A Sample of Knowledge Management Projects at George Washington University

Carl Moore, Seung Ik Baek, and Jay Liebowitz*
Dept. Of Management Science
George Washington University
Washington, D.C. 20052
Tel: 202-994-0554
Fax: 202-994-4930
JAYL@GWIS2.CIRC.GWU.EDU
*RWD Technologies Distinguished Professor, University of Maryland-Baltimore County, as of Aug. 17, 1998

ABSTRACT

Artificial intelligence (AI) has been looking for a "killer application" since its inception. Now, with the world wide web, intranets, intelligent agents, and the emerging knowledge management trends, AI may have found its much needed killer application. Through the use of web-based intelligent agent technology, knowledge management systems can be developed in very effective ways. This paper will highlight several of the knowledge management projects that have been developed in the School of Business and Public Management at George Washington University (GWU) which incorporate the use of intelligent agents for making AI an integral part of knowledge management.

1.0 An Intelligent-Agent Based Information Warfare Advisor Developed for the US Army War College (USAWC)

The capstone, and perhaps most innovative, activity at the US Army College is the "Strategic Crisis Exercise (SCE). The SCE is a two week exercise/simulation which serves as the laboratory for applying the students' knowledge in strategic leadership, military doctrine, policy, and decision making gained over the student's tenure at the War College. The SCE involves all the War College students (320+) role playing certain organizations, (e.g., State Department, Joint Chiefs of Staff, Treasury, Military Commands, etc.) whereby a myriad of activities occur in which the students need to react and formulate decisions. Typically, over 120 subject matter experts are used throughout the SCE to bring some measure of reality to this exercise.

One of the themes that is integrated throughout the SCE is the influence of information warfare (IW). In the SCE, certain civilian and defense-related IW events occur, and the students need to consult the "IW subject matter expert" in order to help resolve some of these issues. To help in this regard, an on-line IW advisor, using intelligent agent technology, was built by George Washington University for USAWC student use during the SCE. Instead of bringing in many experts from around the country to participate in the SCE, the use of intelligent agents for on-line advisors may be a cost-effective and efficient technique for allowing the students to interact, pose questions, and engage in discussions with these "surrogate experts" over the SCE intranet.

1.1 System Goals

In exploratory discussions held at the US Army War College with experts in the field of information warfare (IW) and its use in exercises held at College, the key goals for the Information Warfare Advisor (IWA) are:

- Provide an on-call advisor in IW for students at the USAWC taking courses or participating in real-time exercises. This advisor would draw on both public and private knowledge bases about IW and IW related topics to provide students with requested information on IW.

- Provide a monitoring advisor which would analyze email and other message traffic collected during the AWC annual Strategic Crisis Exercise (SCE) and issue a real-time alert of IW activity or instances detected in

the email or message threads. This alert would be part of the SCE simulations.

- Provide an assessment advisor which would analyze email threads and other message traffic collected during the USAWC annual SCE and develop a post exercise set of cases on the IW activity during the SCE. These cases would be used to de-brief students after the exercise, design future exercises, as well as add to the IWA knowledge base on IW.

The initial research project focused only on developing a portion of the IWA as a proof of concept demonstrator. Initially, a limited scope advisor agent system was developed, providing the user callable advisor capability for students at the USAWC. The on-call advisor advises students on a limited number of topics via the user agent, a single broker agent, and one or two specialist agents in selected IW topics.

The computing environment in which the IWA operates is a heterogeneous network of server and client computers. The servers are a mix of Sun workstations and Windows NT servers used as mail, file, and web servers. The clients are end-user workstations consisting of mostly Intel 486 and Pentium class desktop or laptop PCS. All servers and clients are networked using 10 Mbps Ethernet and the TCP/IP protocol. The Relational Database Management System (RDBMS) installed is the Oracle system. The network is set-up as a private intranet. There is no external access to the Internet at this time, although such access is held out as a future potential and must be accounted for in the architecture for expansion.

The IW knowledge bases which the IWA uses are a combination of existing legacy knowledge bases and databases, as well as new knowledge bases created as part of this research effort and subsequent development efforts for the IWA. These knowledge bases and databases are distributed throughout the computing environment, residing on different servers.

1.2 General Approach

The distributed nature of the knowledge bases in the IWA environment led to the decision to base the IWA conceptual architecture on a distributed AI approach using a collection of cooperating intelligent agents. This approach offers both several advantages, as well as certain disadvantages. However it was felt the advantages of this approach, in lieu of developing a single monolithic expert system advisor, more than outweighed the disadvantages.

The multi-agent approach is a more complex design, and consequently was more difficult to implement than the monolithic approach. However the multi-agent approach is better suited to the distributed environment for the IWA. A multi-agent system allows better shared use of the computing resources available on the network for load leveling, and avoids the potential bottleneck of a single point of access for the advisor. The multi-agent system is also easily extended by adding agents with different specialties. In this manner the capability of the overall IWA system can be increased by re-using the developed agent architectures over and over, but each time with a different knowledge base.

1.3 Conceptual Architecture--Brokered Agents

The broker agents in this architecture act as generalists in their field. Each broker maintains knowledge of specialist agents in the environment which had problem solution capabilities within the broker's general domain. Other agents seeking assistance with a problem then need only to contact the brokers to receive information on what specific agents would be of assistance, and then contact those agents only, rather than placing a universal query. Optionally, the system could be devised so that the problem could be divided into larger pieces by the originating agent and passed to the brokers. The brokers would then handle the job of further breaking down the portion of the problem within their own domain, using domain specific knowledge about problem solving, contact the agents within their brokerage for solutions, synthesize their portions and hand off to the originating agent. The originating agent would then synthesize the partial answers from a few brokers rather than from a larger number of specialists. This second variant of the broker process, using the brokers as domain generalist for

sub-problem planning and solution through subcontracting to specialists is the approach used for the IWA. The broker concept of a cooperative system of agents is called an agency.

