경영정보학연구
제10권 제3호
2000년 9월

분산인공지능 모델을 이용한 효과적인 팀 의사결정에 관한 연구*

강 민 철**

A Study of Effective Team Decision Making Using A Distributed AI Model

Kang, Mincheol

The objective of this paper is to show how team study can be advanced with the aid of a current computer technology, that is distributed Artificial Intelligence (DAI). Studying distributed problem solving by using groups of artificial agents, DAI can provide important ideas and techniques for the study of team behaviors like team decision making.

To demonstrate the usefulness of DAI models as team research tools, a DAI model called "Team-Soar" was built and a simulation experiment done with the model was introduced. Here, Team-Soar models a naval command and control team consisting of four members whose mission was to identify the threat level of aircraft. The simulation experiment was performed to examine the relationships of team decision scheme and member incompetence with team performance. Generally, the results of the Team-Soar simulation met expectations and confirmed previous findings in the literature. For example, the results support the existence of main and interaction effects of team decision scheme and member competence on team performance.

Certain results of the Team-Soar simulation provide new insights about team decision making, which can be tested against human subjects or empirical data.

* The present research has been conducted by the Bisa Research Grant of Keimyung University in 1998.
** 계명대학교 경영학부

Ⅰ. Introduction

Teams are becoming increasingly important to modern organization. Nowadays, many organizations are using teams for different reasons, such as in order to solve exceptional problems or to carry out special projects.

The popularity of team is somewhat due to a common belief that making employees work in teams will achieve high performance. However, many cases like the incident of the U.S.S Vincennes [Hollenbeck et al., 1997] have showed that assigning a job to a team doesn't in and of itself guarantee success of the work.

Therefore, it is necessary to study team behaviors systematically.

Distributed Artificial Intelligence (DAI), a subfield of AI, can provide important ideas and techniques for the study of team behaviors like team decision making. DAI studies distributed problem solving by using groups of artificial agents, whereas team decision making can be considered to be distributed problem solving by a group of humans. Therefore, DAI has potential as a framework for modeling the behavior of teams.

Up to now, however, only a small number of computer team models, for example, TEAMWORK model [Doran, 1985], i-AGENTS model [Jin and Levitt, 1993], SYBORG model [Yoshimura et al., 1994], have been used for team study. In particular, relatively few, if any, distributed AI models have been used for team research due to the interdisciplinary nature that requires knowledge both about the scientific area (i.e., management, especially group dynamics) and the engineering area (i.e., distributed AI). The objective of this paper is to show how team study can be advanced with the aid of the current computer technology. Especially, by introducing a simulation experiment that uses a distributed AI model of team called "Team-Soar", this paper illustrates how beneficial to adopt the interdisciplinary approach in studying team behaviors.

The remainder of this paper is organized as follows: In the next section, we examine what distributed AI is and why the technique is good for team study. Section III introduces Team-Soar model and explains how it models a naval carrier team. Section IV describes a simulation experiment of Team-Soar and an analysis of the result. In Section V, contribution of this research is discussed as a conclusion.

Ⅱ. Distributed AI as a Paradigm for the Modeling of Teams

2.1 Definition

DAI is a subfield of AI that has focused on how a collection of artificially intelligent agents in a problem-solving situation can interact effectively to achieve a common set of global goals [Weiss, 1999; Ferber, 1999; Kraus, 1997; Chaib-Draa et al., 1992; Bond and Gasser, 1988b]. As computers have become more sophisticated, the demands for coordination and cooperation have become more critical. To cope with the demands, research in the area of DAI has studied computer mechanisms by which multiple intelligent and autonomous agents can coordinate and cooperate effectively [Sen, 1997].

DAI are divided into the following three areas: parallel AI (PAI), distributed problem solving

(DPS), and multi-agent systems (MAS). Parallel AI achieves a linear speed-up by applying multiple processors concurrently to a given problem [Griffiths and Purohit, 1991; Bond and Gasser, 1988b]. While studies in distributed problem solving consider how the work of solving a particular problem can be divided among a number of agents, studies in multi-agent systems are concerned with coordinating intelligent behavior among a collection of autonomous intelligent agents [Bond and Gasser, 1988b; Griffiths and Purohit, 1991]. The autonomous intelligent agents must reason about the processes of coordination among themselves. The area of multi-agent systems has been won more popularity than the other two areas owing to that multi-agent systems are able to operate without direct human intervention, able to improve over time, and able to communicate with other agents. To reflect the focus of interest of the active researchers, the area of DAI often uses the name of multi-agent systems to indicate the filed as a whole [Sen, 1997].

2.2 Analogies between Human Teams and DAI Systems

A DAI system is a loosely coupled network of individual artificial agents or processes interacting together to solve problems that are beyond their individual capabilities. There are analogies between human teams and DAI systems. Both human teams and DAI systems are arrangements of parallel distributed intelligence for multi-agent problem solving [Masuch, 1992; Gasser and Hill, 1990]. Like a DAI system, a team can be thought of as a network that consists of agents as nodes and communication channels as connections between these nodes. Furthermore, the properties of both human teams and DAI systems are not derivable or representable solely on the basic properties of their component members or agents [Chaib-Draa et al., 1992]. They both display social behavior [Carley and Newell, 1994]. They even face the same problem of allocating tasks, resources, and information to sets of intelligence [Fox, 1981].

