

CMUNITED-97

RoboCup-97

Small-Robot World Champion Team

Manuela Veloso, Peter Stone, and Kwun Han

■ Robotic soccer is a challenging research domain that involves multiple agents that need to collaborate in an adversarial environment to achieve specific objectives. In this article, 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 that successfully tracks multiple moving objects and predicts trajectories. The article then focuses on the agent behaviors, ranging from low-level individual behaviors to coordinated, strategic team behaviors. CMUNITED won the RoboCup-97 small-robot competition at the Fifteenth International Joint Conference on Artificial Intelligence in Nagoya, Japan.

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 relevant to the development of complete autonomous agents (Asada et al. 1998; Kitano, Tambe, et al. 1997).

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 toward 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 (Kitano, Kuniyoshi, et al. 1997), 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 (Stone and Veloso 1998a, 1998d). 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 article describes the overall architecture of our small-size robotic soccer team. The combination of robust hardware, real-time vision, and intelligent control represented a significant challenge that we were able to successfully meet. The work described in this article is fully implemented as our CMUNITED-97 RoboCup team. CMUNITED-97 won the RoboCup-97 small-robot competition at the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97) in Nagoya, Japan. Our team scored a total of 13 goals and only suffered 1. Figure 1 shows a picture of our robotic agents.¹

The specific contributions of the CMUNITED-97 robot team include the following:

First is a demonstration of a complete integration of perception, action, and cognition in a team of multiple robotic agents.

Second is a set of robust behaviors for individual agents. Each agent is equipped with skills that enable it to effectively perform individual and collaborative actions.

Third is reliable perception through the use

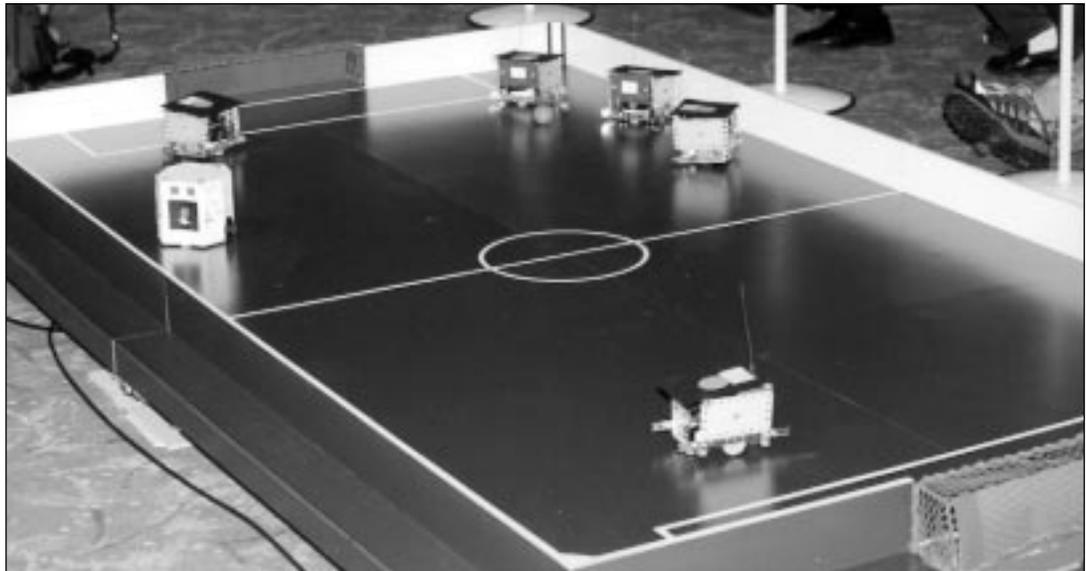


Figure 1. The CMUNITED-97 Robot Team That Competed in RoboCup-97.

of a Kalman-Bucy filter. Sensing through our vision-processing algorithm allows for color-based tracking of multiple moving objects and prediction of object movement, particularly the ball, even when inevitable sharp trajectory changes occur.

