

Building the Foundation for “Generation Robot”

Signal processing is helping to make robots a part of everyday life

Leaping off the pages of fiction and movie screens, robots are poised to create the most profoundly disruptive technological shift since the Industrial Revolution.

While robots have labored in numerous industries for many years—particularly in the automotive and manufacturing sectors—experts predict that a tipping point in robotic deployment is now at hand and that robots will soon join the business and consumer mainstream. To put it plainly, robots are set to become an essential part of everyday life, and life itself will never be the same.

Market research firm Tractica predicts that the global robotics industry will expand from US\$34.1 billion in 2016 to US\$226.2 billion by 2021, representing a compound annual growth rate (CAGR) of 46%. “The definition of a robot is in flux and traditional robot manufacturers that have been building and supplying robots for decades have seen this industry undergo a dramatic transformation in the past few years,” observes Tractica Research Director Aditya Kaul.

As robots transition out of the fictional world and into everyday life, signal processing is playing a major role in their development. In key areas, such as control, sensing, and modeling, signal processing is helping robots become our trusted new home, road, and workplace partners.

Digital Object Identifier 10.1109/MSP.2017.2715599
Date of publication: 6 September 2017

Aiming for easier control

Most current-generation robots possess only a limited amount of autonomy. In many industrial, health care, military, and even consumer applications, humans still need a direct way of instructing robots on how to move and act, as well as to take full control of a robot when things go awry. Yet existing robot control interfaces leave much to be desired. Most require a computer and mouse to independently control a range known as the “six degrees of freedom,” which enables a rigid body to operate within a three-dimensional space.

Yet, for most users, turning three virtual rings and adjusting arrows on a screen to move a robot into position to grab items or perform a specific task is as tricky as it sounds. For someone who isn’t an expert, the ring-and-arrow system is cumbersome and error prone. The approach certainly isn’t ideal for people who aren’t technically skilled, such as older people trying to control assistive robots at home.

A new interface designed by Georgia Institute of Technology researchers promises simpler and more efficient remote robot operation from almost any location without requiring significant training (Figure 1). The user just points and clicks on an item and then chooses an appropriate task grasp. “Instead of a series of rotations, lowering and raising arrows, adjusting the grip and guessing the correct depth of field, we’ve shortened the process to just two clicks,” says David Kent, the Georgia Tech doctoral robotics student who

led the project. Sonia Chernova, a Georgia Tech assistant professor in robotics, advised the research team.

Kent says that the new interface was developed to allow nonexpert users to effectively control a robot through their web browser. “There are many potential applications—basically anywhere you would have a remote human operator controlling a robot with a manipulator, including controlling robots for space exploration, deploying robots in dangerous environments, such as disaster response and search-and-rescue, and controlling assistive robots,” he says. Kent notes that besides being more intuitive for novice operators, intelligent interfaces can potentially reduce the workloads of experienced operators.

“Signal processing, coupled with the interface itself, is what makes our interface more effective than other commonly used robot teleoperation interfaces,” Kent states. “By analyzing and incorporating information from a depth image, or a point cloud, our interface can focus the operator’s interaction around relevant commands for the robot.”

The point-and-click control interface’s key element is an algorithm that calculates and ranks grasps based on their effectiveness for general manipulation. The process consists of the following steps.

- Based on an input point provided by the user clicking on a camera feed, the algorithm filters a point cloud of the environment down to a small region of interest.
- The algorithm next fits a plane to the region of interest using a sample