These analogies and others support the argument that DAI systems could serve as models of human teams for supporting theoretical work [Masuch, 1992; Huberman, 1992; Gasser and Hill, 1990]. In this sense, DAI is the experimental branch of organizational science [Crowston, 1992].

2.3 DAI and Team Decision Making

A DAI system can be used to model the distributed problem-solving process of a team like team decision making, just as an AI system can be used to model the problem-solving process of an individual human [Shaw et al., 1991; Griffiths and Purohit, 1991]. DAI models provide symbolic-level frameworks that can capture the distribution of processing or computation embedded in team decision making [Chandrasekaran, 1981].

Team decision making is a macro-level phenomenon that emerges from micro-level interactions. DAI modeling is a promising approach for studying the micro-macro link, because DAI systems can explore the global behavior of a collection of agents based on the local knowledge and local procedures of each agent [Huberman, 1992]. One of the most important features of DAI models is that each independent computational unit has its own local knowledge and controls its own execution. Therefore, DAI systems can model the problem-solving behavior of a team based on the local knowledge and local problem-solving activities of each individual member.

DAI can provide a means of studying a multiplicity of factors that influence team decision making.

Traditional research on teaming typically employs surveys or other techniques for gathering anecdotal evidence [Katzenbach and Smith, 1993], or uses experimental methods in quasi-laboratory settings [Hollenbeck et al., 1997]. Both approaches have their benefits: surveys can help identify research questions or hypotheses, while experimental methods can provide in-depth understanding of one or two factors in team decision making. However, simulations using artificial agents provide a setting where a multiplicity of theoretic constructs, such as the characteristics of the agents and their interactions, can be incorporated into models and assessed by simulation. In addition, DAI models of team decision making require detailed descriptions of members' cognition, interaction, group design, and task [Carley and Prietula, 1994; Jin and Levitt, 1993]. Hence, DAI models enable researchers to study the complex relations that characterize teams analytically.

So far in this section of the paper, we have reviewed the literature of DAI field. For further review of DAI literature including an architecture of DAI systems, see Weiss [1999], Ferber [1999], Sen [1997], Bond and Gasser [1992, 1988a], Gasser and Huhns [1989], and Huhns [1987].

Ⅲ. Team-Soar

"Team-Soar" is a distributed AI model of a naval command and control team consisting of four members who have different expertise and cooperate interactively to accomplish aircraft identification tasks (i.e., identifying the threat level of aircraft). For the team model, four AI agents are realized on a SUN machine by using a multi-agent Soar technique developed for distributed problem solving [Laird et al., 1993]. In Team-Soar, the four individual AI agents are interconnected to represent a communication channel between team members.

The team being modeled by Team-Soar consists of Commanding Officers (CO) of four units in a naval carrier group. The leader is the CO of the Aircraft Carrier (modeled by CARRIER-Soar). The other members are the CO of a Coastal Air Defense unit (modeled by CAD-Soar), the CO of an AWACs air reconnaissance plane (modeled by AWAC-Soar), and the CO of an Aegis Cruiser (modeled by CRUISER-Soar). Figure 1 shows Team-Soar that models four roles of the naval command and control team and their major activities. In the model, aircraft are tracked by radar and evaluated in terms of nine attributes as shown on the bottom of the picture in Figure 1. In the figure, the lines from members to attributes show which members are expertise on what attributes.

To participate in a team decision, each member first makes its own judgment about the best course of action by using the information available to it, then recommends this judgment to the leader, that is, CARRIER. To make a judgment, a member first interprets the raw data for the attributes and evaluates

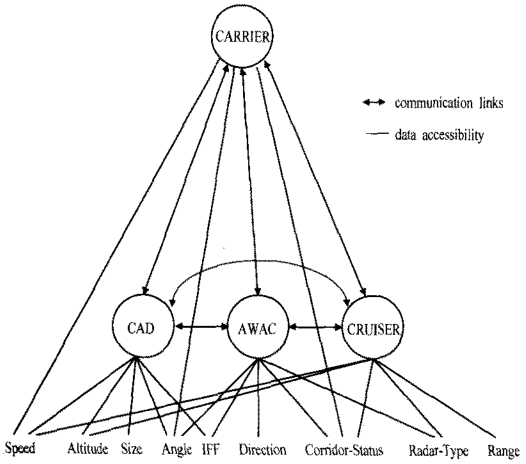

<Figure 1> Team-Soar that models a naval carrier team of four experts

each attribute on a scale of zero to two. In the evaluation, the numbers, zero, one, and two mean "non-threatening," "somewhat-threatening," and "very-threatening" respectively. A member also may ask other

members for their evaluations of certain attributes. When a member has made the required evaluations, it then makes a judgment about which of seven possible courses of action to recommend to the leader. The seven possible courses of action range in degree from Ignore, which has a value of zero, to Defend, which has a value of six [Hollenbeck, et al., 1995]. Intermediate actions on this scale are Review (1), Monitor (2), Warn (3), Ready (4), and Lock-on (6). Upon receiving all other members' judgments, the leader makes a team decision based all members' judgments, including its own (CARRIER's).

