

The Machines from Our Future

Daniela Rus

While the last sixty years have defined the field of industrial robots and empowered hard-bodied robots to execute complex assembly tasks in constrained industrial settings, the next sixty years will usher in our time with pervasive robots that come in a diversity of forms and materials and help people with physical tasks. The past sixty years have mostly been inspired by the human form, but the form diversity of the animal kingdom has broader potential. With the development of soft materials, machines and materials are coming closer together: machines are becoming compliant and fluid-like materials, and materials are becoming more intelligent. This progression raises the question: what will be the machines from our future?

Today, telepresence enables students to meet with tutors and teachers and allows doctors to treat patients thousands of miles away. Robots help with packing on factory floors. Networked sensors enable the monitoring of facilities, and 3D printing creates customized goods. We are surrounded by a world of possibilities. And these possibilities will only get larger as we start to imagine what we can do with advances in artificial intelligence and robotics. Picture a world where routine tasks are taken off your plate. Fresh produce just shows up on your doorstep, delivered by drones. Garbage bins take themselves out, and smart infrastructure systems support automated pick-up. AI assistants – whether embodied or not – act as guardian angels, providing advice to ensure that we maximize and optimize our lives to live well and work effectively.

The field of robotics has the potential to greatly improve the quality of our lives at work, at home, and at play by providing people with support for cognitive and physical tasks. For years, robots have supported human activity in dangerous, dirty, and dull tasks, and have enabled the exploration of unreachable environments, from the deep oceans to deep space. Increasingly more-capable robots will be able to adapt, learn, and interact with humans and other machines on cognitive levels. The objective of robotics is not to replace humans by mechanizing and automating tasks, but rather to find new ways that allow robots to collaborate with humans more effectively. Machines are better than humans at tasks such as crunching numbers and moving with precision. Robots can lift much heavier objects. Humans are better than machines at tasks like reasoning, defining abstractions, and generalizing or specializing, thanks to our ability to draw on prior expe-

riences. By working together, robots and humans can augment and complement each other's skills.

Imagine riding in your flying car, which is integrated with the information technology infrastructure and knows your needs, so it can tell you, for example, that you can buy the plants you have been wanting at a store nearby, while computing a small detour. You can trust your home to take care of itself when you are away. That is what the smart refrigerator is for: it tracks everything you put in and take out so it can automatically send a shopping list to your favorite delivery service when it is time to restock. This automated household can help take care of everything from your new plants to your elderly parents. The intelligent watering system monitors the soil and ensures each type of plant gets the right level of moisture. When your elderly parents need help with cooking, the kitchen robot can assist. The new robotic technologies can also be carried with us, knitted in our sweaters, blended in our garments, or embedded in our accessories. We could have our own wearable computer assistants, like Ironman, with superpowers focused on improving and optimizing our health and everyday lives. The smart exosuit can provide an extra set of eyes that monitors the environment and warns of threats when we walk home at night. This exosuit, shaped as a knitted robot, could become an individual coach to help us perfect a tennis stroke or an assembly sequence. This is just a snapshot of a machine-enhanced future we can imagine. There are so many ways in which our lives can be augmented by robots and AI.

This positive human-machine relationship, in which machines are helpful assistants, is closer to my vision of the future than the scenarios in which the machines either take over as maniacal overlords or solve all of humanity's problems. This vision is starting to mature inside my lab, and in the labs of my friends and colleagues at other universities and institutions and some forward-thinking companies. This future does not resemble the dystopia depicted in so many books, movies, and articles. But none of us expects it to be a perfect world, either, which is why we design and develop the work with potential dangers in mind.

While AI is concerned with developing the science and engineering of intelligence for cognitive tasks, robotics is concerned with physical-world interactions by developing the science and engineering of autonomy. Specifically, robots are made of a body (hardware) and a brain (algorithms and software). For any task that requires machine assistance, we need bodies capable of doing the task and brains capable of controlling the bodies to do the task. The main tasks studied in robotics are mobility (navigating on the ground, in air, or underwater), manipulation (moving objects in the world), and interaction (engaging with other machines and with people).

