Lab Notes

NEWS FROM AROUND LINCOLN LABORATORY

COLLABORATIVE ROBOTICS

From Tools to Teammates: Integrating Robots on Human Teams

Advanced capabilities and algorithms developed for autonomous systems could streamline human-robot teaming in military operations

Consider a scenario in which a six-member dismounted Marine squad is tasked with raiding a compound that recent intelligence reports indicate is harboring a terrorist group. While it is unlikely that the Marines will be directly confronted by hostiles, it is possible they could be targeted by snipers. As the squad assesses the terrorist presence in the area, a companion robot equipped with a novel infrared sensor scans the horizon and ground for unexpected heat signatures and activity. Meanwhile, a swarm of unmanned aerial vehicles generates three-dimensional (3D) colored maps of the compound in real time.

Compound raid vignettes such as the one described above are being played out in Lincoln Laboratory's Autonomous Systems Laboratory as part of a basic research program in collaborative robotics, a new trend in the robotics industry that seeks to advance robot-human synergy. The program, which is being carried out by a team of technical staff, military fellows, and interns from Lincoln Laboratory's Control Systems Engineering, Informatics and Decision Support, and Embedded and Open Systems Groups, is motivated by the Department of Defense's growing interest in utilizing autonomous systems to enhance warfighter situational awareness. Begun in February 2014, the program is funded by the Office of Naval Research (ONR), which is investing in autonomy and unmanned systems research and development under their Naval Science and Technology Strategic Plan.

"The aim of the program is to make robots that can autonomously and seamlessly collaborate with humans as part of a team," says Mark Donahue, program manager. Current autonomous systems neither share a common language with humans nor operate within the same cognitive context (beliefs, knowledge base, cultural perspective, and mental state) or cognitive load (i.e., the amount of information an individual can process and retain at any one time). Because of these factors, human-robot communication and interaction have been limited. Robots lack the social intelligence to understand human goals and intentions, to adapt their behavior when circumstances or perspectives change, and to recognize human emotion; thus, they are often controlled and supervised by human operators. For robots to be seen as partners instead of sophisticated tools, robotics technology must be matured to a point at which robots are capable of learning, reasoning, and making decisions as human beings.

The Laboratory's program in collaborative robotics focuses on developing (1) enabling technologies for autonomous systems and (2) algorithms and cognitive models that autonomously present warfighters with mission-relevant data acquired by autonomous systems. With the goal of enhancing situational awareness among squad members, the program comprises four research objectives:

- Develop techniques for translating raw, calibrated sensor data from autonomous systems into actionable information (e.g., two-dimensional [2D] navigation maps, visual and semantic labels for objects within a map)
- Develop machine-learning algorithms to identify and prioritize information that is relevant given

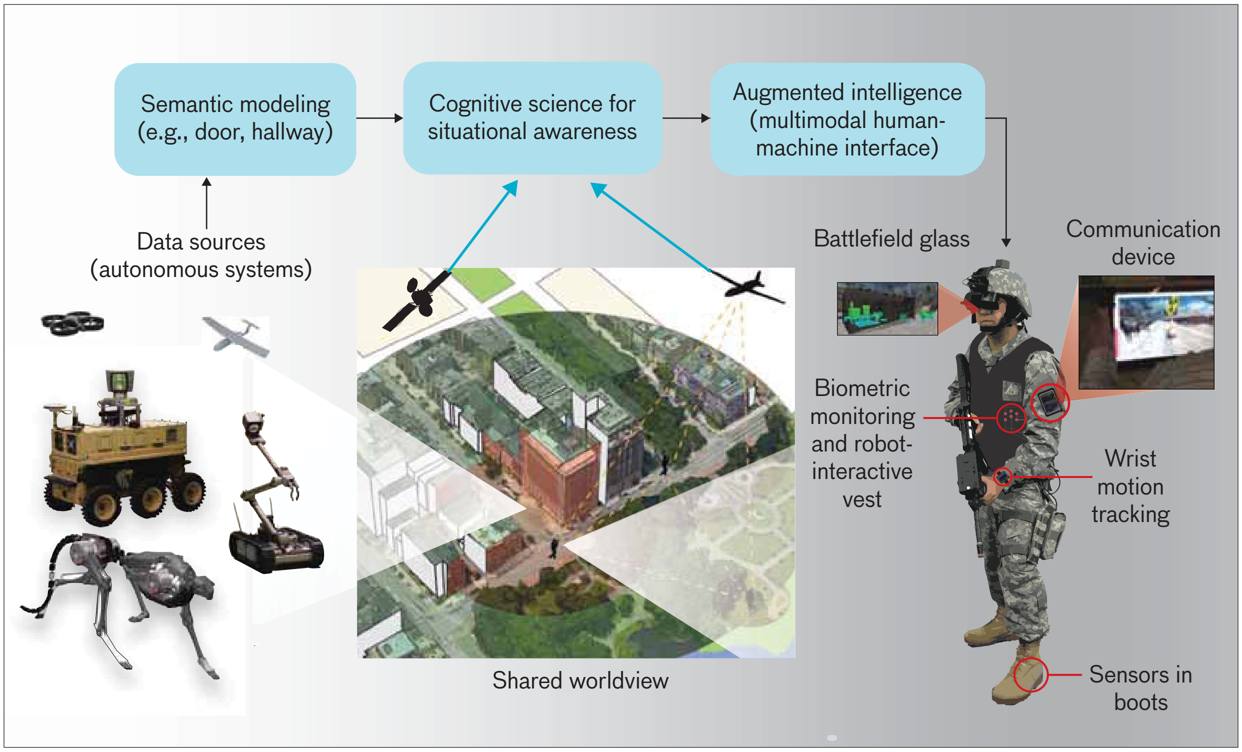

The future soldier, outfitted with technologies such as a heads-up display lens and a biometric vest that monitors vital signs, receives actionable intelligence based on data from autonomous systems, including unmanned ground vehicles equipped with specialized sensors and teams of unmanned aerial vehicles actively mapping an area throughout a mission.

