A conversation with Russell Stewart, July 29, 2015

Participants

- Russell Stewart – PhD Student, Stanford University
- Nick Beckstead – Research Analyst, Open Philanthropy Project
- Holden Karnofsky – Managing Director, Open Philanthropy Project

Note: These notes were compiled by the Open Philanthropy Project and give an overview of the major points made by Russell Stewart.

Summary

The Open Philanthropy Project spoke with Mr. Russell Stewart of Stanford University as part of its investigation into the field of robotics. Conversation topics included current applications of robotics, recent progress, potential future areas of progress, and current challenges and limitations.

Applications of robotics that currently work well

Navigation with Simultaneous Localization and Mapping technology

Simultaneous Localization and Mapping (SLAM) technology uses a sensor, such as cameras or LIDAR, to build a map of the environment. This allows robots to navigate and understand their position in the environment at all times. Current applications include robots driving around warehouses and quadcopters flying around airports.

Navigation is essentially a solved problem, but the sensors that are currently in use are expensive. For example, the 3D laser technology used in Google Cars costs $30,000-$50,000. Advances in 3D cameras such as the Microsoft Kinect have reduced this cost more than 100X to $200, but this technology does not work outdoors in sunlight. Recent research has shown promising results for high quality purely optical SLAM with garden-variety cameras. These results may soon be readily applied in real-world scenarios, reducing cost to the point where it becomes feasible to deploy SLAM technology in mobile robots on a large scale.

Industrial robotics

Arguably, the most important applications of robotics today are in manufacturing. In an industrial setting, robotics refers to manufacturing processes that are completed entirely without human involvement, such as parts being cut by a computer numerical control (CNC) machine or assembled using robotic arms.

The role of robotics in manufacturing is growing. Once the initial investment to automate a manufacturing process with robotics is made, that process generally does not transition back to using human labor. The primary factor inhibiting the automation of more manufacturing processes is that the initial investment often exceeds the cost of human labor for products that do not enjoy large economies of scale.

Building and programming industrial robots currently requires a fair amount of engineering labor and some trial and error to determine how best to program the robot to make specific movements to perform the task at hand. In the future, the high initial investment needed for robotic automation could be addressed by the introduction of more intelligent robots that provide higher-level primitives to software and mechanical engineers. Such a robot would be intelligent enough that it can be told to do something without needing to program the specific steps. These robots reduce not only the engineering costs of the buyer, but also the production costs of the robot itself. This is because a more intelligent robot can be sold in larger volumes to many different types of manufacturers.

For example, high quality pick and place robots are currently mass-produced with great economies of scale, costing as little as $20,000, and can easily be customized to a specific assembly line process. Mr. Stewart expects new categories of general-purpose robots for robotic part assembly or quality inspection applications to become widely deployed for in the next 5-10 years.

Economic benefit

The increasing use of robotics in manufacturing in the US, combined with falling energy prices, is making it more competitive for manufacturers to operate domestically.

Cost

Industrial robots are expensive to build and program to do specific tasks, but inexpensive to operate once the technology is in place. With current technology, it is rarely cost-effective to consider automation when the current manufacturing labor costs are under $100,000-$200,000 per year. Making targeted philanthropic investments in engineering specific industrial robotic applications could potentially have a high impact in areas of manufacturing that are dangerous for humans and/or have a significant negative effect on health, but have not yet been automated because manufacturing costs are not high enough to justify the cost of building robots and writing software.

CNC machines

CNC machines such as lathes and mills increase product quality and, after an initial investment of about $200,000, reduce the manufacturing costs of everyday items. CNC machines are able to interpret computer-aided design (CAD) models and cut them from a block of aluminum. It takes some time to set up the machine to cut a particular part, so this is more efficient if multiple identical parts are being cut in a batch. A CNC machine takes a few minutes to cut each part, and one operator can oversee multiple machines.

The transition to industrial robotics

Mr. Stewart's father runs AGM Container Controls, Inc. (AGM), a manufacturing company in Arizona. As a small company, it does not have sufficient labor costs to justify a specialized robotics team to focus on manufacturing automation, but it has nonetheless gradually incorporated robots into its manufacturing process. In the last 20 years, AGM has transitioned from having parts cut with hand-guided machine tools at a local machine shop to using 7-8 CNC machines to cut parts that are designed by mechanical engineers. AGM continues to assemble its products by hand in most cases because it is not yet cost-effective to program robots to assemble products.

