

# The Learning Web: A System View and an Agent-Oriented Model

Douglas H. Norrie and Brian R. Gaines  
University of Calgary  
Alberta, Canada T2N 1N4

## Abstract

The notion of a learning society has long been promoted as the next stage of evolution of the educational system from the inculcation of skills and knowledge to the development of capacities which enable people to learn continuously for the rest of their lives. This evolutionary step is vital to the continued health of human societies in a post-modern age of rapid and unceasing change where the certainties of the past have been replaced by the uncertainties of a future that requires continuous learning. The move to a learning society requires changes in personal attitudes and the educational infrastructure. It also requires major technological support to provide open access to a learning environment for all people, in all places at all times. This paper presents progress in the development of the Learning Web, an agent-oriented network supporting knowledge access, collaboration and simulation, in order to provide a widely accessible learning environment. The essence of the Learning Web is an open architecture supporting integration of heterogeneous subsystems in which there is a natural symbiosis between human and computer agents.

## 1 Introduction

We are about to see the emergence of what will become a distributed, closely-integrated, intelligent 'learning environment', within which humans will interactively and symbiotically learn from early childhood to the close of life. The term *environment* is used here, initially, rather than *system*, because this learning environment will become part of the accepted surroundings and no longer viewed as an artificial tool, mechanism, or system developed to aid learning. Although not usually perceived as such by the user, this learning environment will, of course, be artificially created, involve numerous 'tools', and will indeed be a *system*.

What makes this 'learning environment' possible is an emerging paradigm which will transform society as radically as did the coming of steam power and the industrial revolution. This is the paradigm of the computer-based *intelligent agent* developing from research in distributed artificial intelligence. By integrating this new approach with recent developments in telecommunications, distributed computing, knowledge databases, language translators, and multimedia systems, a transparent 'learning environment' can be created which will allow the often-foretold *learning society* (Boshier, 1980) to become a reality. This environment can be regarded as a *virtual learning environment*, since the user interacts with visualizations and other simulacra creating an apparent 'physical' reality. The intelligent interconnected network through which the learning environment can be created will here be called the *Learning Web*. This type of intelligent network is already being foreshadowed by the new agent-based 'smart networks' being developed commercially (Reinhardt, 1994).

What will evolve can best be envisaged by extrapolation from the present. Consider the so-called information highway as a starting point. With all its numerous links and connections, this can be seen as an information network, which at its best, allows communication, information and knowledge to transparently flow among its users, reducing the physical distance between them to an inconsequential. In a further metaphorical stage, users see the concept of the network fade with its maplike notions of relatively fixed nodes and interconnections, and the concept of a knowledge web emerge with linking filaments or strands spun across wherever needed. In a later stage, the visibility of the communication linkages disappears and the user interacts with a 'knowledge environment', which responds as appropriate. Requests, for example, lead to responses which may be textual presentations, visual displays (graphics, videos), or spoken explanations. One can envisage the user, at a future time, surrounded by virtual realities that he or she interacts with.

The kind of system outlined above will be used for education, work, recreation, and entertainment. The particular system or aspect of the system which this paper considers is that for education and learning. To avoid a focus too far in the future and too speculative, this paper considers the prospective next-generation stage, namely, a 'web of learning' system. For convenience, this will be called the Learning Web. It will be considered from both the system perspective and component entity point of view, in other words, looked at from the 'top down' and from the 'bottom up'.

## **2 The Learning Web: Structure, Communities, Function**

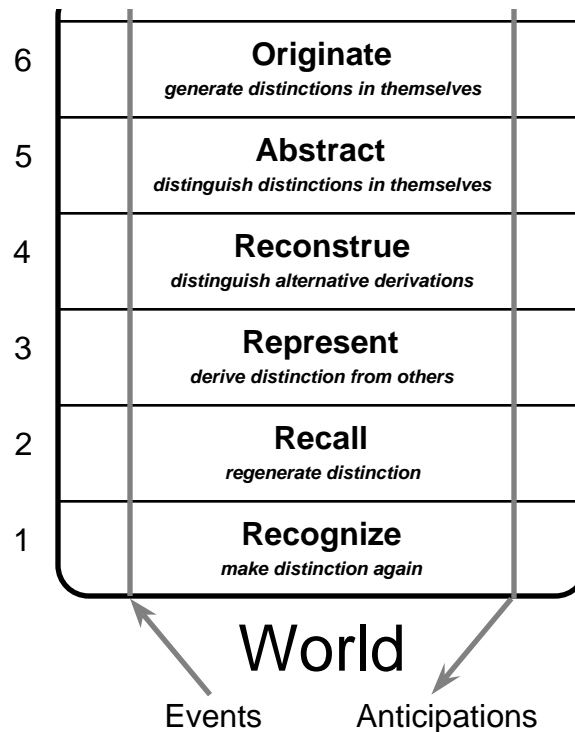
Like all systems, the system being considered is recursively comprised of sub-systems, down to the lowest level of constituency, the component entities. As a 'communications network', it has a base level of physical wires, cables, or EMF transmissions interconnecting the transmitter/receiver nodes. Above this, is the data protocol level, and so on, up to the application level. Each of these levels can be considered as a system, with its component sub-systems and entities. The concern in this paper is with the higher functional and application levels. From this high-level perspective, the system will be considered as an *autopoietic* one and the component entities will be set within an *agent-oriented model*.