external contexts (mission objective and environmental factors, including weather and terrain) and internal contexts (e.g., human cognitive load, health, skill level)

- Assess the effectiveness of different augmented reality devices—technologies that enhance human perception of and interaction with the real world through computer-generated sensory input such as video, graphics, audio, or tactile data (e.g., heads-up displays, vibrating vests, earbuds, armbands)
- Evaluate whether the technology can be developed into a prototype and deployed in the field

"The motivation is to think several years into the future for how we ideally would interact with robots," says Donahue. As illustrated in the figure above, the larger vision of the program is to acquire data from multiple autonomous systems and transform those data into actionable intelligence that can provide situational awareness among the squad for effective mission execution.

To function in the real world, autonomous systems must interact with people and physical elements in their environment. Because of the complexity of these interactions, robotics researchers often conduct initial research by investigating a constrained scenario in a theater setting. This test bed helps to inform the design and development of autonomous systems. For the collaborative robotics program, the scenario is a raid on a compound conducted by a six-person dismounted squad aided by zero to many unmanned aerial and ground vehicles. The Laboratory team used the open-source Gazebo simulator to create a virtual reality in which this scenario could be acted out. Gazebo supports the simulation of any desired number of robots in complex 3D indoor and outdoor environments and provides realistic motion, sensor noise, and ground-truth data for object locations that are useful for benchmarking algo-

The Lincoln Laboratory Interactive Virtual Environment (L-Live) superimposes the virtual world on three sidewalls in the 3D infrared tracking theater (large photo) for a 270° displayable area (inset).

rithms. Gazebo is integrated with the Robot Operating System (ROS), a framework of software libraries for developing robot applications.

Simulations are conducted in Lincoln Laboratory's Autonomous Systems Laboratory 3D infrared (IR) tracking theater, which enables real-time human interaction with the virtual world. In this Lincoln Laboratory Interactive Virtual Environment, or L-Live as it is called, the team can play out different vignettes of their scenario. The room's ceiling-mounted projectors and several wall-mounted cameras track markers placed on moving and static people and objects; these markers enable the team to determine the position and orientation of people and objects to within a millimeter. The motion-capture area overlaps a region in the virtual world, which is projected on the room's walls over a 270° displayable space for 1:1 scaling.

Events that are tracked in motion capture are re-represented in the virtual world. In-room hardware such as the Turtlebot, an autonomous platform for developing robot applications, is simulated in the virtual world. (Note: The Turtlebot is a research-only surrogate for future combat robots.) A "player" drives the demonstration of the scenario, communicating with L-Live, a tablet (the current augmented reality display), and Turtlebots.

L-Live is proving to be effective in meeting the research needs for this program. To date, two novel capabilities have been developed utilizing commercial off-the-shelf components:

- IR depth sensor. An IR camera was added to the motion-sensing ASUS Xtion (similar to the Microsoft Xbox Kinect), which features a color (red, green, and blue [RGB]) camera and depth sensor. The IR data stream is fused with depth data from the Xtion.
- RGB- and IR-colored OctoMaps. OctoMap, software that generates 3D models of environments by recursively partitioning 3D space into eight equal pieces, was adapted to extend mapping capabilities beyond occupancy grid mapping, which only models occupied areas and free space. Algorithms for assigning color were developed to create RGB- and IR-colored OctoMaps that are based on RGB and IR depth data.

For the data produced by autonomous systems to be truly useful to the warfighter, they need to be presented in human terms. Semantic modeling requires that

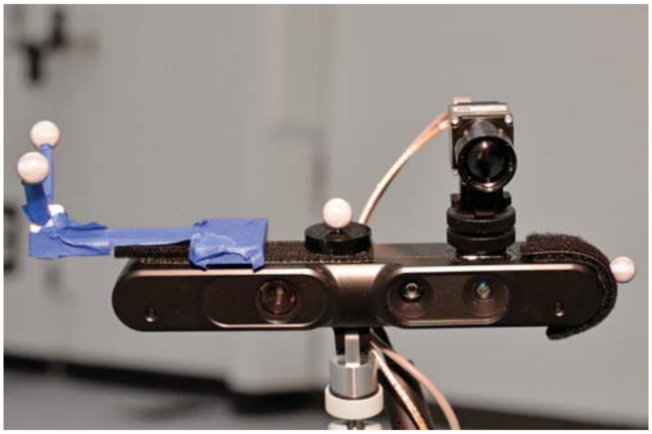

The IR depth sensor (left) can be mounted on the Turtlebot or other mobile robot platforms to generate real-time IR depth streaming data of an environment. Based on these data, a 3D thermal point cloud can then be projected onto an augmented reality display in the user's perspective (below).

