ELSEVIER

Artificial Intelligence 110 (1999) 193–214

Artificial Intelligence

www.elsevier.com/locate/artint

RoboCup: Today and tomorrow—What we have learned

Minoru Asada [a,*], Hiroaki Kitano [b,c], Itsuki Noda [d], Manuela Veloso [e]

[a] *Graduate School of Engineering, Department of Adaptive Machine Systems, Osaka University, Suita, Osaka 565-0871, Japan*
[b] *Sony Computer Science Laboratories, Inc., 3-14-13 Higashi-Gotanda, Shinagawa, Tokyo 141-0022, Japan*
[c] *ERATO Kitano Symbiotic Systems Project, JST, 6-31-15, M-31 Suite 6A, Jinguu-Mae, Shibuya, Tokyo 150-0001, Japan*
[d] *Electrotechnical Laboratory, Tsukuba 305-8568, Japan*
[e] *Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA*

Received 1 February 1998; received in revised form 1 November 1998

Abstract

RoboCup is an increasingly successful attempt to promote the full integration of AI and robotics research. The most prominent feature of RoboCup is that it provides the researchers with the opportunity to demonstrate their research results as a form of competition in a dynamically changing hostile environment, defined as the international standard game definition, which the gamut of intelligent robotics research issues are naturally involved. This article describes what we have learned from the past RoboCup activities, mainly the first and the second RoboCups, and overviews the future perspectives of RoboCup in the next century. First, the issue on what and why RoboCup is addressed, and a wide range of research issues are explained. Next, the current leagues are introduced and the research achievements are reviewed from a viewpoint of system architecture. Some of these achievements are included in this special issue. Finally, prospects for future activities are discussed.

Keywords: RoboCup; Grand challenge; Simulation league; Real robot small-size league; Real robot middle-size league

---

1. Introduction

RoboCup (The Robot World Cup Initiative) is an attempt to promote intelligent robotics research by providing a common task for evaluation of various theories, algorithms, and

---

[*] Corresponding author. Email: asada@ams.eng.osaka-u.ac.jp.

PII: S0004-3702(99)00024-7

agent architectures [9]. RoboCup has currently chosen soccer as its standard task. In order for a robot (a physical robot or a software agent) to play a soccer game reasonably well, many technologies need to be integrated and a number of technical breakthroughs must be accomplished. The range of technologies spans the gamut of intelligent robotics research, including design principles for autonomous agents, multi-agent collaboration, strategy acquisition, real-time reasoning and planning, robot learning, and sensor fusion.

The First Robot World Cup Soccer Games and Conferences (RoboCup-97) was held during the International Joint Conference on Artificial Intelligence (IJCAI-97) at Nagoya, Japan with 37 teams around the world, and the Second Robot World Cup Soccer Games and Conferences (RoboCup-98) was held on July 2–9, 1998 at La Cite des Sciences et de l'Industrie (La Cite) in Paris with 61 teams. RoboCup-99 Stockholm will be held in conjunction with IJCAI-99 participated in by over 120 teams. A series of technical workshops and competitions have been planned for the future. While the competition part of RoboCup is highlighted in the media, other important RoboCup activities include technical workshops, the RoboCup Challenge program (which defines a series of benchmark problems), education, and infrastructure development. As of December 1998, RoboCup activity involves thousands of researchers from over 36 countries. Further information is available from the web site: `http://www.robocup.org/`.

1.1. Why RoboCup?

It's obvious that building a robot to play a soccer game is an immense challenge; readers might therefore wonder why even bother to propose RoboCup. It is our intention to use RoboCup as a vehicle to promote robotics and AI research, by offering a publicly appealing but formidable challenge. The idea of using soccer for robotics research is not new. In 1993, Alan Mackworth proposed in a paper titled "On Seeing Robots" [14] that soccer can be a good testbed of robotics and AI research. Independently, several researchers have been working on the soccer domain. These efforts merged into RoboCup.

A unique feature of RoboCup is that it is a systematic attempt to promote research using a common domain, mainly soccer. Also, it is perhaps the first to explicitly claim that the ultimate goal is to beat a human world cup champion team. One of the effective ways to promote engineering research, apart from specific application developments, is to set a significant long term goal. When the accomplishment of such a goal has significant social impact, we call this kind of goal a *grand challenge project*. Building a robot to play soccer is not such a project. But its accomplishment would certainly be considered as a major achievement in the field of robotics, and numerous technology spin-offs can be expected during the course of the project. We call this kind of project a *landmark project*, and RoboCup is definitely a project of this kind.

In the case of RoboCup, the ultimate goal is:

> "By the mid-21st century, a team of autonomous humanoid robots shall beat the human World Cup champion team under the official regulations of FIFA."

(A more modest goal is "to develop a robot soccer team which plays like a human player".)

In order for the landmark project to be successful, the goal has to be

(1) publicly appealing,
(2) difficult enough so that a great number of innovations need to be made for the accomplishment of the goal,
(3) within feasible challenge so that researchers can start their first step now, and
(4) such that types of technologies to be pursued can form the foundation of next generation industries.

The most successful landmark project in the history was the Apollo project, which has a clear and appealing goal of "landing a man on the moon and returning him safely to earth" as declared by John F. Kennedy in 1961, innovations are required, but feasible within a certain period of time, and technologies developed were transferred to aviation, electronics, material, and information industries, which created a basis of US industries in the post-Apollo era.

In RoboCup, we chose soccer as our main target after a range of feasibility studies, and believe that this is one of the best targets to promote research. This does not imply that soccer requires *all* the elements of technologies we wish to promote. In fact, RoboCup-Rescue is now proposed to complement features which are missing in soccer. However, soccer is considered to have overall merit being the main target.

RoboCup is also regarded as a "standard problem" so that various theories, algorithms, and architectures can be evaluated. Computer chess is a typical example of a standard problem. Various search algorithms were evaluated and developed using this domain. With the recent accomplishment by the Deep Blue team, which beat Kasparov, a human grand master, using the official rules, the challenge of computer chess is close to over. One of the major reasons for the success of computer chess as a standard problem is that the evaluation of progress was clearly defined. The progress of research can be evaluated as the strength of the system, which is indicated as a US chess rating. However, as computer chess is about to complete its original goal, we need a new challenge. The challenge needs to foster a set of technologies for the next generation of industries. We think that RoboCup fulfills such a demand. Table 1 illustrates the difference between the domain characteristics of computer chess and RoboCup.

