

ROBOTICS

The grand challenges of *Science Robotics*

Guang-Zhong Yang,^{1*} Jim Bellingham,² Pierre E. Dupont,³ Peer Fischer,^{4,5} Luciano Floridi,^{6,7,8,9,10} Robert Full,¹¹ Neil Jacobstein,^{12,13} Vijay Kumar,¹⁴ Marcia McNutt,¹⁵ Robert Merrifield,¹ Bradley J. Nelson,¹⁶ Brian Scassellati,^{17,18} Mariarosaria Taddeo,^{7,8,9} Russell Taylor,¹⁹ Manuela Veloso,²⁰ Zhong Lin Wang,²¹ Robert Wood^{22,23}

Copyright © 2018
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim
to original U.S.
Government Works

One of the ambitions of *Science Robotics* is to deeply root robotics research in science while developing novel robotic platforms that will enable new scientific discoveries. Of our 10 grand challenges, the first 7 represent underpinning technologies that have a wider impact on all application areas of robotics. For the next two challenges, we have included social robotics and medical robotics as application-specific areas of development to highlight the substantial societal and health impacts that they will bring. Finally, the last challenge is related to responsible innovation and how ethics and security should be carefully considered as we develop the technology further.

INTRODUCTION

Just over a year ago, we published the first issue of *Science Robotics*. Even within this relatively short period of time, remarkable progress has been made in many aspects of robotics—from micromachines for biomedicine (1) to large-scale systems for robotic construction (2) and from robots for outer space to those involved in deep-sea exploration (3). We have seen the evolution of soft robots and how new materials and fabrication schemes have led to deformable actuators that are compliant, versatile, and self-healing (4–6). We have also seen many examples of bioinspired designs, from the power-modulated jumping robot with agility and power that approach those of galagos (the animal with the highest vertical jumping agility) (7) to a biomimetic robotic platform to study flight specializations of bats (8) and a biorobotic adhesive disc for underwater hitchhiking inspired by the remora suckerfish (9). We also celebrated the 10th anniversary of the Robot Operating System (ROS) (10), the open-source robotics middleware that is making

great strides in realizing its mission of powering the world's robots, from space robot challenges to autonomous driving, industrial assembly, and surgery.

Given all these advances, what does the future hold for the field of robotics? Recently, we conducted an open online survey on major unsolved challenges in robotics. On the basis of the feedback and submissions received, an invited online expert panel was convened, and the panel shortlisted the 30 most important topics and research directions. These are further grouped into 10 grand challenges (Fig. 1) that may have major breakthroughs, significant research, and/or socioeconomic impact in the next 5 to 10 years:

(i) **New materials and fabrication schemes** for developing a new generation of robots that are multifunctional, power-efficient, compliant, and autonomous in ways akin to biological organisms.

(ii) **Biohybrid and bioinspired robots** that translate fundamental biological principles into engineering design rules or integrate liv-

ing components into synthetic structures to create robots that perform like natural systems.

(iii) **New power sources, battery technologies, and energy-harvesting schemes** for long-lasting operation of mobile robots.

(iv) **Robot swarms** that allow simpler, less expensive, modular units to be reconfigured into a team depending on the task that needs to be performed while being as effective as a larger, task-specific, monolithic robot.

(v) **Navigation and exploration in extreme environments** that are not only unmapped but also poorly understood, with abilities to adapt, to learn, and to recover and handle failures.

(vi) **Fundamental aspects of artificial intelligence (AI) for robotics**, including learning how to learn, combining advanced pattern recognition and model-based reasoning, and developing intelligence with common sense.

(vii) **Brain-computer interfaces (BCIs)** for seamless control of peripheral neuroprostheses, functional electric stimulation devices, and exoskeletons.

(viii) **Social interaction** that understands human social dynamics and moral norms and that can be truly integrated with our social life showing empathy and natural social behaviors.

(ix) **Medical robotics** with increasing levels of autonomy but with due consideration of legal, ethical, and technical challenges, as well as microrobotics tackling real demands in medicine.

(x) **Ethics and security** for responsible innovation in robotics.

The field of robotics is broad and covers many underpinning and allied technological areas. The identification of these challenges was a difficult task, and there are many sub-topics not listed that are equally important to future development. The above list is therefore neither exclusive nor exhaustive.

One of the ambitions of *Science Robotics* is to deeply root robotics research in science

¹Hamlyn Centre for Robotic Surgery, Imperial College London, London, UK. ²Center for Marine Robotics, Woods Hole Oceanographic Institution, Woods Hole, MA 02543, USA. ³Department of Cardiovascular Surgery, Boston Children's Hospital, Harvard Medical School, Boston, MA 02115, USA. ⁴Institute of Physical Chemistry, University of Stuttgart, Stuttgart, Germany. ⁵Micro, Nano, and Molecular Systems Laboratory, Max Planck Institute for Intelligent Systems, Stuttgart, Germany. ⁶Centre for Practical Ethics, Faculty of Philosophy, University of Oxford, Oxford, UK. ⁷Digital Ethics Lab, Oxford Internet Institute, University of Oxford, Oxford, UK. ⁸Department of Computer Science, University of Oxford, Oxford, UK. ⁹Data Ethics Group, Alan Turing Institute, London, UK. ¹⁰Department of Economics, American University, Washington, DC 20016, USA. ¹¹Department of Integrative Biology, University of California, Berkeley, Berkeley, CA 94720, USA. ¹²Singularity University, NASA Research Park, Moffett Field, CA 94035, USA. ¹³MediaX, Stanford University, Stanford, CA 94305, USA. ¹⁴Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA 19104, USA. ¹⁵National Academy of Sciences, Washington, DC 20418, USA. ¹⁶Institute of Robotics and Intelligent Systems, Department of Mechanical and Process Engineering, ETH Zürich, Zurich, Switzerland. ¹⁷Department of Computer Science, Yale University, New Haven, CT 06520, USA. ¹⁸Department Mechanical Engineering and Materials Science, Yale University, New Haven, CT 06520, USA. ¹⁹Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218, USA. ²⁰Machine Learning Department, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA. ²¹School of Materials Science and Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA. ²²John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA. ²³Wyss Institute for Biologically Inspired Engineering, Harvard University, Cambridge, MA 02138, USA.

*Corresponding author. Email: g.z.yang@imperial.ac.uk




Fig. 1. Ten grand challenges of Science Robotics.

while developing novel robotic platforms that will enable new scientific discoveries. Of the 10 grand challenges listed here, the first seven represent underpinning technologies that have a wider impact on all application areas of robotics. For the next two challenges, we have included social robotics and medical robotics as application-specific areas of development to highlight the substantial societal and health impacts that they will bring. Finally, the last challenge is related to responsible innovation and how ethics and security should be carefully considered as we develop the technology further.

CURRENT STATE OF THE ART AND 10 GRAND CHALLENGES

New materials and fabrication schemes
Gears, motors, and electromechanical actuators are fundamental to many of the robotic platforms in use today, but laboratories around the world have begun to explore new materials including artificial muscles (11), compliant materials for soft robots (12), and emerging advanced manufacturing and assembly strategies (13). As illustrated in Fig. 2, these promise a new generation of robots that are power-efficient, multifunctional, compliant, and autonomous in ways that are similar to biological organisms. However, most demonstrations using new materials and fabrication strategies have been “one-offs” and must still overcome basic hurdles to achieve wide-scale adoption. These hurdles include improved portable energy storage and harvesting, new materials with

tunable properties, and new fabrication strategies to embody these functional materials as new capabilities for future robots, including the robot building and repairing itself.

New materials that combine sensing and actuation challenge the physical limitations of traditional mechatronic systems and offer a range of opportunities for the design of new robots (14). Many of the design principles draw inspiration from nature. In vertebrates, one finds a wide range of material properties from soft tissue to bone—over seven orders of magnitude in modulus—that is mediated by a continuous gradient of compliance. As opposed to the more “nuts-and-bolts” assembly approaches currently used to combine basic components into complete robots, a seamless integration of dissimilar material properties (e.g., rigid with soft, conductive with dielectric, etc.), spatially patterned with resolution several orders of magnitude smaller than the characteristic dimension of the robot, could obviate the need for complex assembly and lead to distributed function.

Similar to functionally graded materials, multifunctional materials can increase the efficiency of robot design and simultaneously offer distributed networks of hierarchically structured sensors and actuators. Opportunities exist to leverage breakthroughs in folding-based metamaterials that have demonstrated tunable electromagnetic (15) or mechanical (16) properties beyond what is possible with the base material itself. Similarly, multiphase composites may be used for simultaneous fluidic actuation or sensing (17, 18). Textiles are

a promising material for soft and wearable robotics, generating significant interest in embedding electrical functionality into fabrics. Finally, bidirectional transducers can enable sensors and actuators to behave as materials for energy harvesting or storage. While developing new materials for the future of robotics, it will be important to consider the biodegradability issues or as part of the circular economy paradigm to ensure their eco-sustainability. This is particularly relevant given the ubiquitous nature of robotic platforms in future (19, 20).

Fabrication and assembly is typically a serial process that is slow and difficult to scale to very large or very small scales. The 2016 Nobel Prize in Chemistry was awarded to three pioneers in the field of mechanochemistry who created the first synthetic molecular machines. A major remaining challenge that has thus far not been realized, despite Feynman’s prophecy (21), is to develop materials by integrating these molecular machines, or other force-generating molecules or biological motor proteins, into hierarchical materials. Substantial opportunity exists in the convergence of additive and subtractive methods, with emerging technologies involving two-dimensional (2D) to 3D transformations to generate new architectures that can simplify the need for specialized hardware and enhance the robot’s function. For example, 3D printing (or similar techniques such as multiphoton lithography or selective laser sintering) can create features and structures over nine orders of magnitude in size. However, there is no single technique

or machine that can cover this range—the best additive manufacturing strategy covers roughly three or four orders of magnitude in scale range—and none offers more than a handful of materials choices. Alternative methods should be explored that combine techniques from micro-/nanoscale fabrication (e.g., surface and bulk micromachining; physical and chemical deposition; and microscale molding, stamping, and functionalization used in soft lithography), mesoscale methods such as layering and lamination common in multilayer printed circuit boards, and the myriad macroscale multi-axis subtractive methods. Another challenge that requires much more investigation is the development of multiscalar materials able to adapt and heal over time, thus providing 4D robots that achieve the complexity found in natural systems (22).

