

The RoboCup Synthetic Agent Challenge 97

Hiroaki Kitano*

Sony Computer Science Laboratory
kitano@csl.sony.co.jp

Manuela Veloso

Carnegie Mellon University
mmv@cs.cmu.edu

Hitoshi Matsubara

ElectroTechnical Laboratory
matsubar@etl.go.jp

Milind Tambe

ISI/USC

tambe@isi.edu

Silvia Coradeschi

Linköping University
silco@ida.liu.se

Itsuki Noda

ElectroTechnical Laboratory
noda@etl.go.jp

Peter Stone

Carnegie Mellon University
pstone@cs.cmu.edu

Eiichi Osawa

Sony Computer Science Laboratory
osawa@csl.sony.co.jp

Minoru Asada

Osaka University
asada@mech.eng.osaka-u.ac.jp

Abstract

RoboCup Challenge offers a set of challenges for intelligent agent researchers using a friendly competition in a dynamic, real-time, multi-agent domain. While RoboCup in general envisions longer range challenges over the next few decades, RoboCup Challenge presents three specific challenges for the next two years: (i) learning of individual agents and teams; (ii) multi-agent team planning and plan-execution in service of teamwork; and (iii) opponent modeling. RoboCup Challenge provides a novel opportunity for machine learning, planning, and multi-agent researchers — it not only supplies a concrete domain to evaluate their techniques, but also challenges researchers to evolve these techniques to face key constraints fundamental to this domain: real-time, uncertainty, and teamwork.

1 Introduction

RoboCup (The World Cup Robot Soccer) is an attempt to promote AI and robotics research by providing a common task, Soccer, for evaluation of various theories, algorithms, and agent architectures [Kitano, et al., 1995; Kitano et al., 1997a; 1997b]. Defining a standard problem in which various approaches can be compared and progress can be measured provides fertile grounds for engineering research. Computer chess has been a symbolic example of the standard challenge problems. A salient feature of computer chess is that progress can be measured via actual games against human players.

For an agent (a physical robot or a synthetic agent) to play soccer reasonably well, a wide range of technologies need to be integrated and a number of technical breakthroughs must be made. The range of technologies spans both AI and robotics research, such as design principles of autonomous agents, multi-agent collaboration, strategy acquisition, real-time reasoning and planning, intel-

ligent robotics, sensor-fusion, and so forth. RoboCup consists of three competition tracks:

Real Robot League: Using physical robots to play soccer games.

Software Agent League: Using software or synthetic agents to play soccer games on an official soccer server over the network.

Expert Skill Competition: Competition of robots which have special skills, but are not able to play a game.

RoboCup offers a software platform that forms the basis of the software or synthetic agent league. The goal is to enable a wider range of research in synthetic (or “virtual reality”) environments, that are today proving to be critical in training, entertainment, and education [Tambe et al., 1995]. The software agent league also promotes research on network-based multi-agent interactions, computer graphics, and physically realistic animation — a set of technologies which potentially promotes advanced use of internet.

2 Technical Challenges in RoboCup

RoboCup offers significant long term challenges, which will take a few decades to meet. However, due to the clarity of the final target, several subgoals can be derived, which define mid term and short term challenges. One of the major reasons why RoboCup is attractive to so many researchers is that it requires the integration of a broad range of technologies into a team of complete agents, as opposed to a task-specific functional module. The long term research issues are too broad to compile as a list of specific items. Nevertheless, the challenges involve a broad range of technological issues ranging from the development of physical components, such as high performance batteries and motors, to highly intelligent real time perception and control software.

