

# CMUnited-97: RoboCup-97 Small-Robot World Champion Team

Manuela Veloso, Peter Stone, and Kwun Han  
Computer Science Department  
Carnegie Mellon University  
Pittsburgh, PA 15213  
{veloso,pstone,kwunh}@cs.cmu.edu

February 16, 1998

## Abstract

Robotic soccer is a challenging research domain which involves multiple agents that need to collaborate in an adversarial environment to achieve specific objectives. In this paper, we describe CMUnited, the team of small robotic agents that we developed to enter the RoboCup-97 competition. We designed and built the robotic agents, devised the appropriate vision algorithm, and developed and implemented algorithms for strategic collaboration between the robots in an uncertain and dynamic environment. The robots can organize themselves in formations, hold specific roles, and pursue their goals. In game situations, they have demonstrated their collaborative behaviors on multiple occasions. We present an overview of the vision processing algorithm which successfully tracks multiple moving objects and predicts trajectories. The paper then focusses on the agent behaviors ranging from low-level individual behaviors to coordinated, strategic team behaviors. CMUnited won the RoboCup-97 small-robot competition at IJCAI-97 in Nagoya, Japan.

## 1 Introduction

Problem solving in complex domains often involves multiple agents, dynamic environments, and the need for learning from feedback and previous experience. Robotic soccer is an example of such complex tasks for which multiple agents need to collaborate in an adversarial environment to achieve specific objectives. Robotic soccer offers a challenging research domain to investigate a large spectrum of issues of relevance to the development of complete autonomous agents [1, 6].

The fast-paced nature of the domain necessitates real-time sensing coupled with quick behaving and decision making. The behaviors and decision making processes can range from the most simple reactive behaviors, such as moving directly towards the ball, to arbitrarily complex reasoning procedures that take into account the actions and perceived strategies of teammates and opponents. Opportunities, and indeed demands, for innovative and novel techniques abound.

One of the advantages of Robotic Soccer is that it enables the direct comparison of different systems: they can be matched against each other in competitions. We have been pursuing research in the robotic soccer domain within the RoboCup initiative [5], which, in 1997, included a simulator league and small-size and medium-size robot leagues. We have been doing research extensively in the simulator league, developing learning techniques and team strategies in simulation [12, 14]. Many of these team strategies were directly incorporated into the robotic system described

here. We are currently also applying machine learning techniques to acquire hard to tune boundary behaviors for the real robots.

This paper describes the overall architecture of our robotic soccer team. The combination of robust hardware, real-time vision, and intelligent control represented a significant challenge which we were able to successfully meet. The work described in this paper is fully implemented as our CMUnited-97 RoboCup team. CMUnited-97 won the RoboCup-97 small-robot competition at IJCAI-97 in Nagoya, Japan. Our team scored a total of thirteen goals and only suffered one. Figure 1 shows a picture of our robotic agents.<sup>1</sup>

The specific contributions of the CMUnited-97 robot team include:

- Demonstration of a complete integration of perception, action, and cognition in a team of multiple robotic agents.
- A set of robust behaviors for individual agents. Each agent is equipped with skills that enable it to effectively perform individual and collaborative actions.
- Reliable perception through the use of a Kalman-Bucy filter. Sensing through our vision processing algorithm allows for (i) color-based tracking of multiple moving objects; (ii) and prediction of object movement, particularly the ball, even when inevitable sharp trajectory changes occur.
- Multiagent strategic reasoning. Collaboration between robots is achieved through: (i) a flexible role-based approach by which the task space is decomposed and agents are assigned subtasks; (ii) a flexible team structure by which agents are organized in *formations*, and homogeneous agents flexibly switch roles within formations; and (iii) alternative plans allowing for collaboration (e.g. passing to a teammate or shooting at the goal directly), are controlled by pre-defined metrics and are evaluated in real-time.




Figure 1: The CMUnited-97 robot team that competed in RoboCup-97.

---

<sup>1</sup>For the hardware description of our robots, see [15].

## 2 Real-Time Perception for Multiple Agents

The small-size robot league setup is viewed as an overall complete autonomous framework composed of the physical navigational robotic agents, a video camera over-looking the playing field connected to a centralized interface computer, and several clients as the minds of the small-size robot players. Figure 2 sketches the building blocks of the architecture.