Fourth is multiagent strategic reasoning. Collaboration between robots is achieved through (1) a flexible role-based approach by which the task space is decomposed, and agents are assigned subtasks; (2) a flexible team structure by which agents are organized in formations, and homogeneous agents flexibly switch roles within formations; and (3) alternative plans allowing for collaboration (for example, passing to a teammate or shooting at the goal directly) that are controlled by predefined metrics and are evaluated in real time.

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 overlooking 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.

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 the robot. This complete system is fully implemented.

Although it might 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 opted for a global vision system. It is part of our ongoing research to also investigate and develop teams of robots capable of local perception (Shen et al. 1998; Mataric 1995). 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 overlooks 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 on the opposing team.

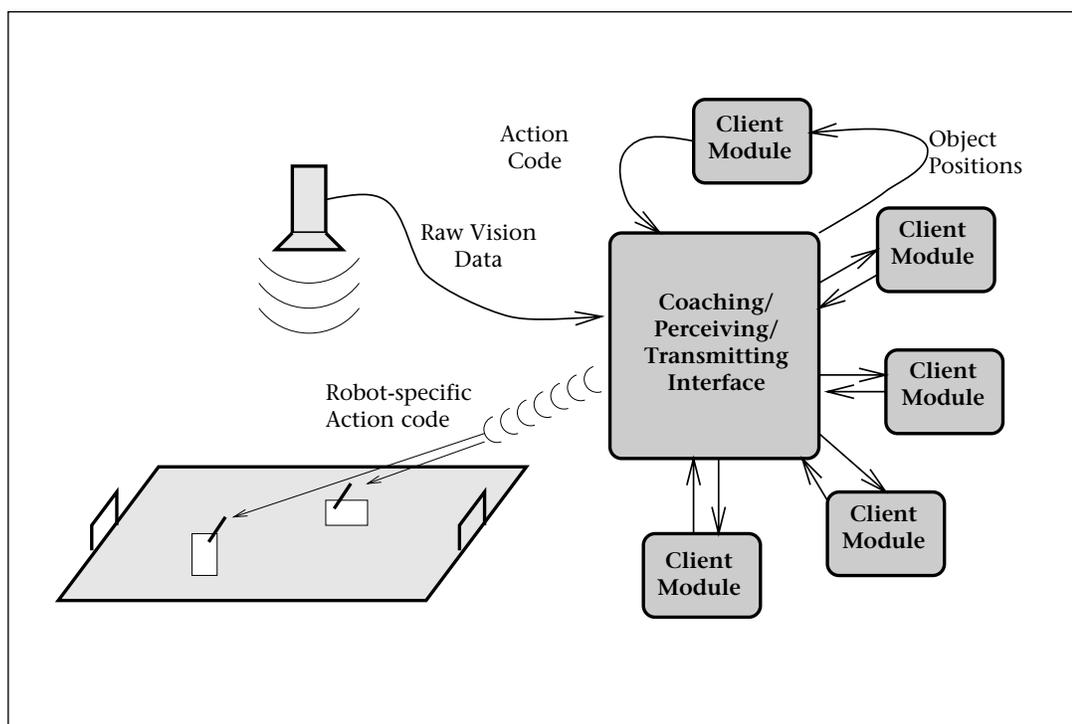


Figure 2. CMUNITED Architecture with Global Perception and Distributed Reaction.

This section focuses on presenting our vision-processing algorithm, whose accuracy makes it a major contribution toward the success of our team.

Detection

The vision requirements for robotic soccer have been examined by different researchers. Small-size and medium-size robotic soccer researchers investigate on-board and off-board vision processors, respectively (Shen et al. 1998; Sargent et al. 1997; Asada et al. 1996; Sahota et al. 1995). Because of the reactivity of soccer robots, both frameworks require a high-perception processing cycle time, and because of the rich visual input, dedicated processors or even digital signal processors have been used.

The vision system we successfully used at RoboCup-97 was surprisingly simple, consisting of a frame grabber with frame-rate transfer from a three-CCD camera.

