

On Cooperative Learning Teams for Multiagent Team Formation

Leen-Kiat Soh

Computer Science and Engineering
University of Nebraska-Lincoln
115 Ferguson Hall, Lincoln, NE 66588-0115
lksoh@cse.unl.edu

Abstract

In this paper, we propose a team formation methodology based on *cooperative learning teams*, adopted from the area of educational research. Cooperative learning is a type of learning where students work in teams and learn through team-based interactions. In education, research in assigning students to appropriate teams and enforcing fair assessment of student performance in a team have generated useful policies and rules. In our multiagent systems project, we use these policies and rules as the underlying framework to evaluate and form teams. We have built a system called I-MINDS as an infrastructure to support cooperative learning among remote and in-class students.

Introduction

In this paper, we propose a team formation methodology based on the cooperative learning model in education. Cooperative learning is basically an instructional strategy where students form small groups or teams and work together to maximize their own and each other's learning (Johnson et al. 1991). Cohen (1994) further added that students are expected to carry out their task without direct and immediate supervision of the instructor in the cooperative learning approach. Research in cooperative learning have provided guidelines on how to form good teams and how to evaluate teams to ensure fairness when rewarding team members. We treat such guidelines as useful policies and rules for multiagent team formation.

The cooperative learning model is actually very similar to team and coalition formation in multiagent systems (e.g., Cohen et al. 1997; Tambe 1997).

First, Johnson et al. (1991) distinguished two educational paradigms into old and new along six dimensions: knowledge, students, faculty purpose, relationships, context, and assumption.

- Traditionally, knowledge is transferred from faculty to students. Also, students are assumed to be passive learners, absorbing what is taught them. The instructor is re-

sponsible for judging and evaluating students—classifying and sorting them. In addition, relationships among students and between students and instructor are usually impersonal. The students learning context is also competitive and individualistic.

- The new paradigm, however, claims the opposite. Both the instructor and students jointly construct the knowledge. Students are active learners and discoverers of own knowledge. Students are assumed to be intelligent to do construct and transform knowledge. As an instructor, he or she is responsible for developing students' competencies and talents. Due to the close interaction, relationships among students are more personal, and so is that between faculty and students. Under the auspices of cooperative learning, students help each other learn and work in teams.

Based on these two paradigms, we see that the new paradigm fits the description of a multiagent system where agents are capable of generating and sharing information, and they work together to improve the system performance. With this differentiation, we see that the cooperative learning model promotes knowledge distributedness and exchange, and autonomy in the students (Bingman and Koutnik 1970), which correspond to our expectation of a multiagent system.

Second, Fellers (1996) explored the use of cooperative learning teams and identified five key elements of the cooperative learning model:

- *positive interdependence* where students depend on each other for goal satisfaction, rewards, resources, division of labor, roles, and so on.
- *face-to-face promotive interaction* where students need to communicate directly and to do so effectively
- *individual accountability* where each student must contribute to stay in a team and to obtain a good score
- *social skills* where students must learn about each other's habits, tendencies, opinions, and ways of doing things and adapt to the differences
- *cooperative process* where problem characteristics are defined, responsibilities are assigned, and steps are taken to solve the problem cooperatively

Equipped with these five elements, the cooperative learning model offers evaluation parameters that can be naturally designed and implemented in a multiagent system to oversee the formation and refinement of teams. These notions lead to procedures such as goal communication, rewards, role assignments, and member accountability for their individual performance (Springer et al. 1999). This, thus, lends itself to multiagent system research naturally.

For our research project, we have built a system called Intelligent Multiagent Infrastructure for Distributed Systems (I-MINDS) for Education where agents interact with the instructor and students in real-time (Liu et al. 2003a, 2003b). In I-MINDS, each student agent forms a “buddy” group for the student that it interacts with. The student agent forms and refines the “buddy group” based on its student’s activities and behavior, and those of the buddies.

In this following, we first briefly discuss some frameworks in team and coalition formation. Then we outline the cooperative learning model, focusing on the types of positive interdependence and self-efficacy evaluation. We also show how to apply the model to multiagent team formation. Subsequently, we describe I-MINDS and its current status and experimental results. Finally we conclude.

