EXECUTIVE INSIGHT

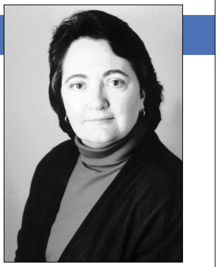

By Sara Hedberg
Emergent, Inc.
sara@hedberg.com

# Robots playing soccer? RoboCup poses a new set of AI research challenges

With the defeat of world chess champion Garry Kasparov by IBM's Deep Blue machine, one of AI's long-standing research challenges was at last achieved. A special award ceremony at this year's National Conference on AI (AAAI '97) honored the achievements of not only the Deep Blue programmers, but also the researchers whose work blazed the trail for Deep Blue.

On stage at AAAI '97 to receive awards for their seminal work in chess was an impressive gathering of computer giants such as Ken Thompson (yep, the Unix progenitor), John McCarthy, and Hans Berliner. Clearly, 40 years of innovation by many great researchers prepared the way for Deep Blue's crowning achievement. Now a new AI challenge is on the horizon: RoboCup, the robot world soccer cup.

## The RoboCup challenge

RoboCup is intended to stimulate innovations in a wide range of technologies and the integration of technologies into a fully functional robotic soccer team. Given how primitive today's robots are compared to a good soccer team of 11 players, the inherent technical challenges are plentiful.

To build a world-class roboteam will require

- capable single players (autonomous agents, robotics, vision, and real-time sensor fusion),
- teamwork (multiagent collaboration, context recognition),
- understanding the competition (cognitive modeling),
- the ability to develop and execute plays and strategies in real time (strategy acquisition, real-time reasoning and planning, and reactive behavior), and
- pre- and post-game training (machine learning),

to name a few items on the heady list of necessary functionality.

RoboCup-97, held in conjunction with the 1997 International Joint Conference on AI, in Nagoya, Japan, at the end of August, was a preliminary step toward meeting this challenge. For the most part, play level was more akin to that of first and second graders learning how to play than that of a proficient (let alone world-class) college team. Much like chess, achieving the scientific challenges posed by RoboCup will take years of work. "Chess took us 40 years," notes Hiroaki Kitano, founder of RoboCup and a senior researcher at the Sony Computer Science Lab. "With real robots, it will take us at least that long or longer."

## Three different leagues

This year's RoboCup featured three different leagues. The Middle-Size and Small Leagues involved physical robots; the Simulation League was for virtual, synthetic teams.

RoboCup organizers posited three specific challenges for the physical robots:

- moving the ball to the specified area (shooting, passing, and dribbling) with no, stationary, or moving obstacles;
- catching the ball from an opponent or a teammate (receiving, goal-keeping, and intercepting), and
- passing the ball between two players.

The first two involve individual agent skills, and the third introduces simple cooperative behavior. All are appropriate for beginner-level soccer.

Likewise, for the Simulation League, RoboCup issued a three-part Synthetic Agents Challenge '97, including

- *Learning*: offline skill learning by individual agents, offline collaborative learning by teams of agents, online skill and collaborative learning, and online adversarial learning.
- *Teamwork*: contingency planning for multiagent adversarial game playing, plan decomposition and merging, and executing team plans.
- *Opponent modeling*: online tracking of opponents' behavior and intentions, online strategy recognition by sideline coach agents, and offline review after the game.

## Limited results this year

These are lofty challenges indeed. In reality, the physical robots that competed were still quite brittle. They were sensitive to even the slightest changes in lighting and color conditions from their practice fields at home to those in the Nagoya competition field. Many bots had difficulty even finding the ball, and spent a good deal of time just standing around looking for it. Robots ran into each other, pushing each other around. "It looked more like a sumo wrestling match than soccer," observed Yumi Iwasaki, a senior research associate from Stanford University, who watched one match. Indeed, all the pushing led to some robots starting to smoke as motors overheated and began to melt plastic chassis.