RoboCup is designed to require the handling of real-world complexities, though in a limited world, while maintaining an affordable problem size and research cost. RoboCup offers an integrated research task covering broad areas of intelligent robotics. Such areas include: real-time sensor fusion, reactive behavior, strategy acquisition, learning, real-

Table 1
Comparison of chess and RoboCup

	Chess	RoboCup
Environment	Static	Dynamic
State change	Turn taking	Real time
Information accessibility	Complete	Incomplete
Sensor readings	Symbolic	Non-symbolic
Control	Central	Distributed

time planning, multi-agent systems, context recognition, vision, strategic decision-making, motor control, intelligent robot control, and many more.

1.2. *What's RoboCup?*

RoboCup has a series of activities such as competitions, conferences, RoboCup challenges, education, infrastructure, and secondary domain. Among them, however, competition remains the most well-known component. We think competition has unique value in testing robots and software teams in environments outside the laboratory. It also forces participants to build robot platforms which reliably perform the task, instead of showing superb performance once in a hundred times. And of course, competition is fun. It motivates students and appeals to spectators.

Currently, RoboCup consists of three competition tracks:

(i) *Simulation league*: Each team consists of eleven programs, each controlling separately each of the eleven team members. The simulation is run using the Soccer Server (see Fig. 1) developed by Noda et al. [16]. Each player has distributed sensing capabilities (vision and auditory) and motion energy both of which are resource bounded. Communication is available between players and strict rules of the soccer game are enforced (e.g., offsides). This league is mainly for researchers who may not have the resources for building real robots, but are highly interested in complex multi-agent reasoning and learning issues.

```
(referee before_kick_off)
(referee kick_off)
```

```
Recv: brazil_7:(turn 60)
Recv: italy_11:(turn 60)
Recv: brazil_9:(turn 60)
Recv: brazil_8:(turn 60)
Recv: italy_5:(turn 60)
Recv: brazil_3:(dash 13.5)
Recv: italy_1:(dash 18.4)
```

Fig. 1. A screen of Soccer Server.

(ii) *Small-size real robot league*: The field is of the size and color of a ping-pong table (see Fig. 2), and up to five robots per team play a match with an orange golf ball. The robot size is limited to approximately 15 cm^3. Typically robots are built by the participating teams and move at speeds of up to 2 m/s. Global vision is allowed, offering the challenge of real-time vision-based tracking of five fast moving robots in each team and the ball.

(iii) *Middle-size real robot league*: The field size is of the size and color of three by three ping-pong tables (see Fig. 3), and up to five robots per team play a match with a Futsal-4 ball. The size of the base of the robot is limited to approximately 50 cm diameter. Global vision is not allowed. Goals are colored and the field is surrounded by walls to allow for possible distributed localization through robot sensing.

There are several reasons to create two robot leagues with different size and regulations. First, there are needs for having leagues with different size due to space and budget

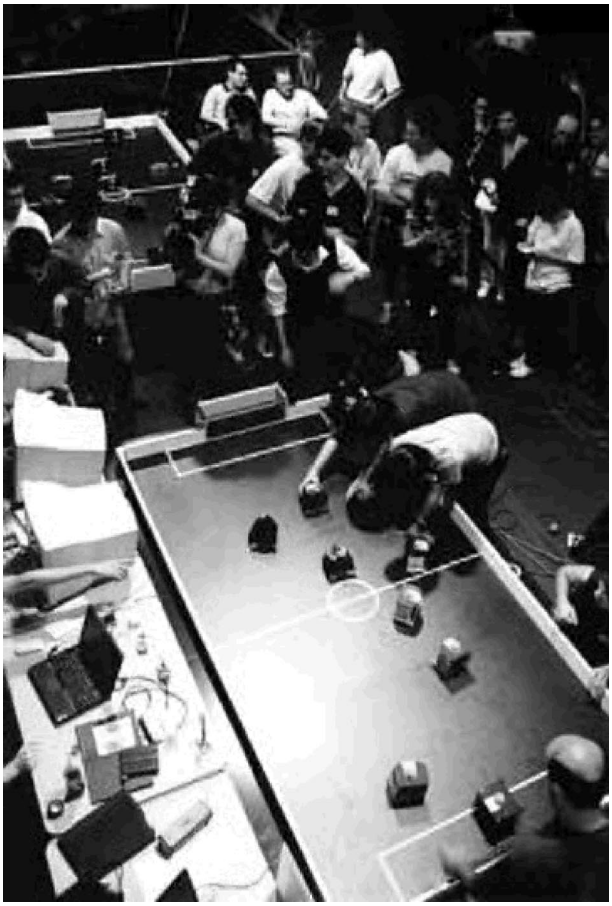

Fig. 2. Real robot small-size league competition site.

Fig. 3. Real robot middle-size league competition site.

constraints of each laboratory. Second, there are different technical issues involved for designing small-size robot and middle-size robots, as well as for potential application of technologies for each size. Third, there is difference in regulations. In the small-size league, we allow cameras and sensors to be located above and along the field, whereas in the middle-size league we require all sensing systems to be on-board. While there are several reasons for this difference, the main reason is that we need research on how to use sensors embedded in the environment to be integrated with on-board sensing and computing in real-time autonomous systems. For example, intelligent traffic systems (e.g., [22]) obviously involve a wide range of sensing systems placed along the road to detect traffic and control each vehicle. Placing sensing systems off-board is not a compromise, but it is a setup designed for research having practical applications in mind. Because the sensing systems are not on-board, it involves issues that are not involved in fully on-board systems, such as delay of sensing and processing, robustness against communication failure, etc.

On the other hand, the middle-size league is designed to promote research on fully autonomous systems where all sensing systems must be on-board.

Aside from the winner of each league, RoboCup awards the Scientific Challenge Award and Engineering Challenge Award for the team which made a major challenge with some success. This award was established to foster challenging scientific and engineering research in RoboCup. In general, the safest approach to winning the competition is to use conventional and reliable technologies well-tuned for the specific domain. The RoboCup domain is challenging enough so that any successful team must use some challenging technologies. However, these awards are given for truly high-risk and high-impact design.