Background

In the area of multiagent systems, there has been much work in coalition formation, where agents typically work together to solve joint problems, based on utility of joining a coalition given possible costs and rewards. For example, Kahan and Rapoport (1984) proposed a coalition game based on the total utility that the member of the coalition can achieve by coordinating and acting together, assuming that information is complete. Zlotkin and Rosenschein (1994) described a coalition driven by task-oriented utilities. In a task-oriented domain, a coalition can coordinate by redistributing their tasks among themselves.

Further, researchers look into bounded rationality, incorporating computation costs into their reasoning process. For example, Sandholm and Lesser (1995) introduced a bounded rationality in which agents are guided by performance profiles and computation costs in their coalition formation process. Shehory et al. (1997) relaxed some of the restrictive assumptions of theoretical coalition formation algorithms for a real-world system. In their model, each agent has a vector of real non-negative capabilities. Each capability is a property of an agent that quantifies its ability to perform a specific type of action and is associated with an evaluation function. Shehory and Kraus (1998) further extended their work to incorporating negotiations, computational and communication costs.

There have also been investigations into optimal coalition formation. For example, Tohme and Sandholm (1999) studied coalition formation among self-interested agents that cannot make sidepayments, pinning the evaluation of a coalition solely on its utility. Sen and Dutta (2000) proposed an order-based genetic algorithm as a stochastic search process to identify the optimal coalition structure. The authors’ algorithm searches for an optimal coalition structure, which consists of all the agents in the environment grouped into one or more coalitions.

Also, there have been work in team formation in MAS, where a team is a set of agents having a shared objective and a shared mental state—without either, there is no unified activity and hence no team (Cohen et al. 1997). Work in this area has focused on joint intentions and the formation process. For example, Grosz and Kraus (1996) constrained an agent’s intentions by requiring that the agent not holding conflicting intentions with others sharing plans. Tambe (1997) emphasized the communication necessary for team formation—team members must flexibly communicate to achieve coherence. To do so, teams must be able to monitor performance and reorganize and reallocate resources to meet any contingencies.

Our team or coalition formation framework is inspired by educational research based on team-based learning (or cooperative learning) among students. Over the years, educational researchers have identified key models to motivate students in a team, to characterize or model each student in a team, to facilitate peer rating of team members, and to evaluate individual members. In terms of utility, each student receives points for his or her role in a team and each is expected to contribute. In terms of joint intentions, each student is motivated to cooperate with other team members as his or her grade is dependent on how well the team performs.

Cooperative Learning Team For Multiagent Team Formation

In this section, we propose and describe a multiagent team formation framework based on methodologies and paradigms in cooperative learning teams, studied in educational research. Particularly, our framework:

- addresses the joint intentions for joining and wanting to do well in a team (Fellers 1996),
- determines self-efficacy of members in a team environment (Fellers 1996),
- evaluates efficacy of teams (Fellers 1996),
- defines a peer rating mechanism (Martinazzi 1998a), and
- outlines an evaluation system for individual team members (Wilkins and Lawhead 2000).

Our multiagent system consists of a group of peer agents (i.e., student agents) and a *judge* (i.e., a teacher agent).

In terms of team formation, each peer agent is capable of (a) forming its team based on perceived goodness of other agents, (b) evaluating the observed goodness of a team member, and (c) revising its team by removing team members of poor utility to its team and adding other agents perceived to be good. Each peer agent thus uses the individual self-efficacy measures to form a team, uses the peer rating mechanism to evaluate its members, and uses the team-based efficacy to measure the team performance

The judge shoulders the responsibility of giving out rewards (or penalties) to each agent, based on the self-reported team performance and peer rating. With this, each agent will be kept in check—agents that do not do well as a team member will be scored low and vice versa. Thus, the judge agent impacts the team formation process by giving out penalties and rewards. It is also capable of comparing and ranking the performances among teams.

