Google unveils progress in robot that learns how to pick up objects using neural-network

March 10 2016, by Bob Yirka

(Tech Xplore)—Google, the search engine giant, as most are aware, has been working on many technology projects, from self-driving cars to augmented reality eyewear. One area that has not received as much press is robotics. Over the past several years, Google has purchased several robotics companies, most notably Boston Dynamics, maker of Big Dog and other mobile robots. But the company has also invested heavily in

the development and enhancement of neural networks—computer and software systems that gain capabilities by learning, rather than having everything programmed in from the start. In this latest update, Google has posted an entry on its Google Research Blog (along with videos) by Sergey Levine and has uploaded a paper to the *arXiv* prepress server outlining progress made with robot arm/graspers guided by a neural network.

As Levine notes, the conventional approach to getting a robot arm and hand to reach out and grab something is to program in code that has the robot go through several steps which include things like scanning an object to be picked up and then using a previously written routine to manipulate the arm and hand movement in a way that will result in the object being grabbed and moved. But, as he also notes, a much better approach would be to follow the animal model—we do not do a lot of scanning and analyzing when we want to pick something up, we just reach over a grab it.

This ability comes about due to the learning process we go through as toddlers. To that end, Google has been building robots and testing them with a neural network approach—instead of telling a robot how to grasp an object and pull it up and out of a bin, it is given a general idea of what it means to grab and pick something up, imagery and a bin full of objects to pick up, one at a time. Noting that children take years to master this process, the team has taken a blunt force approach—they set up several robot arms with graspers next to one another, and had each reach down into a bin of assorted objects and pick up a designated object, over and over. The robots have all been connected together so whatever one learns in picking something up, the others learns as well—sort of like the Borg from Star Trek.

The result, the team reports, is progress. After 800,000 practice grasps the robots are not able to grab and pick up objects any better than other

robots, but they have demonstrated some abilities that until now, have been strictly limited to living creatures—moving another object out of the way, for example, to get at the one they want, or learning to grab sponges by pinching them in the middle and end, as opposed to spreading the fingers the entire width of an object. The researchers cannot say for sure how far the training will go, but the hope is that after more practice, the robot graspers, like those of a child, will grow ever more agile, perhaps at some point equaling our own skill.

Abstract
We describe a learning-based approach to hand-eye coordination for robotic grasping from monocular images. To learn hand-eye coordination for grasping, we trained a large convolutional neural network to predict the probability that task-space motion of the gripper will result in successful grasps, using only monocular camera images and independently of camera calibration or the current robot pose. This requires the network to observe the spatial relationship between the gripper and objects in the scene, thus learning hand-eye coordination. We then use this network to servo the gripper in real time to achieve successful grasps. To train our network, we collected over 800,000 grasp attempts over the course of two months, using between 6 and 14 robotic manipulators at any given time, with differences in camera placement and hardware. Our experimental evaluation demonstrates that our method achieves effective real-time control, can successfully grasp novel objects, and corrects mistakes by continuous servoing.