In Team-Soar, each AI agent of the model, which represents an individual team member, is a theory-based cognitive model of human called Soar [Laird et al., 1993]. Soar is especially used for modeling the individual team members because decision-making involving aircraft identification is a cognitive work. A candidate model for Newell's (1990) unified theories of cognition, Soar models the human cognitive capabilities of knowledge-based problem solving, learning, and interacting with the external environment.

Soar also incorporates into its model the view of the human as a general problem solver (Newell and Simon, 1972), a symbol system (Newell, 1990; Newell and Simon, 1976), and a knowledge system (Newell, 1990).

As modeled by Soar, each member maintains both long-term memory and working memory. All knowledge, including expertise, is stored in each member's long-term memory in the form of a production, that is, an "if-then" rule. Working memory, on the other hand, keeps only the knowledge that is relevant to the current cognitive activity of the member. The content of working memory is decided by a decision mechanism using the preference concept and is selected from the knowledge in long-term memory. The preference concept is part of the Soar architecture. See work by Laird and his colleagues (1993) for a detailed description of Soar's architecture.

Soar implements the idea of the problem space hypothesis (Newell, 1993; Newell, 1980) arguing that all human symbolic goal-oriented behavior can be conceived of as a search in a problem space. According to the problem space paradigm, the Soar model can be described in terms of goals, problem spaces, states, and operators (Laird, et al., 1993). Being modeled by a group of Soars, the Team-Soar model incorporates the problem space paradigm as follows. During the team decision making, each member in Team-Soar develops two problem spaces: a *team* *problem space and a member problem space.* In the team problem space, the member tries to achieve the team goal; in member problem space, the member tries to achieve the member goal. The team goal is to make the team decision correctly, whereas the member goal is to make a good recommendation to the leader. These goals are achieved by applying appropriate operators to states in the corresponding problem spaces.

When a team member first perceives an unidentified target, the member develops its own version of a team problem space in its working memory. Then the initial state within the member's team problem space is refined with *group knowledge* (i.e., the metalevel knowledge about group factors), other members, and other member's expertise. After refining its team problem space, each member develops an individual member goal as a subgoal of the team goal in accordance with its position on the team, and then develops its member problem space. The member who first spotted the unidentified target announces the appearance of that target to the other members (announce-target operator). When receiving this announcement, the other members follow the same process of problem space development as the first member, except for the announcing activity.

In its member problem space, each member reads the raw values of attribute data (read-attribute operator), evaluates the values (evaluate-attribute operator), asks other members their evaluations of some attributes (ask-evaluation operator), or reports its evaluations to other members upon request (report-evaluation). After collecting all required information, the member makes its judgment (make-member-judgment operator), and then returns to the team problem space.

When returning to its team problem space, the leader asks the other members for their judgments (ask-member-judgment operator), and the other members report their judgments to the leader (report-member-judgment operator). Upon receiving all others' judgments, the leader makes a team decision by using a team decision scheme (make-team-decision operator), and announces the decision to all other members (announce-team-decision operator). Then all members close their team problem spaces and wait for a new task. Refer to "Team-Soar: A Model for Team Decision Making" [Kang et al., 1998] for more description of Team-Soar model.

Ⅳ. Simulation Experiment

Using a DAI modeling technique, Team-Soar permits researchers to explore different combinations of modeling factors and allows them to analyze the modeling subject systematically. In order to show how DAI models can provide insights about team decision making, a simulation experiment done with Team-Soar will be introduced in this section.

4.1 Experimental Design

Researchers have recognized from human experiments that the use of different group decision schemes and / or the presence of incompetent team members may result in different performance levels [Hollenbeck et al., 1995; Miller, 1989]. However, no research has been reported that examines the interaction effects of group decision schemes and member competence on group performance. Therefore, it is meaningful to examine the effects on team performance systematically by using a simulation experiment. For the reason, a simulation experiment was performed to examine the role of different team decision schemes in alleviating the effect of member incompetence on team performance. In this simulation experiment, team performance was measured by two standards of team effectiveness: *decision deviation and disaster rate*. Decision deviation refers to the deviation of the decision that the team made from the correct decisions, which are predetermined by the Team-Soar Table 1 A Summary of the Simulation Experiment Design, the better the team's performance, that is, the better the team's decision accuracy. Disaster rate was a function of the frequency of decisions that were off by four or more points from correct decisions [Hollenbeck, et al., 1995]. The simulation experiment used a 4×5 design where team decision scheme and member competence were manipulated. Each of the twenty team models made 10,000 team decisions.