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 pro-

vide orientation information. Thus, we added an additional colored patch (pink) on top of each robot. The ball is an orange golf ball. We are able to differentiate these colors in a robust 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 to each frame. The system attempts to eliminate false detection using two different methods: First, color patches of a size not matching the ones on the robots are discarded. This technique filters off most "salt and pepper" noise. Second, by adding the distance-matching mechanism briefly described earlier, all false detections are eliminated.

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

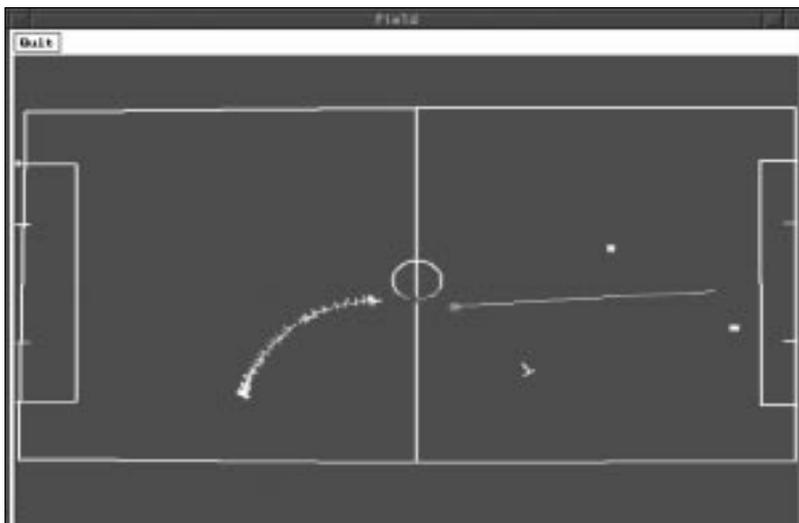


Figure 3. Field View from the Vision-Processing Module.

five robustly distinguishable colors because 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 may make a detection and association mechanism fully based on separable colors unreliable. Therefore, we fit each of the robots with the same color tops, and no attempts are made to differentiate them by color.

Our data-association approach solves the problem of retaining robot identification in subsequent frames, given that all 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 criterium. 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 (Han and Veloso 1998). However, in practice, our detection and association approach is highly reliable.

Tracking and Prediction

In the setting of a robotic soccer game, the ability to detect merely the locations of objects on the field is often not enough. Just as 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 (Kalman and Bucy 1961).

The EKF is a recursive estimator for a possibly nonlinear 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 nonlinear function that represents the behavior of the nonlinear 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 nonlinear: $z_k = h(x_k, v_k)$, where z_k is the vector of observations; $h(\cdot)$ is the nonlinear 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: (1) *update* and (2) *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 modeling the system, we were able to achieve successful tracking and, in particular, prediction of the ball trajectory, even when sharp bounces occur (Han and Veloso 1998). 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 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 because of miscalculation of the predicted ball position.

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.

Single-Agent Behaviors

To be able to successfully collaborate, agents require robust basic skills, including 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 these skills must be executed while the robot avoids obstacles such as the walls and other robots.

The robot's hardware includes two motors that allow it 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° .

If a robot is to accurately direct the ball toward 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: (1) *ball collection*, moving behind a ball and knocking it toward the target, and (2) *ball interception*, waiting for the ball to cross its path and then intercepting the moving ball toward 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 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 4a illustrates this behavior.

When using the ball-interception behavior (figure 4b), however, 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.