The pace of play was more akin to walking, or even handicapped walking, than running. Nor did the bots have much agility. They had no appendage with which to kick the ball. All bots simply pushed the ball along. Plus, the teams in the physical

## Sean Luke wins Scientific Challenge Award: coevolves soccer teams using genetic programming

RoboCup's organizers established the Scientific Challenge Award with the intention that it be just as prestigious as actually winning the championship. The award is for scientific and technical innovation. This year's winner was Sean Luke, a graduate student and research assistant at the University of Maryland.

"Our approach is a very odd one," says Luke. But it was clearly one that caught the eye of the judges. Using hours and hours of CPU time on 16 nodes of an Alpha supercomputer (that included 10 machines with 4 nodes each) at the University of Maryland, Luke was able to evolve a competitive virtual soccer team of 11 players.

RoboCup organizer Hiroaki Kitano believes Luke's work is highly significant. "In the past, most genetic programming-based agent evolution was limited to a single agent, tested on rather simple tasks such as wall following or light-source tracking, and was in non-time-critical tasks," explains Kitano. Kitano is also the associate editor for *IEEE Transactions on Evolutionary Computing*. "So, I have a good feel for what is the current state of the art in the area. This is clearly the most complicated task of evolutionary computing/genetic programming ever undertaken. The potential impact of this research is very wide."

Luke spent about a year preparing for the event. He wanted to play other teams. "If we found out we were even remotely as good as others, we'd be happy." His team won a few and lost a few. "We were not looking to win. We wanted to push the boundaries of what evolutionary computation could do in a hard domain."

Evolutionary programming works much like genetic evolution. Out of a primordial soup of building blocks, the program builds generation after generation of individuals, assessing each individual's fitness against some goal. Individuals with better fitness are kept, and the next generation is formed by breeding, mutating, and reproducing the best individuals in the population.

"We made random individuals—each individual was a soccer team of 11 players," explains Luke. "They would play against each other in tournaments. The teams that did better based on the goals scored got higher fitness." Luke did not know much about soccer going into this work. At the outset, Luke and his colleagues didn't specify behavior, only the rules, and then ran experiments.

"We got some really hilarious teams," recounts Luke. "We got one team that was always going away from the ball. The other team was going away from the other team." Luke laughs as he describes one team inching away from the ball and the other team inching away from its competitors. Such avoidance behavior had low fitness and was taken out of the gene pool.

As teams began to evolve, early behavior was akin to first graders learning to play soccer. When six-year-olds start to play the game, it is a herd sport. Everyone follows the ball. Sometimes the ball pops out of the herd; then the herd moves to it. Occasionally it pops out and makes it into the goal by chance. It might be the opponent's goal or your own goal. There's no distinguishing at this level of play.

As the evolutionary process continued, teams began to exhibit behavior where players would initially assume positions not in a standard soccer arrangement, but symmetrically, such as a line stretching vertically down the field. Eventually, the players figured out how to disperse themselves. Luke wishes that some of the later-evolved teams had had "longer to cook," as he puts it. "We would have had better results yet."

As it was, he was able to run some 80,000 generations of teams, each composed of about 400 teams. The process started with about 400 teams. The program assessed fitness and bred the better performers. After 50 generations, the run would stop and pick the best of that run for breeding. Evaluation time alone was more than two months of CPU time, plus three to four months of the soccer simulation server.

"One of the things that made the project interesting was how the individual (teams) grow and learn. We found if they played against a really good team, no one got better. If they ramped up slowly, they got better. In this coevolutionary approach, they were initially all bad. As they got better, they were more able to play. It got tougher. A team got better relative to other teams. The whole league bootstrapped itself until it got quite good."

"But," he opines, "how amazingly better little kids are at soccer than even the best programs."

leagues showed very limited teamwork. There was much more sophisticated play in the smokeless Simulation League—an environment where researchers could focus on some of the higher-level tasks without worrying about lower-level issues such as robot controls and short battery life (see the sidebar on Sean Luke).

For such a young, immature event, RoboCup elicited quite a bit of interest from the carbon-based world. More than 6,000 humans were on hand to watch the matches, including approximately 1,400 IJCAI attendees and 4,500 of the general public. In addition, there was a live satellite broadcast, plus a live Internet broadcast with more than 400 accesses per minute, more than the server could handle.

