

# The RoboCup Physical Agent Challenge: Phase I

Minoru Asada  
Osaka Univ., Japan  
Yasuo Kuniyoshi  
ETL, Japan

Peter Stone  
CMU, USA  
Alexis Drogoul  
LAFORIA/CNRS, France  
Hajime Asama  
Riken, Japan

Hiroaki Kitano  
Sony CSL, Japan  
Dominique Duhaut  
LRP, France  
Sho'ji Suzuki  
Osaka Univ., Japan

Barry Werger  
Brandeis Univ., USA  
Manuela Veloso  
CMU, USA

September 23, 1997

## Abstract

Traditional AI research has not given due attention to the important role that physical bodies play for agents as their interactions produce complex emergent behaviors to achieve goals in the dynamic real world. The RoboCup Physical Agent Challenge provides a good testbed for studying how physical bodies play a significant role in realizing intelligent behaviors using the RoboCup framework [Kitano, *et al.*, 95]. In order for the robots to play a soccer game reasonably well, a wide range of technologies needs to be integrated and a number of technical breakthroughs must be made. In this paper, we present three challenging tasks as the RoboCup Physical Agent Challenge Phase I: (1) moving the ball to the specified area (shooting, passing, and dribbling) with no, stationary, or moving obstacles, (2) catching the ball from an opponent or a teammate (receiving, goal-keeping, and intercepting), and (3) passing the ball between two players. The first two are concerned with single agent skills while the third one is related to a simple cooperative behavior. Motivation for these challenges and evaluation methodology are given.

## 1. Introduction

The ultimate goal in AI, and probably in robotics, is to build intelligent systems capable of displaying complex behaviors to accomplish the given tasks through interactions with a dynamically changing physical world. Traditional AI research has been mainly pursuing the methodology of symbol manipulations to be used in knowledge acquisition and representation and reasoning about it with little attention to intelligent behavior in dynamic real worlds [Brooks, 1991]. On the other hand, in robotics much more emphasis has been put on the issues of designing and building hardware systems and their controls. However, recent topics spread over the two areas include design principles of autonomous agents, multi-agent collaboration, strategy acquisition, real-time reasoning and planning, intelligent robotics, sensor-fusion, and behavior learning. These topics expose new aspects with which traditional approaches seem unable to cope.

In coping with these issues and finally achieving the ultimate goal, physical bodies play the important role of bringing the system into *meaningful* interaction with the physical environment – complex and uncertain, but with an automatically consistent set of natural constraints. This facilitates the correct agent design, learning from the environment, and rich meaningful agent interaction. The meanings of “having a physical body” can be summarized as follows:

1. Sensing and acting capabilities are not separable, but tightly coupled.
2. In order to accomplish the given tasks, the sensor and actuator spaces should be abstracted under resource-bounded conditions (memory, processing power, controller etc.).

3. The abstraction of the sensor and actuator spaces depends on both the fundamental embodiments inside the agents and the experiences (interactions with their environments).
4. The consequence of the abstraction is the agent-based subjective representation of the environment, and it can be evaluated by the consequences of behaviors.
5. In the real world, both inter-agent and agent-environment interactions are asynchronous, parallel and arbitrarily complex. There is no justification for adopting a particular top-down abstraction level for simulation such as a global clock 'tick', observable information about other agents, modes of interaction among agents, or even physical phenomena like slippage, as any seemingly insignificant parameter can sometimes take over and affect the global multi-agent behavior.
6. Natural complexity of physical interaction automatically generates reliable sample distributions of input data for learning, rather than from an *a priori* Gaussian distribution in simulations which does not always correctly capture the characteristics of the system.

Even though we should advocate the importance of “having a physical body,” it seems required to show that the system performs well coping with new issues in a concrete task domain. In other words, we need a standard problem which people regard as a new one that expose various various aspects of intelligent behaviors in real worlds.