The mid term technical challenges, which are the target for the next 10 years, can be made more concrete, and a partial list of specific topics can be compiled. Following is a partial list of research areas involved in RoboCup, mainly targeted for the mid term time span: (1) agent architecture in general, (2) combining reactive approach and modeling/planning approach, (3) real-time recognition, planning, and reasoning, (4) reasoning and action

*Corresponding Author: Hiroaki Kitano, Sony Computer Science Laboratory, 3-14-13 Higashi-Gotanda, Shinagawa, Tokyo 141, Japan. kitano@csl.sony.co.jp RoboCup Home Page: <http://www.robocup.org/RoboCup>. RoboCup Mailing List: robocup@csl.sony.co.jp

in dynamics environment, (5) sensor fusion, (6) multi-agent systems in general, (7) behavior learning for complex tasks, (8) strategy acquisition, and (9) cognitive modeling in general.

In addition to these technologies, providing a network-based soccer server with high quality 3D graphics capabilities requires advancement of technologies for the real time animation of simulated soccer players and network-based interactive multi-user server system. These are key technologies for network-based services in the coming years.

The RoboCup Challenge shall be understood in the context of larger and longer range challenges, rather than as a one-shot challenge. Thus, we wish to provide a series of short term challenges, which naturally leads to the accomplishment of the mid term and long term challenges. RoboCup challenge is organized into three major classes; (1) Synthetic Agent Challenge, (2) Physical Agent Challenge, and (3) Infrastructure Challenge. The RoboCup Synthetic Agent Challenge deal with technologies which can be developed using software simulator, which is described in this paper. The RoboCup Physical Agent Challenge intends to promote research using real robot, and thus requires longer-time frame for each challenge to be accomplished. Details of this challenge is described in [Asada et al., 1997], and carried out together with the RoboCup Synthetic Agent Challenge but in more moderate timeframe. The Infrastructure Challenge will be presented to facilitate research to establish infrastructure aspect of RoboCup, AI, and robotics in general. Such challenge includes education programs, common robot platforms and components standard, automatic commentary systems and intelligent studio systems for RoboCup games.

3 Overview of The RoboCup Synthetic Agent Challenge

For the RoboCup Synthetic Agent Challenge 97, we offer three specific targets, critical not only for RoboCup but also for general AI research. These challenges will specifically deal with the software agent league, rather than the real robot league.

The fundamental issue for researchers who wish to build a team for RoboCup is to design a multiagent system that behaves in real-time, performing reasonable goal-directed behaviors. Goals and situations change dynamically and in real-time. Because the state-space of the soccer game is prohibitively large for anyone to hand-code all possible situations and agent behaviors, it is essential that agents learn to play the game strategically. The research issues in this aspect of the challenge involve:

(1) machine learning in a multiagent, collaborative and adversarial environment, (2) multiagent architectures, enabling real-time multiagent planning and plan execution in service of teamwork, and (3) opponent modelling.

Therefore, we propose the following three challenges as areas of concentration for the RoboCup Synthetic Agent Challenge 97:

- Learning challenge

- Teamwork challenge
- Opponent modeling challenge

Evaluating how well competing teams meet these challenges in RoboCup is clearly difficult. If the task is to provide the fastest optimization algorithm for a certain problem, or to prove a certain theorem, the criteria are evident. However, in RoboCup, while there may be a simple test set to examine basic skills, it is not generally possible to evaluate the goodness of a team until it actually plays a game. Therefore, a standard, highly skilled team of opponents is useful to set an absolute basis for such evaluation. We hope to use hand-coded teams, possibly with highly domain-specific coordination, to provide such a team of opponents. Indeed, in a series of preliminary competitions such as PreRoboCup-96 held at the IROS-96 conference, and several other local competitions, teams with well-designed hand-coded behaviors, but without learning and planning capabilities, have performed better than teams with learning and planning schemes. Of course, these hand-coded teams enjoyed the advantage of very low game complexities in initial stages of RoboCup — increasingly complex team behaviors, tactics and strategies will necessitate agents to face up to the challenges of learning, teamwork and opponent modeling.

