

# The Social Life of Robots

*Researchers are trying to build robots capable of working together with minimal human supervision. But will they ever learn to get along?*


**R**OBBY THE ROBOT never had many robot friends. Nor did HAL, the Terminator, or most other robots of popular lore. In the public mind, robots have almost always been solitary creatures, carrying out their allotted tasks with single-minded purpose.

In the real world, robots have largely kept to themselves as well. To date, most robotics research has focused on building individual, autonomous machines. But the era of the lone robot may be drawing to a close. As the robotics field has matured, researchers have started to explore the possibilities of “social” machines capable of working together with minimal human supervision.

In theory, collaborative robots hold enormous potential. They could augment human workers in high-risk situations like firefighting or search and rescue, boost productivity in construction and manufacturing, and even help us explore other planets. But teaching robots to collaborate is proving to be a tricky business, raising thorny conceptual problems that go far beyond the largely mechanical challenges of designing single-purpose robots.

“Going from one robot to many increases analytical and computational complexity in a way that becomes unmanageable very quickly,” says Bert Tanner, an assistant professor at the University of Delaware who has been building robots designed to work together amid dangerous conditions such as fires or natural disasters.

Tanner’s team is working on a framework that would allow robots to gather data about their environment and alter their collective behavior in response to changing conditions. Rather than trying to build a universal



**The TRESTLE project at Carnegie Mellon University focuses on developing the architectural framework to coordinate robotic assembly teams.**

robot that tries to do everything, the Delaware team is creating a group of robots with complementary skills. For example, one robot might be good at opening doors, while another might be good at flying through doorways en route to fight a fire.

**To function as a team, robots must learn to negotiate decision-making processes in a distributed, multi-agent environment.**

Designing for this kind of close collaboration poses special challenges for robotics engineers, who have traditionally focused on the relatively more straightforward programming challenges of perception, cognition, and movement. “The multi-robot paradigm brings a new dimension to these problems,” notes professor Manuela Veloso of Carnegie Mellon University’s Department of Computer Science.

At the most basic level, collaborative robots need access to each others’ sensory data, so they cannot only “see” via their fellow robots, but in some cases reconcile perceptual differences as well. They then must learn to merge that shared spatial data into a unified whole, so that the robots can converge effectively in a physical space.

These are nontrivial problems, but researchers face a far more difficult computational challenge in developing systems to support the distributed cognition necessary to support effective collaboration. To function as a team, robots must learn to negotiate decision-making processes in a distributed, multi-agent environment.

In principle, it would be easy enough to solve that problem by putting one robot in charge of the others. But such a simplistic command-and-control architecture would greatly limit the potential of collaborative robots, especially in dangerous situations where any one of the robots could easily become damaged. Just as the Internet is designed to withstand outages by re-routing data dynamically around the network, an effective robot team needs to adapt on the fly if one or more of its members becomes inoperable.

## Distributed Robot Cognition

Some of the earliest work on distributed robot cognition started in the late 1990s, when Veloso’s CMU colleague

Reid Simmons began working with NASA on designing cooperative robots for planetary exploration.

“It became clear that sending up a single monolithic robot to do everything was less than ideal,” says Simmons. By building smaller, special-purpose robots, the team could reduce the overall payload size while mitigating the risks of technical failure. “From the NASA perspective it had to with mass and redundancy,” he says. “The argument was that if you wanted to do it all it would be bigger and bulkier, even though you had multiple robots.”

Simmons and his NASA colleagues began exploring how to create teams of special-purpose robots that could work

**The Distributed Robot Architectures framework tries to permit as much autonomy as possible for each individual robot, while providing a multi-layered structure that helps the extended group function as a team.**

independently and come together, as needed, to accomplish common goals. Rather than put a single robot in charge of the entire operation, the robots would take turns overseeing particular tasks and issuing instructions to the other robots as needed. “The authority could be passed around,” Simmons explains.

