

Can robots make good teammates?

Carolyn Beans, *Science Writer*

In every hospital labor and delivery department, a resource nurse decides which patients go to which room, which nurses care for which patients, and much more. “It’s a lot to keep in your head,” says Kristen Jerrier, a resource nurse in the labor and delivery department at Beth Israel Deaconess Medical Center in Boston.

“Her job is actually more computationally complex than that of an air traffic controller,” adds Julie Shah, who directs the Interactive Robotics Group in the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology.

Shah is developing a decision-support system embodied in a small desktop robot that can help Jerrier and other nurses make easy decisions, freeing up their attention for more difficult or complex decisions. Jerrier was skeptical that a robot could offer real value. But then it made suggestions similar to ones she herself would make. That’s when she thought “this could actually be helpful to me.”


A growing group of human–robot interaction researchers is developing robot teammates for humans working in a range of jobs, from manufacturing cars to exploring outer space. Rather than being remotely controlled, these robotic teammates, whether machines with robotic arms or human-like forms, are designed to collaborate with humans directly.

But being a good teammate is hard. “We need to know what our partner is thinking, anticipate what they will do next, and make fast adjustments when things don’t go according to plan,” explains Shah. Teammates also must learn from one another, communicate clearly, and know when it’s appropriate to interrupt. And robots still lack the full situational awareness that humans possess.

The Robot Apprentice

Shah and her team developed a machine learning algorithm that could learn the skills of a resource nurse from an experienced nurse in action, just as human nurses in training do (1). Starting in late 2015, seven resource nurses at Beth Israel Deaconess Medical Center each worked individually with a computer simulation of activity in the labor and delivery department. The simulation asked them which rooms and nurses to assign to patients as the patients’ statuses changed and as new patients arrived unscheduled or for planned cesareans and inductions. The program recorded more than 3,000 decisions made by the nurses.

Seventeen different physicians and nurses, including Jerrier, then used the same simulation. This time, the machine learning model suggested decisions to them, either through the computer or through a small humanoid robot designed by SoftBank Robotics that stood on the desktop in the nurses’ station. The machine based its advice on an algorithm that assumed that each of the decisions that the nurses made in the learning phase of the study was equivalent or better than all other choices. Drawing on these examples, it could rank the suitability of all possible decisions as each new patient arrived in the simulation. It recommended the top-ranked decision, and the experts accepted that recommendation 90% of



Julie Shah of the Massachusetts Institute of Technology is developing robot teammates for a range of workplace settings, from a robotic arm that could deliver parts on a factory floor to this decision-support robot that could help resource nurses decide which rooms and nurses to assign to patients. Image courtesy of Matthew Gombolay (photographer).

the time, regardless of whether the recommendation was delivered by the robot or the computer.

Shah then conducted a 3-hour pilot study in the labor and delivery department where the robot successfully fielded seven requests for suggestions from nurses at the nurses' station but failed to respond to three because of background noise. "This machine learned very quickly," says Jerrier. "I expected it to take a lot longer to do what I do." The work will soon be published in the *Journal of Artificial Intelligence Research*.

But there are clearly limitations. In the case of the hospital system, some nurses find it bothersome to repeatedly hear suggestions that don't align with their own, even if technically correct, says Jerrier. "Sometimes there are two ways to get from A to B." Shah is exploring methods to tailor recommendations to specific individuals or hospitals.

Robot Pas de Deux

One way to foster productive interactions among humans and their robot teammates is to use the same sort of teambuilding strategies that have proven successful in human-only teams. Cross-training, for example, which involves swapping roles in a shared task, is commonly used in aviation and the military. "Taking someone else's role, physically doing their job, gives you a better idea of what the partner needs to be successful," Shah explains. Her group was the first to try this technique in human-robot teams. In their 2015 study, human-robot pairs repeatedly swapped roles as one teammate placed screws and the other drilled them in (2). The cross-training resulted in a 71% increase in the amount of time that the human and robot were both in motion and a 41% decrease in human idle time compared with a more traditional training practice where the human and robot each had fixed roles and the human offered performance feedback to the robot. The findings suggest that cross-training allows humans and robots to learn to work more fluently together. The study, however, did not distinguish whether the enhanced teamwork stemmed more from the robot improving its ability to work with the human or vice versa.