Therefore, responses to this challenge will be evaluated based on (1) their performance against some standard hand-coded teams as well as other teams submitted as part of the competition; (2) behaviors where task specific constraints are imposed, such as probabilistic occurrence of unexpected events, (3) a set of task specific sequences, and (4) novelty and technical soundness of the approach.

4 The RoboCup Learning Challenge

4.1 Objectives

The objectives of the RoboCup Learning Challenge is to solicit comprehensive learning schemes applicable to the learning of multi-agent systems which need to adapt to the situation, and to evaluate the merits and demerits of proposed approaches using standard tasks.

Learning is an essential aspect of intelligent systems. In the RoboCup learning challenge, the task is to create a learning and training method for a group of agents. The learning opportunities in this domain can be broken down into several types:

1. Off-line skill learning by individual agents;
2. Off-line collaborative learning by teams of agents;
3. On-line skill and collaborative learning;
4. On-line adversarial learning.

The distinction between off-line and on-line learning is particularly important in this domain since games last for only 20 minutes. Thus on-line techniques, particularly if they are to learn concepts that are specific to an individual game, must generalize very quickly. For example, if a team is to learn to alter its behavior against an individual opponent, the team had better be able to improve its performance before the game is over and a

new opponent appears. Such distinctions in learning can be applied to a broad range of multi-agent systems which involve learning capabilities.

4.2 Technical Issues

Technical issues anticipated in meeting this challenge are the development of novel learning schemes which can effectively train individual agents and teams of agents in both off-line and on-line methods. One example of possible learning scheme for meeting this challenge is as follows:

Off-line skill learning by individual agents: learning to intercept the ball or learning to kick the ball with the appropriate power when passing.

Since such skills are challenging to hand-code, learning can be useful during a skill development phase. However, since the skills are invariant from game to game, there is no need to relearn them at the beginning of each new game [Stone and Veloso, 1997].

Off-line collaborative learning by teams of agents: learning to pass and receive the ball.

This type of skill is qualitatively different from the individual skills in that the behaviors of multiple agents must be coordinated. A "good" pass is only good if it is appropriate for the receivers receiving action, and vice versa. For example, if the passer passes the ball to the receiver's left, then the receiver must at the same time move to the left in order to successfully complete a pass. As above, such coordination can carry over from game to game, thus allowing off-line learning techniques to be used [Stone and Veloso, 1997].

On-line skill and collaborative learning: learning to play positions.

Although off-line learning methods can be useful in the above cases, there may also be advantages to learning incrementally as well. For example, particular aspects of an opposing teams' behavior may render a fixed passing or shooting behavior ineffective. In that case, the ability to adaptively change collaborative or individual behaviors during the course of a game, could contribute to a team's success.

At a higher level, team issues such as role (position) playing on the field might be best handled with adaptive techniques. Against one opponent it might be best to use 3 defenders and 8 forwards; whereas another opponent might warrant a different configuration of players on the field. The best teams should have the ability to change configurations in response to events that occur during the course of a game.

On-line adversarial learning: learning to react to predicted opponent actions.

If a player can identify patterns in the opponents' behaviors, it should be able to proactively counteract them. For example, if the opponent's player number 4 always passes to its teammate number 6, then player 6 should always be guarded when player 4 gets the ball.

4.3 Evaluation

For challenge responses that address the machine learning issue (particularly the on-line learning issue), evaluation should be both against the publicly available teams and against at least one previously unseen team.

First, teams will play games against other teams and publicly available teams under normal circumstances. This evaluates the team's general performance. This involves both AI-based and non-AI based teams.

Next, teams will play a set of defined benchmarks. For example, after fixing their programs, challengers must play a part of the game, starting from the defined player positions, with the movement of the opponents pre-defined, but not disclosed to the challengers. After several sequences of the game, the performance will be evaluated to see if it was able to improve with experience. The movement of the opponents are not coded using absolute coordinate positions, but as a set of algorithms which generates motion sequences. The opponent algorithms will be provided by the organizers of the challenge by withholding at least one successful team from being publicly accessible.