At this point, we consider a structural view of the users of the Learning Web. The users constitute a collective or more specifically a society, having the overall purpose of learning about their world and particular aspects of it, both from the heritage of available knowledge and the current inflow of new information. Some of the users, in their exploration and use of available knowledge, will add further contributions for others to use (for example, essays, state of the art reviews, structured data sets).

The society or community of users will be comprised of smaller societal groups or sub-communities, structured according to interest or ability level. For example, there would be groups for history, science, mathematics, or health, with sub-divisions according to sub-fields and ability level. These future patterns in the user community can already be seen in embryonic form in the forums, news groups, and special interest sub-nets of Internet. The World-Wide Web (Vetter, Spell and Ward, 1994) shows the (transparent) cross-linking of knowledge which will characterize the Learning Web. The hypertext, multimedia, and other capabilities (Gaines & Shaw, 1994b) of The World-Wide Web will be early functionalities implemented on the

Learning Web. There are already University courses being taught via Mosaic server over the Internet (Orsak, 1994).

From the users' perspective, the function of the Learning Web is to provide an 'environment' for participatory learning. *Learning in a basic sense is the acquiring of expertise.* This definition, even if thought simplistic, allows a useful transference of knowledge about expertise in individuals, organizations and societies, to the learning sphere. Fundamental to developing expertise are the capabilities shown in Figure 1.



**Fig. 1 Capabilities Fundamental for Learning (Gaines, 1994a)**

These capabilities relate to the structuring, classification, and re-use of knowledge by the individual. Interestingly, organizations can collectively exhibit these capabilities also.

The development of expertise is postulated to be a positive feedback phenomenon, for a society of adaptive agents constituting an autopoietic system (Gaines, 1994a).

### 3 The Learning System as Autopoietic

If one extends the future 'Learning Web' to a 'learning system' by including not only the users but also those who maintain, support, and enhance it, this will be seen to be an autopoietic system.

Autopoietic systems (Maturana, 1981):

*“Are systems that are defined as unities as networks of production of components that (1) recursively, through their interactions, generate and realize the network that produces them; and (2) constitute in the space in which they exist, the boundaries of this network as components that participate in the realization of the network ...”*

Autopoietic systems have components whose recursive activities of self-replication and creative interaction generate the conditions for continued existence. Living systems are autopoietic systems and societal organizations can also be.

The 'learning system' will initially be a controlled artifact of the perhaps loosely-organized grouping of educational, governmental, and business organizations and individuals that establish it. However, at some stage, its component entities will begin the recursive activities of perpetuation and the adaptive behaviours that characterize autopoietic systems. This will begin in the human agents of the system, the individuals and groups that use, develop, and maintain it. Committees and other functional groups will spawn dedicated and specialized sub-groups, which will interact (and sometimes merge) in shifting patterns but always under an overall goal of continued ('persistent') existence. The artificial agents in the system will initially replicate and interact according to preplanned 'intelligent' patterns, but as adaptive mechanisms of increasing sophistication are grafted into the agents a stage can be envisaged when self-adaptive mechanisms begin to be spontaneously generated. At this point, the system will begin to show all of the characteristics of a fully autopoietic system.

#### **4 The Learning System as Agent-Based**

In a convergence of activities from several disciplinary areas, the new disciplinary field of 'multi-agent systems' has recently emerged. This development began in computer science more than a decade ago under the then new heading of 'distributed artificial intelligence' with contributions coming in subsequently from mathematics, the physical sciences, the social sciences (especially psychology), engineering and other areas.

Here, the context for multi-agent systems will be restricted temporarily to computer-based agent systems. Each agent in such a system is a software module capable of sending and receiving messages, and making 'intelligent' decisions and taking actions based on its inherent knowledge and the information in received messages. The inherent knowledge is typically stored in the 'knowledge base' of the agent, structured according to any suitable knowledge representation scheme. Rules, frames, fuzzy logic, and formal logic knowledge bases are commonly used. A recent survey of agent architectures and languages provides further details (Wooldridge and Jennings, 1994). An earlier survey with a more applied focus provides also some developmental background (Norrie and Kwok, 1992). Of particular relevance to this paper, agent-based Wide Area Networks are already being developed (Reinhardt, 1994).