Finally, to encourage the agents to form teams, we use a set of positive interdependence. These relationships supply the joint intentions for the agents to work together in a team.

In the following framework, to apply the cooperative learning model to our multiagent team formation, a student is analogous to a student agent or a peer agent, and the instructor or teacher is analogous to the judge agent.

Positive Interdependence and Joint Intentions

In this section, we look to a set of positive interdependence, six of them shown in Table 1 (Fellers 1996), to form the joint intentions for the peer agents. In our framework, the judge agent announces the different types of interdependencies to the peer agents for each particular team-based activity. Each peer agent processes each announcement and motivates itself accordingly. For example, if the activity is of the positive goal interdependence type, then each agent will help each other to achieve their goals. Each peer agent will have its own goal for the particular agent that it interacts with. Here are some policies derived from the above interdependencies:

- *Positive Goal Policy*: Each peer agent monitors each member in its team and, upon identifying a member who is not reaching its goals, helps that agent achieve its goal. A joint intention exists.
- *Outside Enemy Policy*: In this policy, each peer agent can be a member of only one team to avoid conflict of interests. Each team member tries to optimize its utility or reward, and that contributes to the overall performance of the team. A joint intention exists.
- *Positive Reward Policy*: Each peer agent receives the same reward. Thus, it is possible for some students to choose to contribute minimally to the team and receive greater-than-minimal rewards. Thus, individual scoring

has to take place in this setup. A joint intention may only exist among a subset of the team.

- *Positive Resource Policy*: Peer agents have to share resources in order to accomplish global tasks. A joint intention exists.
- *Positive Task Policy*: Similar to the above policy, peer agents have to allocate sub-tasks among them to accomplish global tasks. A joint intention exists.
- *Positive Role Policy*: In this policy, peer agents assume different roles in their respective teams. Because of the different roles (team leader, team driver, etc.), there are different but consistent expectations and goals for each. A peer agent following this policy can be part of multiple teams simultaneously, playing a different role in each. A relaxed joint intention exists.

With the interdependence, our framework is able to motivate each peer agent to join a team and to do well. Next, we look at how a peer agent perceives the goodness of a potential team member.

PI Type	Description
Positive Goal	Exists when students perceive that they can achieve their learning goals, if and only if, all other members of their group also attain their goals. Success depends on all members reaching the goal.
Outside Enemy	Exists when teams (and not individual students) are placed in competition with each other.
Positive Reward	Exists when each team member receives the same reward for completing the assignment. A joint reward is given for successful teamwork. Every one is rewarded or no one is rewarded.
Positive Resource	Exists when each member has only a portion of the information, resources, or materials necessary for the task to be completed and members' resources have to be combined in order for the team to achieve its goal.
Positive Task	Exists when a division of labor is created so that the actions of one team member have to be completed if the next team member is to complete his or her responsibilities
Positive Role	Exists when each member is assigned complementary and interconnected roles that specify responsibilities that the group needs in order to complete a joint task

Table 1. Types of positive interdependence (PI) (Feller 1996) used in our framework.

Self-Efficacy and Perceived Goodness

A peer agent needs to form a team for the student that it serves after finding out what is expected of teamwork. To form a team, the peer agent needs to obtain the perceived goodness of a potential team member. In our framework, this is done based on two sets of measurements: (a) self-efficacy scores (Fellers 1996) and (b) profile of previous teamwork activities that the peer agent keeps of other agents. Note that the actual scores that each peer agent has obtained is not available publicly—just like student grades are not—and thus a peer agent can only perceive the goodness of the potential team members. The profile of previous teamwork activities is based on peer rating (Section 3.4. later). The self-efficacy scores are based on a set of questionnaire before the teamwork activities start (Table 2).

