NETWORKS OF INTELLIGENT ROBOTS WILL SOMEDAY TRANSFORM WARFARE-BUT SIGNIFICANT HURDLES REMAIN

BY LORA G.WEISS

Two small planes fly low over a village, methodically scanning the streets below. Within minutes, they spot their target near the edge of town. With no way to navigate through the streets, they radio for help. Soon after, a metallic blue SUV begins moving cautiously but purposefully along the dirt roads leading to town, seeking out the target's GPS coordinates. Meanwhile, the planes continue to circle overhead, gathering updated information about the target and its surroundings. In less than half an hour after the planes take to the sky, the SUV.

Last fall, my research team fielded these vehicles at Fort Benning, Georgia, during the U.S. Army's Robotics Rodeo. That's right, the two quarter-scale Piper Cub aircraft and the Porsche Cayenne operated without any humans at the controls. Instead, each robot had an onboard computer running collaborative software that transformed the three machines into an autonomous, interoperable system.

The demonstration may sound simple—the target was just a tarp staked to the ground—but had this been the streets of Kabul or Baghdad, where any pile of debris can conceal a deadly improvised explosive device, such autonomous tracking robots in the future could help keep soldiers out of harm's way. Indeed, military leaders have increasingly embraced the use of unmanned aerial vehicles (UAVs) and other robotic systems over the past decade, to handle the "three D's": the dull, dirty, and dangerous tasks of war. Back in 2000, the U.S. Department of Defense (DOD) had fewer than 50 UAVs in its inventor y; by early 2 010, it had more than 7000. In 2009, the U.S. Air Force started training more pilots to operate unmanned systems than to fly fighters and bombers. And according to market research firm ABI Research, 65 countries now use military robots or are in the process of acquiring them.

The ranks of battlefield robots will only grow: The U.S. Congress has mandated that by the year 2015, one-third of ground combat vehicles will be unmanned, and the DOD is now developing a multitude of unmanned systems that it intends to rapidly field. Meanwhile, thousands of robotics researchers worldwide are making impressive gains in networking robots and boosting the sophistication and autonomy of these systems.

Despite the advances in both their performance and safety, these robots are still far from perfect, and they routinely operate in situations for which they may not have been designed and in which their responses cannot always be anticipated. Some of the DOD's most advanced UAVs carry dozens of sensors, including high-resolution night-vision cameras, 3-D imagers, and acoustic arrays. Yet most cannot distinguish a sleeping dog from a bush, even at high noon. Humans are still needed to operate the vehicles, interpret the data, and coordinate tasks among multiple systems. If we are ever to see fully autonomous robots enter the battlefield—those capable of planning and carrying out missions and learning from their experiences—several key techno logical advances are needed, including improved sensing, more agile testing, and seamless interoperability. Even then, a basic question will remain: How can we equip these robots to make critical decisions on their own?

A reporter on the phone asks me what will happen when robots become so smart that they can clone themselves. He seems to assume it's a given: Robots will someday be

agile enough to create exact copies of their mechanical bodies and of the software code comprising their “brains.” All he wants to know from me is when—not if—this great day will arrive. I suggest that he not hold his breath.

As a researcher at the Georgia Tech Research Institute and a board member of the world’s largest association for unmanned systems—the Association for Unmanned Vehicle Systems International—I’ve been working with robots for more than two decades, starting with underwater vehicles, then moving to air and ground vehicles, and most recently addressing collaborations among robots like those we demonstrated at the Robotics Rodeo. I can attest that while robots are definitely getting smarter, it is no easy task to make them so smart that they need no adult supervision. And call me a skeptic, but I doubt they’ll be cloning themselves anytime soon.

That said, I’m amazed at the pace of progress in the field. With thousands of researchers now engaged in advancing the intelligence and autonomy of unmanned systems, new breakthroughs are announced seemingly every week. Both the variety and the number of unmanned systems now deployed are breathtaking. UAVs run the gamut from the 1-metric-ton MQ-1 Predator drone made by General Atomics to AeroVironment’s tiny 430-gram Wasp micro air vehicle. There are unmanned ground vehicles that roll on treads like tanks, walk like dogs, and slither like snakes. Unmanned maritime vehicles include submarine-like vessels that can cruise underwater for kilometers and boat-like craft that patrol for pirates, smugglers, and other criminal types.