Press interest was also high. The local Japanese press reported the progress through each round of play on daily news broadcasts. The BBC, the German press, CNN, ABC, and others covered the event. "The media didn't portray it as just silly entertainment," said a gratified Kitano. "It was a new, significant research project. It was really nice. They clearly understood what it's about."

The tournament also received a fair amount of industry support. Sponsors included such industry giants as Sony, Nihon Sun Microsystems, Citoh Techno-Science, Nihon Silicon Graphics-Cray, and Fujita.

## Rules of the game

RoboCup-97 had modified soccer rules. Players received no penalties for being offside, for example. The physical robot leagues allowed a maximum of only five robots per team. Middle-size robots had to be within 50 centimeters in diameter, but there was no height restriction, so many

## Related Web sites

| | |
|---|---|
| CMUnited | http://www.cs.cmu.edu/afs/cs/project/robosoccer/www/index.html |
| DreamTeam | http://www.isi.edu/isd/dreamteam/ |
| Humboldt | http://www.ki.informatik.hu-berlin.de/RoboCup97/index_e.html |
| ISIS | http://www.isi.edu/soar/tambe/socteam.html |
| RoboCup | http://www.csl.sony.co.jp/RoboCup/ |
| Sean Luke | http://www.cs.umd.edu/users/seanl/soccerbots/ |
| Soar | http://www.isi.edu/soar/soar-homepage.html |

were roughly the size of R2D2 in Star Wars. Weight ranged from 100 pounds to the winning DreamTeam robots, which were less than a foot high and weighed under two pounds each. Small robots could be as large as 15 centimeters in diameter.

The middle-size field was roughly the size of a tennis court, and the small field was a Ping-Pong table (see Figure 1). Both had walls. Middle-size robots used a size 4 soccer ball painted red, and the small bots an orange golf ball. For the Middle-Size League, the game consisted of a five-minute first half, a 10-minute break, and a five-minute second half. For the Small League, each half and the break were 10 minutes.

RoboCup provided the soccer simulator for the Simulation League. A full 11 players were on a simulation team. The simulator is available at the RoboCup Web site (see the sidebar, "Related Web sites"). It was designed with lots of obstacles to make it more realistic. For instance, it was difficult to figure out where a player was, and if a player ran fast, it got tired.

### Middle-Size League

Although the Middle-Size League comprised fewer than half a dozen competitors, the teams exhibited a diversity of technologies. The cowinners, for instance, were on opposite ends of the robot-control spectrum. The DreamTeam from the University of Southern California's Information Sciences Institute had autonomous robots, each with an on-board camera and computer. The robots from Osaka University had cameras on board, but they were controlled by a central computer that issued commands to the robots. Despite problems with radio interference, the Osaka team managed to tie for first.

The team from USC/ISI began to prepare for the event in February of this year. That gave them only six months. In spite of the short time frame, they built and programmed their team of small, lightweight, autonomous agents. Each robot consisted of a small 80x86-based motherboard, camera, and batteries on board a four-wheel-drive model car (see Figure 2). The camera was mounted up front, with a plastic bumper below for pushing the ball.

"Our autonomous agent architecture was key to our win," according to DreamTeam developer Rogelio Adobbati, a research assistant at USC/ISI. This included a vision module to process visual input, a drive controller to steer the bot, and a decision engine to help the bot decide what to do.

The decision engine had several components: a model manager to keep track of the field and nearby objects, a strategy planner to react to situations appropriately, and specialized components for goalkeeper, forward, and defender positions. The strategy was hard-coded, using C++. "We had very little time this year," explains Adobbati. "Next year we can work more on high-level strategies."

When the DreamTeam humans arrived in Nagoya, they weren't sure their approach would work. "Before the first match, we were very nervous. We were playing very big robots and were intimidated," says Adobbati. "We thought they'd push us out of the game since they were about 50 times heavier." As play unfolded, the large robots didn't move and were more like pillars than players. "Our robots went around them and scored. We won by four goals."

### Small League

CMUnited from Carnegie Mellon University won this league. The developers used machine-learning techniques and achieved some multiagent innovations in team formations, position switching, and passing. The team distinguished itself from many of the physical teams in both leagues by being able to pass and kick the ball and by showing some teamwork. For example, at least two of the goals the team scored involved one robot passing the ball to another, which shot the ball into the goal.

