

Soccer-playing robots could lead to completely autonomous intelligent machines.

Autonomous Robot Soccer Teams



Manuela Veloso is professor of computer science at Carnegie Mellon University.

Manuela Veloso

The idea of autonomous robot soccer teams invariably inspires images and expectations that, ironically, remove us somewhat from the real concept they embody. Indeed, the underlying research goes well beyond entertaining soccer fans to the creation of completely autonomous intelligent robots. I am in the fortunate position of pursuing research in artificial intelligence (AI), a fascinating field of research started by Allen Newell and Herb Simon at Carnegie Mellon. In the late 1980s, Allen Newell announced that it was time for the subareas of AI to merge and create “complete intelligence agents” capable of perception, action, and cognition. I fully embraced this challenge as the subject of my research.

Robot soccer teams compete in matches called RoboCup, which set itself a challenge of creating a robot team that could beat a human soccer team in the World Cup in 2050. RoboCup competitions are organized in a way that advances the state of the art of AI and robotics. Every year, the leagues are revised and moved closer to reality. The final goal is for robots to coexist with humans in a common physical environment.

The research platforms defined for the RoboCup international competitions present many challenges: (1) the environment is only partially observable; (2) the effects of a player’s actions in the presence of opponents are uncertain and difficult to model; and (3) the cycle of perception, cognition, and action must run in real time. Soccer differs from other adversarial

scenarios in fundamental ways. In chess, for example, there are no uncertainties about the effects of a player's actions. When Kasparov played chess against Deep Blue, there was no uncertainty in the execution of the moves—no tables were shaken; no pieces fell accidentally. Real-time response required only a combination of deliberative planning and reactive execution.

Autonomous playing robots face many technical challenges. The robots function today in a color-coded world. The floor is green; the goals are yellow and blue; the ball is orange; the uniforms are red and blue; the field is marked with unambiguous colored landmarks. Unlike the real world, the entire environment is customized. The real world in all of its complexity will be incrementally addressed over time. Each time the leagues are revised, the teams move a bit closer toward realization of the final goals. This year, for example, black and white balls will replace orange balls to determine whether robots can cope without the color.

Eventually robots will be able to cope with the real world—run in grass, function in rain or shine, and do all kinds of beautiful kicks through the air. The objective is to develop robots that can perceive and model the environment with which they are interacting and then respond to problems or changes in that environment in real time. Even in the color-coded world, we had to develop new segmentation and object-recognition algorithms capable of reliably and continuously processing images in real time.

To function as a true approximation of human intelligence, AI must encompass the idea of thinking forever, of deliberative planning, of when to stop thinking and start executing, and of assessing and learning from an execution to improve future executions. In other words, the goal is the integration of thinking, perceiving, and acting. Realization of the goal is decades away, but significant progress has been made in many dimensions, ranging from hardware and strategic teamwork to intelligent response to the world and other robots.

In the 2002 RoboCup competitions, in the laboratory, and in many demonstrations, our robots were active all the time, which represents a huge leap forward in terms of their reliability and robustness. The robots also move much more quickly today than they did even a few years ago. They are capable of maneuvering around obstacles, scoring goals, and localizing themselves, all autonomously (i.e., without remote control). Our robot team can cope with a large degree of uncertainty. Indeed, unlike real-world sports teams, they must

because they do not see their opponents prior to play. They have no videos of games, do not know what their opponents look like, and do not know if their opponents are adept at finding the ball, blocking, or any other aspect of the game. It's easy to get caught up in the excitement of the game, but we can also appreciate their performance from a technical view—their ability to play different roles as a team, to search continuously for the ball and chase it, to localize themselves and navigate in the field without getting lost, even if they are occasionally picked up by a referee.

A soccer robot's view of the world is utterly different from ours.

In 2002, we were able to gather a complete sequence of images of a robot's view of the world. The images left the research team at Carnegie Mellon speechless. The bouncing vision camera on the four-legged body of the robot captures an utterly different view of the world than the one we see: the field is upside down; the ball is not always round; objects change position and size radically with the motion of the robot; the ball actually disappears from view when it is near the robot. Overall, the images illustrate the challenge of processing perceptual data to be used by intelligent robots. Through our CMVision processing algorithm, robots effectively process such images and act based on the objects they recognize (Bruce et al., 2000).

One of the main functions of the robot is localizing itself on the field. Localization involves determining its position in the world. Human beings take for granted that they know where they are. Robots have no clue where they are unless algorithms have been written to help them filter the world and try to predict where they are going. How is this accomplished? The classical approach to localization uses a probability distribution for the robot's belief of its position. The distribution includes an a priori model of the robot's movement and a model of the field as a function of the sensory input (e.g., the fixed colored landmarks). When the robot moves, it updates its belief using the a priori model of its own movement. When the robot senses the landmarks,

it further updates the distribution according to the a priori model of the environment. When we tried this classical approach, however, the robot could not handle the large errors in the robot movement model.