To make a team decision, the leader uses a decision rule (i.e., decision scheme) that provides a mechanism of combining member judgments into a team decision. Four team decision schemes were examined in this experiment: *majority win (tie breaker: leader/CARRIER), majority win (tie breaker: CAD), average win (fixed weight), and average win (dynamic weight)* schemes. In majority win schemes, the team's decision was determined by which recommendation was made by the majority of the team members. When there was a tie, the leader (CARRIER) either followed its own judgment (majority win with tie breaker leader) or followed the judgment of the member CAD, who had accessed more information than the leader (majority win with tie breaker CAD). In the average win schemes, a rounded-off average of all the weighted member judgments determined the decision. To compute this average, for both fixed and dynamic versions of the average win scheme, member judgments were converted into quantitative values and weighted appropriately according to the leader's assessment of the quality of each member's judgment.

This assessment was based on the member's expertise or past performance. In the fixed weight version of the average win scheme, the leader assigned a fixed and equal weight to each member's judgment; this weight was used for all decision tasks. In the average win scheme with dynamic weights, the leader, for each task, assigned weights to each member's judgment by considering past performance history of the member up to the decision point. Therefore, in the decision scheme, the dynamic weight assigned to a member's judgment was derived from the member's performance rating at that point in time. A member's performance rating consisted of points accrued from the correctness of its judgments on previous decision tasks: each member received seven points whenever its judgment was correct, six points whenever its judgment was off by one from the correct decision, and so forth.

Two different types of members can be identified according to their competence; competent and incompetent members. Competence members are agents who use available information rationally to make their judgments. They combine evaluations of attributes to produce a score, then make member judgments based on that score. Incompetent members act exactly the same as competent members except that they make

Table 1 A Summary of the Simulation Experiment Design

Classification	Description
experimental subjects	Team-Soar consisting of 4 AI agents
measures	decision deviation and disaster rate
design	4×5 (= 4 levels of factor 1×5 levels of factor 2)
factor 1	team decision scheme (4 levels) - majority win (tie breaker: leader/CARRIER) - majority win (tie breaker: CAD) - average win (fixed weight) - average win (dynamic weight)
factor 2	member competence (5 levels) - team with no incompetent member - team with one incompetent member - team with two incompetent members - team with three incompetent members - team with four incompetent members

their judgments by randomly selecting one simulation experiment, member competence was manipulated at five levels by changing the number of incompetent members from zero to four in the four-member teams. The design of the simulation experiment is summarized in Table 1.

4.2 Hypothesis Tests and Results

To statistically examine the effects of team decision scheme and member competence on team decision deviation, the following three null hypotheses were tested:

(1) the four team decision schemes affect team decision deviation equally;
(2) the five different levels of member competence affect team decision deviation equally; and
(3) there are no interaction effects of the decision scheme and the level of member competence on team decision deviation.

Two-way ANOVA was used for the hypothesis tests. All three null hypotheses were rejected at the .0001 level of significance. Thus, the test results strongly support the existence of main and interaction effects of team decision scheme and member competence on team's mean decision deviation.

Since interaction effects were significant, Tukey's studentized range test was performed at the 0.05 level of significance in order to determine which combinations of decision scheme and member competence differ significantly. Note that because each team model of the simulation experiment was configured with a combination of the two decision variables, the total number of observations for both ANOVA procedure and Tukey test was 200,000 (= 4×5×10,000). Tukey's test results are shown in the second column of Table 2. According to the Tukey test, within a team decision scheme, team decision deviation increased as more incompetent members were added to teams.

By inspecting members' performances in Table 2, we can tell which team positions in each configuration were occupied with incompetent members because performances by any incompetent members result in values above 2.25. According to Table 2, for teams with one incompetent member, AWAC is the position occupied by the incompetent member. For teams with two incompetent members, the AWAC and CRUISER positions are incompetent. For teams with three incompetent members, the AWAC, CRUISER and CAD positions are incompetent. For teams with four incompetent members, all positions are filled with incompetent members.

To examine the effects of team decision scheme and member competence on disaster rate, the