We have already come a long way. Today's state of the art in robotics, AI, and machine learning is built on decades of advancements and has great potential for positive impact. The first industrial robot, called The Unimate, was introduced in

1961. It was invented to perform industrial pick and place operations. By 2020, the number of industrial robots in operation reached around twelve million, while the number of domestic robots reached thirty-one million.¹ These industrial robots are masterpieces of engineering capable of doing so much more than people can, yet these robots remain isolated from people on factory floors because they are large, heavy, and dangerous to be around. By comparison, organisms in nature are soft, safe, compliant, and much more dexterous and intelligent. Soft-bodied systems like the octopus can move with agility. The octopus can bend and twist continuously and compliantly to execute many other tasks that require dexterity and strength, such as opening the lid of a jar. Elephants can move their trunks delicately to pick up potato chips, bananas, and peanuts, and they can whip those same trunks with force enough to fight off a challenger. If robots could behave as flexibly, people and robots could work together safely side by side. But what would it take to develop robots with these abilities?

While the past sixty years have established the field of industrial robots and empowered hard-bodied robots to execute complex assembly tasks in constrained industrial settings, the next sixty years will usher in soft robots for human-centric environments and to help people with physical and cognitive tasks. While the robots of the past sixty years have mostly been inspired by the human form, shaped as industrial arms, humanoids, and boxes on wheels, the next phase for robots will include soft machines with shapes inspired by the animal kingdom and its diversity of forms, as well as by our own built environments. The new robot bodies will be built out of a variety of available materials: silicone, wood, paper, fabric, even food. These machines of our future have a broader range of customized applications.

Today's industrial manipulators enable rapid and precise assembly, but these robots are confined to operate independently from humans (often in cages) to ensure the safety of the humans around them. The lack of compliance in conventional actuation mechanisms is part of this problem. In contrast, nature is not fully rigid; it uses elasticity and compliance to adapt. Inspired by nature, soft robots have bodies made out of intrinsically soft and/or extensible materials (such as silicone rubbers or fabrics) and are safe for interaction with humans and animals. They have a continuously deformable structure with muscle-like actuation that emulates biological systems and provides them with a relatively large number of degrees of freedom as compared with their hard-bodied counterparts. Soft robots have capabilities beyond what is possible with today's rigid-bodied robots. For example, soft-bodied robots can move in more natural ways that include complex bending and twisting curvatures that are not restricted to the traditional rigid body kinematics of existing robotic manipulators. Their bodies can deform continuously, providing theoretically infinite degrees of freedom and allowing them to adapt their shape to their environments (such as by conforming to natural terrain or forming enveloping power grasps). However, soft robots have also been

shown to be capable of rapid agile maneuvers and can change their stiffness to achieve a task- or environment-specific impedance.

What is soft, really? Softness refers to how stretchy and compliant the body of the robot is. Soft materials and electromechanical components are the key enablers for creating soft robot bodies. Young's modulus, which computes the ratio of stress to strain of a material when force is applied, is a useful measure of the rigidity of materials used in the fabrication of robotic systems. Materials traditionally used in robotics (like metals and hard plastics) have Young's moduli on the order of 10^9 to 10^{12} pascals (a unit of pressure), whereas natural organisms are often composed of materials (like skin and muscle tissue) with moduli on the order of 10^4 to 10^9 pascals, orders of magnitude lower than their engineered counterparts. We define soft robots as systems capable of autonomous behavior that are primarily composed of materials with moduli in the range of soft biological materials.