RoboCup (The World Cup Robot Soccer Games:[Kitano, *et al.*, 95; Kitano, *et al.*, 97]) is an attempt to promote AI and robotics research by providing a common task for evaluation of various theories, algorithms, and agent architectures, and was proposed as a new standard problem. Not only the integration of a wide range of technologies but also accomplishment of technical breakthroughs are required for the robots to play a soccer game reasonably well. RoboCup consists of three competition tracks: (1) **Real Robot League**: using physical robots to play soccer games, (2) **Software Robot League**: using software agents to play soccer games on an official soccer server over the network, and (3) **Expert Robot Competition**: competition of robots which have special skills, but are not able to play a game.

In this paper, we propose the RoboCup Physical Agent Challenges as new research issues of physical agents. First, we show how the challenge is significant for physical agent research with long range issues. Then, we show mid- and short-term issues which spans from simple skill acquisition to a simple teamwork behavior. In particular, we pick two single agent skills (ball moving and ball catching) and one cooperative skill (passing the ball between two players) with different situations (no obstacles, stationary or moving obstacles) as the RoboCup physical agent challenge Phase I. We describe how these techniques can be evaluated in terms of various kinds of design issues.

## 2. Research Issues of RoboCup Physical Agent Track

RoboCup physical agent challenges are summarized as follows:

- **Perception**: The player should observe in real-time the behaviors of other objects which it cannot completely predict or control in order to take an action to deal with them. Such objects include a ball, an opponent, and in some sense, a common-side player, and further a judge. Capabilities of wide range perception, discrimination of other agents, and estimation of their locations and motions coping with occlusions are needed. Such perception is a basic technology to expand robotic applications.
- **Actions**: Controlling a ball in a multi-agent situation introduces a new set of problems even in robotics, where traditionally the manipulated objects are mostly stationary or follow predictable trajectories. Some of the immediate challenges are discussed in depth in the following sections. Long term challenges including: a) controlling a 3D (bouncing or flying) motion of a ball, b) soccer playing with limbed (quadrupeds, hexapods, or even bipeds) robots, and c) generating "faint" actions assuming that the opponent has a capability of action observation and prediction.

- **Situation and Behavior:** The task domain itself is simple, but infinitely many situations can occur in accordance with dynamic changes of the relationships in position and relative motion between the ball, the goal, and the players, and the context of the game. The optimal behavior changes from one situation to another. Since our goal is a soccer playing team, we need abilities beyond simple reflexive behaviors. For example, we need situation understanding, tactics selection and modification, minimum communication with teammates, teamwork behaviors acquired through practical training. These issues are closely related to the cognitive issues such as organization of spatio-temporal memory of the world and categorization of sets of motor behaviors into skills (symbols) grounded by the physical bodies.
- **Real-Time:** Since the situation rapidly changes according to motions of the ball and other players, there is no time to carefully analyze the situation and deliberate a plan. Therefore, the player should take an action such as kicking a ball immediately or dribbling it if surrounded by no opponents in order to wait for a better situation, or moving to a certain position in order to facilitate future or collaborative actions.
- **Platform:** The initial barrier for physical agent track challengers would be to build a team of robotic platforms for soccer playing. A platform should have a power source, high speed and quick response mobility, a ball kicking capability, sensors for a ball, the goal, the soccer court, self position, and other players' position and movements, on-board and/or remote computers, and a wireless communication for at least accepting the judge's commands. Integrating all these into a size-constrained physical structure while achieving entire real-time performance is already a challenge which requires constrained optimization of a large number of design parameters. Participating in RoboCup requires producing and maintaining multiple such copies. A valuable short range challenge would be to propose a standard robotic platform design for robot soccer. Although this is not exactly an AI issue, such a standard platform will surely facilitate many AI researchers potentially wanting to test their ideas on real robots.

These challenges described above are significantly long-term ones along the way to realizing a good soccer-playing robot team. Some of them 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 that 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 physical agent track, mainly targeted for the mid term time span: (1) agent architecture in general, (2) implementation of real-time and robust sensing, (3) realization of stable and high-speed robot control, (4) sensor fusion, (5) behavior learning for multi agent environments, and (6) cooperation in dynamic environments.