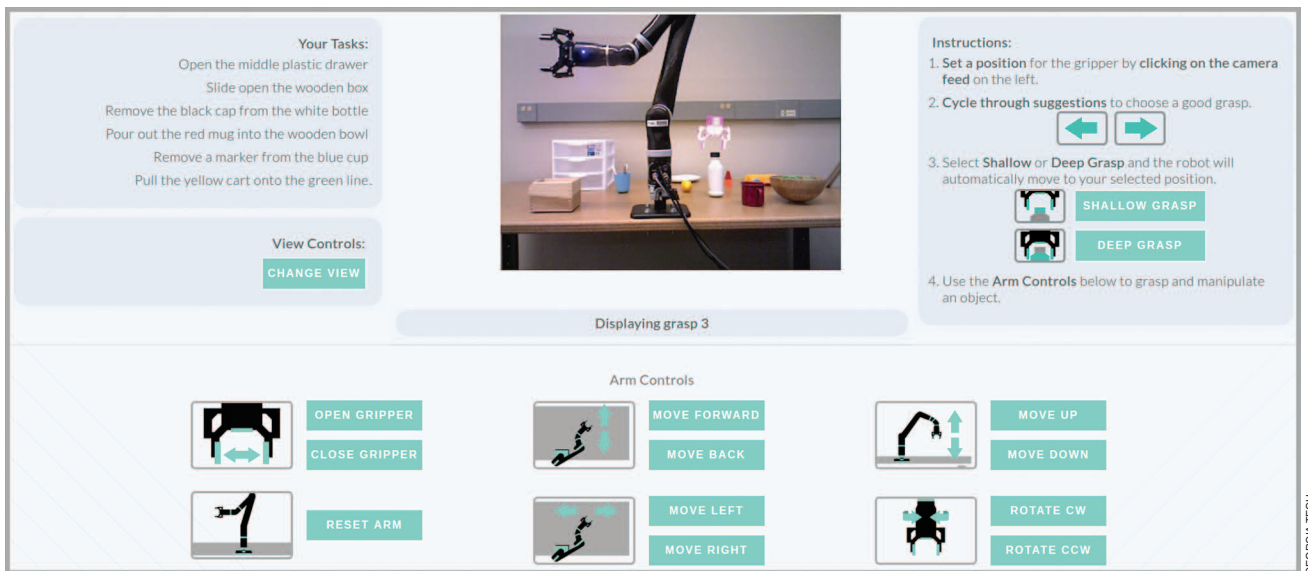


FIGURE 1. The test version of a display interface developed by Georgia Tech researchers that promises to allow nonexpert users to easily and effectively control a robot through a web browser.

consensus method. “We use the plane to calculate an orientation that’s perpendicular to the region of interest,” Kent says.

- The algorithm uses clustering, based on Euclidean distance between points, to identify individual components in the region of interest. “We use principal component analysis to extract a principal direction, or essentially an orientation, for each cluster,” Kent explains.
- Information calculated from the previous point cloud processing—the plane normal and the principal direction—is used to inform heuristics that ranks a set of grasp hypotheses calculated from AGILE (http://wiki.ros.org/agile_grasp), prioritizing grasps that align with the plane normal and principal direction.
- The best grasps are presented to the operator, who can then select a grasp for the robot to execute in a supervisory manner.

A control interface designed for non-expert users must be as foolproof as possible. “We have to make sure our system can safely execute any commands given by a user, which involves careful design and thorough testing,” he says.

Kent notes that the most significant remaining technical challenge the team faces involves the point clouds. “They’re constructed by observing a surface from

a certain viewing angle and, as such, they are missing data due to occlusion,” he says. “We’re essentially trying to extract useful information from a signal that’s missing a lot of data from the real-world object it is trying to represent, and, on top of that, the data is potentially noisy, depending on the reflectiveness and opacity of the objects in the environment.”

After testing college students on both existing ring-and-arrow systems and the new interface, the researchers discovered that the point-and-click method created significantly fewer errors, allowing participants to perform tasks more quickly and reliably. Overall, the point-and-click approach resulted in approximately one mistake per task, compared to nearly four per task for the ring-and-arrow interface. The study also found that point-and-click users could perform a specific task about 2 min faster than when using the traditional interface.

The project’s next step is deploying an online version of the interface that will allow authorized Internet users to remotely control a robot and teach it complex tasks.

Skin deep

The stereotypical robot is a rigid metal machine capable of tolerating extreme environments and external forces. Yet,

in many situations, it’s actually more beneficial to build a robot with a flexible, skin-like exterior that’s embedded with multiple sensors. Such a robot, like Co-mau’s recently introduced AURA (www.comau.com/EN/media/news/2016/06/aura), can act on its current environment or adjust its activities to meet a specific physical command, such as being tapped by a human to immediately stop or nudged to move a certain way or perform a specific action.

