INNIS LIBRARY
NON-CIRCULATING

McMaster University
MICHAEL G. DeGROOTE SCHOOL OF BUSINESS

RESEARCH AND WORKING PAPER SERIES

A METHODOLOGY FOR THE ASSESSMENT OF BUDDY-AGENTS

By

Xiaoqing Li and **Ali R. Montazemi**

Michael G. DeGroote School of Business
McMaster University
Hamilton, Ontario

Working Paper # 452

March, 2001

Innis
[redacted].
HB
74.5
.R47
no.452

TE
n Stre
, Ontario, Canada L8S 4M4
ie (905) 525-9140

A Methodology for the Assessment of Buddy-Agents

Xiaoqing Li and Ali R. Montazemi
Michael G. DeGroote School of Business
McMaster University
Hamilton, ON
Canada L8S 4M4
Email: montazem@mcmaster.ca & lix2@mcmaster.ca

A Methodology for the Assessment of Buddy-Agents

Abstract

Computer-based information systems connected to high-speed communication networks provide increasingly rapid access to a wide variety of data resources. However, this connectivity to data resources burdens decision-makers the need to access and analyze a large volume of data to support their decision making processes. Without effective decisional guidance, access to data resources provides only a minor benefit to decision-makers. Intelligent agents are expected to act like human-assistants in support of complex decision processes by anticipating the information requirements of the decision-makers or by autonomously performing a specific set of tasks. In this article, we provide a methodology for assessment of buddy-agents in a multi-agent information system environment in support of complex decision problems. Our findings from an empirical assessment of the methodology that was used to support common stocks selection among investors support the viability of the proposed methodology.

1. Introduction

Information systems increasingly are used to help deal with complex decision problems to the extent that a user-friendly human-computer interface has become crucial to their success. Intelligent agent, an artificial intelligence (AI) technique, is developed in support of human-computer interface to implement a complementary style of interaction (Lieberman, 1997; Maes, 1994; Minsky, 1994 and 2000; Montazemi and Gupta, 1997). An intelligent agent is expected to reduce the complexity of dialogue by understanding user goals and assisting their interaction with the system (Lewis, 1998; Montazemi and Gupta, 1996; Pilkington, 1992). We can identify application of agent technology in such diverse areas as: information retrieval systems to help retrieve relevant documents (Lewis, 1998; Montazemi and Gupta, 1996), and in electronic commerce to help buy and sell products and services (Maes et al., 1999).

The objective of this research is to present a model of how a group of agents can plan their interaction, and use learning techniques to improve their operational performance in support of such complex decision making processes as selecting music or a common stock. For example, we may have to sift through hundreds of music CDs to find one we have not heard and matches our taste. A knowledgeable assistant in a music store may help us narrow our search by asking about our musical taste (e.g., artist, composer, type, rhythm and price). As for selecting a common stock, many financial institutions (e.g., www.ETRADE.COM) enable investors to purchase common stocks on-line. However, the onus is on the investors to have complete knowledge of thousands of common stocks traded in different exchange centers (e.g., The New York Stock

Exchange). This creates information overload, renders the on-line information market somewhat inefficient, and sets the stage for the emergence of information "intermediaries" in the market. In either case, we are faced with a significantly large database of possible selections that may satisfy our needs. In addition, the selection criteria are complex and usually not Boolean. In these situations, the information system must act intelligently, taking into account the knowledge of the user and of the decision environment to provide the user with information at the right level of detail.

Intelligent agents are expected to embody some of the key capabilities of a human assistant: observing and forming models of the decision environment; inferring the decision-maker's intentions based on these observations; and formulating plans and taking action in support of the decision-maker's decision making processes. This is a challenging task but there have been major breakthroughs in the design and development of agent-based information systems in support of complex decision problems (Bordetsky and Mark, 2000; March et al., 2000; Montazemi, 1999).

There are almost as many opinions on the definition of agents as there are agents themselves. The diversity of definitions can be attributed to the range of applications that can use this technology to enhance decision-making processes. In this paper, agent and intelligent agent are used interchangeably. Section 2 of this paper elaborates on the characteristics of agents. Agents have to interact with each other as well as with environmental entities (e.g., human decision-makers and databases) to achieve their goals. One of the basic problems facing designers of multi-agent systems for open and complex information environments such as the Internet is that of connection: finding buddy-agents who might have the needed information or other capabilities. To this end, Section 3 describes the characteristics of multi-agent systems. The objective of this research and proposed methodology to assess buddy-agent membership is described in Section 4. Section 5 describes our empirical investigation to test the effectiveness of the proposed methodology; Section 6 presents the results of our empirical investigation. Concluding remarks closes the paper.

2. Characteristics of Intelligent Agents

Intelligent agents are expected to work in open and complex information environments. The central idea underlying software agents is that of delegation. The user delegates a task to the agent and the agent autonomously performs that task on behalf of the user. Researchers have described the characteristics and classified agents in numerous ways. For example, Hayes-Roth (1995) views intelligent agents as having the capability to perform three necessary functions:

- To perceive dynamic conditions in the environment.
- To take action to affect conditions in the environment.
- To reason in order to interpret perceptions, solve problems, draw inferences, and determine actions.

Nwana (1996) provides a typology that defines different types of agents based on their abilities to cooperate, learn, and act autonomously. Autonomy refers to the principle that agents can operate on their own without the need for human guidance. With cooperation capability, agents can interact with each other and possibly with humans via some communication language media. The key attribute of any intelligent being is its ability to learn. Intelligent agents have to *learn* as they react and/or interact with their surrounding operational environment.

