CONCEPT MAPS INDEXING MULTIMEDIA KNOWLEDGE BASES

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

Concept maps have long been used as a method to structure argument forms, and as visual languages for expressing relationships between ideas. An open architecture concept mapping tool is described for indexing and accessing multimedia material. The tool allows visual languages to be defined for specific purposes, and the user interface to concept maps in the languages to be programmed to provide access to other material including multimedia display tools and concept maps. The tool has been used in a wide variety of applications from very informal concept maps, to formal representations of knowledge structures with precisely defined semantics. It has been used to link the visual language material with multimedia material, text, diagrams, pictures, sound and video. It is suggested that a set of linked maps in a variety of visual languages is a natural basis for indexing and navigating a complex database of multimedia material.

INTRODUCTION

The work reported in this paper has originated from the development of systems to support large volumes of heterogeneous multimedia material generated in a variety of contexts such as knowledge acquisition (Gaines and Shaw, 1992), sports coaching (Vickers and Gaines, 1988) and large-scale project support (Gaines and Norrie, 1994). The primary requirement has been to give end users access to multimedia material through indexing mechanisms that are simple and natural to use. A further requirement has been to support shared access to material through a variety of interfaces that reuse the same material for different purposes.

These requirements have been addressed by the development of a general visual language technology supporting customizable interactive concept maps (Gaines and Shaw, 1993c) and semantic networks (Gaines, 1991). The maps may be used as stand-alone documents or embedded as interactive pictures in active documents (Gaines and Shaw, 1993b). The technology is open architecture, and user interaction with the concept maps may be programmed to initiate any activity feasible on the host system. Thus maps may be linked to other maps for retrieval purposes, and may be used to retrieve, play and edit multimedia material either through the host system or under its control. The technology is designed to be shared and operate over local or wide area networks so that multiple users can access the same maps, and maps can be used to retrieve material from remote sites.

Informal concept maps may be used primarily for indexing and retrieval of material, providing an attractive, meaningful and easy to use interface. Concept maps with formal semantics may be used to provide a programming interface for the control of multimedia material. For example, Petrinets may be used to specify the essential synchronizations in the playback of a mix of multimedia material, and to index through the time relations in such material. Semantic networks may be used to represent knowledge bases that are operational and can be used for knowledge-based access to material.

It is suggested that concept mapping technology provides a powerful general-purpose technology for the indexing and reuse of multimedia material.

KMAP: AN OPEN ARCHITECTURE VISUAL LANGUAGE TOOL

KMap is an open architecture visual language system developed support graphic interaction with computers on a wide range of topics including knowledge representation, concept mapping, Petrinets, bond graphs, and so on. The system was developed as an alternative to textual interfaces to take advantage of modern graphic workstations through a simple and natural visual language of great generality. The generality is achieved because the underlying data structure is a sorted digraph which can represent typed binary relations and hence virtually any mathematical structure. The provision of simple graphic expression for equality, inequality and subgraph operations allow complex conditions on graphs to be represented in a natural way. An attractive and comprehensible user interface is provided through the capability to reflect the type structure in node decorations such as fonts, shading and color.

Concept maps have long been used as a method to structure argument forms, and as a visual language for expressing relationships between ideas (Gaines and Shaw, 1993c). Concept maps have been used as tools to support the interviewing process in knowledge acquisition from experts, for example in the Wright-Patterson development of the pilot's associate (McNeese, Zaff, Peio, Snyder, Duncan and McFarren, 1990). In management, Axelrod (1976) proposed cognitive maps as a means of representing the conceptual structures underlying decision making, and these have been used empirically to analyze organizational decision making (Eden, Jones and Sims, 1979) social systems (Banathy, 1991) and the policies of political leaders (Hart, 1977).

In linguistics, Graesser and Clark (1985) have developed an analysis of argument forms in text in terms of structured concept maps with eight node types and four link types, and Woodward (1990) has developed tools to extract such concept maps from text. They have also been used extensively in education, mainly to investigate a student's understanding of some topic such as concepts in science (Novak and Gowin, 1984), and there are many different forms that have been applied in this field (Lambiotte, Dansereau, Cross and Reynolds, 1989). In the history of science, the dynamics of concept maps have been used to represent the processes of conceptual change in scientific revolutions (Nersessian, 1989; Thadgard, 1992). In the philosophy of science, Toulmin (1958) developed a theory of scientific argument based on typed concept maps that is regarded as one of the major models of the rhetoric of western thought (Golden, Berquist and Coleman, 1976).

Visual presentation of knowledge structures has been an attractive feature of semantic networks since their inception (Quillian, 1968). Computer production of visual forms of knowledge represented in a computer has been a topic of research since the early days of knowledge representation research (Schmolze, 1983) and a feature of many research systems (Kindermann and Quantz, 1989) and commercial products. In knowledge acquisition, in particular, the presentation of formal knowledge to those from whom it has been elicited is important to its validation (Gaines and Shaw, 1993a). There are many techniques for such acquisition but they ultimately result in an operational knowledge base with formal semantics (Gaines and Shaw, 1992). However, the expression of this knowledge base in the formal language used by the system is usually not very comprehensible to non-programmers (Nosek and Roth, 1990). A visual language that is both comprehensible and formal offers attractive possibilities not only for comprehension but also for editing, and for parts of the knowledge acquisition process itself (Gaines, 1991).