Applied to a multiagent system in our framework, the self-efficacy scores are measured empirically by each agent:

- *Level of experience working in teams* is proportional to the number of teams that the agent has been a member of,
- *Degree of positive experiences* is proportional to the number of successful teams that the agent has been a member of where success is based on team-based efficacy, peer rating, and rewards (both individual and team),
- *Desire of working alone* is proportional to the ratio of individual rewards over team rewards, to the average size of teams that the agent is a member of, and to the number of failed teams that the agent has been a member of,
- *Motivation to make teamwork successful* is proportional to the ratio of the reward over the effort that the agent contributes to its team, where the effort can be modeled as the computation, communication, and resource costs,
- *Expectation of team success* is proportional to the number of failed activities and the peer rating of the team members,
- *Expected significance of personal contribution* is proportional to the evaluation of the judge of the peer agent and the given rewards, and to the ratio of the actual costs over the expected costs as part of the team,
- *Expected individual contributions* are based on the profiling by the agent of its team members, and
- *Concern about team performance* is based on the perceived unfairness by a peer agent of the way reward is distributed among the team members. For example, given the positive reward policy, a peer agent knows exactly what reward each team member receives and if it rates some team members very poorly, and some very highly, then it perceives unfairness if everybody receives the same reward.

Note that in our framework, the above self-efficacy score can be updated (dynamically and automatically by the agents) from time to time as students change their be-

havior. Also, it is possible for the self-efficacy to be exchanged between peer agents that are interested in working as a team.

Questions
Level of experience working in teams (Seldom → Often)
I like to participate in teams (Strongly Disagree → Strongly Agree)
I have had positive experiences thus far working in teams in the environment (Strongly Disagree → Strongly Agree)
I would rather work in teams than on my own (Strongly Disagree → Strongly Agree)
How motivated are you to make this team successful? (Not Very → Very)
How successfully do you expect the team to accomplish the required outcomes? (Not Very → Very)
How significant do you expect your personal contribution will be to the team's outcomes? (Not Very → Very)
How equal do you feel the individual contributions of the team members will be? (Unequal Participation → Equal Participation)
Are you concerned about having a majority of the points being tied to team performance rather than individual performance? (Not Very → Very)

Table 2. Self-efficacy questionnaire, modified from (Fellers 1996), to help determine the perceived goodness of a potential team member.

Note also that depending on the different interdependence policies, a peer agent may not be allowed to recruit only the most motivated peers to join its team. For example, in the positive role policy, more motivated peers will become team leaders and usually that multiple leaders on a team is not a good option. In the positive resource policy, for example, a peer agent must identify the peer agents that can provide the necessary resources first, and then consider the motivations second.

In the above, we touch upon the success of a team and team-based efficacy, which is based on the peer agent's perception of how well the team works together and the final rewards. We will talk about the perception of team performance in the next section.

Team-Based Efficacy

According to Fellers (1996), at the end of a teamwork process, each team member can evaluate how well the team works together based on a set of questions (Table 3). Note that this team-based efficacy is used in two ways in our framework. First, a peer agent uses it to help revise its self-efficacy to be part of a team (Section 3.2). Second,

the judge agent uses it to evaluate the performance of a team. Note that the idea of cooperative learning teams in education hinges upon not only the final outcome of the team but also on how each team member participate in the teamwork activities. For example, suppose that there is a team with 2 students. One of them slacks off and does not contribute while the other works relentlessly to get the programming assignment done. In terms of the final outcome, the team should receive full points. However, in terms of working as a team, the team has failed. The questionnaire in Table 3 is designed to address this.

Questions
Over the course of the teamwork, our team was very successful in accomplishing the outcomes required of us. (Strongly Disagree → Strongly Agree)
Over the course of the teamwork, our team was very successful in working together as a team. (Strongly Disagree → Strongly Agree)
Over the course of the teamwork, we had little problem with conflict within our team. (Strongly Disagree → Strongly Agree)
As the teamwork process draws to a close, I feel more comfortable having a majority of my points based on my team's performance. (Strongly Disagree → Strongly Agree)
I believe that working on the team has been a valuable learning experience for me. (Strongly Disagree → Strongly Agree)
I would like to participate as a team member in the future. (Strongly Disagree → Strongly Agree)
Cooperative teams should continue to be a required element of this environment. (Strongly Disagree → Strongly Agree)
What percentage of the work done by your team was completed : Working together as a team: Working individually:

Table 3. Team-based efficacy questionnaire, modified from (Fellers 1996), to help determine the perceived success of a team.