Bioinspired and biohybrid robots

As human technologies take on more of the characteristics of nature, nature becomes a more useful teacher (23).

By bioinspired robotics, we mean the use of fundamental biological principles that are translated into engineering design rules for creating a robot that performs like a natural system. If the biological understanding results in the direct use of biological material to design synthetic machines, then we refer to this as a biohybrid robot. Specific grand challenge lists for biorobotics have remained largely unchanged for the past 30 years—a battery to match metabolic conversion, muscle-like actuators, self-healing material that manufactures itself, autonomy in any environment, human-like perception, and, ultimately, computation and reasoning. For recent progress on these and other specific challenges, we refer readers to a few of the many outstanding perspectives and reviews (4, 24–31). Here, we identify major goals that, if met, would accelerate the design and implementation of bioinspired and biohybrid robots at an unprecedented pace.

Major challenges remain for nearly all component technologies (Fig. 3) that could enable bioinspired behavior. Materials that couple sensing, actuation, computation, and




Fig. 2. Multifunctional materials. New materials and fabrication schemes promise a new generation of robots that are power-efficient, multifunctional, compliant, and autonomous in ways that are similar to biological organisms.

communication are critical and must be shared as developed (32). Novel designs of heterogeneous, anisotropic, hierarchical, multifunctional materials have used differing designs of structural elements to increase material strength, stiffness, and flexibility; fracture toughness; wear resistance; and energy absorption (33). These advances promise to provide robots with features such as body support, weight reduction, impact protection, morphological computation, and mobility. Techniques newly available to fabricate architectures at the micro-, meso-, and macroscales include recombinant technologies, biomineralization, layer-by-layer deposition, ori- and kirigami, self-assembly, bio-templating, magnetic manipulation, freeze-casting, vacuum-casting, extrusion and roll compaction, laser engraving, additive manufacturing, actuator-embedded molding, and soft lithography (33).

For biohybrid and bioinspired robots, actuation and energy remain major bottlenecks compared with performance seen in animals (34). Electromagnetic motors are adequate actuators for large robots but inefficient at small scales or in soft systems. New artificial muscles could revolutionize bioinspired robots; current versions that have muscle-like function and operate by shrinkage or expansion of material—such as shape-memory materials and electro-active polymers—lack robustness, efficiency, and energy and power

density. No battery can yet match metabolic energy generation in organisms, so highly miniaturized, biohybrid robots actuated by biological muscle become advantageous (28). Biohybrid robots can exploit the unique features of living cells that include self-healing (35), embedded sensors, dynamic response to changing environmental conditions, and use of inexpensive and eco-friendly fuel (28).

A major challenge remains as to how these components are effectively integrated and embodied to perform system-level behaviors (Fig. 3). The field of bioinspired robotics must address different challenges, mainly due to the synthesis/fabrication of efficient and scalable artificial components. However, biology has made progress toward providing principles, especially for mobility and manipulation. New discoveries in hydro-, aero-, and terradynamics have led to an impressive “robo-zoo” of bio-

inspired robots (24, 25) benefiting from the nonlinear, unsteady, self-stabilizing, energy storage, and return principles quantified in animals. Further development is required to understand transitions and multimodal performance (36) within the same platform. Significant progress has been made in bioinspired, quasi-static, pick-and-place manipulation, and grasping, but no system has integrated components sufficiently to match the flexibility and dexterity of human hands (37).

As bioinspired robots venture beyond the laboratory, models of real-world, unstructured environments will be required, but none can yet adequately represent our complex, dynamic world. Although first-principle models exist for hydro- and aerodynamic systems (i.e., the Navier-Stokes equations), a similar framework for terradynamics (38) is required to understand how bioinspired robots effectively interact with the ground. Because of their staggering complexity, one of the greatest challenges to extracting fundamental principles from biological systems involves model abstraction and dimensional reduction (39). Internal models can allow us to test hypotheses and simplify control, especially when placed into a dynamical systems theory framework (40). These models become even more important as we require simple representations for use in reinforcement, supervised,

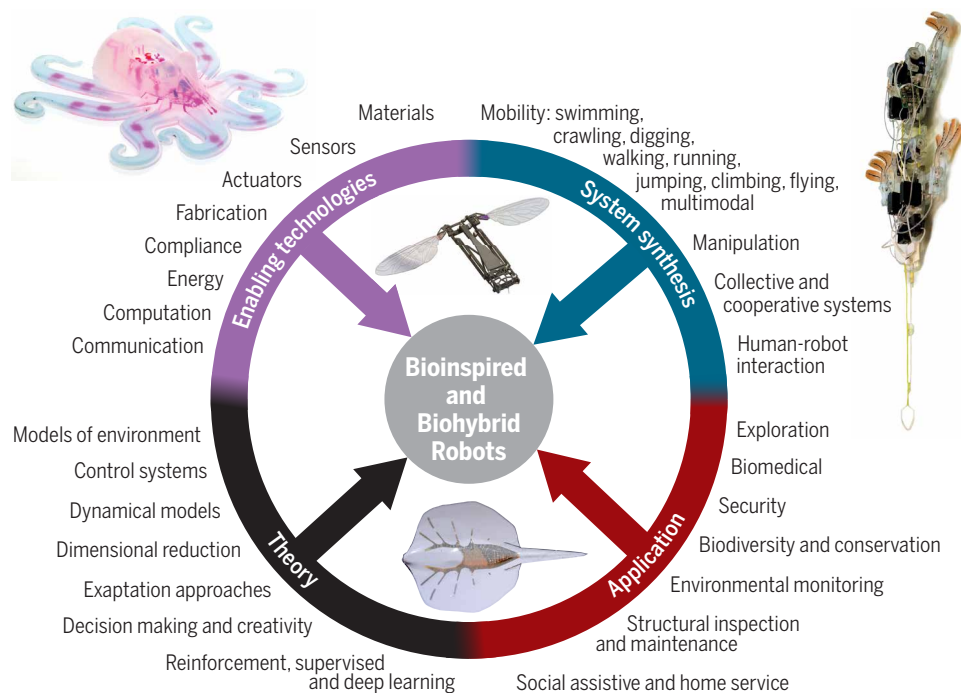


Fig. 3. Convergence of conditions accelerating opportunities for design of bioinspired and biohybrid robots. Enabling technologies, development of theory, a synthesis of systems, and application drivers all provide a foundation for a frontier. [Adapted by N. Cary/*Science Robotics*. Image credits: Octobot (4); RoboBee, Wyss Institute at Harvard University; StickyBot, PolyPEDAL Laboratory-Pauline Jennings, robot from M. Cutkosky, Stanford University; ray (105)].

and deep learning for adaptability, decision-making, and even creativity.

Power and energy

As for any electronic system, power and energy sources represent one of the most challenging areas of robotics research and deployment, especially for mobile robotics (Fig. 4). Underwater and particularly deep-sea exploration requires compact, stable, high-energy density batteries to support robots working in challenging conditions and extreme environments. The increasing adoption of drones and autonomous vehicles is fueling the development of new battery technologies that are safe and affordable, with longer cycle lives, robust temperature tolerance, higher energy densities, and relatively low weight. Beyond the currently available commercial technologies such as lead-acid, nickel-metal hydride, and lithium-ion batteries, there has already been extensive research on developing next-generation technologies, such as fuel cells and supercapacitors. These new areas include the development of silicon anodes with smart electrodes through conductive nanoporous structures and binder designs, which greatly enhances cyclability and minimizes pulverization. Other emerging electrode designs for achieving enhanced ca-

pacities use Ni-, Li-, and Mn-rich, layered materials (41). Although many new ideas are being developed, the fundamental issues being addressed remain the same for many historical technologies: irreversible phase transitions of active materials and/or unstable electrode-electrolyte solution interfaces (41).

Metal-oxygen, lithium-sulfur, aluminum-ion, and sodium-ion batteries are some of the key technologies being actively pursued. The potential of lithium-sulfur batteries combined with solar panels has already been demonstrated with the Zephyr-6 unmanned aerial vehicle in its record-setting, high-altitude, long-endurance flights (42). Although most battery research is focused on liquid electrolytes because of high ion conductivity and good surface-wetting properties, they often suffer from electrochemical and thermal instabilities, as well as low ion selectivity. Advanced battery systems based on solid electrolytes could bring advantages because of their safety, excellent stability, long cycle lives, and low cost (43). The advent of flow-based, lithium-ion, lithium-sulfur, and lithium-organic batteries also promises new opportunities (44). The future will also see new improvements to the current radioisotope power systems used for space exploration.

In practice, the operational longevity of a mobile, autonomous system is typically dictated by the battery power, its size, and its weight. Efforts continue to minimize power utilization through development of power-efficient electronics and actuators, but for robots to operate wirelessly for appreciable times in unstructured environments, they must extract useful energy from their surroundings and use radical new solutions for highly energy-dense storage, such as solar light, vibration, and mechanical movement. Compared with biological machines at any scale, robots are typically very energy-inefficient [e.g., the 100-horsepower (75 kW) consumption of Boston Dynamic's horse-sized LS3]. Whereas the quintessential robot arm bolted to the factory floor and tethered to an unlimited power supply works well in industrial settings, mobile robots lack a standard fuel source, storage, and distribution system. Batteries, of course, are ubiquitous, although their energy density remains low compared with hydrocarbons (about 1 MJ/kg and 50 MJ/kg, respectively). One benchmark comes from biology, where carbohydrates (about 17 MJ/kg) power the effective running, swimming, and flying of organisms over a huge range of physical scales (45).