KMap has been used in the wide variety of applications outlined in this section, from very informal concept maps, to formal representations of knowledge structures with precisely defined semantics. What is significant in the context of this workshop is that in all applications it has been found appropriate to link the visual language material with multimedia material, text, diagrams, pictures, sound and video. A set of linked maps in a variety of visual languages is a natural mechanism for indexing and navigating a complex database of multimedia material.

KMAP IN ACTION

User interaction with the concept maps takes place through the creation of statements in the visual language, and through interaction with existing statements through popup menus whose content is specific to node type. The actual action initiated is context-sensitive: to the node selected for the popup, to nodes linked to it, and to other nodes preselected by clicking on them in the graph. This allows both simple and complex activities to be initiated by simple and comprehensible user actions.

KMap is itself programmable through the Apple high-level object event protocol (Apple, 1993a), and can be driven by any of the scripting languages in Apple's open scripting architecture such as Applescript (Goodman, 1993), Frontier (Winer, 1992) and TCL (Ousterhout, 1994). It can also run scripts triggered by user interaction. The combination of these two capabilities enables KMap to be integrated with other applications, and user interaction with graphical structures in the visual language to be used to control any activity supported on the host computer or network. The development of a specific system involving writing scripts to provide the required functionality by drawing upon existing applications.

Figure 1 shows a user interacting with a concept map of multimedia materials concerned with the field identification of North American birds. The map itself gives an overview of the topic and an index to available material. As the user mouses over a node a popup menu symbol appears. Clicking on this displays a menu that can be used to access the material, display text, play a sound, run a movie, and so on. The menu is generated from an underlying script, and user actions are reported to the script. The material accessed can be retrieved using either standard facilities for sound and movie playing within the application, or through commands which initiate another application and appropriate dataset.

Note that the concept map shown is not a formal knowledge structure, although it has some degree of hierarchical organization. It is basically a visual index designed to provide users with access to a collection of materials through a meaningful diagram.

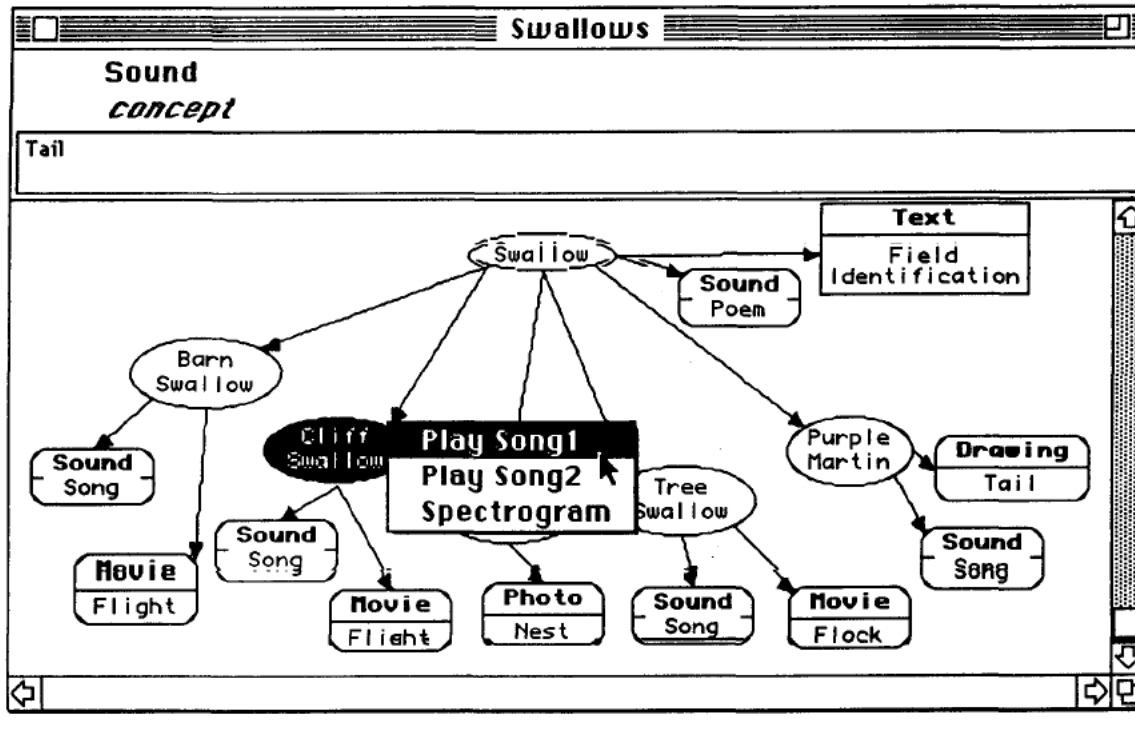

Figure 1 Concept map accessing multimedia materials

The actions on the popup menu can include the opening of other concept maps so that is possible to index a large body of material through layers of connected maps. It can also include opening files in other applications which may themselves make calls back to the concept mapping tool. Hence complex heterogeneous systems may be developed if needed for particular applications.

Figure 2 shows the generation of the visual language being used in Figure 1. The system designer generates nodes of different types appropriate to the domain of materials to be accessed. The visual features of each type can be designed to be distinctive and attractive. The script can associate simple retrieval activities for each type dependent on the data stored with each node, or, since it has full access to the concept map, can take context-sensitive actions as appropriate.