<Table 2> Results of Team-Soar simulation experiment

Team Configuration	Mean of Team Decision Deviation (Tukey Grouping at α= 0.05)	Average Member Decision Deviation	Deviation Difference	Mean of CARRIER Deviation (Leader)	Mean of CAD Deviation	Mean of AWAC Deviation	Mean of CRUISER Deviation	Mean of Disaster Rate (Tukey Grouping at α= 0.05)
4poor, ldr	2.260 (A)	2.281	0.021	2.281	2.273	2.269	2.301	0.2405 (A)
4poor, cad	2.286 (A)	2.280	-0.006	2.281	2.293	2.291	2.254	0.2482 (A)
4poor, fix	1.853 (C)	2.287	0.434	2.273	2.301	2.291	2.284	0.1120 (D)
4poor, dyn	1.950 (B)	2.274	0.324	2.267	2.274	2.280	2.275	0.1406 (C)
3poor, ldr	1.229 (E)	1.875	0.646	0.714	2.253	2.271	2.264	0.0805 (E)
3poor, cad	1.949 (B)	1.893	-0.056	0.711	2.300	2.268	2.295	0.1919 (B)
3poor, fix	1.488 (D)	1.885	0.397	0.711	2.264	2.290	2.277	0.0315 (F)
3poor, dyn	0.782 (G)	1.886	1.104	0.711	2.279	2.275	2.281	0.0016 (H)
2poor, ldr	0.758 (G)	1.444	0.686	0.704	0.482	2.323	2.267	0.0190 (G)
2poor, cad	0.665 (H)	1.440	0.775	0.722	0.482	2.282	2.276	0.0179 (G)
2poor, fix	1.116 (F)	1.441	0.325	0.707	0.482	2.289	2.284	0.0011 (H)
2poor, dyn	0.534 (I, J)	1.433	0.899	0.713	0.489	2.278	2.251	0 (H)
1poor, ldr	0.575 (I)	1.005	0.430	0.720	0.480	2.271	0.549	0.0009 (H)
1poor, cad	0.502 (J, K)	1.008	0.506	0.719	0.487	2.279	0.548	0.0003 (H)
1poor, fix	0.764 (G)	1.007	0.243	0.721	0.478	2.272	0.555	0 (H)
1poor, dyn	0.456 (K, L)	1.007	0.551	0.720	0.501	2.252	0.554	0 (H)
0poor, ldr	0.465 (K)	0.558	0.093	0.717	0.483	0.487	0.544	0.0003 (H)
0poor, cad	0.414 (L, M)	0.553	0.139	0.706	0.479	0.481	0.548	0 (H)
0poor, fix	0.404 (M)	0.560	0.156	0.724	0.489	0.482	0.546	0 (H)
0poor, dyn	0.368 (M)	0.562	0.194	0.717	0.489	0.488	0.553	0 (H)

Note:
- 0poor: no incompetent members in the team, 1poor: one incompetent member in the team, and so on.
- ldr: majority win scheme (tie breaker: leader), cad: majority win scheme (tie breaker: CAD)
  fix: average win scheme (fixed weight), dyn: average win scheme (dynamic weight).
- In Tukey grouping, teams with the same letter are not significantly different.

hypothesis tests were repeated, substituting disaster rate for team decision deviation. The main and interaction effects of team decision scheme and member competence on mean disaster rate were found to be significant at the 0.0001 level. Tukey's test performed at the 0.05 level of significance (see the last column of Table 2) showed that within a team decision scheme, disaster rate increased as more incompetent members were added to teams.

4.3 Discussion

The simulation experiment generated same kinds of data as human experiments, which can be analyzed in the same ways as analyzing human data. See Table 2 in order to appreciate the data from the simulation experiment. Figure 2 is provided to aid interpretation of the data.

Figure 2 part (a) shows that the two majority win schemes worked relatively well until incompetent members comprised the majority of the team. In particular, the performance of teams using the majority win scheme with CAD as a tie breaker declined dramatically after the position of CAD was taken by an incompetent member. This decline occurred because leaders followed CAD's judgments in making team decisions whenever there was a tie between members' judgments. The results indicate that the choice of the right team decision scheme is contingent upon the level and distribution of member competence. That is, the effectiveness of different schemes depends not only on the number of incompetent members, but also on the positions occupied by the incompetent members.

According to part (a) of Figure 2, decision deviation for teams using the average win scheme with fixed weight increased linearly as incompetent members increased one by one. Further, except for teams consisting of all incompetent members, this team decision scheme was not effective with the presence of incompetent members, because this scheme assigns incompetent members the same weight as other members. However, for teams with all incompetent members, this scheme is more effective at reducing the decision deviations from correct answers than any other scheme because it makes team decisions by taking an average member's judgments.

The gaps in mean decision deviation between teams using different team decision schemes at the same member competence level increased as the number of incompetent members increased until all members in the teams became incompetent members (see part (a) of Figure 2). The same is true with disaster rate (see part (b) of Figure 2). This reflects that the use of an appropriate team decision scheme becomes more important as the number of incompetent members increases, provided that at least one of its members is competent. The results also indicate that a good team decision scheme can provide a buffer for the effect of incompetent members on team decision deviation and disaster rate. For example, the performances of teams using the average win scheme with dynamic weight did not decline significantly as the presence of incompetent members increased until all members became incompetent.

Except for the extreme case in which all members are incompetent, the average win scheme with dynamic weight is the best team decision scheme in the presence of incompetent members when team performance is measured in terms of both team decision deviation and disaster rate (see Figure 2). Even in the extreme case, the performance of the teams using the dynamic weight scheme are close to the performance of teams using the average win scheme with fixed weight, which is the best scheme in this extreme case. Therefore, team decision schemes similar to the dynamic weight scheme that consider members' past performance history are recommended, especially when the number and position of incompetent members in a team change dynamically. Theoretically, however, to make team decisions when all members are incompetent, it is better to use the fixed weight scheme instead of the dynamic weight scheme because the fixed weight scheme does not require the extra memory and calculations needed to consider past performances.

