Dr. Asimov's Automatons Take on a Life of their Own

Twenty years after his death, author Isaac Asimov's robot fiction offers a blueprint to our robotic future...and the problems we could face

by Alan S. Brown

THIS PAST April, the University of Miami School of Law held We Robot, the first-ever legal and policy issues conference about robots. The name of the conference, which brought together lawyers, engineers, and technologists, played on the title of the most famous book ever written about robots, *I, Robot*, by Isaac Asimov.

The point was underscored by Laurie Silvers, president of Hollywood Media Corp., which sponsored the event. In 1991, Silvers founded SyFy, a cable channel that specializes in science fiction. Within moments, she too had dropped Asimov's name.

Silvers turned to Asimov for advice before launching SyFy. It was a natural choice. Asimov was one of the greatest science fiction and science popularizers of his generation. As the author of short stories and more than 500 books, he ranks among the most prolific authors ever.

The Three Laws of Robotics defined by Isaac Asimov, shown in a 1965 portrait, are:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.

2. A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.

3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

Greatest Legacy

Yet Asimov believed people would ultimately remember him for his robot stories, which he began writing when he was still a teenager. "He told me that he thought his Three Laws of Robotics would be his greatest legacy," Silvers recalled.

This is no small claim. Asimov's 500 books covered topics as diverse as science, history, religion, and literature. His fiction inspired Nobel Prize winning liberal economist Paul Krugman; conservative politician Newt Gingrich; the inventor of the industrial robot, George Devol; and the pioneering founder of MIT's Artificial Intelligence Laboratory, Marvin Minsky. His science works educated a nation. According to the United Nations, he is the world's twentieth most widely translated author, ahead of Tolstoy, Dickens, and even Plato.

Yet twenty years after his death, Asimov's robot stories are clearly his most influential accomplishment. He published his first story about intelligent robots in 1939, well before the invention of computers or software. Yet by the time he was 21, he had worked out a set of principles that guided how his fictional robots should interact with people. These became the Three Laws. Today, they are the starting point for any serious conversation about how humans and robots will behave around one another.

As the mere fact of lawyers discussing robot law shows, the issue is no longer theoretical. If robots are not yet intelligent, they are increasingly autonomous in how they carry out tasks. At the most basic level, every day millions of robotic Roombas decide how to navigate tables, chairs, sofas, toys, and even pets as they vacuum homes.

At a more sophisticated level, autonomous robots help select inventory in warehouses, move and position products in factories, and care for patients in hospitals. South Korea is testing robotic jailers.

Life And Death

The United States, South Korea, and Israel all operate sophisticated military drones and robotic land vehicles. Today, human operators control their operation, but militaries around the world are investing heavily in more autonomous systems. The day is approaching when life and death decisions could balance on a mathematical algorithm.

"Robots are approaching a take-off point," said A. Michael Froomkin, the Miami Law professor behind We Robot. He foresees a world in which technology moves so rapidly, we will suddenly find ourselves surrounded by millions of autonomous robots serving and working along side us.

It sounds like a stretch, but Froomkin has seen it before. He began working in Internet law in the early 1990s, when few outside university and military teams had heard the term. Yet even then, engineers had already deployed many of the key standards.

"Some of the early choices, such as how we ignored privacy and security in our standards, had already been made. We could have avoided a significant fraction of today's problems if the engineers who made those choices had been thinking about those issues.

"In robotics, it's still early. The industry has no standards

Graphic design by Roland Liwag/ University of Miami School of Law. All rights reserved.

yet. Thanks to events like this one, people can get in on the ground floor and make our concerns known," he said.

The conference was an early attempt to do just that. For example, University of Washington law professor Neil Richards and computer scientist Bill Smart delivered a paper on how the law should think about robots. It quickly evolved into a debate over the nature of autonomous robots themselves.

One group argued that robots are like tools, wielded by users like hammers. Another compared them to horses, which we can train and contain, but which have minds of their own and can escape and cause damage that is not the fault of the owner.

Law Enforcement

The discussions veered in unexpected directions. A group of scholars submitted a paper about robotic law enforcement and asked if we really wanted "perfect" law enforcement. Policing today, whatever its faults and strengths, relies on judgment. Not everyone speeding gets a ticket, but police pull enough of them over to remind everyone else to drive carefully.