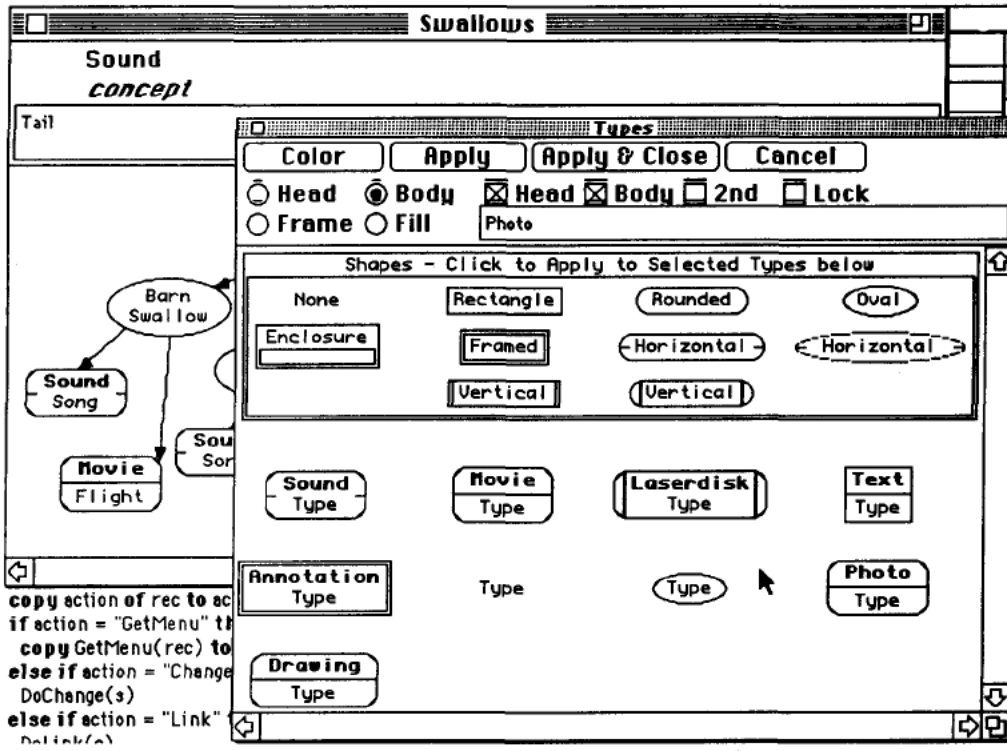

Figure 2 Entering concept map types to define a visual language

KWRITE: A DOCUMENT APPLICATION

KMap is implemented in a C++ class library as a sub class of a general 'pane' class that may be embedded as a component of any window. Such panes can be embedded as 'picture' components in the word processing and page makeup class also provided in the class library. The word processing class allows panes to be embedded within a document in such a way that each pane component is relatively independent of the text functionality. The text is laid out to run around or hop over pane areas but is not otherwise affected by them.

The embedded panes may be used to display graphic material embedded in the document including making them available to other applications as if each were a drawing pane in an arbitrary window. This enables them to be used to support visual activity ranging from bit map and line graphics to QuickTime (Drucker and Murie, 1992) or laserdisk videos, through simulation and animation, to graphic editors for visual languages representing programs, concept maps, and formal knowledge structures. Mouse down clicks within a pane are reported to the associated application rather than to the document software, and hence user interaction can be supported in a completely different environment. Such functionality is now becoming commonly available through the OLE 2 (Microsoft, 1994) and OpenDoc (Apple, 1993b) operating system extensions and allows many applications to be combined in a compound document.

Figure 3 shows a concept map of a knowledge structure developed for coaching ice hockey (Vickers, 1990) embedded in a document together with a QuickTime video showing the relevant activities. The material relates to some earlier studies providing a detailed comparison of laserdisks and books as media through which to present the coaching material (Vickers and Gaines, 1988). When the multimedia document technology became available some of the video material was digitized and incorporated together with text and active knowledge structures in documents that emulate sections of the existing books and laserdisk material. At the top of Figure 3 is the relevant part of the overall knowledge structure as an active concept map. As the user mouses over a node in the map a popup menu symbol appears. Clicking on this sends a message to a script associated with the document that determines the menu items, in this case the capability to play a video clip. The selection of an item sends a message to the script which then plays the appropriate section of the QuickTime movie below the knowledge structure.

