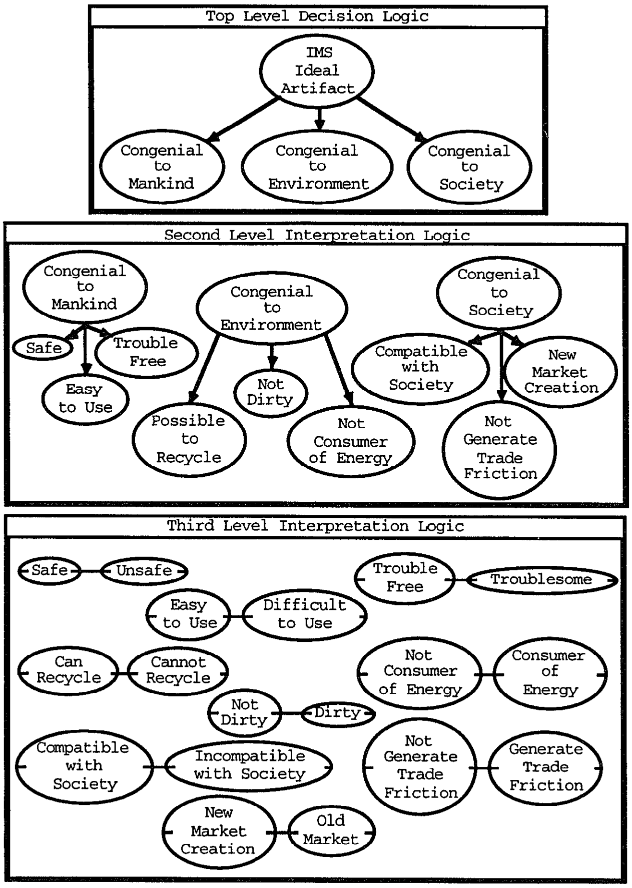

Figure 4 Knowledge structures for concept of IMS

At the top of Figure 4 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 [Sergot 1988]. 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 4 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 Tomiyama's 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 [Lenat and Guha 1990].

The pattern of reasoning when the knowledge structure of Figure 4 is loaded into an inference engine is interesting because it does not follow the common stereotype of an 'expert system'. If an artifact is asserted to be an *ideal 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 [Devanbu, Selfridge, Ballard and Brachman 1989]. The pattern also characterizes scientific research in which conjectures are made that may be refuted [Popper 1963]. The conjectures are never verified and remain forever fallible, and new knowledge as it is added may lead to the refutation of past conjectures.

This pattern of reasoning overcomes one of the problems often noted for the representation and use of rules in CLASSIC-like systems, that the converse of a rule is not used in inference, that is that '$A \supset B$' is not used to infer that '$\neg B \supset \neg A$' [Buchheit, Donini and Schaerf 1993]. The use of contradiction to cause retraction is equivalent to this inference, but it is only used with concrete assertions about individuals—it is a feature of the A-box not the T-box. It does not lead to intensional inferences of a theorem-proving nature in current implementations of CLASSIC-like systems, but it is very significant to practical reasoning about concrete assertions and conjectures.

The dimensions at the lowest level in Figure 4 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.

Groupware Knowledge Support Systems

The knowledge structures presented in this paper are from a groupware coordination tool, Mediator, designed to share knowledge structures across local and wide area networks. One application is to knowledge systematization through the product life cycle, from needs through requirements, design, engineering, production, maintenance, reuse and recycling. Another application is to project management, including distributed research activities. Figure 2 is a screen dump from Mediator in a project management role.

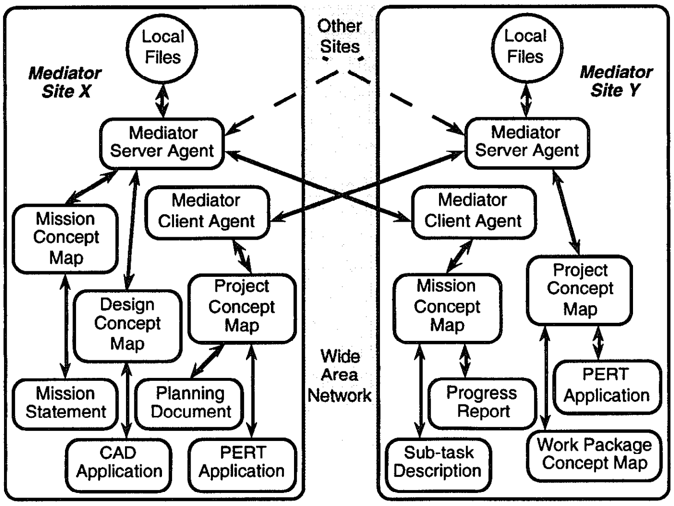

Figure 5 Mediator knowledge-based coordinator

Figure 5 shows the architecture of Mediator. A server agent at a site manages a knowledge base consisting of a heterogeneous set of files from different applications. Concept maps are used to represent the files and relations between them. Files may be opened from the maps in the appropriate applications. Since the maps and hypermedia documents of Mediator are also files, the system can be used to support large-scale linked knowledge structures.

Client agents at remote sites connect to server agents across the network and allow files to be accessed remotely in the same way as they are locally. A write token for each file is passed around the network allowing collaborative development of knowledge structures. The data structures for the visual languages are very compact allowing real-time updates with network data rates as low as 1 Kbyte/sec.