"Do we really want a system that is relentless and never looks the other way?" asked Mary Ann Franks, a Miami Law associate professor who analyzed the paper for the conference.

She worries about the social costs of such a system. "It would degrade the ideal of responsible citizenship. We would be following the law due to fear of reprisal, rather than internalizing the law and acting like good citizens," Franks argued. She worried that robotic law enforcement would make citizens more docile.

For two days, discussions spun off in many different directions, from how autonomous military robots might operate to such traditional concerns as liability and contracts. One particular illuminating exchange looked at whether asking a robot to spy on its master would be similar to wiretapping, and if commanding the same robot to ask its owner questions whose answers would incriminate him or her crosses the line when it comes to invasion of privacy.

It was exactly the kind of debate Isaac Asimov would have loved.

Asimov was born in Russia in 1920, the same year Czechoslovakian playwright Karel Capek coined the term "robot". It comes from the term "serf," and Czechs used it to describe toil or drudgery. Along with the word robot, Capek also gave the world one of its most enduring plot lines. In his play, *R.U.R.*, the robots rebel against their masters, who doomed them to a life of cruel servitude.

Asimov understood this plot. When he was three, his family moved to Brooklyn and his father eventually opened a candy store. When he was not going to school or doing homework, Asimov was learning about robots—and a lot else—from the pulp fiction magazines in his father's candy store.

Pulp fiction was not high art. Printed on cheap paper with untrimmed edges, it combined fast-paced stories, square-jawed heroes, menacing villains, and lurid artwork prominently featuring scantily clad women. Science fiction pulps were similar, except the villains were usually aliens or mechanical monsters that turned on their creators with alarming frequency.

Asimov's father wanted his children studying so they could attend college, not reading pulps. But young Isaac would point to titles that prominently featured the word "science" and argue that these pulps were educational.

Yet even as a teenager, Asimov had a problem with robots turning on their masters. This was because he wanted to be a scientist, and also read widely about science and engineering.

"I didn't think a robot should be sympathetic just because it happened to be nice," he later explained in "The Word I Invented," a 1980 essay. "It should be engineered to meet

Rebelling robots attack their masters in this photo of a 1922 production of Karel Capek's play *R.U.R.*

Digital Predator Prowls The Skies

A Hellfire missile is launched from a Predator unmanned aerial vehicle, top. U.S. Air Force Captain Richard Koll, left, and Airman 1st Class Mike Eulo, above, perform function checks after launching a Predator at Balad Air Base, Iraq recently. Captain Koll, the pilot, and Airman Eulo, the sensor operator, will handle the Predator in a radius of approximately 25 miles around the base before handing it off to personnel stationed in the United States to continue its mission. Both are assigned to the 46th Expeditionary Reconnaissance Squadron. Below, a Hellfire missile impacts an armored vehicle during a test firing exercise.

certain safety standards as any other machine should in any right-thinking technological society. I therefore began to write stories about robots that were not only sympathetic, but were sympathetic because they couldn't help it."

Just as screwdrivers have hilts, electrical wires have circuit breakers, and elevators have friction brakes, Asimov thought robots should have built-in mechanisms to keep them from harming the people.

By the time Asimov was 19, he had earned a master's degree in chemistry from Columbia University. (He received his Ph.D. in biochemistry after World War II.)

He was also an established science fiction author. He had published his first story at 17, and soon came under the influence of John W. Campbell, the new editor of th magazine *Astounding Science Fiction.*

Before Campbell, the "science" in science fiction was tenuous at best. Spaceships sailed between stars as if the speed of light posed no barrier. John Carter fell sleep and woke up naked on Mars. Aliens had giant brains, but no one ever explained how their puny bodies could generate enough energy or blood flow to keep them thinking.

Campbell demanded something different. He wanted real people, often scientists and engineers, and real—or at least believable—science. Since starships could not exceed the speed of light, his writers had them "fold" space or detour through higher dimensions. His aliens had psychologies. Asimov's robots had "positronic" brains.