A stimulus to the development of multi-agent systems was provided by the object-oriented programming paradigm, in which objects are software modules which send and receive messages and take action (typically the sending of a message) according to 'hard-coded' programmed instructions. By providing each object with a knowledge base and the capability to infer or reason from this knowledge base, the object-oriented paradigm becomes extended to the agent-oriented one (Shoham, 1993; Kwok and Norrie, 1992). Object-oriented languages and architectures can thus be extended into agent-oriented ones (Kwok and Norrie, 1994a).

The architecture or structure of a multi-agent system or sub-system - essentially defines the allowable interactions (e.g. who can communicate with whom, send what information, request or control co-operation). This structure can vary from the hierarchical to the heterarchical, with the degree of control varying from highly authoritarian to 'free will' cooperation.

The restriction on agents being only software modules will now be removed, since the learning system being considered will be constituted from both human and artificial (computer-based) agents. The structure and organization of this multi-agent 'learning system' will, initially, be based on the patterns found useful in human societies and organization. Differentiation and specialization of function are characteristics of human systems and similar patterns will also be found in the learning system. For example, in some cases, an agent will control or manage a resource (e.g. a database), and in others, an agent will direct or coordinate problem-solving behaviour.

### ***Agents and Resources in the Learning System***

The following are some of the agent and resource types that will exist in the learning system:

#### ***Knowledge Agent***

An 'expert' in a particular area of knowledge. This can be either a human or an artificial agent. The knowledge agent can not only answer questions and provide information but also will be able to direct a user to other relevant knowledge agents and knowledge servers.

#### ***Knowledge Server***

An artificial agent with major capabilities for storing and retrieving knowledge, and for answering queries and providing information by inferring or reasoning using the stored knowledge bases. It can be expected that the various knowledge bases will use different knowledge representation schemes, so it will be important that the server be able to convert knowledge in one type of representation to another. Thus a query coming in, framed as a formal logic statement (e.g. a Prolog query), could be converted to equivalent rule form to initiate inferencing across a rule-type knowledge base, with the 'answer' being converted back to logic form and transmitted as the reply. The knowledge server would also often be providing, as required, even very large chunks or sets of knowledge as knowledge bases in any required representation to knowledge agents and other servers for subsequent knowledge processing locally. Knowledge servers have already been developed with the above kinds of functionality. Gaines (1994b) has detailed the principles, architecture and object-oriented implementation of knowledge servers with such capabilities.

#### ***Interface Agent***

An agent associated with the user interface which monitors and learns from the user's actions and then functions as an intelligent assistant. A very simple illustration would be the completion of reformatting a document or set of data, by the agent, as soon as it had 'learnt' from the user's manual reformatting of the first portion. Recent research (Maulsby, 1994; Roesler and Hawkins, 1994) has shown that considerable capability can be obtained for such agents, without excessive complexity.

#### ***Coach or Tutor Agent***

'Scaffolding' is an educational technique to support a learner who is learning a new task (Unsworth, 1993; Kirschner, 1991). As the learner is working through the task, structure and guidance is provided by the coach or tutor to assist in the successful completion of the task. Suggestions, additional information, and correction of mistakes are some of the supporting 'scaffolding' provided. As the learner acquires the knowledge and skill to carry through the task,

the scaffolding progressively is removed. There is considerable evidence for the effectiveness of both human and ‘intelligent tutoring system’ coaching.

In word processing and other software packages, the form of coach described as a ‘Wizard’ is becoming already quite common (Ditlea, S., 1994).

### ***Mediator Agent***

An agent which acts to coordinate the activities of other agents and to resolve conflicts between them. Some of the activities requiring coordination within distributed agent systems are planning, task building and contracting, negotiation, task allocation, assessment of solutions. A monograph by Durfee (1988) describes a number of approaches to the coordination of distributed problem solvers. Approaches to coordination and conflict in multi-agent systems are also considered by de Mazeau and Muller (1990). Coordination of planning in multi-agent systems is considered in detail by von Martial (1992).

Until recently, most of the agent systems developed in industry used a ‘centralized single-level organization of agents built around a powerful controller agent’ (Adler, Durfee, Huhns, et al, 1992). This simple approach to coordination through a powerful mediator is now giving way to more sophisticated approaches. At The University of Calgary, research is in progress on a generic and recursive approach to creating mediator agents for sub-problem levels. These agents both coordinate their sub-tasks and coordinate their activities with other mediators, and dissolve when their work is done.

