

Ken Goldberg

The New York Times

In the Future, Warehouse Robots Will Learn on Their Own

By Cade Metz
September 10, 2017

SF

248 Utah Street
SF, CA 94103
+ 415 399 1439

NY

313 W 14th Street 2F
New York, NY
By appointment only

WEB

www.cclarkgallery.com



Jeff Mahler, left, and Ken Goldberg have studied ways to help robots figure out tasks on their own at the University of California, Berkeley. Credit Jason LeCras for The New York Times

BERKELEY, Calif. — The robot was perched over a bin filled with random objects, from a box of instant oatmeal to a small toy shark. This two-armed automaton did not recognize any of this stuff, but that did not matter. It reached into the pile and started picking things up, one after another after another.

“It figures out the best way to grab each object, right from the middle of the clutter,” said Jeff Mahler, one of the researchers developing the robot inside a lab at the University of California, Berkeley.

For the typical human, that is an easy task. For a robot, it is a remarkable talent — something that could drive significant changes inside some of the world’s biggest businesses and further shift the market for human labor.

Today, robots play important roles inside retail giants like Amazon and manufacturing companies like Foxconn. But these machines are programmed for very specific tasks, like moving a particular type of container across a warehouse or placing a particular chip on a circuit board. They can't sort through a big pile of stuff, or accomplish more complex tasks. Inside Amazon's massive distribution centers — where sorting through stuff is the primary task — armies of humans still do most of the work.

The Berkeley robot was all the more remarkable because it could grab stuff it had never seen before. Mr. Mahler and the rest of the Berkeley team trained the machine by showing it hundreds of purely digital objects, and after that training, it could pick up items that weren't represented in its digital data set.

"We're learning from simulated models and then applying that to real work," said Ken Goldberg, the Berkeley professor who oversees the university's automation lab.

The robot was far from perfect, and it could be several years before it is seen outside research labs. Though it was equipped with a suction cup or a parallel gripper — a kind of two-fingered hand — it could reliably handle only so many items. And it could not switch between the cup and the gripper on the fly. But the techniques used to train it represented a fundamental shift in robotics research, a shift that could overhaul not just Amazon's warehouses but entire industries.

Rather than trying to program behavior into their robot — a painstaking task — Mr. Mahler and his team gave it a way of learning tasks on its own. Researchers at places like Northeastern University, Carnegie Mellon University, Google and OpenAI — the artificial intelligence lab founded by Tesla's chief executive, Elon Musk — are developing similar techniques, and many believe that such machine learning will ultimately allow robots to master a much wider array of tasks, including manufacturing.

"This can extend to tasks of assembly and more complex operations," said Juan Aparicio, head of advanced manufacturing automation at the German industrial giant Siemens, which is helping to fund the research at Berkeley. "That is the road map."

Physically, the Berkeley robot was nothing new. Mr. Mahler and his team were using existing hardware, including two robotic arms from the Swiss multinational ABB and a camera that captured depth.

What was different was the software. It demonstrated a new use for what are called neural networks. Loosely based on the network of neurons in the human brain, a neural network is a complex algorithm that can learn tasks by analyzing vast amounts of data. By looking for patterns in thousands of dog photos, for instance, a neural network can learn to recognize a dog.

Over the past five years, these algorithms have radically changed the way the internet's largest companies build their online services, accelerating the development of everything from image and speech recognition to internet search. But they can also accelerate the development of robotics.

The Berkeley team began by scouring the internet for CAD models, short for computer-aided design. These are digital representations of physical objects. Engineers, physicists and designers build them when running experiments or creating new products. Using these models, Mr. Mahler

and his team generated many more digital objects, eventually building a database of more than seven million items. Then they simulated the physics of each item, showing the precise point where a robotic arm should pick it up.

That was a large task, but the process was mostly automated. When the team fed these models into a neural network, it learned to identify a similar point on potentially any digital object with any shape. And when the team plugged this neural network into the two-armed robot, it could do the same with physical objects.

When facing a single everyday object with cylindrical or at least partly planar surfaces — like a spatula, a stapler, a cylindrical container of Froot Loops or even a tube of toothpaste — it could typically pick it up, with success rates often above 90 percent. But percentages dropped with more complex shapes, like the toy shark.

What's more, when the team built simulated piles of random objects and fed those into the neural network, it could learn to lift items from physical piles, too. Researchers at Brown University and Northeastern are exploring similar research, and the hope is that this kind of work can be combined with other methods.

Like Siemens and the Toyota Research Institute, Amazon is helping to fund the work at Berkeley, and it has an acute need for this kind of robot. For the past three years, the company has run a contest in which researchers seek to solve the “pick and place” problem. But the promise of machine-learning methods like the one used at Berkeley is that they can eventually extend to so many other areas, including manufacturing and home robotics.

“Picking an object up is the first thing you want a manipulator robot to do,” said Stefanie Tellex, a professor at Brown. “A lot of more sophisticated behavior begins with that. If you can't pick it up, game over.”

The research demonstrated how a task learned in the digital world can be transferred to the physical. Since the camera on Berkeley's robot could see depth, it captured three-dimensional images that were not unlike the CAD models the team uses to train its neural network.

Other researchers are developing ways for robots to learn directly from physical experience. For example, at Google, using an algorithmic technique called reinforcement learning, robots are training themselves to open doors through extreme trial and error. But this kind of physical training is both time consuming and expensive. Digital training is more efficient.

For this reason, some organizations are hoping to train robots using complex virtual worlds — digital recreations of our physical environment. If a system can train itself to navigate a car racing game like Grand Theft Auto, the thinking goes, it can navigate real roads.

This is still largely theory. But at places like Berkeley and Northeastern, researchers are showing that digital learning can indeed make the leap into the real world.

“This is a challenge,” said Rob Platt, a professor at Northeastern. “But it's a challenge we're dealing with.”