Active Documents Combining Multimedia and Expert Systems

Mildred L G Shaw and Brian R Gaines

Knowledge Science Institute
University of Calgary
Calgary, Alberta, Canada T2N 1N4
{mildred, gaines}@cpsc.ucalgary.ca

Abstract: An active multimedia document publication system is described which integrates a number of different representation technologies to provide a medium offering a wide spectrum of usage, including new forms of scholarly communication. The capabilities range from emulation of current paper publication, through electronic document delivery on the Internet or CD-ROM, multimedia inclusion of video and sound, hypermedia linkage, to formal knowledge representation for simulation and inference. Aspects of the implementation are described, and examples are given of applications, including one in which an active document forms the knowledge base of an expert system.

Introduction

A problem with the concept of multimedia systems is that the term covers a diversity of technologies from fiber-optic ISDN and interactive cable television through laser light show presentations to office document architectures. Even in the specific context of multimedia documents, Herzner and Hocevar at the Eurographics Workshop in Stockholm explicitly avoided the issue of defining the term *multimedia*, noting that "its definition is possibly a major issue to be solved" [20]. This is not unreasonable in that every new type of document component one makes available to authors and readers may be regarded as an extension of the medium. Even the restrictions of paper allow for major differences in media type, such as text, diagrams and photographs. An in-depth analysis reveals significant subtypes of text, such as bulleted lists, mathematical and chemical formula, and semiotic analysis of diagrams results in a large number of subtypes [2, 28]. The emphasis of much multimedia research has been on sound and video, but it is important to realize that any operational computational structure incorporated in a medium in such a way as to affect its presentation may be providing new multimedia capabilities. This is obvious when a mathematical model is used to simulate a physical system and drive an animation, but it is less so when, for example, a semantic net is used to represent knowledge and present deductive inferences.

The approach to multimedia systems adopted in the work reported in this article is pragmatic and eclectic. We do not need to define 'media' because the term is already well defined as a contraction of 'representational media', and we can see the term 'multi' as having a obvious connotation when applied to 'representational.'

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

Any new representational capability offered in a medium extends its multimedia capabilities. The key question is 'what difference does it make to users?' This approach is important when it comes to establishing a multimedia production environment. There will always be a population of users interested in experimenting with new technology, adding color, sound and video to documents just because it can be done and is not yet routine. Such users are significant in championing new technology but may not provide insights into its long-term significance. The important users are those with significant goals in their disciplines that are being seriously impeded by existing media, and where new techniques may make a major difference to progress in those disciplines. A problem with finding such users is that they may be unaware of the new technologies, and, worse, may have persuaded themselves that the limitations impeding them are something they have to live with to the extent that they have forgotten them. Explicit requirements specifications rarely exist in real user communities. They have to be generated through examples and experience.

The table in Figure 1 is useful in presenting multimedia document capabilities to potential user communities. It makes a two-way classification of potential media. On the vertical dimension, Popper's [27] classification of knowledge structures into three 'worlds' is used to separate the direct representation of the physical world, from its modeling in human minds, and from its overt modeling in scientific data and theories.

The central, world 2, component is important since the difference between direct representation and scientific modeling is taken for granted, but the role of personal expertise and discourse between scholars is downplayed in scientific publication. Its natural communication 'medium' is through visits to laboratories and through seminars and workshops.

	Passive	Active	Interactive
World 3 Overt Knowledge	Scholarly writing, Mathematics, Diagrams	Animation, Calculation, Deduction	Simulation, Reasoning, Operational & accessible models
World 2 Mentality, Discourse, Expertise	Photos, transcripts, protocols	Movies of Discourse and Expertise processes	Groupware, Networked discourse
World 1 Physical Reality	Phographs, Datasets	Sound, Movies, Dynamic graphs	Exploratory model building from datasets

Figure 1 Media analysis of active documents

However, it may well be past media limitations that blind us to the value of seeing a scholar in action—a photograph of a scientist in his or her laboratory has little value, but a movie of the discussion surrounding a critical experiment or theoretical issue may be highly informative.

The horizontal dimension classifies media by activity and interactivity. Multimedia supports both processes that are active in time such as sounds and videos, and also processes that are interactive in enabling the reader to explore variants on the material supplied.

Document Architecture

The typographic facilities of modern word processors, the availability of QuickTime [5] technology for the digitization of sound, movies and animations, and the support of its embedding in standard word processors, has made it possible to develop documents encompassing the *passive* and *active* columns in Figure 1. Support for the *interactive* functionality in the right column is becoming available in commercial document systems as heterogeneous multi-part document standards such as OLE 2 [24] and OpenDoc [1] become used in commercial products.

For research purposes, an open-architecture document production system has been developed that supports embedded panes which in the paper form appear as normal diagrams and pictures, but which in the electronic form are fully interactive windows into a heterogeneous and extensible range of computer applications. KWrite [14] provides a word processing and page makeup environment similar in functionality and quality to that of the commercial packages with which our users are familiar. As shown in Figure 2, it supports standard multimedia and hypermedia capabilities in this way, but also animation, simulation, active concept maps providing informal knowledge representation, and active semantic nets providing formal knowledge representation.

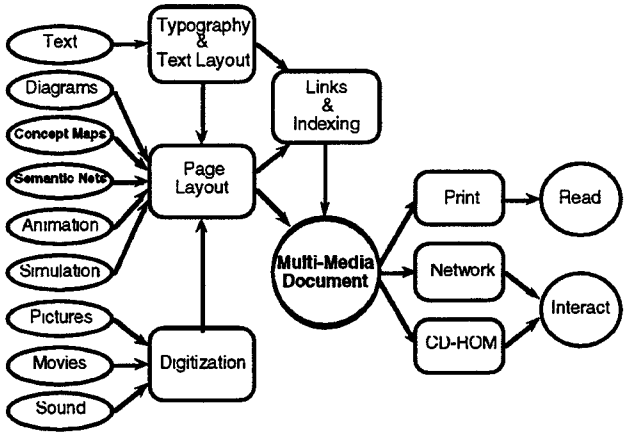

Figure 2 Multimedia document architecture

Interaction with documents is supported through a variety of mechanisms. An application controlling a pane in a document can provide its own interface, typically through popup menus and dialogs that do not appear when the document is printed. Script languages are also used to make documents 'programmable' independent of specific applications. KWrite combines these methodologies, providing script language access at a document level which can be used directly through embedded controls or indirectly by calls from embedded applications. This makes it simple to support common forms of hypertext functionality, and enables embedded applications, such as visual languages, to be provided in a generic form that is specialized as required through the script language. AppleScript [19] is supported on Macintosh systems and tcl [26] on Unix systems. A delegation mechanism is supported so that documents without script may be given active functionality by linking them to documents with scripts that act as 'agents' managing a corpus of active documents [10].