It needs noting that there is, in the literature another sense in which mediator agents are defined: “--- a software module that exploits encoded knowledge about certain sets or sub-sets of data to create information for a higher layer of application” (Wiederhold, 1992). This usage would include interface agents as mediators and is quite different to the usage in this paper.

### ***Knowledge Management Agent***

The user will interact with the Learning Web through a variety of interfaces. In the early stages of the Learning Web, there will be graphics, text, visual languages and hypermedia, with more sophisticated vertical reality interfaces being later incorporated. Through these interfaces they will access various ‘packages of functionality’ or ‘applications’ which will allow desired functions to be carried out such as directory access; design; graphics creation; knowledge searches; reasoning on specified knowledge bases; visualization (not only as graphics but as other spatial representations); group work of several kinds. These functions involve the creation, assembly, manipulation and interpretation of knowledge. A Knowledge Management Agent provides the high-level coordination of these knowledge activities within either an individual or collective ‘project’. It is a particular kind of high-level Mediator Agent for project coordination of human and computer-based agents within a collaborative framework. A Knowledge Management Agent can also be considered as a system and is exemplified by The Mediator described in Section 5.

### ***Information Search Agent***

This is an agent which is dispatched into a system to search out specific information and to return back with it. The agent ‘travels’ through the system by transmitting itself from location to location, checking at each location or node for the desired information. An architecture for a set of such information retrieval agents was proposed by Kahn and Cerf (1988), who called these

agents knowbots. Agents of this kind are also known as infobots and have already been sent roaming around portions of the Internet on test missions (Woelk & Tomlinson, 1994). A Telescript agent can also move around an appropriate network on assigned missions (Wagner, 1994).

### ***Directory Agent***

As the name implies, a directory agent ‘points’ (supplies the address of) to an appropriate agent, service, or resource when queried ‘Where is XXX?’. Since data, information, and knowledge will be extensively distributed across the system in a variety of data-, information-, and knowledge-bases, effective directory agents will be essential. It can be expected that over time there will be duplicated, outdated, or contradictory elements in these bases and mechanisms for resolving the consistency problem are already being developed (Haggith, 1994).

### ***Mentor Agent***

In business or the professions, a mentor is a more experienced person who acts as a guide in areas such as higher level expertise, strategy, and career development. Here, a mentor agent is envisaged as acting in a rather analogous way in the learning environment, as a kind of coach for the higher level strategies of learning.

### ***Other Agents and Resources***

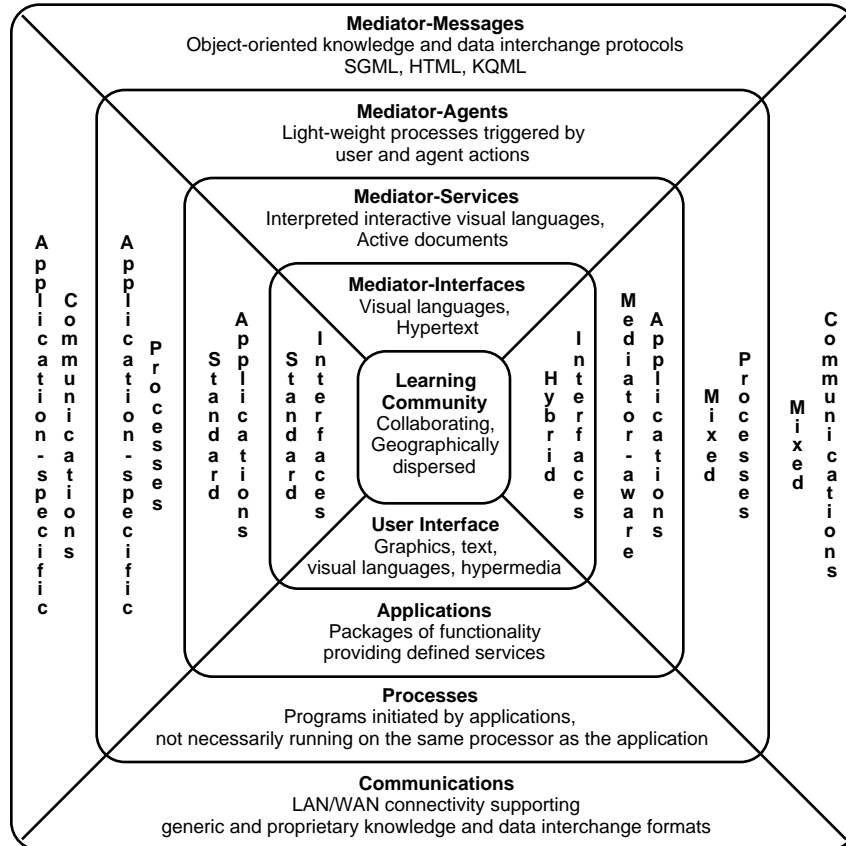
The above listing is intended to be exemplary rather than comprehensive. In its formative years, the Learning Web will also contain many agents and resources providing familiar functions, e.g. citation and document retrieval, relational data bases, dictionaries, atlases, and citation data bases. Roesler and Hawkins (1994) provide a summary of agents developed to aid users in a number of common tasks.