In RoboCup-97, the Scientific Challenge Award was given to Sean Luke of the University of Maryland for demonstrating the utility of evolutionary computation by evolving soccer teams. Sean Luke used Genetic Programming to evolve soccer players trained on the RoboCup simulator in a massively parallel machine for a few months. The evolved team beat two hand-coded teams and survived the round-robin round to advance to the single-elimination round. Two engineering challenge awards were given to Uttori-United (a joint team consisting of Utsunomiya University, Toyo University, and RIKEN, Japan) and RMIT Raiders (Royal Melbourne Institute of Technology, Australia), for designing novel omnidirectional driving mechanisms. These teams designed new robot driving mechanisms which use special wheels (Uttori) and balls (RMIT) to enable robots to move in any direction without rotation (giving them, in effect, holonomic movement). Such mechanisms significantly improve a robot's maneuverability and simplify its control system, and their potential impact is far reaching.

In RoboCup-98, the Scientific Challenge Award was given to three research groups (Electrotechnical Laboratory (ETL), Japan, Sony Computer Science Laboratories, Inc., Japan, and German Research Center for Artificial Intelligence GmbH (DFKI)) for their simultaneous development of fully automatic commentator systems for the RoboCup simulator league.

While RoboCup started with the above three leagues, new leagues will be added as technology progresses. For RoboCup-99 Stockholm, a new league is created for legged robots, and a biped humanoid league is planned for RoboCup-2002.

Technical details of teams represented in RoboCup-97 as well as related research results were published in the official publication [10], and those of RoboCup-98 will appear in [2]. For the details of the results of all matches in RoboCup-97 and RoboCup-98, please visit the site: `http://www.robocup.org/`.

2. Research issues and approaches

In this section, we discuss several research issues involved in the development of real robots and software agents for RoboCup. One of the major reasons why RoboCup attracts 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 following is a partial list of research areas which RoboCup covers:

- agent architecture in general;
- combining reactive approaches and modeling/planning approaches;
- real-time recognition, planning, and reasoning;
- reasoning and action in a dynamic environment;
- sensor fusion;
- multi-agent systems in general;
- behavior learning for complex tasks;
- strategy acquisition;
- cognitive modeling in general.

Currently, each league has its own architectural constraints, and therefore research issues are slightly different from each other. We have published proposal papers [4,12] about research issues in the RoboCup initiative. For the synthetic agent in the simulation league, the following issues are considered:

- Teamwork among agents, from low-level skills like passing the ball to a teammate, to higher-level skills involving execution of team strategies.
- Agent modeling, from primitive skills like recognizing agents' intentions to pass the ball, to complex plan recognition of high-level team strategies.

- Multi-agent learning, for on-line and off-line learning of simple soccer skills for passing and intercepting, as well as more complex strategy learning.

For the robotic agents in the real robot leagues, for both the small- and middle-size ones, the following issues are considered:

- Efficient real-time global or distributed perception possibly from different sensing sources.
- Individual mechanical skills of the physical robots, in particular target aim and ball control.
- Strategic navigation and action to allow for robotic teamwork, by passing, receiving and intercepting the ball, and shooting at the goal.

More strategic issues are dealt with in the simulation league and in the small-size real robot league while acquiring more primitive behaviors of each player is the main concern of the middle-size real robot league.

We held the first RoboCup competitions in August 1997, in Nagoya, in conjunction with IJCAI-97 [10]. There were 28, 4, and 5 participating teams in the simulation, small-size robot, and middle-size robot leagues, respectively. The second RoboCup workshop and competitions took place in July 1998, in Paris [1] in conjunction with ICMAS-98 and AgentsWorld. The number of teams increased significantly from RoboCup-97 to 34, 11, and 16 participating teams in the simulation, small-size robot, and middle-size robot leagues, respectively. More than twenty countries participated. Every team had its own features some of which have been exposed during their matches with different degrees of success.

2.1. Architectural analysis

There are two kinds of aspects in designing a robot team for RoboCup:

(i) Physical structure of robots: actuators for mobility, kicking devices, perceptual (cameras, sonar, bumper sensor, laser range finder) and computational (CPUs, microprocessors) facilities.
(ii) Architectural structure of control software.

In the simulation league, both of the above issues are fixed, and therefore more strategical structure as a team has been considered. On the other hand, in the real robot leagues, individual teams have devised, built, and arranged their robots. Although the small league and the middle one have their own architectural constraints, there are variations of resource assignment and control structure of their robots. Table 2 shows the variations in architectural structure in terms of number of CPUs and cameras, and their arrangement.

Three types A, B, and C indicate a variation adopted in the real robot small-size league. Type A is a typical structure many teams used in this league: the centralized control of multiple bodies through a global vision. Type B is a kind of multi-agent system in which decision-making is distributed and independent from each other although they share the global vision. CMUnited-98 in the small-size league took this sort of architecture. Type C features sensor coordination of global and local views based on the centralized control with multiple bodies. I-space (a joint team of Utsunomiya University and University of Tokyo, Japan) in the small-size league adopted this type architecture.

Table 2
Variations in architectural structure

Type	CPU	Vision	Issues	League
A	1	1 global	Strategy	Small-size
B	n	1 global	Sharing of information	Small-size
C	1	1 global $+n$ local	Sensor fusion; coordination	Small-size
D	$1+n$	n local	Multiple robots	Middle-size
E	n	n local	Sensor fusion; teamwork	Middle-size and simulation

On the other hand, type E is a typical architecture adopted in both the simulation league and the real robot middle-size league: a completely distributed multi-agent system. Type D used by the C.S. Freiburg team in the middle-size league adopted a combination of types A and E utilizing laser range finders mounted on players which make it possible to reconstruct the global view and to localize observed objects (teammates, opponents, and ball) in the field. That is, they changed the problem in the middle-size league into one in the small-size league.

Communication between agents is possible in all of the leagues. The simulation league is the only one that uses it except for one team Uttori in the middle-size league. In the following, we attempt to analyze the achievements in RoboCup-97 and RoboCup-98 in terms of each league.

3. Simulation league

The simulation league continues to be the most popular part of the RoboCup leagues, with 34 teams participating in RoboCup-98, which is a slight increase over the number of participants at RoboCup-97. As with RoboCup-97, teams were divided into leagues. In the preliminary round, teams played within leagues in a round-robin fashion, and that was followed by a double-elimination round to determine the first three teams.

3.1. Changes of regulation

For RoboCup-98, the regulation has been changed as follows from RoboCup-97:

- *Offside* rule is introduced. As in human's soccer, offside players[1] cannot receive the ball. By this rule, defending players can use strategic plays called "offside trap", and offensive players should be aware of such strategies of opponent teams. This rule causes "compact soccer" in which most of the players gather into a relatively small area near the ball. This condition requires more flexible teamwork. (See below.)