But none of these systems are fully autonomous. The RQ-4 Global Hawk UAV, made by Northrop Grumman, is guided by satellite waypoint navigation, yet it still requires a human pilot sitting in a remote ground station, plus others to operate the drone’s sensors and analyze the data being sent back. iRobot’s PackBot tactical robot is teleoperated by means of a video-game-style controller, complete with joystick. Even the driverless vehicles that participated in the Defense Advanced Research Projects Agency’s Grand Challenge competitions in 2004, 2005, and 2007 weren’t entirely autonomous, as the courses they had to negotiate were tightly controlled.

So why haven’t we seen a fully autonomous robot that can sense for itself, decide for itself, and seamlessly interact with people and other machines? Unmanned systems still fall short in three key areas: sensing, testing, and interoperability. Although the most advanced robots these days may gather data from an expansive array of cameras, microphones, and other sensors, they lack the ability to process all that information in real time and then intelligently act on the results. Likewise, testing poses a problem, because there is no accepted way to subject an autonomous system to every conceivable situation it might encounter in the real world. And interoperability becomes an issue when robots of different types must interact; even more difficult is getting manned and unmanned systems to interact.

To appreciate the enormous challenge of robotic sensing, consider this factoid, reported last year in The Economist: “During 2009, American drone aircraft…sent back 24 years’ worth of video footage. New models…will provide ten times as many data streams…and those in 2011 will produce 30 times as many.” It’s statistics such as those that once prompted colleagues of mine to print up lanyards that read “It’s the Sensor, Stupid.”

But a robot is more than just a platform of sensors. Let’s say an unmanned jeep is traveling down a city street. Its cameras may detect a parked car along the curb, an open

manhole in the middle of the road, and a knot of school kids crossing at the intersection. But unless the jeep can correctly classify the car as a car, the manhole as a manhole, and the children as children, it won't have sufficient information to avoid those obstacles.

So the sensing problem in robotics extends well beyond just designing sophisticated new sensors. An autonomous robot needs to be able to automatically process the data from those sensors, extract relevant information from those data, and then make decisions in real time based on that information and on information it has gathered in the past. The goal is to achieve what researchers call situational understanding.

And with no humans in the loop to help interpret the data, reason about the data, and decide how to respond, situational understanding gets even trickier. Using current technology, no robot has all the onboard sensors needed to precisely decipher its environment. What's more, decisions have to be made based on uncertainties and incomplete or conflicting information. If a robo-sentry armed with a semiautomatic rifle detects some- one running from a store, how can it know whether that person has just robbed the store or is simply sprinting to catch a bus? Does it fire its weapon based on what it thinks is happening?

Humans, too, may struggle to read such a situation, but per- haps unsurprisingly, society holds robots to a higher standard and has a lower tolerance for their errors. This bias may create a reluctance to take the leap in designing robots for full autonomy and so may prevent the technology from moving ahead as quickly as it could. It should not take five people to fly one UAV; one soldier should be able to fly five UAVs.

On the other hand, because military robots typically operate in geopolitically sensitive environments, some added caution is certainly warranted. What happens, for example, if a faulty sensor feeds a UAV erroneous data, causing it to cross a border without authorization? What if it mistakenly decides that a "friendly" is a target and then fires on it? If a fully autonomous, unmanned system were to make such a grave mistake, it could compromise the safety of other manned and unmanned systems and exacerbate the political situation.

The Predator UAV, developed in the 1990s, went from concept to deployment in less than 30 months, which is extremely fast by military procurement standards.