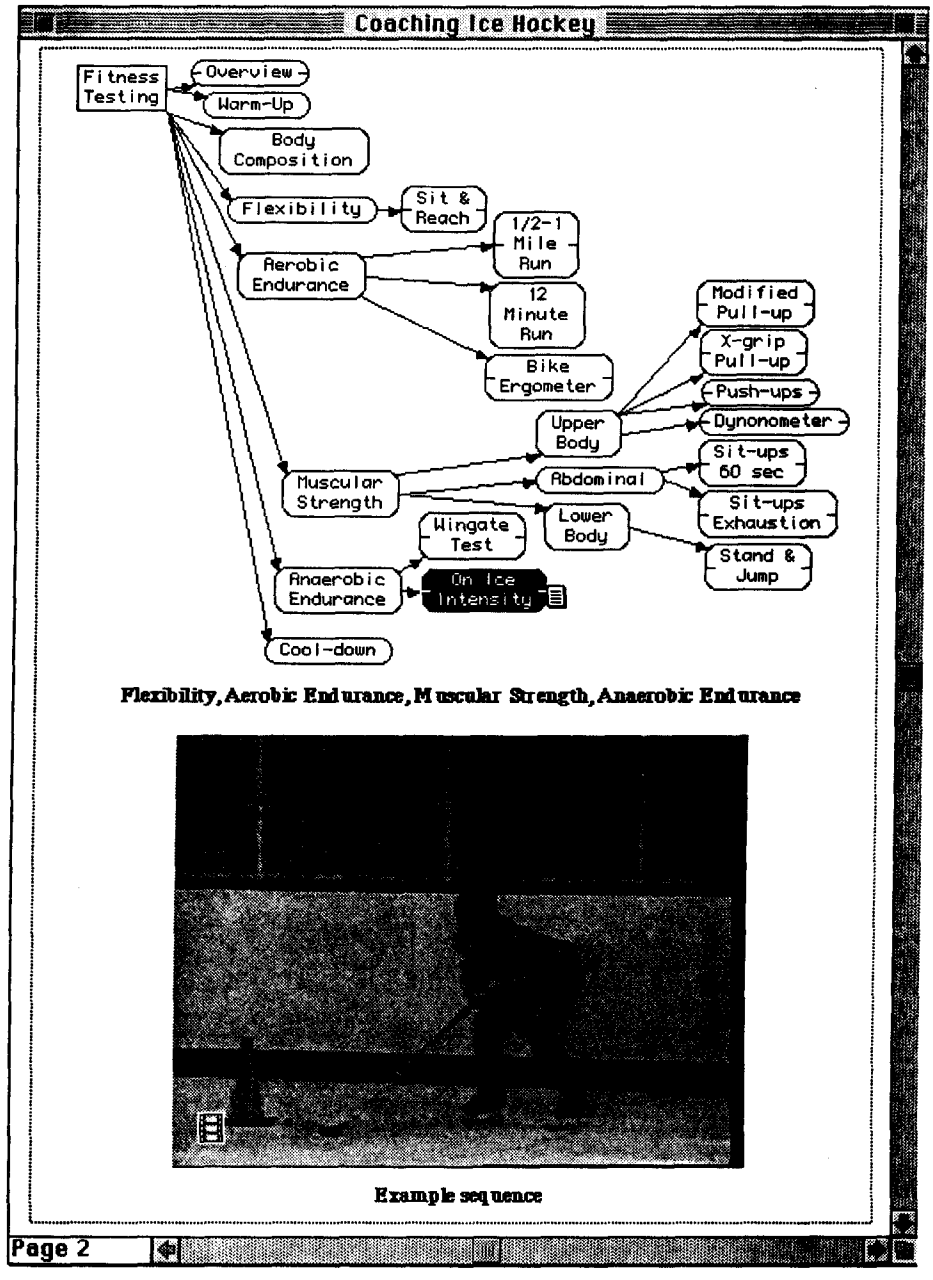

Figure 3 Concept map of coaching knowledge structure linked to video exemplars

Producing the document shown in Figure 3 is a simple matter of selecting the 'Place' dialog in the 'File' menu and selecting the files for the concept map and the video respectively. The concept map is automatically active and can be edited in place, including entering the links to the appropriate sections of the movie. The underlying script for the multimedia documents provides a generic linkage facility and does not have to be edited by the document producers. However, it is accessible to them if they wish to enhance the facilities in some way.

SUPPORT OF SCIENTIFIC COMMUNITIES

KMap was designed to support collaborative editing of concept maps over a network, and to be able to manage the archives of geographically distributed communities. As part of the GNOSIS test case in the international Intelligent Manufacturing Systems research program, a version of KMap is being developed to manage manufacturing activities throughout the product life cycle from initial needs through design, engineering, production, to reuse and recycling (Gaines and Norrie, 1994) More interesting in general terms, is that the same technology is being used to manage the international research project itself, its communications, project management, and document archives. The 1993 archives of GNOSIS have been issued on CD-ROM with KMap indexing the material involved, including some 2,000 pages of documents, 200 pictures, 10 movies, and several software demonstrations. The KMap system is currently being ported under Unix to operate over the Internet in order to support the long-term GNOSIS project commencing in 1995. A study using KMap to define and analyze major policy and technical issues in the GNOSIS project itself is being reported at AAAI94 (Gaines and Shaw, 1994).

Generating Concept Maps by Text Analysis

The GNOSIS studies provide some interesting examples of generating concept maps automatically rather than by manual entry. Figure 4 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 IMS 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 4 but are neutral to the major issues of the IMS program on the right.

The text analysis package provides links back to the relevant sections of the document from which the map was derived. Links may also be added to other documents such as the progress reports on the associated projects.