3. Characteristics of Multi-Agent Systems (MAS)

With the development of agent technology, the need for a system of multiple agents communicating in a peer-to-peer fashion is becoming apparent (Sycara et al., 1996) . Characteristics of MASs are as follows (Sycara, 1998):

- Each agent has incomplete information or ability to solve the decision problem.
- There is no global system control.
- Data are decentralized.
- Computation is asynchronous.

MAS is expected to solve problems that are beyond the individual capabilities of a single agent. The motivations for the increasing interest in MAS research include the ability of MAS to do the following:

- Solve problems that are too large for a centralized agent to solve.
- Allow for the interconnection and interoperation of multiple existing legacy systems.
- Provide solutions to problems that naturally can be regarded as a society of autonomous interacting component-agents.
- Provide solutions that efficiently use information sources that are spatially distributed (e.g., information gathering from the Internet).
- Provide solutions in situations where expertise is distributed (e.g., health care, stock market).

To communicate effectively, each agent in a multi-agent system needs to know the characteristics of the other agents that best can serve its requirements. A popular model to facilitate communication among agents is by means of a middle-agent (also called "matchmaker" or "broker"). All agents register with the middle-agent. For example, let us assume that agent A1 has a request (i.e., ASK(X)). To perform this request, the middle agent (M) can use one of the following two procedures.

(1) Recommend performative: A1 ask M to "recommend" an agent to whom it would be appropriate to send the performative ASK(X). Once M learns that A2 is willing to accept ASK(X) performatives, it replies to A1 with the name of agent A2. A1 is then free to initiate a dialog with A2 to answer this and similar queries (See Figure 1) (Finin et al., 1997).

Figure 1

The recommend performative is used to ask the facilitator agent to respond with the "name" of another agent who is appropriate to send a particular performative

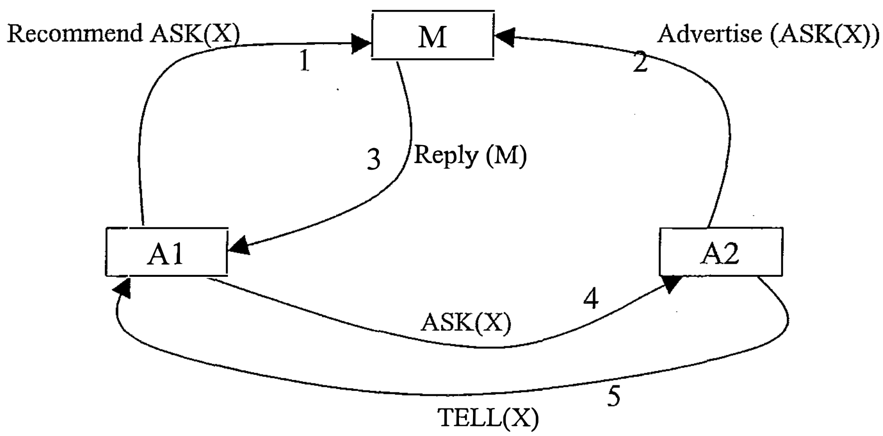

(2) Broker performative: A1 asks M to find an agent that can process an ASK(X) performative. A2 independently informs M that it is willing to accept performatives matching ASK(X). Once F has both of these messages, it sends A2 the query, gets a response and forwards it to A1 (See Figure 2) (Finin et al., 1994).

Figure 2

The broker performative is used to ask a facilitator agent to find an agent who can process a given performative and forward the reply

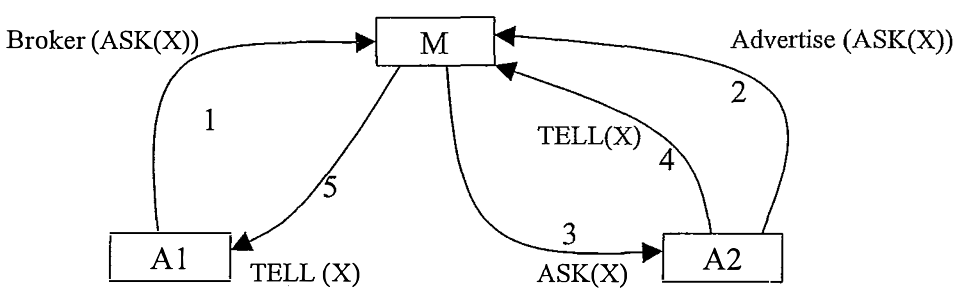

The above methodology is applicable when the characteristic of the decision environment represented by each agent is simple, and the number of agents in need of cooperation is small. This structure becomes difficult to effectively apply in complex decision environments such as common stock selection.

In the stock selection decision environment, investors are distributed globally. This community is not controlled centrally, and each investor independently chooses his/her portfolio. Each investor has a local database storing a personal stock portfolio representing his/her personal knowledge and judgment of selected stocks. In addition, investors tend to share their knowledge to improve the quality of their decision processes in making portfolio selection (e.g., see investor community at www.ETRADE.COM). Thus, the stock selection decision environment can be regarded as a society of autonomous investors that tend to share knowledge to improve decision-making performance. Let us assume that there are N agents (serving N investors), and that agent A1 wants to know if others could recommend a stock similar to IBM. Thus, A1 sends a message to other agents. Other Agents (A2 ... An) search their own portfolio (case-base) and select similar stocks. These selections are sent back to agent A1. After receiving all the responses, A1 is in a position to assess them and select those that best match IBM stock. This is an acceptable process as long as there is no cost involved in sending, receiving, and processing data. Let us assume that A1 would like to send its message to agents who are most likely to give a good response within the shortest time. We call these agents "buddy-agents of A1". In the next section, we will develop a methodology for assessing buddy-agent membership.