Other benchmarks which will clearly evaluate learning performance will be announced after discussion with challenge participants.

5 The RoboCup Teamwork Challenge

5.1 Objectives

The RoboCup Teamwork Challenge addresses issues of real-time planning, re-planning, and execution of multi-agent teamwork in a dynamic adversarial environment. Major issues of interest in this specific challenge for the 97-99 period are architectures for real-time planning and plan execution in a team context (essential for teamwork in RoboCup). In addition, generality of the architecture for non-RoboCup applications will be an important factor.

Teamwork in complex, dynamic multi-agent domains such as Soccer mandates highly flexible coordination and communication to surmount the uncertainties, e.g., dynamic changes in team's goals, team members' unexpected inability to fulfil responsibilities, or unexpected discovery of opportunities. Unfortunately, implemented multi-agent systems often rely on preplanned, domain-specific coordination that fails to provide such flexibility. First, it is difficult to anticipate and preplan for all possible coordination failures; particularly in scaling up to complex situations. Thus, it is not robust enough for dynamic tasks, such as soccer games. Second, given domain specificity, reusability suffers. Furthermore, planning coordination on the fly is difficult, particularly, in domains with so many possible actions and such large state spaces. Indeed, typical planners need significantly longer to find even a single valid plan. The dynamics of the domain caused by the unpredictable opponent actions make the situation considerably more difficult.

A fundamental reason for these teamwork limitations is the current agent architectures. Architectures such as Soar [Newell, 1990], RAP [Firby, 1987], IRMA [Pollack,

1991], and BB1 [Hayes-Roth et al., 1995] facilitate an individual agent’s flexible behaviors via mechanisms such as commitments and reactive plans. However, teamwork is more than a simple union of such flexible individual behaviors, even if coordinated. A now well-known example (originally from [Cohen and Levesque, 1991]) is ordinary traffic, which even though simultaneous and coordinated by traffic signs, is not teamwork. Indeed, theories of teamwork point to novel mental constructs as underlying teamwork, such as team goals, team plans, mutual beliefs, and joint commitments [Grosz, 1996; Cohen and Levesque, 1991], lacking in current agent architectures. In particular, team goals, team plans or mutual beliefs are not explicitly represented; furthermore, concepts of team commitments are absent. Thus, agents cannot explicitly represent and reason about their team goals and plans; nor flexibly communicate/coordinate when unanticipated events occur. For instance, an agent cannot itself reason about its coordination responsibilities when it privately realizes that the team’s current plan is unachievable — e.g., that in the best interest of the team, it should inform its teammates. Instead, agents must rely on domain-specific coordination plans that address such contingencies on a case-by-case basis.

The basic architectural issue in the teamwork challenge is then to construct architectures that can support planning of team activities, and more importantly execution of generated team plans. Such planning and plan execution may be accomplished via a two tiered architecture, but the entire system must operate in real-time. In RoboCup Soccer Server, sensing will be done in every 300 to 500 milli-seconds, and action command can be dispatched every 100 milli-second. Situation changes at milli-second order, thus planning, re-planning, and execution of plans must be done in real-time.

5.2 Technical Issues

We present a key set of issues that arise assuming our particular two tiered planning and plan-execution approach to teamwork. Of course, those who approach the problem from different perspective may have different issues, and the issues may change depending on the type of architecture employed.

The following is the envisioned teamwork challenge in this domain: (i) a team deliberately accumulates a series of plans to apply to games with different adversarial teams; (ii) game plans are defined at an abstract level that needs to be refined for real execution; (iii) real-time execution in a team-plan execution framework/architecture that is capable of addressing key contingencies. Such an architecture also alleviates the planning concerns by providing some “commonsense” teamwork behaviors — not all of the coordination actions are required to be planned in detail as a result. The key research tasks here are:

Contingency planning for multiagent adversarial game playing: Before a game starts, one would expect the team to generate a strategic plan for the game that includes contingency plan segments that are to be recognized and eventually slightly adapted in real-time. Two main challenges can be identified in this task:

- Definition of strategic task actions with probabilistic applicability conditions and effects. Uncertainty in the action specification is directly related to the identification of possible probabilistic disruptive or favorable external events.
- Definition of objectives to achieve. In this domain, the goal of winning and scoring should be decomposed in a variety of more concrete goals that serve the ultimate final scoring goal. Examples are actions and goals to achieve specific attacking or defending positioning.