---

[1] A player is offside when he/she is in the area that satisfied the following conditions:
- between the opponent goal-line and the second opponent defender (including a goalie);
- between the opponent goal-line and the ball;
- between the opponent goal-line and the center line.

- A player in a team can become the *goalie* who can *catches* the ball in his penalty area. When the goalie catches the ball, the match restarts by a free-kick of the goalie's team from the catching position. On the other hand, other defending players are inhibited to stay in their own goal area. [2]
- The *ball speed* is reduced.
  In RoboCup-97, a player can kick the ball three times continuously and the ball gets such a speed that a player can shoot the ball to the goal from the center line. Moreover, there is no way for defenders to stop such a fast ball because the ball moves so much during the delay between recognition and action of the defenders. In order to inhibit such nonsense plays, the limitation of the ball speed is introduced. By this limit, the ball can reach about half of the width of the field.
- The control of player's *stamina* is modified so that a player cannot run continuously through a match. In RoboCup-97, the stamina mechanism was already introduced, but the setup was not good because a player could dash by full power through a match. By the new stamina mechanism for RoboCup-98, the stamina exhausts after about 50 m dash of full power, and also, the speed of recovering the stamina decays when a player uses stamina so much (long-term stamina control). The new mechanism forces participants to take care of resource (stamina) management in the dynamic environment.

3.2. Changes of strategies and research issues

3.2.1. Offside and compact soccer

Because of the offside rule, most of the matches are carried out in the "compact soccer" style like human soccer, in which most of the players except the goalies are placed in the narrow band whose width is about 30–40 m. By the offside rule, the defenders use the offside trap in which they go up toward the opponent side in order to let the opponent offensive players offside. In order to avoid this trap, offensive players must go back to their own side if necessary. By such tactics, most of the players are pushed into the narrow area near of the ball.

The compact soccer provides two research issues:

- *Dynamic Formation*: Because players are pushed into the narrow area and the narrow area moves following the ball, the players must change their position dynamically. In RoboCup-97, most of the teams used a static formation strategy in which each player has his/her own home position in the field and stays there unless he/she is chasing the ball. In RoboCup-98, players cannot have such static home position, and change the position according to the status of a match. Because of the dynamic formation, the strategy of pass-work also changes. Under the static formation, the simplest way to pass the ball to the teammate is to kick the ball toward one of the home positions of the teammates. Because the home positions do not change during a match, a player can kick the ball without looking up to confirm the home position.

---

[2] The players can be in the goal area temporally. For example, a player can run into the goal area in order to chase the ball, but cannot stay too long in front of the goal in order to cover the goal intensively. This judgment is done by a human referee.

Under the dynamic formation, the player needs to confirm where the teammates are standing. Therefore, each team took care how to manage dynamic formation using communication and/or pre-defined team strategies that should be applied dynamically by each player independently.

- *Opponent Monitoring*: Because of tactics between offside traps and avoiding the traps, it becomes more important to monitor movement of opponent players. In order to set an offside trap, defenders need to recognize position and movement of all opponent offenders. On the other hand, offensive players should take care of the positions of all opponent defenders not to be caught in the trap. Though the strategy of the opponent monitoring used in RoboCup-98 is not so complicated, it will become more sophisticated in future, because this tactics is end-less game.

3.2.2. Man marking

Strong teams, especially CMUnited-98, use explicit man-marking defense. In RoboCup-97, most of the teams use space-marking defense, in which, each defense player is assigned a role to cover a certain space in their own side of the field. In other words, the formation and position of defense player do not change by positions of opponent players during a match. In RoboCup-98, on the other hand, each defender of CMUnited-98 is assigned a role to mark a certain opponent forward player, or a role of sweeper.[3] Therefore, two or three defenders are completely following the movement of opponent offensive players.

As for compact soccer described before, man-marking strategy is a starting point of the research of opponent-monitoring and opponent-modeling. While defenders were just following the opponents' movements in this year, it will be easy to extend this strategy to predict the movements and to expect the intention of opponents in future RoboCups.

3.2.3. Variation of plays

The variation of plays of ball-players increased. Especially, the following are new styles of plays in RoboCup-98.

- *Short-range Pass*: Because the ball speed is limited, the reach of a pass is shortened. As a result, a team should perform systematic pass-works in short range in order to break a defense line of opponent teams.
- *Dribbling and Pass*: Because of the short reach of a pass, dribbling became one of the effective plays in RoboCup-98. In RoboCup-97, a pass reach was too long, so that dribbling had fewer advantages than a pass. In addition to it, the "offside" rule forbids forward players to be in front of the opponent goal. Therefore, it became more reasonable for mid-fielders to dribble rather than to pass the ball to the forward players directly.
- *Through Pass and Back Pass*: Because of compact soccer, a ball player often cannot find any suitable pass-receiver in front of him. In such case, good teams in RoboCup-98 can perform back pass or through pass. These plays require an important research issue, that is, how to expect the next play of a teammate. In case of a through pass, the ball player should expect the receiver will run to the "no man's land" where he will kick the ball to. In the case of a back pass, the ball player should recognize that

---

[3] A sweeper is a free player placed behind other defense players who covers their fails of defense.

the receiver who is standing behind him can receive the ball freely and find suitable receivers of the next pass. Otherwise, such plays are bad strategies for the team.

Dribble and other pass works give players variations of plays they can select. For example, a mid-fielder can select one of a pass to a forward player, dribbling, a back pass to another mid-fielder, and a through pass to a forward player. The variation brings an important issue of research, dynamic and real-time decision-making in multi-agent systems.

3.3. RoboCup challenge in simulation

Teams in the RoboCup simulation league are faced with three strategic research challenges: multi-agent learning, teamwork and agent modeling. All three are fundamental issues in multi-agent interactions. The learning challenge has been categorized into on-line and off-line learning both by individuals and by teams (i.e., collaborative learning). One example of off-line individual learning is learning to intercept the ball, while an example of on-line collaborative learning is to adaptively change player positions and formations based on experience in a game.

The RoboCup Teamwork Challenge addresses issues of real-time planning, re-planning, and execution of multi-agent teamwork in a dynamic adversary environment. A team should generate a strategic plan, and execute it in a coordinated fashion, monitoring for contingencies and select appropriate remedial actions. Stone and Veloso introduced a concept of periodic team synchronization (hereafter, PTS) to emphasize this domain characteristics and proposed a general team member architecture suitable in the PTS domain (for the details, see [18]).