Current research on device-level and algorithmic aspects of soft robots has resulted in a range of novel soft devices. But how do we get to the point where soft robots deliver on their full potential? The capabilities of robots are defined by the tight coupling between their physical bodies and the computation that makes up their brains. For example, a robot fish must have both a body capable of swimming and algorithms to control its movement in water. Today's soft-bodied robots can do basic locomotion and grasping. When augmented with appropriate sensors and computation, they can recognize objects in restricted situations, map new environments, perform pick and place operations, and even act as a coordinated team.

Figure 1 shows SoFi, the soft robotic fish.² SoFi is an autonomous soft robot developed for close observations and interactions with marine life. SoFi enables people to observe and monitor marine life from a distance, without interference. The robot swims continuously at various depths in a biomimetic way by cyclic undulation of its posterior soft body. The fish controls the undulating motion of its tail using a hydraulically actuated soft actuator with two internal cavities separated by an inextensible constraint. The fish tail has two chambers with ribbed structure for pressurization, and the inextensible constraint is in the middle. Maneuvering is accomplished by moving water from one chamber to the other using a pump. When the pump moves water equally between the left and right chambers of the tail, the tail moves back and forth evenly, and the fish exhibits forward swimming. It is possible to make right-hand turns by pumping more water in the right chamber than the left and doing the reverse for left-hand turns. The swimming depth is controlled by two dive planes that represent the robot's fins. SoFi has onboard capabilities for autonomous operation in ocean environments, including the ability to move along 3D trajectories by adjusting its dive planes or by controlling its buoyancy. Onboard sensors perceive the surrounding environ-

Figure 1

SoFi, the Soft Robotic Fish for Underwater Observatories



SoFi, the soft robotic fish swimming in a coral reef. Source : Photo by Joseph DelPreto. See Robert K. Katzschmann, Joseph DelPreto, Robert MacCurdy, et al., “Exploration of Underwater Life with an Acoustically Controlled Soft Robotic Fish,” *Science Robotics* 3 (16) (2018).

ment, and a mission control system enables a human diver to issue remote commands. SoFi achieves autonomy at a wide range of depths through 1) a powerful hydraulic soft actuator; 2) a control mechanism that allows the robot to adjust its buoyancy according to depth, thus enabling long-term autonomous operation; 3) onboard sensors to observe and record the environment; 4) extended ocean experiments; and 5) a mission control system that a human diver can use to provide navigation commands to the robot from a distance using acoustic signals. SoFi has the autonomy and onboard capabilities of a mobile underwater observatory, our own version of Jules Verne’s marine observatory in *Twenty Thousand Leagues Under the Sea*. Marine biologists have long experienced the challenges of documenting ocean life, with many species of fish proving quite sensitive to the underwater movements of rovers and humans. While multiple types of robotic

instruments exist, the soft robots move by undulation and can more naturally integrate in the undersea ecosystems. Soft-bodied robots can move in more natural and quieter ways.


The body of a soft robot like SoFi may consist of multiple materials with different stiffness properties. A soft robot encases in a soft body all the subsystems of a conventional robot: an actuation system, a perception system, driving electronics, and a computation system, with corresponding power sources. Technological advances in soft materials and subsystems compatible with the soft body enable the autonomous function of the soft robot.

At the core of any robot is actuation. One of the primary focus areas to date for soft robots has been the exploration of new concepts for compliant yet effective actuators. Researchers have made progress on several classes of soft actuators, most prominently with fluidic or various electrically activated tendon actuators. Fluidic elastomer actuators (FEAs) are highly extensible and adaptable, low-power soft actuators. FEAs were used to actuate SoFi's tail. Figure 2 shows the actuation principle. A silicone chamber has an inextensible constraint. When it is pressurized – for example, with air or liquid – the skin expands and forms a curvature. By controlling this curvature, we can control the movement of the robot.

The soft actuator in Figure 2 can move along one axis and is thus called a one-degree-of-freedom actuator. Such an actuator can be composed in series and in parallel to create any desired compliant robotic morphology: a robotic elephant trunk, a robotic multifinger hand, a robotic worm, a robotic flower, a robotic chair, even a robotic lamp.