It is easy to think, mistakenly, that the selection of a team decision scheme does not matter when all members are incompetent, that is, when all members make their judgments by random selection. Surprisingly, this experiment revealed that the selection of the right team decision scheme is still important even in this extreme situation. In this extreme situation, the use of the average win scheme with fixed weight significantly improved the team's decision accuracy (i.e., reduced the team's decision deviation) and disaster rate when compared to the results of using the two majority win schemes (see Figure 2). Particularly in comparison to the disaster rates produced by the majority win schemes, the average win scheme with fixed weight decreased the mean disaster rate by about half. However, it is not obvious immediately why this happens. Further research will be necessary to explain why.

In general, the results analyzed in terms of disaster rate coincided with the results analyzed in terms of team decision deviation. However, the results analyzed with respect to team decision deviation did not always match with the disaster rate results. For example, the mean decision deviations for teams using the average win scheme with fixed weight were the worst, or at least worse than teams using the majority win scheme with leader as a tie breaker when the teams contained one to three incompetent members; nevertheless, the mean disaster rates of the teams using the average win scheme

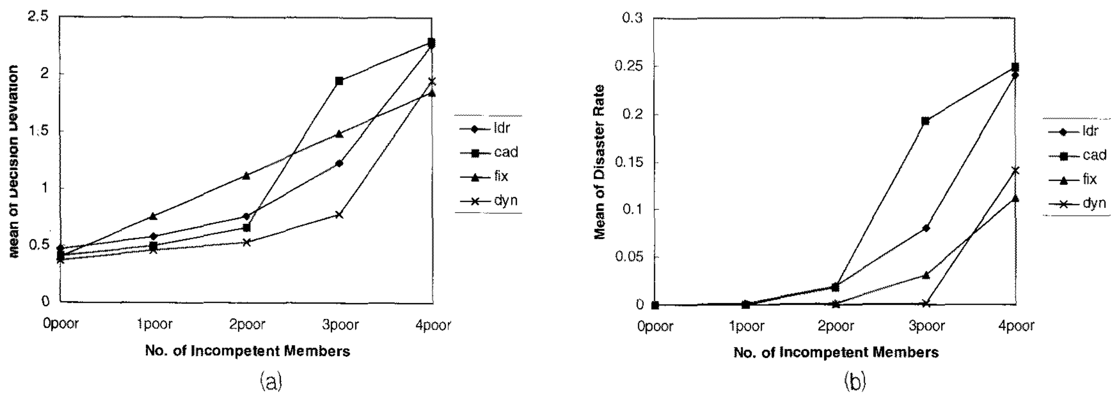

<Figure 2> The impact of member competence on (a) team decision accuracy and (b) disaster rate at different team decision schemes

with fixed weight were significantly lower than the latter teams' (see Figure 2). In fact, unlike the cases of team decision deviation where none of the average win schemes or the majority win schemes were dominant, the two average win schemes outperformed the two majority win schemes for all levels of member competence in the case of disaster rate. The reason is that the average win schemes try to approach the correct decisions by reducing the deviations of their decisions from the correct ones as much as possible by taking an average of members' judgments. On the other hand, the majority win schemes try to find out the correct decisions directly by following the majority judgments, while they do not give much attention to the deviations if they fail. Therefore, teams using the majority win schemes tend to produce more disasters than teams using the average win schemes.

It is necessary to be aware of a limitation when interpreting or applying the results of the Team-Soar simulation. Team-Soar models a particular type of team, which is a decision-making team with one level of hierarchy and distributed expertise. Therefore, some of the results may be only applicable to this special type of team.

V. Conclusion

The major purpose of this paper was to demonstrate the usefulness of DAI models as team research tools.

To this end, a distributed AI test bed called Team-Soar was built and results of a simulation experiment was introduced.

Generally, the results of the Team-Soar simulation met expectations and confirmed previous findings in the literature. Researchers have recognized that the use of different group decision schemes may lead to making different decisions, and, therefore, result in different performance levels [Miller, 1989]. The results of the Team-Soar simulation confirmed this result. Generally, the average win schemes outperformed the majority win schemes with respect to team decision accuracy and disaster rate. Researchers have also found that, in general, organizations perform better when given feedback if the agents can learn [Carley, 1991]. In the Team-Soar simulation, the average win scheme with dynamic weight is the only team decision scheme that utilizes an agent who learns from experience: the leader makes decisions based on team members' past performance history. The results of the Team-Soar simulation support Carley's finding by displaying that team models using the average win scheme with dynamic weight generally outperformed teams using other team decision schemes, which did not incorporate an agent who can learn. Researchers using human experiments have found that teams that contain an incompetent member perform worse than teams that contain no incompetent members [Hollenbeck et al., 1995]. The same results were found in the Team-Soar simulation. In the simulation, a team's decision accuracy deteriorated (i.e., a team's decision deviation increased) as the number of incompetent members in the team increased.

Certain results of the Team-Soar simulation provide new insights (i.e., new research propositions) about team decision making. For example, the simulation results indicate the interrelation of team decision scheme with member competence. According to results, the choice of the right team decision scheme becomes more critical when team contains more incompetent members. These new research propositions can be tested against human subjects or empirical data.

For some results of the Team-Soar simulation are counterintuitive, that is, explanations are not immediately obvious. For example, it cannot be explained easily why some team decision schemes (i.e., average win schemes) are still more effective than others' (i.e., majority win schemes) even when all members make their judgments by random selection. More research needs to be done for explaining these kinds of simulation results.