4. Objective of This Research

In this research, we propose a methodology to assess buddy-agents in a distributed multi-agent system in support of complex ill-structured decision problems. Our basic premise is the desire to understand and share knowledge among decision-makers (Bordetsky and Mark, 2000). This can be as simple as asking the address of a restaurant that is famous for its special Chinese seafood in Metropolitan of Toronto. Another example could be tapping into a review of books such as those offered by Amazon.com. An equally complex task is sharing knowledge of our stock portfolio. Internet sites such as www.ETRADE.COM try to support this sharing through a chat group facility.

Our basic assumption is that a message sent by agent A1 to find stocks similar to IBM is best answered by agents of those investors whose portfolio (case-base) is "more" similar to the portfolio represented by agent A1. The objective of this research is to identify agents (buddies) who can best respond to the request of another agent. This is based on the assumption that in order to solve a new problem, one should first try using methods similar to those that have worked on similar problems. This is called "reinforcement learning" model (Minsky, 1995).

Reinforcement learning occurs when some aspects of the behavior of a system are caused to become more (or less) prominent in the future as a consequence of the application of a "reinforcement operator" Z (See Figure 3). In response to a stimulus from the environment, the "reinforcement model" chooses one of several possible responses. It remembers what decisions were made in selecting this response. Shortly thereafter, the Trainer (decision-maker) sends the model positive or negative reinforcement signals.

Figure 3

Reinforcement Learning Model

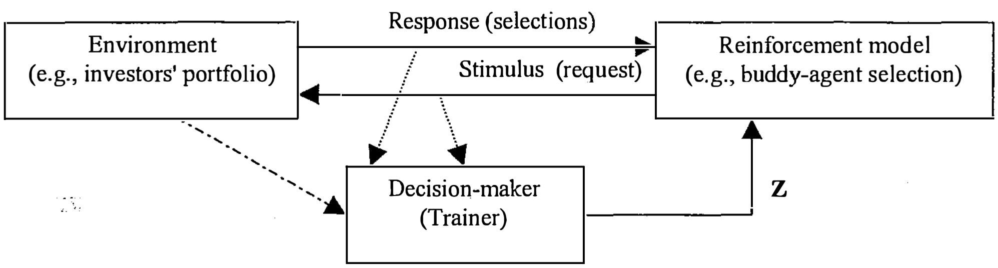

Based on the reinforcement learning model, in response to a request from the decision-maker D1, agent A1 sends messages to a number of other agents seeking to satisfy the request from D1. Next, the responses of other agents are presented to decision-maker D1. The decision-maker Dl sends agent A1 his/her degree of satisfaction (i.e., positive or negative reinforcement signal) for each response received. The agent (reinforcement operator) does not initiate behavior, but merely selects what the decision-maker likes from what has occurred.

Our proposed framework used to assess the degree of membership of buddy-agents is based on fuzzy set modeling. The objective is to select buddy-agents that are expected to meet a set of criteria in responding to a request. Let us assume that the two criteria related to a stock selection request are as follows:

(1) Response time (T): Length of time it takes for each agent to respond to a request. Therefore, an agent tends to select buddies that respond quickly to its requests (i.e., T close to 0).

(2) Response quality (Q): The quality of the response (recommendation) received. This is the degree of match between the requested stock and the recommendations offered by an agent. We use a range 0-1 where 1 indicates a perfect match and 0 represents no match at all. Thus, the objective is to select agents as buddy-agents with Q close to 1.

We use a variation of Yager fuzzy intersection (Cox, 1999) to assess the value of goal attainment by each agent as follows:

- Identify goal attainment μ by agent x for each criteria using the following formula:

$$\mu\ (x) = \frac{1}{\left(1 + (x - a)^2\right)}$$

is used to calculate the membership function of the vicinity $\boldsymbol{x}$ to its desired limit $\boldsymbol{a}$. Where $\boldsymbol{x}$ is a possible buddy-agent and $\boldsymbol{a}$ is the value of each criteria. Thus, the goal attainment for T (in which $\boldsymbol{a}$ is desired to take the lower limit of zero) and Q (in which $\boldsymbol{a}$ is desired to take the upper limit of 1) are derived as follows:

$$\mu\ (T) = \frac{1}{\left(1 + T^2\right)}$$

$$\mu\ (Q) = \frac{1}{\left(1 + (Q - 1)^2\right)}$$

Goal attainment j for all the agents is computed as follows:

$$(G_t(x_i))^{wt} = \{(x_1, \mu(t_1)), (x_2, \mu(t_2)), (x_3, \mu(t_3))\}^{wt}$$

$$(G_q(x_i))^{wq} = \{(x_1, \mu(q_1)), (x_2, \mu(q_2)), (x_3, \mu(q_3))\}^{wq}$$

Where wt and wq are the weights assigned by the decision-maker as to the significance of buddy-agents' response time and quality of response. For example, a decision-maker may assign timeliness of response as wt=2 and quality of response as less significant with a value of wq=1.2. The final membership value (D) for each agent is computed by the intersection of all the criteria that they should attain as follows:

$$D = \{[x_i, \min_j (G_j(x_i)^{wj})] \text{ where } i = 1,\ldots,n; j = t, q\}$$