In practice, the robot chooses 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 toward 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 to even-

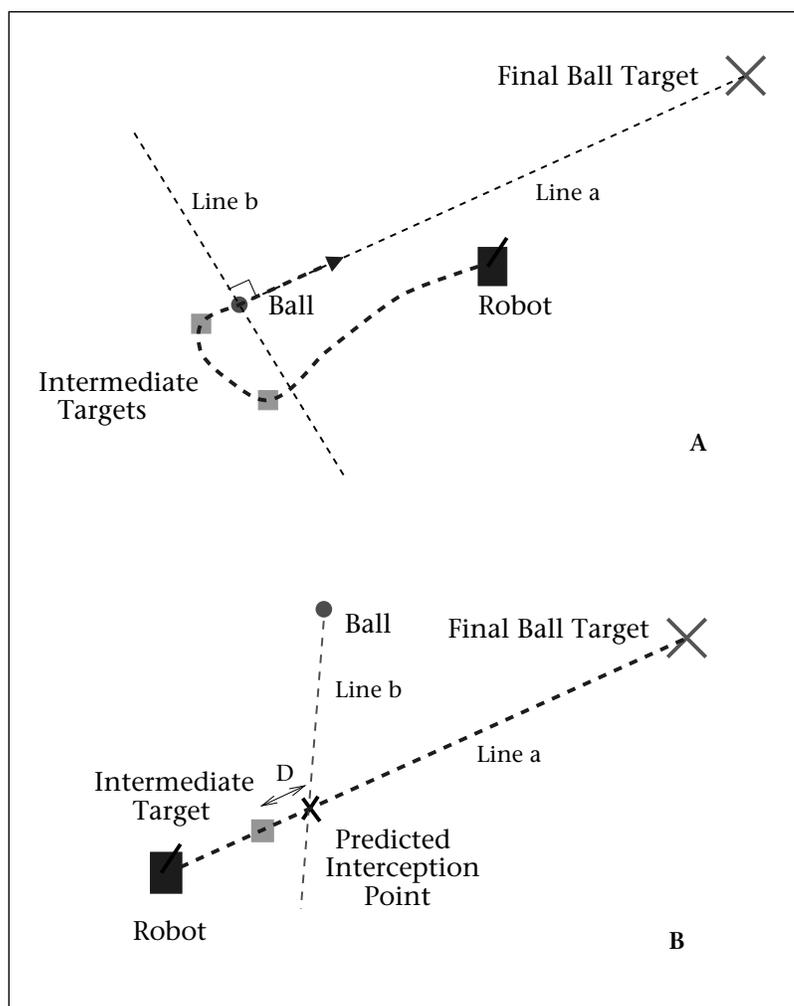


Figure 4. Single-Agent Behaviors to Enable Team Collaboration.

- A. Ball collection (aiming for a pass or the goal).
 B. Ball interception (receiving a pass).

tually position itself in a place from which it can intercept the ball toward the target.

Multiagent Behaviors

Although the single-agent behaviors are 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. 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 that the robots should move to in the default situation. There are also different active

*CMUNITED-97
successfully
demonstrated
the feasibility
and
effectiveness
of teams of
multiagent
robotic
systems.*

modes that 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 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 Stone and Veloso (1998b, 1998c).

Positions, Formations, and Active Modes

Positions are defined as flexible regions within which the player attempts to move toward 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 toward it. Positions are classified as defender, midfielder, or 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 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 predefined positions known to all players are collected into formations. 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.

Run-Time Evaluation of Collaborative Opportunities

One of CMUNITED-97's main features is the robots' ability to collaborate by passing the ball. The robots use an evaluation function that takes into account teammate and opponent positions to determine whether to pass the ball or 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, but a forward would consider passing to other forwards but not backward to a midfielder or defender.

For each such position that is occupied by a teammate, the robot evaluates the pass to the position as well as 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 6). 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 indexes. Robots can be biased toward 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 figure 5 (variable names correspond to those in figure 6).

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.

The Goalkeeper The goalkeeper robot has both special hardware and special software. Thus, it does not switch positions 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 (18 centimeters) 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 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, because 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 about when to move to the predicted ball position is both crucial and difficult, as illustrated in figure 7. Our goalkeeper robot currently takes into account the predicted velocity and direction of the ball to select its moves.

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, that is, to integrate all the different modules, ranging from perception to strategic multiagent reasoning. CMUNITED is an example of a fully implemented multiagent system in

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, that is, the end at which the ball will be kicked by the robot (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*.

Figure 5. Algorithm for the Run-Time Evaluation of Collaboration Opportunities (Pass or Shoot).