Good teammates also anticipate a partner's next move. Today, for safety reasons, robots on factory floors stop whenever a human is near. In principle, robots could traverse factory floors without constantly shutting down if they could predict where humans were moving. "If the robot knows a person is going to turn right, the robot can cut a straight line," Shah explains. In 2014, Shah's team used motion capture technology to track the head orientation and body velocity of individuals as they walked a straight line and then turned (3). Using this data, her group developed an algorithm that predicts, two steps in advance, when and in which direction a human is going to turn.

Last year, building on this predictive ability, the group tested a robotic system that assists in manufacturing (4). The robot moved along a straight line, picking up parts with a robotic arm and delivering the parts to humans building automotive engines. A Kinect motion sensor



NASA expects that Astrobees, a free-flying robot, will begin working with humans aboard the International Space Station by the end of this year. This illustration shows what two Astrobees might look like as they cooperate to move cargo. Image courtesy of NASA.

and a set of algorithms allowed the robot to predict the trajectory of humans based on a combination of factors—for example, the person's velocity and where his or her current movement fell within a typical sequence of human activities on the factory floor. The robot then planned its own trajectory—when to move forward, backward, or stop—around the human's movement. The robot adjusted its plans every tenth of a second based on changes in the person's activity. In both a simulation and a demonstration in a BMW Group test facility in Munich, the system enabled tasks to be completed with fewer safety-related stoppages, an important safety metric. But before actually implementing any robot, a manufacturing facility would perform its own risk and safety analyses.

The BMW Group does not expect that such a robot's reduction in safety-related stoppage time would have a major impact on productivity in the coming years, says Andreas Hemmerle, a communications officer with the BMW Group. But the company is nevertheless committed to such research because robots that can plan around human movement do offer potential safety improvements. Still, it's too early to say whether technology like Shah's will be implemented on a large scale, says Hemmerle.

Jeremy Marvel, a researcher at the National Institute of Standards and Technology in Gaithersburg, MD, who specializes in human-robot collaborations and safety, calls Shah's work at the BMW Group a "proof-of-concept" study that "opens up possibilities." Although he says that major improvements in motion-sensing technologies will be necessary before robot teammates can be used on a large scale in

manufacturing, he anticipates that Shah's work will help inspire the industry to meet this need.

Shah agrees that there are still technical challenges to address, especially in advancing the sensing technology. But assuming those issues are overcome, she anticipates that robots such as the one she tested at the BMW Group will be used in production environments in the next 2 to 3 years.

The tight confines of the International Space Station, however, might see such technologies sooner. Terry Fong at NASA's Ames Research Center is developing a system of free-flying robots named Astrobees that is scheduled to operate in the International Space Station starting in November (5). The robots will conduct surveys for environmental variables inside the space station, such as sound levels and air quality, and search for supplies that drift in the microgravity environment. Fong is collaborating with Shah and her graduate student Przemyslaw Lasota to endow Astrobees with a capacity similar to the factory robots so that it anticipates human motion and adjusts its plans mid-task. The robot is able to avoid running into walls or equipment, says Fong. "But it's much harder to avoid mobile obstacles like people."

In some respects, human-robot teammates are already a reality in the form of semiautonomous cars made by Tesla, Cadillac, and others. According to computer scientist Christopher Crick of Oklahoma State University, under a broad definition of teamwork, these cars do function as teammates. Some driving systems, for example, monitor drivers' alertness before accepting control. Crick, though, suggests that the automobile industry is more intent on minimizing the human component as they work toward fully autonomous vehicles.