Applied to a multiagent system in our framework, the team-based efficacy scores can also be measured empirically by each agent. Due to the page limit and the similarity with the measurements discussed in Section 3.2, we will not elaborate further and will let the readers make the analogy. However, one point worth noting is how the judge agent makes use of this team-based efficacy information. To measure whether a team has worked well, in our framework, the judge agent computes a weighted score of each efficacy question. Each member's score is tallied and multiplied with a peer-based weight (Section 3.4). The score of a member that is highly rated by its peer will be weighted more, and vice versa.

Given this setup, each peer agent now is able to monitor and evaluate its own motivation, its contribution to a team, and the goodness of teamwork. The judge agent also has a mechanism to evaluate how well team members have work in a team. Next, we look at how each agent rates its peers as team members.

Peer Rating

Martinazzi (1998a, 1998b) proposed a peer evaluation model for student learning team (Table 4). It is based on a set of questions rating a peer on his or her ability to function as a member of a team, using the following scale:

- 5 Always/Positive Contributor
- 4 Most of the time
- 3 Sometimes
- 2 On occasion but not consistently
- 1 Rarely/Negative contributor

Questions
Has a "sharing" attitude toward team members?
Has a positive attitude towards team?
Is this team member truly earning the reward being received (assuming positive reward interdependence)?
Willing to help other team members anytime?
Eagerly accepts and shares all team responsibilities?
How well does team member attempt to accomplish team's mission and goals?
Did team member participate in establishing team's mission and goals?
Participates in the team's discussions?
Level of contribution to the team?

Table 4. Peer rating questionnaire, modified from (Fellers 1996), to help evaluate individuals and how well a team has worked together.

Basically, Table 4 identifies the parameters that an agent should keep track of its team members to profile them (Section 3.2). Once again, the profile can be automatically updated empirically by an agent:

- *Sharing attitude* is proportional to the amount of result, information, data, and knowledge shared by the peer with the team members,
- *Positive attitude* is specific to what type of positive interdependence that is involved and proportional to sharing attitude, willingness to help, eagerness, and others listed in this section,
- *Reward fairness* is proportional to the number of resources shared, sub-tasks accomplished, and roles successfully filled by the peer,
- *Willingness to help* is based on the number of voluntary acts such as initiating messages and unsolicited sharing of resources,

- *Eagerness* is based on the responsiveness of the peer to a team request such as the number of minutes between a request and a response,
- *Tenacity* measures how well the peer attempts to solve a problem especially when the team's goals and tasks change, based on the activities logged by and recorded by the peer, after a change in the task specifications, and
- *Leadership* measures the degree of the peer participating in establishing team's mission and goals and is proportional to the number of plans and the utility of the plans that the peer recommends to the team.

There are also other measures such as timeliness in accomplishing sub-tasks and in sharing resources, reliability in carrying out promises or commitments, the number of meaningful messages sent to the team members, and so on. With these measurements, an agent can evaluate each of its team members and submit the rating to the judge agent to assign the final team and individual rewards.

Evaluating Individuals

In this section, we outline an evaluation scheme for each individual team member, based on the peer rating system discussed in Section 3.4 and (Wilkins and Lawhead 2000). In our framework, we propose to use the following assessment strategies (Wilkins and Lawhead 2000):

- *matching* such as "director, opinion leader, scorekeeper, consumer" (House 1988),
- *fill-in-the-blank* such as "The strengths of X are ____." "____ was the Team Leader." "____ worked the most hours." "____ does the best quality work." "____ dominated the group." "____ nurtured the group." "____ never got a chance to express his/her opinion." (Wilkins and Lawhead 2000).
- *short answer* such as "The greatest obstacle overcome by the team was ____." "Our team was often sidetracked by ____." "Who would you not want to work with? Why?" "What was the most difficult problem in dealing with the team?" "During the teamwork activity, the five best things this team did were: ____." (Wilkins and Lawhead 2000).