## **5 Intelligent Knowledge Management in the Learning Web**

The Learning Web will be *a distributed system of intelligent autonomous entities dynamically collaborating*. Its implementation will require a distributed intelligent knowledge management system. Such a system is being developed at The University of Calgary. Although the original focus for this work was a different application area, this knowledge management system has the basic functionality needed for an initial Learning Web prototype. The system is known as The Mediator (Gaines and Norrie, 1994). Mediator is a system designed to coordinate an overall application project over the whole life cycle from requirements to completion. In particular, it can support the integration of computer-based sub-systems within a mutually collaborative framework. It can be used as the top-level coordination framework for an initial Learning Web prototype. It provides facilities for recording and tracing decisions taken at each stage of the project life cycle, particularly the dependencies between knowledge, decisions, datasets, and so on. Within the Learning Web, there could effectively exist a ‘network’ of Mediators, each of which acts as the major knowledge management agent for a sub-web of activities (e.g. a collective learning project).

As shown in Figure 2, the Mediator architecture is a four-layer distributed client, distributed server design, in which multiple users can collaborate synchronously or asynchronously through processes running anywhere on the network. A heterogeneous environment is assumed in which there are multiple protocols and multiple forms of user interface. The collaborating and

geographically dispersed user community interacts with the system through a variety of forms of user interface, typically graphics, text, visual languages and hypermedia. They access a variety of 'applications' defined as packages of functionality providing defined services. The functionality is made operational by initiating processes which may run anywhere on the network; that is, remote procedure calls are expected to be common. Inter-process communications is provided through local and wide area networks supporting a range of generic and proprietary knowledge and data interchange formats.



**Fig. 2 Mediator Architecture Supporting the Learning Web**

On the left of Figure 2, the four layers are shown instantiated in terms of conventional software packages designed independently of Mediator: as standard interfaces, standard applications, application-specific processes and application-specific communications. Such standard applications can play a role in a Mediator-coordinated systems, minimally by the application and its datasets being registered in Mediator with its application windows open on the Mediator desktop, and maximally by Mediator controlling its inputs, outputs and operations by job-control scripts. On the right, the four layers are shown instantiated in terms of separately designed applications that are 'Mediator-aware' to some extent, for example in using the Mediator interface technology, applications, agents or protocols as part of their normal operation. The Mediator shell technology is being designed to be highly modular and readily integrated in whole or in part with existing applications.

The current Mediator implementation uses an open architecture visual language tool (Gaines & Shaw, 1994a) for representing conceptual schema in a way that is comprehensible to users, and can be easily tailored to different application domains. Mediator supports semantic networks that



can be used to represent rules, procedures and constraints with formal, operational semantics such that an inference system can be used to give advice and check for constraint violations. The representation is domain independent and the system can reason with legal constraints and corporate procedures as well as with design constraints. Less formal concept maps are also supported where the role of the system is primarily information retrieval rather than reasoning, and the formal and informal systems can be combined.

A semantic network visual language for Mediator is already available (Gaines, 1991a) that compiles into a KL-ONE-like knowledge representation system (Gaines, 1991b; Gaines, 1993a) and may be used to represent computational knowledge structures (Gaines, 1993b). This has been used to operationalize a corporation's procedures manuals (Kremer, 1991) in a way that is directly applicable, for example, to educational project management. Procedures can be triggered through the user interface to the visual language, and the procedures can include other conventional or agent-based applications or an 'agent' system that supports active integration between applications, machines, and between sites. Data translation and knowledge interchange are supported through such agents.

The primary user interface to Mediator is through an open architecture visual language system developed to 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 effect the type structure in node decorations such as fonts, shading and color. User interaction with Mediator 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.

Figure 3 is a screen dump from a user interacting with the current implementation of Mediator accessing material relating to an international manufacturing collaborative research project. The user has activated a Mediator "Client Agent" and connected to a "Server Agent" managing an archive of documents, photographs, digital movies, programs and agents. She is first shown the concept map at the top left which provides an overview of the archive. It 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*. The node at the top left gives access to a series of slides on the project displayed using *Replica*. A similar node in the original concept map at the top left gives access to a movie on the project that will be opened in Apple's *MoviePlayer*.