Peter Stone, a graduate research assistant in CMU's Computer Science Department, has been working on machine learning in multiagent systems for the past two years, so the RoboCup challenge fit right in with his interests. Preparations for the competition began in earnest in November 1996

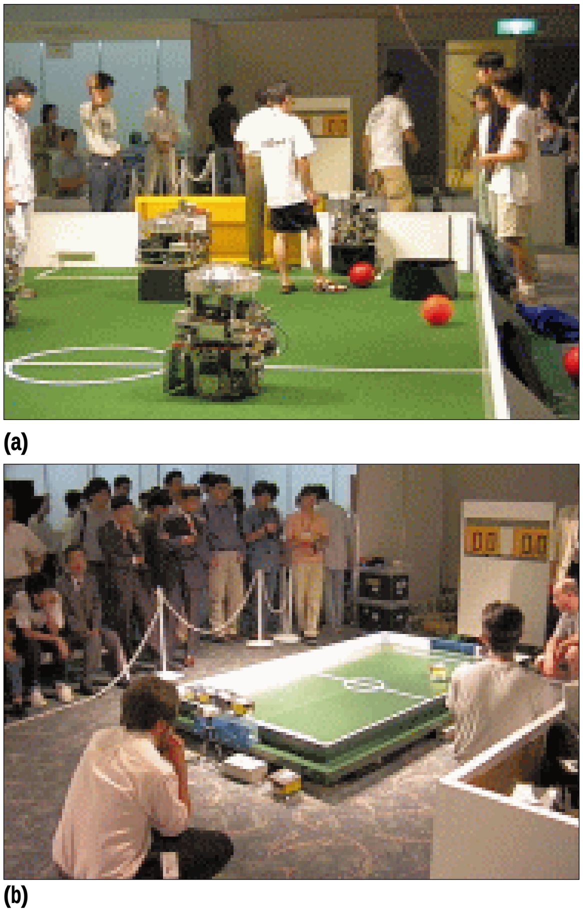
(a)

(b)

Figure 1. RoboCup playing fields: (a) Medium-Size League; (b) Small League (photo courtesy of LAC of RoboCup-97).

with a small team of three researchers (Manuela Veloso, Sorin Achim, and Kwun Han). “We chose the Small League to test a full version of our team in limited lab space using an overhead camera, and to focus on strategy and AI innovations instead of low-level control issues,” explains Stone.

The CMUnited team included five robots, plus one on the bench that was used when a first-string player’s batteries were low. Each bot was 12 cubic centimeters and weighed about two pounds, with approximately 60% of the weight being the batteries. With a single overhead camera, image processing was done off board and relayed via a wireless radio link. The vision system was able to determine both the position and orientation of 10 robots, and the position and velocity of the ball.

As with any team, much of the success was in the training and preparation. “We carefully tested situations that would come up repeatedly, such as a ball moving towards the goaltender. We threw the ball towards the goalie over and over again until we were satisfied with the performance,” says Stone. “We also worked from the bottom up, first making sure that the robots could hit a stationary ball into an empty goal, then working on moving balls, and finally reasoning about the positions of teammates and opponents. The reasoning was accomplished with evaluation functions that determined a free path from the ball to the goal.”

Why did CMUnited do so well? “Our vision system, unlike any other, was almost noiseless,” according to Stone. “That was a big advantage. We also had very good strategic teamwork: the players were able to switch positions and pass the ball to each other. We were the only team to meet the challenge of getting five robots to work as a team in a noisy, fast environment.”

Stone is already thinking about next year. “The biggest need for improvement is in robot speed. They actually moved fairly slowly (a walking pace). We could have moved them faster,” he continues, “but control and reliability would have suffered. We were focused on exhibiting reliable, intelligent behaviors.”

## Freed from physical bodies

CMUnited also competed in the Simulation League (they came in fourth, losing in the semifinals to the eventual champions). Stone found distinct advantages in the simulator. “With real robots, we spent most of our time just getting the sensors and actuators to work,” he explains. “Only at the end were we able to focus on strategy. However, the simulator abstracts away the perception and action programs. Therefore it is possible to focus on more high-level issues. For that reason, my research on layered learning is going on mostly in the simulator.”