Campbell expected science fiction stories to probe technology's implications. As Asimov's contemporary, Frederick Pohl, explained, "A good science fiction story should be able to predict not the automobile but the traffic jam." Asimov, an early Campbell favorite, excelled in just those kinds of logical gymnastics.

Although Campbell rejected Asimov's first robot story, "Robbie," he later remembered that a character had said, "He just can't help being faithful and loving and kind. He's a machine—made so." This later morphed into the First Law of Robotics, that robots must protect human life.

Campbell bought Asimov's second story, "Reason," two years later. In it, a character said, "Those robots are guaranteed to be subordinate." That suggested the Second Law, that robots must follow human commands.

"Liar!", published one month after "Reason," laid the foundations for the robotic laws saying: "On no conditions is a human being to be injured in any way, even when such injury is directly ordered by another human."

Came Together

Everything came together in "Runaround," published in October 1941 when Asimov was 21. It concerned two engineers on Mercury who desperately needed selenium to cool their living quarters. They send a robot to get it. Instead of returning, they find the robot circling a pool of selenium and acting drunk.

The reason? The engineers had not told the robot that they would die without the selenium. The robot followed their command (Second Law) until it got close to the selenium, which was in a dangerous location. It then backed off to preserve itself (Third Law). When it was safe, it turned around to get the selenium, then backed off again.

The robot was caught in a loop and could not accept new instructions. One engineer broke the loop by exposing himself to danger. This caused the robot to default to the First Law in order to protect his life. They then gave the robot instructions that emphasized their need, and the robot retrieved the selenium.

The Three Laws were implicit in "Runaround," but it took Campbell to spell them out in a letter to Asimov. "It always seemed to me that John invented those Laws, but whenever I accused him of that, he always said that they were in my stories and I just hadn't bothered to isolate

them. Perhaps he was right," Asimov wrote.

"Liar!" and "Runaround" solidified the Three Laws, but more importantly, they anticipated "the traffic jam"—the contradictions inherent in trying to control intelligent robots through a simple set of rules.

The ability of Asimov's robots to think for themselves often led to conflict. In "Liar!", for example, a robot develops the ability to read people's minds and lies to everyone to keep from hurting their feelings. In "The Evitable Conflict," robots running the economy decide they must hurt some humans in order to benefit the majority, effectively taking over the world in order to help us.

Asimov's stories explore these contradictions. In them, robots act bizarrely when under stress, develop emotional attachments, and sometimes even rationalize the right to harm humans. One robot develops its own religion. Another sues to be declared a person. Asimov's robot mysteries, starting with *The Caves of Steel* in 1953, pair a robotic investigator with a detective who dislikes robots.

Asimov's stories probed the Three Laws' limitations and failures. These contradictions have moved from theory to fact as the gap between fiction and smart and/or autonomous robots narrowed over the past decade.

Spoken Commands

Artificial intelligence has certainly come a long way. Apple's Siri, which understands spoken commands well enough to run iPhone applications, is just the latest example. Other AI applications summarize news, trade stocks and bonds, unearth paper trails among legal documents, and manage aircraft check-in and cargo—all without human intervention.

AI is driving autonomous robots. Nowhere is this more evident than in driverless vehicles. Eight years ago, in *The New Division of Labor,* economists Frank Levy and Richard Murnane analyzed which tasks humans and computers did best. Humans, they argued, topped machines when rules were difficult to understand. One example was driving a truck while reacting to road and traffic conditions.

Their argument made sense. That same year, the U.S. Defense Advanced Research Projects Agency (DARPA) sponsored a driverless vehicle challenge. No contestant went further than seven miles on the straight, 150-mile course.

Three years later, six autonomous vehicles completed the DARPA Urban Challenge, a 60-mile course through an abandoned military base with street signs, stoplights, and merging traffic.

Last year, Google unveiled its fleet of autonomous cars. It has logged more than 200,000 miles on city streets and highways with only minor human intervention. More recently, Google launched an autonomous racing car team, and Mercedes-Benz and Cadillac disclosed that they are developing vehicles to navigate stop-and-go traffic autonomously.

