


A step closer to self-aware machines—engineers create a robot that can imagine itself

30 January 2019



An image of the deformed robotic arm in multiple poses as it was collecting data through random motion. Credit: Robert Kwiatkowski/Columbia Engineering

Robots that are self-aware have been science fiction fodder for decades, and now we may finally be getting closer. Humans are unique in being able to imagine themselves—to picture themselves in future scenarios, such as walking along the beach on a warm sunny day. Humans can also learn by revisiting past experiences and reflecting on what went right or wrong. While humans and animals acquire and adapt their self-image over their lifetime, most robots still learn using human-provided simulators and models, or by laborious, time-consuming trial and error. Robots have not learned simulate themselves the way humans do.

Columbia Engineering researchers have made a major advance in robotics by creating a [robot](#) that learns what it is, from scratch, with zero prior knowledge of physics, geometry, or motor dynamics. Initially the robot does not know if it is a spider, a snake, an arm—it has no clue what its shape is. After a brief period of "babbling," and within about a day of intensive computing, their

robot creates a self-simulation. The robot can then use that self-simulator internally to contemplate and adapt to different situations, handling new tasks as well as detecting and repairing damage in its own body. The work is published today in *Science Robotics*.


To date, robots have operated by having a human explicitly model the robot. "But if we want robots to become independent, to adapt quickly to scenarios unforeseen by their creators, then it's essential that they learn to simulate themselves," says Hod Lipson, professor of mechanical engineering, and director of the Creative Machines lab, where the research was done.

For the study, Lipson and his Ph.D. student Robert Kwiatkowski used a four-degree-of-freedom articulated robotic arm. Initially, the robot moved randomly and collected approximately one thousand trajectories, each comprising one hundred points. The robot then used deep learning, a modern machine learning technique, to create a self-model. The first self-models were quite inaccurate, and the robot did not know what it was, or how its joints were connected. But after less than 35 hours of training, the self-model became consistent with the physical robot to within about four centimeters. The self-model performed a pick-and-place task in a closed loop system that enabled the robot to recalibrate its original position between each step along the trajectory based entirely on the internal self-model. With the closed loop control, the robot was able to grasp objects at specific locations on the ground and deposit them into a receptacle with 100 percent success.

Even in an open-loop system, which involves performing a task based entirely on the internal self-model, without any external feedback, the robot was able to complete the pick-and-place task with a

44 percent success rate. "That's like trying to pick up a glass of water with your eyes closed, a process difficult even for humans," observed the study's lead author Kwiatkowski, a Ph.D. student in the computer science department who works in Lipson's lab.

The self-modeling robot was also used for other tasks, such as writing text using a marker. To test whether the self-model could detect damage to itself, the researchers 3-D-printed a deformed part to simulate damage and the robot was able to detect the change and re-train its self-model. The new self-model enabled the robot to resume its pick-and-place tasks with little loss of performance.



An image of the intact robotic arm used to perform all of the tasks Credit: Robert Kwiatkowski/Columbia Engineering

Lipson, who is also a member of the Data Science Institute, notes that self-imaging is key to enabling robots to move away from the confinements of so-called "narrow-AI" towards more general abilities. "This is perhaps what a newborn child does in its crib, as it learns what it is," he says. "We conjecture that this advantage may have also been the evolutionary origin of self-awareness in humans. While our robot's ability to imagine itself is still crude compared to humans, we believe that this ability is on the path to machine self-awareness."

Lipson believes that robotics and AI may offer a

fresh window into the age-old puzzle of consciousness. "Philosophers, psychologists, and cognitive scientists have been pondering the nature of self-awareness for millennia, but have made relatively little progress," he observes. "We still cloak our lack of understanding with subjective terms like 'canvas of reality,' but robots now force us to translate these vague notions into concrete algorithms and mechanisms."

Lipson and Kwiatkowski are aware of the ethical implications. "Self-awareness will lead to more resilient and adaptive systems, but also implies some loss of control," they warn. "It's a powerful technology, but it should be handled with care."

The researchers are now exploring whether robots can model not just their own bodies, but also their own minds, whether robots can think about thinking.

More information: R. Kwiatkowski et al., "Task-agnostic self-modeling machines," *Science Robotics* (2019). robotics.sciencemag.org/lookup.../scirobotics.aau9354

Provided by Columbia University School of Engineering and Applied Science

APA citation: A step closer to self-aware machines—engineers create a robot that can imagine itself (2019, January 30) retrieved 22 September 2019 from <https://techxplore.com/news/2019-01-closer-self-aware-machinesengineers-robot.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.