At RoboCup'98, in Paris, the robots often became entangled with each other, so the referee lifted them up and put them down in a different location. At that point, the robots became completely lost because their perceptions of a different location did not match their locale belief because they had been moved but had not moved themselves (Veloso et al., 1998).



FIGURE 1 Robot soccer players on the playing grid.

In the classical approach to localization, robots being lifted up, pushed, or falling down are not accounted for. Earlier robots were big and were not moved around manually, so there was no need to localize algorithms. Robot soccer created the first small robots that execute a complete task. We devoted a great deal of research time to devising a new localization algorithm, called sensor resetting localization (SRL) that is capable of detecting “failure” in localization updates when the sensory information contradicts belief above a set threshold (Lenser and Veloso, 2000). SRL then abruptly creates a new hypothesis for the robot’s position based on the sensory data. With SRL, the robots can localize themselves despite inevitable errors in their movement models.

Once the robots have been equipped with robust vision and localization, they need to act to achieve their goals. Now that they see the world and they know where they are, they need to kick the ball in the right direction. How do they know what they are supposed to do? We call this a planning-behavior-based approach,

which our research has revealed must be a function of the robot’s confidence in its world model. This discovery led to “multifidelity behaviors,” in which a robot scores with different procedures as a function of how much it trusts its sensors (Winner and Veloso, 2000). If, for example, the robot has low confidence in its position on the field, it approaches the ball by a straight path. If it knows its position well, it can vary its approach to the ball to set itself behind the ball facing the opposing goal. Multifidelity behaviors are an innovation; no previous behavior architectures had explicit procedures for behavior as a function of a robot’s confidence in its world model.

The behavioral states transition among each other upon verification of conditions that test the visual perceptual input for specific environment states. For example, the robot transitions from searching for the ball to approaching it, if it can see the ball. Any image other than the ball is ignored. General perception becomes, therefore, “purposeful perception,” as the robot’s behavior-state machine focuses its attention only on specific perceptual conditions. Considering the images the robot actually sees, this purposeful perception explains how the robot can perform well. Even if it sees a series of apparently confusing images, it ignores everything except the specific conditions set at each state (e.g., the presence of the ball in the searching state).

Individual robots with real-time object recognition, effective SRL, and multifidelity behaviors can function as autonomous individual creatures. The next question to address is forming a team of robots. The first step in team organization is assigning roles, different behaviors to different members of the team (e.g., goalie, midfielders, offensive players, and defensive players). Robots can then be organized in formations. A team member, as a single robot, executes a particular role through a behavioral-state machine. Coordination among team members during real-time execution may require communication among them. However, communication may be expensive or not available, so we have devised coordination approaches that do not depend on real-time communication. We introduced predefined team plans, which we call “locker room agreements,” that encode coordination plans the robots can carry out as a team, triggered by universal world features that all of the robots can detect without communication (Stone and Veloso, 1999). Time and score, for example, are special world features the robots can

perceive without communicating with other team members. So, if a team is winning by more than two goals and there is only one minute left in the game, then the team moves to a defensive formation. The robots have been equipped with alternative, predefined plays they can execute as a team. They can actually assess the success of each play in the presence of different opponents and adapt to using the play most likely to succeed against a particular opponent.

Before 2002, the robots could not talk to each other and could see each other only in terms of recognition of the colors of their uniforms, which they could perceive. In 2002, the robots acquired wireless communication. Communication among members of a team creates opportunities for sharing information and for dynamic coordination. In a communicating team, the model of the environment does not have to be inferred from one robot's view of the world. Team members can share their views of the world to create a global world model. Therefore, even if one robot cannot see the ball, perhaps because the ball is too far away or is occluded, the robot may know the position of the ball through communication with its teammates.

Asynchronous communication in a highly dynamic environment like robot soccer inevitably leads to inconsistencies in the information shared. Two robots may communicate different ball positions, for example. Therefore, we developed an approach in which each individual robot keeps two separate world models, one that corresponds to its view of the world and one that merges information received from its teammates in terms of their positions and the position of the ball, as well as of the confidence in the shared information (Roth et al., 2003). The robot relies mostly on its individual world model and invokes the shared world model only when its confidence in its model is below a preset threshold.

Except for the goalie, which has a fixed role, the robots are prepared to switch roles dynamically and opportunistically during the course of a game. For example, the CMPack'02 team of Sony legged robots consists of four robots. One is a goalie, and the other three play the roles of primary attacker, offensive supporter, and defensive supporter. The robots coordinate in two separate phases. First, they assign roles to each other; then they position themselves on the field according to their assigned roles. For example, a primary attacker would move toward the ball; the offensive supporter would position itself in a supportive

attacking position; the defensive supporter would move closer to its own goal. Role assignment is achieved by the introduction of values functions computed based on the world model. Each robot can compute the value of each role for all of the robots as a function of their distance to the ball and their positions on the field. Roles are hence assigned through local computations based on the shared world model, thus eliminating the need for additional negotiations.

