'Robobarista' can figure out your new coffee machine

April 15 2015, by Bill Steele

In the near future we may have household robots to handle cooking, cleaning and other menial tasks. They will be teachable: Show the robot how to operate your coffee machine, and it will take over from there.

But suppose you buy a new, different coffee maker. Will you have to start over?

"The robot already has seen two or three coffee machines; it should be able to figure out how to use this one," said Ashutosh Saxena, assistant professor of computer science. In robotics work up to now, he noted, a robot must be trained for each task and always positioned in the same relationship to the machine and its controls.

In his Robot Learning Lab, Saxena is making robots more adaptable. A new deep-learning algorithm developed by Saxena and graduate student Jaeyong Sung enables a robot to operate a machine it has never seen before, by consulting the instruction manual – probably available online – and drawing on its experience with other machines that have similar controls. Saxena will describe their work in a woekshop on "learning from demonstraion"at the 2015 Robotic Science and Systems conference July 16 in Rome.

One thing that makes this hard is the "noise" in natural language instructions. Do you turn on the machine with a "knob" or a "switch?" Do you dispense coffee by pulling a "handle" or a "lever?" And then, where is that control on the machine, and what's the proper way to manipulate it? For this, the robot draws on a database of recorded actions.

"We use a deep learning neural network that can tell the robot which action in a database is the closest to the one it has to perform," Sung explained. Deep learning classifies information in layers, moving down from general to specific. It works best with a very large database, so the researchers turned to crowdsourcing to collect a large library of actions. Through the Amazon Mechanical Turk service, where people can be recruited to do simple online tasks for small payments, they invited hundreds of visitors to guide a robot through the motions to perform various tasks described in a set of printed instructions. To make crowdsourcing possible, the researchers created a Web interface that enables a user to guide an imaginary robot arm, almost like playing a

video game.

From the database the robot also learns to identify various kinds of controls by their shape rather than location, and to relate them to the various labels that might be used in the instructions.

Using an espresso machine as their first test bed, the researchers refer to their robot as a robobarista. The Italian word barista refers to a server in a coffee shop. (They later learned that there is a self-contained coffee machine on the market with that name.) But they have trained and tested their robot with 116 different appliances, including juice makers, lamps, a soda machine and even bathroom sinks. Actions such as "turning a knob" can be transferred from one kind of knob to another. Turning a knob on a sink is a similar action to turning a timer control on a slow cooker, and pressing a lever on a water cooler can be transferred to pressing a lever on a coffee machine.

The robot, equipped with a 3-D camera (in the lab, a Microsoft Kinect that combines a video image with a laser rangefinder), begins with a "point cloud" – a list of the X, Y and Z coordinates of every point in an image. After translating the label in the instruction manual, it locates the control in the point cloud and consults the crowdsourced model to plan the trajectory the robot arm will follow to manipulate the control.

"Instead of trying to figure out how to operate an [espresso machine](#), we figure out how to operate each part of it," Sung explained. In tests so far on a variety of different machines, the robot has performed with 60 percent accuracy when operating a machine it has never seen before. One problem is that glossy parts that reflect complex patterns of light are sometimes hard to identify by shape.

In the future, the researchers said, robots may need tactile feedback to operate controls properly and visual monitoring during execution to

avoid collisions. Robots may also learn to use trial and error with unfamiliar controls and have routines to recover from failure. The action model, Sung said, eventually will be available in the online RoboBrain database Saxena has created for robots everywhere to consult.

  More information: You can try teaching a robot yourself at the Robobarista website: robobarista.cs.cornell.edu/

Provided by Cornell University