Reluctant Teammates

To perfect teammate set-ups, roboticists will have to replicate another quality of a human teammate: social awareness. Crick is designing robots that know when it's appropriate to ask for help. "Robots need to know that this human that I have on my team is a resource for me and can help me out of a jam, but I am not the only thing on that human's mind," he says.

Crick developed a robot that can sense when a human teammate is too frazzled to provide quality

advice (6). He and his team first asked humans to give a robot directions to navigate a maze. The humans usually gave good directions. The task then became more challenging with time pressure or an additional robot to direct. The people made mistakes, and the robots, because they had their own internal map of the maze, knew it. The robots looked for associations between bad directions and human behaviors—for example, how often participants became overwhelmed and gave one command, only to quickly override it with another. The humans then directed robots through a maze toward prizes that the robots couldn't see. If the robots detected those same signs of an overwhelmed human, they could then decide to reduce their reliance on the humans, says Crick. They couldn't go after the prizes without listening to the humans, but they could consult their internal map to avoid running into walls.

Despite these gains, most robots today cannot work collaboratively, in part, because they lack situational awareness. "We are still far away from a robot that understands the full variety and nuance of how human behavior changes in different contexts," Shah says. One of the challenges, she says, is that we don't always know each of the myriad factors that affect someone's behavior and, hence, can't provide that information to the robot. Shah and her team are beginning to develop machine learning techniques that enable the robot to hypothesize factors that could be affecting a person's behavior.

There are also barriers to humans accepting robots as teammates. "Especially in Western cultures, we automatically come to the table with fear," says Leila Takayama, a cognitive and social scientist at the University of California, Santa Cruz, who studies human-robot interaction. "People need to feel like they can anticipate what the robot is going to do if they are going to feel safe," says Takayama.

Then there is the job-replacement concern. "In places where people feel that they don't have a lot of job security, it can make the introduction of new technology scary," says Takayama. Jerrier, however, isn't worried. "There are definitely aspects of nursing," she says, "that a robot and a machine could never understand."

- 1 Gombolay M, et al. (2016) *Robotic Assistance in Coordination of Patient Care*, Proceedings of Robotics: Science and Systems (Robotics Science and Systems, Ann Arbor, MI). Available at www.roboticsproceedings.org/rss12/p26.html. Accessed August 22, 2018.
- 2 Nikolaidis S, Lasota P, Ramakrishnan R, Shah J (2015) Improved human-robot team performance through cross-training, an approach inspired by human team training practices. *Int J Rob Res* 34:1711–1730.
- 3 Unhelkar VV, Pérez-D'Arpino C, Stirling L, Shah JA (2015) *Human-Robot Co-navigation Using Anticipatory Indicators of Human Walking Motion*, IEEE International Conference on Robotics and Automation (IEEE, New York). Available at [https://ieeexplore.ieee.org/document/71140067/?tp=&number=71140067&refinements%3D4224266977%26filter%3DAND\(p_IS_Number:7138973\)](https://ieeexplore.ieee.org/document/71140067/?tp=&number=71140067&refinements%3D4224266977%26filter%3DAND(p_IS_Number:7138973)). Accessed August 22, 2018.
- 4 Unhelkar VV, et al. (2018) Human-aware robotic assistant for collaborative assembly: Integrating human motion prediction with planning in time. *IEEE Robot Autom Lett* 3:2394–2401.
- 5 Bualat MG, et al. (2018) Astrobees: A New Tool for ISS Operations, 2018 SpaceOps Conference (Aerospace Research Council, Reston, VA). Available at <https://arc.aiaa.org/doi/abs/10.2514/6.2018-2517>. Accessed August 22, 2018.
- 6 al Mahi SM, Atkins M, Crick C (2017) *Learning to Assess the Cognitive Capacity of Human Partners*. Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction (Association for Computing Machinery, New York). Available at <https://dl.acm.org/citation.cfm?id=3038430>. Accessed August 22, 2018.