The multimedia document production environment is best illustrated in action through specific projects. The following sections give examples selected to address critical requirements in different disciplines and hence to provide exemplars and role models for other potential use.

Project 1: Rapid Workshop Documentation

A major KSI project in 1993 was to support the knowledge processes of an international manufacturing research project, IMS Test Case 7 'GNOSIS', involving over 100 participants in 31 industry and university organizations in 14 countries in Japan, Europe and North America. List servers and document archives were set up to minimize coordination costs, but physical meetings were still a major expense costing over $1M in the first year [18]. The typical structure of the initial workshops was a 'show and tell' session in which participants demonstrated the existing work in their research groups, gave an overview presentation, and distributed background documents.

Multimedia digital document technology was used to capture the contents of a workshop by video-taping demonstrations and major presentations, OCR'ing the documents circulated, and digitizing slides. This proved to be fast and effective in enabling structured workshop proceedings to be developed from relatively informal meetings. For example, a 3-day workshop in Tokyo in March 1993 involving some 30 participants resulted in a 300-page proceedings containing 30 articles and 8 QuickTime movies. The production time to OCR the source material, edit it into a uniform format, digitize the slides and movies, and issue the proceedings in print and CD-ROM was one week. The effort required was some 40 person-hours.

Figure 3 shows a page from the proceedings in Microsoft Word with the heading, comments, and movies of computer software demonstrations preceding a technical paper.

This application is significant to the support of the semi-formal discourse in scholarly communities that takes place through seminars, workshops and informal conferences. Much of what occurs at such meetings in currently lost and not disseminated to others, and yet digital document technology now makes it simple to capture the essence of a meeting without imposing constraints on the participants.

Project 2: Large QuickTime Animation

Digital video technology enables sound and movies to be embedded in multimedia documents. The publicized usage is for the digitization of movies as illustrated in the previous section. Without use of special hardware decoders the digital bandwidth of current general-purpose computers limits such movies to sizes of some 180 by 240 pixels with playback rates of about 15 frames a second. However, the underlying digital video technology is a simple and general way of showing a sequence of images with synchronized sound, and it can be used to support slide sequences and animations of much larger size because the video playback rate required is lower.

To demonstrate this capability a QuickTime movie was prepared for the MultiMedia'93 CD-ROM distributed at SIGGraph'93 which was the full size of the smallest Macintosh screen (342 by 512 pixels) and yet played effectively from a normal speed CD-ROM drive.