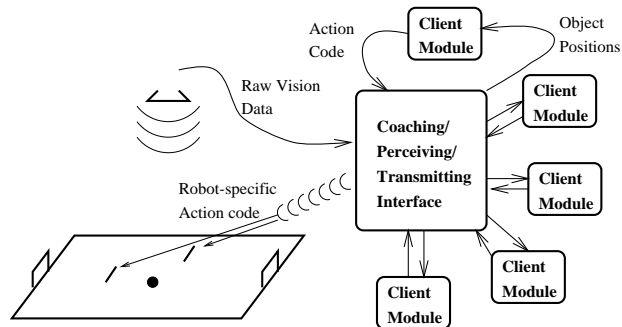


Figure 2: CMUnited architecture with global perception and distributed reaction.

The complete system is fully autonomous consisting of a well-defined and challenging processing cycle. The global vision algorithm perceives the dynamic environment and processes the images, giving the positions of each robot and the ball. This information is sent to an off-board controller and distributed to the different agent algorithms. Each agent evaluates the world state and uses its strategic knowledge to decide what to do next. Actions are motion commands that are sent by the off-board controller through radio communication. Commands can be broadcast or sent directly to individual agents. Each robot has an identification binary code that is used on-board to detect commands intended for that robot. This complete system is fully implemented.

Although it may be possible to fit an on-board vision system onto robots of small size, in the interest of being able to quickly move on to strategic multiagent research issues, we have opted for using a global vision system. It is part of our on-going research to also investigate and develop teams of robots capable of local perception [7, 10]. Part of our challenge in developing approaches to individual robot autonomy will consist of combining different sources of perception, namely local sensing, and targeted and broadcasted communication.

The fact that perception is achieved by a video camera that over-looks the complete field offers an opportunity to get a global view of the world state. Although this setup simplifies the sharing of information among multiple agents, it presents a challenge for reliable and real-time processing of the movement of multiple moving objects—in our case, the ball, five agents on our team, and five agents in the opponent team.

This section focusses on presenting our vision processing algorithm whose accuracy makes it a major contribution towards the success of our team.

### 2.1 Detection

The vision requirements for robotic soccer have been examined by different researchers. Small-sized and medium-sized robotic soccer researchers investigate on-board and off-board vision processors respectively [2, 8, 9, 10]. Due

to the reactivity of soccer robots both frameworks require a high perception processing cycle time. And, due to the rich visual input, dedicated processors or even DSPs have been used.

The vision system we successfully used at RoboCup-97 was surprisingly simple, consisting of a framegrabber with framerate transfer from a 3-CCD camera. The processing is done by a relatively slow processor (166MHz Pentium).

The detection mechanism was kept as simple as possible. The RoboCup rules have well defined colors for different objects in the field and these were used as the major cue for object detection. The RoboCup rules specify a green color field with white markings at the side. Also, it specifies a yellow or blue colored circular area on the top of the robots, one color for each team. A single color patch on the robot is not enough to provide orientation information. Thus, we added an additional second colored patch (pink) on top of each robot. The ball is an orange golf ball. We are able to differentiate these colors in a straightforward manner in color-space.

The set of detected patches is unordered. The detected color patches on the tops of the robots are then matched by their distance. Knowing the constant distance between the team-color and the pink orientation patch, we match patches that are this distance apart. Two distance-matched patches are marked as a robot capturing its position and orientation.

Noise is inherent in all vision systems. False detections in the current system are often of a magnitude of 100 spurious detections per frame. The system attempts to eliminate false detection using two different methods. First, color patches of size not matching the ones on the robots are discarded. This technique filters off most “salt and pepper” noise. Second, adding the distance matching mechanism briefly described above, all false detections are eliminated.

## **2.2 Data Association**

Data association addresses the problem of retaining robot identification in subsequent frames. One obvious approach to differentiate a number of robots using color-based detection is to use that number of different colors. However, with five robots, it is not simple to find five robustly distinguishable colors, as several colors are assigned to shared objects, such as green for the field, orange for the ball, white for the field markings, and blue and yellow for the team colors. Furthermore the inevitable variations in lighting conditions over the area of the field are enough to make a detection and association mechanism fully based on separable colors unreliable. We fit therefore each of the robots with the same color tops and no attempts are made to differentiate them via color.