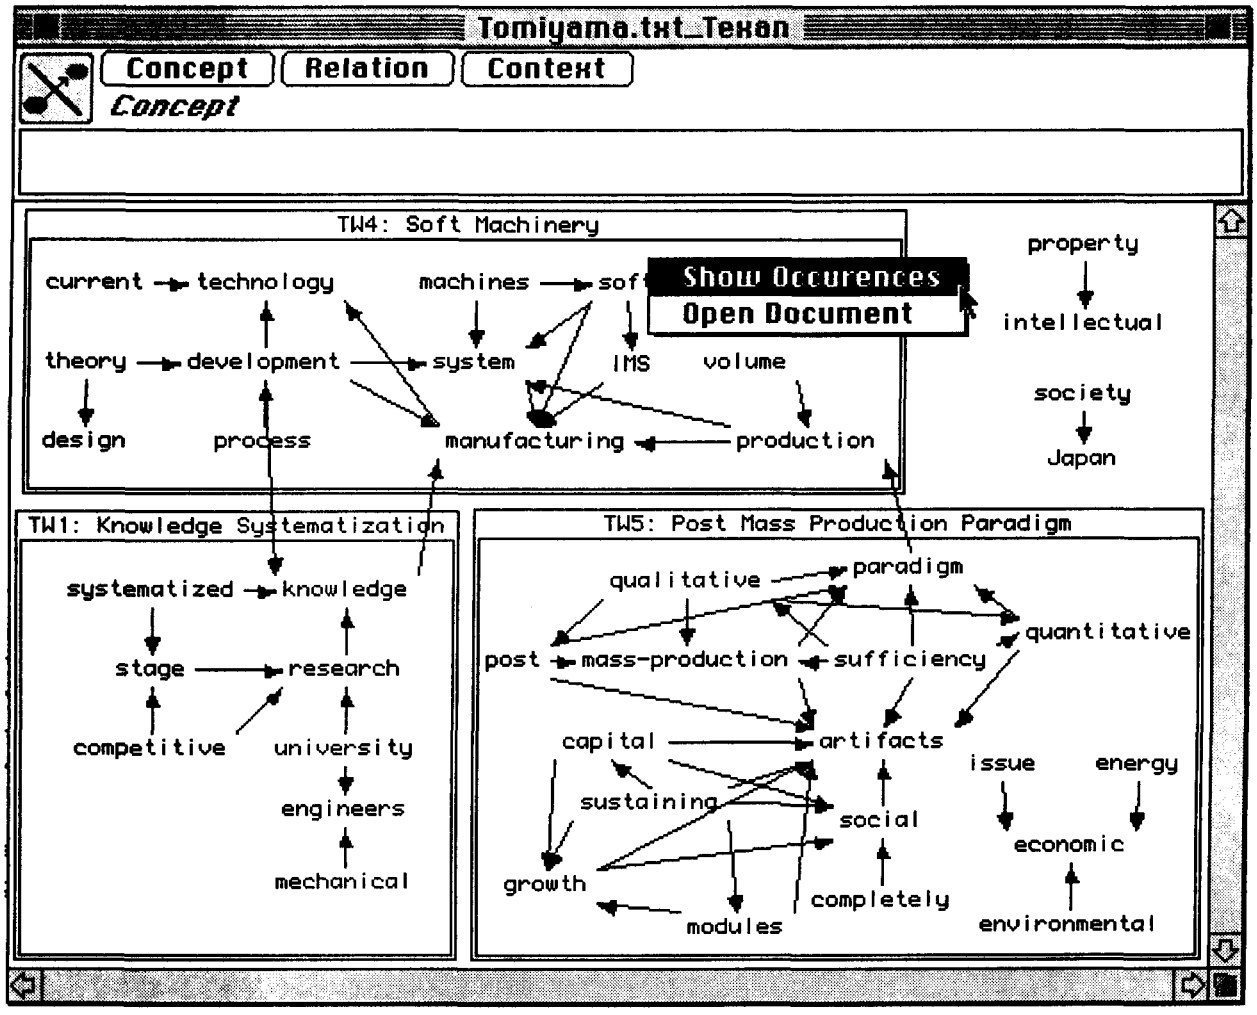

Figure 4 Concept map of GNOSIS project derived by text analysis

Indexing Project Archives through Concept Maps

Figure 5 shows the GNOSIS project archives being accessed through layered concepts maps. The map in the window at the upper left is a top level "Server Agent" that manages a particular collection of material. In the example shown a local user is accessing material directly through this agent. Remote users connect to the server agent over the network using client agents that give them the same functionality through calls to the server.

The concept map at the top left is currently write-disabled, and the cursor has changed to a button as the user mouses over the "Group Photo" node. Clicking at this point will display the photograph in a separate window.

The user has already clicked on the node "GNOSIS Final Reports" to open the concept map shown at the lower right. This has a node for each report, and clicking on one will open the appropriate report, in this application using Farallon's *Replica* tool to make the material available in a platform-independent form. The node at the top left gives access to a series of slides on the project displayed using Persuasion's GIF *Player* tool. A similar node in the original concept map at the top left gives access to a movie on the GNOSIS project that will be opened in Apple's *MoviePlayer*.

The user has already clicked on the "TW4 Workshop Papers" node in the concept map at the top left, and opened the relevant concept map at the bottom left. She has then clicked on the node

"Click here to see the report in Replica", and opened the report visible at the back on the right of Figure 5. She has also clicked on the "Click here to see all the movies" node and opened the KWrite document visible behind the concept maps. This displays eight QuickTime movies of various demonstrations given at the Workshop, any of which can be played by double clicking upon it.