The screenshot displays the Mediator Server Agent interface. At the top, there's a title bar 'Mediator Server Agent' and a toolbar with 'Open' and 'Link' buttons. Below this, a text box reads: 'This is a demonstration of a prototype of the "Mediator" intelligent network information system being used to manage GNOSIS information sources. Click on a box to open the associated document.' This text points to a concept map with nodes like 'Picture Group Photo', 'ConMap TH1 Technical Reports', 'ConMap TH4 Workshop Papers', 'Report Final Report to EEC', 'Movie Gnosis Overview', 'Document Slide Sequence Overview', and 'ConMap Gnosis Partner Materials'. To the right, a window titled 'TW4WS2.RPL' shows the Gnosis logo and the text 'in intelligent manufacturing systems', 'IMS TEST CASE 7', and 'GNOSIS'. Below this, another window titled 'GNOSIS Reports' shows a concept map with nodes like 'Report Overview Slides', 'Report Final Report of IMS Test Case 7 "GNOSIS"', 'Report Evaluation Questionnaire', 'Report TH2 Configuration Management Systems', 'Report TH4 Soft Machinery', 'Report TH1 Systematization of Knowledge and Supporting Information for Design and Manufacturing', 'Report TH5 Post Mass Production Paradigm', 'Report DW1 Terminology Guide', and 'Report TH3 Configurable Production Systems'. A video player window titled 'Workshop Videos Tokyo, March 18 to 21, 1993' is also visible, showing a man speaking. A 'TW4 WS2 Report' window at the bottom left contains text: 'Within one week after the second TW4 workshop in Tokyo in March 1993 all the papers, slides and movies of the demonstrations presented were digitized, made into a workshop proceedings, and issued in print and on CD-ROM.' It points to nodes like 'Movie Click here to see a movie of Professor Tomiyama's introduction' and 'Document Click here to see all the movies'.

**Fig. 3 Accessing Learning Materials through Mediator Agents and Layered Concept Maps**

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 5.1. She has also clicked on the “Click here to see all the movies” node and opened the document visible behind the concept maps. This displays eight QuickTime movies of various demonstrations given at a project workshop, any of which can be played by double clicking on it.

## 6 Agent Development Systems

For creating intelligent agent systems and applications, an Intelligent Agent Development System (Kwok and Norrie, 1993a, 1993b, 1994b) has been created based on Intelligent Agent Objects (IAO's). This can be used to develop agent sub-systems operating under Mediator. The

following describes this existing IAO software development environment (i.e. language, editor, browsers, inspectors, graphical tools, etc.). The basic entities created in this system are intelligent agents. Each such agent can contain not only data and procedural methods but also a knowledge base and can moreover initiate an inferencing process operating on its own knowledge base and those of other agents. The following knowledge representation schemes and associated inferencing are currently implemented: forward, backward, and hybrid production-rules; confidence-factor rules; fuzzy logic. Formal logic, and other representations can be subsequently incorporated without changing the basic architecture. A unique and special feature is that there are no conventional style inference engines in the system; *all inferencing is carried out by a newly developed message inferencing technique*. Other key features are: incorporates fully the object-oriented concepts of abstraction, encapsulation, inheritance, and polymorphism; provides propagation-oriented, access-oriented, and group-oriented programming paradigms; inheritance applies to knowledge-base rules; variables can be extended to being compound structures such as organization charts or design cases and rules can operate on these; provides seamless integration of procedural and inferencing execution processing.

This new IAO approach for inferencing uses a coordination pattern of cogent messages between interacting intelligent agents. Under the guidance of a Messenger-Manager, a knowledge base (KB) rule in an intelligent agent object triggers the creation of a messengerObject to trace the needed premise or conclusion in whichever other agent KB rule it may be. If the premise or conclusion is compound, more than one messengerObject is needed and as many as necessary are created. Each successful link found is recorded so that the full chain can be subsequently presented or used for tracing. As this linking (chaining) propagates (and backtracks if necessary) so the inferencing is accomplished. Different MessengerManagers control different types of inferencing, searching and reasoning. The IAO System is a large and comprehensive development system for creating intelligent agent applications and represents an advance in architecture for this area.