The teamwork challenge interacts also with the third challenge in the RoboCup simulation league, that of agent modeling. Agent modeling refers to modeling and reasoning about other agent’s goals, plans, knowledge, capabilities, or emotions. The RoboCup opponent modeling challenge calls for research on modeling a team of opponents in a dynamic, multi-agent domain. Such modeling can be done on-line to recognize a specific opponent’s actions, as well as off-line for a review by an expert agent.

At least some researchers have taken these research challenges to heart, so that teams at RoboCup-97 and RoboCup-98 have addressed at least some of the above challenges. In particular, out of the three challenges outlined, researchers have attacked the challenge of on-line and off-line learning (at least by individual agents). Thus, in some teams, skills such as intercept, and passing are learned off-line. The two final teams, namely CMUnited simulation (USA) as winner of the first place and AT-Humboldt-98 (Germany) as runner-up, included an impressive combination of individual agent skills and strategic teamwork.

Research in teamwork has provided concepts such as exhibiting reusability of domain-independent teamwork skills (i.e., skills that can be transferred to domains beyond RoboCup), about roles and role reorganization in teamwork. Tambe et al. proposed a teamwork modeling and learning method and have shown some results based on the proposed method (for the details, see [21] in this special issue). RoboCup opponent modeling, in terms of tracking opponents’ mental state, has, however, not received significant attention by researchers. There are, however, some novel commentator agents that have used statistical and geometric techniques to understand the spatial pattern of play.

4. Small-size real robot league

The RoboCup-98 small-size real robot league provides a very interesting framework to investigate the full integration of action, perception, and high-level reasoning in a team of multiple agents. Therefore, three main aspects need to be addressed in the development of a small-size RoboCup team:

(i) hardware of physical robots;
(ii) efficient perception; and
(iii) individual and team strategy.

Although all of the eleven RoboCup-98 teams included distinguishing features at some of these three levels, it showed crucial to have a *complete* team with robust hardware, perception, and strategy, in order to perform well overall. This was certainly the case for the four top teams in the competition, namely CMUnited-98 (USA), Roboroos (Australia), 5DPO (Portugal), and Cambridge (UK), who classified in first, second, third, and fourth place, respectively.

Fig. 4 shows a scene from the final match between CMUnited-98 and Queensland Roboroos (Australia). We overview now the characteristics of the RoboCup-98 teams and the research issues addressed.

Hardware. All of the eleven RoboCup-98 participating teams consisted of robots built by each participating group. The actual construction of robots within the strict size limitations offered a real challenge, but gave rise to a series of interesting physical and mechanical devices. Remarkably, the robots exhibited sensor-activated kicking devices (iXs and J-Star, Japan, Paris-6, France, and CMUnited-98, USA), sophisticated ball holding and shooting tools for the goalie robot (Cambridge, UK), and impressive compact and robust designs

Fig. 4. Real robot small-size final match.

(Roboroos, Australia, and UVB, Belgium). All of the robots were autonomously controlled through radio communication by off-board computers.

Perception. Ten out of the eleven teams used a single camera overlooking the complete field. The ISpace (Japan) team included one robot with an on-board vision camera.

Global perception simplifies the sharing of information among multiple agents. However global perception presents at the same time a real challenging research opportunity for reliable and real-time detection of the multiple mobile objects—the ball, and five robots on each team. In fact, both detection of robot position and orientation and robot tracking need to be very effective. The frame rate of the vision processing algorithms clearly impacted the performance of the team. Frame rates reached 30 frames/s as in the CMUnited-98 team.

In addition to the team color (blue or yellow), most of the teams used a second color to mark their own robots and provide orientation information, hence only about their own robots. Robot identification was achieved in general by greedy data association between frames. The 5DPO (Portugal) and the Paris-6 (France) teams had a robust vision processing algorithm that used patterns to discriminate among the robots and to find their orientation.

The environment in the small-size league is highly dynamic with robots and the ball moving at speeds between 1 and 2 m/s. An interesting research issue consists of the prediction of the motion of the mobile objects to combine it with strategy. It was not clear which teams actually developed prediction algorithms. In the particular case of the CMUnited-98 team, prediction of the movement of the ball was successfully achieved and highly used for motion (e.g., ball interception) and strategic decisions (e.g., goaltender behavior and pass/shoot decisions).

Motion. In this RoboCup league, a foul should be called when robots push each other. This rule offers another interesting research problem, namely obstacle avoidance and path planning in a highly dynamic environment. The majority of the teams in RoboCup-98 successfully developed algorithms for such difficult obstacle avoidance and the semi-final and final games showed smooth games that demonstrated impressive obstacle avoidance algorithms.

Strategy. Following up on several of the research solutions devised for RoboCup-97 both in simulation and in the small-size robot league, at RoboCup-98, all of the small-size teams showed a role-based team structure. As expected, the goaltender played a very important role in each team. Similarly to the goaltender of CMUnited-97, the goaltender of most of the teams stayed parallel to the goal line and tried to stay aligned with or intercept the ball. The goaltender represented a very important and crucial role. Notable were the goaltenders of Roboroos, CMUnited-98, and Cambridge.

Apart for CMUnited-98 which had a single defender and three attackers, most of the other teams invested more heavily on defense, assigning two robots as defenders. In particular, defenders in the Belgium and in the Paris-8 teams occupied key positions in front of the goal making it difficult for other teams to plan a path around them and to try to devise shots through the reduced open goal areas. Defending with polygonally-shaped robots proved to be hard, as the ball is not easily controlled at a fine grain. In fact a few goals for different teams were scored into their own goals due to small movements of the defenders or goaltender very close to the goal. It is clearly still an open research question how to control the ball more accurately.

Finally, it is interesting to note that one of the main features of the winning CMUnited-98 team is its ability to collaborate as a team. Attacking robots continuously evaluate (30 times per second) their actions, namely either to pass the ball to another attacking teammate or to shoot directly at the goal. A decision-theoretic algorithm is used to assign the heuristic and probabilistic based values to the different possible actions. The action with the maximum value is selected. Furthermore, in the CMUnited-98 team, a robot who was not the one actively pursuing the ball is not merely passive. Instead each attacking robot *anticipates* the needs of the team and it positions itself in the location that maximizes the probability of a successful pass. CMUnited-98 uses a multiple-objective optimization algorithm with constraints to determine this strategic positioning. The objective functions maximize repulsion from other robots and minimize attraction the ball and to the attacking goal.