Many small to medium size American manufacturing companies share a similar profile. Even if they are a market leader in an important product category, they may lack the volume to justify automation costs with the current generation of robotics technology. Improvements in high tech robotics could lead to significant reductions in up-front programming costs, spawning a dramatic increase in the automation of manufactured products with sales smaller than $1 million per year.

Active areas of research in robotics

Teaching by example

Some researchers are trying to make more intelligent industrial robots that can be more efficiently and cheaply taught to do new things by interacting with objects rather than being programmed to perform specific tasks. For example, one goal might be to make robots that would be able to perform tasks such as picking up all objects in a room to tidy up.

Coping with uncertainty

The capabilities of robots are already superior to that of humans when they are able to fully discretize a state space and understand how exactly how different motors impact transitions through the state space. Self-parking cars are one such example. However, this approach is not feasible in larger and more complex state spaces, such as humanoid walking robots. A lot of work is being done to find algorithms that require fewer assumptions and can be used in state spaces that are too large to explicitly model.

Underactuated robotics

Fully actuated robots are able to control all of their movements at all times, and methods for controlling fully actuated are well understood. For example, a manufacturing robot with an arm can be told where to move, and be in complete control of its motion throughout the trajectory. An underactuated robot does not have total control over its movement. For example, a walking robot is actuated in all places except in the interactions between the foot and the ground. It is necessary to do extensive advance planning to prevent the robot from falling over.

The Advanced Step in Innovative Mobility (ASIMO) robot designed by Honda Motor Co., Ltd. is an underactuated robot that is capable of doing interesting tricks and remaining stable as it walks around. When watching ASIMO, it appears that much of humanoid robotics is solved. But under the hood, ASIMO uses about 20 times more power than an optimal robot could and requires very strong motors to function. It does not have a spring in its step, which would make it more power efficient and capable of moving faster.

Many researchers are currently working on finding smarter algorithms to keep such humanoid robots upright. Current research into making fully actuated robots is inspired by animal movements and the human brain. Russ Tedrake, a robotics professor at Massachusetts Institute of Technology (MIT) with a background in neuroscience, is one researcher that is solving these kinds of problems. He is using algorithms that are inspired by biological behavior and enable much more human-like humanoid movement than the ASIMO robot.

Planning

Professor Leslie Kaelbling, a researcher at MIT, works on enabling robots to plan actions such as moving objects out of the way when crossing a cluttered room, or traveling from Boston to San Francisco without building a complete map of the United States. Hand engineered solutions can be compiled for a specific route, but more research needs to be done on how to scale these skills up into a more general ability. Many researchers are working on this, and Mr. Stewart thinks that progress may involve the combination of techniques from machine learning and the symbolic AI developed in the 1970s.

Low-power robotics

While most progress in robotics has involved improvements to software, Joshua Smith, a professor of computer science and electrical engineering at the University of Washington, is working on improving hardware with his research on low-power robotics. Professor Smith has developed robots with increasingly small power footprints. Recently, he has even developed some robots that have no battery at all! Instead, these robots are powered by ambient energy in radio waves, and can be used as sensors in a room without worrying about battery lifetime.

Applying machine learning to robotics

Early robotics had a lot of success using direct reasoning over discretized state spaces with dynamic programming, but it is becoming clear that computers are not going to get fast enough to continue to use this method to solve new problems. There is now an effort to find creative ways to control motion without explicitly modeling the desired motion. One way to do this is to have the robot learn by example.

Mr. Stewart has not done a lot of work on deep learning, which shows significant promise for robotics applications. Deep Learning is a somewhat nascent field, but it may be able to solve problems that have traditionally escaped the capabilities of robots by simply showing a robot enough examples and having it teach itself. For example, a robot can be trained to fold a towel by feeding it the camera image of towel as it is being folded, and training it issue the same motor commands as those the human operator originally executed. Deep Learning is promising because it does not require the programmer to specify exactly how the robot should represent the world. This is valuable because robots that learn by teaching themselves seem to be smarter than robots that are told exactly how to think by humans

Quadcopters

There has not been a lot of recent progress on quadcopter navigation, but in the early 2000s, machine learning enabled quadcopters to learn to navigate efficiently and facilitated Professor Andrew Ng of Stanford's work on the Stanford Autonomous Helicopter Project.