The RoboCup Physical Agent Challenge should 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.

### 3. Overview of The RoboCup Physical Agent Challenge Phase I

For the RoboCup Physical Agent Challenge Phase I, we offer three specific challenges, essential not only for RoboCup but also for general mobile robotics research.

The fundamental issue for researchers who wish to build real robot systems to play a soccer game in RoboCup is how to obtain basic skills to control a ball in various kinds of situations. Typical examples are to shoot the ball into the goal, to intercept the ball from an opponent, and to pass the ball to a teammate. These

skills are needed to realize cooperative behaviors with teammates and competitive ones against opponents in soccer game. Among basic skills to control the ball, we selected three challenges as the RoboCup Physical Agent Challenge Phase I:

1. moving the ball to the specified area with no, stationary, or moving obstacles,
2. catching the ball from an opponent or a teammate, and
3. passing the ball between two players.

These three challenges have many variations in different kinds of situations such as passing, shooting, dribbling, receiving, and intercepting the ball with/without opponents whose defensive skills vary from amateur to professional levels. Although they seem very specific to RoboCup, these challenges can be regarded as very general tasks in the field of mobile robotics research in a flat terrain environment. Since target reaching, obstacle avoidance, and their coordination are basic tasks in the area, the task of shooting the ball while avoiding opponents that try to block the player should be ranked as one of the most difficult challenges in the area. Once the robot succeeds in acquiring these skills, it can move anything to anywhere.

In another aspect, these three challenges can be regarded as a sequence of one task which leads to an increase of the complexity of the internal representation according to the complexity of the environment [Asada, 1996].

In the case of visual sensing, the agent can discriminate the static environment (and its own body if observed) from others by directly correlating the motor commands the agent sent and the visual information observed during the motor command executions. In other words, such observation can be classified as a self and stationary environment. In contrast, other active agents do not have a simple and straightforward relationship with the self motions. In the early stage, they are treated as noise or disturbance because of not having direct visual correlation with the self motor commands. Later, they can be found as having more complicated and higher correlation (cooperation, competition, and others). As a result, the complexity is drastically increased especially since between the two there is a ball which can be stationary or moving as a result of self or other agent motions.

The complexities of both the environment and the internal representation of the agent can be categorized as a cognitive issue in general, and such an issue is naturally involved in this challenge. In the following, we describe the challenges more concretely.

## 4. Task-I: Ball Moving

### 4.1. Objectives

The objective of this challenge is to check how the most fundamental skill of moving a ball to the specified area under several conditions with no (**Level I**), stationary (**Level II**), or moving (**Level III**) obstacles in the field can be acquired in various kinds of agent architectures, and to evaluate merits and demerits of realized skills using the standard tasks.

The specifications of the ball and the surface of the field is an issue common to all the challenges in the physical agent track. In order to emerge various behaviors, the field surface should not be so rough as to prevent the ball from rolling, but not so smooth that there is no friction. The former would rule out kicking or passing, and the latter, dribbling.

Since there are few variations of the task environment in **Level I**, agent architecture, and sensing in particular, is a major focus. In **Level II** motion control is the central issue, and in **Level III** prediction of the motion of obstacles is the key issue.

## 4.2. Technical Issues

### 4.2.1. Vision

General computer vision and robot vision issues are too broad to deal with here. Finding and tracking independently moving objects (ball, players, judges) and estimating their motion parameters (2-D and further 3-D) from complicated background (field lines, goals, corner poles, flags waved by the supporters in the stadium) is too difficult for the current computer and robot vision technologies to completely perform in real-time.