The underlying design principle of these surveys is to check for consistency in the response among the team members. Usually, when a member has not participated in the teamwork activities, his or her response will be observably different from other team members who have participated. Thus, this effectively serves as a self-rating mechanism. Combining the peer rating (Section 3.4) with this, the judge agent is able to assign the most points to the member who scores the highest in its peer ratings and who produces a consistent response to the Wilkins and Lawhead (2000) survey.

Note that in our current framework, we have not made the application of the survey to our multiagent team forma-

tion process as we are still not clear yet on how to embed automated interpretation of open-ended responses (e.g., short answer) in an agent. However, natural language processing and pattern matching are plausible solutions to check for consistency in the open-ended responses, and in view of that, we believe that the above survey can be designed and implemented in autonomous agents.

I-MINDS

In this section, we describe our research project called I-MINDS. The long-term goals of the project are two-fold: (1) developing a distributed computing infrastructure specifically for education, addressing the interaction issues in real-time classroom scenarios, distance learning, and so on, and (2) developing an intelligent multiagent information system, built atop the infrastructure developed in (1), that is able to monitor the activities, recognize patterns, and interact with students and instructors alike to improve the quality of teaching and learning. Thus, our objectives include real-time data gathering, information dissemination, and decision making in a classroom setting, utilizing the advancements in both hardware and software.

Most agent-based education systems use software agents without fully utilizing the flexibility (or intelligence) of an agent such as reactivity, pro-activeness, and social ability (Wooldridge and Jennings 1995). For example, though some intelligent tutoring systems such as ANDES (Gertner and VanLehn 2000) have achieved some level of success in classrooms, criticisms of the current state of tutoring systems (Graesser et al. 2001) stem from the lack of sufficient intelligence in the tutoring system necessary to monitor and detect a student's pedagogical behavior. Students may simply keep guessing until they find an action that gets positive feedback and thus learn to do the right thing for the wrong reasons, and the tutoring system will never detect such shallow learning (Aleven et al. 1999).

Moreover, most agent-based education systems are simply a group of non-collaborative, individual agents. Thus, our objective is to also exploit "multiagent system intelligence" to help the transfer of information towards helping teachers teach better and students learn better.

Design and Implementation

The development and implementation of I-MINDS is a unique and innovative approach to computer-aided instruction and learning because of the incorporation of interacting agents. It is based on the team formation framework discussed in Section 3. I-MINDS includes both student agents (peer agents) and a teacher agent (the judge agent). Currently, we have built the system based on two tasks: asking good questions in the classroom and participating in real-time forums with team members (buddies). Thus, the

perceived goodness and the success of a team, and other parameters delineated in Section 3, are based on how well each team member performs the above two tasks. The type of positive interdependence used for the two tasks is positive reward.

The student agents collaborate to autonomously form student *buddy groups* and provide services to their respective student users. The teacher agents monitor classroom activities and analyze student behavior to help the teacher respond to questions and to assess student comprehension and interest. These agents are designed to assess their own performance based on the observed impact of the buddy groups and the agent-initiated interventions, such as question ranking, on student learning. It is this buddy group feature that uses multiagent intelligence and team formation strategies to empower I-MINDS as an infrastructure for educational applications. For details on I-MINDS, please refer to (Liu et al. 2003a, 2003b).

The I-MINDS prototyping process was initiated in September 2002 using a National Center for Information Technology in Education (NCITE) Seed Grant, which allowed us to build a “proof-of-concept” software package and conduct preliminary experiments to evaluate the technical correctness and educational feasibility of I-MINDS. The multiagent prototype was implemented in Java, along with a host of multimedia processing and interface technologies (Table 5) to support the agents.

Technology
Live audio (2-way)
Live video (1-way)
Superimposition of handwriting (MIMIOS) on lecture screen, superimposition of text on archived lecture pages
Automated archival
Multicast/broadcast
Multicharacter forum (chatroom) (colors and fonts)
Multicharacter e-whiteboard (colors and fonts, exclusive tokens)
Annotation and asynchronous review of archived lectures
Rich and flexible control of system features

Table 5. Technology implemented for I-MINDS.