Plan decomposition and merge: A correspondence between team actions and goals and individual actions and goals must be set. The team plan decomposition may create individual goals that are not necessarily known to all the team players. Furthermore, within the contingency team plan, it is expected that there may be a variety of adversary-independent and adversary-dependent goals. The decomposition, coordination, and appropriate merge of individual plans to the service of the main team plan remain open challenging research tasks. RoboCup provides an excellent framework to study these issues.

Executing Team Plans: Team plan execution during the game is the determining factor in the performance of the team. It addresses the coordination contingencies that arise during the execution, without the need for detailed, domain-specific coordination plans. Execution also monitors the contingency conditions that are part of the global contingency team plan. Selection of the appropriate course of action is driven by the state information gathered by execution.

5.3 Evaluations

The Teamwork Challenge scenario described above has been idealized by several AI researchers, at least in the planning and multiagent communities. RoboCup, both in its simulated and real leagues, provides a synergistic framework to develop and/or test dynamic planning multiagent algorithms.

Specifically, we are planning to evaluate the architecture and teams in the following evaluation scheme:

Basic Performance: The team must be able to play reasonably well against both the best hand-coded teams, which has no planning, and against other planning-based systems. Relative performance of the team can be measured by actually playing a series of games against other unknown teams. Thus, basic performance will be measured by:

- Performance against hand-coded teams.
- Performance against other teams.

Robustness: The robustness in teamwork means that the team, as a whole, can continue to carry out the mission even if unexpected changes, such as accidental removal of the players in the team, sudden change of team composition, or changes in operation environment. For example, if one of players in the

team was disabled, the team should be able to cope with such accidents, by taking over the role of disabled players, or reformulating their team strategy. Thus, this evaluation represents a set of unexpected incidents during the game, such as:

- Some players will be disabled, or their capability will be significantly undermined by these accidents. Also, some disabled players may be enabled later in the game.
- Opponent switch their strategy, and the team must cope with their new strategy in real time.
- Some of opponent's players will be disabled, or their performance will be significantly undermined. These disabled players may come back to the game later.
- Teammate changes during the game.
- Weather factor changes.

The RoboCup Teamwork Challenge therefore is to define a general set of teamwork capabilities to be integrated with agent architectures to facilitate flexible, reusable teamwork. The following then establish the general evaluation criteria:

General Performance: General performance of the team, thus the underlying algorithms, can be measured by a series of games against various teams. This can be divided into two classes (1) normal competitions where no accidental factors involved, and (2) contingency evaluation where accidental factors are introduced.

Real-Time Operations: The real-time execution, monitoring, and replanning of the contingency plan is an important factor of the evaluation. For any team to be successful in the RoboCup server, it must be able to react in real time: sensory information arrives between 2 and 8 times a second and agents can act up to 10 times a second.

Generality: Reuse of architecture in other applications: Illustrate the reuse of teamwork capabilities in other applications, including applications for information integration on the internet, entertainment, training, etc.

Conformity with Learning: Finally, given the premises above and the complexity of the issues, we argue and challenge that a real-time multiagent planning system needs to have the ability to be well integrated with a learning approach, i.e., it needs to refine and dynamically adapt and refine its complete behavior (individual and team) based on its past experience.

Other issues such as reuse of teamwork architecture within the RoboCup community, and planning for team players that are not yet active in order to increase their probability of being useful in future moves, such as role playing and positioning of the team players that *do not* have the ball, will be considered, too.