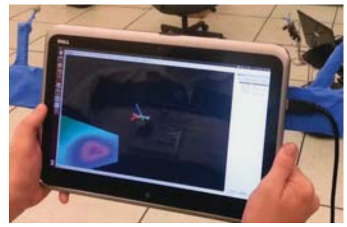

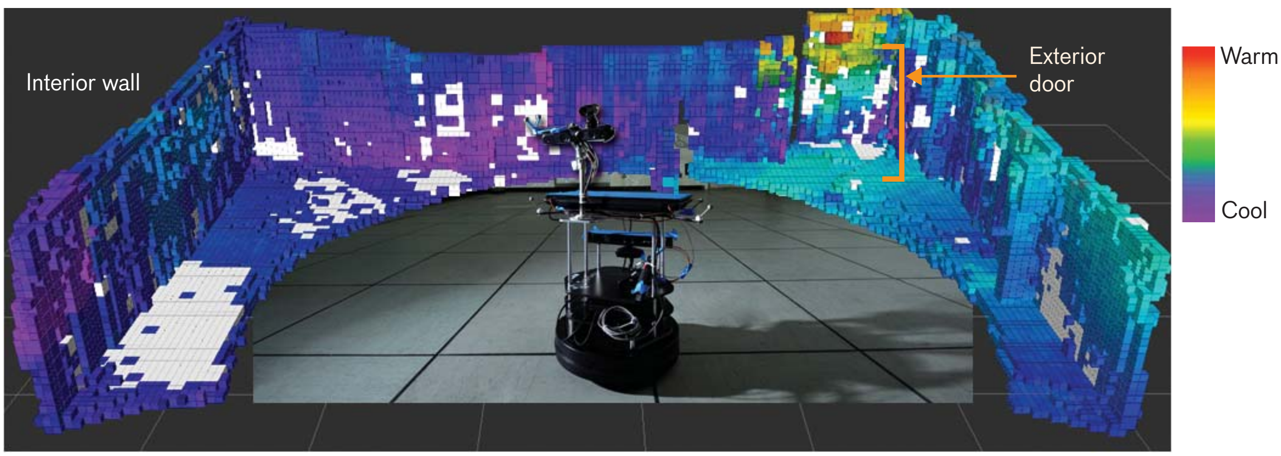

This IR-colored map of the L-Live area was created with data from the novel Lincoln Laboratory–developed IR depth sensor, mounted on a Turtlebot, and the Lincoln Laboratory–adapted OctoMap software. Note the difference in temperature between the interior wall and the exterior door as indicated by their respective colors on the temperature gradient.

robots understand the space in which they operate such that they can identify and classify features of the environment. Within the mission context, these classifications must distinguish friend from foe and asset from threat. Lincoln Laboratory researchers are working with military liaisons, who are supplying ground-truth data, to develop machine-learning algorithms for identifying mission-relevant scene objects and for conveying this information to the warfighter. So far, the focus has been on generating data to create a relevant world model. The researchers are defining data structures on the basis of what they think is relevant to the mission, given what is known about operations. The machine-learning algorithms can then be tested and redefined if need be.

Once the model is built, the next question is how the information should be presented to the soldier. Determining which type of display mode—visual, audio, or tactile—to use involves several factors: The time of day, weather, and noise levels in the environment, and the operator's cognitive load, stress levels, health, skills, and preferences must all be considered within the mission context. If a soldier is running, a tactile warning sent via a vibrating vest or pressure wristband may be ineffective because the soldier may not feel the signal. If the environment is noisy, an auditory message conveyed through earbuds is not the best interface. For situations in which a squad member's attention is needed elsewhere, perhaps a subtle visual icon rather than a startling audio alarm is better. Working backwards from these kinds of considerations enables the team to incorporate key features in the machine-learning algorithms. In the future, real-time commands and feedback from operators as well as physiological sensor data (e.g., heart rate, perspiration) could be fed back into the information flow system to make the models more robust.

The team is also researching improved algorithms to present streaming sensor data from autonomous systems in the human's frame of reference. Known as viewpoint transformation, this technique requires that the position and orientation (i.e., pose) of objects and people are tracked over time. For data to be properly rendered on a visual display, pose tracking and viewpoint transformation must be computed in real time (10–100 ms for a human interface) and to high precision. While the team plans to explore auditory and tactile interfaces, their current focus is on visual augmented reality devices that superimpose computer-generated graphics and video, based on sensor data from autonomous systems, on the physical world (see figure below). "We [the team] chose to focus on vision because it is the highest-bandwidth, most natural method of human absorption of data, and visualization is a popular method of human-computer interaction," says Evan Krause, a former Laboratory researcher who had been investigating how to display augmented intelligence in the human's frame of reference. For now, a tablet interface is being used; eventually, it will be replaced by a device more suited to field use, such as BAE Systems' Q-Warrior. The team is also looking into Osterhout Design Group's Smart Glasses, a 3D stereoscopic, see-through, high-definition display.