However, while achieving compliance, this FEA actuator structure has not achieved muscle-like or motor-like performance in terms of force, displacement, energy density, bandwidth, power density, and efficiency. In order to create muscle-like actuation, we can leverage the idea of combining soft bodies with compliant origami structures to act as “flexible bones” within the soft tissue. The idea of fluidic origami-inspired artificial muscles (FOAM) provides fluidic artificial muscles with unprecedented performance-to-cost ratio.³ The FOAM artificial muscle system consists of three components: a compressible solid skeletal structure (an origami structure), a flexible fluid-tight skin, and a fluid medium. When a pressure difference is applied between the outside and the inner portion, a tension is developed in the skin that causes contraction that is mediated by the folded skeleton structure. In a FOAM system, the skin is sealed as a bag covering the internal components. The fluid medium fills the internal space between the skeleton and the skin. In the initial equilibrium state, the pressures of the internal fluid and the external fluid are equal. However, as the volume of the internal fluid changes, a new equilibrium is achieved. A pressure difference between the internal and external fluids induces tension in the flexible skin. This tension will act on the skeleton, driving a trans-

Figure 2
Soft Fluidic Actuation



Source : Robert Katzschmann, “Building and Controlling Fluidically Actuated Soft Robots: From Open Loop to Model-based Control” (Ph.D. diss., MIT, 2013).

formation that is regulated by its internal skeletal geometry. These artificial muscles can be programmed to produce not only a single contraction, but also complex multi-axial actuation and even controllable motion with multiple degrees of freedom. Moreover, a variety of materials and fabrication processes can be used to build the artificial muscles with other functions beyond basic actuation. Experiments reveal that these muscles can contract over 90 percent of their initial lengths, generate stresses of approximately 600 kilopascals, and produce peak power densities over 2 kilowatts per kilogram: all equal to, or in excess of, natural muscle. For example, a 3 gram FOAM actuator that includes a zig-zag pattern for its bone structure can lift up to 3 kilograms! This architecture for artificial muscles opens the door to rapid design and low-cost fabrication of actuation systems for numerous applications at multiple scales, ranging from miniature medical devices to wearable robotic exoskeletons to large deployable structures for space exploration.

The soft FOAM grippers shown in Figure 3 are made from a soft origami structure, encased by a soft balloon.⁴ When a vacuum is applied to the balloon, the origami structure – a design based on a folding pattern – closes around the object, and the gripper deforms to the geometric structure of the object. While this motion lets the gripper grasp a much wider range of objects than ever before, such as soup cans, hammers, wine glasses, drones, even a single broccoli floret or grape, the greater intricacies of delicacy – in other words, how hard to squeeze – require adding sensors to the gripper. Tactile sensors can be made from latex “bladders”

Figure 3
FOAM Grippers (*top*) and Objects that can be Handled with the Magic Ori-
gami Ball (Tulip) Gripper (*bottom*)



Source : Photos by Shuguang Li.

(balloons) connected to pressure transducers. The new sensors let the gripper not only pick up objects as delicate as potato chips, but it also classifies them, providing the robot with a better understanding of *what* it is picking up, while also exhibiting that light touch. When the embedded sensors experience force or strain, the internal pressure changes, and this feedback can be used to achieve a stable grasp.

In addition to such discrete bladder sensors, we can also give the soft robot bodies sensorized “skin” to enable them to see the world by feeling the world. The sensorized skin provides feedback along the entire contact surface, which is valuable for learning the type of object it is grasping and exploring the space of the robot through touch. Somatosensitive sensors can be embedded in the silicone body of the robot using 3D printing with fugitive and embedded ink. Alternatively, elec-

trically conductive silicone can be cut using a variety of stretchable kirigami patterns and used for the sensor skin of the robot. Machine learning can then be used to associate skin sensor values with robotic deformations, leading to proprioceptive soft robots that can “see” the world through touch.