For example, if

$$\left(\tilde{G}_t(x_i)\right)^2 = \{(x_1, 0.7^2), (x_2, 0.5^2), (x_3, 0.4^2)\}$$
$$= \{(x_1, 0.44), (x_2, 0.2), (x_3, 0.12)\}$$

$$\left(\tilde{G}_q(x_i)\right)^{1.2} = \{(x_1, 0.3^{1.2}), (x_2, 0.8^{1.2}), (x_3, 0.6^{1.2})\}$$
$$= \{(x_1, 0.24), (x_2, 0.76), (x_3, 0.54)\}$$

Then

$$D = \{(x_1, 0.24), (x_2, 0.2), (x_3, 0.12)\}$$

This indicates the degree of membership for $x_1 > x_2 > x_3$. Thus, based on the number of possible agents, we can limit the choice to the top n selection as the buddy-agent.

The goal attainment μ for all goals having little importance (Wt, Wq <1) becomes larger, and those with objectives having more importance (Wt, Wq >1) become smaller. This has the effect of making the membership function of the decision subset D, which is the minimum value of each X over all objectives, being determined by the most important objectives.

5. Empirical Evaluation of the Methodology

We performed an empirical test to determine the effectiveness of the proposed methodology for the design and development of buddy-agent membership in support of complex decision environments. Section 5.1 describes the decision environment; Section 5.2 describes the tool developed for use in our investigation; and Section 5.3 presents the experimental design used to assess the effectiveness of our proposed methodology.

5.1 Decision Environment

We developed a multi-agent system (MAS) to assist investors receiving (and providing) advice about stock market securities. The decision environment faced by an investor is highly ill-structured, so much so that security prices are posited to follow the random walk hypothesis, which stated that at any point in time the size and direction of the next price change is random with respect to the state of knowledge available at that point in time (Dyckman et al., 1975). The major cause of this random behavior is caused by:

(1) the large number of causal variables,

(2) the fact that variables are highly stochastic, and

(3) the unknown significance of causal relationships among the variables.

Consequently, in such an extremely complex and rapidly changing environment, the forecast of security prices is expected to rely heavily on the analyst's (investor's) cognitive efforts (e.g., intuition, training). Therefore, one would expect that the pertinent decision processes involved should not only vary among different investors, but also should be context dependent (Simon and Hayes, 1976; Tversky and Kahneman, 1982). In addition, due to the presence of a significantly large number of stocks, an investor can cover only a small subset of stocks. This makes it highly desirable for investors to share their knowledge of specific stocks (See www.ETRADE.COM investor community). This decision environment resembles the society of minds as hypothesized by Minsky (1994), suitable for support by MAS because (1) decisions are distributed, (2) each decision-maker is autonomous, and (3) decision-makers need to share their knowledge to improve their own decision performance.

5.2 Tool

Using the framework of methodologies proposed in this research (see Appendix 1) we have developed a multi-agent system (MAS). Our MAS performs the following functions: (1) enables selection of buddy-agents, (2) broadcasts the requirements of an investor for new stocks to other investors' agents, (3) facilitates local comparison of stocks for selection of most similar stocks by means of distributed case-based reasoning systems (CBR), and (4) ranks and presents stock information received from other agents. The procedures used in the development of the CBR systems was adopted from previous research (Gupta and Montazemi, 1996; Montazemi & Gupta,1997). AGENTBUILDER software was used in support of communication protocol among agents. An overview of the proposed system is depicted in Figure 4.