As a computational model, the Team-Soar model needs to be validated. In order to validate the model, we need to identify the goal of the model first because a computational model should always be validated for a particular goal or purpose (Burton and Obel, 1995). According to the purposes for which the models are formulated, Cohen and Cyert (1965) classify computational models of organizational behavior into the following four categories: descriptive, illustrative, normative, and man-machine. The goals of descriptive models are to examine why existing organizations have behaved in particular ways and to predict how they will behave in the future. The purpose of an illustrative model is to explore the implications of reasonable assumptions about organizational behavior in order to determine what the world is like when these assumptions are true. The purpose of a normative model is to allow researchers to determine which of several possible forms of organizations are best suited to particular goals, whereas the goal of a man-machine model is to train people to function better in organizational settings. According to the classification, Team-Soar corresponds to a descriptive model, because the goal of the model is to understand team behavior (i.e., team decision making). To validate a descriptive model, one can compare the model's results with the actual results made in the organization (Burton and Obel, 1995). With Team-Soar, this means matching the simulation results with the behavior of human teams. As described before, some results of the Team-Soar experiments correspond with the findings from previous research on human subjects. The comparison of observed human team phenomena with the simulation results provides some degree of validation of the Team-Soar model. Actually, a number of computational models has been validated by comparing their simulation results with built-in theories or empirical data (Doran, 1985; Lin, 1993).

In conclusion, this paper argues that DAI is a valuable research tool for studying teams by illustrating how a DAI model (i.e., Team-Soar) can provide deep insights into team decision making. Indeed, the study of teams can be progressed by adopting technologies and ideas from the area of distributed AI.

〈References〉

[1] Bond, A.H. and L. Gasser, *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann, San Mateo, CA., 1988.

[2] Bond, A.H. and L. Gasser, "An Analysis of Problems and Research in DAI," In Bond, A.H. and Gasser, L. (Eds.), *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann, San Mateo, CA., 1988, pp. 3-35.

[3] Bond, A.H. and L. Gasser, "A Subject-Indexed Bibliography of Distributed Artificial Intelligence," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 22, No. 6 (1992), pp. 1260-1281.

[4] Burton, R.M. and B. Obel, The Validity of Computational Models in Organization Science: From Model Realism to Purpose of the Model, *Computational and Mathematical Organization Theory*, Vol. 1, No. 1 (1995), pp. 57-71.

[5] Carley, K.M., "Designing Organizational Structures to Cope with Communication Breakdowns: A Simulation Model," *Industrial Crisis Quarterly*, vol. 5 (1991), pp. 19-57.

[6] Carley, K.M. and A. Newell, "The Nature of the Social Agent," *Journal of Mathematical Sociology*, Vol. 19, No. 4 (1994), pp. 221-262.

[7] Carley, K.M. and M.J. Prietula, "ACTS Theory: Extending the Model of Bounded Rationality," In Carley, K.M. and Prietula, M.J. (Eds.), *Computational Organization Theory*, Lawrence Erlbaum Associates, Hillsdale, NJ, 1994, pp. 55-88.

[8] Chaib-Draa, B., B. Moulin, R. Mandiau, and P. Millot, "Trends in Distributed Artificial Intelligence," *Artificial Intelligence Review*, Vol. 6 (1992), pp. 35-66.

[9] Chandrasekaran, B., "Natural and Social System Metaphors for Distributed Problem Solving: Introduction to the Issue," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC-11, No. 1 (1981), pp. 1-5.

[10] Cohen, K.J. and R.M. Cyert, "Simulation of Organizational Behavior," In March, J.G. (Ed.), *Handbook of Organizations*, Rand McNally, Chicago, IL, 1965, pp. 305-334.

[11] Crowston, K., "Modeling Cooperation in Organizations," In Masuch, M. and Warglien, M. (Eds.), *Artificial Intelligence in Organization and Management Theory*, North-Holland, Amsterdam, New York., 1992, pp. 215-234.

[12] Doran, J., "The Computational Approach to Knowledge, Communication and Structure in Multi-Actor Systems," In Gilbert, G.N. and Heath, C. (Eds.), *Social Action and Artificial Intelligence*, Gower, Aldershot, 1985, pp. 160-171.

[13] Drazin, R. and A.H. Van de Ven, "Alternative Forms of Fit in Contingency Theory," *Administrative Science Quarterly*, Vol. 30 (1985), pp. 514-539.

[14] Ferber, J. *Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence*, Addison Wesley, 1999.

[15] Fox, M.S., "An Organizational View of Distributed Systems," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC-11, No. 1 (1981), pp. 70-80.

[16] Gasser, L. and R.W. Hill, "Coordinated Problem Solvers," *Annual Review of Computer Science*, Vol. 4 (1990), pp. 203-253.

[17] Gasser, L. and M.N. Huhns, *Distributed Artificial Intelligence*, Volume II, Morgan

Kaufmann, San Mateo, CA., 1989.