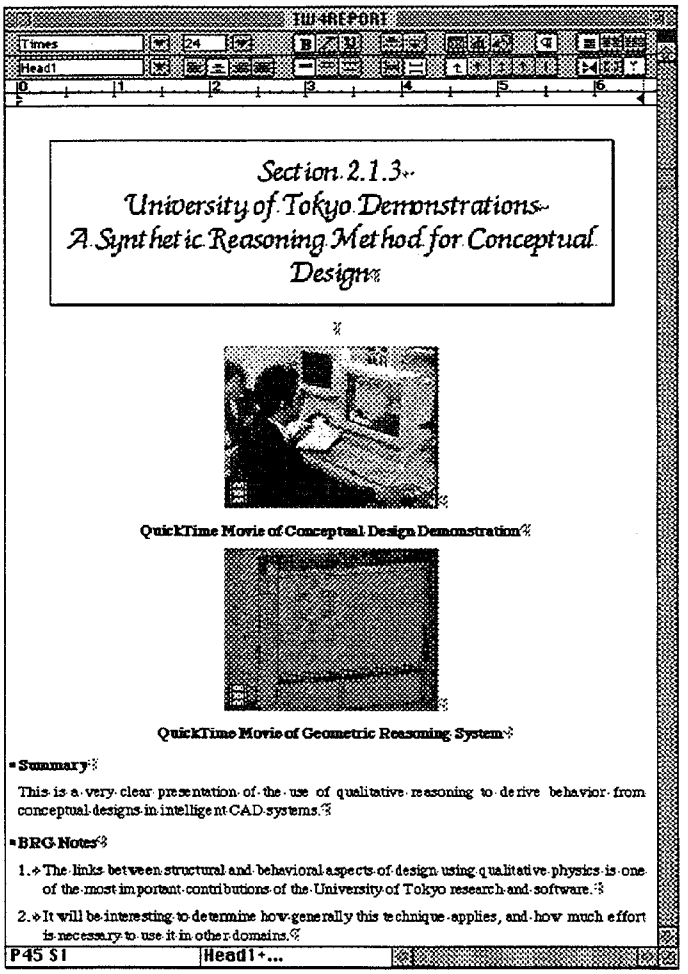

Figure 3 Workshop proceedings with embedded demonstrations

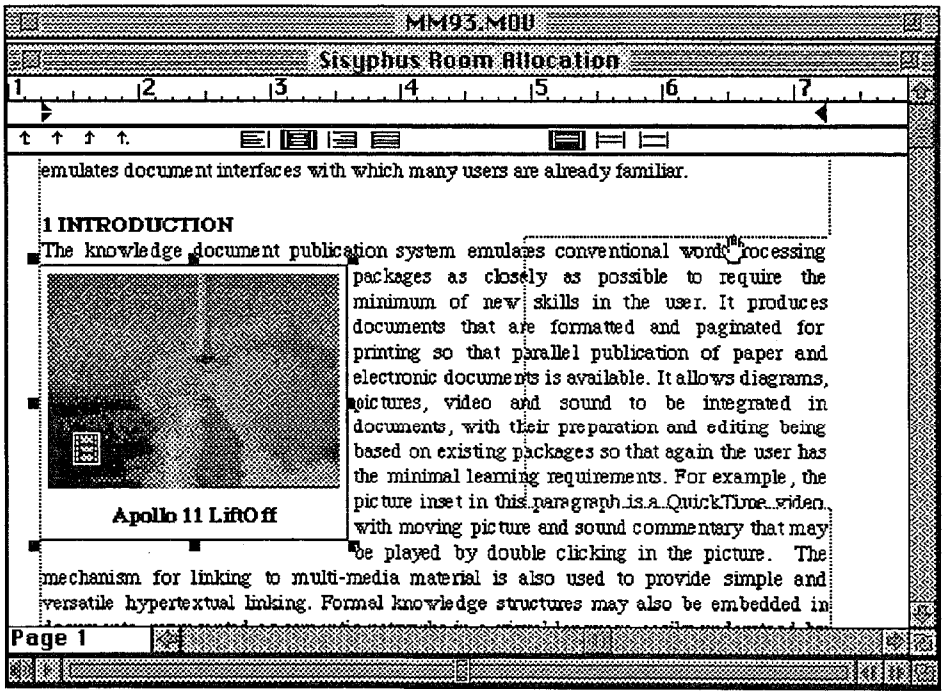

Figure 4 QuickTime capture of animation and software demonstration

The movie illustrates a paper [14] on KWrite, and consists of a commentary over an animated slide sequence on the technology and sequences of screen dumps of the knowledge acquisition and representation tools in use. Figure 4 shows a screen dump from this movie playing in MoviePlayer. The sequence was captured in CameraMan from a KWrite document and illustrates the way in which a picture pane may be dragged into place with text automatically wrapping around it. The picture in this case is itself an active QuickTime movie, and the sequence goes on to play this movie, capturing its playing and the accompanying sound within the larger movie—an example of recursion at work!

It is a simple process to develop animation sequences using standard drawing tools such as Deneba *Canvas* combined with QuickTime editors such as Adobe *Premiere*, and this provides a pathway to multimedia for users not prepared to deal with more complex tools such as MacroMind *Director*. Current personal computers support the image production very effectively. The greatest effort needs to be put into the sound capture system since, as usual, sound quality has a major influence on the effectiveness of the commentary. The technique used in the movie described was to record a spontaneous voice commentary as the movie was played, transcribe this as a script, and then record it again using a DAT recorder in a quiet environment. This enabled a natural commentary to be developed while ensuring high-quality sound.

Project 3: Combining Animation and Movies

Playing back a movie within a movie as illustrated above, shows how animation and movies can be combined. This is particularly effective in training material, tutorials and active manuals. For example, we have developed knowledge acquisition tools for direct use by experts [13]. They are based on psychological methodologies for concept elicitation, and there is significant skill in the way in which they are presented to users and the way in which user interaction is facilitated by discussion, particularly in group situations. If the tools are just given to experts to use it is unlikely they will be effective. It is very difficult to explain in words the procedures to be followed, yet that are easily acquired through experience working alongside a skilled facilitator.

Previous experience showed that movies of interaction with the computer are not very effective in either showing the detail of what is happening, or of putting it into the context of a complete knowledge elicitation process. What is required is a mix of text, diagrams, live video, and high resolution screen replays with voice commentary. Figure 5 shows a frame of a QuickTime movie produced by dumping out a sequence of interaction with an elicitation tool using CamerMan, and then using Premiere to insert a video clip of the users in the lower right corner. Commentary and the recorded discussion between the users is interspersed to capture a tutorial explanation of an actual application session.

The integration of animations and short movie sequence is a powerful technique for presenting tutorial examples in an application context. The screen dumps have precisely the same resolution and dynamics as the original interaction, and the inserted video clips are adequate to situate the use of the tool in the psychological and social context.

Project 4: Controlling Embedded Movies

QuickTime movies embedded in documents can be randomly accessed and played using the normal movie controller. In a programmable active document they can also be controlled from other applications. KWrite provides tools for concept map [15, 25] design supporting nodes of different shapes and colors, and linkage from nodes to scriptable document actions, such as playing QuickTime movies.

One project using these features is a continuation of some previous media studies undertaken with the KSI by Dr Joan Vickers, Head of the Neuro Motor Control Laboratory in the Faculty of Physical Education at the University of Calgary. Vickers has developed an approach to sports coaching based on knowledge structures for various sports expressed as hierarchical concept maps. Knowledge structures for specific sports have been developed by senior international coaches for those sports and have been used as the basis for an extensive book series on sports coaching [29].

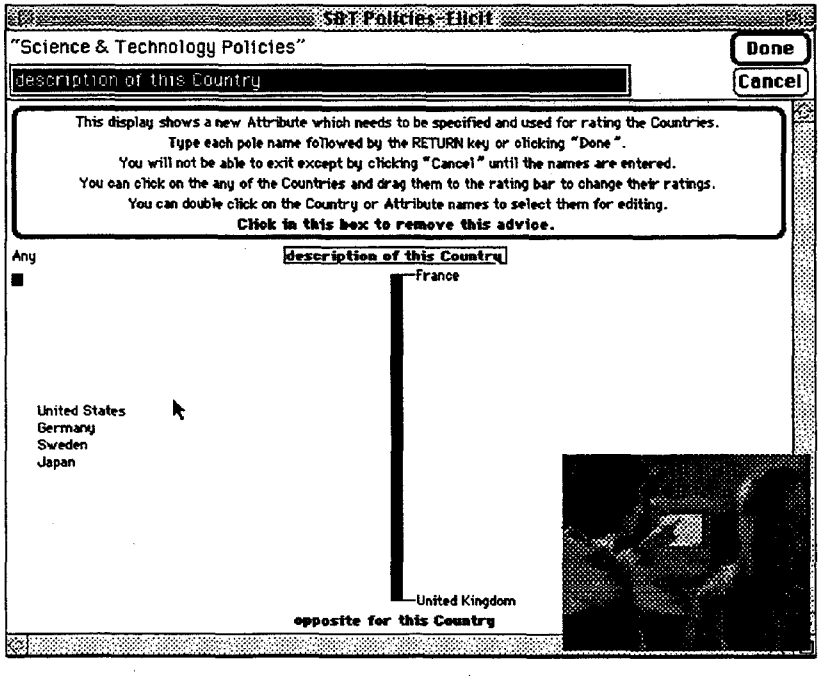

Figure 5 Active screen sequence as a QuickTime movie with inserted film clip

The original project took place in the mid-1980s as the knowledge structures were created and involved the production of laserdiscs incorporating video material illustrating athletes of various levels of skill in actual situations performing activities illustrating the knowledge structures. These discs were accessed through programs graphically representing the knowledge structures and managing the lesson sequences, drill and practice in class situations [16, 17]. One outcome of these studies was a detailed comparison of laserdiscs and books as media through which to present the coaching material [30].