This example gives some feeling for how the tools may be used to index a large heterogeneous collection of multimedia material. The network capabilities allow multiple users at geographically dispersed sites to collaborate in maintaining the archives. In this particular example the annual report has been issued on CD-ROM to allow large movies to be included that it would not be feasible to access over the Internet.

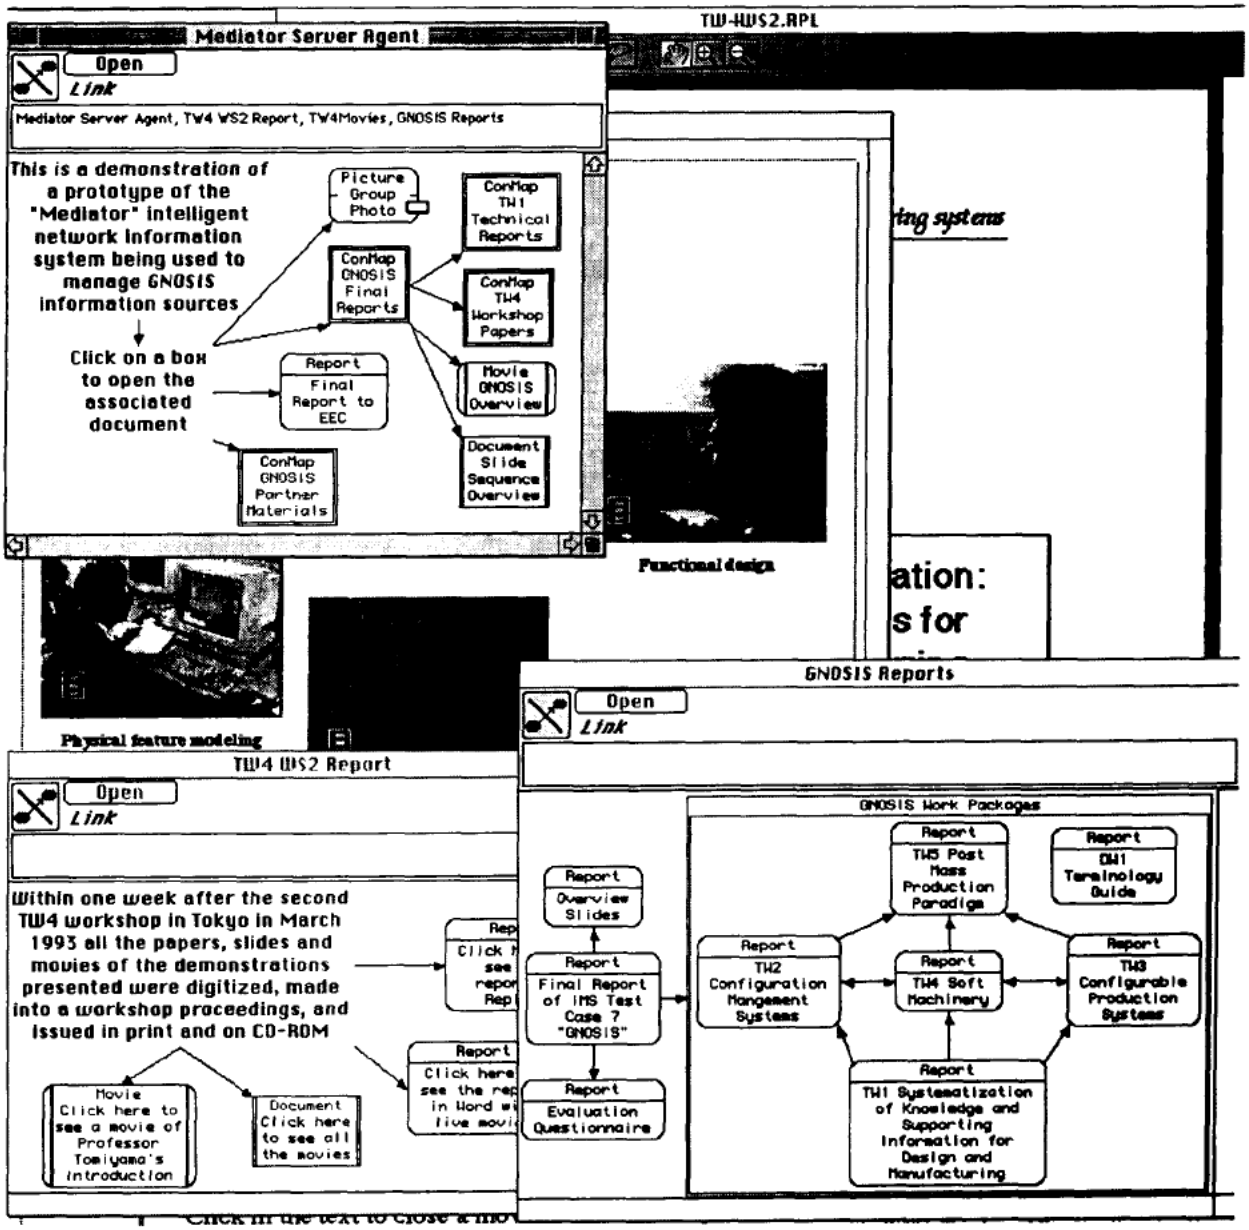

Figure 5 Accessing the GNOSIS archives through layered concept maps

"Click here to see the report in Replica", and opened the report visible at the back on the right of Figure 5. She has also clicked on the "Click here to see all the movies" node and opened the KWrite document visible behind the concept maps. This displays eight QuickTime movies of various demonstrations given at the Workshop, any of which can be played by double clicking upon it.