Our data association approach solves the problem of retaining robot identification in subsequent frames given that all of the robots have the same color marker. We devised a greedy algorithm to retain association based on the spatial locations of the robots. During consecutive frames, association is maintained by searching using a minimum distance criterion. Current robot positions are matched with the closest positions from the previous frame, taking into account the size of the robots and an estimate of their velocity. The algorithm is robust to noisy detections, but in theory it is not guaranteed to find the optimal correct matches [3]. However in practice our detection and association approach is highly reliable.

## 2.3 Tracking and Prediction

In the setting of a robot soccer game, the ability to detect merely the locations of objects on the field is often not enough. Like for real soccer players, it is essential for robots to predict future locations of the ball (or even of the other players). We have used an Extended Kalman filter (EKF) for such a purpose [4]. The Kalman-Bucy filter is very suitable for such a purpose since the detection of the ball's location is noisy.

The EKF is a recursive estimator for a possibly non-linear system. The goal of the filter is to estimate the state of a system. The state is usually denoted as an  $n$ -dimensional vector  $x$ . A set of equations is used to describe the behavior of the system, predicting the state of the system as:  $x_{k+1} = f(x_k, u_k, w_k)$ , where  $f(\cdot)$  is a non-linear function which represents the behavior of the non-linear system,  $u_k$  is the external input to the system and  $w_k$  is a zero-mean, Gaussian random variable with covariance matrix  $Q_k$ ;  $w_k$  captures the noise in the system and any possible discrepancies between the physical system and the model; and  $k$  denotes time.

The system being modeled is being observed (measured). The observations can also be non-linear:  $z_k = h(x_k, v_k)$ , where  $z_k$  is the vector of observations and  $h(\cdot)$  is the non-linear measurement function, and  $v_k$  is another zero-mean, Gaussian random variable with covariance matrix  $R_k$ , which captures any noise in the observation process.

The EKF involves a two-step iterative process, namely *update* and *propagate*. The current best estimate of the system's state  $\hat{x}$  and its error covariance is computed on each iteration. During the update step, the current observations are used to refine the current estimate and recompute the covariance. During the propagate step, the state and covariance of the system at the next time step are calculated using the system's equations. The process then iteratively repeats, alternating between the update and the propagate steps.

Through a careful adjustment of the filter parameters modelling the system, we were able to achieve successful tracking and, in particular prediction of the ball trajectory, even when sharp bounces occur [3]. Figure 3 shows a screen shot of the field as generated by the vision processing module. The ball predicted trajectory is shown as the white line off of the ball, the teammates are displayed as "Ts" to represent their orientation, and opponents are circles. The figure shows the trajectory of a teammate robot.

Our vision processing approach worked perfectly during the RoboCup-97 games. We were able to detect and track 11 moving objects (5 teammates, 5 opponents and the ball). The prediction of the movement of the ball provided by the EKF is used by several agent behaviors. In particular, it allows the goalkeeper to look ahead in time and predict the best defending position. During the game, no goals were suffered due to miscalculation of the predicted ball position.

## 3 Multiagent Strategy Control

We achieve multiagent strategy through the combination of accurate individual and collaborative behaviors. Agents reason through the use of persistent reactive behaviors that are developed to aim at reaching team objectives.