Robotics will require a shift in energy storage technologies to produce similar behavior. Electrochemical storage technologies are attractive for numerous reasons, although many autonomous robots leverage combustion (13) or monopropellant decomposition (46) as alternatives.

Developments in energy-harvesting techniques (e.g., mechanical, thermoelectric, photovoltaic, and electrochemical) and wireless power transmission (47) are expected to play a key role in addressing the power and energy challenges of robotics. Different mechanisms have been established for harvesting mechanical energy, including electromagnetic and electrostatic generators, as well as piezoelectric nanogenerators and triboelectric nanogenerators (based on the coupled effect of contact electrification and electrostatic induction) (48). Besides serving as a small power supplies, nanogenerators can be self-powered sensors and flexible actuators with the use of a range of materials from functional polymers, fabrics, and nanomaterials to traditional metal foils and ceramic thin films (49). The most important characteristic of a nanogenerator is its high response to low-frequency mechanical triggering, with complementary applications




Fig. 4. A summary of different energy sources for robotics. Power generators, which include fuel cells, classical electromagnetic generators, and solar cells. Energy storage, including batteries and capacitors/supercapacitors. Power harvesting and newly developed nanogenerators, as micro-/nano-energy sources, self-powered sensors, and flexible transducers.

with an electromagnetic generator that usually works well at high operation frequency. In the working environment of a robotics, low-frequency mechanical stimulation is fairly popular, which can be effectively converted into electric output using a triboelectric nanogenerator.

As stated in the previous section, no battery can yet match metabolic energy generation in organisms. Biohybrid robots could use the unique features of living cells for potential solutions (28).

Robot swarms

Robot swarms allow simpler, less expensive, modular robotic units to be reconfigured into a team (Fig. 5) depending on the task at hand while being as effective as a larger, task-specific, monolithic robot, which may be more expensive and have to be rebuilt depending on the task. Nature provides a repertoire of examples that illustrate this idea (50). Independently acting organisms cannot achieve a goal by themselves but, in coordination with other organisms, can solve complex problems and complete a mission. This “force multiplication” requires individuals to sense not only the environment but also their neighbors and to communicate with other individuals in their team while acting independently. This paradigm has been seen in fish, birds, and insects and is

fundamental to navigating as a flock or horde, foraging, hunting, building nests, and surviving harsh environments. Similarly, a group of relatively unsophisticated robots can form a networked team that realizes a range of behaviors well beyond the capabilities of the individuals by communicating and cooperating with team members. The swarm principle can be used at macro-, micro-, and nanoscales with a plethora of application areas.

There are three technology drivers suggesting that robot swarms will have an impact in the next 5 to 10 years that stem from falling prices and increasing performances of sensors, processors, storage devices, and communications hardware. First, the integration of components for computation and storage is resulting in a software-centric architecture that tightly couples computation, storage, networking, and virtualization resources—a framework that is being called “hyper-convergence” (51). Soon, sensors and wireless communication devices will be part of this hyper-convergence. Second, we are seeing the convergence of the hardware for consumer electronics (smart phones, tablets, and virtual reality devices) and intelligent autonomous systems (drones, robots, and self-driving vehicles), with concurrent advances in 5G wireless technologies. Third, cognitive systems relying on statistical machine learning and AI are becoming

mainstream. Tools from data science, machine learning, and predictive analytics are now being routinely used to extract information from text and speech and to recognize objects from imagery (pictures and videos).

As we think about swarms, it is useful to consider different forms of collective behavior. Coordination and cooperation can be seen in groups that are homogeneous, but heterogeneity is powerful because it allows for collaboration (52, 53). For example, a large robot may be able to carry more powerful sensors or have more powerful computational resources or radios, but it may be less agile than its smaller counterparts. Scale is particularly important in robot swarms where small groups lend themselves to centralized control, and information across the group can, in principle, be shared via communication and sensing. The analysis of group behaviors in these settings or the synthesis of group behaviors for a given task is easier for smaller groups with centralized architecture than for larger groups like swarms, where it is impractical to efficiently share information across the swarm and architecture because these systems are necessarily decentralized. From a mathematical perspective, the state space, which is the Cartesian product of the individual state spaces, grows linearly, and the types of interactions that can occur across individuals grow combinatorially. Thus, it is necessary to develop stochastic models for predicting collective behavior in large-scale swarms. However, we lack mathematical models of flock- or herd-like groups that elude the enumeration in small-scale groups yet do not justify ensemble-averaged models seen in large-scale swarms.

Robotic systems are equipped with sensors that allow them to perceive the environment. They reason about the environment and take actions, forming a feedback loop that is called a perception-action loop. Designing perception-action loops is fundamental to creating autonomous robots that function in unstructured environments. Robot swarms require their communication ability to be embedded in this feedback loop. Thus, perception-action-communication loops are key to designing multifunctional, adaptive robot swarms. There are currently no systematic approaches for designing such multidimensional feedback loops across large groups.

Whether we think of smaller robot groups, in which the combinatorics do not pose formidable challenges, or larger swarms, much of the literature addressing the problem of coordination makes use of simpler mathematical models; algorithms for perception, estimation, planning, and control; and robot deployments (54). The dynamics and control of cooperation have been addressed in cooperative manipulation, multi-fingered grasping, and legged locomotion, but systematic approaches to questions of synthesis do not exist. Similarly, although there is interesting work on collaboration between humans and robots (55) and between aerial and ground robots (56), a general framework for modeling heterogeneity and the design of heterogeneous groups and desired behaviors does not exist.

As we develop robot swarms, one must also develop the tools to create teams that can be responsive to human commands, can adapt to changing conditions, are robust to disturbances (to the extent that is possible given the constraints on resources), and are resilient to adversarial, disruptive changes caused by large-scale failures or damage to the swarm infrastructure. Responsiveness is generally characterized by the time a system takes to respond to input or meet input-output (task) specifications. Robustness is the property of the system to be responsive even in the presence of disturbances and modeling errors (and failures), although the majority of the literature addresses robustness with carefully constructed bounds on those disturbances and modeling errors. As pointed out by Rodin (57) in the context of similar challenges that confront urban societies today, resilience is a fundamentally different property that is about systems that can bend without breaking. Resilient systems are self-aware and self-regulating and can recover from large-scale disruptions to the network. Thus, a science of resilient robot swarms must focus beyond robust individual agents to resilient integration across diverse elements of the group that leverage new mechanisms (e.g., mobility, reconfiguration, sensing, communication, platform diversity, and involvement of human peers) for achieving macroscale resilience.

Robot networks integrated with our infrastructure have tremendous potential for solving the most pressing problems facing



Fig. 5. Robot swarms. New opportunities and research challenges.

human civilization. They can provide solutions to feed an ever-increasing population with limited resources by increasing the efficiency of food production and decreasing water consumption by an order of magnitude (58). They can respond to natural disasters and adversarial attacks by enabling resilience in our infrastructure (59). They are a part of any practical solution to space colonization. We are poised to see great advances and impacts in this area in the next 5 to 10 years.

Navigation and exploration

Path planning, obstacle avoidance, localization, and environment mapping are ubiquitous requirements of robot navigation and exploration. Advances in sensing, machine vision, and embedded computation have underpinned the remarkable progress of autonomous vehicles roaming complex terrains at speed, drones forming swarms for completing collaborative tasks, and surgical robots delivering targeted therapy while negotiating complex, delicate anatomical structures. Many robots we deploy are intrinsic explorers that we send to the far reaches of the planet—the deep oceans, under the Arctic ice pack, into volcanoes—and go where no human has yet tread, often under unknown and extreme conditions. The associated challenges are therefore much greater than those encountered today.

Foremost, the robots must operate in environments that are not only unmapped, but, at times, their very nature is not understood. Adding to this are challenges associated with communications and navigation. Robots

in tunnels or mines must cope with rough terrain, narrow passageways, and degraded perception. Robots undertaking nuclear decommissioning must withstand radiation and restricted access, and those used to construct and assemble infrastructure must be able to resist chemicals and materials used in the construction process as well as being resistant to dirt, dust, and large impact forces. Undersea robots operate in an environment where radio does not penetrate and our usual forms of communication and navigation disappear; untethered undersea vehicles must be autonomous. As robotic spacecraft take on tasks like roaming distant planetary surfaces and, in the not-so-distant

future, possibly landing on the icy moons of the outer planets, they enter a regime where long latency and low bandwidths of communications not only greatly reduce productivity but also put the survival of the robot itself at risk.

Undoubtedly, current mapping and navigation techniques will continue to evolve. For example, techniques such as SLAM (simultaneous localization and mapping) will go beyond the current rigid and static assumptions of the world and will effectively deal with time-varying, dynamic environments with deformable objects (60). With resource constraints, specific challenges include how to learn, forget, and associate memories of scene content both qualitatively and semantically, similar to how human perception works; how to surpass purely geometric maps to have semantic understanding of the scenes; how to reason about new concepts and their semantic representations and discover new objects or classes in the environment through learning and active interactions; and how to evolve through online, prospective, and lifelong continuous learning.

For navigation, the grand challenge is to handle failures and being able to adapt, learn, and recover (Fig. 6). For exploration, it is developing the innate abilities to make and recognize new discoveries. From a system perspective, this requires the physical robustness to withstand harsh, changeable environments, rough handling, and complex manipulation. The robots need to have significant levels of autonomy leading to complex self-monitoring, self-reconfiguration, and repair such that there is no single point of complete failure but rather




Fig. 6. Intelligent explorers. Handling failures and being able to adapt, learn, and recover are major challenges for navigation and exploration, especially for robots operating in extreme environments. [Reproduced from (106) with permission].

graceful system degradation. When possible, solutions need to involve control of multiple heterogeneous robots; adaptively coordinate, interface, and use multiple assets; and share information from multiple data sources of variable reliability and accuracy.