5. Middle-size real robot league

The performance of robot behaviors in RoboCup-98 was better than in RoboCup-97 although the number of teams in the middle-size league drastically increased from 5 to 16, more than three times. However, the level of skills is under development, mainly putting more focus on individual behavior acquisition than cooperative teamwork. Engineering issues such as precise robot control and robust object detection are still main issues in this league.

5.1. Technological state-of-the-art

(i) *Platforms*: Many of the new teams used off-the-shelf platforms, such as Activmedia's Pioneer-1 and Pioneer-AT robots (one team in RoboCup-97 and six teams in RoboCup-98) or Nomadics' Scout robot (one team in RoboCup-98). There is a trade-off between the use of well-equipped off-the-shell platforms and the building of originally designed ones such as RMIT, Uttori and so on. The former seems easy to use immediately but less flexible while the latter seems more flexible but to consume much more time to build. Several teams such as Osaka, USC, and NAIST used the same cheap platforms, radio-controlled toy cars, but fairly modified them in different ways.

(ii) *Sensors*: Vision remains a central research topic in RoboCup. Recent progress of PC-based image processors enabled many teams to install them on-board, owing to which better performance of robot behaviors were shown in RoboCup-98 than in RoboCup-97. In addition to image processors, sonar and bumper sensors, a laser range finder was introduced to reconstruct a global view of the field. However, there are still many problems with the perceptual capabilities of the robots, especially detecting other agents.

(iii) *Kicking mechanisms*: Some teams such as Freiburg, Ulm, and Uttori introduced their kicking devices on their robots based on pneumatic or solenoid-based activation in RoboCup-98. The kicking devices produced much higher ball accelerations than pushing the ball, and as a result, such robots could move the ball significantly better, which is one of the research issues in this league.

5.2. Research results

Since it seems difficult to survey the research issues attacked by all the teams in the middle-size league, we show some of them which seem essential for RoboCup.

5.2.1. Behavior acquisition

Since the engineering issues such as precise motion control and robust vision are still main concerns in this league, many teams explicitly specify robot behavior as a form of IF-THEN rules from the viewpoint of the designer. Trackies (Osaka University, Japan) has been focusing on learning and evolutionary approaches, and the work in this special issue [5] has made reinforcement learning applicable to a dynamically changing multi-agent environment after state vector estimation based on a method of system identification. Further, they applied genetic programming [24] to emerge turning around the ball behavior which is difficult to acquire by reinforcement learning because of loss of ball sight.

Werger [26] (Ullanta Performance Robotics) adopted a simple behavior coordination based on the subsumption architecture and well-considered physical constraints on the relationship between robot motions and the environment. During the matches, both teams showed some good performance based on their methods.

5.2.2. Vision

All teams used vision as a main source of external sensing for individual robots. Most on-board cameras are fixed to the body, that is no active vision. Therefore, many robots rotate their bodies to expand robots' views. Two alternatives are considered. One is to use active vision, that is, panning a camera. NAIST's players have on-board panning cameras, and their searching behaviors (observation) apart from body motions seem to have been explicitly programmed. One of Osaka's players also has a panning camera and they attempted to emerge a behavior combined with camera and body motions based on reinforcement learning, but still under development except for simple situations. Active camera control is one of the essential issues because it is closely related to the attention mechanism. Although this problem has not been intensively addressed yet, RoboCup can provide a good testbed on this issue, that is, the role of attention will be made clear in the global context.

The other one is to use multiple or omni-directional vision. USC's Dream Team-98 [17] put two cameras oriented to opposite directions of each other onto one player to extend the single robot's view. Omni-directional vision was first used by Osaka in RoboCup-97 in order for a goal keeper to track both its position in front of the goal as well as the ball position [19]. In RoboCup-98, more teams such as Italian and Australian adopted this sort of vision system. Again, Osaka attacked to emerge goal-keeping behavior by reinforcement learning, and showed an aggressive behavior by not simply goal saving but also pushing the ball ahead expecting to pass the ball to one of its teammates in future.

5.2.3. Environment model and localization

Currently, many teams used the geometric field model to localize their robots. For example, USC's Dream Team used encoder of the robot to roughly estimate its position in the field to switch its behavior. However, there is no feedback, which results in accumulation of position errors. Then, visual information may correct these errors by capturing landmarks such as goals.

C.S. Freiburg [8] used laser range finders which provide fast and accurate range data. Each robot sends its range data to the host computer, which reconstructs the global field view including teammates and opponents based on these data and makes a plan, then informs such global information with action plan to each robot. As a result, their approach seems to be much more based on global positioning by LRF and centralized control (Type D in Table 2) although each player has its own CPU to detect a ball and to control its body.

Aside from the use of the geometric model of the field, one can approach the issue of how to represent the world as one of cognitive model. From sensorimotor mapping to a higher-level cognition process, the robot may not need to have an explicit world model like a geometric map. Instead, some sort of internal representation can be obtained, which implicitly represents the relationship between a sensorimotor mapping and a robot environment [3,20].

5.2.4. Communication

Communication between players is definitely a very important part of multi-agent coordination. Explicit communication lines between players were only used by the Uttori team [13] via WEB-LAN. Players exchange information and action commands to cooperate with each other such as passing or saving the goal together based on the pre-specified communication protocol.

A more interesting approach can be considered following human players behavior. Observation and action can be regarded as message receiving and passing if the robot can predict other robot actions through visual information. In order to acquire such capability, Uchibe et al. [23] proposed a state vector estimation method so that reinforcement learning can be applied to multi-agent behavior learning. That is, implicit communication.

6. Future issues

6.1. Current leagues

6.1.1. Simulation league

The major progress from RoboCup-97 to RoboCup-98 has been shown in the aspect of more dynamic and systematic teamworks. Especially, introduction of the offside rule and improvement of individual plays force flexible team plays. However, the stage in RoboCup-98 is still in the preliminary level. For example, tactics to escape the offside trap were still passive even in champion teams. In future RoboCup, such tactics will require recognition of intention of opponent players/teams. In this stage, opponent modeling and management of team strategies would become more important. Similarly, on-line learning will become more important, because team strategies should be changed during a match according to the strategies of opponent teams.

6.1.2. Real robot small-size league

The small-size RoboCup league provides a very rich framework for the development of multi-agent real robotic systems. We look forward to better understanding several issues, including the limitations imposed by the size restrictions on on-board capabilities; the robustness of global perception and radio communication; and strategic teamwork. One of the main interesting open questions is the development of algorithms for on-line learning of the strategy of the opponent team and for the real-time adaptation of one's strategy in response. Finally, similarly to the simulation and middle-size leagues, we want to abstract from our experience algorithms that will be applicable beyond the robotic soccer domain.