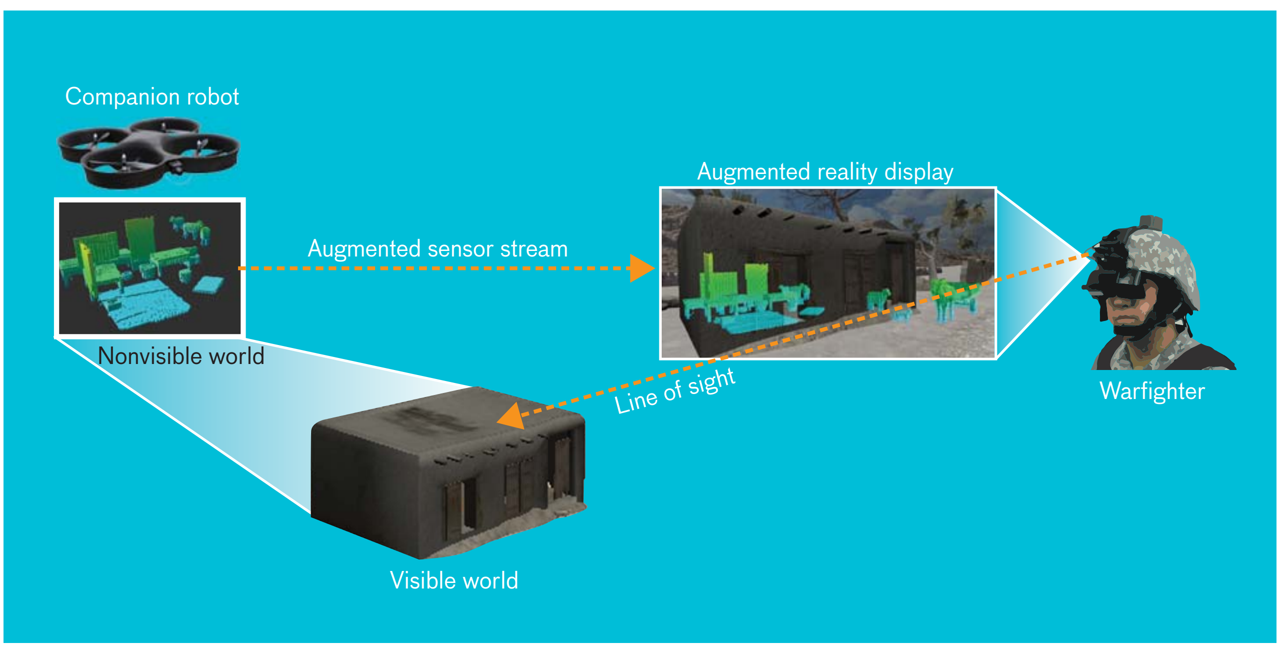

The warfighter receives intelligence about the nonvisible world (i.e., the world beyond his line of sight) to augment his perception of the visible world. In this example, sensor data from an unmanned aerial vehicle are used to generate a 3D OctoMap of a hut's interior in real time. These data are then rendered in the warfighter's frame of reference, giving the warfighter the ability to "see through" walls.

In L-Live, Evan Krause wears BAE Systems' Q-Warrior, a helmet-mounted transparent display designed for soldiers. The display shows waypoints, tracks friendlies and foes, and links to sensor feeds from autonomous systems. An eye-motion box enables users to make relatively large movements with their helmets while maintaining view of their displays.

Enabling visual augmented reality involves two key challenges: streaming data and pose tracking in real time. To address the first challenge, the team is using a system-level approach to display only relevant information and intelligent data structures, such as OctoMaps, to represent the environment. The extensions that are currently supported by OctoMap—visible and IR color—are useful in segmenting maps and determining scene meaning. The software will be further extended with two capabilities that together will help users to assess map staleness: a visual indexing feature for assigning semantic labels to voxels (3D pixels) will show not only if a space is occupied but also what is occupying that space; a time-tagging function will provide better motion detection between scenes.

The team is working to solve the pose-tracking challenge by leveraging the motion-capture capabilities of L-Live and researching various approaches to enable tracking in the wild. One option is to use collaborative multiagent simultaneous localization and mapping (SLAM), an algorithm-based technique in which mobile robots and soldiers fitted with sensors build their own maps of an unknown environment from a sequence of landmark measurements while navigating through that environment and localizing within that map. These individual maps can then be combined to produce a global map. Another option is mobile motion capture. The team is looking into Project Tango, a smartphone equipped with sensors that enable the device to track its position and orientation in real time and to make more than a quarter-million 3D measurements every second; these measurements are then combined to produce a single 3D map of the surrounding space. For the collaborative robotics program, a robot equipped with Project Tango–like sensors could track, relative to itself, the pose of fiducial markers placed on a soldier.

Because the pose-tracking research is in its initial stages, it is not yet clear which approach is best. "The SLAM approach seems like a better long-term solution, but whether it is feasible is unknown at this point," explains Krause. A notoriously difficult problem in the robotics community, SLAM has been traditionally researched for single agents in constrained environments (e.g., indoor laboratories). The problem has not been solved for many practical settings (e.g., outdoors under various weather and lighting conditions). SLAM suffers from high computational and memory requirements, especially in unconstrained open-world scenarios in which there are hundreds of landmarks.

Future work will focus on quantitative and qualitative evaluations of human task performance with and without augmented intelligence provided by autonomous systems via augmented reality displays. Quantitative parameters to be evaluated include time to complete tasks, effectiveness in responding to environmental threats (e.g., avoiding an improvised explosive device, detecting a sniper), and biometrics (e.g., heart rate). Participants' opinions on the usefulness of augmented intelligence and its usability in the field will be recorded. The results from this analysis will help determine if augmented intelligence is an effective, intuitive form of human-machine interaction—one that can enhance warfighter situational awareness and mission execution without imposing significant physical and mental burdens on the warfighter. Once perfected, the collaborative robotics technology could be applied to any circumstance in which autonomous systems data may be useful to humans. The team envisions the technology eventually being used by first responders to assess situations at emergency scenes, by emergency services professionals to find survivors following natural or man-made disasters, and by security personnel to protect borders or critical infrastructures.

AUTONOMOUS GEOLOCATION

Improving Aerial Searches

An unmanned aircraft system combines ideas from the animal kingdom and nonlinear programming to locate objects

Sensor-equipped, small unmanned aerial vehicles (UAV) are increasingly popular tools for accomplishing tasks that require wide-area surveillance. Their size and agility make UAVs suitable for surveying places that are difficult for large aircraft to navigate. Their autonomy shields humans from encountering risks in dangerous locations. Compared to pilot-operated airplanes, UAVs are inexpensive to buy, operate, and maintain; readily available; and, if necessary, expendable.

Emergency responders are using UAVs to facilitate search and rescue missions. Researchers have employed them to monitor the natural environment and to map terrain. The military has deployed UAVs for intelligence, surveillance, and reconnaissance activities. These users need the capability to efficiently find objects of interest, for example, a person lost in the wild, a topographical feature, or a terrorist's van. Lincoln Laboratory researchers have developed a technique that can increase the speed of such searches by enabling an unmanned aircraft system to exploit its own sensors to make decisions about where to focus the search.