AI for robotics

As the underpinning technology for robotics, AI is undergoing a renaissance after more than 60 years of ongoing development. There is a widespread myth that AI did not work for the first 50 years, but the truth is that for certain classes of domain- and task-specific problems, given enough development time as well as computing and data resources, the applications could be made to work. The advent of deep learning methods resulted in remarkable levels of object recognition accuracy (61) using hierarchical pattern recognition that retained information coherence at each level of the hierarchy. The new machine-learning algorithms were combined with unprecedented access to data, as well as inexpensive and powerful computing hardware. The resulting progress in solving narrow classes of AI problems has led many to think that we are on the verge of solving intelligence—in all its multifaceted and (still) poorly understood dimensions.

However, we still have a long way to go to replicate and exceed all the facets of intelligence that we see in humans. Combining advanced pattern recognition and model-based reasoning is critical for building systems that can go beyond statistical correlation and

begin to reason about underlying interdisciplinary mechanisms and systems dynamics. Meta-learning, or learning how to learn new things, is a critical new AI capability not only with large training data sets but also with limited data. The challenge is to be able to learn on the fly, adapting to dynamic and uncertain environments. One promising approach in this area has been developed based on neuroscience insights about the human hippocampus as a predictive map of novel situations (62).

AI systems that know their own limitations and know how to seek help could go beyond the current methods of training and knowledge acquisition. These systems will know how to interact, how to seek help, how to recover from failure, and how to become smarter. AI systems and robots that can model their own components and operations are critical for adaptation and evolution. We need AIs that are able to detect their own subcomponents, model their operations, and modify those models if their structure changes. Work by Bongard *et al.* (63) provides an early example of this type of robotic system, which can discover its own components and learn to use them dynamically in locomotion.

AI that can learn complex tasks on its own and with a minimum of initial training data will prove critical for next-generation systems. Most machine-learning systems are data-intensive and require massive data in order to learn complex tasks. DeepMind's AlphaGo Zero system that taught itself to play Go significantly better than the world champion in Go (64) was an impressive example of this.

However, we do not yet have systems that can do this easily across heterogeneous tasks and domains. AI systems that can comprehend deeply and synthesize across complex texts and narratives will prove useful in a variety of applications. We have already seen some initial examples, but the real world is both interdisciplinary and complex, and building robust systems of this class will prove extremely challenging.

One of the enduring grand challenges in AI is to provide a coherent and comprehensive mapping of the key mechanisms of human intelligence in a software system. The first key step in doing this is to produce a thorough account of how the neocortex actually works, including learning to learn and learning from limited examples. A recent paper on this provides some detailed and testable predictions concerning how columns in the neocortex provide location signals that enable learning the structure of the world (65). We need to test theories of this type rigorously, both in terms of neuroscience data and in the operation of AI software (66). In addition, much progress has been made recently in building AI systems that understand natural language. A key set of targets is to build systems that maintain coherent conversations and deal with unknown environments and contexts.

Ambient intelligence and ubiquitous and networked AI and robotics (cloud robotics) will be critical in the development of integrated heterogeneous AI and robotic services. There are many initial examples of cloud AIs that update situation assessments and share knowledge but few working examples of heterogeneous AI or robotic services that integrate smoothly and reliably over time. DeepMind's PathNet architecture points to systems that allow for new contexts to be learned at the same time, leveraging knowledge of training in other contexts to learn much faster.

One of the big questions for AI is its ability to perform deep moral and social reasoning about real-world problems. As AI and robotic systems undergo accelerating growth in power and capabilities, there will be an increasing premium on systems that can demonstrate moral and social reasoning. Although human-in-the-loop may be a preferred design constraint for systems that touch life-or-death situations, in autonomous driving and aerospace applications, the relevant decision loops may well be too fast for the human brain, hence the need for embedded moral and

social reasoning. These challenges need to be framed in the context of baseline risks that humans have already habituated to, such as 1.2 million people dying worldwide as a result of largely avoidable driver errors committed by humans. We can expect to see considerable and rapid operational progress on this front.

Brain-computer interfaces

A BCI forges a direct, online communication between brain and machine, independent from the user's physical abilities, and represents a new way to augment human capabilities and restore patient function (Fig. 7). Direct use of brain activity to control a computer or external device without the mediation of the peripheral, somatomotor nervous system has major applications in enabling paralyzed patients to communicate and control robotic prosthetics and in rehabilitation for restoring neural function (67–71). BCIs translate the user's intentions into outputs or actions by means of machine-learning techniques, operating by either presenting a stimulus to the operator and waiting for his/her response (synchronous) or continuously monitoring the operator's cognitive activity and responding accordingly (asynchronous). Beyond their clinical use, BCIs also have emerging applications in neuroergonomics, communication and control, education and self-regulation, as well as games and entertainment (72). Despite being a relatively new field, recent advances in BCIs have been accelerated by allied technologies, including neuroscience, sensor technologies and com-

ponent miniaturization, biocompatibility of materials, and embedded computing.

For practical use, a BCI can be classified as active, reactive, or passive (73). Active BCI derives its outputs from brain activity, which is directly and consciously controlled by the user, not necessarily depending on external events, for controlling an application. In reactive BCI, the outputs are derived from brain activities arising in response to specific external stimuli. Passive BCI is a relatively newer concept, which derives its outputs from arbitrary brain activity arising without the purpose of voluntary control, for enriching human-machine interaction with implicit information on the actual user state.

Both invasive and noninvasive methods are used to record brain activity. Invasive approaches measure the neural activities of the brain by either intracortical neural interfaces with microelectrode arrays, which capture spike signals and local field potentials, or cortical surface electrocorticography, providing both high temporal and spatial resolution with good immunity to artifacts (70). Noninvasive BCIs require no surgical implantation; typical signals used include slow cortical potentials, sensorimotor rhythms, P300 event-related potentials, steady-state visual evoked potentials, error-related negative evoked potentials, blood oxygenation levels, and cerebral hemodynamic changes. Common assessment methods include fMRI (functional magnetic resonance imaging), fNIRS (functional near-infrared spectroscopy), MEG (magnetoencephalography), and EEGs (electroencephalograms) (70).

Despite the success of BCI for patients with amyotrophic lateral sclerosis (also known as motor neuron disease), spinal cord injury, and rehabilitation of motor function after stroke, there remain significant challenges for the wider adoption of BCI (74). The first is in sensing and data acquisition because current modalities are expensive and cumbersome. Parallel developments in implantable sensing with new microfabrication, packaging, and flexible electronics, combined with ultralow-power local processing and wireless data paths, would bring new opportunities for completely untethered implants, providing improved patient experience and uptake in both clinical and home environments. For noninvasive techniques, newly emerging, low-cost, and ergonomically designed wireless EEG and fNIRS systems have shown promise for general BCI-based robotic control.

The second challenge is in data processing and dealing with artifacts of noncerebral origin, particularly for wearable approaches. The data-processing challenge is also associated with the fact that cortex folding differs between individuals, as do relevant functional maps. Furthermore, sensor locations may differ across different recording sessions, and brain dynamics can be intrinsically nonstationary. Current methods often involve extended periods of training, calibration, learning, and adaptation, thus making it prohibitive for general use.

Third, it remains to be seen whether BCI will always outperform simpler techniques, such as those using eye tracking or muscle-based devices. The development of hybrid BCIs may represent a viable way forward by combining with other, more mature assistive technologies. This would allow more reliable and seamless interfacing with peripheral neuroprostheses, functional electric stimulation devices, and exoskeletons.

A further challenge is dealing with tasks with high degrees of freedom. Current multiclass BCI classification generalizes poorly across individuals and tasks. In such cases, it may be more appropriate to rely on BCI for intention detection and task initiation and on autonomous robot manipulation for task completion.

Continuing development of BCIs will bring exciting new research opportunities not only in robot control and functional rehabilitation but also in knowledge exchange and cross-fertilization between neuroscience and robotics. It will also play an important role in

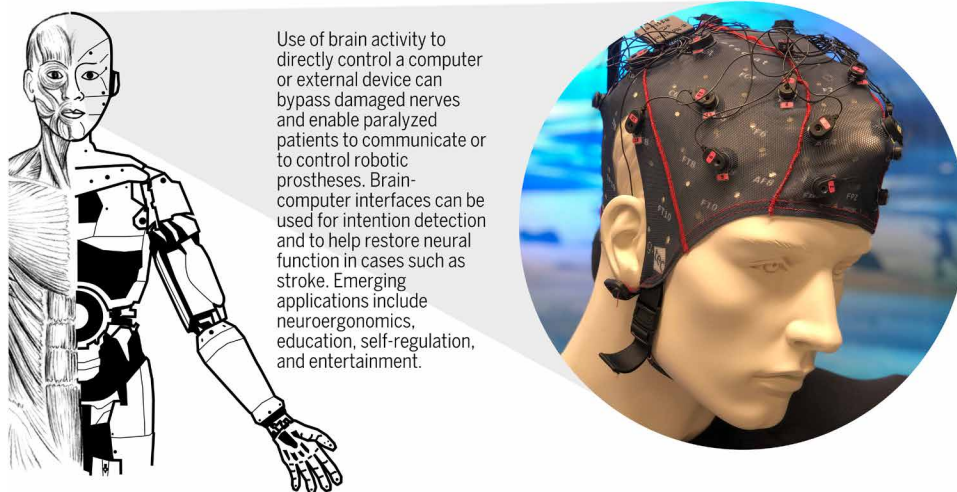


Fig. 7. Brain-computer interfaces. BCIs have extensive applications in enabling paralyzed patients to communicate with and control robotic prosthetics and in rehabilitation for restoring neural function. Continuing development of BCIs will also see applications in performing mission- or safety-critical tasks.

performing mission- or safety-critical tasks, whereby adaptive levels of automation, context-sensitive decision support, and motion constraints are provided depending on mental workload, task engagement, hypovigilance, mood or emotion, and precursors to human errors (e.g., hesitation and disorientation).