6 RoboCup Opponent Modeling Challenge

Agent modeling – modeling and reasoning about other agent's goals, plans, knowledge, capabilities, or emo-

tions — is a key issue in multi-agent interaction. The RoboCup opponent modeling challenge calls for research on modeling a team of opponents in a dynamic, multi-agent domain. The modeling issues in RoboCup can be broken down into three parts:

On-line tracking: Involves individual players' real-time, dynamic tracking of opponents' goals and intentions based on observations of actions. A player may use such tracking to predict the opponents' play and react appropriately. Thus if a player predicts that player-5 is going to pass a ball to player-4, then it may try to cover player-4. Such on-line tracking may also be used in service of deception. The challenges here are (i) real-time tracking despite the presence of ambiguity; (ii) addressing the dynamism in the world; (iii) tracking teams rather than only individuals – this requires an understanding of concepts involved in teamwork.

On-line tracking may feed input to the on-line planner or the on-line learning algorithm.

On-line strategy recognition: "Coach" agents for teams may observe a game from the sidelines, and understand the high-level strategies employed by the opposing team. This contrasts with on-line tracking because the coach can perform a much higher-level, abstract analysis, and in the absence of real-time pressures, its analysis can be more detailed.

The coach agents may then provide input to its players to change the team formations, or play strategy.

Off-line review: "Expert" agents may observe the teams playing in an after-action review, to recognize the strengths and weaknesses of the teams, and provide an expert commentary. These experts may be trained on databases of human soccer play.

These issues pose some fundamental challenges that will significantly advance the state of the art in agent modeling. In particular, previous work has mostly focused on plan recognition in static, single-agent domains, without real-time constraints. Only recently has attention shifted to dynamic, real-time environments, and modeling of multi-agent teamwork [Tambe, 1996b].

A realistic challenge for IJCAI-99 will be to aim for on-line tracking. Optimistically, we expect some progress towards on-line strategy recognition; off-line review will likely require further research beyond IJCAI-99.

For evaluation, we propose, at least, following evaluation to be carried out to measure the progress:

Game Playing: A team of agents plays against two types of teams:

- One or two unseen RoboCup team from IJCAI-97, shielded from public view.
- The same unseen RoboCup teams from IJCAI-97 as above, but modified with some new behaviors. These teams will now deliberately try out new adventurous strategies, or new defensive strategies.

Disabled Tracking: Tracking functionality of the agents will be turned off, and compared with normal performance.

Deceptive Sequences: Fake teams will be created which generates deceptive moves. The challenger's agent must be able to recognize the opponent's deceptive moves to beat this team.

For each type of team, we will study the performance of the agent-modelers. Of particular interest is variations seen in agent-modelers behaviors given the modification in the opponents' behaviors. For each type of team, we will also study the advice offered by the coach agent, and the reviews offered by the expert agents, and the changes in them given the changes in the opponents' behaviors.

7 Managing Challenges

In order to facilitate technical progress based on the RoboCup challenge, we offer basic resources and opportunities.

The RoboCup Challenge Committee: The RoboCup Challenge Committee will be formed to execute the challenge initiative. The committee will include members of the international executive committee for RoboCup and distinguished researchers not directly involved in RoboCup. The committee will create specific tasks and criteria for evaluation, as well as providing technical advice for the challengers.

Resources: In the RoboCup home page, basic software resources and technical information can be obtained. (<http://www.robocup.org/RoboCup>) Software includes the Soccer Server system, which is a server system for the simulation track, and various sample teams. In addition, sample test sequences will be provided. The home page also provides a set of papers and technical documents related to RoboCup.

Competitions: A series of RoboCup competitions are planned to provide opportunities to test ideas. As international events, we are planning to have RoboCup-98 in Paris (The Official Event of the World Cup), RoboCup-98 Victoria (as a part of IROS-98 conference), and RoboCup-98 Singapore (as a part of PRICAI-98 Conference). Several local competitions will be organized by local committee in each region. The final evaluation and exhibit of the results will be made at IJCAI-99.

