

Teaching Robots Lateral Thinking

Larry Hardesty, MIT News Office

New algorithms could help household robots work around their physical shortcomings.



Many commercial robotic arms perform what roboticists call “pick and place” tasks: The arm picks up an object in one location and places it in another. Usually, the objects — say, automobile components along an assembly line — are positioned so that the arm can easily grasp them; the appendage that does the grasping may even be tailored to the objects’ shape.

General-purpose household robots, however, would have to be able to manipulate objects of any shape, left in any location. And today, commercially available robots don’t have anything like the dexterity of the human hand.

At this year’s IEEE International Conference on Robotics and Automation, students in the Learning and Intelligent Systems Group at MIT’s Computer Science and Artificial Intelligence Laboratory will present a pair of papers showing how household robots could use a little lateral thinking to compensate for their physical shortcomings.

One of the papers concentrates on picking, the other on placing. Jennifer Barry, a PhD student in the group, describes an algorithm that enables a robot to push an object across a table so that part of it hangs off the edge, where it can be grasped. Annie Holladay, an MIT senior majoring in electrical engineering and computer science, shows how a two-armed robot can use one of its graspers to steady an object set in place by the other.

Colliding Approaches

Most experimental general-purpose robots use a motion-planning algorithm called the rapidly exploring random tree, which maps out a limited number of collision-free trajectories through the robot’s environment — rather like a subway map overlaid on the map of a city. A sophisticated-enough robot might have arms with seven different joints; if the robot is also mounted on a mobile base — as was the Willow

Teaching Robots Lateral Thinking

Published on Wireless Design & Development (<http://www.wirelessdesignmag.com>)

Garage PR2 that the MIT researchers used — then checking for collisions could mean searching a 10-dimensional space.

Add in a three-dimensional object with three different axes of orientation, which the robot has to push across a table, and the size of the search space swells to 16 dimensions, which is too large to search efficiently. Barry's first step was to find a concise way to represent the physical properties of the object to be pushed — how it would respond to different forces applied from different directions. Armed with that description, she could characterize a much smaller space of motions that would propel the object in useful directions. "This allows us to focus the search on interesting parts of the space rather than simply flailing around in 16 dimensions," she says. Finally, after her modification of the motion-planning algorithm, she had to "make sure that the theoretical guarantees of the planner still hold," she says.

By contrast, Holladay's algorithm in some sense inverts the ordinary motion-planning task. Rather than identifying paths that avoid collisions and adhering to them, it identifies paths that introduce collisions and seals them off. If the robot is using one hand to set down an object that's prone to tipping over, for instance, "I might look for a place for the other hand that will block bad paths and kind of funnel the object into the path that I want," Holladay says.

Like Barry, Holladay had to find a simple method of representing the physical properties of the object the robot is manipulating. In addition to the placement of tall, tippy objects, her algorithm can also handle cases in which the robot is setting an object on a table, but the object sticks to the rubber sheath of the robot's gripper. With Holladay's algorithm, the robot can use its free gripper to prevent the object from sliding as it withdraws the other gripper.

Independent Learning

Both Barry and Holladay allow modification of their algorithms, through application programming interfaces that would allow other researchers to plug in parameters describing the physical behavior of new types of objects. But the ultimate goal is for the robot itself to infer the relevant properties of objects by lifting, shoving, or otherwise manipulating them.

Nor are the researchers concerned that hardware improvements will render their algorithmic research obsolete. "The thought is that we're unlikely to get hands that are as flexible and dexterous as human hands, and even if we did, it would be hard to figure out the AI and planning for those," Barry says. "So we'll always have to think about interesting ways to grasp things."

"You see a lot of demos where a robot might do something like slide plates, but it's usually hard-coded for the demo: The robot knows that at this point, it needs to do this action for this particular thing," says Kaijen Hsiao, a research scientist and manager at Willow Garage, the company that manufactures the PR2. Barry and Holladay's research, by contrast, is "a framework for incorporating behaviors like that as a more general motion-planning problem," she says. "Which is a very difficult thing, because it's very high-dimensional. I think it's really important

Teaching Robots Lateral Thinking

Published on Wireless Design & Development (<http://www.wirelessdesignmag.com>)

research, and it's very novel."

For more information visit www.mit.edu [1].

Source URL (retrieved on *01/26/2015 - 11:36am*):

http://www.wirelessdesignmag.com/articles/2013/02/teaching-robots-lateral-thinking?qt-digital_editions=0

Links:

[1] <http://www.mit.edu>