Social interaction

Robotics and AI have often underestimated the difficulty of replicating capabilities that humans find particularly easy. Perhaps most notorious was the early belief that computer vision was a simple problem suitable for an undergraduate research project (what could be simpler than seeing a table as a table?) (75), but similar stories can be told for locomotion, manipulation, and understanding language. Social interaction has the same status: Because humans are so adept at recognizing and interpreting social behavior, we often underestimate the complexity of the challenge that this represents for a robot (Fig. 8).

As common as social interactions are in our daily lives, we have very few comprehensive, quantitative analyses of human social responses; our understanding of human social behavior is not nearly as advanced as our knowledge of Newtonian mechanics or even human visual perception. Although this alone might make some believe that building social interactions for robots is premature, the practicality of putting robots into our human environments—our schools, hospitals, shops, and homes—necessitates addressing social interaction. The three most significant challenges that stem from building robots that interact socially with people are modeling social dynamics, learning social and moral norms, and building a robotic theory of mind.

Social interaction is a major challenge for robotics in part because the perceptual demands are so significant. Social cues—such as gaze direction, facial expressions, or vocal intonation—are often extremely detailed, rapid, and nuanced signals that are embedded within other activity; the difference between an enthusiastic greeting and a sarcastic scolding might depend on a single wink, or rising inflection on just one phoneme. The temporal patterning of these signals is also frequently significant—a small delay when answering a question may be inter-




Fig. 8. Social robotics. Social interaction requires building and maintaining complex models of people, including their knowledge, beliefs, goals, desires, and emotions.

preted as a sign of uncertainty or mistrust. Although we have made substantial advances in machine perception in the last decade, especially in object recognition (76), action recognition (77), and human gaze analysis (78), we still lack systems that operate under the diverse natural conditions and real-world time constraints that social interactions demand. Next-generation systems will need to richly mix elements from multiple input modalities and combine these perceptual systems with predictive models of social intention to more fully capture the rich, dynamic nature of social interactions.

Social signals are also very context-dependent and culturally determined. Two individuals standing nearly nose to nose in a conversation might be typical in Argentina, but could be an indication that they are either close friends or about to have an argument in the United States. Robots that are deployed in human environments must be able not only to adapt to these cultural differences but also to learn the appropriate social and moral norms for their setting. A robot that expresses excitement when the death of a family member is being discussed, one that shouts at inappropriate times, or one that takes a coffee mug before it is empty will not find itself welcome in home or workplace. The development of robots capable of understanding empathy, ownership, and the need to keep a promise will be essential to building the long-term trust and relationships necessary for operating side by side with people. To take the next step in this domain, new tools are required for modeling the expectations of the people around the robot and expanding the robot's understanding of the consequences of its own actions.

Social interaction also requires building and maintaining complex models of people, including their knowledge, beliefs, goals, desires, and emotions. We routinely simplify our language based on what we know our partners understand, coordinate our actions to match the preferences of our collaborators, and interpret the actions of others as representing their inner goals. These “hidden” states allow us not only to understand why someone has taken a particular action but also to predict their likely future behavior. Modern work on intent recognition (79), user modeling in intelligent tutoring systems

(80), collaboration models in human-machine interaction (81), emotion recognition via facial feature analysis (82), and other domains touch on single aspects of this problem, but none of these domains has yet produced comprehensive or integrated models that allow robots to begin to have rich, usable models of human mental states (83). Advancing the state of the art in this domain will require integration of models of episodic memory, hierarchical models of tasks and goals, and robust models of emotion to create detailed cognitive models that capture the naive psychology that we effortlessly apply to understanding human behavior.

Solutions should also work for long-term interactions and relationships: A joke told once might be funny, but the same joke told every day for a month is not. Most of our current social robots have been designed for interaction that lasts on the order of a few minutes or hours, but our human-human social interactions often span months, years, and even decades. Just as machine learning struggles to scale to continuous, long-term adaptation models (84), social robotics must expand from moment-to-moment engagements to long-term relationships. This expansion will require models of robot behavior and personality that distinguish between changes that are appropriate at different time scales, the capability to autonomously generate interaction content (both verbal and nonverbal) rather than relying on prescribed components, and personalization and adaptation mechanisms that adjust both short-term responses and long-term tendencies based on current interactions.

Medical robotics

From minimally invasive surgery, targeted therapy, hospital optimization, to emergency response, prosthetics, and home assistance, medical robotics represents one of the fastest growing sectors in the medical devices industry (85).

The impact of robotics on medicine is undeniable. The therapeutic and commercial success of Intuitive Surgical's da Vinci system has spurred a number of commercial ventures targeting surgical applications, which echo the emerging trend in precision surgery, focusing on early malignancies with minimally invasive intervention and greater consideration of patient recovery and quality of life (86, 87). These efforts will continue to improve healthcare in terms of both outcomes and cost. Other research and commercial efforts are focusing on what many see as an inevitable future in which intelligent robotic devices assist healthcare workers in a variety of ways. As we move toward this future, however, many grand challenges remain. One of the primary challenges in surgical and interventional robotics is a move toward systems that exhibit increasingly higher degrees of autonomy (85). A second grand challenge is the creation of fully implantable robots that replace, restore, or enhance physiological processes. A third grand challenge is in the realization of micro- and nanorobotic devices of clinical relevance (Fig. 9).

In those industries in which robots are most successful (e.g., manufacturing and warehouse

automation), teleoperation has been replaced by semiautonomous or autonomous operation. Autonomy in medical robotics is incredibly challenging (88); whereas products and assembly lines can be designed to fit the capabilities of robots, this is not possible with the human body. Consequently, autonomy in existing medical robots remains limited. In most cases, the contribution of the robot has been to enhance the skill level of the surgeon. For example, Intuitive Surgical's da Vinci robot makes laparoscopy easy (89); routine procedures can be performed at a higher level of proficiency, and difficult cases that would otherwise be treated with open surgery can be performed laparoscopically. Similarly, Stryker's Mako robotic arm enhances hip and knee replacement by enabling more precise bone drilling than the surgeon can perform on his or her own. In both these examples, the robot acts as an extension of the surgeon's hand, and its motion is continuously under the surgeon's control. Other systems, such as Think Surgical's Robodoc system, execute precomputed and surgeon-approved cutting paths based on medical images. All these systems exercise some degree of "autonomy" in translating a surgeon's intentions (expressed in joystick motions or in preoperative planning) into the actual motions of the robot's actuators. The challenge arises when the controller needs to make more complex decisions in interpreting the clinician's intentions. Thus, we anticipate that the development of autonomy in medicine satisfying regulatory and ethical concerns will progress in stages. Two examples are described below.

Although medical robot autonomy is often discussed within the context of surgery, emergency medicine provides another set of challenges and opportunities. In this case, an emergency medical technician (EMT) needs to assess the condition of a patient quickly, prioritize problems, and often take time-urgent steps to stabilize the patient. Intelligent robotic systems that could assist with such tasks as placing and monitoring sensors, inserting intravenous lines or breathing tubes, and preparing a patient for transport could significantly improve the ability of an EMT to provide urgent care. In addition to obvious challenges in dexterity and device development, there are also difficult computational challenges. The robot assistant will need to recognize relevant patient anatomy in what is often a highly unstructured environment. It will need to use its situational understanding to perform tasks appropriately under direction of the EMT, who is likely to rely primarily on spoken commands, supplemented with gestures, to explain what needs to be done.

A long-term challenge is to enable one surgeon to supervise a set of robots that can perform routine procedure steps autonomously and only call on the surgeon to take control during critical, patient-specific steps. For example, intracardiac interventions involve navigating through percutaneous entry in the peripheral vasculature to specific locations inside the heart using a combination of pre- and intraoperative imaging. The theory of image-based robot navigation is well developed, so developing safe navigation algorithms seems quite feasible. As clinical experience with intracardiac devices (e.g., transcatheter valves) grows, the deployment of these devices may become sufficiently standardized to enable automated deployment. Furthermore, miniaturized and multifunctional fully implantable robots represent an emerging area of development (90, 91). Issues related to biocompatibility, packaging, power efficiency, and harvesting are important to be addressed (92). Perhaps the most significant challenge of automating any clinical task is to be able to anticipate, detect, and respond to all possible failure modes. Medical device regulation of autonomous robots will likely need to develop in a manner that balances the requirements for provably safe algorithms with compliance costs.

An emerging area of medical robotics is implantable robotic devices. These




Fig. 9. Medical robotics. From macro to micro and from large systems to small, smarter devices that can support the future development of precision medicine.

bionic systems are being proposed as replacement organs, e.g., for the pancreas (91); as assist devices for damaged organs, e.g., for the heart (90); and to induce organ growth, e.g., of the esophagus and bowel (93). There are a number of challenges that must be addressed to advance this field. These include biocompatibility, reliability, adaptability, security, and providing power. Full biocompatibility is important in order to maintain long-term functionality. Furthermore, for those implants that provide temporary physiological support, designing the implant to be resorbable could eliminate the need for surgery to remove the device. Implants must also be designed to react to changing conditions, such as exercise, and extreme reliability is a necessity because malfunction could quickly lead to death. Although remote programming to provide software updates is advantageous, security is critically important to prevent one's organ from being hacked. Last, because the power requirements of a robotic device are high in comparison to, e.g., a pacemaker, the capability for wireless power transfer will be crucial.