Covering a robot with flexible sensors requires a technology that’s both cost-effective and easy to manufacture in bulk quantities. Researchers at Massachusetts Institute of Technology’s (MIT’s) Computer Science and Artificial Intelligence Laboratory believe that three-dimensional (3-D) printing promises the best approach to this challenge. To demonstrate the feasibility of a flexible robot that can interpret and act on local conditions, the researchers have designed and built a device that responds to external mechanical stresses by changing the color of a spot on its surface.

The device was inspired by the golden tortoise beetle, or “goldbug.” Native to the Americas, the insect’s exterior is usually a golden tone. Yet the insect quickly switches to a reddish orange hue if poked or prodded or, technically speaking, mechanically stressed.




FIGURE 2. A prototype device that MIT researchers claim could lead to external robot sensors that can be inexpensively printed in bulk quantities.

According to Subramanian Sundaram, an MIT graduate student in electrical engineering and computer science, who led the project, organisms in nature are internally connected with a dense sensorimotor network that allows them to sense vast amounts of data and, more importantly, process and respond to the most relevant signals. “So we want to replicate the ‘sense-process-respond’ pipeline we see in a simple creature—the golden tortoise beetle,” he says. The goldbug senses external mechanical disturbances and changes the color/transmission of its reflective shell. “Replicating functions like these are extremely relevant for robots; a majority of contemporary robots operate without a high density of sensors on their external surfaces,” Sundaram explains. “In that sense, they are blind.”

Collaborating with Sundaram on the project is his advisor, Wojciech Matusik, an associate professor of electrical engineering and computer science, and Marc Baldo, a professor of electrical engineering and computer science and director of MIT’s Research Laboratory of Electronics. Other participants include Pitchaya Sitthi-Amorn, a former post-doctoral researcher in Matusik’s lab; Ziwen Jiang, an undergraduate electrical engineering and computer science student; and David Kim, a technical assistant in Matusik’s Computational Fabrication Group.

Printable electronics—flexible circuitry deposited onto a plastic substrate—has

been a major research area for decades. The MIT project, however, marks the first demonstration of printed electronics and printed substrates combined. “We use a custom-built multimaterial 3-D printer to print the entire composite,” Sundaram says. “We show devices with up to six materials printed together.”

The researchers’ prototype device (Figure 2) is approximately T-shaped, featuring a short, wide base and an elongated crossbar with a strip of silver running across its length. The base includes a pair of printed transistors and a circle of semiconducting polymer—dubbed a *pixel* by the researchers—that changes its color when the crossbar stretches, modifying the silver strip’s electrical resistance.

Signal processing plays a key role in the project. “Overall, we work on printing low-level signal processing elements,” Sundaram says. The main signal processing element in the printed composite is a single stage common source amplifier with a diode connected load. “We use it for thresholding as well,” Sundaram says. “However, at the high level, we believe that signal selection is an equally important problem when a vast number of sensors are present.”

Mechanical strain is measured by printed strain sensors connected in a resistive ladder configuration. “The signal is amplified using our single stage amplifier, and then the output is displayed

using an electrochromic pixel,” Sundaram says.

According to a paper published by Sundaram and his coresearchers, the monolithic integration of sensing, processing and response mechanisms allows transducing signals across mechanical, electrical and optical domains using low-power organic processors and sensors that can be powered by 1.5 V. Controlling multidomain properties with uniform resolution and without any external processing should enable advances in biologically inspired autonomous multifunctional systems with increased local signal processing efficiencies and levels of self-sufficiency currently only seen in nature.

“We believe that arranging multiple electrical devices in freeform 3-D is still a challenge,” Sundaram says. “Constructing the supporting polymer composite along with the electronic devices is useful for application in robotics and, more broadly, flexible and nonplanar electronics.”

Into the air

Signal processing lies at the heart of a novel new approach to robotic unmanned aerial vehicles (UAVs) developed by a pair of doctoral students at the National University of Singapore’s Unmanned System Research Group.