This example gives some feeling for how the tools may be used to index a large heterogeneous collection of multimedia material. The network capabilities allow multiple users at geographically dispersed sites to collaborate in maintaining the archives. In this particular example the annual report has been issued on CD-ROM to allow large movies to be included that it would not be feasible to access over the Internet.

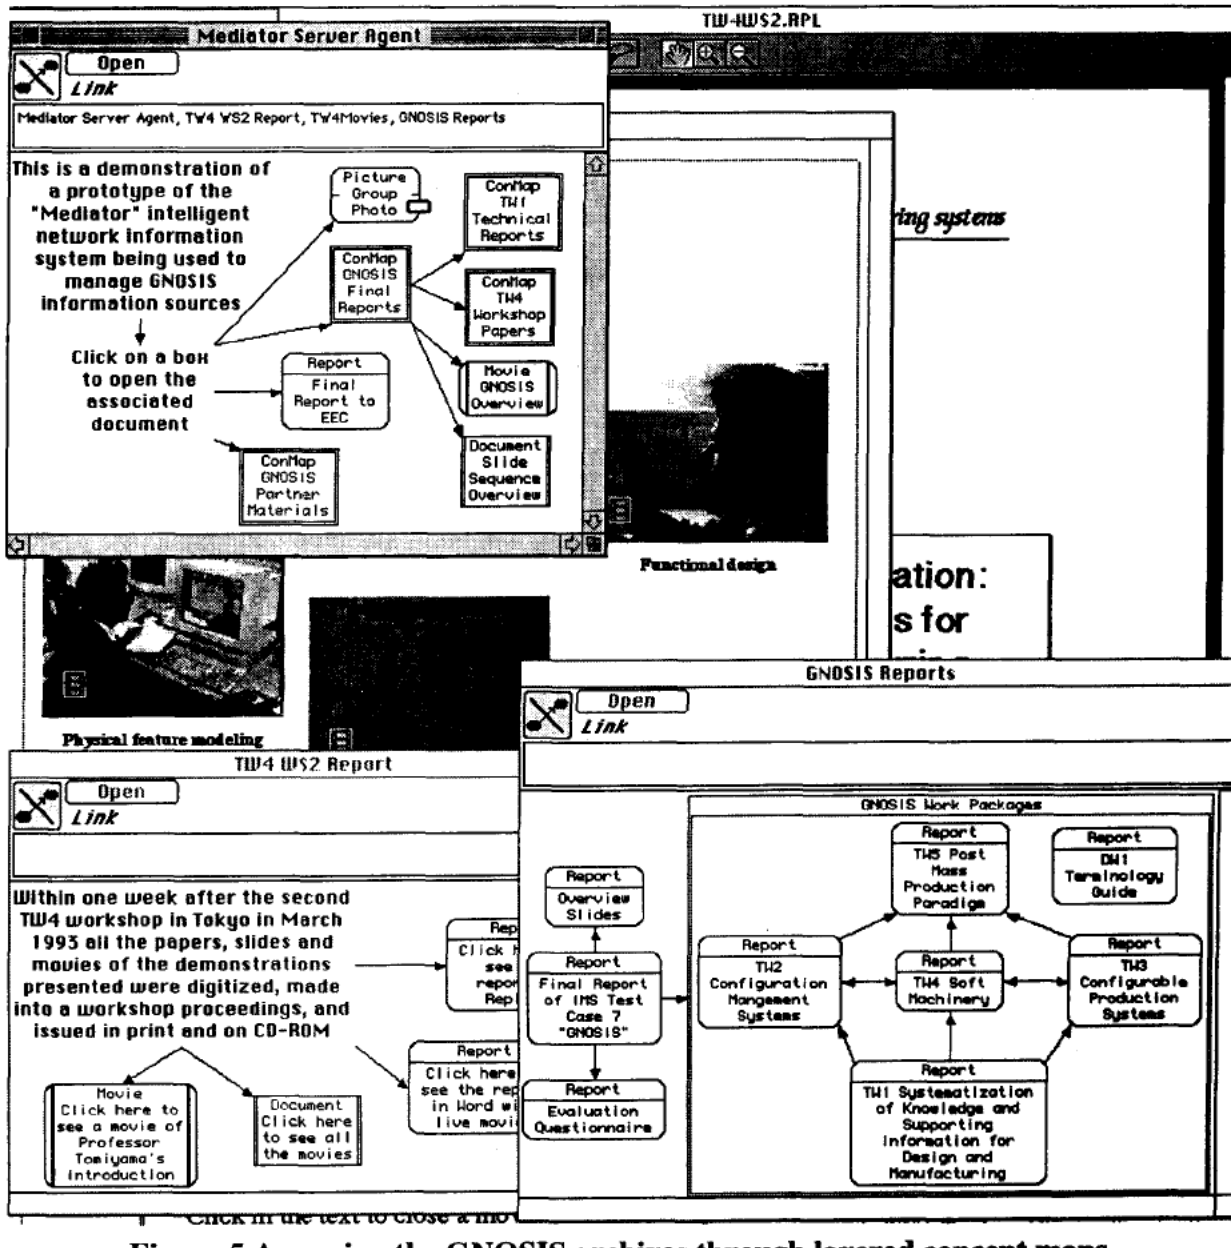

Figure 5 Accessing the GNOSIS archives through layered concept maps

CONCLUSIONS

Concept maps have long been used as a method to structure argument forms, and as visual languages for expressing relationships between ideas. An open architecture concept mapping tool has been described for indexing and accessing multimedia material. The tool allows visual languages to be defined for specific purposes, and the user interface to concept maps in the languages to be programmed to provide access to other material including multimedia display tools and concept maps. The tool has been used in a wide variety of applications from very informal concept maps, to formal representations of knowledge structures with precisely defined semantics. It has been used to link the visual language material with multimedia material, text, diagrams, pictures, sound and video. It is suggested that a set of linked maps in a variety of visual languages is a natural basis for indexing and navigating a complex database of multimedia material.

The system shown currently operates on the Macintosh platform. A version of it has already been ported to Windows, and a port is currently taking place to Unix to operate in collaboration with NCSA Mosaic .

ACKNOWLEDGMENTS

This work was funded in part by the Natural Sciences and Engineering Research Council of Canada and by Industry, Science and Technology, Canada.

REFERENCES

Apple (1993a). **Inside Macintosh Interapplication Communication**. Reading, Massachusetts, Addison-Wesley.

Apple (1993b). **OpenDoc Developer Overview**. Cupertino, California, Apple Corporation.

Axelrod, R. (1976). **Structure of Decision**. Princeton, New Jersey, Princeton University Press.

Banathy, B.H. (1991). Cognitive mapping of educational systems for future generations. **World Futures 30**(1) 5-17.

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

Drucker, D.L. and Murie, M.D. (1992). **QuickTime Handbook**. Carmel, Indiana, Hayden.

Eden, C., Jones, S. and Sims, D. (1979). **Thinking in Organizations**. London, Macmillan.

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, California, Morgan Kaufmann.

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

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). Open architecture multimedia documents. **Proceedings of ACM Multimedia 93**. pp.137-146.