An other emerging area of medical robotics is micro- and nanorobotics, with increasing numbers of groups maintaining high-profile research efforts. The field has made impressive strides over the past decade as researchers have created a variety of small devices capable of locomotion within liquid environments (94). Robust fabrication techniques have been developed, some devices have been functionalized for potential applications (95), and therapies are being actively considered (96). Although excitement remains high for this field, it faces a number of significant challenges that must be addressed head-on to make continued progress toward clinical relevance. The primary roadblocks to overcome include the development of biodegradable and noncytotoxic microrobots, development of autonomous devices capable of self-directed targeting, catheter-based delivery of microrobots near the target, tracking and control of swarms of devices in vivo, and the pursuit of clinically relevant therapies.

Robot ethics and security

With increasing levels of autonomy and human-robot interaction, there needs to be careful consideration of potential regulatory, ethical, and legal barriers and the context of how robots are deployed. Because robotics and AI are fueled by data, some challenges are rooted in human-environment interactions and data governance (97), especially consent, discrimination, fair-




Fig. 10. Ethical and security risks of robotics and AI developments.

ness, ownership, privacy, surveillance, and trust (98). In terms of ethics, robotics and AI pose five increasingly pressing topics (Fig. 10).

First, excessive reliance on robotics and AI may lead to the delegation of sensitive tasks to autonomous systems that should remain at least partly subject to human supervision, either “in the loop” for monitoring purposes or “post-loop” for redressing. Thus, it is problematic that the European Union (EU) General Data Protection Regulation does not include an explicit right to an explanation when decisions affecting people are reached “solely” algorithmically (99).

Second, robotics and AI may de-responsibilize people whenever an autonomous system could be blamed for a failure. A recent EU proposal to treat forms of AI as “electronic persons” would only exacerbate this problem. Instead, new forms of distributed responsibility need to be developed, learning from the legal analysis of strict liability (100).

Third, unemployment is an ethical problem, not just an economic one. Robotics and AI could change the workforce structure, cause a shift in the skills base, and potentially facilitate a complete de-skilling of the work force even in safety-critical contexts; however, this would be imprudent. Radiologists need to keep studying images for the same reason pilots need to keep landing airplanes so that they still can even if the AI cannot, or if the AI gets it wrong. According to a recent report, AI could displace between 400 and 800 million jobs. Fairness dictates sharing the economic benefits of this huge and rapid transformation,

thus lowering inequality, whereas social solidarity should ensure that AI's costs are shouldered by future generations, too, because they will profit enormously from it.

Fourth, AI may erode human freedom, because it may lead to unplanned and unwelcome changes in human behaviors to accommodate the routines that make automation work and people's lives easier. AI's predictive power and relentless nudging, even if unintentional, should foster and not undermine human dignity and self-determination.

Finally, there is straightforward misuse. Strictly speaking, this is not a problem with AI's smart agency, but with the unethical application of AI by those who control it. The issues under this heading refer to “the human use of human beings,” to cite the title of Wiener's farsighted book (101). Examples range from scanning citizens' faces in illiberal regimes to discriminating among applicants for a job or punishing law offenders unfairly. In this case, Kant provides the right antidote: AI should be designed and used to treat every human being always as an end and never only as a means.

In terms of security, AI can improve security by increasing systems' resilience (enduring attacks) and robustness (averting attacks) and combining both with counterthreat strategies. Thanks to its autonomy, fast-paced threat analysis, and decision-making capabilities, AI can enable systems verification and patching and counter incoming threats by exploiting the vulnerabilities of antagonist systems. However, two challenges may hamper AI's potential for security. One is escalation: Robotics and AI can refine strategies and launch more aggressive counteroperations. This may snowball into an intensification of attacks and responses, which, in turn, may threaten key infrastructures of our societies (102). The solution may be to use AI to strengthen deterring strategies and discourage opponents before they attack, rather than mitigating the consequences of successful attacks afterward. The other challenge is lack of control. Pervasive distribution, multiple interactions, and fast-paced execution will make control of AI systems progressively less effective while increasing the risks for unforeseen consequences and errors. Regulations may mitigate the lack of control by ensuring proportionality of responses, the legitimacy of targets, and a higher degree of responsible behavior, but it is crucial to start shaping and enforcing

policies and norms for the use of AI in security as soon as possible while the technology is still nascent.

DISCUSSION AND CONCLUSIONS

The general field of robotics is quickly evolving, which makes the identification of key grand challenges particularly difficult. In this article, we have focused mainly on underpinning technologies that may have wider impacts across different application domains in the next 5 to 10 years. There are, of course, many domain-specific robotics challenges that need to be addressed, such as those related to space and marine sciences, digital architecture and construction, humanoids, human assistance, rehabilitation, agrifood, infrastructure, and robots designed for emergency response and disaster relief. However, truly addressing these grand challenges requires scientists and researchers from many disciplines to form ongoing collaborations.

When Scott, the legendary polar explorer, died of exhaustion in the Antarctic, he and his team were within sight of their supply tent. Their ponies had died early in the expedition, and his team had to pull their heavy sleds across the frozen landscape acting as human pack animals. What did they carry that was so important it could not be left behind? Buried under the canvas of their sled were rocks containing fossils of leaves, showing that the barren Antarctic continent had at one time been much warmer and had once had forests. Although Scott and his team lost the race to be first to the South Pole, they made one of the great discoveries of Antarctic exploration. What is notable, besides their determination to bring back the fossils, is that they recognized their significance. Such is the spirit we should bear in mind while pursuing these challenges: The ability to recognize discoveries as we progress is as important as conquering these academic missions.

Addressing these grand challenges also requires a major cultural shift. For example, to meet the challenges of bioinspired and biohybrid robot design, engineers, physicists, applied mathematicians, and biologists must form mutually beneficial interdisciplinary collaborations. To extract principles, understand a biological design, and use biological material effectively, it is first necessary to understand that evolution is not engineering. Evolution works on the principle of sufficiency, not optimality, and organisms are severely constrained by their complex histories, development, and multifunctionality. Therefore, engineered mim-

icry in the absence of guiding principles is discouraged. Breathtaking progress is being made on relevant grand challenges in organismal biology, but much remains unknown given the complexity and constraints. Biologists should not only share these advances but also reveal how direct, comparative, and phylogenetic experiments using biodiversity are used to extract a principle. Particularly important for robotics is the development of a synergy where biological principles inspire novel robot or component design, and these robots (or their parts) are then used by biologists as physical models to further test hypotheses of biological structure–function relationships. This realization in biology—that bioinspired robots are invaluable physical models for pursuing further advances in understanding structure–function, ecology, neuroethology, etc.—is also found in physics: The term “robophysics” first emerged (103) for the use of robots as tools to study concepts in the terramechanics of locomotion, particularly on complex granular media.

If bioinspired and biohybrid robots are to move beyond proofs of concept and one-off laboratory demonstrations into real-world applications, then we must match robot capability with need while not compromising curiosity-based research. Bioinspired and biohybrid robots will be uniquely situated for exploration, environmental monitoring, biodiversity conservation, structural inspection and maintenance, security, social assistance and home service, and a wide range of biomedical applications. Market estimates forecast that bioinspired designs could account for a substantial part of U.S. and global gross domestic product (GDP) and result in millions of future jobs. If we can meet the grand challenge of developing bioinspired and biohybrid robots—and if we can establish a strong partnership between basic research in bioinspired engineering and industry—then the impact will be felt far beyond consumers and affect many areas of engineering, science, and social science as our human and natural technologies converge.

In this article, we have also highlighted the importance of robot ethics and security. Given the rapid pace of development in robotics and general public concerns, it is timely that this challenge is addressed in synchrony by basic science researchers, engineers, legal professionals and policy makers. Initiatives like AI4People, the IEEE (Institute of Electrical and Electronics Engineers) Global Initiative on Ethics of Autonomous and Intelligent Systems, and the Partnership on Artificial In-

telligence to Benefit People and Society are working on the ethics of robotics and AI.

As with any technological innovation, the advantages of robotics and AI enable us to not do (or do less of) things that we do not want to do, like driving a car, and to do (either at all or better) things that we want to do, like enjoying a safe and secure life. In both cases, robotics and AI can help us tackle the many concrete evils oppressing humanity and our planet, from environmental disasters to financial crises and from crime, terrorism, and war to famine, poverty, ignorance, inequality, and appalling living standards. Robotics and AI can and will help us manage the increasing complexity of our societies, from megacities to industrial production. Yet, the risk remains that we may misuse or underuse robotics and AI. We should be worried about real human ignorance, not fanciful artificial superintelligence. Churchill once said that “we shape our buildings and afterwards our buildings shape us” (104); this applies to robotics and AI as well. We must design and use robotics and AI ethically and securely and do so now. Humans, not technology, are both problem and solution and shall remain so for any foreseeable future.