Ten years ago, military robots existed only in laboratories. By 2009, the United States had deployed more than 7,000 drones and 12,000 ground robots. South Korea and Israel use robots to patrol their borders. Humans guided all

They've Come A Long Way

The Japanese "karakuri" tea-serving automaton, above, was built around 1800, above. "Karakuri" means "mechanism" and they were typically spring-powered. The da Vinci surgical system below is controlled by a surgeon from a console. It is commonly used for prostatectomies, and increasingly for cardiac valve repair and gynecologic surgical procedures

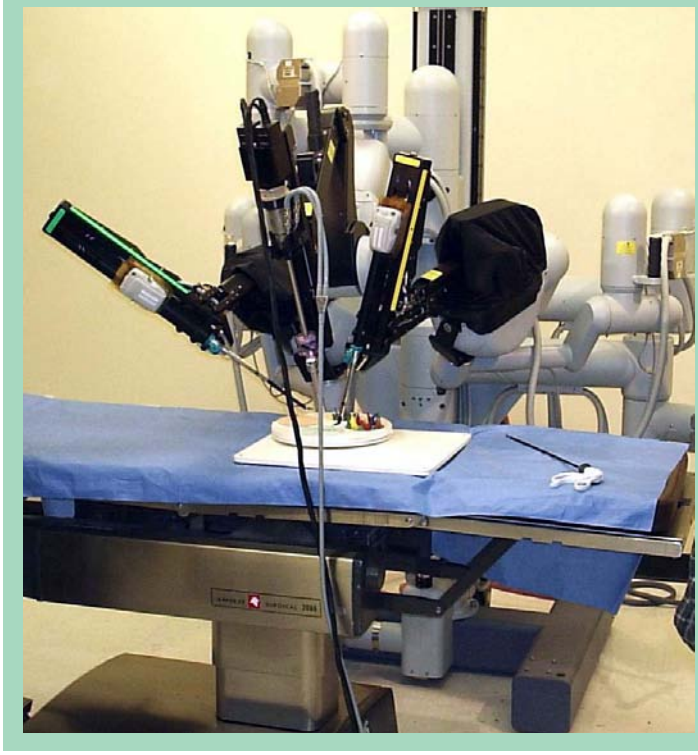

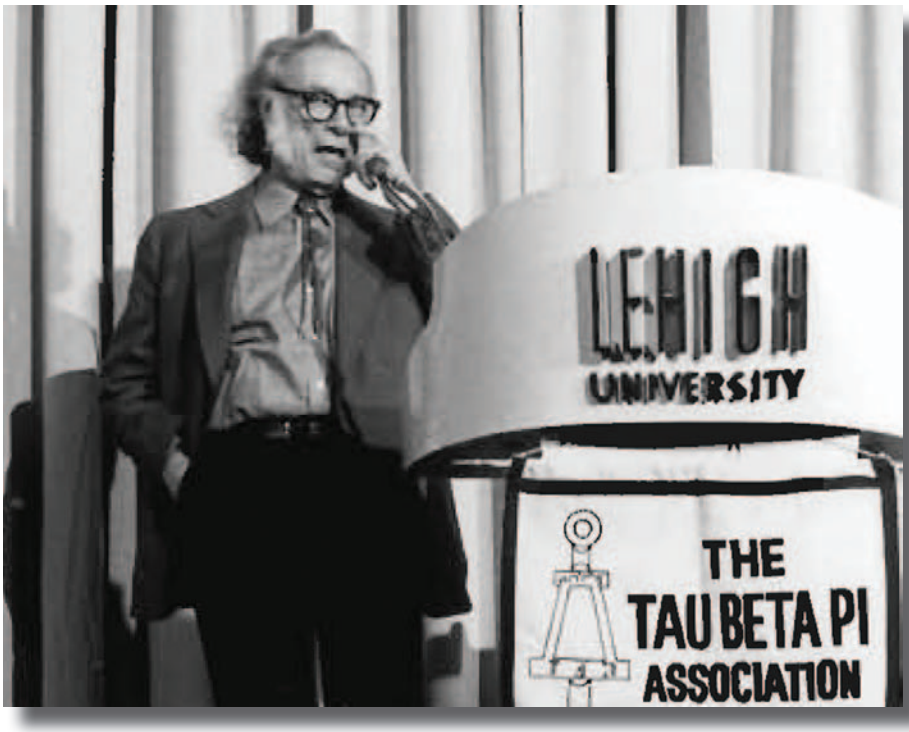

Guest speaker Isaac Asimov speaks at the ΤΒΠ 100th anniversary celebration in 1985.

these robots, but militaries around the world are developing more autonomous warrior robots. Some will pack weapons.