Figure 4

An Overview of the MAS System in Support of Stock Selection

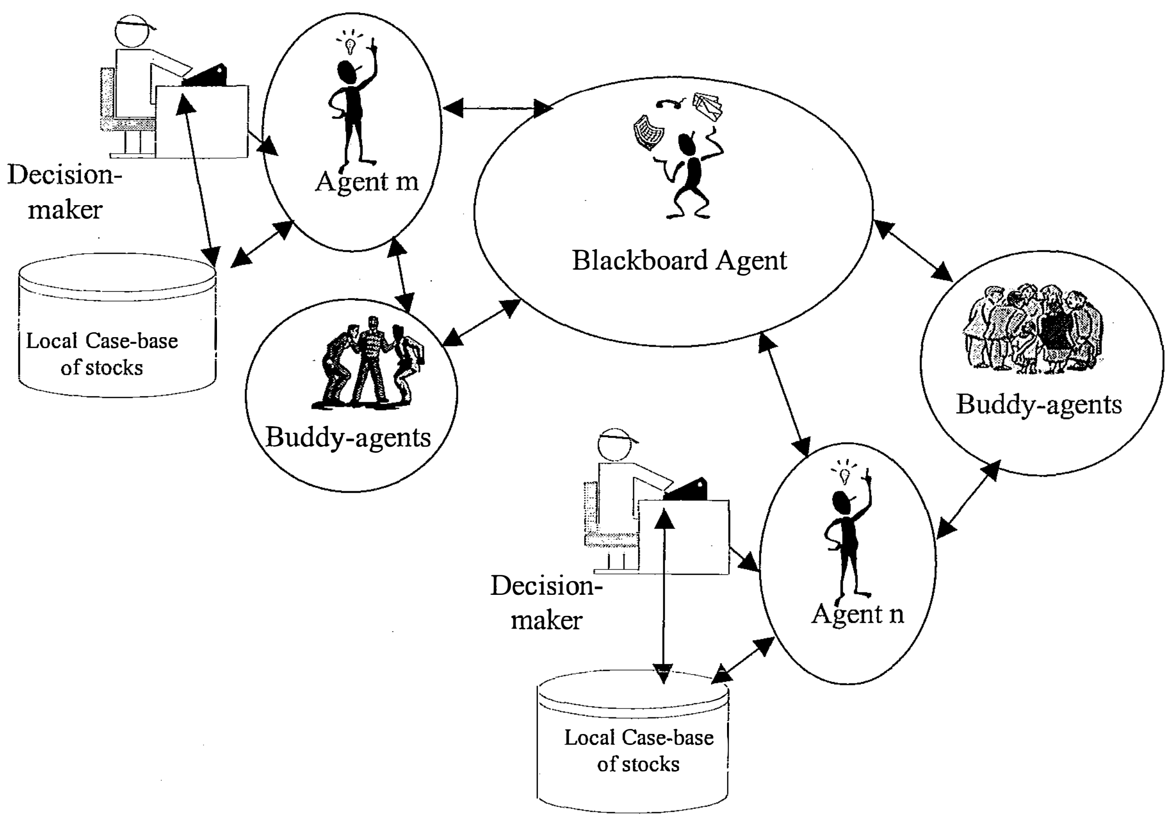

5.3 Experimental Design

The effectiveness of the methodologies used for selecting the degree of membership of the buddy-agents was assessed as follows:

I. Information about 5000 stocks was collected. Each stock was represented by the 17 financial attributes that are generally used to select stocks (See Table 1). Next, cluster analysis was used to identify groups of stocks with similar characteristics. Each cluster represented the portfolio of an investor. There were eight major clusters. The stocks in each cluster were entered in the case-base of a CBR system. Thus, we created eight case-bases representing eight investors.

Table 1

Attributes used to represent selected stocks

Abbreviation	Definition
Ask	Latest asking price.
Bid	Latest bid price.
Dividend	Annual dividend payment representing either the latest fiscal year or indicated annual rate based on the most recently announced dividend payments.
Earning	Annual earnings per share representing either the latest fiscal year or indicated annual rate based on the most recent published earnings.
Last	Last trade price or value.
Net change	Difference between latest trading price or value, and the historic closing value or settlement price.
Opening Price	Today's opening price or value.
PE Ratio	Ratio of stock price to earnings per share.
Percent change	Percentage changes from the latest trade price or value from the historic close.
Today's high	Today's highest transaction value.
Today's low	Today's lowest transaction value.
Year high	Highest value of the year.
Year low	Lowest value of the year.
Yield	For equities, dividend per share expressed as a percentage of the price.
Last	Last trade price or value.
Historic close	Most recent non-zero closing value or settlement price.
Trade volume	Transactional volume of the trade price reported in the Last field.

II. Cluster analysis is a multivariate procedure for detecting groupings in the data. Cluster analysis begins with no knowledge of group membership and often without knowing just how many clusters there are. The distance between clusters (e.g., portfolio of stocks) indicates how far apart clusters are from each other.

This can be taken as the closeness of portfolios to each other or the degree of membership of one portfolio with others. Thus, we can use distance between clusters as a good benchmark for assessing the goodness of our proposed model of agent membership assignment.

III. The multi-agents system was used to assess the degree of membership of agents. This was achieved by having each of the eight agents request recommendations about 48 stocks similar to those in its case-base. The quality of the response received from agents was used to compute the degree of membership of remaining agents (i.e., seven buddies) for each agent.

IV. Degree of membership of the buddy-agents for each agent should be similar to the distance between clusters. Kendal Tau was computed to assess this similarity.

6. Analysis

The objective of testing the MAS was to determine the effectiveness of the proposed methodology to assess the degree of membership of the buddy-agents. To this end, 5000 stocks were selected at random from the NYSE and financial data for each stock was collected. Cluster analysis, based on 17 financial attributes, identified eight stock-cluster among the 5000 selected stocks. The data for each cluster was saved in the case-base of a CBR system to represent the portfolio of an investor. Next, 48 requests were generated from each case-base. The MAS was responsible for sending the requests to other agents and returning the responses to the requesting agents. Here, we used the quality of response as our only criteria for assessing the degree of membership of the agents. The reason for ignoring other attributes such as response time is that our basis for assessing the goodness of computed membership was its comparison with the distance between clusters (see Table 2). The rank-order of the portfolios (clusters) in relation to the distance between each is presented in Table 3. Table 4 presents the degree of membership of agents based on 48 requests sent to other agents and the feedback received. Our contention is that the goodness of our proposed buddy-agent membership model can be assessed by comparing it with the distance among clusters.

The correlation between the membership ranking of the two methods (the distance between clusters and our proposed method) show that the two types of assessments are highly correlated (Kendall's tau = 0.469 ($p < 0.001$)) as shown in Table 5. This indicates that our proposed model for assessing membership of the buddy-agents works well.