U-Lion is a reconfigurable hybrid UAV offering both vertical takeoff and landing (VTOL) and cruise flying capabilities (Figure 3). VTOL allows U-Lion to take off and land in small, tight spaces lacking a conventional runway. Cruise flying permits the aircraft to perform long-range and duration missions. “With the ability to fully open and retract the wings, its flying performance is optimized for both flying modes,” says coresearcher Kangli Wang, who worked on the project with fellow doctoral student Yijie Ke and advisor Ben M. Chen. “We have utilized advanced signal processing, modeling as well as control techniques to enable U-Lion with fully autonomous flying capabilities.”

Hybrid UAV technology is a hot research topic both academically and commercially. Several companies, such as Google (Project Wing), have devoted significant resources to developing hybrid

UAVs. Yet due to the radical structural differences between VTOL and fixed-wing UAVs, the researchers faced a huge challenge in combining the two functionalities into a single aircraft.

Wang believes that he and Ke have surmounted the fundamental obstacles blocking the development of a reliable hybrid UAV. Advanced modeling and control algorithms are used to overcome the uncertainties in transition between configurations, allowing the aircraft to achieve autonomous full-envelope flight. “Compared to conventional rotorcraft, our platform is much more efficient with much longer flying endurance,” he explains. “It can transit to a VTOL mode anytime for effective surveillance and monitoring.” Additionally, in comparison to other currently available hybrid UAV platforms, almost all of which are a simple combination of a fixed-wing aircraft and a rotorcraft, Wang and Ke designed U-Lion to use only a single propulsion system. “It is thus more compact, optimal, and efficient,” Wang says.

An autonomous hybrid UAV can be used for long-range and long-endurance flying missions. Yet the aircraft’s VTOL capability allows it to be used almost anywhere. “U-Lion has the potential to be employed for forestation and precision agriculture, geographical mapping, sea surveillance, powerline inspection, disaster reactions, and many other applications,” Wang says.

Signal processing makes the U-Lion possible. “We need to process signals



FIGURE 3. The U-Lion is a hybrid UAV that can take off and land vertically and fly horizontally at high speed.

NATIONAL UNIVERSITY OF SINGAPORE


from onboard sensors, including the accelerometer, gyroscope, magnetometer, barometer, airspeed sensor, and global positioning system (GPS) sensor to obtain and estimate the orientation, velocity and position of the aircraft for its whole flight envelope,” Wang continues. “Advanced signal processing techniques, such as extended Kalman filter (EKF) and sensor fusion, are of vital importance to the autonomous control of the UAV.”

Sampling data is retrieved from the accelerometer and gyro sensors at a frequency of 1 kHz. Magnetometer and barometer data are collected at a sampling rate of 200 Hz. “Some very carefully selected low-pass filters are applied to process the obtained raw data to get rid of high-frequency noises,” Wang says. An EKF is applied to estimate the orientation of the U-Lion. A sensor fusion technique is used to fuse orientation with GPS and barometer data for the estimation of the aircraft position and velocity.

The researchers opted to use an EKF for its relatively lower computational intensity. “As our platform is highly non-linear in nature, the best suitable choice for signal processing and sensor fusion would perhaps be using a UKF instead,” Wang says, adding, “It would certainly be the direction of our future development once we have integrated a more powerful processing hardware unit onboard.”

Wang says that he and his partner are now researching and developing platforms featuring an optimized mechanical structures and energy supply systems, as well as advanced signal processing and flight control techniques. After that, the next step will be creating platforms capable of carrying larger payloads on longer flights.

Author

John Edwards (jedwards@johnedwardsmedia.com) is a technology writer based in the Phoenix, Arizona, area. 

PRESIDENT’S MESSAGE *(continued from page 5)*

Consulting), Jill Nelson (George Mason University), Meng Wang (Rensselaer Polytechnic Institute), and Ervin Sejdic (University of Pittsburgh and 2016 recipient of the U.S. Presidential Early Career Award).

These distinguished professionals are committed to advancing society and mak-

ing a positive change through their work and expertise in signal processing. The IEEE SPS offers endless opportunities—providing educational avenues, career development, networking, and volunteerism—to help the bright minds of tomorrow begin their careers and make an impact. Through our collaborative efforts,

we can bring signal processing to the forefront as a viable and exciting career path for young professionals. Let’s get started and find our future colleagues!