> *Will there ever be a world-class match between humans and robots like the chess match between Kasparov and Deep Blue? Not in the foreseeable future.*

Stone’s layered-learning research is his approach to building multiagent systems using machine learning. Much as a child learns to play soccer, this approach layers increasingly complex learned behaviors. In the world of soccer, for instance, first a player (agent) learns low-level skills to control the ball. Then, building on this learned skill, it learns the higher-level skills of playing with teammates (multiple cooperating agents).

To train the virtual player, Stone’s group provided the agent with a large number of training examples and used neural networks. Once the player could control the ball, it was ready to be out on the field playing with the team. Passing the ball from one player to another is a fundamental component of team play. The developers successfully employed decision trees to help the player decide whether or not to pass.

## Simulation League

The CMUnited team effort illustrates the intent of the RoboCup organizers to push and test new technologies. For USC/ISI researchers already deeply involved in multiagent collaboration work for the military, RoboCup offered a different testbed for their innovations. The Soar system developed at USC/ISI has been used for some time to build synthetic agents for military exercises. (See “Intelligent Agents: The First Harvest of Softbots Looks Promising,” *IEEE Expert*, Aug. 1995, pp. 6–9.) Milind Tambe, a research scientist at USC/ISI, has been working on building a layer on top of Soar for multi-

## RoboCup-97 participants

Teams from leading research centers around the world competed in this year’s RoboCup. Entrants included teams from

- *Europe*: Chalmers Univ. of Technology, Sweden; Humboldt Univ., Germany; Univ. Carlos III de Madrid, Spain; Univ. of Girona, Spain; Univ. of Oulu, Finland; Univ. of Padua, Italy; Univ. of Paris.
- *North America*: Carnegie Mellon Univ., US (2 teams); Colorado School of Mines, US; Georgia Tech., Mobile Robot Laboratory, US; NASA Ames Research Center, US; Stanford Univ., US; Univ. of British Columbia, Canada; Univ. of Maryland, US; Univ. of Southern California, Information Sciences Inst., US (2).
- *The Pacific Rim*: Aoyama Gakuin Univ., Japan; Chuubu Univ., Japan (2); ElectroTechnical Laboratory, Japan (2); Justsystem, Japan; Kinki Univ., Japan (2); Nagoya Inst. of Technology, Japan; Nara Inst. of Science and Technology, Japan; Royal Melbourne Inst. of Technology, Australia (3); Tokyo Inst. of Technology, Japan (2); Toyo Univ., Japan (2); Toyohashi Gijutsu Univ., Japan; Univ. of Tokyo, Japan; Waseda Univ., Japan.

Most participants were in the Simulation League (32 teams); the physical leagues each had four or five teams competing.

agent collaboration and teamwork. He was pleased to find that approximately 30% of the code was general enough to transfer directly from a military application to the RoboCup soccer domain.

To prepare for RoboCup, Tambe worked intensively over four months with a group of graduate students, and a year before that at a leisurely pace by himself. Their soccer team, ISIS (*ISI S*ynthetic), was the top US team in the Simulation League, coming in third in the competition. Tambe is already looking toward next year. "This year we defended well. Our approach worked well," asserts Tambe. "As for next year, we need a structured offensive strategy. We need better agent modeling and plan-recognition abilities to understand what our opponents are up to."

Humboldt was the winning team for the Simulation League. Harking from Humboldt University in Germany, the developers used agent-oriented programming in the belief-desire-intention style to design the virtual team. As team leader Hans-Dieter Burkhard, a professor at the Institute of Informatics at Humboldt University in Berlin, describes their approach, agents (players) have beliefs about the world that are updated according to incoming sensor information. The agents then compute the coordinates of objects. The agents also have action sequences for intercepting the ball, reaching certain positions on the field, and manipulating and kicking the ball.