"We wanted to develop a system that could find a known target in an unknown location within a specific region and find it as fast as possible," says Michael Park, a former technical staff member in the Embedded and Open Systems Group and one of the principal developers of the technique.

> The Laboratory's technique is fully autonomous, takes advantage of the maneuverability of a very small UAV, and provides humans with a good starting point for a more refined search.

The typical process for acquiring the data needed to locate a person, place, or thing requires multiple steps: the UAV is sent out to gather data along a route that operators estimate will yield a discovery; the onboard sensor then streams those data back to a human analyst; after processing those data, the analyst determines if the target was found. If the UAV fails to find its quarry, it is directed to another route and the process repeats until success is achieved or the search is halted.

Lincoln Laboratory demonstrated a two-stage method that takes the human analyst's steps out of the process. In the

The commercial quadrotor (3D Robotics' ArduCopter) was modified to include a commercial receiver and Lincoln Laboratory–developed software.

first stage, the UAV flies what is known as a Lévy flight to survey the region of interest. A Lévy flight, also called a Lévy walk, is a random pattern of movements. Envision a foraging creature, say a shark, honeybee, or deer, seeking food. The animal explores a region, making numerous small, random movements and, if food is not discovered, making a long move to a new place where it will refine its search with small moves. Much like the shark or bee, the UAV used in the Lincoln Laboratory trials, a commercially available 3D Robotics' ArduCopter with an integrated, onboard processor, flew a random pattern of short hops and longer runs to locate a radio-frequency (RF) transmitter in a rural field.

The second stage of the Laboratory's technique is the employment of simplex minimization to home in on the object of interest. The UAV's onboard processor takes data collected during the Lévy flight and identifies three data points that carry the minimum amount of information necessary to indicate the presence of the object (in the Laboratory's demonstrations the information is an RF signal stronger than a predetermined threshold). Using the values from three chosen points, the UAV's processor estimates a fourth point at which the UAV should encounter a value—for example, an RF signal—greater than the least strong value. The UAV then flies to that fourth point, and if the data there indicate a greater likelihood of the object's presence than do the data at any of the first three chosen points, that fourth point replaces the least conclusive point of the original triad, creating a new search area. Through successive rounds of flight and triangulation of points that increasingly meet thresholds for pertinent data, the UAV finds the object.

"The simplex is just a triangle," says Charles Coldwell, a technical staff member in the RF Technology Group who worked on the firmware and communications protocol for the system. "The RSSI [received signal strength indication] is measured on every vertex. For each simplex, the vertex with the lowest RSSI is moved by reflecting it across the line joining the other two RSSI vertices, and the UAV is flown to the new vertex and measures the RSSI there." If the RSSI at the new vertex is the highest of the data points, then the vertex is moved even farther in the same direction. If the RSSI at the new vertex is lower than the RSSI at any of the three points, the vertex is contracted halfway back the distance it had been moved out.

The Laboratory's technique was demonstrated in 20 fully autonomous test flights over Davis Field, a radio-controlled-aircraft flying site in Sudbury, Massachusetts. Unlike typical Lévy walks, which cover unbounded territory, the flights undertaken by the ArduCopter, flying at an altitude of about 25 m, were constrained to an area with a radius of 90 m from a ground designation. The ArduCopter was modified to include a commercial signal-strength-indication receiver and peripherals for communication

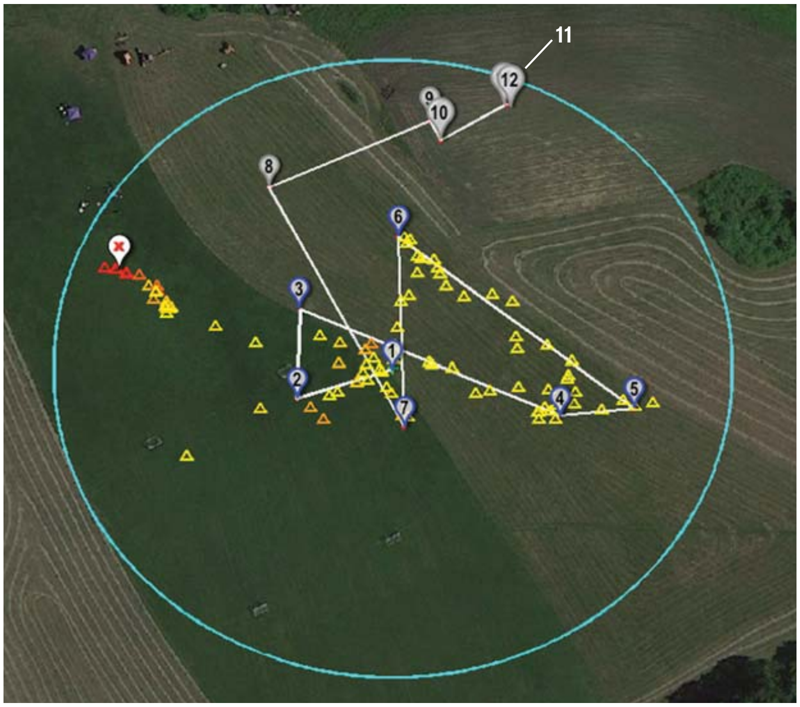

The plot shows the results of the ArduCopter's Lévy flight for locating a transmitter. Flight begins in the center of the circle at marker 1. The yellow and orange triangles denote the actual GPS flight path that the ArduCopter reported. It was doing its best to follow the instructions generated by the algorithm (indicated by the white line path). Wind and other factors pushed the copter off its course at times. The color gradient goes from yellow (distant from the target, the red X in the white marker) through orange (nearer to target) all the way to red (at target). Therefore, as the ArduCopter got closer to the target, the triangles turned from yellow to red. The numbers indicate the sequence of random-flight waypoints the ArduCopter was instructed to fly before switching over to the simplex routine. The copter only flew to points 1–7 (blue markers); at waypoint 7, the copter breaks from its Lévy routine because it found three points above the threshold value and then, using the simplex routine, moved toward the target. Waypoints 8–12 (gray markers) were thus ignored.

and control capabilities. A laptop was used to provide ground control for the ArduCopter; to ingest Global Positioning System (GPS), altitude, and signal-strength data; and to generate the simplex minimization points and Lévy flight instructions. The laptop's Internet connection allowed the system to query and exploit U.S. Geological Survey data, thereby making the system usable anywhere there is GPS coverage and an Internet connection.