Gaines, B.R. and Shaw, M.L.G. (1993c). Supporting the creativity cycle through visual languages. **AAAI Spring Symposium: AI and Creativity**. pp.155-162. Menlo Park, California, AAAI.

Gaines, B.R. and Shaw, M.L.G. (1994). Using knowledge acquisition and representation tools to support scientific communities. **AAAI'94: Proceedings of the Twelfth National Conference on Artificial Intelligence**. pp.to appear. Menlo Park, California, AAAI Press/MIT Press.

Golden, J.L., Berquist, G.F. and Coleman, W.E. (1976). **The Rhethoric of Western Thought.** Dubuque, Idaho, Kendall/Hunt.

Goodman, D. (1993). **The Complete AppleScript Handbook.** New York, Random House.

Graesser, A.C. and Clark, L.F. (1985). **Structures and Procedures of Implicit Knowledge.** New Jersey, Ablex.

Hart, J.A. (1977). Cognitive maps of three latin american policy makers. **World Politics 30**(1) 115-140.

Kindermann, C. and Quantz, J. (1989). **Graphics-oriented user interfaces for KL-ONE. KIT Internal Report 23.** Berlin, Technical University of Berlin.

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

McNeese, M.D., Zaff, B.S., Peio, K.J., Snyder, D.E., Duncan, J.C. and McFarren, M.R. (1990). An Advanced Knowledge and Design Acquisition Methodology for the Pilot's Associate. Harry G Armstrong Aerospace Medical Research Laboratory, Wright-Patterson Air Force Base, Ohio. AAMRL-TR-90-060.

Microsoft (1994). **OLE 2 Programmer's Reference Volume 1: Working with Windows Objects.** Redmond, Washington, Microsoft Press.

Nersessian, N.J. (1989). Conceptual change in science and in science education. **Synthese 80**(1) 163-184.

Nosek, J.T. and Roth, I. (1990). A comparison of formal knowledge representations as communication tools: predicate logic vs semantic network. **International Journal of Man-Machine Studies 33**(227-239)

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

Ousterhout, J.K. (1994). **Tcl and the Tk Toolkit.** Reading, Massachusetts, Adison-Wesley.

Quillian, M.R. (1968). Semantic memory. Minsky, M., Ed. **Semantic Information Processing.** pp.216-270. Cambridge, Massachusetts, MIT Press.

Schmolze, J. (1983). KLONEDRAW—a facility for automatically drawing pictures of KL-ONE networks. **Research in Knowledge Representation for Natural Language Understanding. Report No.5421.** pp.41-44. Cambridge, Massachusetts, Bolt Beranek and Newman.

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.

Vickers, J.N. (1990). **Instructional Design for Teaching Physical Activities: A Knowledge Structures Approach.** Champaign, Illinois, Human Kinetics.

Vickers, J.N. and Gaines, B.R. (1988). A comparison of books and hypermedia for knowledge-based sports coaching. **Microcomputers for Information Management 5**(1) 29-44.

Winer, D. (1992). **UserLand Frontier User Guide.** Palo Alto, California, UserLand Software.

Woodward, B. (1990). Knowledge engineering at the front-end: defining the domain. **Knowledge Acquisition** 2(1) 73-94.