As the new multimedia document technology became available it was interesting to investigate it as a third option in which the book and laserdisc material could be combined. 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 laserdisc material.

Figure 6 shows a screen dump of a page being accessed. At the top is the relevant part of the overall knowledge structure as a concept map. The popup menu accessible at the right of each node activates the relevant section of the QuickTime movie below the knowledge structure.

The coaching material is of interest not only because of the availability of comparisons with both paper and laserdisc material, but also because its presentation is driven by the knowledge structures presented as a concept maps. Such conceptual structures for documents have been proposed as a way of managing multimedia documents [23], and their acquisition is also related to the elicitation of semantic networks as a basis for structuring hypermedia documents [21]. Thus, the material is representative of some of the possibilities for new document architectures that continue to support parallel publication in paper and digital form yet have greater structure than conventional books and papers.

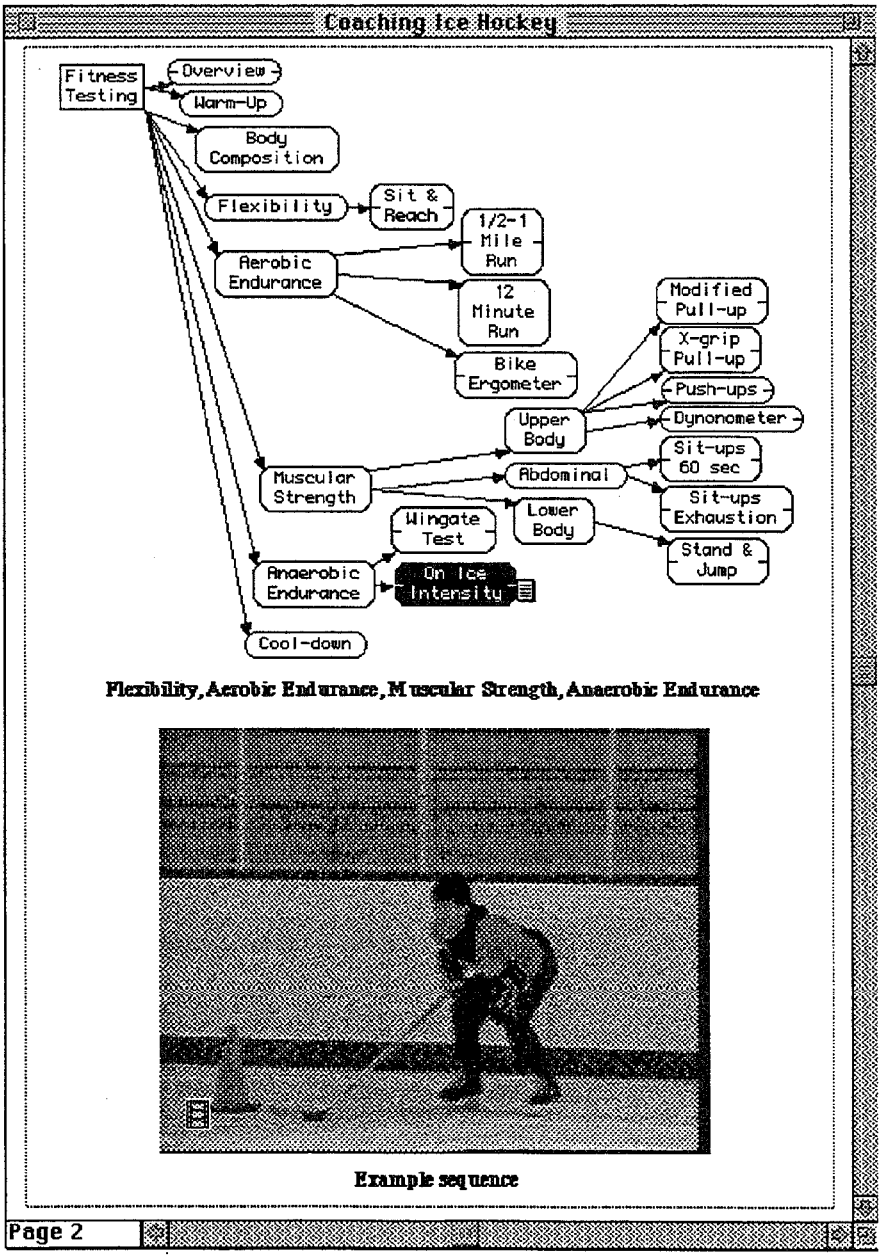

Figure 6 Knowledge structure linked to video exemplars

Project 5: A Document as a Knowledge Base

One reason for interest in multimedia scholarly publication is to increase the credibility of the results presented and reduce the possibility of misinterpretation or fraud. For example, Claerbout [3, 4] published a geological research book in conventional form and on an interactive CD-ROM with full datasets and analysis programs, giving the rationale:

> "The scandal of Nobelist Dr. David Baltimore being accused of fraud and ultimately stepping down as President of Rockefeller University carries many interesting lessons for laboratory scientists. There are lessons too for us whose research is heavily computer dependent. From a journal article alone, the task of reproducing the computations and data analysis leading to computer-generated illustrations is often formidable. Typically, the authors themselves, if stripped of their files, starting over again from only their data and their published paper, might require weeks or many months to reproduce the illustrations on which their conclusions are based. Referees and readers rarely attempt to reproduce anything because the effort could be comparable to the first two years of work on a dissertation. I thought this reproducibility problem hopelessly mired in human nature when, much to my amazement, I solved it by accident using paper-plus-electronic publication technology."

One KSI objective has been to show how multimedia publication could be used to present scientific material in operational form fully accessible to the reader. An available example from KSI research was that of knowledge-based systems in artificial intelligence [9, 12, 13]. One develops complex deductive engines, associated knowledge acquisition tools, and example applications, but when one publishes the results all one can present is screen dumps, timings, results and discussion. The reader has to take it on trust that the system exists, let alone works, and has no capability to see if the example is carefully chosen to show off the system to best advantage. Making the system separately available to other researchers overcomes some of these problems, but few people take advantage of such availability. Those who do so often obtain the system some time after having read the paper and many of the questions that they had in mind at the time of reading are forgotten.

