

Robots learning to work with humans

April 2 2014, by Eric Brown



Julie Shah

With the advent of "inherently safe" robots, industrial designers are changing their ideas about the factory of the future. Robots such as ABB's Frida and the Baxter robot from MIT spinoff Rethink Robotics are working "elbow to elbow with people," says Julie Shah, an assistant

professor in MIT's Department of Aeronautics and Astronautics and director of the MIT Interactive Robotics Group. "They're designed so that if they hit a person they don't significantly harm them."

Working in the Interactive Robotics Group at MIT's Computer Science and Artificial Intelligence Laboratory, Shah is taking the next step: teaching these inherently safe robots how to work together in teams with people, and vice versa. "We're focused on [robot](#) learning, planning, and decision making, and how they interact with humans in high intensity and safety critical environments," Shah says. "We're looking to develop fast, smart tasking algorithms so robots can work interdependently with people."

Despite the rapid spread of robotics in manufacturing, many final assembly tasks, especially in building airplanes, automobiles, and electronics, still depend largely on human labor. With the availability of more intelligent, adaptable, and inherently safe robots, there are new opportunities for automation.

"In most factories, robots and the people are kept very separate," Shah says. "But factories of the near future are going to look very different. We're beginning to see safety standards and technology that lets us put some of these large, dangerous industrial robots onto mobile bases and rails so that they can safely work with people."

With most of the safety issues solved, the main focus for Shah is in training robots and people to work together more productively. "How do we program the robots to work in teams in a very dynamic environment where you have people coming and going?" Shah says.

The current state of the art for training robots depends on demonstration and interactive rewards. "If the robot does something good, we tell them it's good, and if not, we say it's not good, and the robot learns through

that reinforcement process," Shah says.

Yet when Shah considered that "these reward methods are documented as among the most inefficient ways to help humans work together," she imagined they might be even less effective in human/robot teams. Indeed, her research showed that robots are often unclear what the reward is referring to. "Are we rewarding the robot based on what it just did, or what it did a few steps ago, or what we think the robot is going to do in the future?" Shah says. "It's hard to train someone how to apply these rewards."

Cross-training to the rescue

To improve robot training methods, Shah studied how flight crews, medical teams, military tactical units, and other human teams train to work together effectively. Again and again, she found that one of the most effective approaches involved cross-training: people taking turns doing each other's job. "There's evidence that by doing someone else's job, you take that information back with you when you do your own job, since you can better anticipate what your partners need," Shah says. "The outcome is more effective, especially when responding to errors and disturbances."

Shah and her research team modified reinforced learning techniques and algorithms so that instead of the robot receiving input as a positive and negative reward, it receives input by switching roles with the person. They performed a simple experiment in a virtual environment in which a person performs the robot's role and vice versa.

The outcomes were "surprising and exciting," Shah says. "We saw improvements after cross-training in objective measures of team performance and statistically significant increases in concurrent motion between human and robot. We also saw significant reductions in idle

time, as well as subjective improvements. People agreed more strongly that they trusted the robot and that the robot worked according to their preference."

In the control group, which instead used active reinforcement learning, "you can see the person hesitate and wonder what the robot will do next," Shah says. "When the robot moves, the person pulls their hand out of the space. But with the cross-training, the person is more confident about what the robot will do, and is more likely to leave their hand in the shared space."

By switching roles, the human is teaching the robot more explicitly what it thinks the robot should do. This is more straightforward and intuitive for both human and robot, Shah says. "Through switching roles the robot learns the person's preferences better, and develops an increased certainty of what the person will do," she says. "And by watching the robot do what it thinks the person should be doing, the person benefits as well."

Shah is now looking into alternative training approaches. "Cross-training is the gold standard for training, but it's inherently limiting because if the robot could be doing the person's job maybe it would already be doing it," she says. For example, she notes that it would not make sense to train a robotic surgery assistant by trying to teach it to do surgery when only a select group of humans currently possess that skill.

Humans: the ultimate uncontrollable entity

In addition to training human/robot teams, Shah is also looking into optimizing task planning and implementation in hybrid teams. Choreographing human and robot movements in the same workspace is a challenge.

"When we have multiple robots working together in a factory cell, their motions and timing are pre-planned, and they often use a centralized controller, so it's really like one big robot," Shah says. "When you have a person in that space, pre-planned motions are difficult because you don't know exactly where the person will be and when. Humans are the ultimate uncontrollable entity. Robot decision-making algorithms need to be very fast in order to respond."

One challenge is that safety measures inherently slow productivity. For example, when a person nears a robot, the robot is programmed to slow down or stop. Yet, if a person stops in front of a robot while talking to somebody else, they impede the robot from working. If many people are working in the space, the robot is always stopping, reducing any efficiency benefit.

To address this issue, the researchers have built a statistical model of what a person is likely to do. "We're looking at how we can re-sequence the motion plans so the robot maneuvers further away from the person," Shah says. "It may be a longer motion path, but ultimately it's more efficient than being stopped."

In a project with BMW, Shah and her team are attempting to install mobile robotic assistants on final car assembly lines. This is still primarily a manual process, but there are places for robot partners to help out.

"People waste time walking back and forth to pick up the next piece to install," Shah says. "A mobile robotic assistant can fetch the right tools and parts at the right time."

The challenge is that the humans and robots are working in a very confined space. "The robot needs to maneuver around many people, and may need to straddle a moving conveyor belt," Shah says. "It has to move

on and off the line seamlessly."

To help robots negotiate in this dynamic environment, the researchers are teaching them how to interpret anticipatory signals in human motion. "Biomedical studies show people can anticipate whether a person will turn left or right about a step or two before they do," Shah says. "If we can teach the robot to anticipate which way the person will move, and modify its motion paths and speed accordingly, we could improve efficiency while maintaining safety."

There are several practical hurdles to human/robot team deployments that are beyond the scope of Shah's research. For example, in order for the robot to track human movements, the worker must wear an expensive motion capture body suit. Shah expects that this problem will soon be solved by cheaper, less intrusive, and more accurate sensing equipment.

"Our goal is to translate anticipatory signals with a few cameras rather than relying on body sensors," she says. "There are researchers working on vision technology that can sense within a millimeter where a person is," she says. "Those advancements are coming along in parallel with our research. Sensing and computation are large enablers for us."

Disaster response and beyond

Another new area of research is in cross-training robots and humans in disaster response situations. Shah is working to extract domain knowledge from the Web-based tools that are increasingly used in disaster response planning. Algorithms based on the knowledge could "help unmanned aerial or autonomous ground vehicles respond more intelligently in an uncertain environment," she says.

As robots spread out into new areas such as medical care and home

assistance, some of these insights into human/robot cross-training should still prove effective. "Potentially some of this research could translate to a robot that helps cook our dinner," Shah says.

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: Robots learning to work with humans (2014, April 2) retrieved 27 April 2024 from <https://phys.org/news/2014-04-robots-humans.html>

<p>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.</p>
--