In order to focus on skill acquisition, visual image processing should be drastically simplified. Discrimination by color information such as a red ball, a blue goal, a yellow opponent makes it easy to find and track objects in real-time [Asada *et al.*, 1996b]. Nevertheless, robust color discrimination is a tough problem because digitized signals are so naive against the slight changes of lighting conditions. In the case of remote (wireless) processing, increased noise due to environmental factors causes fatal errors in image processing. Currently, human programmers adjust key parameters used in discriminating colored objects on site. Self calibration methods should be developed, which will be able to expand the general scope of image processing applications.

Visual tracking hardware based on image intensity correlation inside a window region can be used to find and track objects from the complicated background by setting the initial windows [Inoue, *et al.*, 92]. Currently, a color tracking version is commercially available. As long as the initialized color pattern inside each window does not change much, tracking is almost successful. Coping with pattern changes due to lighting conditions and occlusions is one of the central issues in applying this type of vision hardware.

As long as the vision system can cope with the above issues, and capture the images of both the specified area (the target) and the ball, there might be no problem [Nakamura and Asada, 1995, Nakamura and Asada, 1996]. To prevent the agent from losing the target, and/or the ball (in **Level II** and **III**, obstacles, too), an active vision system with panning and tilting motions seems preferable, but this makes the control system more complicated and introduces the spatial memory organization problem for keeping track of lost objects. A more practical way is to use a wider-angle lens. One extreme of this sort is to use the omni-directional vision system to capture the image all around the agent. This sort of lens seems very useful not only for acquiring the basic skills but also for realizing cooperative behaviors in multi agent environments. Currently this type of lens is commercially available as spherical and hyperboloidal ones [Ishiguro, 96].

### 4.2.2. Other perception

In the case of other sensing strategies, the agent should find the ball, (in **Level II and III**, obstacles, too) and know what the target is. Beside vision, typical sensors used in mobile robot research are range finders (e.g., sonars) and contact sensors (e.g., bumpers). However, it seems difficult for each or any combination among them to discriminate the ball (obstacles, too in higher levels) and the target unless special equipment such as transmitter is positioned inside the ball or the target, or a global positioning system besides on-board sensing and communication lines are used to inform the positions of all agents. The simplest case is no on-board sensing but only a global positioning system, which is adopted in the small robot league in the physical agent track because on-board sensing facilities are limited due to its size regulation.

In **Level II** and **III**, requirements include an obstacle avoidance behavior and the coordination of this behavior with a ball-carrying (or passing/shooting) behavior. One good strategy is assign the sensor roles in advance. For example, sonar and bumper sensors are used for obstacle avoidance while vision sensor is used for the target reaching. One can make the robot learn to assign the sensor roles [Nakamura *et al.*, 1996].

### 4.2.3. Action

As described in section 2, total balance of the whole system is a key issue to the robot design. In order for the system to facilitate various kinds of behaviors, a more complicated mechanical system and its sophisticated control techniques are necessary. We should start with a simpler one and then step up. The simplest case is to use just a car-like vehicle which has only two DOFs (degrees of freedom, for example forward and turn), and pushes the ball to the target (dribbling).

The target can be just a location, the goal (shooting), and one of the teammates (passing). In the case of location, a dribbling skill to carry the ball to the location might be sufficient. In the latter case, the task is to kick the ball into the desired direction without caring about final position of the ball. To discriminate it from a simple dribbling skill, we may need more DOFs to realize a kick motion with one foot (or what we may call arm). In the case of passing, velocity control of the ball might be a technical issue because one of the teammates to be passed to is not stationary but moving.

In **Levels II** and **III**, requirements include an obstacle avoidance behavior and the coordination of this behavior with a ball-carrying (or passing/shooting) behavior. To smoothly switch two behaviors, the robot should slow down, which increases the possibility that the opponent will take the ball. To avoid these situations, the robot can quickly switch behaviors, which causes instability of the robot motion. One might use the omni-directionally movable vehicle based on a sophisticated mechanism [Asama *et al.*, 1996]. For example, the vehicle could move to any direction anytime. In addition to the motion control problem, there are more issues to be considered such as how to coordinate these two behaviors (switching conditions) [Uchibe *et al.*, 1996].