As noted previously, one of the procedures linked through the Mediator user interface to the visual language can be an agent system which supports active integration between applications, agents and other sites. Such a Mediator Agent System (MAS) is under development at The University of Calgary. Agent applications within the Learning Web can be integrated under the Mediator through a Mediator Agent System of this kind. One particular application would be an active 'mediating system' which 'knows' the organizational structure of a learning community and which handles efficiently the communications traffic across and within this sub-web. It is possible, in principle, to incorporate organizational knowledge within an agent network and to efficiently provide the needed communications facilities through this network. Such an agent network can provide the 'mediating system' described above, by having an agent associated with each active entity of the sub-web. Each such agent can provide intelligent 'mediating' functions as required, for example, from the simpler functions such as message protocol translation, through to more sophisticated functions such as evaluating the knowledge transfer to and from modules to recommend or assist in organizational restructuring. In later incarnations, a Mediator Agent System which learns and reasons in very sophisticated ways could provide higher level functionalities. The importance of both individual and multi-agent learning has already been noted (Huhns, 1994). Experimental work on techniques for self-adaptive changing of an agent's decision strategies has been reported (Grefenstette, 1992). It is not unreasonable, therefore, to envisage a Mediator Agent System with sophisticated learning capabilities.

## 7 Conclusions

Recent developments in distributed artificial intelligence and the convergence of new technologies from telecommunications, distributed computing, multimedia, and databases now make possible a Learning Web of diverse but interconnected educational and learning facilities. From a system point of view the Learning Web can be viewed as an evolving autopoietic system. From an architectural perspective, the Learning Web can be represented by a Multi-Agent model. The Mediator system being developed at The University of Calgary could function as the high-level information management system for the Learning Web. Applications and sub-systems can be developed using an agent development system such as the Intelligent Agent Object Development System.

In making the case for a learning society, Boshier (1980) remarks:

*“The major purpose of education cannot remain the inculcation of knowledge and skills, but should become the development of intellectual and psychological capacities which enable people to learn continuously for the rest of their lives.”*

One must add to this list, the need for a technological capacities to provide access to learning resources to all people, at any place and at any time, and the need for the human and machine capacities to be integrated symbiotically in a unified learning environment. The Learning Web is necessary to make the learning society a reality.

## Acknowledgements

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada.

## References

- Boshier, R., Towards a Learning Society. Vancouver: Learningpress, 1980.
- de Mazeau, Y. and Muller, J.P., “Decentralized AI”, North Holland, 1990.
- Ditlea, S., “Silent Partners”, PC Computing, pp. 160-171, 1994.
- Durfee, E.H., “Coordination of Distributed Problem Solvers”, Kluwer Academic Publishers, Boston, 1988.
- Gaines, B.R., “An Interactive Visual Language for Term Subsumption Visual Languages”, IJCAI ‘91, Proceedings of the Twelfth International Joint Conference on Artificial Intelligence, San Mateo, California, Morgan Kaufmann, pp. 817-823, 1991a.
- Gaines, B.R., “Empirical Investigations of Knowledge Representation Servers: Design Issues and Applications Experience with KRS”, AC SIGART Bulletin, 2(3), pp. 45-56, 1991b.
- Gaines, B.R., “A Class Library Implementation of a Principled Open Architecture Knowledge Representation Server with Plug-in Data Types”, IJCAI ‘93, Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence, San Mateo, California, Morgan Kaufmann, pp. 504-509, 1993a.
- Gaines, B.R., “Experience With a Class Library for Organizational Modeling and Problem Solving”, Integrated Computer-Aided Engineering, 1(2), pp. 93-107, 1993b.