The robot body needs a robot brain to command and coordinate its actions. The robot brain consists of the set of algorithms that can get the robot to deliver on its capabilities. These algorithms typically map onto computation for physically moving the components of the robot (also called low-level control) and computation for getting the robot to perform its assignment (also called high-level or task-level control).

While we have a surge in developing soft bodies for robots, the computational intelligence and control of these robots is more challenging. Results from rigid robots do not immediately translate to soft robots because of their inherent high dimensionality. The state of a rigid robot can be described compactly with a finite set of degrees of freedom: namely, the displacement of each of its joints as described in their local coordinate frames. Their bodies are constrained by the inflexible nature of their rigid links. Fully soft robots, by contrast, may not have a traditional joint structure, relying on their flexible body to solve tasks. Soft robots have a dramatically different interaction with the environment through rich compliant contact. There is currently a divide in the approach to control: rigid robots control contact forces/contact geometry while soft robots rely almost entirely on open-loop interactions, mediated by material properties, to govern the resulting forces/geometry. One strategy for bridging this gap lies in optimization-based control via approximate dynamic models of the soft interface: models with a fidelity that is customized to the task. The governing equations of the soft robots are complex continuum mechanics formulations that are typically approximated using high-dimensional finite-element methods. The dynamics are highly nonlinear, and contacts with the environment make them nonsmooth. These models are too complex for state-of-the-art feedback design approaches, which either make linearity assumptions or scale badly with the size of the state space. The challenge is to find models simple enough to be used for control, but complex enough to capture the behavior of the system.

For low-level control of soft robots, we can often identify a sequence of actuated segments, in which torques are dominant, so it is possible to assume the curvature to be constant within each segment, leading to a finite-dimensional Piecewise Constant Curvature (PCC) kinematic description. We can then describe the PCC of the soft robot through an equivalent rigid robot with an augmented state space.

Task-level control of soft robots is often achieved in a data-driven way using machine learning. Some of today’s greatest successes of machine learning are due to a technique called *deep learning*. Deep learning uses data – usually millions of

hand-labeled examples – to determine the weights that correspond to each node in a convolutional neural network (CNN), a class of artificial neural networks, so that when the network is used with new input, it will classify that input correctly. Deep learning has been successfully applied to soft robots to provide them with capabilities for proprioception (sensitivity to self-movement, position, and action), exteroception (sensitivity to outside stimuli), and grasping.

But deep learning faces a number of challenges. First among them is the data. These techniques require data availability, meaning massive data sets that have to be manually labeled and are not easily obtained for every task. The quality of that data needs to be very high, and it needs to include critical corner cases – that is, cases outside the training distribution or outside usual operations – for the application at hand. If the data are biased, the performance of the algorithm will be equally bad. Furthermore, these systems are black boxes: there is no way for users of the systems to truly “learn” anything based on the system’s workings. It is difficult to detect behavior that is abnormal from a safety point of view. As a result, it is hard to anticipate failure modes tied to rare inputs that could lead to potentially catastrophic consequences. We also have robustness challenges and need to understand that the majority of today’s deep-learning systems perform pattern matching rather than deep reasoning. Additionally, there are sustainability issues related to data-driven methods. Training and using models consume enormous amounts of energy. Researchers at the University of Massachusetts Amherst estimated that training a large deep-learning model produces 626,000 pounds of carbon dioxide, equal to the lifetime emissions of five cars. The more pervasive machine learning becomes, the more of these models will be needed, which in turn has a significant environmental impact.

Today’s machine learning systems are so costly because each one contains hundreds of thousands of neurons and billions of interconnections. We need new ideas to develop simpler models, which could drastically reduce the carbon footprint of AI while gaining new insights into intelligence. The size of a deep neural network constrains its capabilities and, as a result, these networks tend to be huge and there is an enormous cost to running them. They are also not interpretable. In deep neural networks, the architecture is standardized, with identical neurons that each compute a simple thresholding function. A deep neural network that learns end-to-end from human data how to control a robot to steer requires more than one hundred thousand nodes and half a million parameters.