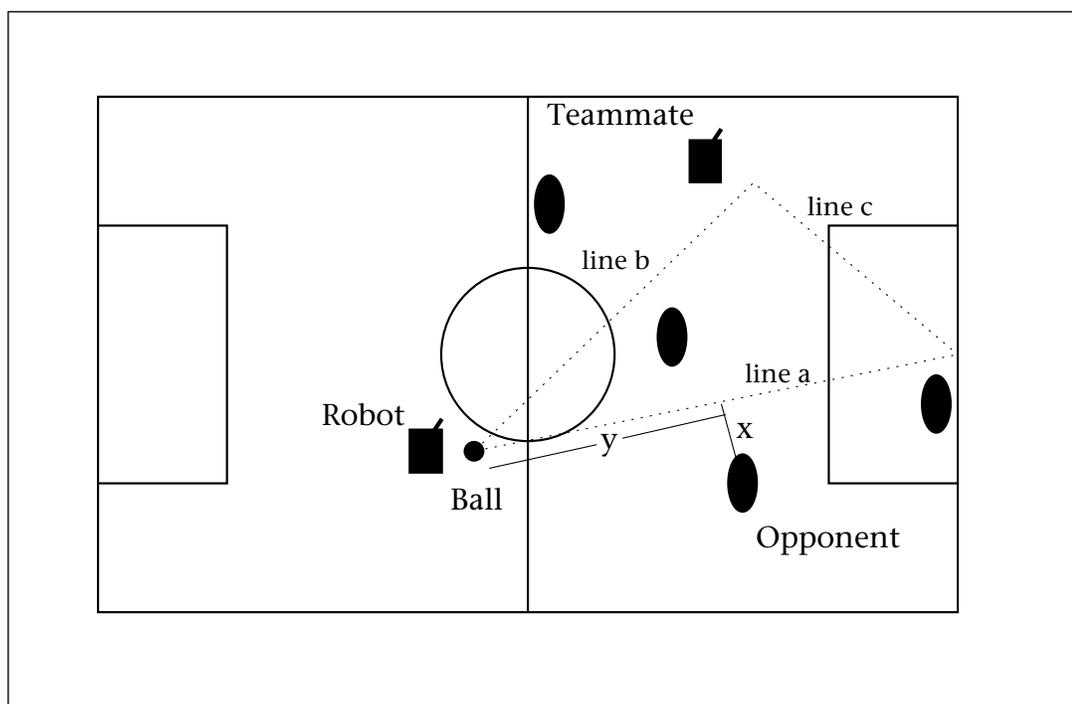


Figure 6. Run-Time Pass Evaluation Is Based on the Position of the Opponents.

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 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, to date, we have only had the

resources to build a single team of five robots, with one spare. Therefore, we offer a restricted evaluation of CMUNITED based on the results of 4 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 the Nara Institute of Science and Technology (NAIST), Japan; Uni-