The Humboldt agents also have a desire/goal component. Based on their beliefs, the players decide which desire to adopt (such as intercept, kick, dribble, or run). Each agent is so well-articulated that it can evaluate various plans based on the world model and choose a course of action. Although the programmers hoped to use case-based reasoning to adapt playing to the opponents' behavior, it was not ready to use this year.

Given this approach, the Humboldt team has not used any machine-learning techniques. "We have learned," says Burkhard, "that learning should not start with simple behavior like single kicks, but it should be used to improve the behavior. First, make a careful analysis and implement a raw skill; then tune it by learning methods. For manipulating the ball, for instance, you need several steps that are very difficult to learn as a sequence from scratch." Once the basic skills were in place, the developers found that AI learning and planning methods worked well to tune behavior.

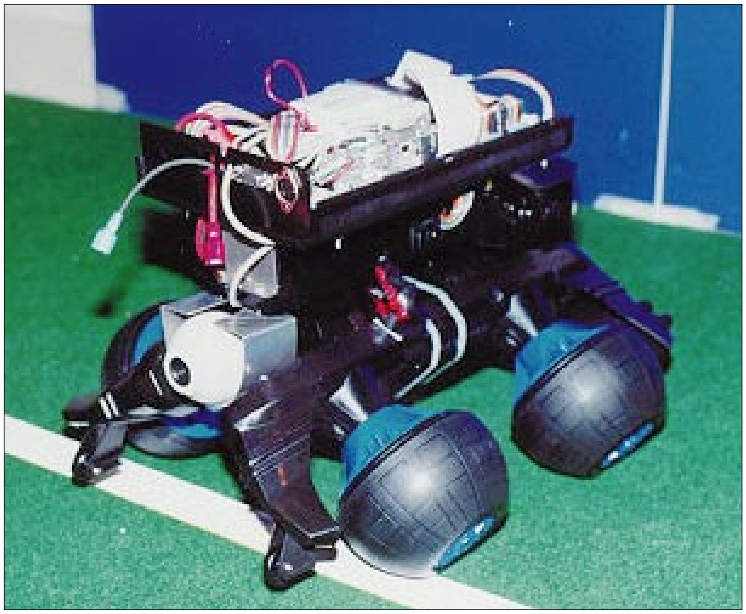
Figure 2. DreamTeam robot (photo courtesy of USC/ISI).

## Getting better

Although still quite seminal, robot soccer promises to be a challenging testbed for AI. "Already, we're seeing the beginnings of teamwork. In a domain such as soccer, where it is very difficult to program teamwork, we're seeing AI systems doing a better job than hard-coded, procedural code systems. This is encouraging to the AI community," notes Kitano.

"Now people are back in their offices understanding what to do for next year," he continues. "I expect very rapid progress in the next five to 10 years in this area." The challenges are set. Future visions are of legged robots, humanoid robots, and a telepresence competition. Robots will cooperate in space, on terra firma, and in cyberspace in a variety of endeavors.

Will there ever be a world-class match between humans and robots like the chess match between Kasparov and Deep Blue? Not in the foreseeable future. Who would want to collide with a metal robot on the soccer field? Or kick it in the shins and break a foot? Many bot innovations must occur before such a day: softer materials for robots, a different design for agility, better batteries for endurance, and so on.

Remembering what Kitano said about the AI community taking 40 years to rise to the grand challenge in chess, we'd be wise to adopt an evolutionary view of the field. For now, competitors are planning for RoboCup-98 in Paris.

**Sara Reese Hedberg** is a technology journalist who has been closely following AI for 14 years. She is a columnist for *IEEE Expert* and *IEEE Concurrency* and has written extensively about emerging technology for various publications. She can be reached at sara@hedberg.com.

## RoboCup-97 World Champions

And the winners were:

- Middle-Size League: DreamTeam (Univ. of Southern California's Information Sciences Inst.) tied with Trakies (Osaka Univ.)
- Small League: CMUnited (CMU)
- Simulator League: Humboldt (Humboldt Univ., Germany)
- RoboCup Scientific Challenge Award: Sean Luke (Univ. of Maryland), for demonstrating the utility of the evolutionary approach by coevolving soccer teams for the Simulator League (see the sidebar on Sean Luke).