In 17 of the 20 flights, the system successfully located the transmitter and autonomously landed the ArduCopter within 10 m of this target. While a manually driven search performed by a UAV with GPS geolocation capability may locate a target within a range of 4 m if the human operator has visual sight of the target, the Laboratory's technique is fully autonomous, takes advantage of the maneuverability of a very small UAV, and provides humans with a good starting point for a more refined search. The flight times to locate and land the UAV ranged from approximately 6 to 8 minutes, a timespan well within the flight time achievable by a battery-powered UAV. In the three flights that failed to find the transmitter, the ArduCopter's battery could not supply power long enough to respond to the RSSI measurements.

An autonomous UAV Lévy search algorithm with simplex minimization has not been documented in the current literature. The Laboratory's technique could open a new research direction for engineers seeking to advance the search capabilities of unmanned systems. The sensor input for the system's search is not limited to RF and can be extended to object feature detection with a camera or lidar. In addition, multiple platforms may be deployed in parallel in a "divide-and-conquer" search method. Future research could also look at applying the technique to larger, more powerful UAVs to enable faster and longer-range searches than those executed by the ArduCopter.

"The search algorithm can be applied to tasks for which humans need systems to autonomously locate an object with statistically minimal time and energy," says Park. "I can imagine it being best used in open regions, such as the desert or ocean, where a human operator could use assistance locating a missing comrade or an adversary's small UAV in the vicinity."

ROBOTICS

Quick-Reaction Autonomous Driving

Students apply advanced autonomy algorithms to race miniature self-driving cars around MIT

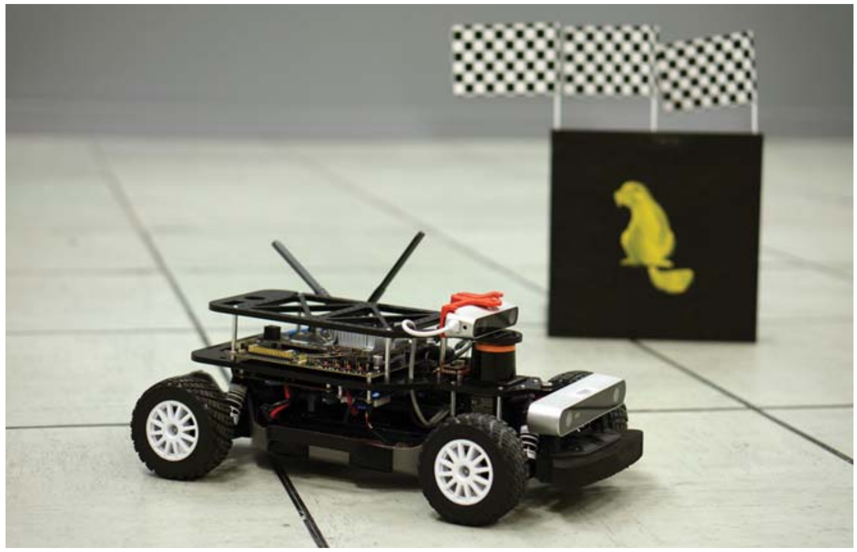

The vehicles used in the RACECAR IAP course sported an array of sensors and an embedded microcomputer. For the Beaver Works Summer Institute, the vehicles navigated a track constructed from stands like that above decorated with the yellow Beaver Works logo and checkered flags.

For more than a century, students, faculty, and staff have used the underground tunnels and basement hallways of MIT to escape the weather as they hustle between classrooms and laboratories. But for a few weeks out of the year, students jogging to class are easily outpaced by small robotic vehicles zipping through the tunnel network. The miniature cars employ student-developed algorithms to accelerate, brake, and nimbly turn to autonomously navigate and avoid obstacles at speeds topping 10 mph in a timed race. The self-driving vehicles, called RACECARs for Rapid Autonomous Complex-Environment Competing Ackermann-steering Robots, are platforms for a series of robotics education workshops and classes at MIT. Student teams design, code, and test autonomy algorithms before pitting their cars against one another to be the fastest robot to successfully complete a circuit in the basement tunnels beneath MIT's Stata Center.

First introduced as a three-week activity during the 2015 MIT Independent Activities Period (IAP), a special January term with course offerings distinguished by variety, innovative spirit, and fusion of fun and learning, RACECAR has since become the cornerstone of a full-term undergraduate course and a summer workshop for talented high school students. The classes focus specifically on robotics software engineering and advanced algorithms. In contrast with build-centric activities that start with a kit of parts, students are provided with a complete robot with integrated sensing, computation, and a basic teleoperation software stack. The students are tasked with developing the artificial intelligence needed to perceive and react to the environment.