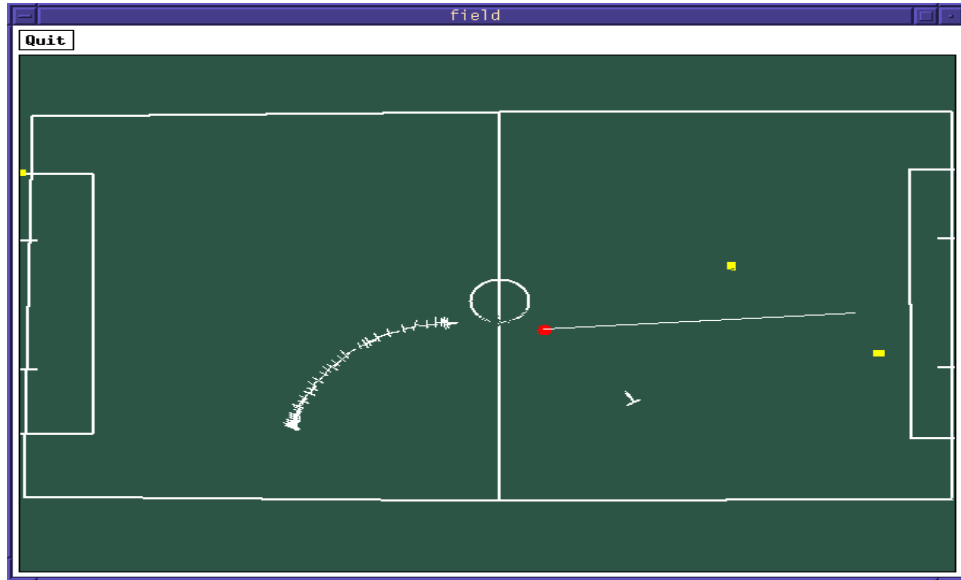


Figure 3: Field view from by the vision processing module.

### 3.1 Single-agent Behaviors

In order to be able to successfully collaborate, agents require robust basic skills. These skills include the ability to go to a given place on the field, the ability to direct the ball in a given direction, and the ability to intercept a moving ball. All of these skills must be executed while avoiding obstacles such as the walls and other robots.

The robot's hardware includes two motors which allow a robot to turn on itself. The front and the back of the robots are also absolutely equivalent in terms of navigation. Through these two features, the robots can therefore efficiently switch direction by turning at most  $90^\circ$ .

If a robot is to accurately direct the ball towards a target position, it must be able to approach the ball from a specified direction. Using the ball prediction from the vision system, the robot aims at a point on the far side of the target position. The robots are equipped with two methods of doing so:

- **Ball collection:** Moving behind a ball and knocking it towards the target.
- **Ball interception:** Waiting for the ball to cross its path and then intercepting the moving ball towards the target.

When using the ball collection behavior, the robot considers a line from the target position to the ball's current or predicted position, depending on whether or not the ball is moving. The robot then plans a path to a point on the line and behind the ball such that it does not hit the ball on the way and such that it ends up facing the target position. Finally, the robot accelerates to the target. Figure 4(a) illustrates this behavior.

When using the ball interception behavior (Figure 4(b)), on the other hand, the robot considers a line from *itself* to the target position and determines where the ball's path will intersect this line. The robot then positions itself along this line so that it will be able to accelerate to the point of intersection at the same time that the ball arrives.

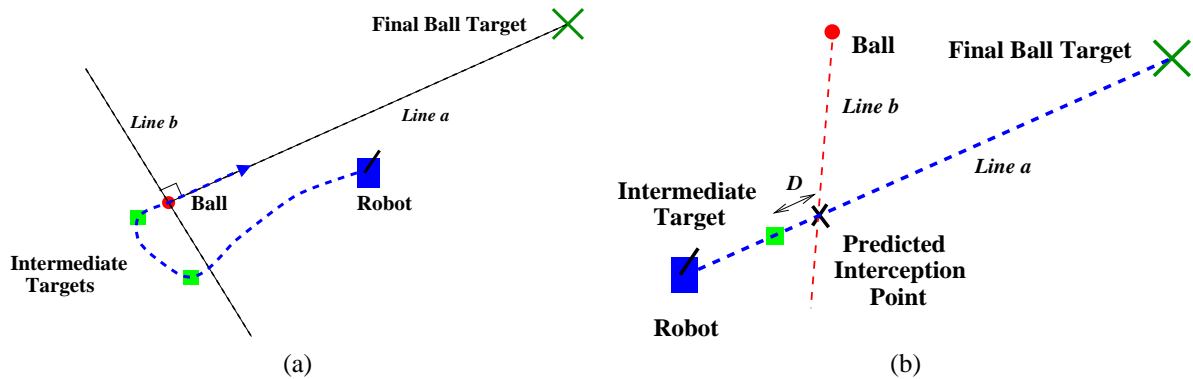


Figure 4: Single-agent behaviors to enable team collaboration (a) Ball collection (aiming for a pass or to the goal); (b) Ball interception (receiving a pass).

In practice, the robot chooses from between its two ball handling routines based on whether the ball will eventually cross its path at a point such that the robot could intercept it towards the goal. Thus, the robot gives precedence to the ball interception routine, only using ball collection when necessary. When using ball collection, it actually aims at the ball's predicted location a fixed time in the future so as to eventually position itself in a place from which it can intercept the ball towards the target.

## 3.2 Multiagent Behaviors

Although the single-agent behaviors are very effective when just a single robot is on the field, if all five robots were simultaneously chasing the ball and trying to shoot it at the goal, chaos would result. In order to achieve coordinated multiagent behavior, we organize the five robots into a flexible team structure.