Proof-of-Concept Experiments

Here we briefly report on our proof-of-concept experiments. For details on the experiments, please refer to (Soh et al. 2004).

To determine the potential impact of I-MINDS on student learning, a pilot study was conducted in May 2003 where the tool was used by subjects in a controlled experiment to assess what impact it had on student learning of Global Information Systems (GIS) content.

There were two groups: I-MINDS group and control group. Each group took two days of lectures, separately. On Day 1, subjects in both groups completed a 109-point pretest of the content that was to be taught during the two sessions. At the conclusion of the class on Day 1 for both groups, the subset of 60 items that related to the content of that class was included on the posttest. After the Day 2 instruction, the subset of 49 items that related to the content of that class constituted the posttest. Subjects in the control group learned the identical content during each of the two sessions, as did subjects in the I-MINDS group. The difference was that the control group students were in the same room as the instructor. Their class was taught in a very traditional manner with the professor using PowerPoint slides identical to those used for the I-MINDS group to teach the content.

In general, the subjects in the I-MINDS group commented about how they had considered ways to use the software to assist their learning. Also, *the amount that the I-MINDS group improved from the pretest to the posttest was nearly twice that of the control group*. This result was very promising as were comments from the subjects in the I-MINDS group related to their comfort level in using the tool. Comments from the instructor who used I-MINDS in teaching both of the content lessons were also encouraging. He indicated that the tool could enhance distance learning, especially by making it possible for building an archive of information that could be accessed “on-demand” by students. The instructor also noted that questions asked of him via I-MINDS tended to be of a higher quality, reflect a deeper understanding, and demand a richer response than those questions posed during the control sessions (Soh et al. 2004).

Conclusions

We have proposed and described a framework for multi-agent team formation based on the cooperative learning model from the educational research. Within this framework, peer agents (or team members) will be able to empirically perceive team performance and peer-rate their team members. Also, a *judge* can be designed to monitor, observe, and evaluate the performance of each team and assign fair grades (or rewards) to each individual team member. The combination of these two features motivate the team members to refine their teams.

We have implemented part of the framework in our I-MINDS project where the teacher agents assume the role of the judge, and the student agents are team members. Currently, in our first prototype, the performance of a team is based on the quality of the questions asked and its forum discussion. We are now implementing other team formation strategies discussed in this paper. For future work, we

will investigate how the various strategies impact the team formation process among the student agents and how such multiagent intelligence impact teaching and learning in a classroom.

Acknowledgment

The author would like to acknowledge the I-MINDS members: Xuli Liu, Xuesong Zhang, and Hong Jiang. I-MINDS is supported in part by a National Center for Information Technology in Education (NCITE) seed grant.