In *Wired for War*, Peter Singer, a military analyst at the Brookings Institute, worried that robots reduced the moral hazard—the personal risk—of going to war. Without casualties, memories, and losses, Singer fears that war could morph into entertainment. In fact, videos of drone attacks have become very popular on the web. Soldiers call them "war porn."

The lawyers and technologists at We Robot struggled with similar concerns. Today's robots exercise only limited autonomy, and even then, only when executing such simple, constrained tasks as bringing medicines to patients or moving warehouse inventory. Tomorrow's robots are likely to have far greater choice about how they perform a broader range of tasks.

That is why some thinkers have developed their own laws of robotics. They range from the earnest (robots must always know they are robots) to the provocative (robots should seek better sources of power and reproduce other robots) to the cynical (military robots should kill only the enemy).

More Pragmatic

Others take a more pragmatic view. Texas A&M computer scientist Robin Murphy and Ohio State Cognitive Systems Engineering Lab director David Woods argue that we need laws that apply to increasingly autonomous but not yet intelligent robots.

Their first law is that since humans deploy robots, any discussion of safety and ethics must consider human-robot systems rather than robots alone. Second, robots must obey only appropriate commands, and only from a limited number of people. Third, robots should protect themselves, but only after they transfer control of whatever they are doing (like driving or surgery) to humans.

As robots grow more capable, the line separating autonomy from intelligence is likely to grow fuzzier. After all, a robot may not have any sense of its own "personhood," but then, neither does a horse. Yet both robots and horses can modify their behavior based on what they learn from experience. In this way, a robot is more like a horse than a hammer, and no one would argue that horses have no intelligence.

As robots become smarter and more independent, what kind of rules might bind them to humans? Michael Anissimov, director of advocacy at Singularity Institute for Artificial Intelligence, which was founded to develop safe AI software, grappled with that problem in 2004. He asserted that it is "not so straightforward to convert a set of statements into a mind that follows or believes in those statements."

A robot, he explained, could easily misapply laws in complex situations if it did not understand their original purpose. Instead of laws, Anissimov believes we must create "friendly AI" that loves humanity.

He may be right. Yet as the participants showed, it is not easy to develop hard and fast rules for a rapidly changing technology. This is especially true for robotics, which relies on open-source software modified by dozens and perhaps hundreds of users and components made by scores of different companies.

In Spain, Qbo, a small robot company, is tapping all these sources to build mobile utility robots that speak and respond to voice commands, play music, change television channels, and find information on the Internet. They can also upload information to cloud computers, so a Qbo that learns to recognize an object in Madrid can share that information with a Qbo in Barcelona.

Right now, the system is new and there are not a lot of Qbo robots in circulation. But imagine a world in the not-to-distant future where millions of autonomous robots share what they have learned that day over the cloud. On powerful servers, machine learning programs will sift through the information, finding patterns, drawing lessons, and downloading what they have discovered.

Autonomous robots will learn and modify their behavior in wholly new ways. The results are likely to be far-reaching and unpredictable. We are creating something truly different in the world.

Perhaps that is why We Robot's logo borrows from Michelangelo's famous depiction at the Sistine Chapel of God imparting life to Adam. Only on the logo, a metallic robotic hand has replaced God's finger. The point is clear: Although man invented robots, our robots are going to change us.

Somewhere, Isaac Asimov is smiling.

Alan S. Brown has been an editor and freelance writer for more than 25 years and lives in Dayton, NJ (*insight01@verizon.net*). A member of the National Association of Science Writers and former co-chair of the Science Writers in New York, he graduated magna cum laude from New College at Hofstra University in 1974. He is an associate editor of *Mechanical Engineering* and contributes to a wide range of engineering and scientific publications.