Proper publication of material about an operational system should include immediate accessibility of that system. CD-ROM publication offers the opportunity to make programs and data available as auxiliary material. Ideally that material should be linked to the arguments in the paper in such a way that it is interactively accessible as part of the process of reading the book or paper.

To illustrate what can be achieved a knowledge-based system was put into active document form that had been developed as a solution to a 'challenge' problem circulated by Project Sisyphus, an initiative of the European Knowledge Acquisition Workshop. The problem is one of room allocation derived from an ESPRIT project [31] and a major part of the EKAW'91 program was devoted to reports on the solution of these problems using different approaches and techniques [22].

An article was published in a Cambridge University Press book containing the proceedings of the British Computer Society Expert Systems Conference in December 1992 in which the abstract states:

> "This paper is written in a document production tool that appears to a user as a word processor but also acts as an expert system shell with frame and rule representations supporting deductive inference. The electronic version of the document is active, providing typographic text and page layout facilities, versioning, hypermedia sound and movies, hypertext links, and knowledge structures represented in a visual language. It can be read as a hypermedia document and also interrogated as a knowledge-based system for problem-solving. The paper version of the document, which you are now reading, is produced by printing the electronic version. It loses its active functionality but continues to act as a record of the knowledge in the document. The overall technology has been developed as an alternative approach to the dissemination of knowledge bases. It also provides a different interface to knowledge-based systems that emulates document interfaces with which many users are already familiar." [11]

The electronic version of the article was made available through anonymous ftp and as a CD-ROM. Thus parallel publication was demonstrated of the paper version as a camera-ready copy book copy and the electronic version, identical in appearance, as a full working demonstration of the problem solution. What is particularly significant is that there were *no* hidden data structures. The semantic networks in the paper were the complete knowledge structures and operated directly in the inference engine to solve the problem. Thus, the document was also the knowledge base.

Figure 7 shows the user interaction with the active document. The knowledge base was represented through semantic networks in the formal visual language KDraw [7]. An interface to use the knowledge base as an expert system was provided through a HyperCard stack that interrogates the document. The stack has no computational capability. It would be equally feasible to provide an interrogation interface within the document. This did not seem appropriate in this application since it would have no particular function in the paper document. That is, it is not proposed that the document be the sole interface to active functionality. It serves the purpose of parallel publication, and beyond that particular requirement there may be alternative modes of interaction more appropriate to other requirements.

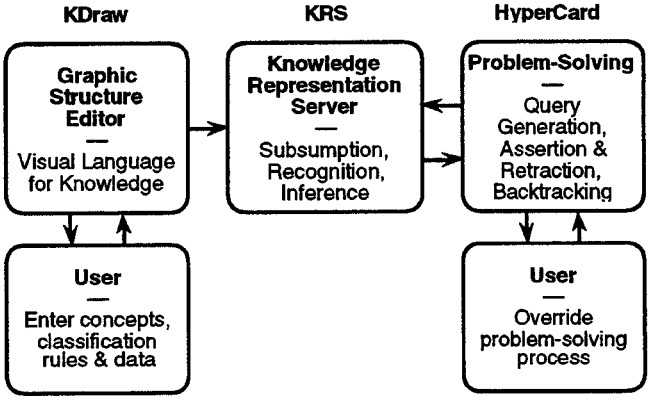

Figure 7 User interaction with the knowledge base

The knowledge representation server KRS, shown at the center of Figure 7 is invisible to the user. It is an inference engine that is loaded with the knowledge structures embedded in the document, and reasons with them when interrogated from the HyperCard stack. This is typical of the use of computational structures within a document for reasoning or simulation. The user interface provides access to the knowledge structures completely represented in some visible form. In this example it is a visible language representing semantic networks that can be interpreted by an inference engine. In another example it could be an arithmetic or algebraic expression that can be interpreted by a symbolic mathematics package such as Mathematica.

Figure 8 is a screen dump of some interaction with the active document. The paper has been opened in KWrite and scrolled to a page containing the top level knowledge structures, the ontology of the domain. These are represented as semantic networks in a visual language with well-defined formal semantics [7], that load directly into a KL-ONE-style deductive engine [6]. These networks may be edited in a separate graphic editor but they may also be edited directly within the document. Double-clicking on one of the nodes brings up the floating dialog box shown in the center of Figure 8 which allows nodes to be edited for type and content, and new nodes to be added. When the cursor is moved to the right side of a node it changes to a menu symbol and mousing down then produces the popup menu shown just below the dialog. This may be used to align nodes and create arrows between them. It may also be used to access an annotation stack in HyperCard through the System 7 IAC.