### 4.2.4. Mapping from perception to action

There are several approaches to implementing the control mechanisms which perform the given task. A conventional approach is first to reconstruct the geometrical model of the environment (ball, goal, other agents etc.), then deliberate a plan, and finally execute the plan. However, this sort of approach is not suitable for the dynamically changing game environment due to its time-consuming reconstruction process.

A look-up table (LUT) indicating the mapping from perception to action by whatever method seems suitable for quick action selection. One can make such an LUT by hand-coding given a priori, precise knowledge of the environment (the ball, the goals, and other agents) and the agent model (kinematics/dynamics). In a simple task domain, a human programmer can do that to some extent, but seems difficult to cope with all possible situations completely. An opposite approach is learning to decide action selection given almost no *a priori* knowledge. Between exist several variations with more or less knowledge. The approaches are summarized as follows: (1) complete hand-coding (no learning), (2) parameter tuning given the structural (qualitative) knowledge (self calibration), (3) Subtask learning fit together in a layered fashion [Stone and Veloso, 1997] (4) typical reinforcement learning such as Q-learning with almost no a priori knowledge, but given the state and action spaces, (5) action selection from the state and action space construction [Asada *et al.*, 1996a, Takahashi *et al.*, 1996], and (6) tabula rasa learning. These approaches should be evaluated in various kinds of viewpoints.

## 4.3. Evaluation

In order to evaluate the achieved skills, we set up the following standard tasks with some variations.

1. move the ball to a specified circle of radius 25cm in the middle size class, and 8cm in the small size. Variations are different configurations of the ball near the robot and the target area.
2. move the ball to the goal whose size is specified in the regulations. Variations are the same as the above.

3. 1 and 2 with stationary obstacles (the number of obstacles changes from task to task) and moving obstacles (the number of obstacles is fixed, but the policy including speed can change from task to task).

The speed and the accuracy are the main issues in **Level I**.

## 5. Task-II: The Ball Catching

### 5.1. Objectives

The objective of this challenge is to check how the most fundamental skill of catching a ball under several conditions such as pass receiving (**Situation A**), goal keeping (**Situation B**), or intercepting (**Situation C**) can be acquired, and to evaluate merits and demerits of realized skills using the standard tasks.

### 5.2. Technical Issues

In addition to issues in the challenge (1), several other issues remain:

- **Situation A:** Prediction of the ball speed and direction is a key issue to receive the ball. To receive the passed ball while moving, the relationship between the moving ball and the self motion should be made clear [Stone and Veloso, 1997].
- **Situation B:** In addition to the above issue, goal protection is important. To estimate the goal position, the agent may have to watch the goal area lines and the penalty area line. Again, the omni directional lens is much better to see the coming ball and the goal position simultaneously. In the goal area line, the agent can receive and keep the ball while outside this line it may have to kick the ball (not receiving but just protecting the goal). Discrimination of these lines might cause the vision to be much more complicated.
- **Situation C:** It seems similar to **Situation A**, but the main difference is to get the ball from the pass between opponents. This requires more accurate prediction of motions of not only the ball but also opponents (both passer and receiver). Also, an additional load in perception module.

### 5.3. Evaluation

Although we can test Situations A and B with human players, we would like to test with the skill obtained in challenge (1). In Situation C, there are no opponents to pass the ball to each other, therefore we can increase the speed of the ball. After the following challenge, we can check both skills: passing and intercepting. Of course, both should be evaluated separately in advance.

## 6. Task-III: The Cooperative Behavior (Passing the Ball)

### 6.1. Objectives

The objective of this challenge is to check how the most fundamental cooperative skill (passing a ball between two players) can be acquired, and to evaluate merits and demerits of realized skills using the standard tasks.

This task focuses on a basic skill of cooperative behavior between two agents while task-I and -II focus on the basic skills of one single agent even though the environment includes other agents (possible opponents). If task-I and -II are successfully achieved, passing the ball between two players might become easy. That would mean a combination of passing and receiving skills. However, from a viewpoint of cooperative behaviors, there might be more issues.

