

Enabling human-robot rescue teams

February 17 2016, by Larry Hardesty



Communication system could make it easier to design systems that enable humans and robots to work together in emergency-response teams. Credit: Jose-Luis Olivares/MIT

Autonomous robots performing a joint task send each other continual updates: "I've passed through a door and am turning 90 degrees right."
"After advancing 2 feet I've encountered a wall. I'm turning 90 degrees

right." "After advancing 4 feet I've encountered a wall." And so on.

Computers, of course, have no trouble filing this information away until they need it. But such a barrage of data would drive a human being crazy.

At the annual meeting of the Association for the Advancement of Artificial Intelligence last weekend, researchers at MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) presented a new way of modeling robot collaboration that reduces the need for communication by 60 percent. They believe that their model could make it easier to design systems that enable humans and robots to work together—in, for example, emergency-response teams.

"We haven't implemented it yet in human-robot teams," says Julie Shah, an associate professor of aeronautics and astronautics and one of the paper's two authors. "But it's very exciting, because you can imagine: You've just reduced the number of communications by 60 percent, and presumably those other communications weren't really necessary toward the person achieving their part of the task in that team."

The work could also have implications for multirobot collaborations that don't involve humans. Communication consumes some power, which is always a consideration in battery-powered devices, but in some circumstances, the cost of processing new information could be a much more severe resource drain.

In a multiagent system—the computer science term for any collaboration among autonomous agents, electronic or otherwise—each agent must maintain a model of the current state of the world, as well as a model of what each of the other agents takes to be the state of the world. These days, agents are also expected to factor in the probabilities that their models are accurate. On the basis of those probabilities, they have to

decide whether or not to modify their behaviors.

Communication costs

In some scenarios, a robot's decision to broadcast a new item of information could force its fellows to update their models and churn through all those probabilities again. If the information is inessential, broadcasting it could introduce serious delays, to no purpose. And the MIT researchers' work suggests that 60 percent of communications in multiagent systems may be inessential.

The state-of-the-art method for modeling multiagent systems is called a decentralized partially observable Markov decision process, or Dec-POMDP. A Dec-POMDP factors in several types of uncertainty; not only does it consider whether an agent's view of the world is correct and whether its estimate of its fellows' worldviews is correct, it also considers whether any action it takes will be successful. The robot may plan, for instance, to move forward 20 feet but find that crosswinds blow it off course.

Dec-POMDPs generally assume some prior knowledge about the environment in which the agents will be operating. Because Shah and Vaibhav Unhelkar, a graduate student in aeronautics and astronautics and first author on the new paper, were designing a system with emergency-response applications in mind, they couldn't make that assumption. Emergency-response teams will usually be entering unfamiliar environments, and the very nature of the emergency could render the best prior information obsolete.

Adding the requirement of mapping the environment on the fly, however, makes the problem of computing a multiagent plan prohibitively time consuming. So Shah and Unhelkar's system ignores uncertainty about actions' effectiveness and assumes that whatever an

agent attempts to do, it will do.

Balancing act

When an agent acquires a new item of information—that, for instance, a given passage through a building is blocked—it has three choices: it can ignore the information; it can use it but not broadcast it; or it can use it and broadcast it.

Each of these choices has benefits but imposes costs. In Shah and Unhelkar's model, communication is a cost. But if an agent incorporates new information into its own model of the world and doesn't broadcast it, it also incurs a cost, as its worldview becomes more difficult for its fellows to estimate correctly. For every new item of information an agent acquires, Shah and Unhelkar's system performs that cost-benefit analysis, based on the agent's model of the world, its expectations of its fellows' actions, and the likelihood of accomplishing the joint goal more efficiently.

The researchers tested their system on more than 300 computer simulations of rescue tasks in unfamiliar environments. A version of their system that permitted extensive communication completed the tasks at a rate between 2 and 10 percent higher than the version that reduced communication by 60 percent.

In the experiments, however, all the agents were electronic. "What I'd be willing to bet, although we have to wait until we do the human-subject experiments, is that the human-robot team will fail miserably if the system is just telling the person all sorts of spurious information all the time," Shah says. "For human-robot teams, I think that this algorithm is going to make the difference between a team that can function effectively versus a team that just plain can't."

In a separate research project, members of Shah's group have asked teams of human subjects to execute similar virtual rescue missions that computer systems did in the experiments reported in the new paper. Using machine-learning algorithms, the researchers have mined the results for statistics on human communication patterns, which can be incorporated into the new model to more explicitly accommodate human-robot teams.

"It is well-understood that in human teams, when one team member gains new information, broadcasting this new information to all team members is generally not a good solution, especially when the cost of communication is high," says Tim Miller, an assistant professor of computing and information systems at the University of Melbourne in Australia. "This work has applications outside of multiagent systems, reaching into the critical area of human-agent collaboration, where [communication](#) can be costly, but more importantly, human team members are quickly overloaded if presented with too much [information](#)."

More information: ConTaCT : Deciding to Communicate During Time-Critical Collaborative Tasks in Unknown, Deterministic Domains: interactive.mit.edu/contact-de...eterministic-domains

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Enabling human-robot rescue teams (2016, February 17) retrieved 5 May 2024 from <https://phys.org/news/2016-02-enabling-human-robot-teams.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.