Making matters more complicated, the robots needed to work together in a remote environment with long communication delays. In addition to the 20 minutes it takes a signal to travel from Earth to Mars, orbiting conditions and thermal limitations meant the robots could communicate with Earth for at most three hours a day. Given these constraints, it would have been terribly inefficient for the robots to wait on continual instructions from Earth-based coordinators. The team instead began exploring a model of mixed autonomy, in which robots could function largely by themselves, but determine when they needed to stop and ask for help from the Earth-bound human staff.

The communications constraints prompted the engineers to consider another thorny issue: how to program the robots to determine when to ask for help. “Robots are really, really bad at detecting when they are outside the bounds of their accepted behaviors,” says Simmons. “You find that people tend to realize pretty quickly that they’re in a situation that’s completely unfamiliar, but robots are bad at determining that they need help.”

For example, a robot might try to insert an object into a container over and over again without realizing that

the object is turned backward. Without the ability to diagnose the problem, the robot might attempt to repeat the procedure ad infinitum. And while it might be possible to control for a particular known error condition in advance, it is far more difficult to abstract the problem to account for the vast array of potential unforeseen circumstances that could occur on a remote planetary surface.

“At some point the robot needs to understand this isn’t working,” Simmons says. “We looked a lot into how you can make predictions about how well the task is progressing in order to proactively repair multi-agent, multi-robot plans. For instance, instead of waiting until something fails and trying to do something [about it], maybe I can pull this resource off this task.”

Simmons’ work with NASA eventually gave rise to the Distributed Robot Architectures (DIRA) project, a framework that allows robots to react to changing and unexpected conditions by replanning and renegotiating their working relationships.

Simmons is now taking what he learned with NASA and applying it to more terrestrial endeavors. His team is currently trying to apply the DIRA framework to projects for General Motors on next-generation robotics for the manufacturing floor.

DIRA tries to permit as much autonomy as possible for each individual robot, while providing a multi-layered structure that helps the extended group function as a team. Simmons is now developing communication protocols that allow the robots to negoti-

## Awards

# 2012 ACM Fellows Nominations

The ACM Fellows program was established by the ACM Council in June 1993 to recognize outstanding ACM members for technical, professional, and leadership contributions that advance the arts, sciences, and practices of information processing; promote the free interchange of ideas and information in the field; develop and maintain the integrity and competence of individuals in the

field; and advance the objectives of ACM.

Each candidate is evaluated as a whole individual and is expected to bring honor to the ACM. A candidate’s accomplishments are expected to place him or her among the top 1% of ACM members. In general, two categories of accomplishments are considered: achievements related to information technology and

outstanding service to ACM or the larger computing community. A person selected as an ACM Fellow should be a role model and an inspiration to other members.

Nominations and endorsements must be submitted online no later than Sept. 5, 2012. For Fellows Guidelines, go to [http://awards.acm.org/html/fellow\\_nom\\_guide.cfm/](http://awards.acm.org/html/fellow_nom_guide.cfm/).

Nomination information organized by a principal

nominator should include excerpts from the candidate’s current curriculum vitae, listing selected publications, patents, technical achievements, honors, and other awards; a description of the work of the nominee, drawing attention to the contributions which merit designation as Fellow; and supporting endorsements from five ACM members. For the list of 2011’s ACM Fellows, see p. 23.

**“You find that people tend to realize pretty quickly that they’re in a situation that’s completely unfamiliar,” says Reid Simmons, “but robots are bad at determining that they need help.”**

ate their tasks and monitor progress, distribute sensory data, and cope with reliability issues when one member of the group malfunctions. By distributing the decision-making responsibilities, these machines will likely prove more robust and able to adapt to unforeseen adverse conditions.

### Learning From Mistakes

In a similar vein, Manuela Veloso is exploring how to help robots reflect on their own experiences and learn from their successes and mistakes to improve future decision making.