REFERENCES AND NOTES

1. J. Li, B. E.-F. de Ávila, W. Gao, L. Zhang, J. Wang, Micro/nanorobots for biomedicine: Delivery, surgery, sensing, and detoxification. *Sci. Robot.* **2**, eaam6431 (2017).
2. S. J. Keating, J. C. Leland, L. Cai, N. Oxman, Toward site-specific and self-sufficient robotic fabrication on architectural scales. *Sci. Robot.* **2**, eaam8986 (2017).
3. N. Jacobstein, J. Bellingham, G.-Z. Yang, Robotics for space and marine sciences. *Sci. Robot.* **2**, eaan5594 (2017).
4. C. Laschi, B. Mazzolai, M. Cianchetti, Soft robotics: Technologies and systems pushing the boundaries of robot abilities. *Sci. Robot.* **1**, eaah3690 (2016).
5. E. W. Hawkes, L. H. Blumenschein, J. D. Greer, A. M. Okamura, A soft robot that navigates its environment through growth. *Sci. Robot.* **2**, eaan3028 (2017).
6. S. Terryn, J. Brancart, D. Lefeber, G. Van Assche, B. Vanderborght, Self-healing soft pneumatic robots. *Sci. Robot.* **2**, eaan4268 (2017).
7. D. W. Haldane, M. M. Plecnik, J. K. Yim, R. S. Fearing, Robotic vertical jumping agility via series-elastic power modulation. *Sci. Robot.* **1**, eaag2048 (2016).
8. A. Ramezani, S.-J. Chung, S. Hutchinson, A biomimetic robotic platform to study flight specializations of bats. *Sci. Robot.* **2**, eaal2505 (2017).
9. Y. Wang, X. Yang, Y. Chen, D. K. Wainwright, C. P. Kenaley, Z. Gong, Z. Liu, H. Liu, J. Guan, T. Wang, J. C. Weaver, R. J. Wood, L. Wen, A biorobotic adhesive disc for underwater hitchhiking inspired by the remora suckerfish. *Sci. Robot.* **2**, eaan8072 (2017).
10. L. Zhang, R. Merrifield, A. Deguet, G.-Z. Yang, Powering the world's robots—10 years of ROS. *Sci. Robot.* **2**, eaar1868 (2017).
11. T. Mirfakhrai, J. D. W. Madden, R. H. Baughman, Polymer artificial muscles. *Mater. Today* **10**, 30–38 (2007).

12. C. Majidi, Soft robotics: A perspective—Current trends and prospects for the future. *Soft Robot.* **1**, 5–11 (2013).
13. N. W. Bartlett, M. T. Tolley, J. T. B. Overvelde, J. C. Weaver, B. Mosadegh, K. Bertoldi, G. M. Whitesides, R. J. Wood, A 3D-printed, functionally graded soft robot powered by combustion. *Science* **349**, 161–165 (2015).
14. G.-Z. Yang, P. Fischer, B. Nelson, New materials for next-generation robots. *Sci. Robot.* **2**, eaap9294 (2017).
15. S. Yao, X. Liu, S. V. Georgakopoulos, M. M. Tentzeris, A novel tunable origami accordion antenna. *Proceedings of the 2014 IEEE Antennas and Propagation Society International Symposium (APSURSI)*, Memphis, TN, 6 to 11 July 2014 (IEEE, 2014).
16. J. L. Silverberg, A. A. Evans, L. McLeod, R. C. Hayward, T. Hull, C. D. Santangelo, I. Cohen, Using origami design principles to fold reprogrammable mechanical metamaterials. *Science* **345**, 647–650 (2014).
17. P. Polygerinos, N. Correll, S. A. Morin, B. Mosadegh, C. D. Onal, K. Petersen, M. Cianchetti, M. T. Tolley, R. F. Shepherd, Soft robotics: Review of fluid-driven intrinsically soft devices; manufacturing, sensing, control, and applications in human-robot interaction. *Adv. Eng. Mater.* **19**, 1700016 (2017).
18. H.-W. Huang, M. S. Sakar, A. J. Petruska, S. Pané, B. J. Nelson, Soft micromachines with programmable motility and morphology. *Nat. Commun.* **7**, 12263 (2016).
19. J. Rossiter, J. Winfield, I. Ieropoulos, Here today, gone tomorrow: Biodegradable soft robots. *Proc. SPIE* **9798**, 97981S (2016).
20. M. Irimia-Vladu, “Green” electronics: Biodegradable and biocompatible materials and devices for sustainable future. *Chem. Soc. Rev.* **43**, 588–610 (2014).
21. R. Feynman, Infinitesimal machinery. *J. Microelectromech. Syst.* **2**, 4–14 (1993).
22. A. S. Gladman, E. A. Matsumoto, R. G. Nuzzo, L. Mahadevan, J. A. Lewis, Biomimetic 4D printing. *Nat. Mater.* **15**, 413–418 (2016).
23. S. Vogel, *Cats’ Paws and Catapults: Mechanical Worlds of Nature and People* (WW Norton & Company, 1998).
24. R. Pfeifer, M. Lungarella, F. Iida, Self-organization, embodiment, and biologically inspired robotics. *Science* **318**, 1088–1093 (2007).
25. A. J. Ijspeert, Biorobotics: Using robots to emulate and investigate agile locomotion. *Science* **346**, 196–203 (2014).
26. S. Kim, C. Laschi, B. Trimmer, Soft robotics: A bioinspired evolution in robotics. *Trends Biotechnol.* **31**, 287–294 (2013).
27. D. Rus, M. T. Tolley, Design, fabrication and control of soft robots. *Nature* **521**, 467–475 (2015).
28. L. Ricotti, B. Trimmer, A. W. Feinberg, R. Raman, K. K. Parker, R. Bashir, M. Sitti, S. Martel, P. Dario, A. Menciassi, Biohybrid actuators for robotics: A review of devices actuated by living cells. *Sci. Robot.* **2**, eaaq0495 (2017).
29. Y. Mengüç, N. Correll, R. Kramer, J. Paik, Will robots be bodies with brains or brains with bodies? *Sci. Robot.* **2**, eaar4527 (2017).
30. R. Raman, R. Bashir, Biomimicry, biofabrication, and biohybrid systems: The emergence and evolution of biological design. *Adv. Healthc. Mater.* **6**, 1700496 (2017).
31. S. Kim, M. Spenko, S. Trujillo, B. Heyneman, V. Mattoli, M. R. Cutkosky, Whole body adhesion: Hierarchical, directional and distributed control of adhesive forces for a climbing robot, 2007 *IEEE International Conference on Robotics and Automation*, Rome, Italy, 10 to 14 April 2007 (IEEE, 2007).
32. M. A. McEvoy, N. Correll, Materials that couple sensing, actuation, computation, and communication. *Science* **347**, 1261689 (2015).
33. S. E. Naleway, M. M. Porter, J. McKittrick, M. A. Meyers, Structural design elements in biological materials: Application to bioinspiration. *Adv. Mater.* **27**, 5455–5476 (2015).
34. M. H. Dickinson, C. T. Farley, R. J. Full, M. A. R. Koehl, R. Kram, S. Lehman, How animals move: An integrative view. *Science* **288**, 100–106 (2000).
35. D. L. Taylor, Self-healing hydrogels. *Adv. Mater.* **28**, 9060–9093 (2016).
36. K. H. Low, T. Hu, S. Mohammed, J. Tangorra, M. Kovac, Perspectives on biologically inspired hybrid and multi-modal locomotion. *Bioinspir. Biomim.* **10**, 020301 (2015).
37. F. J. Valero-Cuevas, M. Santello, On neuromechanical approaches for the study of biological and robotic grasp and manipulation. *J. Neuroeng. Rehabil.* **14**, 101 (2017).
38. C. Li, T. Zhang, D. I. Goldman, A terradynamics of legged locomotion on granular media. *Science* **339**, 1408–1412 (2013).
39. R. J. Full, D. E. Koditschek, Templates and anchors: Neuromechanical hypotheses of legged locomotion on land. *J. Exp. Biol.* **202**, 3325–3332 (1999).
40. P. Holmes, R. J. Full, D. E. Koditschek, J. Guckenheimer, The dynamics of legged locomotion: Models, analyses, and challenges. *SIAM Rev.* **48**, 207–304 (2006).
41. J. W. Choi, D. Aurbach, Promise and reality of post-lithium-ion batteries with high energy densities. *Nat. Rev. Mater.* **1**, 16013 (2016).
42. X. Zhu, Z. Guo, Z. Hou, Solar-powered airplanes: A historical perspective and future challenges. *Prog. Aerospace Sci.* **71**, 36–53 (2014).
43. A. Manthiram, X. Yu, S. Wang, Lithium battery chemistries enabled by solid-state electrolytes. *Nat. Rev. Mater.* **2**, 16103 (2017).
44. M. Park, J. Ryu, W. Wang, J. Cho, Material design and engineering of next-generation flow-battery technologies. *Nat. Rev. Mater.* **2**, 16080 (2016).
45. V. A. Tucker, The energetic cost of moving about: Walking and running are extremely inefficient forms of locomotion. Much greater efficiency is achieved by birds, fish—and bicyclists. *Am. Sci.* **63**, 413–419 (1975).
46. M. Wehner, R. L. Truby, D. J. Fitzgerald, B. Mosadegh, G. M. Whitesides, J. A. Lewis, R. J. Wood, An integrated design and fabrication strategy for entirely soft, autonomous robots. *Nature* **536**, 451–455 (2016).
47. M. Boyvat, J.-S. Koh, R. J. Wood, Addressable wireless actuation for multijoint folding robots and devices. *Sci. Robot.* **2**, eaan1544 (2017).
48. Z. L. Wang, Triboelectric nanogenerators as new energy technology for self-powered systems and as active mechanical and chemical sensors. *ACS Nano* **7**, 9533–9557 (2013).
49. Z. L. Wang, L. Lin, J. Chen, S. Niu, Y. Zi, *Triboelectric Nanogenerators* (Springer, 2016).
50. J. K. Parrish, L. Edelstein-Keshet, Complexity, pattern, and evolutionary trade-offs in animal aggregation. *Science* **284**, 99–101 (1999).
51. https://en.wikipedia.org/wiki/Hyper-converged_infrastructure.
52. J. P. Desai, J. P. Ostrowski, V. Kumar, Modeling and control of formations of nonholonomic mobile robots. *IEEE Trans. Rob. Autom.* **17**, 905–908 (2001).
53. A. Jadbabaie, J. Lin, A. S. Morse, Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Trans. Automat. Contr.* **48**, 988–1001 (2003).
54. S. Berman, Q. Lindsey, M. S. Sakar, V. Kumar, S. C. Pratt, Experimental study and modeling of group retrieval in ants as an approach to collective transport in swarm robotic systems. *Proc. IEEE* **99**, 1470–1481 (2011).
55. A. Bauer, D. Wollherr, M. Buss, Human–robot collaboration: A survey. *Int. J. Human. Robot.* **5**, 47–66 (2008).
56. A. Prorok, M. A. Hsieh, V. Kumar, The impact of diversity on optimal control policies for heterogeneous robot swarms. *IEEE Trans. Robot.* **33**, 346–358 (2017).
57. J. Rodin, *The Resilience Dividend: Being Strong in a World Where Things Go Wrong* (Public Affairs, 2014).
58. T. Kozai, G. Niu, M. Takagaki, *Plant Factory: An Indoor Vertical Farming System for Efficient Quality Food Production* (Academic Press, 2015).
59. Q. Zhu, T. Başar, Robust and resilient control design for cyber-physical systems with an application to power systems, 2011 50th *IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, Orlando, FL, 12 to 15 December 2011 (IEEE, 2012).
60. C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, J. J. Leonard, Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Trans. Robot.* **32**, 1309–1332 (2016).
61. A. Krizhevsky, I. Sutskever, G. Hinton, Imagenet classification with deep convolutional neural networks, in *Advances in Neural Information Processing Systems 25*, P. Bartlett, K. Q. Weinberger, C. J. C. Burges, L. Bottou, F. C. N. Pereira, Eds. (Neural Information Processing Systems, 2012), pp. 1097–1105.
62. K. L. Stachenfeld, M. M. Botvinick, S. J. Gershman, The hippocampus as a predictive map. *Nat. Neurosci.* **20**, 1643–1653 (2017).
63. J. Bongard, V. Zykov, H. Lipson, Resilient machines through continuous self-modeling. *Science* **314**, 1118–1121 (2006).
64. D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van den Driessche, T. Graepel, D. Hassabis, Mastering the game of Go without human knowledge. *Nature* **550**, 354–359 (2017).
65. J. Hawkins, S. Ahmad, Y. Cui, A theory of how columns in the neocortex enable learning the structure of the world. *Front. Neural Circuits* **11**, 81 (2017).
66. B. M. Lake, T. D. Ullman, J. B. Tenenbaum, S. J. Gershman, Building machines that learn and think like people. *Behav. Brain Sci.* **40**, e253 (2017).
67. J. P. Donoghue, Connecting cortex to machines: Recent advances in brain interfaces. *Nat. Neurosci.* **5**, 1085–1088 (2002).
68. M. D. Serruya, N. G. Hatsopoulos, L. Paninski, M. R. Fellous, J. P. Donoghue, Brain-machine interface: Instant neural control of a movement signal. *Nature* **416**, 141–142 (2002).
69. M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, A. B. Schwartz, Cortical control of a prosthetic arm for self-feeding. *Nature* **453**, 1098–1101 (2008).
70. U. Chaudhary, N. Birbaumer, A. Ramos-Murguialday, Brain–computer interfaces for communication and rehabilitation. *Nat. Rev. Neurol.* **12**, 513–525 (2016).
71. J. N. Mak, J. R. Wolpaw, Clinical applications of brain–computer interfaces: Current state and future prospects. *IEEE Rev. Biomed. Eng.* **2**, 187–199 (2009).
72. L. F. Nicolas-Alonso, J. Gomez-Gil, Brain computer interfaces, a review. *Sensors* **12**, 1211–1279 (2012).
73. T. O. Zander, C. Kothe, S. Jatzev, M. Gaertner, Enhancing human–computer interaction with input from active and passive brain–computer interfaces, in *Brain–Computer Interfaces*, D. S. Tan, A. Nijholt, Eds., (Springer, 2010), pp. 181–199.
74. E. Vaadia, N. Birbaumer, Grand challenges of brain computer interfaces in the years to come. *Front. Neurosci.* **3**, 151–154 (2009).
75. S. A. Papert, The summer vision project (1966); <http://hdl.handle.net/1721.1/6125>
76. P. Loncomilla, J. Ruiz-del-Solar, L. Martínez, Object recognition using local invariant features for robotic