[18] Griffiths, D.G. and B.K. Purohit, "Fundamentals of Distributed Artificial Intelligence," *British Telecommunication Technology Journal*, Vol. 9, No. 3 (1991), pp. 88-96.

[19] Hollenbeck, J.R., D.J. Sego, et. al. Team Decision-Making Accuracy Under Difficult Conditions: Construct Validation of Potential Manipulations Using TIDE^2 Simulation, In M. T. Brannick, E. Salas, and C. Prince (Eds.), *Team Performance, Assessment and Measurement: Theory, Method, and Applications*, Hillsdale, NJ: Erlbaum, 1997, Chapter 6.

[20] Hollenbeck, J.R., D.R. Ilgen, D.J. Sego, J. Hedlund, D.A. Major, and J. Phillips, "Multilevel Theory of Team Decision Making: Decision Performance in Teams Incorporating Distributed Expertise," *Journal of Applied Psychology*, Vol. 80, No. 2 (1995), pp. 292-316.

[21] Huberman, B.A., "The Value of Cooperation," In Masuch, M. and Warglien, M. (Eds.), *Artificial Intelligence in Organization and Management Theory*, North-Holland, Amsterdam, New York., 1992, pp. 235-243.

[22] Huhns, M.N., *Distributed Artificial Intelligence*, Morgan Kaufmann, San Mateo, CA., 1987.

[23] Jin, Y. and R.E. Levitt, "i-AGENTS: Modeling Organizational Problem Solving in Multi-Agent Teams," *Intelligent Systems in Accounting, Finance and Management*, Vol. 2 (1993), pp. 247-270.

[24] Kang, M., L. B. Waisel, and W. A. Wallace, "Team-Soar: A model for team decision making," In M. Prietula, K. Carley, and L. Gasser (Eds.), *Simulating Organizations: Computational Models of Institutions and Groups*, AAAI Press / MIT Press, Menlo Park, CA, 1998, pp. 23-45.

[25] Katzenbach, J.R. and D.K. Smith, "The Discipline of Teams," *Harvard Business Review*, Vol. 71, No. 2 (1993: March-April), pp. 111-121.

[26] Kraus, S., Negotiation and cooperation in multi-agent environments, *Artificial Intelligence*, Vol. 94 (1997), pp. 79-97.

[27] Laird, J., C.B. Congdon, E. Altmann, and R. Doorenbos, *Soar User's Manual: Version 6*, Electrical Engineering and Computer Science Department, University of Michigan and School of Computer Science, Carnegie Mellon University, 1993.

[28] Lin, Z., *Organizational Performance - Theory and Reality*, Doctoral Dissertation, School of Public Policy and Management, Carnegie Mellon University, 1993.

[29] Masuch, M., "Artificial Intelligence in Organization and Management Theory," In Masuch, M. and Warglien, M. (Eds.), *Artificial Intelligence in Organization and Management Theory*, North-Holland, Amsterdam, New York, 1992, pp. 1-19.

[30] Miller, C.E., "The Social Psychological Effects of Group Decision Rules," In Paulus, P.B. (Ed.), *Psychology of Group Influence*, 2nd ed., Erlbaum, Hillsdale, NJ, 1989, pp. 327-355.

[31] Sen, S., Multiagent systems: milestones and new horizons, *Trends in Cognitive Sciences*, Vol. 1, No. 9 (1997), pp. 334-340.

[32] Shaw, M.J., B. Harrow, and S. Herman, "Distributed Artificial Intelligence for Multi-Agent Problem Solving and Group Learning," *Proceedings of the Twenty-Fourth Annual Hawaii International Conference on System Sciences*, Vol. IV, IEEE Computer Society Press, Los Alamitos, CA, 1991, pp. 13-26.

[33] Weiss, G., *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, MIT Press, 1999.

[34] Yoshimura, S., K. Sasou, and K. Takano, "Team Behaviour Simulation Model - Proposal of the Conception," *Journal of Scientific and Industrial Research*, Vol. 53 (1994: July), pp. 574-578.

◈ 이 논문은 1999년 12월 18일 접수하여 1차 수정을 거쳐 2000년 8월 14일 게재확정 되었습니다.

◆ 저자소개 ◆

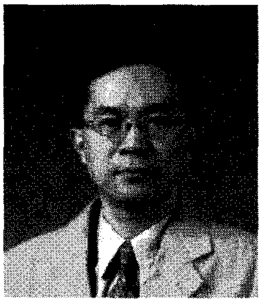

강민철 (Kang, Mincheol)
현재 계명대학교 경영정보학과에 전임강사로 재직중이다. 한국항공대학교 항공전자공학과에서 학사(1984), 미국 뉴욕주립대 (SUNY at Albany)에서 전산학 석사(1989), 미국 렌슬러 공과대학 (Rensselaer Polytechnic Institute)에서 공학박사(1996) 학위를 취득하였다. 삼보컴퓨터에서 연구원, 삼성SDS에서 경영컨설턴트로 근무한 바 있으며, 주요 연구분야는 Computational Organization Theory, Multi-Agents Systems (MAS) 및 Internet Business이다.