After a role has been assigned, the robots must position themselves as a function of their roles. We have devised two similar solutions for strategic positioning: a constraint-based objective optimization and a gradient-based potential field. For the supportive attacker, the objective function finds a position that maximizes the distance to the opponents and teammates and minimizes the distance to the ball and to the goal, under constraints (e.g., do not block the goal, do not compromise passes, etc.) (Veloso et al., 1999). The primary attacker goes to the ball, and the supporter moves to a good open position trying to maximize the chances of an emerging pass. Recently, we developed a similar potential-field-based approach that combines multiple repulsion and attraction points and allows the robots to navigate in the direction of the gradient of the field (Vail and Veloso, in press). Using this approach, our CMPack'02 team successfully coordinated and positioned itself, becoming the RoboCup'02 World Champions.

Communication among team members creates opportunities for sharing information and dynamic coordination.

Other research we are pursuing includes dynamic multirobot path planning, coaching, and multiagent learning. The algorithms we devised for path planning probabilistically combine past plans into the generation of new plans and allow for a smooth real-time execution of planned trajectories (Bruce and Veloso, 2002). Coaching addresses the challenging question of providing and following advice (Riley and Veloso, 2002). Multiagent learning enables an agent to learn in the

presence of other learning agents. We have introduced a learning principle that changes the learning rate as a function of whether the learner is winning or losing (Bowling and Veloso, 2002).

Human responses to robot behavior can be fascinating. Spectators, as well as researchers, cheer the robots on and get truly caught up in the game. This year, we wired a victory dance into the robots. The dance, which of course exists because humans programmed it to exist, does not represent consciousness on the part of the robots, nothing they see, perceive, or plan. Nevertheless, people respond to the victory dance because the robots appear to be expressing their emotions. By 2003, we plan to have five or six different dances the robots can select randomly, which will increase the illusion that they are creative and will elicit a stronger response. At a demonstration last year, a child asked if the robots wonder why people pick them up. Based on their autonomous behavior, people often infer that robots can do much more than they actually can. Indeed, they are currently only little soccer-playing robots. But in time they will surely become much more efficient.

The first RoboCup American Open for all the Americas will be held April 30 thru May 4, 2003, at Carnegie Mellon in Pittsburgh. The cognition and action involved in competitions between multirobot teams continues to be challenging scientifically and at the engineering level and will provide opportunities for research and development for years to come.

Acknowledgment

This research is part of Cooperate, Observe, Reason, Act and Learn (CORAL), a large research project at Carnegie Mellon. Videos and publications are available at <http://www.cs.cmu.edu/~coral>.

References

- Bowling, M., and M. Veloso. 2002. Multiagent learning using a variable learning rate. *Artificial Intelligence* 136: 215–250.
- Bruce, J., T. Balch, and M. Veloso. 2000. Fast and inexpensive color image segmentation for interactive robots. Pp. 2061–2066 in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*. Piscataway, N.J.: IEEE.
- Bruce, J., and M. Veloso. 2002. Real-time randomized path planning for robot navigation. Pp. 2383–2388 in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*. Piscataway, N.J.: IEEE.
- Lenser, S., and M. Veloso. 2000. Sensor resetting localization for poorly modeled mobile robots. Pp. 1225–1232 in *Proceedings of the International Conference on Robotics and Automation*. Piscataway, N.J.: IEEE.
- Riley, P., and M. Veloso. 2002. Planning for distributed execution through use of probabilistic opponent models. Pp. 72–81 in *Proceedings of the Sixth International Conference on Artificial Intelligence Planning Systems*, M. Ghallab, J. Hertzberg, and P. Traverso, editors. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Roth, M., D. Vail, and M. Veloso. 2003. A world model for multi-robot teams with communication. Submitted for publication.
- Stone, P., and M. Veloso. 1999. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence* 110(2): 241–273.
- Vail, D., and M. Veloso. In press. Dynamic multi-robot coordination. In *Multi-Robot Systems*, A. Schultz, L. Parker, and F. Schneider, editors. New York: Kluwer Academic Publishers.
- Veloso, M., P. Stone, and M. Bowling. 1999. Anticipation as a key for collaboration in a team of agents: a case study in robotic soccer. Pp. 134–141 in *Proceedings of SPIE Sensor Fusion and Decentralized Control in Robotic Systems II* (Volume 3839), G.T. McKee and P.S. Schenker, editors. Bellingham, Wash.: SPIE Press.
- Veloso, M., W. Uther, M. Fujita, M. Asada, and H. Kitano. 1998. Playing soccer with legged robots. Pp. 437–442 in *Proceedings of the IEEE International Conference on Robotics and Automation*. Piscataway, N.J.: IEEE.
- Winner, E., and M. Veloso. 2000. Multi-fidelity behaviors: acting with variable state information. Pp. 872–877 in *Proceedings of the National Conference on Artificial Intelligence of the American Association for Artificial Intelligence*. Menlo Park, Calif.: American Association for Artificial Intelligence.