Advances in machine learning also reduced the demand for the development of new heuristics as robots began to be able to learn from examples.

Learning by example has already been used to write proportional-integral-derivative controllers (PID controllers) for quadcopters. Quadcopters have four motors and four directions of motion (roll, pitch, yaw, and vertical thrust), and PID controllers adjust each of these parameters to keep the quadcopter stable. The best way to develop an effective PID controller is to give it access to a lot of data and let it write its own algorithms.

Quadcopters made at the University of Pennsylvania have a vision system that gives them complete odometry, which allows them to learn the dynamics of how they move and ultimately results in a quadcopter that has greater control.

Private sector progress in robotics

Compared to academics, the private sector spends less effort trying to maximize the intelligence of robots and more effort trying to make robots faster and cheaper to build. Robotics companies have made some progress in this area, but their impact is still smaller than academics in many areas, in part because they are less likely to publish papers about their work.

Potential future progress in robotics

In the next 5-10 years, Mr. Stewart hopes to see progress in the following areas: using different types of sensors, learning by example, and higher-level reasoning. It is also possible that machine learning and symbolic logic may be combined into a joint approach.

Using new types of sensors

Most robots currently use highly calibrated force sensors to detect objects in their environment, but new deep learning algorithms may enable robots to use a wider variety of sensors with less rigid profiles, which would facilitate new applications of robotics. As an undergraduate, Mr. Stewart studied how mice use their whiskers to map their environment. Whiskers are nothing like current force sensors used in robots as they are highly flexible and measure force in the cheek of the mouse, rather than at the point of contact. Mr. Stewart would like to see algorithms that can integrate this type of information in the same way that mice do.

Learning by example

Algorithms that allow a robot to learn by interacting with its environment are getting better, which makes it easier to train robots to do new tasks. For example, there are currently algorithms that enable robots to fold towels, but these algorithms require hundreds of examples. Ideally, it would be possible to teach a robot to fold towels by demonstrating the task only 3-4 times. Learning from fewer examples would be a critical step in making these types of robots easily applicable in real-world scenarios.

Combining machine learning with symbolic logic

Machine learning has begun to dominate the artificial intelligence community and reduce the popularity of symbolic logic algorithms. There is some interest in combining the two approaches, but this is not yet possible because currently robots that use machine learning only learn by example, while robots that use symbolic logic require precise representations. However, as machine learning matures, it may become possible to combine the two. This seems to be a promising area for research.

Current challenges and limitations

Interacting with objects

Robots have trouble picking up objects and interacting with non-rigid objects. There may be breakthroughs in this area using the same type of algorithms that were used to achieve breakthroughs in object detection in 2012. Learning to interact with objects and avoid breaking things while walking around may be the final step necessary to bring robotics into a more consumer-oriented domain.

Research capacity

Problems such as how to use low-level sensors to map the environment require creative solutions, but it can be difficult to predict which ideas will be successful, so it is best to test a wide range of ideas. However, there are a limited number of researchers who have the right training and experience to work on solving problems in robotics, which limits the speed of progress.

It can be difficult to improve upon successful algorithmic approaches, because often these approaches have already been implemented to their full extent and leave little room for further research. For this reason, a lot of future robotics research will require coming up with entirely new approaches.

Other people to talk to

- Joshua Smith – Associate Professor, Department of Computer Science and Engineering and Department of Electrical Engineering, University of Washington
- Russ Tedrake – X Consortium Associate Professor of Electrical Engineering and Computer Science, Aeronautics and Astronautics, and Mechanical Engineering, MIT; Director, Center for Robotics at the Computer Science and Artificial Intelligence Lab
- Dieter Fox – Professor, Department of Computer Science and Engineering, University of Washington
- Peter Abbeel – Associate Processor, Department of Electrical Engineering and Computer Sciences, UC Berkeley