The team structure, or *formation*, defines a set of roles, or *positions* with associated behaviors. The robots are then dynamically mapped into the positions.

Each robot is equipped with the knowledge required to play any position in each of several formations. The positions indicate the areas of the field which the robots should move to in the default situation. There are also different *active modes* which determine when a given robot should move to the ball or do something else instead. Finally, the robot with the ball chooses whether to shoot or pass to a teammate using a passing evaluation function.

These high-level, multiagent behaviors were originally developed in simulation and then transferred over to the robot-control code. Only the run-time passing evaluation function was redefined. Further details, particularly about the flexible team structures, are available in [11, 13].

### 3.2.1 Positions, Formations, and Active Modes

Positions are defined as flexible regions within which the player attempts to move towards the ball. For example, a robot playing the “right-wing” (or “right forward”) position remains on the right side of the field near the opponents’ goal until the ball comes towards it. Positions are classified as defender/midfielder/forward based on the locations of

these regions. They are also given behavior specifications in terms of which other positions should be considered as potential pass-receivers.

At any given time each of the robots plays a particular position on the field. However, each robot has all of the knowledge necessary to play any position. Therefore the robots can—and do—switch positions on the fly. For example, robots A and B switch positions when robot A chases the ball into the region of robot B. Then robot A continues chasing the ball, and robot B moves to the position vacated by A.

The pre-defined positions known to all players are collected into formations, which are also commonly known. An example of a formation is the collection of positions consisting of the goalkeeper, one defender, one midfielder, and two attackers. Another possible formation consists of the goalkeeper, two defenders and two attackers.

### 3.2.2 Run-time Evaluation of Collaborative Opportunities

One of CMUnited-97's main features is the robots' ability to collaborate by passing the ball. When in active mode, the robots use an evaluation function that takes into account teammate and opponent positions to determine whether to pass the ball or whether to shoot. In particular, as part of the formation definition, each position has a set of positions to which it considers passing. For example, a defender might consider passing to any forward or midfielder, while a forward would consider passing to other forwards, but not backwards to a midfielder or defender.

For each such position that is occupied by a teammate, the robot evaluates the pass to that position as well as evaluating its own shot. To evaluate each possible pass, the robot computes the *obstruction-free-index* of the two line segments that the ball must traverse if the receiver is to shoot the ball (lines b and c in Figure 5). In the case of a shot, only one line segment must be considered (line a). The *value* of each possible pass or shot is the product of the relevant obstruction-free-indices. Robots can be biased towards passing or shooting by further multiplying the values by a factor determined by the relative proximities of the active robot and the potential receivers to the goal. The robot chooses the pass or shot with the maximum value. The obstruction-free-index of line segment  $l$  is computed by the algorithm shown in Table 1 (variable names correspond to those in Figure 5).

Thus the obstruction-free-index reflects how easily an opponent could intercept the pass or the subsequent shot. The closer the opponent is to the line and the farther it is from the ball, the better chance it has of intercepting the ball.

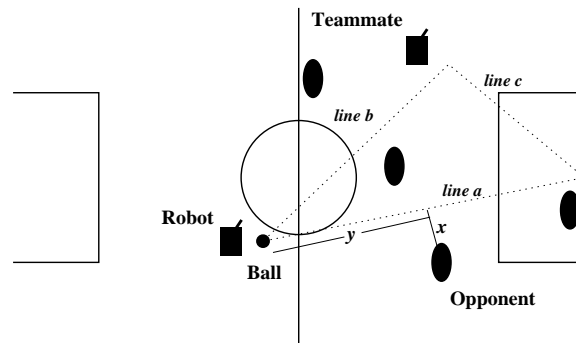


Figure 5: Run-time pass evaluation is based on the position of opponents.



1. *obstruction-free-index* = 1.
2. For each opponent *O*:
  - Compute the distance  $x$  from  $O$  to  $l$  and the distance  $y$  along  $l$  to  $l$ 's origin, i.e. the end at which the ball will be kicked by the robot (See Figure 5).
  - Define constants *min-dist* and *max-denominator*. Opponents farther than *min-dist* from  $l$  are not considered. When discounting *obstruction-free-index* in the next step, the  $y$  distance is never considered to be larger than *max-denominator*. For example, in Figure 5, the opponent near the goal would be evaluated with  $y = \text{max-denominator}$ , rather than its actual distance from the ball. The reasoning is that beyond distance *max-denominator*, the opponent has enough time to block the ball: the extra distance is no longer useful.
  - if  $x < \text{min-dist}$  and  $x < y$ ,  
 $\text{obstruction-free-index} *= x / \text{MIN}(\text{max-denominator}, y)$ .
3. return *obstruction-free-index*.