References

- Aleven, V., K. R. Koedinger, and K. Cross 1999. Tutoring Answer Explanation Fosters Learning with Understanding, in S. P. Lajoie and M. Vivet (eds.) *Artificial Intelligence in Education*, 199-206, Amsterdam: IOS.
- Bingman, R. and P. Koutnik 1970. A Small Group Study Approach for Biology Based Inquiry Objectives, *The American Biology Teacher*, 32:11-16.
- Cohen, E. G. 1994. Restructuring the Classroom: Conditions for Productive Small Groups, *Rev. Educational Research*, 64(1):1-35.
- Cohen, P. R., H. J. Levesque, and I. A. Smith 1997. On Team Formation, in G. Holmstrom-Hintikka and R. Tuomela (eds.). (1997). *Contemporary Action Theory, Volume 2: Social Action*, pp. 87-114, Dordrecht: Kluwer Academic Publishers.
- Fellers, J. W. 1996. Teaching Teamwork: Exploring the Use of Cooperative Learning Teams in Information Systems Education, *The DATA BASE for Advances in Information Systems*, 27(2):44-60.
- Gertner, A. S. and K. VanLehn 2000. ANDES: A Coached Problem-Solving Environment for Physics, in *Intelligent Tutoring Systems: 5th Int. Conf., ITS 2000*, eds. G. Gauthier, C. Frasson, and K. VanLehn, 133-142.
- Graesser, A. C., K. VanLehn, C. P. Rosé, P. W. Jordan, and D. Harter 2001. Intelligent Tutoring Systems with Conversational Dialogue, *AI Magazine*, 22(4):39-51.
- Grosz, B. and S. Kraus 1996. Collaborative Plans for Complex Group Action, *Artificial Intelligence*, 86(2):269-357.
- House, R. S. 1988. *The Human Side of Project Management*, Addison-Wesley, Reading.
- Johnson, D. W., R. T. Johnson, and K. A. Smith 1991. Cooperative Learning: Increasing College Faculty Instructional Productivity, *ASHE-ERIC Higher Education Report No. 4*, Washington, D.C.: The George Washington Univ., School Edu. & Human Development.
- Kahan, J. P. and A. Rapoport 1984. *Theories of Coalition Formation*, Lawrence Erlbaum.
- Liu, X., X. Zhang, J. Al-Jaroodi, P. Vemuri, H. Jiang, and L.-K. Soh 2003a. I-MINDS: An Application of Multiagent System Intelligence to On-Line Education, in *Proc. IEEE Int. Conf. on Systems, Man, and Cybernetics*, Washington, D.C., October, 4864-4871.
- Liu, X., X. Zhang, L.-K. Soh, J. Al-Jaroodi, and H. Jiang 2003b. A Distributed, Multiagent Infrastructure for Real-Time, Virtual Classrooms, accepted to *International Conference on Computers in Education (ICCE2003)*, Hong Kong, China.
- Martinazzi, R. 1998a. Design and Development of a Peer Evaluation instrument for "Student Learning Teams", *Proc of 1998 FIE Conference*, 784-789.
- Martinazzi, R. 1998b. Student Learning Teams in Engineering Economic, *Proc. 1998 ASEE Annual Conference*, Seattle, WA, June 18-July 1, 1998.
- Sandholm, T. W. and V. R. Lesser 1995. Coalition Formation amongst Bounded Rational Agents, *Proc. IJCAI'95*, Montreal, Canada, 662-669.
- Sen, S. and P. S. Dutta 2000. Searching for Optimal Coalition Structures, *Proc. ICMAS'00*, Boston, MA, July, 286-292.
- Shehory, O. and S Kraus 1998. Methods for Task Allocation via Agent Coalition Formation, *Artificial Intelligence*, 101(1-2):165-200.
- Shehory, O. M., K. Sycara, and S. Jha 1997. Multi-Agent Coordination through Coalition Formation, *Proc. ATAL'97*, Providence, RI.
- Soh, L.-K., H. Jiang, and C. Ansorge 2004. Agent-Based Cooperative Learning: A Proof-of-Concept Experiment, to appear in *SIGCSE'2004*, Norfolk, Virginia, April.
- Springer, L., M. E. Stane, and S. S. Donovan 1999. Effects of Small-Group Learning on Undergraduates in Science, Mathematics, Engineering, and Technology: A Meta-Analysis, *Rev. Educational Research*, 69(1):21-51.
- Tambe, M. 1997. Towards Flexible Teamwork, *J. Artificial Intelligence Research*, 16:389-423.
- Tohme, F. and T. Sandholm 1999. Coalition Formation Process with Belief Revision among Bounded-Rational Self-Interested Agents, *J. Logic & Computation*, 9(6):793-815.
- Wilkins, D. E. and p. B. Lawhead 2000. Evaluating Individuals in Team Projects, *SIGCSE 2000*, March, Austin, Texas, 172-175.
- Wooldridge, M. and N. R. Jennings 1995. Intelligent Agents: Theory and Practice, *The Knowledge Engineering Review*, 10(2):115-152.
- Zlotkin, G. and J. S. Rosenschein 1994. Coalition, Cryptography and Stability: Mechanisms for Coalition Formation in Task Oriented Domains, *Proc. AAAI'94*, Seattle, WA, July, 432-437.