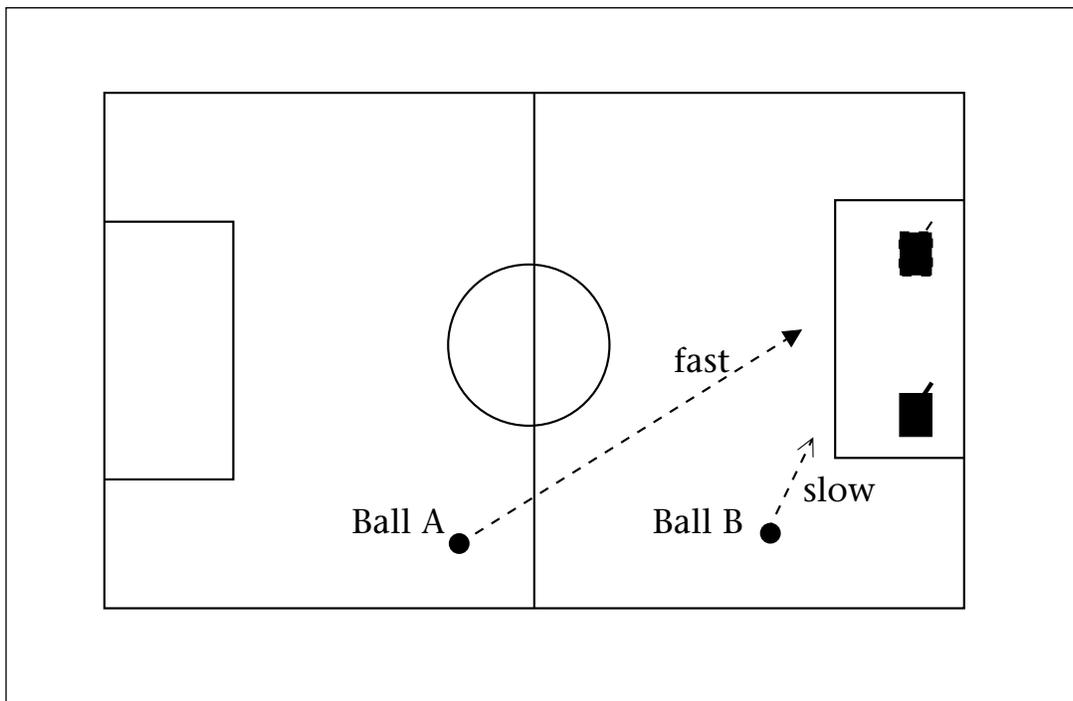


Figure 7. Goalkeeping.

Opponent	Score
NAIST, Japan	5-0
Paris 6, France	3-1
Girona, Spain	2-0
NAIST, Japan (finals)	3-0

Table 1. The Scores of CMUNITED-97's Games in the Small-Robot League of RoboCup-97.

CMUNITED-97 won all four games.

versity of Paris 6, France; and University of Girona, Spain. The results of the games are given in table 1.

In total, CMUNITED-97 scored 13 goals, allowing only 1 against. The one goal against was scored by the CMUNITED goalkeeper against itself, although under an attacking situation from France. We refined the goalkeeper's goal behavior, as presented in The Goalkeeper, following the observation of our goalkeeper's error.

As the matches proceeded, spectators noticed the team behaviors described in Multiagent Behaviors. 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 four-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 when we let them try controlling a single robot with a joystick program that we developed. All these people (several children and a few adults) found it 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 for our next team version.

Acknowledgments

We thank Sorin Achim for designing and building the robots. This research is sponsored in part by the Defense Advanced Research Projects Agency and Rome Laboratory, Air Force Materiel Command, United States Air Force, under agreement F30602-95-1-0018 and in part by the Department of the Navy, Office of

Naval Research, under contract 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, the Department of the Navy, the Office of Naval Research, or the U.S. government.

Notes

1. For the hardware description of our robots, see Veloso et al. (1998).

References

Asada, M.; Kuniyoshi, Y.; Drogoul, A.; Asama, H.; Mataric, M.; Duhaut, D.; Stone, P.; and Kitano, H. 1998. The RoboCup Physical Agent Challenge: Phase 1. *Applied Artificial Intelligence* 12:251–264.

Asada, M.; Noda, S.; Tawaratumida, S.; and Hosoda, K. 1996. Purposive Behavior Acquisition for a Real Robot by Vision-Based Reinforcement Learning. *Machine Learning* 23:279–303.

Han, K., and Veloso, M. 1998. Reactive Visual Control of Multiple Nonholonomic Robotic Agents. Paper presented at the International Conference on Robotics and Automation, May, Leuven, Belgium.

Kalman, R. E., and Bucy, R. S. 1961. New Results in Linear Filter and Prediction Theory. *Journal of Basic Engineering*, March, 95–108.

Kitano, H.; Kuniyoshi, Y.; Noda, I.; Asada, M.; Matsubara, H.; and Osawa, H. 1997. RoboCup: A Challenge Problem for AI. *AI Magazine* 18(1): 73–85.