Table 1: Algorithm for the run-time evaluation of collaboration opportunities (pass or shoot).

### 3.2.3 The Goalkeeper

The goalkeeper robot has both special hardware and special software. Thus, it does not switch positions (or active modes) like the other robots. The goalkeeper's physical frame is distinct from that of the other robots in that it is as long as allowed under the RoboCup-97 rules (18cm) so as to block as much of the goal as possible. The goalkeeper's role is to prevent the ball from entering the goal. It stays parallel to and close to the goal, aiming always to be directly even with the ball's lateral coordinate on the field.

Ideally, simply staying even with the ball would guarantee that the ball would never get past the goalkeeper. However, since the robots cannot accelerate as fast as the ball can, it would be possible to defeat such a behavior. Therefore, the goalkeeper continually monitors the ball's trajectory. In some cases it moves to the ball's predicted destination point ahead of time. The decision of when to move to the predicted ball position is both crucial and difficult, as illustrated in Figure 6. Our goalkeeper robot currently take into account the predicted velocity and direction of the ball to select its moves.

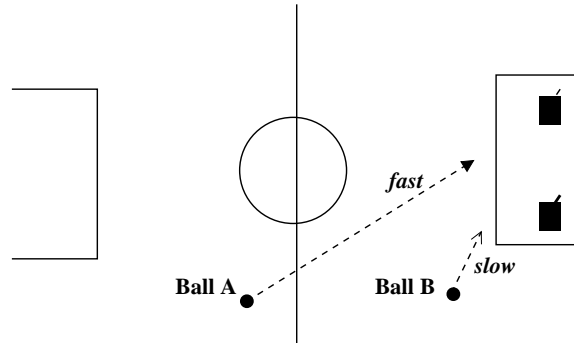


Figure 6: Goalkeeping.

## 4 Discussion and Conclusion

CMUnited-97 successfully demonstrated the feasibility and effectiveness of teams of multiagent robotic systems. Within this paradigm, one of the major challenges was to “close the loop,” i.e., to integrate all the different modules, ranging from perception to strategic multiagent reasoning. CMUnited is an example of a fully implemented multiagent system in which the loop is closed. In addition, we implemented interesting strategic behaviors, including agent collaboration and real-time evaluation of alternative actions.

It is generally very difficult to accumulate significant scientific results to test teams of robots. Realistically, extended runs are prohibited by battery limitations and the difficulty of keeping many robots operational concurrently. Furthermore, we only had the resources to build a single team of five robots, with one spare so far. Therefore, we offer a restricted evaluation of CMUnited based on the results of four effective 10-minute games that were played at RoboCup-97. We also include anecdotal evidence of the multiagent capabilities of the CMUnited-97 robotic soccer team.

The CMUnited-97 robot team played games against robot teams from Nara Institute of Science and Technology (NAIST), Japan; University of Paris VI, France (team name “MICROB”); and University of Girona, Spain. The results of the games are given in Table 2.

Opponent	Score
NAIST	5-0
MICROB	3-1
U. of Girona	2-0
NAIST (finals)	3-0
TOTAL	13-1

Table 2: The scores of CMUnited’s games in the small robot league of RoboCup-97. CMUnited-97 won all four games.

In total, CMUnited-97 scored thirteen goals, allowing only one against. The one goal against was scored by the CMUnited goalkeeper against itself, though under an attacking situation from France. We refined the goalkeeper’s goal behavior, as presented in Section 3.2.3, following the observation of our goalkeeper’s error.

As the matches proceeded, spectators noticed many of the team behaviors described in Section 3.2. The robots switched positions during the games, and there were several successful passes. The most impressive goal of the tournament was the result of a 3-way passing play: one robot passed to a second, which passed to a third, which shot the ball into the goal.