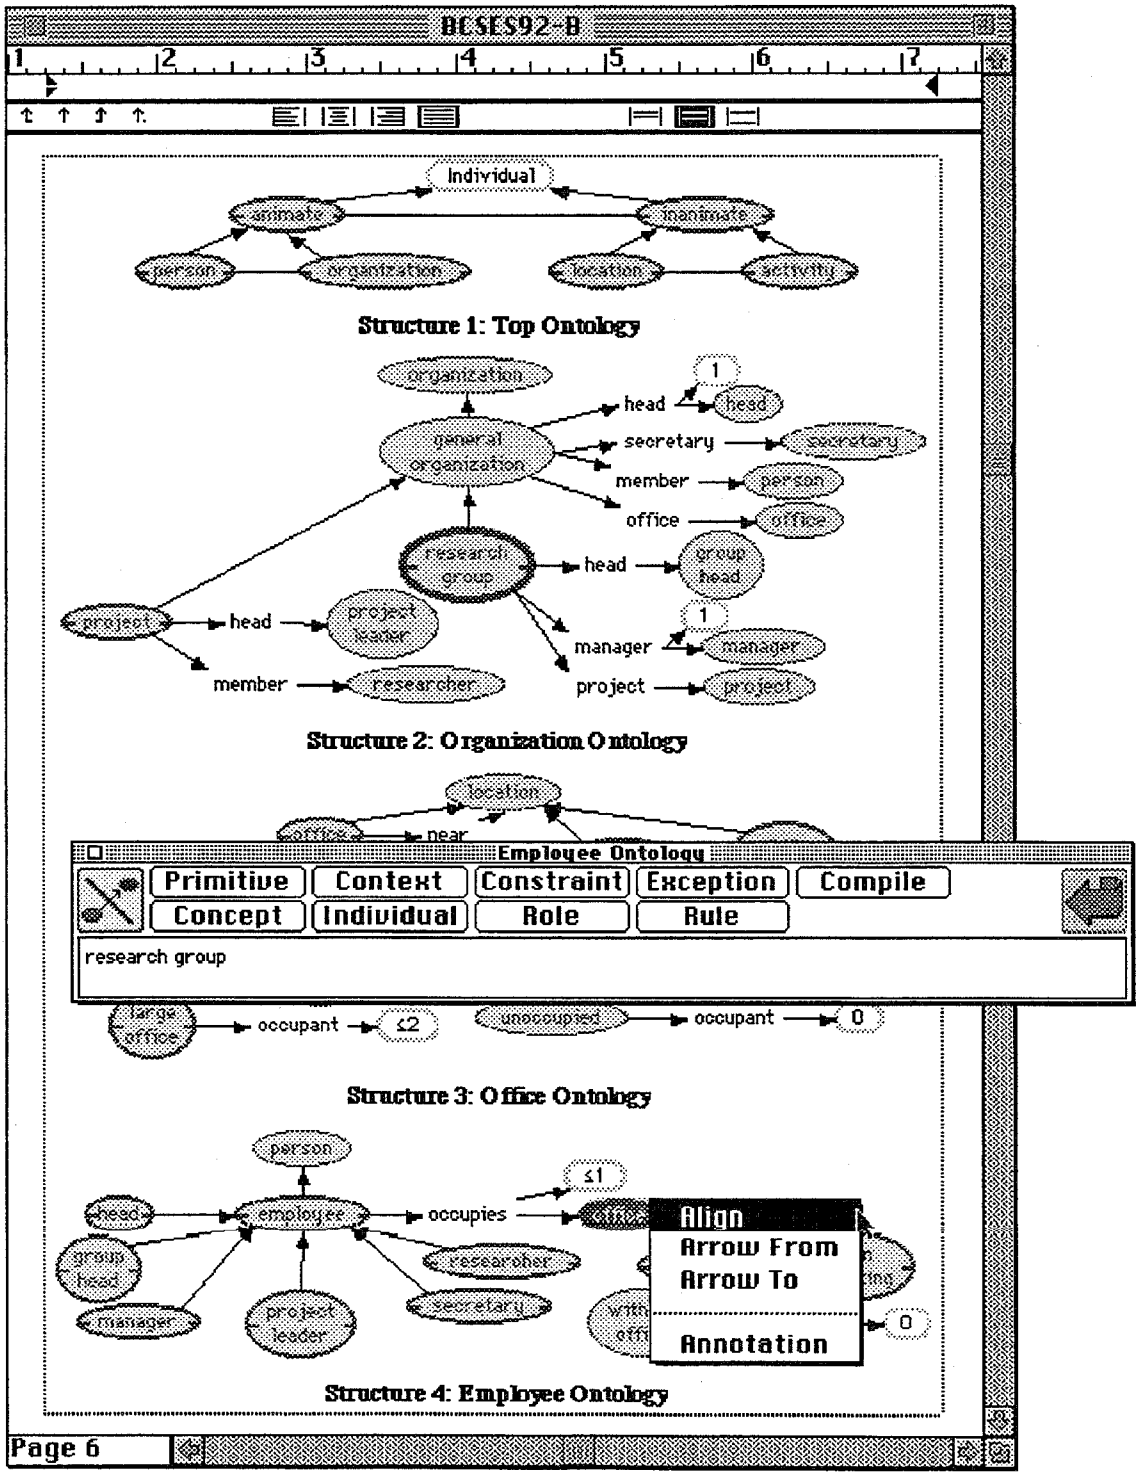

Figure 8 Interaction with knowledge structures in an active document

Project 6: Indexing Document Archives

The projects described so far have involved single documents having compound structures. It is also important to be able to treat complete document archives as a single entity and provide user access to multiple, linked documents using the same interface mechanisms as for individual documents. At the end of the first year of the GNOSIS project described in Project 1, the project reports, slide sequences, movies, and so on, were issued on a CD-ROM. To enable widespread access to the reports, a hybrid format was used that could be read on Macintosh, Windows and Unix platforms, and all the reports were issued in Microsoft Word, Farallon Replica and Adobe PostScript formats. In addition, for Macintosh use the material on the CD was indexed using multi-level linked concept maps accessing reports, slide sequences, movies and software demonstrations.

Figure 9 shows the GNOSIS project archives being accessed through layered concepts maps. The map in the window at the upper left is a top level Mediator "Server Agent" [10] 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.

