

# New 'deep learning' technique enables robot mastery of skills via trial and error

May 21 2015, by Sarah Yang



This team of UC Berkeley researchers has developed algorithms that enable their PR2 robot, nicknamed BRET for Berkeley Robot for the Elimination of Tedious Tasks, to learn new tasks through trial and error. Shown, left to right, are Chelsea Finn, Pieter Abbeel, BRET, Trevor Darrell and Sergey Levine. Credit: UC Berkeley Robot Learning Lab

UC Berkeley researchers have developed algorithms that enable robots to learn motor tasks through trial and error using a process that more closely approximates the way humans learn, marking a major milestone in the field of artificial intelligence.

They demonstrated their technique, a type of [reinforcement learning](#), by having a robot complete various tasks—putting a clothes hanger on a rack, assembling a toy plane, screwing a cap on a water bottle, and more—without pre-programmed details about its surroundings.

"What we're reporting on here is a new approach to empowering a robot to learn," said Professor Pieter Abbeel of UC Berkeley's Department of Electrical Engineering and Computer Sciences. "The key is that when a robot is faced with something new, we won't have to reprogram it. The exact same software, which encodes how the robot can learn, was used to allow the robot to learn all the different tasks we gave it."

The latest developments will be presented on Thursday, May 28, in Seattle at the International Conference on Robotics and Automation (ICRA). Abbeel is leading the project with fellow UC Berkeley faculty member Trevor Darrell, director of the Berkeley Vision and Learning Center. Other members of the research team are postdoctoral researcher Sergey Levine and Ph.D. student Chelsea Finn.

The work is part of a new People and Robots Initiative at UC's Center for Information Technology Research in the Interest of Society (CITRIS). The new multi-campus, multidisciplinary research initiative seeks to keep the dizzying advances in [artificial intelligence](#), robotics and automation aligned to human needs.

"Most robotic applications are in controlled environments where objects are in predictable positions," said Darrell. "The challenge of putting robots into real-life settings, like homes or offices, is that those

environments are constantly changing. The robot must be able to perceive and adapt to its surroundings."

## **Neural inspiration**

Conventional, but impractical, approaches to helping a robot make its way through a 3D world include pre-programming it to handle the vast range of possible scenarios or creating simulated environments within which the robot operates.

Instead, the UC Berkeley researchers turned to a new branch of artificial intelligence known as deep learning, which is loosely inspired by the neural circuitry of the human brain when it perceives and interacts with the world.

"For all our versatility, humans are not born with a repertoire of behaviors that can be deployed like a Swiss army knife, and we do not need to be programmed," said Levine. "Instead, we learn new skills over the course of our life from experience and from other humans. This learning process is so deeply rooted in our nervous system, that we cannot even communicate to another person precisely how the resulting skill should be executed. We can at best hope to offer pointers and guidance as they learn it on their own."


In the world of artificial intelligence, deep learning programs create "neural nets" in which layers of artificial neurons process overlapping raw sensory data, whether it be sound waves or image pixels. This helps the robot recognize patterns and categories among the data it is receiving. People who use Siri on their iPhones, Google's speech-to-text program or Google Street View might already have benefited from the significant advances deep learning has provided in speech and vision recognition.

Applying deep reinforcement learning to motor tasks has been far more challenging, however, since the task goes beyond the passive recognition of images and sounds.

"Moving about in an unstructured 3D environment is a whole different ballgame," said Finn. "There are no labeled directions, no examples of how to solve the problem in advance. There are no examples of the correct solution like one would have in speech and vision recognition programs."

## **Practice makes perfect**

In the experiments, the UC Berkeley researchers worked with a Willow Garage Personal Robot 2 (PR2), which they nicknamed BRETT, or Berkeley Robot for the Elimination of Tedious Tasks.



BRETT is shown here learning how to screw a cap onto a water bottle. Credit: UC Berkeley Robot Learning Lab

They presented BRETT with a series of [motor tasks](#), such as placing blocks into matching openings or stacking Lego blocks. The algorithm controlling BRETT's learning included a reward function that provided a score based upon how well the robot was doing with the task.

BRETT takes in the scene, including the position of its own arms and hands, as viewed by the camera. The algorithm provides real-time feedback via the score based upon the robot's movements. Movements that bring the robot closer to completing the task will score higher than those that do not. The score feeds back through the neural net, so the robot can learn which movements are better for the task at hand.

This end-to-end training process underlies the robot's ability to learn on its own. As the PR2 moves its joints and manipulates objects, the algorithm calculates good values for the 92,000 parameters of the neural net it needs to learn.

With this approach, when given the relevant coordinates for the beginning and end of the task, the PR2 could master a typical assignment in about 10 minutes. When the robot is not given the location for the objects in the scene and needs to learn vision and control together, the learning process takes about three hours.

Abbeel says the field will likely see significant improvements as the ability to process vast amounts of data improves.

"With more data, you can start learning more complex things," he said. "We still have a long way to go before our robots can learn to clean a house or sort laundry, but our initial results indicate that these kinds of [deep learning](#) techniques can have a transformative effect in terms of enabling robots to learn complex tasks entirely from scratch. In the next five to 10 years, we may see significant advances in [robot](#) learning capabilities through this line of work."

Provided by University of California - Berkeley

Citation: New 'deep learning' technique enables robot mastery of skills via trial and error (2015, May 21) retrieved 9 April 2024 from <https://phys.org/news/2015-05-deep-technique-enables-robot-mastery.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>
--