Little wonder, then, that the UAV exhibited quite a few kinks upon entering the field. Among other things, it often failed when flying in bad weather, it was troublesome to operate and maintain, and its infrared and daylight cameras had great difficulty discerning targets. But because commanders needed the drone quickly, they were willing to accept these imperfections, with the expectation that future upgrades would iron out the kinks. They didn't have time to wait until the drone had been thoroughly field-tested.

But how do you test a fully autonomous system? With robots that are remotely operated or that navigate via GPS waypoints, the vehicle's actions are known in advance. Should it deviate from its instructions, a human operator can issue an emergency shutdown command.

However, if the vehicle is making its own decisions, its behavior can't be predicted. Nor will it always be clear whether the machine is behaving appropriately and safely. Countless factors can affect the outcome of a given test: the robot's cognitive information processing, external stimuli, variations in the operational environment, hardware and software failures, false stimuli, and any new and unexpected situation a robot might

encounter. New testing methods are therefore needed that provide insight and introspection into why a robot makes the decisions it makes.

Gaining such insight into a machine is akin to performing a functional MRI on a human brain. By watching which areas of the brain experience greater blood flow and neuronal activity in certain situations, neuroscientists gain a better under- standing of how the brain operates. For a robot, the equivalent would be to conduct software simulations to tap the “brain” of the machine. Subjecting the robot to certain conditions, we could then watch what kinds of data its sensors collect, how it processes and analyzes those data, and how it uses the data to arrive at a decision.

Another illuminating form of testing that is often skipped in the rush to deploy today’s military robots involves simply playing with the machines on an experimental “playground.” The playground has well-defined boundaries and safety constraints that allow humans as well as other robots to inter- act with the test robot and observe its behavior. Here, it’s less important to know the details of the sensor data and the exact sequence of decisions that the machine is making; what emerges on the playground is whether or not the robot’s behavior is acceptably safe and appropriate.

Moving to smarter and more autonomous systems will place an even greater burden on human evaluators and their ability to parse the outcomes of all this testing. But they’ll never be able to assess all possible outcomes, because this would involve an infinite number of possibilities. Clearly, we need a new way of testing autonomous systems that is statistically meaningful and also inspires confidence in the results. And of course, for us to feel confident that we understand the machine’s behavior and trust its decision making, such tests will need to be completed before the autonomous robot is deployed.

A swarm of small robots scatters across the floor of an abandoned warehouse. Each tread-wheeled bot, looking like a tiny tank with a mast-like antenna sticking out of its top, investigates the floor space around it using a video camera to identify windows and doors and a laser scanner to measure distances. Employing a technique called SL A M (for “simultaneous localization and mapping”), it creates a map of its surroundings, keeping track of its own position within the map. When it meets up with another robot, the two exchange maps and then head off to explore uncharted territory, eventually creating a detailed map of the entire floor.

These ingenious mapping robots, designed by researchers through the U.S. Army–funded Micro Autonomous Systems and Technology program, represent the cutting edge of robot autonomy. In future iterations, their designers plan to equip the machines with wall-penetrating radar and infrared sensors, as well as a flexible “whisker” to sense proximity to obstacles. Clever as they are, though, these robots lack a key capability that all future robots will need: They cannot easily interact with other kinds of robots.

Now consider the U.S. Navy’s Littoral Combat Ship. Rather than having a fixed architecture, it will have swappable “mission modules” that include vertical takeoff unmanned aerial vehicles, unmanned underwater vehicles, and unmanned surface vehicles. All these robotic systems will have to operate in concert with each other as well as with manned systems, to support intelligence, surveillance, and reconnaissance missions, oceanographic surveys, mine war- fare, port security, and so on.

Achieving this interoperability will be no small feat. While significant progress has been made on automating a single robot as well as a team of identical robots, we are not yet at the point where an unmanned system built for the Army by one contractor can

seamlessly interact with another robotic system built for the Navy by another contractor. Lack of interoperability isn't exclusively a robotics problem, of course. For decades, developers of military systems of all kinds have tried and often failed to standardize their designs to allow machines of different pedigrees to exchange data. But as different branches of the military continue to add to the ranks of their battlefield robots, the enormous challenge of interoperability among these disparate systems only grows.