Table 2

Distance between the Eight Portfolios (clusters)

	Portfolio1	Portfolio 2	Portfolio 3	Portfolio 4	Portfolio 5	Portfolio 6	Portfolio 7	Portfolio 8
Portfolio 1		1.85	3.39	2.87	3.18	1.83	3.07	3.74
Portfolio 2	1.85		4.39	3.9	3.77	2.01	3.62	4.14
Portfolio 3	3.39	4.39		2.27	2.26	3.21	3.09	4.68
Portfolio 4	2.87	3.9	2.27		2.09	2.4	1.82	2.82
Portfolio 5	3.18	3.77	2.26	2.09		1.95	1.36	3.34
Portfolio 6	1.83	2.01	3.21	2.4	1.95		1.65	2.72
Portfolio 7	3.07	3.62	3.09	1.82	1.36	1.65		1.99
Portfolio 8	3.74	4.14	4.68	2.82	3.34	2.72	1.99	

Table 3

Rank Order of the Distance between Clusters

	Portfolio1	Portfolio 2	Portfolio 3	Portfolio 4	Portfolio 5	Portfolio 6	Portfolio 7	Portfolio 8
Portfolio 1		1	5	6	5	2	5	5
Portfolio 2	2		6	7	7	4	7	6
Portfolio 3	6	7		3	4	7	6	7
Portfolio 4	3	5	2		3	5	3	3
Portfolio 5	5	4	1	2		3	1	4
Portfolio 6	1	2	4	4	2		2	2
Portfolio 7	4	3	3	1	1	1		1
Portfolio 8	7	6	7	5	6	6	4	

Table 4

Rank Order of the Membership of the Buddy-Agents

	Agent1	Agent 2	Agent3	Agent4	Agent5	Agent6	Agent7	Agent8
Agent1		3	2	2	5	4	5	6
Agent2	4		7	7	4	2	4	4
Agent3	5	6		1	6	6	6	7
Agent4	6	7	1		7	7	7	5
Agent5	3	4	3	4		3	2	3
Agent6	2	1	5	5	2		1	2
Agent7	1	2	4	3	1	1		1
Agent8	7	5	6	6	3	5	3	

Table 5
Correlation Between the Agent Rank order of Buddy-Agents
and the Distance Between Portfolio (clusters)

			Agent	Portfolio
Kendall's tau_b	Agent	Correlation Coefficient	1	0.469
		Sig. (2-tailed)		0
		N	56	56
	Portfolio	Correlation Coefficient	0.469	1
		Sig. (2-ailed)	0	
		N	56	56

** Correlation is significant at the .01 level (2-tailed).

7. Concluding Remarks

Agent technology is believed to be an effective way to reduce decision-makers' information overload. By delegating tasks to agent systems, decision-makers save not only time and energy, but also increase their opportunities to access valuable information and work on more complex and creative jobs. To communicate effectively, each agent in a multi-agent system needs to know the characteristics of the other agents that can best serve its requirements. In this paper, we proposed a methodology to assess the degree of membership of buddy-agents. This methodology is based on fuzzy set modeling. The objective is to select buddy-agents that are expected to meet a set of criteria in responding to a request.

Our proposed methodology was empirically tested based on an application of portfolio selection. The empirical test was to determine the effectiveness of the proposed framework to assess buddy-agent membership in support of complex decision problems. The test result validates our proposed model for assessing membership of the buddy-agents.

7.1 Implications for Practice

The Internet has created an enormous amount of business opportunities for electronic commerce. Web-based businesses are becoming increasingly popular and Internet users have become fast-growing groups that form a promising market (Ting-Peng Liang et al., 2000). We can purchase a variety of goods and services such as flowers, books, and automobiles on-line. However, searching through thousands of possible sites is time-consuming and frustrating. The methodology presented in this paper, can be used to facilitate searches among a large number of distributed informational sources. For example, many music lovers share music through Napster (www.NAPSTER.COM). Napster enables a music lover to connect to a community of millions of music lovers. It allows a music lover to search and browse music files in MP3 and WMA formats and to chat with other members of the Napster community. Users search for music by artist and title, and then download the located music from other music lover's sites. With Napster, users may have to sift through hundreds of titles to find one that they have not heard matches their taste. Our proposed buddy-agents-membership methodology can reduce search time by recommending the best match of interests. Under this application, music style preferences are represented by the parameters of artist, composer, rhythm, price, etc. Our buddy-agent-membership methodology can identify the best members (i.e., buddy-group members) that can offer information about music to a person seeking advice. This helps music lovers narrow their selection and learn about new music they are interested in.

7.2 Implications For Research

While the results of our empirical tests are encouraging, we acknowledge that its context is limited. Stock traders identifying other traders with similar investment portfolios provide us with a reasonable experimental context, but we need to know the result in other contexts to assess the complete capability of the model. One drawback of our methodology is that when the number of buddy-agents goes below a threshold, then request has to be sent to all other agents to recruit new buddy-agents. This could be inefficient (e.g., computing workload increases proportionally as the number of agents increases). Future research should identify methodologies that utilize information about the buddy-agents of the requesting agents to assess new buddy-agents' membership. Another issue in need of attention is selection procedures for responding to a request. We used FIFO (first-in-first-out) procedure by a responding agent to serve the requesting agents in the queue. This can be an inefficient procedure since it is based on the assumption that is equally likely to give valuable recommendation to all the requesting agents. Markov decision process modeling (Sutton and Barto, 1998) may provide a better procedure for serving requests from buddy-agents. It is our hope that our findings will generate fruitful discussion and provoke further research.

References

Bordetsky, A. and Mark, G., "Memory-Based Feedback Controls to Support Groupware Coordination," Information Systems Research, Volume 11, Number 4, December 2000, pp.366-385.

Cox, E., "The Fuzzy Systems Handbook," Second Edition, AP Professional, New York, NY, 1999.

Finin, T., Labrou, Y. and Mayfield, "KQML as an Agent Communication Language", In J.M. Bradshaw (Ed.), Software Agents, The MIT Press, Cambridge, MA, 1997, pp.291-316.