Kitano, H.; Tambe, M.; Stone, P.; Veloso, M.; Coradeschi, S.; Osawa, E.; Matsubara, H.; Noda, I.; and Asada, M. 1997. The RoboCup Synthetic Agent Challenge 97. In Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, 24–29. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Mataric, M. J. 1995. Designing and Understanding Adaptive Group Behavior. *Adaptive Behavior* 4(1).

Sahota, M. K.; Mackworth, A. K.; Barman, R. A.; and Kingdon, S. J. 1995. Real-Time Control of Soccer-Playing Robots Using Off-Board Vision: The DYNAMITE Test Bed. In Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, 3690–3663. New York: IEEE Computer Society.

Sargent, R.; Bailey, B.; Witty, C.; and Wright, A. 1997. Dynamic Object Capture Using Fast Vision Tracking. *AI Magazine* 18(1): 65–72.

Shen, W.-M.; Adibi, J.; Adobbati, R.; Cho, B.; Erdem, A.; Moradi, H.; Salemi, B.; and Tejada, S. 1998. Building Integrated Mobile Robots for Soccer Competition. Paper presented at the International Conference on Robotics and Automation, May, Leuven, Belgium.

Stone, P., and Veloso, M. 1998a. A Layered Approach to Learning Client Behaviors in the RoboCup Soccer Server. *Applied Artificial Intelligence* 12(2–3): 165–188.

Stone, P., and Veloso, M. 1998b. Task Decomposition, Dynamic Role Assignment, and Low-Band-

width Communication for Real-Time Strategic Teamwork. *Artificial Intelligence*. Forthcoming.

Stone, P., and Veloso, M. 1998c. The CMUNITED-97 Simulator Team. In *RoboCup-97: Robot Soccer World Cup I*, ed. H. Kitano, 389–397. Berlin: Springer Verlag.

Stone, P., and Veloso, M. 1998d. Using Decision Tree Confidence Factors for Multiagent Control. In *RoboCup-97: Robot Soccer World Cup I*, ed. H. Kitano, 99–111. Berlin: Springer Verlag.

Veloso, M.; Stone, P.; Han, K.; and Achim, S. 1998. The CMUNITED-97 Small-Robot Team. In *RoboCup-97: Robot Soccer World Cup I*, ed. H. Kitano, 242–256. Berlin: Springer Verlag.



Manuela Veloso is associate professor of computer science at Carnegie Mellon University. She received her Ph.D. in computer science from Carnegie Mellon in 1992. She received a B.S. in electrical engineering in 1980 and an M.Sc. in electrical and computer engineering in 1984 from the Instituto Superior Tecnico in Lisbon as well as an M.A. in computer science in 1986 from Boston University. Veloso researches in the area of AI. Her long-term research goal is the effective construction of intelligent agents where cognition, sensors, and action are combined to address planning, execution, and learning tasks. In 1995, she received a National Science Foundation Career Award to pursue her research interests in autonomous robotic agents. Veloso's team of soccer robots won the RoboCup-97 competition. She was awarded the Allen Newell Medal for Excellence in Research in 1997.



Peter Stone is a Ph.D. candidate in the Computer Science Department at Carnegie Mellon University. He received his B.S. from the University of Chicago in 1993. He is on the organizing committee of RoboCup—the Robotic Soccer World Cup. His research interests include planning and machine learning, particularly in multiagent systems. His e-mail address is pstone@cs.cmu.edu.



Kwun Han is currently a Ph.D. candidate in computer science at Carnegie Mellon University. He received his B.Sc. with honors in computer science from Brown University in 1996. His current research interests include multiagent systems, machine learning, mobile robotics, robotic soccer, and computer vision.

Computers & Thought



Edward A. Feigenbaum
& Julian Feldman

EDITORS

ISBN 0-262-56092-5 560 pp., index. \$18.00 softcover

The AAAI Press • Distributed by The MIT Press
Massachusetts Institute of Technology, 5 Cambridge Center, Cambridge, Massachusetts 02142
To order, call toll free: (800) 356-0343 or (617) 625-8569.
MasterCard and VISA accepted.