In research coordination the top-level knowledge structures are concerned with the mission. They link to structures concerned with technical projects and, since the visual language tools support Petrinets, knowledge bases, STEP/EXPRESS graphic representations, bond graphs, and so on, much of the detailed technical material can be captured in Mediator. At the lowest level files represented in Mediator may be opened in unrelated applications. All of the Mediator knowledge structures may also be embedded in active, printable documents [Gaines and Shaw 1993c], so that reports, manuals, and so on, are readily generated.

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 [Swaminathan 1992]. Computer-based concept maps have been used to support scientific knowledge processes [Smolensky, Bell, Fox, King and Lewis 1987; Cavalli-Sforza et al. 1992].

Research on shared ontologies and knowledge interchange formats is relevant to the development of knowledge support systems as envisioned in this paper [Neches, Fikes, Finin, Gruber, Patil, Senator and Swartout 1991]. 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 4. 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 [Lederberg and Uncapher 1989], and on the empirical study of the social and knowledge processes underlying scientific activity [Mitroff 1974; Abelson 1990; McCain 1990]. 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

Knowledge acquisition and representation technology is now mature and useful enough to be used routinely in supporting research communities. Much of what needs to be done does not involve major advances in artificial intelligence research, but rather the integration of what has already been achieved with existing communication systems on the Internet. However, the day to day use of overt knowledge structures by world-class scientific communities collaborating through the Internet is likely to highlight both the achievements of artificial intelligence research to date, and also the deficiencies of current KA and KR systems on different dimensions to current analyses largely motivated by considerations of pure logic and worst-case complexity analysis. The support of human practical reasoning may place lesser requirements than expected on some aspects of KA and KR systems, and greater requirements than expected on others. The development of practical knowledge support systems on the Internet and the resultant human-computer symbiosis at the knowledge level provides a stimulating environment for a new thrust in artificial intelligence research.

References

Abelson, P. 1990. Mechanisms for evaluating scientific information and the role of peer review. *Journal American Society Information Science* 41(3) 216-222.

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

Boose, J.H. 1984. Personal construct theory and the transfer of human expertise. *Proceedings AAAI-84*. pp.27-33. California: AAAI.

Borgida, A., Brachman, R.J., McGuiness, D.L. and Resnick, L.A. 1989. CLASSIC: a structural data model for objects. *Proceedings of 1989 SIGMOD Conference on the Management of Data*. pp.58-67. NY: ACM Press.

Buchheit, M., Donini, F.M. and Schaerf, A. 1993. Decidable reasoning in terminological knowledge representation systems. *Journal of Artificial Intelligence Research* 1 109-138.

Callon, M., Law, J. and Rip, A., Ed. 1986. *Mapping the Dynamics of Science and Technology*. UK: MacMillan.

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

Cole, S. 1992. *Making Science: Between Nature and Society*. Cambridge, MA: Harvard University Press.

Collins, H.M. 1990. *Artificial Experts: Social Knowledge and Intelligent Machines*. Cambridge, MA: MIT Press.

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

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

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

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

Gaines, B.R. and Shaw, M.L.G. 1989. Comparing the conceptual systems of experts. *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*. San Mateo, California: Morgan Kaufmann.

Gaines, B.R. and Shaw, M.L.G. 1992. Integrated knowledge acquisition architectures. *Journal for Intelligent Information Systems* 1(1) 9-34.

Gaines, B.R. and Shaw, M.L.G. 1993a. Basing knowledge acquisition tools in personal construct psychology. *Knowledge Engineering Review* 8(1) 49-85.

Gaines, B.R. and Shaw, M.L.G. 1993b. Eliciting knowledge and transferring it effectively to a knowledge-based systems. *IEEE Transactions on Knowledge and Data Engineering* 5(1) 4-14.

Gaines, B.R. and Shaw, M.L.G. 1993c. Open architecture multimedia documents. *Proceedings of ACM Multimedia 93*. pp.137-146.

GNOSIS, Ed. 1993. *TW4 Soft Machinery Workshop Proceedings*. Canada: KSI, University of Calgary.

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

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

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

Lambiotte, J.G., Dansereau, D.F., Cross, D.R. and Reynolds, S.B. 1989. Multirelational semantic maps. *Educational Psychology Review* 1(4) 331-367.

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

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

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

McCain, K.W. 1990. Mapping authors in intellectual space: a technical overview. *Journal American Society Information Science* 41(6) 433-443.

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

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

Neches, R., Fikes, R., Finin, T., Gruber, T., Patil, R., Senator, T. and Swartout, W.R. 1991. Enabling technology for knowledge sharing. *AI Magazine* 12(3) 36-56.

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

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

Quillian, M.R. 1968. Semantic memory. *Semantic Information Processing*. pp.216-270. Cambridge, Massachusetts: MIT Press.

Sergot, M. 1988. Representing legislation as logic programs. *Machine Intelligence 11*. pp.209-260. Oxford: Clarendon Press.

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

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

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

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

Wan, D. and Johnson, P. 1992. Supporting scientific learning and research review using COREVIEW. *AAAI-92 Workshop on Communicating Scientific and Technical Knowledge*. pp.107-114. Menlo Park, California: AAAI.