Gupta, K.M. and Montazemi, A.R., "Empirical Evaluation of Retrieval in Case-Based Reasoning Systems Using Modified Cosine Matching Function," IEEE Transaction on Systems, Man, and Cybernetics," Volume 27, Number 5, 1997, pp.601-612.

Hayes-Roth, B. (1995), "An architecture for adaptive intelligent systems", Artificial Intelligence, 72, 329 - 365.

Lewis, M., "Designing for Human-Agent Interaction", AI Magazine, Volume 19, Number 2, Summer 1998, pp.67-78.

Lieberman, H., "Autonomous Interface Agents", Proceedings of the ACM Conference on Computers and Human Interface, CHI-97, Atlanta, Georgia, March 1997.

Maes, P., "Agents that reduce work and information overload," Communications of the ACM, Volume 37, Number 7, July 1994, pp.31-40.

Maes, P., Guttman R.H. and Moukas A.G., "Agents that Buy and Sell: Transforming Commerce as we Know It", Communications of the ACM, March 1999.

March, S., Hevner, A., and Ram, S., "Research Commentary: An Agenda for Information Technology Research in Heterogeneous and Distributed Environments," Information Systems Research, Volume 11, Number 4, December 2000, pp.327-341.

Minsky, M., "Common Sense-Based Interfaces," Communications of the ACM, Volume 43, Number 8, August 2000, pp.66-73.

Minsky, M., "Steps Toward Artificial Intelligence," In E.A Feigenbaum and J. Feldman (Eds.), Computers & Thought, AAAI Press, 1995, pp.406-450.

Minsky, M., "A Conversation with Marvin Minsky about Agents," Communications of the ACM, Volume 37, Number 7, July 1994, pp.22-29.

Montazemi, A.R., "Case-Based Reasoning and Multi-Agent Systems in Support of Tacit Knowledge", AAAI Workshop on Exploring Synergies of Knowledge Management & Case-Based Reasoning, AAAI-99 Sixth National Conference on Artificial Intelligence, Orlando, Florida, July 18-22, 1999, pp.60-63.

Montazemi, A.R. and Gupta, K.M., "On the Effectiveness of Cognitive Feedback from an Interface Agent," OMEGA, International Journal of Management Sciences, Volume 25, Number 6, 1997, pp.648-658.

Montazemi, A.R. and Gupta, K.M., "A framework for Retrieval in Case Based Reasoning Systems," Annals of OR on the Interface Between IS and OR, Volume 72, 1997 pp.51-73.

Montazemi, A.R. and Gupta, K.M., "An Adaptive Agent for Case Description in Diagnostic CBR Systems," Computers in Industry, Volume 26, 1996, pp.209-224.

Nwana, H. S. (1996), "Software Agents: An Overview," Knowledge Engineering Review, Volume 11, Number 3, pp.205-244.

Pilkington, R.M., "Intelligent Help: Communicating with Knowledge-Based Systems", Oxford University Press, Oxford, U.K., 1992.

Powell, W.W., "Learning from collaboration: knowledge and networks in the biotechnology and pharmaceutical industries", California Management Review, Volume 40, Number 3, 1998, pp.228-40.

Simon, H.A., and Hayes, J.R., "The Understanding Process: Problem Isomorphs", Cognitive Psychology, Volume 8, Number 2, 1976, pp.165-190.

Sutton, R.S., and Barto, A.G., "Reinforcement Learning," The MIT Press, Cambridge, MA, 1998.

Sycara, K., "Multi -agent Systems," AI Magazine, Volume 19, Number 2, 1998, pp.79-92.

Sycara, K. and Pannu, A. and Williamson, M. and Zeng, D. (1996), "Distributed Intelligent Agents", IEEE EXPERT December 1996, pp.36-46.

Tversky, A., and Kahneman, D., "Causal Schemas in Judgment Under Uncertainty", In D. Kahneman, P. Slovic and A. Tversky (Eds.), Judgment Under Uncertainty: Heuristics and Biases, Cambridge University Press, Cambridge, UK., 1982.

Ting-Peng Liang and Jin-Shiang Huang, "A framework for applying intelligent agents to support electronic trading", Decision Support systems 28 (2000) 305-317.