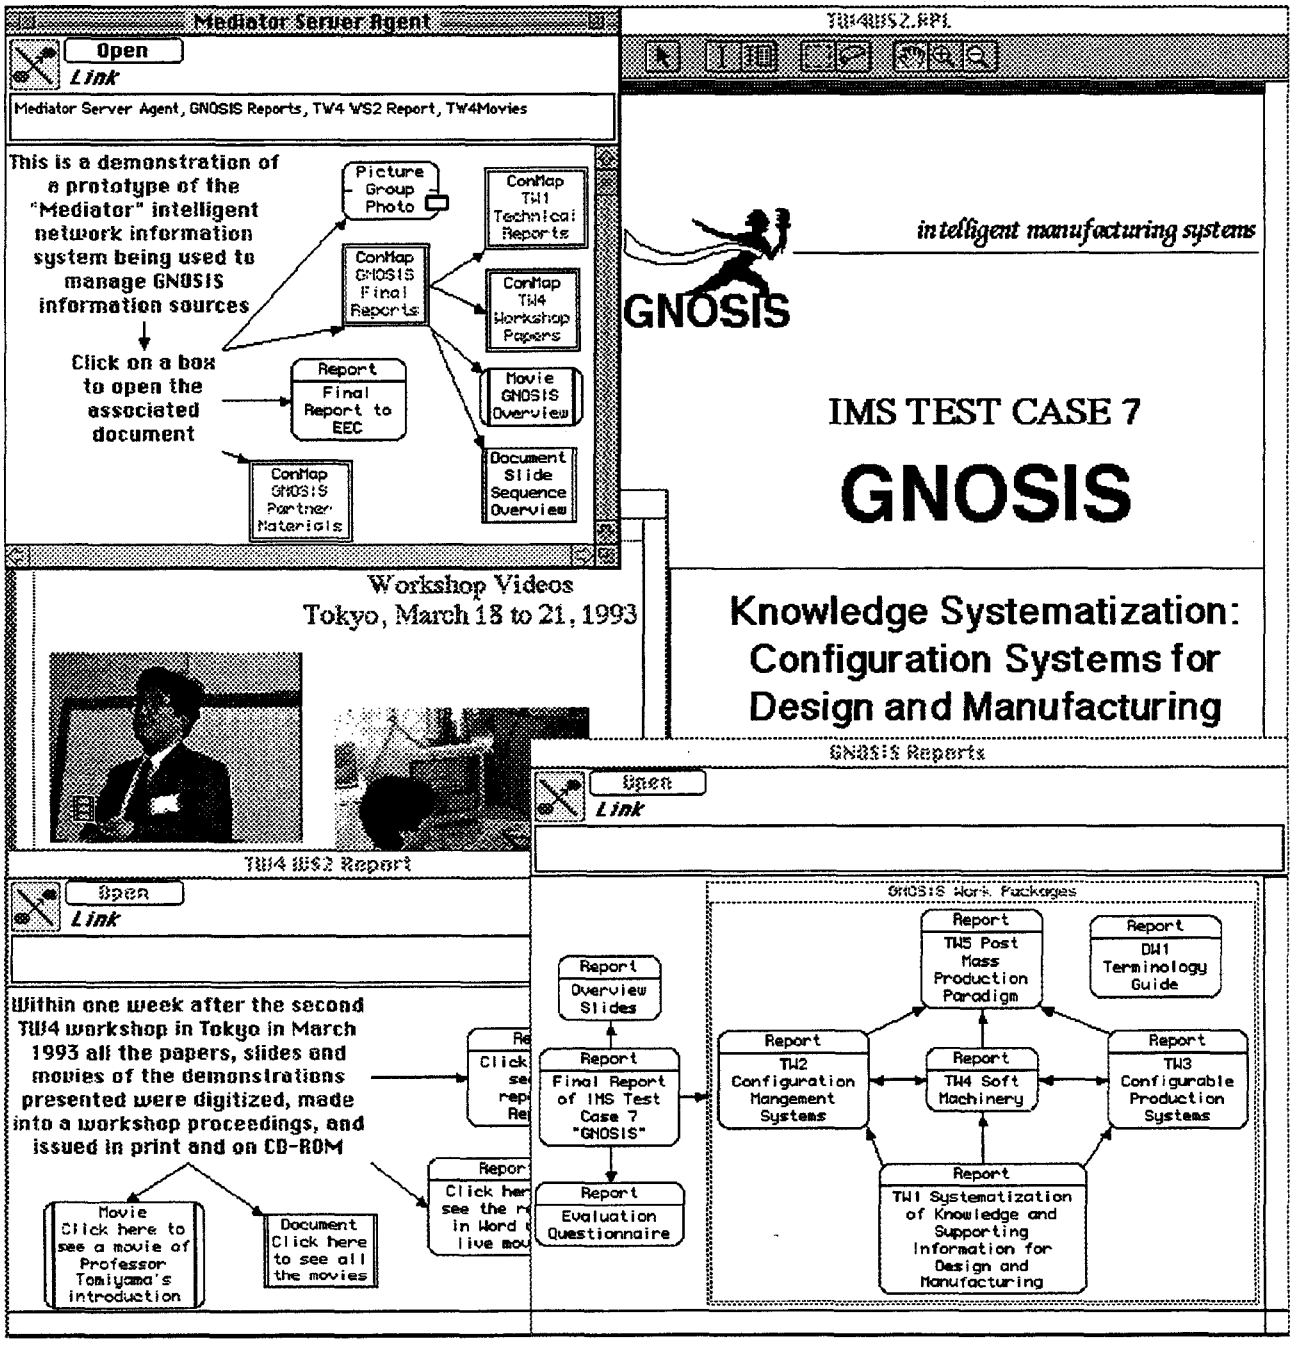

Figure 9 Accessing the GNOSIS archives through layered concept maps

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 9. 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.

The total volume of material on the GNOSIS CD-ROM is: 57 reports totaling 1590 pages, each in 3 formats; 11 movies totaling 70 minutes; plus software and maps of the material. The material from projects 1, 2 and 5 above is included and indexed through the concept maps. In particular, for example, the room allocation expert system described as project 5 may be run from the concept maps. In addition, KWrite documents are provided with embedded concept maps that reference other documents. Thus, the active document technology described is able to achieve seamless integration across bodies of heterogeneous material in different formats managed by different applications. 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.

Production Costs

The total capital cost of a production environment for multimedia documents is remarkably low. A total of US$25,000 was adequate for the hardware and software used in the projects described in this article: a Nikon 801 camera, a Sony TR101 camcorder, an Apple Quadra 900 with two attached 650M drives, SuperMac Video Spigot digitizer, Adobe Premiere and CameraMan video editing and capture software, QuickTopix CD-ROM mastering software, and a Philips CDD521 compact disk recorder. This is low enough for many research groups to have their own facility. If the infrequently used equipment was provided through a central facility, an individual scholar could get into production with an expenditure of some US$5,000.

This does not take into account the one-off cost of developing the document system itself, but such systems should be commercially available from multiple sources in the near future. The figures indicate, as is common nowadays, that labor costs will far exceed capital costs in multimedia document production. The additional effort required to produce a high-quality multimedia document over and above that required for a conventional book or paper is significant.

From past experience in film and laserdisc production it may well be that support from professional media personnel will be required. However, the digital editing capabilities for pictures, sound and video becoming available are already begin to put excellent 'desktop video' capabilities into the hands of authors, and we may see substantial deprofessionalization as has happened with desktop publishing.