Veloso’s team has discovered a great deal about robot learning by way of robot soccer. Her CMU team has built several robot soccer teams that have to play with and against each other in different configurations. To collaborate effectively, the robots must negotiate constraints and build coalitions with each other to solve problems. Veloso has found game theory particularly instructive in this regard, exploring how robots can learn to negotiate winning outcomes in cooperative, adversarial, or semi-cooperative situations.

Robot soccer also raises the question of how groups of robots might begin to negotiate relationships with fellow robots when they meet each other for the first time. As cooperative robots eventually find their way into the real world, there will inevitably be a growing need for robots to find ways to negotiate their initial encounters.

Veloso thinks that such open col-

laboration will likely depend on the development of standards and protocols for robot interaction. “We must develop a model for robots to declare their actuation capabilities,” she explains. “It’s crucial.”

Alas, no such model currently exists. Veloso has been pressing the U.S. National Science Foundation to fund research on open protocol standards, so far to no avail. However, she remains hopeful that the need for such standards will eventually prove self-evident.

Tanner also sees a need for more unifying frameworks to guide future collaborative robot development. “We don’t have a theory that can capture and express clearly how actions of one agent can enhance or inhibit the capabilities of others, and then use that theory to plan and coordinate cooperative behavior,” he notes.

For now, it seems that collaborative robots will continue to muddle along in their developmental stages, trying to do their jobs, improve their relationships, and bridge the communication gaps that continue to keep them apart. If robots ever manage to solve those problems, perhaps they will have something to teach the rest of us. **□**

### Further Reading

Cao, Y.U., Fukunaga, A.S., and Kahng, A.

Cooperative mobile robotics: antecedents and directions, *Autonomous Robots* 4, 1, March 1997.

McCarty, K. and Manic, M.

Adaptive behavioral control of collaborative robots in hazardous environments, *Proceedings of the 2<sup>nd</sup> Conference on Human System Interactions*, Piscataway, NJ, May 21–23, 2009.

Nakanishi, R., Bruce, J., Murakami, K., Naruse, T., and Veloso, M.

Cooperative 3-robot passing and shooting in the RoboCup Small Size League. *RoboCup 2006: Robot Soccer World Cup X*, Lakemeyer, G., Sklar, E., Sorrenti, D.G., and Takahashi, T. (Eds.), Springer-Verlag, Berlin and Heidelberg, Germany, 2007.

Roth, M., Simmons, S., and Veloso, M.

Exploiting factored representations for decentralized execution in multiagent teams, *Proceedings of the 6<sup>th</sup> International Joint Conference on Autonomous Agents and Multiagent Systems*, Honolulu, Hawaii, May 14–18, 2007.

Alex Wright is a writer and information architect based in Brooklyn, NY.

© 2012 ACM 0001-0782/12/02 \$10.00

### Society

## Modeling Ethnic Conflict

New England Complex Systems Institute (NECSI) researchers have developed a computational model that shows social stability among various ethnic groups in the same country does not depend on integrated coexistence, but on well-defined topographical and political boundaries separating the different groups.

The NECSI researchers discovered that civil violence is not likely when diverse communities are so integrated that one group does not control others, or so segregated that geographic and political boundaries match demographic borders.

Led by NECSI president Yaneer Bar-Yam, the researchers tested their computational model on data from Switzerland, which is characterized by a history of social stability and economic prosperity, but contains diverse languages and religions, including French and German and Italian, Catholic and Protestant.

The researchers found that stability is maintained through the existence of natural boundaries, such as lakes and mountains, or political boundaries, such as cantons or quasi-autonomous member states, that often separate ethnic and religious groups, rather than through integrated social coexistence.

The model found that one Swiss area with an increased propensity for violence was the country’s northwest region, where the Jura mountains form a weak boundary between French and German-speaking communities. Violence erupted between the aforementioned two communities in this region during the 1970s.

The study’s most significant finding, says Bar-Yam, is that “it is possible to mitigate ethnic conflict using within-country borders that provide local autonomy.”

The NECSI model might help resolve conflicts in other ethnically diverse countries and regions of the world, according to the researchers.

—Bob Violino