System Architecture of BlackBoard Agent

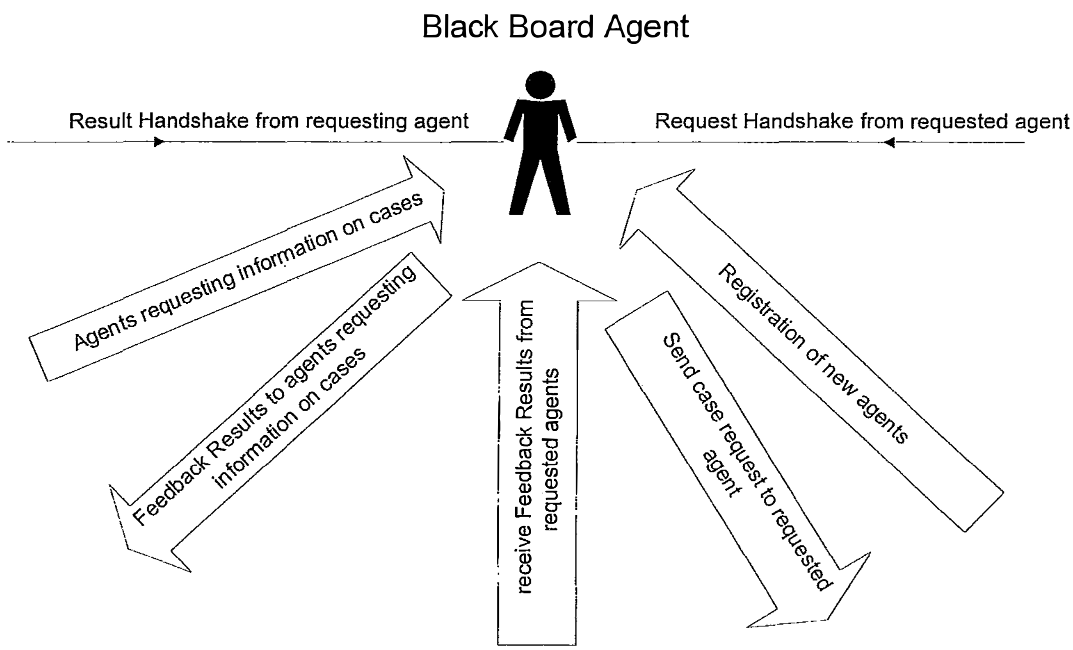

Michael G. DeGroote School of Business
McMaster University

WORKING PAPERS - RECENT RELEASES

424. Harish C. Jain, Michael Piczak, Işik Urla Zeytinoğlu, "Workplace Substance Testing - An Exploratory Study", November, 1997.

425. S. Suarga, Yufei Yuan, Joseph B. Rose, and Norman Archer, "Web-based Collective Bargaining Support System: A Valid Process Support Tool for Remote Negotiation", January, 1998.

426. Pawan S. Budhwar and Harish C. Jain, "Evaluating Levels of Strategic Integration and Development of Human Resource Management in Britain", March, 1998.

427. Halit Üster and Robert F. Love, "Application of Weighted Sums of Order p to Distance Estimation", April, 1998.

428. Halit Üster and Robert F. Love, "On the Directional Bias of the ℓ_{bp}-norm", April, 1998.

429. Milena Head, Norm Archer, and Yufei Yuan, "MEMOS: A World Wide Web Navigation Aid", October, 1998.

430. Harish C. Jain and Parbudyal Singh, "The Effects of the Use of Strike Replacement Workers on Strike Duration in Canada", February, 1999.

431. Parbudyal Singh and Harish C. Jain, "Strike Replacements in the United States, Canada and Mexico: A Review of the Law and Empirical Research", February, 1999.

432. John W. Medcof and Jeremy Boyko, "Reinforcing, Revising and Reconciling Attributions in the Employment Interview", March, 1999.

433. Norm Archer, "World Wide Web Business Catalogs in Business-to-Business Procurement", March, 1999.

434. Diwakar Gupta and Saifallah Benjaafar, "Make-to-order, Make-to-stock, or Delay Product Differentiation? - A Common Framework for Modeling and Analysis", April, 1999.

435. Harish C. Jain, Parbudyal Singh and Carol Agocs, "Recruitment, Selection and Promotion of Visible Minorities and Aboriginals in Selected Canadian Police Services", April, 1999.

436. Harish C. Jain and Angus Bowmaker-Falconer, "Employment Equity/Affirmative Action Codes of Practice and Best Practices in USA, Britain, Canada and Other Selected Countries", May, 1999.

437. Diwakar Gupta, Yavuz Günalay, and Mandyam M. Srinivasan, "On the Relationship Between Preventive Maintenance and Manufacturing System Performance", June, 1999.

438. Jinliang Cheng, George Steiner, and Paul Stephenson, "A Fast Algorithm to Minimize Makespan for the Two-Machine Flow-Shop Problem with Release Times", June, 1999.

439. Jinliang Cheng, George Steiner, and Paul Stephenson, "A Fast Algorithm to Minimize Maximum Lateness for the Two-Machine Flow-Shop Problem with Release Times, June, 1999.

440. Norm Archer and Yufei Yuan, "Electronic Commerce and the Business-to-Business Customer Relationship Life Cycle", August, 1999.

441. Halit Üster and Robert F. Love, "Calculation of Confidence Intervals for Estimated Distances", August, 1999.

442. George O. Wesolowsky, "Detecting Excessive Similarity in Answers on Multiple Choice Exams", October, 1999.

443. Naresh C. Agarwal, "Mandatory Retirement and the Canadian Human Rights Act", November, 1999.

444. Susan Sproule and Norm Archer, "Software Agents in Electronic Commerce", March, 2000.

445. Milena Head and Yufei Yuan, "Privacy Protection in Electronic Commerce – a Theoretical Framework", June, 2000.

446. Susan Sproule and Norm Archer, "Knowledgeable Agents for Search and Choice Support in E-commerce: A Decision Support Systems Approach", July, 2000.

447. Norm Archer and Judith Gebauer, "Managing in the Context of the New Electronic Marketplace", August, 2000.

448. Norm Archer and Yufei Yuan, "Business-to-Business E-commerce and Customer Relationship Management: Trends and Issues", August, 2000.

449. Janet Romaine and I.U. Zeytinoglu, "Are Women Managers Really More Participative? Some Conflicting Evidence from a Field Study", September, 2000.

450. Norm Archer, "The New Economy: Some Issues and Impacts of Electronic Commerce", January 2001.

451. Norm Archer, Milena Head, and Yufei Yuan, "Matching Customer Expectations and Perceptions in e-Commerce", February 2001.

McMASTER UNIVERSITY LIBRARY
3 9005 0232 3444 0

Innis
Ref.
HB
74.5
.R47
no.452