## 6.2. Technical Issues

In addition to issue in the task-I and -II, three issues are left.

- Since the control architecture is not centralized but decentralized, each agent should know capabilities in passing and receiving skills of not only itself but the other. That is, the agent should estimate the level of the other agent skills. This means agent modeling. The problem includes partial observability due to perceptual limitations.
- Even though both agents have reasonable passing and receiving skills, the timing of passing and receiving must be coordinated. If both agents try to learn to improve their own behavior independently, the learning may not converge because the policy of the other changes simultaneously. To prevent this situation, one of the agents should be a coach (fixed policy) and the other, a learner. In this case, modeling of the learner is another issue for good teaching.
- In all these behaviors, and especially in a pass interception, the success rate will drastically increase if the agent can predict the ball holder's ball controlling behavior, specifically, when and in which direction it will kick. For this, the agent should first find a ball holder, track its motion, and predict the oncoming event based on the relative positions of the ball and the surrounding agents, such as the potential pass receivers. A primitive example of vision based interception of a static obstacle from another robot's trajectory has been demonstrated [Kuniyoshi, 95]. However, a general interception in fully dynamic situations like soccer playing is an open problem.
- Selection of passing direction depends on the motions of opponents. This introduces the opponent modeling issue which makes the cooperative behavior much harder to realize.

## 6.3. Evaluation

Since the challenge with many issues is very hard in the current stage, the Phase I challenge will only check cooperative behavior in a benign environment: two players with equal skills and with no opponents.

## 7. Discussion

### 7.1. Managing the Challenge

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

**Web Site:** On the RoboCup home page (<http://www.robocup.org/RoboCup/>), we are planning to show how to design a basic robot platform and technical information. The home page also provides a set of papers and technical documents related to RoboCup.



**Competitions and Conferences:** A series of RoboCup competition is planned to provide opportunities to test their ideas in conjunction with workshops at major international conferences, as well as local workshops, in order to facilitate exchange of information, discussions, and to feedback the status of the challengers to the overall framework of the challenge. As international events, we are planning to have RoboCup-98 Paris (with ICMAS-98 conference), 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.

## 7.2. Infrastructure Issues

One of the most important issues in the physical agent challenge is the design of the standard platform and their supply because the range of the physical agent designs is too wide to evaluate the skills achieved in the same level. If we had a precisely specified platform for RoboCup, the situation would be similar to the current simulation track because the design issues of mechanical structure would disappear. On the other hand, completely free design makes it hard to evaluate the skills achieved on equal terms. As a middle ground, we would like to have a platform that allows us to easily reconfigure the physical structure by which we expect physical agents to develop complex behaviors.

Recently, an open architecture called OPENR has been proposed as such a platform [Fujita and Kageyama, 1997]. OPENR has 1) standard interfaces between physical and software components and a programming framework, 2) configurable physical components with a common interface and information exchangers of their function and configurations, and 3) is constructed as a layered architecture based on object-oriented robot OS.

Expecting that an open architecture such as OPENR will be available for the RoboCup context in the very near future, we will offer basic resources and opportunities in order to facilitate technical progress based on the RoboCup physical agent challenge.

## 8. Conclusion

The RoboCup Physical Agent Challenge Phase I offers a set of three fundamental challenges, focused on ball moving, ball catching, and ball passing. These were chosen as the first set of challenges because they are essential technical issues for RoboCup, as well as for general robot system control and in particular for multi-agent control.