The conceptual process of problem solution in a brokered agency system begins with a user presenting a problem to the system for solution. The presentation of the problem is via a user agent. The user agent consists of both a user interface as well as the underlying logic and reasoning ability to break the problem into portions and pass those portions to a broker or brokers for solution. This passing of problem portions is termed subcontracting. Under the terms of the subcontract the user agent would request assistance, and the broker would accept the problem. In the event no broker was able to solve the problem portion, or no agent was available, then no subcontract would be made and the user is notified of the failure to perform.

If the broker can solve the problem, then the broker would break the subcontract into portions for solution by the specialist agents within its domain. Each specialist agent could proceed independently, or in conjunction with other agents, to solve the problem. There is no restriction on specialist agents possessing the same ability as the user agent to subcontract parts of its assigned problem to other agents through a broker, either the same or a different broker.

The specialist agents also have access to both private databases and knowledge bases of their own, as well as public databases and knowledge bases accessible by all specialist agents. These public databases and knowledge bases may be legacy databases made available to the population of agents in a standard format wrapper agent. The job of the wrapper agent is to act as a translator for the system from the unique format and structure of the legacy database to the general structure and format used by the agent system.

The brokered agency conceptual architecture has advantages and disadvantages. Maintenance and expansion are more difficult since new agents must be somehow registered with one or more broker agents so the brokers know they exist and what their capabilities are. In like manner, new broker agents must also be made known to the user agent. This registration can be made manual at the time of adding or updating agents, increasing the load on maintenance programmers, or can be made automatic through logic included in the design of the agent architecture. However this increases the load on the agent developers, and indirectly on the maintenance programmers since now the agent designs would be more complex and hence more costly to maintain.

However, there are a number of key advantages offered by the brokered agency architecture that outweigh the disadvantages. With a brokerage cooperative agreements, or sub-contracts, communication among the agents is easier to establish. The brokerage concept reduces sharply the number of other agents a user agent must be aware of in order to process a problem. Each broker acts as a generalist in its domain, providing a source of information on the specific agents which can be of assistance, or actually taking on the sub-task itself for solution. Brokers can also call upon other brokers, not just the specialist agents within its domain, thereby widening its field of expertise, as well as providing the system with a hierarchy of domain knowledge. Some brokers would have more general knowledge than others, each level of broker breaking the problem into sub-tasks for subcontracting to other more specific brokers or to specialist agents. This structure provides the system with a large degree of modularity, making expansion easier to plan, execute, test, and manage.

The advantages offered by the brokered agency concept were sufficient to warrant its adoption as the conceptual architecture of choice for the IWA project.

2.0 A Web-Based, Intelligent Agent Approach for Managing Knowledge Relating to Storyboard Development for Multimedia Designers in a Virtual Setting

Another GWU knowledge management project that used AI/intelligent agent technology involved web-based access for jointly developing a storyboard for a multimedia effort. In this context, knowledge management is a method for systematically and actively managing and

leveraging ideas and design decisions among team members while developing storyboards. This research involved developing a knowledge management system to help multimedia team designers create, exchange, and share their storyboards (the storyboard is the script/thread/flow of activities relating to a multimedia project).

There were three intelligent agents used in this knowledge management architecture. These agents were created using Javascript, HTML, and CFML (Cold Fusion Markup Language). The knowledge creating activity was handled by the user agent which remembered all user activities, and dynamically organized a person's agenda. The knowledge securing activity was handled by the knowledge agent to index knowledge, detect inconsistencies and generate recommendations, and save/retrieve/update knowledge. The knowledge distributing and retrieving activities were performed by the knowledge manager to monitor all changes that occurred in a knowledge repository and forward them to the user agent, reformulate queries based on an ontology, and dynamically retrieve annotations and generate hyperlinks for them.

These three agents have well-defined goals and knowledge to perform their tasks (goal-oriented). They perform their tasks without direct human intervention (autonomous and self-starting). The agents interact with each other to perform their tasks (cooperative). The agents interact with a user database, annotation database, and storyboard database.

Through this architecture, the multimedia designers used in the research studies could effectively and efficiently develop their storyboards in a collaborative fashion via the web in a virtual setting.

3.0 Summary

In order for knowledge management to succeed, AI technologies (like intelligent agents) are critical to the main knowledge management functions of creating, securing, capturing, combining, distributing, and retrieving knowledge. The two projects briefly described in this paper highlight some of the work in which AI can greatly contribute to knowledge management systems and activities. At GWU, we are furthering our work in this area by also examining performance and learning among intelligent agents for developing improved knowledge management systems. Lastly, we are looking at developing a methodology for valuating intellectual capital (particularly, human capital) in an organization.

Selected References

Liebowitz, J. (ed.), The Handbook of Applied Expert Systems, CRC Press, Boca Raton, 1998.

Liebowitz, J. and T. Beckman, Knowledge Organizations: What Every Manager Needs to Know, St. Lucie Press/CRC Press, Boca Raton, 1998.

Liebowitz, J. and L. Wilcox (eds.), Knowledge Management and Its Integrative Elements, CRC Press, Boca Raton, 1997.