RACECAR was proposed by Michael Boulet, assistant leader of Lincoln Laboratory's Control and Autonomous Systems Engineering Group, and Prof. Sertac Karaman of MIT's Department of Aeronautics and Astronautics (AeroAstro) as part of the Beaver Works collaboration between the MIT School of Engineering and Lincoln Laboratory. "Thanks to the Beaver Works collaboration, we have gotten a chance to help develop the RACECAR platform and tailor it to our needs," says Karaman. "AeroAstro has made it a priority to lead in the autonomy and embedded systems area, and the platform has been invaluable to advancing our research, particularly in vision-based algorithms for agile autonomous navigation; for teaching major classes for undergraduates; and for building a community focused

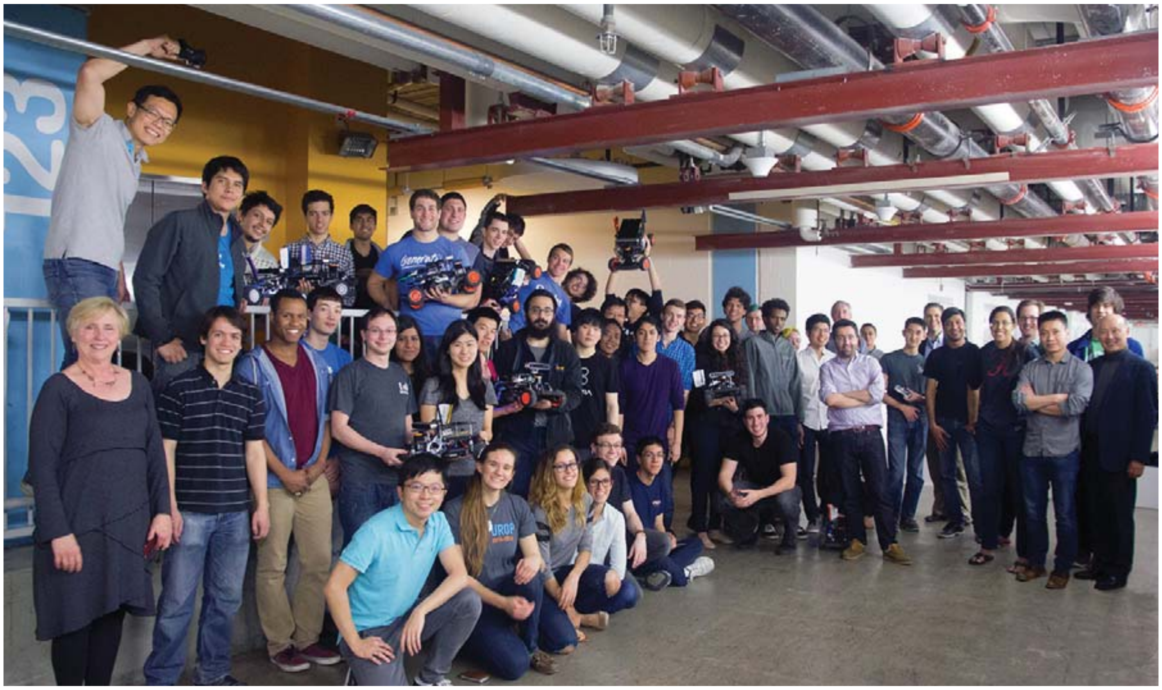

Instructors and students from Robotics Science and Systems (MIT courses 6.141 and 16.405) are shown with their autonomous RACECAR systems following the finale of the class, a car race through the tunnels of MIT's Stata Building.

on autonomous systems and embedded systems."

"RACECAR challenges students to solve problems similar to those faced by professional roboticists developing self-driving cars and autonomous systems for national defense," says Boulet. "For example, students must calibrate their motion model in order to accurately predict the miniature car's future path. Lincoln Laboratory researchers performed analogous tasks to develop a precision autopilot for a nine-ton, seven-meter-long vehicle." (See the article "Automation of Armored Four-Wheel-Counter-Steer Vehicles" later in this issue.) Students also need to contend with imperfect sensor measurements of the environment, complex vehicle dynamics that arise at speeds above 3 meters per second, and finite computational capacity. As in professional practice, there is rarely an obvious best answer to these challenges.

Inspired by robotics competitions sponsored by the Defense Advanced Research Projects Agency (DARPA), the RACECAR activity is a highly parallel approach to problem solving. "There is value in many teams pursuing a diverse array of possible solutions to difficult problems," says Boulet. "Some teams' algorithms work well early in the development cycle, others improve as the class progresses, and a couple might never succeed." Course organizers recognize every team's contribution to searching the solution space, regardless of the final standings, by asking each team to share what in their approach worked well and how their approach could be improved.

Real-time onboard processing is a key element of the challenge. Student algorithms only have fractions of a second to process sensor data, make decisions, and command vehicle speed and steering to avoid oncoming obstacles. "Robots of the recent past might take 10s of seconds to create models of the environment and search through 1000s of possible actions before making a decision. That timeframe doesn't work for RACECAR," says Michael Park, a former technical staff member at the Laboratory who helped direct the IAP workshop. Advances in computation, such as the teraflop-class embedded computer integrated into the vehicle, enable processing of increasingly sophisticated autonomy algorithms. However, the students must still consider the computational efficiency of their approach and code.

RACECAR challenges students to solve problems similar to those faced by professional roboticists.

"Another goal for the course is exploring how robotics development can be achieved on a short timescale," adds Park. Handing students an assembled robotic car and core software infrastructure, along with a virtual vehicle in a simulated world, allows them to focus on the autonomy immediately. The robot is constructed from an off-the-shelf 1:10-scale radio-controlled (RC) electric model rally car, complete with car-like Ackermann steering geometry, four-wheel-drive transmission, and suspension system. Course organizers remove the RC components and wire a high-performance embedded computer to the throttle and steering motors. For perceiving their motion and the local environment, the racecars are outfitted with a rich sensor suite, including a scanning laser range-finder, stereo camera, inertial measurement unit, and visual odometer.