6.1.3. Real robot middle-size league

Despite the encouraging development of the middle-size league, we have to carefully review our current testbed and slowly adapt it to foster research in new directions and new areas. In most cases, this will require a slow evolution of rules.

The focus on colors to visibly distinguish objects exerts a strong bias for research in *color-based* vision methods. It is desirable to permit other approaches as well, such as using *edges*, *texture*, *shape*, *optical flow* etc., thereby widening the range of applicable vision research within RoboCup.

Another issue is the study of a better obstacle avoidance approach. Currently, most robots except NAIST [15] and a few others cannot reliably detect collisions with walls or other robots. Solving the charging problem using a rich set of on-board sensors is another major field of future research for RoboCup teams.

Finally, the use of communication in the different leagues is also an active research topic. Communication allows interesting research in a variety of topics, including multi-robot sensor fusion and control. We want to explore limited communication environments and its relationship to agent autonomy, and learning of cooperative behavior.

6.2. New leagues

6.2.1. Sony Legged Robot league

In RoboCup-98, Sony Legged Robot exhibition games and demonstration were held, and they attracted many spectators, especially boys and girls, for its cute style and behaviors. The four-legged robot has totally 16 DOFs (degrees of freedom): each leg has three joints, the head has pan, tilt, and roll movements, and the tail has one DOF [7].

Fig. 5 shows a scene from their demonstrations. Three teams from Osaka University, CMU, and University of Paris-VI showed their exhibition games provided the fundamental behavior control library and software development environment [25]. In 1999, Sony Legged Robot league will be one of the RoboCup official competitions with more nine teams around the world.

6.2.2. Humanoid league

Currently, except for Sony Legged Robot league, games are played by wheel-based robots. A soccer game by humanoid robots is the next major leap in the field which will lead to the ultimate goal of the RoboCup [11]. We proposed that the ultimate goal of the RoboCup Initiative be stated as follows:

Fig. 5. Sony Legged Robot league competition site.

> By the mid-21st century, a team of fully autonomous humanoid robot soccer players shall win the soccer game, comply with the official rules of the FIFA, against the winner of the most recent World Cup.

We propose that this goal be one of the grand challenges shared by robotics and AI community for the next 50 years. This goal may sound overly ambitious given the state-of-the-art technology today. Nevertheless, we believe it is important that such a long range goal be claimed and pursued. It took only 50 years from the Wright Brothers' first aircraft to the Apollo mission to send men to the moon and safely return them to the earth. Also, it took only 50 years, from the invention of digital computer to the Deep Blue, which beat the human world champion in chess. We recognize, however, that building a humanoid soccer player requires an equally long period and extensive efforts of an broad range of researchers, and that the goal will not be met in any near term.

Before, actually playing soccer with human players, RoboCup will organize humanoid leagues in the following categories, and start the Humanoid league competition from RoboCup-02 (2002).

- *Fully Autonomous Humanoid League*: Soccer games by teams of fully autonomous humanoid biped robots. A regular league will be performed by humanoid robots of height equivalent to a real human. In addition to that, smaller-size leagues are considered to meet the demand of some institutes that have already started humanoid design and its realization.

- *Tele-operation Humanoid League*: Soccer games by teams of tele-operated humanoid robots. The operator is allowed to control the robot only through the information obtained by sensors on-board the robot.
- *Virtual Humanoid League*: Soccer games by teams of simulated humanoid robots, with high quality computer graphics, accurate physics simulation, and vision and sensor simulation.

A great number of AI and robotics issues are involved in building humanoid robots for RoboCup, as pointed out by Kitano and Asada [11], some of them require fundamental innovation in integrated areas of material, chemical, computer science, and robotics. At the same time, as seems apparent from Honda Humanoid, the initial set of technologies is already available.

6.3. RoboCup new activities

6.3.1. RoboCup Jr.

The comprehensive nature of RoboCup makes it an ideal subject for project-oriented AI and robotics courses (e.g., [6]). Already, a few undergraduate and graduate courses are now being planned using RoboCup. Further, an education infrastructure named *RoboCup Jr.* is proposed, reflecting the needs of educational institutions. RoboCup Jr. will use cheaper robots and a much simpler task domain, rather than the highly challenging arrangement seen in the current RoboCup competition. A prototype of platforms will make a debut in 2000 or 2001, and the official league RoboCup Jr. will start from 2002.

As a currently available platform for RoboCup Jr., Lego Mind Storm is able to provide cheap and intellectual toys. In RoboCup-98 Paris, the University of Aarhus has built an exciting soccer stadium using Lego Mind Storm with many figures of supporters that could wave and give great cheers for the play. [4]

6.3.2. RoboCup-Rescue

Disaster rescue is one of the most serious social issues which involves very large numbers of heterogeneous agents in a hostile environment. RoboCup-Rescue intends to promote research and development in this socially significant domain by creating a standard simulator and forum for researchers and practitioners. While the rescue domain is intuitively appealing as large-scale multi-agent domain, it has not yet given thorough analysis on its domain characteristics. RoboCup-Rescue targets search and rescue activities for large-scale disasters like the Kobe earthquake in 1995.

We chose Urban Search and Rescue as a secondary domain for RoboCup, because

(1) it is a socially significant real-world domain,
(2) numbers of features which are missing in soccer exist, and
(3) there are certain commonalities between rescue and soccer, where the essence of autonomous multi-agent systems can be investigated through the use of two domains.

---

[4] Visit www.daimi.au.dk/~hhl.

Table 3
Features of Rescue and Soccer

	Rescue	Soccer
Number of agents	100 or more	11 per a team
Agents in the team	Heterogeneous	Homogeneous (except Goalee)
Logistics	Major issue	No
Long-term planning	Major issue	Less emphasized
Emergent collaboration	Major issue	No
Hostility	Environment	Opponent players
Environment	Different for each time	Same (fixed play field)

Features missing in soccer, but playing important roles are long-term strategy planning, logistics, planning for heterogeneous agents, interaction with human agents, emergent collaborations, etc. (see Table 3).

RoboCup-Rescue consists of a simulator league and a real robot league. The simulator league focuses on strategy planning and team coordination, whereas the focus of the real robot league will be on capability of individual robots in rescue operation. We are planning to hold the first RoboCup-Rescue Forum (exhibition matches and conferences) in Kobe, 2002, when the festival of recovery from the large earthquake in 1995 will be held.