Using inspiration from neuroscience, my colleagues and I have developed neural circuit policies,⁵ or NCPs, a new approach to machine learning. With NCPs, the end-to-end steering task requiring more than one hundred thousand simple neurons can be learned with nineteen NCP neurons in the deep neural network model, resulting in a more efficient and interpretable system. The neuroscience inspiration from the natural world is threefold. First, NCP neurons can compute more than a

step function; each NCP neuron is a liquid time differential equation. Second, NCP neurons can be specialized, such as input, command, and motor neurons. Third, the wiring architecture has organism-specific structure. Many other tasks related to spatial navigation and beyond can be realized with neuroscience-inspired, compact, and interpretable neural circuit policies. Exploring robot intelligence using inspiration from the natural world will yield new insights into life and provide new computational models for intelligence that are especially useful for soft robots.

Novel soft design, fabrication, and computation technologies are ushering in a new era of robots that come in a variety of forms and materials and are designed to help people with physical tasks in human-centric environments. These robots are smaller, safer, easier to fabricate, less expensive to produce, and more intuitive to control.

Robots are complex systems that tightly couple the physical mechanisms (the body) with the software aspects (the brain). Recent advances in disk storage, the scale and performance of the Internet, wireless communication, tools supporting design and manufacturing, and the power and efficiency of electronics, coupled with the worldwide growth of data storage, have helped shape the development of robots. Hardware costs are going down, the electromechanical components are more reliable, the tools for making robots are richer, the programming environments are more readily available, and the robots have access to the world's knowledge through the cloud. Sensors like the LiDAR (light detection and ranging) systems are empowering robots to measure distances very precisely. Tiny cameras are providing a rich information stream. Advances in the development of algorithms for mapping, localization, object recognition, planning, and learning are enabling new robotic capabilities. We can begin to imagine the leap from the personal computer to the personal robot, leading to many applications in which robots exist pervasively and work side by side with humans.

How might these advances in robotics shape our future? Today, if you can think it, you can write it on paper. Imagine a world where if you can think it, you can make it. In this way, the scientific advancement of soft robotics could give every one of us superpowers. Each of us could use our talents, our creativity, and our problem-solving skills to dream up robots that save lives, improve lives, carry out difficult tasks, take us places we cannot physically go, entertain us, communicate, and much more. In a future of democratized access to robots, the possibilities for building a better world are limitless. Broad adoption of robots will require a natural integration of robots in the human world, rather than an integration of humans into the machines' world.

These machines from our future will help us transform into a safer society living on a healthier planet, but we have significant technological and societal challenges to get to that point.

On the technical side, it is important to know that most of today's greatest advances in machine learning are due to decades-old ideas enhanced by vast amounts of data and computation. Without new technical ideas and funding to back them, more and more people will be ploughing the same field, and the results will only be incremental. We need major breakthroughs if we are going to manage the major technical challenges facing the field. We also need the computational infrastructure to enable the progress, an infrastructure that will deliver to us data and computation like we get water and energy today: anywhere, anytime, with a simple turn of a knob. And we need the funding to do this.

On the societal side, the spread of AI and robots will make our lives easier, but many of the roles that they can play will displace work done by humans today. We need to anticipate and respond to the forms of economic inequality this could create. In addition, the lack of interpretability and dependence could lead to significant issues around trust and privacy. We need to address these issues, and we need to develop an ethics and legal framework for how to use AI and robots for the greater good. As we gather more data to feed into these AI systems, the risks to privacy will grow, as will the opportunities for authoritarian governments to leverage these tools to curtail freedom and democracy in countries around the world.

These problems are not like the COVID-19 pandemic: we know they are coming, and we can set out to find solutions at the intersection of policy, technology, and business, in advance, now. But where do we begin?