The Robot Operating System (ROS), which is a collection of open-source drivers, algorithms, tools, and libraries widely used by researchers and industry, serves as the framework for the software development. Students combine their custom-designed algorithms with existing ROS software modules to rapidly configure a complete autonomous system. According to Owen Guldner, a technical staff member in Lincoln Laboratory's Rapid Prototyping Group and RACECAR instructor, the students typically learn the core elements of ROS quickly and are able to leverage it to create successful algorithms.

"In a short time, we had to come up with a working device, capable of racing confidently through MIT tunnels," says Valerio Varricchio, an MIT doctoral degree candidate in AeroAstro, who found the timeframe one of the most challenging aspects of the 2015 IAP course. "Even though we were potentially armed with state-of-the-art localization and planning algorithms, and had several different onboard sensors available, we soon realized that a simpler, reactive approach based on lidar readings was the way to go." In fact, most teams turn to reactive planning approaches that used the car's laser ranger to track its position relative to the tunnel walls.

The three-week turnaround for developing and implementing the prototype algorithms is enabled by what Karaman calls the "inverted lab" structure of the workshop. Instead of being consigned to working on their systems just during class periods at the Beaver Works center near the MIT campus, the students are furnished with their own cars on which they can experiment outside scheduled "lab" time. The teams take advantage of the inverted lab to test their algorithms and to hold informal races against each other. "We're careful not to constrain the students. The ingenuity embedded in some of the solutions has surprised even us, the instructors," says Karaman.

In addition to algorithm building sessions, or "hackathons," the RACECAR activity includes seven lectures that cover instruction on ROS; explanations of sensing, perception, control and planning algorithms; and discussions of case studies in autonomous systems.

Participants have rated RACECAR a great learning experience, as well as fun. Varricchio says the hands-on work was a welcome change from dealing with theory and exams. He predicts career opportunities presented by the experience with an embedded computer and robotics software: "I see an incomparably exciting future for the field of autonomous vehicles and autonomy in general." John Alora, a master's degree student in AeroAstro at MIT and a military fellow at Draper Laboratory, found that the research into autonomy was interesting and has "wide applications, especially in the military."

The staff from Lincoln Laboratory also appreciate the experience. Guldner says he found out "how much you have to learn just to be able to teach it to someone else—you never know what questions will be asked." For Park, the class reaffirmed the importance of clear communication and iterative trials in a rapid development process. He also observes that the Laboratory gained a number of benefits from RACECAR: "We had the opportunity to train staff, increase the Laboratory's exposure to MIT students, foster and develop a relationship with a faculty member, and generally promote the Laboratory's capabilities in autonomous systems."

Boulet admits that he initially had doubts that the students could achieve the autonomous navigation necessary for the racecars in the short time span allotted for the workshop. But, in the first RACECAR IAP activity, three of the four cars completed the 515-foot tunnel course without mishap. The winning entry covered the race route in less than 50 seconds at an average speed of more than 7 mph, a speed surpassed by 3 mph in 2016. According to Boulet, "The students are out of breath trying to keep up with their systems, an unusual occurrence as typical small autonomous vehicles are slow movers."

As autonomy algorithms become even more capable, Boulet and Karaman hope to increase the level of difficulty of MIT's RACECAR courses. Ideas include setting up a longer and more complex tunnel racetrack, introducing moving obstacles, pitting the cars against each other for a true race, and performing the same race with small unmanned multirotor aerial vehicles. Course organizers are also working toward offering Lincoln Laboratory staff a similar course during which the robots would navigate the Laboratory's hallway network.

In 2016, a RACECAR activity geared to talented high school students, primarily those about to enter their senior year, was offered at MIT as the first program of the Beaver Works Summer Institute (BWSI). The BWSI initiative was

At the Mini Grand Prix, a high school student maneuvers his team's RACECAR to the start line in preparation for its run through the championship course. The crowd at the Mini Grand Prix included not only the Beaver Works Summer Institute staff and students but also the families and friends of the participants.

conceived by Robert Shin, director of Beaver Works and head of Lincoln Laboratory's ISR and Tactical Systems Division, as a mechanism for providing intensive science, technology, engineering, and mathematics (STEM) experiences for high-achieving students. Karaman, Shin, and Kenneth Gregson, a member of the technical staff at Lincoln Laboratory, led instructors from MIT, Lincoln Laboratory, NASA's Jet Propulsion Laboratory, and Continental Corporation; a team of student counselors; and Beaver Works personnel in providing the four-week, residential RACECAR program for 46 students from across the country. The students divided into teams to program their RACECARs to navigate a circuitous race route. The finale of the program was the Mini Grand Prix in which nine RACECARs competed to be the fastest to successfully complete a course in MIT's Walker Memorial building.

"I believe we have taught and inspired several young roboticists in a way that is stronger than any other high school STEM program can possibly imagine to do," says Karaman, a claim that is echoed in the thank-you letters sent to the BWSI team. Anusha Datar from Burlington High School in Massachusetts called RACECAR "the most unique, educational, and incredible experience" and George Jeno of the Illinois Mathematics and Science Academy wrote, "I thank you for giving me the opportunity to be part of the future."

Equally impressed with the first BWSI program were the instructors from the Jet Propulsion Laboratory and Continental Corporation. Having seen firsthand the skills the students acquired in just four weeks, these instructors have expressed a strong interest in helping their organizations develop RACECAR educational outreach programs. The BWSI team is already working with participating high schools to create "satellite" RACECAR projects, and the BWSI team is investigating the feasibility of hosting an international challenge race next summer at MIT.

From the graduate students in the IAP course, to the undergraduates in the MIT robotics classes, down to the high school seniors of BWSI, the RACECAR students have acquired knowledge and skills that they may apply to their future research or careers and that may ultimately have impact on pressing global problems. Perhaps it will be RACECAR participants who go on to improve advanced driver-assistance technology that could help save thousands of lives on the roadways and decrease urban traffic congestion.