A particular difficulty is that most automation and control approaches, especially those used for collaborating, assume that all the unmanned systems have the same level of autonomy and the same software architecture. In practice, that is almost never the case, unless the robots have been designed from scratch to work together. Clearly, new approaches are needed so that you can introduce an unknown, autonomous system without having to reconfigure the entire suite of robots.

Interoperability between manned and unmanned systems is even more challenging. The ultimate goal is to have autonomous systems collaborate with humans as equal partners on a team, instead of simply following commands issued by their operators. For that to happen, though, the robots will need to understand human language and intent, and they will need to learn to communicate in a way that is natural for humans.

Interoperability also requires standards, procedures, and architectures that enable effective integration. Today, for instance, unmanned ground and maritime systems use a messaging standard called the Joint Architecture for Unmanned Systems (JAUS). The messaging standard for unmanned air systems, meanwhile, is STANAG-4586, a NATO-mandated format. Within their respective domains, both of these serve their purpose.

But when a UAV needs to communicate with an unmanned ground vehicle, should it use JAUS or STANAG-4586 or some- thing else entirely? The most promising effort in this arena is the JAUS Tool Set, an open, standards-based unmanned vehicle messaging suite that is in beta testing. Using the tool set seems to improve interactions among unmanned vehicles. In the future, the tool set should allow the two message formats to be merged. Ultimately, that should accelerate the deployment of compatible and interoperable unmanned systems.

As robotic systems become more autonomous, they will also need the ability to consider the advice, guidance, and opinions of human users. That is, humans won't be dictating behavior or issuing hard directives, but they should still be able to influence the robot 's planning and decision making. Integrating such information, including its vagaries, nuances, and uncertain- ties, will be a challenge for any autonomous system as its intelligence increases. But attaining these capabilities is within our reach. Of that, I am not skeptical.

SIDEBAR: Robots in Combat

Books have been written about the feasibility and ethics of weaponizing robots, and it's not my intent to explore that topic in any great detail here. The fact is, weaponized robots—missile-launching unmanned combat air vehicles, rifle-toting unmanned combat ground vehicles, and mine-deploying unmanned combat underwater vehicles—are already a reality.

At present, though, the decision of whether these robots attack is still left to humans. But as robots gain more autonomy, will we or won't we allow them to decide to fire weapons on their own? The U.S. Defense Department continues to mull the issue. In 2007, for instance, it released a report called Unmanned Systems Safety Guide for DOD Acquisition, which includes a section on designing weaponized unmanned systems. It lays out a number of ethical, legal, and technical areas of concern that any designer of armed autonomous robots should be prepared to address. These include the inadvertent firing of weapons, erroneous target discrimination, and the possibility of the enemy taking control of the unmanned system.

John Canning of the Naval Surface Warfare Center Dahlgren Division, in Virginia, has pointed out that deploying weaponized robots but then maintaining a human operator to do the actual firing is costly. He's put forth several concepts of operation that might allow autonomous armed robots to coexist on the battlefield with other manned and unmanned systems. One of Canning's key concepts is to "let machines target other machines." That is, design armed unmanned systems so that they can automatically identify, target, and neutralize or destroy the weapons used by adversaries, but not the people using the weapons.

In those instances when it becomes necessary to target humans, Canning proposes that an armed unmanned systems not be allowed to act autonomously but rather be remotely controlled by humans. The machine, he suggests, should be designed with "dial-a-level" autonomy, so that it can switch among operational modes according to its environment and other circumstances. It would also be equipped with both non-lethal and lethal weapons, the former for convincing the enemy to abandon its arms and the latter for actually destroying those weapons.

Ronald C. Arkin, director of the Mobile Robot Laboratory at Georgia Tech, has been looking at ways to imbue robots with a sense of "ethics" and even an artificial "conscience," so that they adhere to international rules of warfare. That should make it possible, he believes, for autonomous robots to conduct themselves on the battlefield at least as well as humans and probably better.

—L.G.W.