In its report on AI ethics, the Defense Innovation Board describes five AI principles. First is responsibility, meaning that humans should exercise appropriate levels of judgment and remain responsible for the development, deployment, use, and outcomes of these systems. Second, equitability, meaning that we need to take deliberate steps to anticipate and avoid unintended bias and unintended consequences. Third is traceability, meaning that the AI engineering discipline should be sufficiently advanced such that technical experts possess an appropriate understanding of the technology, development processes, and operational methods of its AI systems. Fourth is reliability, meaning that AI systems should have an explicit, well-defined domain of use, and the safety, security, and robustness of such systems should be tested and assured. And finally, governance, meaning that AI systems should be designed and engineered to fulfill their intended function, while possessing the ability to detect and avoid unintended harm or disruption.⁶ Beyond these general principles, we also need to consider the environmental impacts of new technologies, as well as what policy actions are needed to stem possible dangers associated with technological advances.

Neural circuit policies may sound like phrases you would only ever hear walking the hallways of places like CSAIL, the Computer Science and Artificial Intelligence Laboratory at MIT, where I work. We do not need everybody to understand in great detail how this technology works. But we do need our policy-makers and

citizens to know about the effects of new technologies so we can make informed decisions about their adoption. Together, we can build a common understanding around five vital questions: First, what *can* we do, or more specifically, what is really possible with technology? Second, what *can't* we do, or what is not yet possible? Third, what *should* we do? Fourth, what *shouldn't* we do? There are technologies and applications that we should rule out. And finally, what *must* we do. I believe we have an obligation to consider how AI technology can help. Whether you are a technologist, a scientist, a national security leader, a business leader, a policy-maker, or simply a human being, we all have a moral obligation to use AI technology to make our world, and the lives of its residents, safer and better, in a just, equitable way.

The optimist in me believes that can and will happen.

AUTHOR'S NOTE

The author gratefully acknowledges the following support: NSF grant No. EFRI-1830901, the Boeing Company, JMPC, and the MIT-Air Force AI Accelerator program.

ABOUT THE AUTHOR

Daniela Rus, a Fellow of the American Academy since 2017, is the Andrew (1956) and Erna Viterbi Professor of Electrical Engineering and Computer Science and the Director of the Computer Science and Artificial Intelligence Laboratory (CSAIL) at MIT, and the Deputy Dean of Research in MIT's Schwarzman College of Computing. She is also a Fellow of the National Academy of Engineering. She has published in such journals as *Nature*, *Proceedings of the National Academy of Sciences*, and *Science*, and in professional journals such as *International Journal of Robotics Research* and *IEEE Transactions on Robotics*.

ENDNOTES

- ¹ International Federation of Robots, <https://ifr.org>.
- ² Daniela Rus and Michael T. Tolley, “Design, Fabrication and Control of Soft Robots,” *Nature* 521 (7553) (2015): 467–475.
- ³ Shuguang Li, Daniel M. Vogt, Daniela Rus, and Robert J. Wood, “Fluid-Driven Origami Inspired Artificial Muscles,” *Proceedings of the National Academy of Sciences* 114 (50) (2017): 13132–13137.
- ⁴ Shuguang Li, John J. Stampfli, Helen J. Xu, et al., “A Vacuum-Driven Origami ‘Magic-Ball’ Soft Gripper,” 2019 *International Conference on Robotics and Automation (ICRA)* (2019): 7401–7408, <https://doi.org/10.1109/ICRA.2019.8794068>.
- ⁵ Mathias Lechner, Ramin Hasani, Alexander Amini, et al., “Neural Circuit Policies Enable Auditable Autonomy,” *Nature Machine Intelligence* 2 (10) (2020): 642–652.
- ⁶ Defense Innovation Board, *AI Principles: Recommendations on the Ethical Use of Artificial Intelligence by the Department of Defense* (Washington, D.C.: U.S. Department of Defense, 2019).