## References

- [Asada *et al.*, 1996a] M. Asada, S. Noda, and K. Hosoda. Action-based sensor space categorization for robot learning. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems 1996 (IROS '96)*, pages 1502–1509, 1996.
- [Asada *et al.*, 1996b] M. Asada, S. Noda, S. Tawaratumida, and K. Hosoda. Purposive behavior acquisition for a real robot by vision-based reinforcement learning. *Machine Learning*, 23:279–303, 1996.
- [Asada, 1996] Minoru Asada. An agent and an environment: A view on “having bodies” - a case study on behavior learning for vision-based mobile robot -. In *Proceedings of 1996 IROS Workshop on Towards Real Autonomy*, pages 19–24, 1996.
- [Asama *et al.*, 1996] H. Asama, M. Sato, N. Goto, H. Kaetsu, A. Matsumoto, and I. Endo. Mutual transportation of cooperative mobile robots using forklift mechanisms. In *Proc. of IEEE Int. Conf. on Robotics and Automation*, pages 1754–1759, 1996.

- [Brooks, 1991] R. A. Brooks. Elephants don't play chess. In P. Maes, editor, *Designing Autonomous Agents*, pages 3–15. MIT/Elsevier, 1991.
- [Inoue, *et al.*, 92] H. Inoue and T. Tachikawa and M. Inaba. Robot vision system with a correlation chip for real-time tracking, optical flow and depth map generation. In *Proc. of IEEE Int. Conf. on Robotics and Automation*, pages 1621–1626, 1992.
- [Ishiguro, 96] H. Ishiguro. <http://www.lab7.kuis.kyoto-u.ac.jp/vision/omni-sensor/omni-sensor.html>.
- [Kitano, *et al.*, 95] Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I. and Osawa, E., "RoboCup: The Robot World Cup Initiative", *IJCAI-95 Workshop on Entertainment and AI/Alife*, 1995
- [Kitano, *et al.*, 97] Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., Osawa, E., and Matsubara, H., "RoboCup: A Challenge AI Problem", *AI Magazine*, Spring, 1997.
- [Kuniyoshi, 95] Y. Kuniyoshi. Behavior Matching by Observation for Multi-Robot Cooperation. In G. Giralt and G. Hirzinger (eds.) *Robotics Research - The Seventh International Symposium*, pages 343–352, Springer, 1996.
- [Nakamura and Asada, 1995] T. Nakamura and M. Asada. Motion sketch: Acquisition of visual motion guided behaviors. In *Proc. of International Joint Conference on Artificial Intelligence*, pages 126–132, 1995.
- [Nakamura and Asada, 1996] T. Nakamura and M. Asada. Stereo sketch: Stereo vision-based target reaching behavior acquisition with occlusion detection and avoidance. In *Proc. of IEEE Int. Conf. on Robotics and Automation*, pages 1314–1319, 1996.
- [Nakamura *et al.*, 1996] T. Nakamura, J. Morimoto, and M. Asada. Direct coupling of multisensor information and actions for mobile robot behavior acquisition. In *Proc. of 1996 IEEE/SICE/RSJ International Conference on Multisensor Fusion and Integration*, pages 139–144, 1996.
- [Stone and Veloso, 1997] P. Stone, M. Veloso. Towards Collaborative and adversarial learning: A case study in robotic soccer. In *International Journal of Human-Computer Systems (IJHCS)*, 1997.
- [Stone and Veloso, 1997] P. Stone, M. Veloso. A Layered Approach to Learning Client Behaviors in the RoboCup Soccer Server In *Applied Artificial Intelligence (AAI) Journal*, 1997.
- [Takahashi *et al.*, 1996] Yasutake Takahashi, Minoru Asada, and Koh Hosoda. Reasonable performance in less learning time by real robot based on incremental state space segmentation. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems 1996 (IROS '96)*, pages 1518–1524, 1996.
- [Uchibe *et al.*, 1996] Eiji Uchibe, Minoru Asada, and Koh Hosoda. Behavior coordination for a mobile robot using modular reinforcement learning. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems 1996 (IROS '96)*, pages 1329–1336, 1996.
- [Fujita and Kageyama, 1997] Masahiro Fujita and Koji Kageyama. An Open Architecture for Robot Entertainment In *Proc. of First International Conference on Autonomous Agents*, pages 435–442, 1997.