One advantage of having control of the document technology is that it avoids licensing fees for multimedia document readers for material issued on CD-ROMs or through the Internet. It is particularly important that the document system technology can be licensed at a sufficiently low cost not to undermine the production economics of CD-ROM. It is irrelevant that a CD-ROM can be produced for a dollar if the presentation technology costs tens or hundreds of dollars.

Conclusions

A major challenge in multimedia research is to develop multimedia document production systems that provide new facilities recognized by significant user communities as emancipating them from the severe limitations of paper-based media. Scholarly research provides a major instances of a communities where the severe limitations of paper-based media have been recognized for many decades, but information technology has yet to have a significant impact on scholarly publishing.

Attempts to develop digital publications in the past twenty years have been largely unsuccessful [8]. Part of the reason for this is that any new technology for publication has to be progressed along its learning curve to a level of maturity supporting major cost reduction before it can compete with the very low costs of print and paper. Multimedia information technology has reached this stage in the 1990s.

However, there are major social factors relating to the reward system in scholarship in general, and how this is supported through publication in high status journals, that have to be taken into account in developing multimedia scholarly publications. These considerations indicate that technology supporting parallel publication in paper and electronic form will have the greatest impact on scholarly publishing in the short term. Again, all the software and hardware technologies necessary to achieve such parallel publication are now available.

What is now necessary to support the process of change from paper to electronic media is to activate those user communities who can make significant gains in achieving their objectives by taking advantage of the capabilities of the new medium. The objective of the research and development program described in this paper is to establish a multimedia production environment as a pilot service that will eventually become a routine facility of the University Communications Media Department.

This article has described experience in establishing a production environment for the parallel publication of documents in paper and electronic form to support the development of multimedia scholarly publication through books and journals. While much of the effort in setting up a production facility has been involved with the technology, a major part has been in analyzing the capabilities and roles of multimedia publication in order that users may understand it in terms of its impact on the argument forms of publications in their disciplines. The article has presented this analysis, described the production environment, and illustrated experience with some initial projects.

Acknowledgments

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada. We are grateful to our colleagues Larry Katz, Mike Mattson and Joan Vickers for many discussions and joint activities over the years involving multimedia development, and to Joan Vickers for the use of her research material in Project 4.

References

[1] Apple, *OpenDoc Developer Overview*, Cupertino, California: Apple Corporation. 1993.

[2] J. Bertin, *Semiology of Graphics*, Madison, Wisconsin: University of Wisconsin Press. 1983.

[3] J.F. Claerbout, *Earth Soundings Analysis: Processing versus Inversion*, Cambridge, Massachusetts: Blackwell. 1991.

[4] J.F. Claerbout and J.A. Dellinger, *Stanford Exploration Project: SEP-CD-1: (1) Earth Soundings Analysis—Processing versus Inversion; (2) Anisotropic Seismic Wave Propagation; (3) Progress Reports SEP-72 and SEP-70*, Palo Alto, California: Stanford University Press. 1991.

[5] D.L. Drucker and M.D. Murie, *QuickTime Handbook*, Carmel, Indiana: Hayden. 1992.

[6] B.R. Gaines, "Empirical investigations of knowledge representation servers: Design issues and applications experience with KRS," *ACM SIGART Bulletin*, vol. 2, no. 3 pp. 45-56, 1991.

[7] 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.

[8] 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.

[9] B.R. Gaines, "Experience with a class library for organizational modeling and problem solving," *Integrated Computer-Aided Engineering*, vol. 1, no. 2 pp. 93-107, 1993.

[10] 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.

[11] B.R. Gaines and M.L.G. Shaw, "Documents as expert systems," in *Research and Development in Expert Systems IX. Proceedings of British Computer Society Expert Systems Conference*, M.A. Bramer and R.W. Milne, Editor. Cambridge University Press: Cambridge, UK. p. 331-349, 1992.

[12] 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.

[13] 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.

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

[15] B.R. Gaines and M.L.G. Shaw, "Supporting the creativity cycle through visual languages," in *AAAI Spring Symposium: AI and Creativity*. AAAI: Menlo Park, California. p. 155-162, 1993.

[16] B.R. Gaines and J.N. Vickers, "Design considerations for hypermedia systems," *Microcomputers for Information Management*, vol. 5, no. 1 pp. 1-27, 1988.

[17] B.R. Gaines and J.N. Vickers, "Hypermedia design," in *Proceedings of RIAO'88 Conference on User-Oriented Content-Based Text and Image Handling*. MIT: Cambridge, Massachusetts. p. 14-23, 1988.

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

[19] D. Goodman, *The Complete AppleScript Handbook*, New York: Random House. 1993.

[20] W. Herzner and E. Hocevar, "CDAM - Compound Document Access and Management. An object-oriented approach," in *Multimedia Systems, Interaction and Applications*, L. Kjelldahl, Editor. Springer: Berlin. p. 17-36, 1992.

[21] D.H. Jonassen, "Semantic net elicitation: tools for structuring hypertext," in *Hypertext: State of the Art*, R. McAleese and C. Green, Editor. Intellect: Oxford, UK. p. 142-152, 1990.

[22] M. Linster, ed. *Sisyphus Working Papers Part 2: Models of Problem Solving*. GMD: Bonn, Germany. 1991.

[23] C. Meghini, F. Rabitti, and C. Thanos, "Conceptual modeling of multimedia documents," *IEEE Computer*, vol. 24, no. 10 pp. 23-30, 1991.

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

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

[26] J.K. Ousterhout, *Tcl and the Tk Toolkit*, Reading, Massachusetts: Adison-Wesley. 1994.

[27] K.R. Popper, "Epistemology without a knowing subject," in *Logic, Methodology and Philosophy of Science III*, B.V. Rootselaar, Editor. North-Holland: Amsterdam. p. 333-373, 1968.

[28] E.R. Tufte, *Envisioning Information*, Cheshire, Connecticut: Graphics Press. 1990.

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

[30] J.N. Vickers and B.R. Gaines, "A comparison of books and hypermedia for knowledge-based sports coaching," *Microcomputers for Information Management*, vol. 5, no. 1 pp. 29-44, 1988.

[31] A. Voß, W. Karbach, U. Drouven, D. Lorek, and R. Schuckey, "Operationalization of a synthetic problem," in *ESPRIT Basic Research Project P3178 REFLECT Task I.2.1 Report*. GMD: Bonn, Germany. 1990.