Workshops: Workshops will be organized at major international conferences, as well as at local workshops, in order to facilitate exchange of information, to have technical discussions, and to get feedback on the status of the challengers in relation to the overall framework of the challenge.

8 Conclusion

The RoboCup Challenge-97 offers a set of three fundamental challenges, focused on learning, real-time planning, and opponent modeling. Learning and real-time planning of multi-agent systems were chosen as the first set of challenges because they are essential technical issues for RoboCup, as well as for general AI systems using a multi-agent approach. Together with the physical

agent challenge, these challenges will be a basis for the RoboCup Challenge-99, and for longer research enterprises.

References

- [Asada et al., 1997] Asada, M., Kuniyoshi, M., Drogoul, A., Asama, H., Mataric, M., Duhaut, D., Stone, P., and Kitano, H., "The RoboCup Physical Agent Challenge: Phase-I," To appear in *Applied Artificial Intelligence (AAI) Journal*, 1997.
- [Cohen and Levesque, 1991] Cohen, P. R. and Levesque, H. J., "Confirmation and Joint Action", *Proceedings of International Joint Conf. on Artificial Intelligence*, 1991.
- [Firby, 1987] Firby, J., "An investigation into reactive planning in complex domains", *Proceedings of National Conf. on Artificial Intelligence*, 1987.
- [Grosz, 1996] Grosz, B., "Collaborating Systems", *AI magazine*, 17, 1996.
- [Hayes-Roth et al., 1995] Hayes-Roth, B. and Brownston, L. and Gen, R. V., "Multiagent collaboration in directed improvisation", *Proceedings of International Conf. on Multi-Agent Systems*, 1995.
- [Jennings, 1995] Jennings, N., "Controlling cooperative problem solving in industrial multi-agent systems using joint intentions", *Artificial Intelligence*, 75, 195-240, 1995.
- [Kitano et al., 1997a] Kitano, H., Asada, M., Osawa, E., Noda, I., Kuniyoshi, Y., Matsubara, H., "RoboCup: The Robot World Cup Initiative", *Proc. of the First International Conference on Autonomous Agent (Agent-97)*, 1997.
- [Kitano et al., 1997b] Kitano, H., Asada, M., Osawa, E., Noda, I., Kuniyoshi, Y., Matsubara, H., "RoboCup: A Challenge Problem for AI", *AI Magazine*, Vol. 18, No. 1, 1997.
- [Kitano, et al., 1995] Kitano, H. and Asada, M. and Kuniyoshi, Y. and Noda, I. and Osawa, E., "RoboCup: The Robot World Cup Initiative", *IJCAI-95 Workshop on Entertainment and AI/Alife*, 1995
- [Newell, 1990] Newell, A., *Unified Theories of Cognition*, Harvard Univ. Press, Cambridge, Mass., 1990.
- [Pollack, 1991] Pollack, M., "The uses of plans", *Artificial Intelligence*, 57, 1992.
- [Stone and Veloso, 1997] Stone, P. and Veloso, M., "A layered approach to learning client behaviors in the robocup soccer server," To appear in *Applied Artificial Intelligence (AAI) Journal*, 1997.
- [Tambe, 1996a] Tambe, M., "Tracking dynamic team activity", *Proceedings of National Conf. on Artificial Intelligence*, 1996.
- [Tambe, 1996b] Tambe, M., "Teamwork in real-world, dynamic environments", *Proc. International Conf. on Multi-agent Systems*, 1996.
- [Tambe et al., 1995] Tambe, M. and Johnson, W. and Jones, R. and Koss, F. and Laird, J. and Rosenbloom, P. and Schwamb, K., "Intelligent agents for interactive simulation environments", *AI Magazine*, 16, 1995.