

Researchers help robots think and plan in the abstract

8 February 2018, by Kevin Stacey



New research shows that robots can learn abstract representations of the world that are useful in planning for multi-step tasks, something that's monumentally difficult for robots to do. Here, a robot learns useful abstractions about the world by executing a set of motor skills. Credit: Intelligent Robot Lab / Brown University

Researchers from Brown University and MIT have developed a method for helping robots plan for multi-step tasks by constructing abstract representations of the world around them. Their study, published in the *Journal of Artificial Intelligence Research*, is a step toward building robots that can think and act more like people.

Planning is a monumentally difficult thing for robots, largely because of how they perceive and interact with the world. A [robot](#)'s perception of the world consists of nothing more than the vast array of pixels collected by its cameras, and its ability to act is limited to setting the positions of the individual motors that control its joints and grippers. It lacks an innate understanding of how those pixels relate to what we might consider meaningful concepts in the world.

"That low-level interface with the world makes it

really hard to do decide what to do," said George Konidaris, an assistant professor of computer science at Brown and the lead author of the new study. "Imagine how hard it would be to plan something as simple as a trip to the grocery store if you had to think about each and every muscle you'd flex to get there, and imagine in advance and in detail the terabytes of visual data that would pass through your retinas along the way. You'd immediately get bogged down in the detail. People, of course, don't plan that way. We're able to introduce [abstract concepts](#) that throw away that huge mass of irrelevant detail and focus only on what is important."

Even state-of-the-art robots aren't capable of that kind of [abstraction](#). When we see demonstrations of robots planning for and performing multistep tasks, "it's almost always the case that a programmer has explicitly told the robot how to think about the world in order for it to make a plan," Konidaris said. "But if we want robots that can act more autonomously, they're going to need the ability to learn abstractions on their own."

In computer science terms, these kinds of abstractions fall into two categories: "procedural abstractions" and "perceptual abstractions." Procedural abstractions are programs made out of low-level movements composed into higher-level skills. An example would be bundling all the little movements needed to open a door—all the motor movements involved in reaching for the knob, turning it and pulling the door open—into a single "open the door" skill. Once such a skill is built, you don't need to worry about how it works. All you need to know is when to run it. Roboticists—including Konidaris himself—have been studying how to make robots learn procedural abstractions for years, he says.

But according to Konidaris, there's been less progress in perceptual abstraction, which has to do with helping a robot make sense of its pixelated

surroundings. That's the focus of this new research. "Our work shows that once a robot has high-level motor skills, it can automatically construct a compatible high-level symbolic representation of the world—one that is provably suitable for planning using those skills," Konidaris said.

about her surroundings," Konidaris said. "Doors need to be closed before they can be opened. You can't get the bottle out of the cupboard unless it's open, and so on. And she was able to learn them just by executing her skills and seeing what happens."

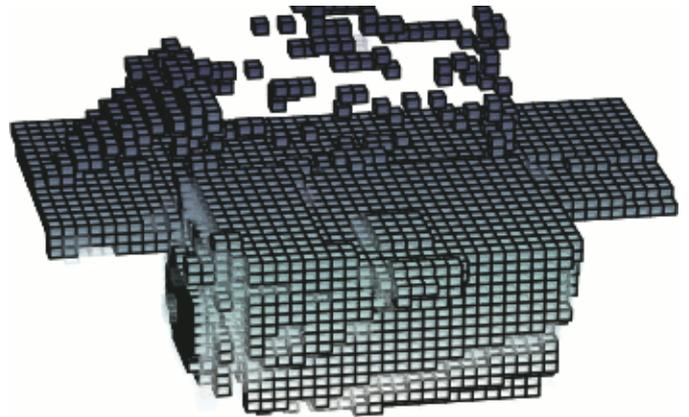
Learning abstract states of the world

For the study, the researchers introduced a robot named Anathema Device (or Ana, for short) to a room containing a cupboard, a [cooler](#), a switch that controls a light inside the cupboard, and a [bottle](#) that could be left in either the cooler or the cupboard. They gave Ana a set of high-level motor skills for manipulating the objects in the room—opening and closing both the cooler and the cupboard, flipping the switch and picking up a bottle. Then they turned Ana loose to try out her motor skills in the room, recording the sensory data from her cameras and actuators before and after each skill execution. Those data were fed into the machine-learning algorithm developed by the team.

The researchers showed that Ana was able to learn a very abstract description of the environment that contained only what was necessary for her to be able perform a particular skill. For example, she learned that in order to open the cooler, she needed to be standing in front of it and not holding anything (because she needed both hands to open the lid). She also learned the proper configuration of pixels in her visual field associated with the cooler lid being closed, which is the only configuration in which it's possible to open it.

She learned similar abstractions associated with her other skills. She learned, for example, that the light inside cupboard was so bright that it whited out her sensors. So in order to manipulate the bottle inside the cupboard, the light had to be off. She also learned that in order to turn the light off, the cupboard door needed to be closed, because the open door blocked her access to the switch. The resulting abstract representation distilled all that knowledge down from high-definition images to a text file, just 126 lines long.

"These were all the important abstract concepts



Just by executing her motor skills, Ana learned the appropriate configuration of pixels associated with a cooler having a closed lid, a necessary condition for her to run her "open the cooler" skill. Credit: Brown University

Planning in the abstract

Once Ana was armed with her learned abstract representation, the researchers asked her to do something that required some planning: take the bottle from the cooler and put it in the cupboard.

As they hoped she would, Ana navigated to the cooler and opened it to reveal the bottle. But she didn't pick it up. Instead, she planned ahead. She realized that if she had the bottle in her gripper, then she wouldn't be able to open the cupboard, because doing so requires both hands. So after she opened the cooler, she navigated to the cupboard. There she saw that the light switch in the "on" position, and she realized that opening the cupboard would block the switch, so she turned the switch off before opening the cupboard, returning to the cooler and retrieving the bottle, and finally placing it in the cupboard. In short, she planned ahead, identifying problems and fixing them before

they could occur.

"We didn't provide Ana with any of the [abstract representations](#) she needed to plan for the task," Konidaris said. "She learned those abstractions on her own, and once she had them, planning was easy. She found that plan in only about four milliseconds."

Konidaris says the research provides an important theoretical building block for applying [artificial intelligence](#) to robotics. "We believe that allowing our robots to plan and learn in the abstract rather than the concrete will be fundamental to building truly intelligent robots," he said. "Many problems are often quite simple, if you think about them in the right way."

Konidaris' coauthors on the paper were Leslie Pack Kaelbling and Tomas Lozano-Perez from MIT. The research was supported by an award from the Defense Advanced Research Projects Agency and by MIT's Intelligence Initiative.

More information: George Konidaris, Leslie Pack Kaelbling and Tomas Lozano-Perez (2018) "From Skills to Symbols: Learning Symbolic Representations for Abstract High-Level Planning", Volume 61, pages 215-289 , [DOI: 10.1613/jair.5575](https://doi.org/10.1613/jair.5575) , jair.org/papers/paper5575.html

Provided by Brown University

APA citation: Researchers help robots think and plan in the abstract (2018, February 8) retrieved 17 January 2022 from <https://techxplore.com/news/2018-02-robots-abstract.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.