In general, the robots’ behaviors were visually appealing and entertaining to the spectators. Several people attained a first-hand appreciation for the difficulty of the task as we let them try controlling a single robot with a joystick program that we developed. All of these people (several children and a few adults) found it quite difficult to maneuver a single robot well enough to direct a ball into an open goal. These people in particular were impressed with the facility with which the robots were able to pass, score, and defend.

We are aware that many issues are clearly open for further research and development. We are currently systematically identifying them and addressing them towards our next team version. In particular, we are enhancing the robot's behaviors by using machine learning techniques.

## Acknowledgements

We thank Sorin Achim for developing and building the robots. This research is sponsored in part by the Defense Advanced Research Projects Agency (DARPA), and Rome Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-95-1-0018 and in part by the Department of the Navy, Office of Naval Research under contract number N00014-95-1-0591. Views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of the Air Force, of the Department of the Navy, Office of Naval Research or the United States Government.

## References

- [1] Minoru Asada, Yasuo Kuniyoshi, Alexis Drogoul, Hajime Asama, Maja Mataric, Dominique Duhaut, Peter Stone, and Hiroaki Kitano. The RoboCup physical agent challenge: Phase-i. *Applied Artificial Intelligence*, 12, 1998.
- [2] Minoru Asada, Shoichi Noda, Sukoya Tawaratumida, and Koh Hosoda. Purposive behavior acquisition for a real robot by vision-based reinforcement learning. *Machine Learning*, 23:279–303, 1996.
- [3] Kwun Han and Manuela Veloso. Reactive visual control of multiple non-holonomic robotic agents. In *Proceedings of the International Conference on Robotics and Automation*, 1998.
- [4] R. E. Kalman and R. S. Bucy. New results in linear filter and prediction theory. *Journal of Basic Engineering*, pages 95–108, March 1961.
- [5] Hiroaki Kitano, Yasuo Kuniyoshi, Itsuki Noda, Minoru Asada, Hitoshi Matsubara, and Eiichi Osawa. RoboCup: A challenge problem for AI. *AI Magazine*, 18(1):73–85, Spring 1997.
- [6] Hiroaki Kitano, Milind Tambe, Peter Stone, Manuela Veloso, Silvia Coradeschi, Eiichi Osawa, Hitoshi Matsubara, Itsuki Noda, and Minoru Asada. The RoboCup synthetic agent challenge 97. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, pages 24–29, San Francisco, CA, 1997. Morgan Kaufmann.
- [7] Maja J. Mataric. Designing and understanding adaptive group behavior. *Adaptive Behavior*, 4(1), December 1995.
- [8] Michael K. Sahota, Alan K. Mackworth, Rod A. Barman, and Stewart J. Kingdon. Real-time control of soccer-playing robots using off-board vision: the dynamite testbed. In *IEEE International Conference on Systems, Man, and Cybernetics*, pages 3690–3663, 1995.

- [9] Randy Sargent, Bill Bailey, Carl Witty, and Anne Wright. Dynamic object capture using fast vision tracking. *AI Magazine*, 18(1):65–72, Spring 1997.
- [10] Wei-Min Shen, Jafar Adibi, Rogelio Adobbati, Bonghan Cho, Ali Erdem, Hadi Moradi, Behnam Salemi, and Sheila Tejada. Building integrated mobile robots for soccer competition. In *Proceedings of the International Conference on Robotics and Automation*, 1998.
- [11] Peter Stone and Manuela Veloso. The CMUnited-97 simulator team. In Hiroaki Kitano, editor, *RoboCup-97: The First Robot World Cup Soccer Games and Conferences*. Springer Verlag, Berlin, 1998. In Press.
- [12] Peter Stone and Manuela Veloso. A layered approach to learning client behaviors in the RoboCup soccer server. *Applied Artificial Intelligence*, 12, 1998. In Press.
- [13] Peter Stone and Manuela Veloso. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Submitted to Artificial Intelligence*, 1998.
- [14] Peter Stone and Manuela Veloso. Using decision tree confidence factors for multiagent control. In Hiroaki Kitano, editor, *RoboCup-97: The First Robot World Cup Soccer Games and Conferences*. Springer Verlag, Berlin, 1998. In Press.
- [15] Manuela Veloso, Peter Stone, Kwun Han, and Sorin Achim. The CMUnited-97 small-robot team. In Hiroaki Kitano, editor, *RoboCup-97: The First Robot World Cup Soccer Games and Conferences*. Springer Verlag, Berlin, 1998.