


'Deep learning'—a giant step for robots

March 7 2016, by Wallace Ravven



BRETT folds a towel under the watchful eyes of Pieter Abbeel and Ph.D. candidate Sandy Huang. BRETT uses deep learning to distinguish corners and edges, as well as how wrinkly certain areas are, and matches what it sees to past

experience to determine what to do next. Credit: Peg Skorpinski

The prospect of robots that can learn for themselves—through artificial intelligence and adaptive learning—has fascinated scientists and movie-goers alike. Films like *Short Circuit*, *Terminator*, *Millennial Man*, *Chappie* and *Ex Machina* flirt with the idea of a machine intelligence beyond the restricted rules of a set program. Pieter Abbeel is on a mission to create architectures to do just that. He is part of a growing cadre of scientists exploring deep machine learning.

Robots today can be programmed to reliably carry out a straightforward task over and over, such as installing a part on an assembly line. But a [robot](#) that can respond appropriately to changing conditions without specific instructions for how to do so has remained an elusive goal.

A robot that could learn from experience would be far more versatile than one needing detailed, baked-in instructions for each new act. It could rely on what [artificial intelligence](#) researchers call deep learning and reinforcement learning.

Deep learning enables the robot to perceive its immediate environment, including the location and movement of its limbs. Reinforcement learning means improving at a task by trial and error. A robot with these two skills could refine its performance based on real-time feedback.

For the past 15 years, Berkeley robotics researcher Pieter Abbeel has been looking for ways to make robots learn. In 2010 he and his students programmed a robot they named BRET (Berkeley Robot for the Elimination of Tedious Tasks) to pick up different sized towels, figure out their shape and neatly fold them.

The key instructions allowed the robot to visualize the towel's limp shape when held by one gripper and its outline when held by two. It may not seem like much but the challenge was daunting for the robot. After as many as a hundred trials—holding a towel in different places each time—BRETT knew the towel's size and shape and could start folding. A YouTube video of BRETT's skills was viewed hundreds of thousands of times.

"The algorithms instructed the robot to perform in a very specific set of conditions, and although it succeeded, it took 20 minutes to fold each towel," laughs Abbeel, associate professor of electrical engineering and computer science.

"We stepped back and asked 'How can we make it easier to equip robots with the ability to perfect new skills so that we can apply the learning process to many different skills?'"

This year in a first for the field Abbeel gave a new version of BRETT the ability to improve its performance through both deep learning and reinforcement learning. The deep learning component employs so-called neural networks to provide moment-to-moment visual and sensory feedback to the software that controls the robot's movements.



Pieter Abbeel and BRET, the Berkeley Robot for Elimination of Tedious Tasks. BRET has mastered a range of skills, including folding laundry, knot-tying, and basic assembly. Credit: Peg Skorpinski

With these programmed skills, BRET learned to screw a cap onto a bottle, to place a clothes hanger on a rack and to pull out a nail with the claw end of a hammer.

Its onboard camera allowed BRETT to pinpoint the nail to be extracted, as well as the position of its own arms and hands. Through trial and error, it learned to adjust the vertical and horizontal position of the hammer claw as well as maneuver the angle to the right position to pull out the nail.

The deep [reinforcement learning](#) strategy opens the way for training robots to carry out increasingly complex tasks. The achievement gained widespread attention, including an article in The New York Times.

BRETT learned to complete his chores in 30 to 40 trials, with each attempt taking only a few seconds. Still, he has more trial and error ahead: Learning to screw a cap on a bottle doesn't prepare him to screw a lid on a jar. Instead, he re-starts learning as if he had never mastered caps and bottles. Abbeel has begun research aimed at enabling robots to do something humans take for granted, generalize from one task to another.

Starting this year, the Bakar Fellows Program will support Abbeel's lab with \$75,000 a year for five years to help him refine the deep-learning strategy and move the research towards commercial viability. In addition to financial support, the Bakar Fellows Program provides mentoring in such crucial areas as the intricacies of venture capital and strategies to secure intellectual property rights.

"The Bakar support will allow us to improve the robot's [deep-learning](#) ability and to apply a learned skill to new tasks," Abbeel says.

Applications for such a skilled robot might range from helping humans with tedious housekeeping chores all the way to assisting in highly detailed surgery. In fact, Abbeel says, "Robots might even be able to teach other robots."

Provided by University of California - Berkeley

Citation: 'Deep learning'—a giant step for robots (2016, March 7) retrieved 19 April 2024 from <https://techxplore.com/news/2016-03-deep-learning-a-giant-robots.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.