7. Conclusion

As a grand challenge, RoboCup is definitely stimulating a wide variety of approaches, and has produced rapid advances in key technologies. With a growing number of participants RoboCup is set to continue this rapid expansion. With its three leagues, RoboCup researchers face a unique opportunity to learn and share solutions in three different agent architectural platforms.

We hope that RoboCup activity promotes AI and robotics research, and is a source of innovation for the 21st century.

Acknowledgement

Initial inputs to this article from Milind Tambe (ISI-USC, USA) and Gerhard K. Kraetzschmar (University of Ulm, Germany) are appreciated.

References

[1] M. Asada (Ed.), Proc. 2nd RoboCup Workshop, The RoboCup Federation, 1998.
[2] M. Asada, H. Kitano (Eds.), RoboCup-98: Robot Soccer World Cup II, Lecture Notes in Artificial Intelligence, Springer, Berlin, 1999, to appear.
[3] M. Asada, S. Noda, K. Hosoda, Action based sensor space segmentation for soccer robot learning, Applied Artificial Intelligence 12 (2–3) (1998) 149–164.

[4] M. Asada, P. Stone, H. Kitano, A. Drogoul, D. Duhaut, M. Veloso, H. Asama, S. Suzuki, The RoboCup physical agent challenge: Goals and protocols for phase I, in: H. Kitano (Ed.), RoboCup-97: Robot Soccer World Cup I, Lecture Notes in Artificial Intelligence, Vol. 1395, Springer, Berlin, 1998, pp. 41–61.

[5] M. Asada, E. Uchibe, K. Hosoda, Cooperative behavior acquisition for mobile robots in dynamically changing real worlds via vision-based reinforcement learning and development, Artificial Intelligence 110 (2) (1999) 275–292 (this issue).

[6] S. Coradeschi, J. Malec, How to make a challenging AI course enjoyable using the RoboCup soccer simulation system, in: RoboCup-98: Proc. 2nd RoboCup Workshop, 1998, pp. 57–62.

[7] M. Fujita, S. Zrehen, H. Kitano, A quadruped robot for RoboCup legged robot challenge, in: RoboCup-98: Proc. 2nd RoboCup Workshop, 1998, pp. 73–88.

[8] S. Gutmann, B. Nebel, The CS Freiburg team, in: RoboCup-98: Proc. 2nd RoboCup Workshop, 1998, pp. 451–458.

[9] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, H. Matsubara, RoboCup: A challenge problem of AI, AI Magazine 18 (1997) 73–85.

[10] H. Kitano (Ed.), RoboCup-97: Robot Soccer World Cup I, Lecture Notes in Artificial Intelligence, Vol. 1395, Springer, Berlin, 1998.

[11] H. Kitano, M. Asada, RoboCup humanoid challenge: That's one small step for a robot, one giant leap for mankind, in: Proc. IROS-98, 1998, pp. 419–424.

[12] H. Kitano, M. Tambe, P. Stone, M. Veloso, S. Coradeschi, E. Osawa, H. Matsubara, I. Noda, M. Asada, The RoboCup synthetic agent challenge 97, in: H. Kitano (Ed.), RoboCup-97: Robot Soccer World Cup I, Lecture Notes in Artificial Intelligence, Vol. 1395, Springer, Berlin, 1998, pp. 62–73.

[13] K. Yokota, K. Ozaki, N. Watanabe, A. Matsumoto, D. Koyama, T. Ishikawa, K. Kawabata, H. Kaetsu, H. Asama, Cooperative team play based on communication, in: RoboCup-98: Proc. 2nd RoboCup Workshop, 1998, pp. 491–496.

[14] A. Mackworth, On seeing robots, in: Computer Vision: Systems, Theory, and Applications, World Scientific Press, Singapore, 1993, pp. 1–13.

[15] T. Nakamura, K. Terada, A. Shibata, J. Morimoto, H. Adachi, H. Tadeka, The RoboCup-NAIST: A cheap multisensor-based mobile robot with on-line visual learning capability, in: RoboCup-98: Proc. 2nd RoboCup Workshop, 1998, pp. 291–304.

[16] I. Noda, H. Matsubara, K. Hiraiki, I. Frank, Soccer server: A tool for research on multiagent systems, Applied Artificial Intelligence 12 (2–3) (1998) 233–250.

[17] W.M. Shen, J. Adibi, R. Adobbati, S. Lanksham, H.M.B. Salemi, S. Tejada, Integrated reactive soccer agents, in: RoboCup-98: Proc. 2nd RoboCup Workshop, 1998, pp. 251–264.

[18] P. Stone, M. Veloso, Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork, Artificial Intelligence 110 (2) (1999) 241–273 (this issue).

[19] S. Suzuki, T. Katoh, M. Asada, An application of vision-based learning for a real robot in RoboCup learning of goal keeping behaviour for a mobile robot with omnidirectional vision and embedded servoing, in: RoboCup-98: Proc. 2nd RoboCup Workshop, 1998, pp. 281–290.

[20] Y. Takahashi, M. Asada, K. Hosoda, Reasonable performance in less learning time by real robot based on incremental state space segmentation, in: Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS96), 1996, pp. 1518–1524.

[21] M. Tambe, J. Adibi, Y. Al-Onaizan, A. Erdem, G.A. Kaminka, S.C. Marsella, I. Muslea, Building agent teams using an explicit teamwork model and learning, Artificial Intelligence 110 (2) (1999) 215–239 (this issue).

[22] C. Thorpe, T. Jochem, D. Pomerleau, Automated highways and the free agent demonstrations, in: Y. Shirai, S. Hirose (Eds.), Robotics Research, 7th International Symposium, Springer, Berlin, 1998, pp. 246–254.

[23] E. Uchibe, M. Asada, K. Hosoda, State space construction for behavior acquisition in multi agent environments with vision and action, in: Proc. ICCV-98, 1998, pp. 870–875.

[24] E. Uchibe, M. Nakamura, M. Asada, Co-evolution for cooperative behavior acquisition in a multiple mobile robot environment, in: Proc. IROS-98, 1998, pp. 425–430.

[25] M. Veloso, W. Uther, M. Fujita, M. Asada, H. Kitano, Playing soccer with legged robots, in: Proc. IROS-98, 1998, pp. 437–442.

[26] B.B. Werger, Cooperation without deliberation: A minimal behavior-based approach to multi-robot teams, Artificial Intelligence 110 (2) (1999) 293–320 (this issue).