- Gaines, B.R., "The Collective Stance in Modelling Expertise in Individuals and Organizations", *Int. J. of Expert Systems*, Vol. 7, No. 1, pp. 19-49, 1994a.
- Gaines, B.R. "Class library implementation of an open architecture knowledge support system." *International Journal Human-Computer Studies* 41(1-2) 59-107, 1994b
- Gaines, B.R., Norrie, D.H., "Mediator: Information and Knowledge Management for the Virtual Factory", *Proc. of SIGMAN AAAI 94 Workshop on 'Reasoning About the Shop Floor'*, 1994.
- Gaines, B.R. and Shaw, M.L.G. "Concept maps indexing multimedia knowledge bases." *AAAI-94 Workshop: Indexing and Reuse in Multimedia Systems*. Menlo Park, California, AAAI, 1994a.
- Gaines, B.R. & Shaw, M.L.G, "Interacting with the Web", *CHI'95*, submitted, 1994b.
- Grefenstette, J.J., "The Evolution of Strategies for Multiagent Environments", *Adaptive Behaviour*, Vol. 1, No. 1, pp. 65-90, 1992.
- Haggith, M., "A Meta-Level Approach to Exploration of Multiple Knowledge Bases", *Draft Proceedings of the Second International Working Conference on Cooperating Knowledge Based Systems*", Vol. 2, 14-17 June 1994, University of Keele, pp. 313-335, 1994.
- Huhns, M., "A DAI Perspective on Co-Operating Knowledge Based Systems", *Draft Proceedings of the Second International Working Conference on Cooperating Knowledge Based Systems*", Vol. 1, 14-17 June 1994, University of Keele, pp. 1-8, 1994.
- Kahn, R.E., Cerf, V.G., "An Open Architecture for a Digital Library System and a Plan for its Development", *Corporation for National Research Initiatives Technical Report*, March 1988.
- Kirschner, D.M., "Using Verbal Scaffolding to Facilitate Conversational Participation and Language Acquisition in Children With Pervasive Developmental Disorders", *Journal of Childhood Communication Disorders*, Vol. 14, No. 1, pp. 81-98, 1991.
- Kremer, R.C., "Experience in Applying KRS to an Actual Business Problem", Boose, J.H. and Gaines, B.R., Ed., *Proceedings of the Sixth AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop*, University of Calgary, Calgary, Canada, pp. 11-1 - 11-2, 1991.
- Kwok, A.D., Norrie, D.H., "IAO: Multi-Paradigm Programming Environment", *Proc. of Int. Conf. on Object-Oriented Manufacturing Systems*", 3-6 May 1992, Calgary, pp. 219-226, 1992.
- Kwok, A.D., Norrie, D.H., "Integrating Multiple Reasoning in Intelligent Agent System", *Integrated Computer-Aided Engineering*, 1(2), pp. 83-90, 1993a.
- Kwok, A.D., Norrie, D.H., "Intelligent Agent Systems for Manufacturing Applications", *Journal of Intelligent Manufacturing*, pp. 285-293, 1993b.
- Kwok, A.D., Norrie, D.H. "Integrating a Rule-Based Object System Within the Smalltalk Environment", *Journal of Object-Oriented Programming*, Vol. 5, No. 9, pp. 48-55, 1994a.
- Kwok, A.D., Norrie, D.H., "A Development System for Intelligent Agent Software", To be published in *International Journal of Integrated Manufacturing Systems*, 1994b.
- Maturana, H.R., "Autopoesis", in M. Zeleny (Ed.) "Autopoesis: A Theory of Living Organization", North-Holland, pp. 21-33, 1981.
- Maulsby, D., "Instructible Agents", *Ph.D. Thesis*, The University of Calgary, 1994.

- Norrie, D.H., Kwok, A.D., "Object-Oriented Distributed Artificial Intelligence", Proc. of Int. Symp. on New Results and New Trends in Computer Science, Graz, Austria, 20-21 June 1994, LCNS 555, Springer-Verlag, Berlin, pp. 225-242, 1992.
- Orsak, G.C., "Teaching Signal Processing on the Information Superhighway", Proc. of the 1994 Sixth IEEE Signal Process Workshop, Yosemite, California, pp. 131-134, October 2-5, 1994.
- Reinhardt, A., "The Network With Smarts", Byte, Vol. 19, No. 10, pp. 51-64, Oct. 1994.
- Roesler, M., Hawkins, D.T., "Intelligent Agents: Software Servants for an Electronic Information World (And More)", Online, V. 18, No. 4, pp. 18-32, July 1994.
- Shoham, Y., "Agent-Oriented Programming", Artificial Intelligence, No. 60, pp. 51-92, 1993.
- Unsworth, L., "Multiple Semiotic Sources as Scaffolding for Young Children's Emergent Reading of Picture-Story Books", Australian Review of Applied Linguistics, Vol. 16, No. 2, pp. 1-14, 1993.
- Vetter, R.J., Spell, C., Ward, C., "Mosaic and the World-Wide Web", Computer, IEEE, pp. 49-57, October 1994.
- von Martial, F., "Coordinating Plans of Autonomous Agents", Springer-Verlag, Berlin, 1992.
- Wagner, P., "Agents Away", Byte, Vol. 19, No. 5, pp. 113-118, May 1994.
- Wiederhold, G., "Mediators in the Architecture of Future Information Systems", Computer, IEEE, pp. 38-49, 1992.
- Woelk, D. & Tomlinson, C., "The InfoSleuth Project: Intelligent Search Management via Semantic Agents", Proc. of World-Wide Web Conference, Chicago, October, 1994. (<http://www.ncsa.uiuc.edu/SDG/IT94/Proceedings/Searching/woelk/woelk.html>)
- Wooldridge, M.J., Jennings, N.R., "Agent Theories, Architectures and Languages: A Survey", Proc. of ECAI 94 Workshop on Agent Theories, Architectures, and Languages (ed. Wooldridge, M.J. and Jennings, N.R.), Amsterdam, Netherlands, pp. 1-32, 1994.