- applications: A survey. *Pattern Recogn.* **60**, 499–514 (2016).
77. S. Herath, M. Harandi, F. Porikli, Going deeper into action recognition: A survey. *Image Vis. Comput.* **60**, 4–21 (2017).
 78. H. Admoni, B. Scassellati, Social eye gaze in human-robot interaction: A review. *J. Hum. Robot Interact.* **6**, 25–63 (2017).
 79. G. Sukthankar, C. Geib, H. H. Bui, D. Pynadath, R. P. Goldman, *Plan, Activity, and Intent Recognition: Theory and Practice* (Elsevier, 2014).
 80. S. Rossi, F. Ferland, A. Tapus, User profiling and behavioral adaptation for HRI: A survey. *Pattern Recognit. Lett.* **99**, 3–12 (2017).
 81. B. Hayes, B. Scassellati, Autonomously constructing hierarchical task networks for planning and human-robot collaboration, 2016 *IEEE International Conference on Robotics and Automation (ICRA)*, Stockholm, Sweden, 16 to 21 May 2016 (IEEE, 2016).
 82. C. A. Corneanu, M. O. Simón, J. F. Cohn, S. E. Guerrero, Survey on RGB, 3D, thermal, and multimodal approaches for facial expression recognition: History, trends, and affect-related applications. *IEEE Trans. Pattern Anal. Mach. Intell.* **38**, 1548–1568 (2016).
 83. B. Scassellati, Theory of mind for a humanoid robot. *Auton. Robot.* **12**, 13–24 (2002).
 84. M. I. Jordan, T. M. Mitchell, Machine learning: Trends, perspectives, and prospects. *Science* **349**, 255–260 (2015).
 85. G.-Z. Yang, J. Cambias, K. Cleary, E. Daimler, J. Drake, P. E. Dupont, N. Hata, P. Kazanzides, S. Martel, R. V. Patel, V. J. Santos, R. H. Taylor, Medical robotics—Regulatory, ethical, and legal considerations for increasing levels of autonomy. *Sci. Robot.* **2**, eaam8638 (2017).
 86. C. Bergeles, G.-Z. Yang, From passive tool holders to microsurgions: Safer, smaller, smarter surgical robots. *IEEE Trans. Biomed. Eng.* **61**, 1565–1576 (2014).
 87. V. Vitiello, S.-L. Lee, T. P. Cundy, G.-Z. Yang, Emerging robotic platforms for minimally invasive surgery. *IEEE Rev. Biomed. Eng.* **6**, 111–126 (2013).
 88. A. Shademan, R. S. Decker, J. D. Opfermann, S. Leonard, A. Krieger, P. C. W. Kim, Supervised autonomous robotic soft tissue surgery. *Sci. Transl. Med.* **8**, 337ra64 (2016).
 89. K. P. Sajadi, H. B. Goldman, Robotic pelvic organ prolapse surgery. *Nat. Rev. Urol.* **12**, 216–224 (2015).
 90. C. J. Payne, I. Wamala, D. Bautista-Salinas, M. Saeed, D. Van Story, T. Thalhofer, M. A. Horvath, C. Abah, P. J. del Nido, C. J. Walsh, N. V. Vasilyev, Soft robotic ventricular assist device with septal bracing for therapy of heart failure. *Sci. Robot.* **2**, eaan6736 (2017).
 91. V. Iacovacci, L. Ricotti, P. Dario, A. Menciasci, Design and development of a mechatronic system for noninvasive refilling of implantable artificial pancreas. *IEEE ASME Trans. Mechatron.* **20**, 1160–1169 (2015).
 92. G.-Z. Yang, *Implantable Sensors and Systems—From Theory to Practice* (Springer, 2018).
 93. D. D. Damian, K. Price, S. Arabagi, I. Berra, Z. Machaidze, S. Manjila, S. Shimada, A. Fabozzo, G. Arnal, D. Van Story, J. D. Goldsmith, A. T. Agoston, C. Kim, R. Jennings, P. D. Ngo, M. Manfredi, P. E. Dupont, In vivo tissue regeneration with robotic implants. *Sci. Robot.* **3**, eaaq0018 (2018).
 94. B. J. Nelson, I. K. Kaliakatsos, J. J. Abbott, Microrobots for minimally invasive medicine. *Annu. Rev. Biomed. Eng.* **12**, 55–85 (2010).
 95. A. Servant, F. Qiu, M. Mazza, K. Kostarelos, B. J. Nelson, Controlled in vivo swimming of a swarm of bacteria-like microbotic flagella. *Adv. Mater.* **27**, 2981–2988 (2015).
 96. O. Felfoul, M. Mohammadi, S. Taherkhani, D. de Lanauze, Y. Zhong Xu, D. Loghini, S. Essa, S. Jancik, D. Houle, M. Lafleur, L. Gaboury, M. Tabrizian, N. Kaou, M. Atkin, T. Vuong, G. Batist, N. Beauchemin, D. Radzioch, S. Martel, Magneto-aerotactic bacteria deliver drug-containing nanoliposomes to tumour hypoxic regions. *Nat. Nanotechnol.* **11**, 941–947 (2016).
 97. C. Cath, S. Wachter, B. Mittelstadt, M. Taddeo, L. Floridi, Artificial Intelligence and the ‘Good Society’: The US, EU, and UK approach. *Sci. Eng. Ethics* **2017**, 1–24 (2017).
 98. L. Floridi, M. Taddeo, What is data ethics? *Philos. Trans. A Math. Phys. Eng. Sci.* **374**, 20160360 (2016).
 99. S. Wachter, B. Mittelstadt, L. Floridi, Transparent, explainable, and accountable AI for robotics. *Sci. Robot.* **2**, eaan6080 (2017).
 100. L. Floridi, Distributed morality in an information society. *Sci. Eng. Ethics* **19**, 727–743 (2013).
 101. N. Wiener, *The Human Use of Human Beings* (Houghton Mifflin, 1954).
 102. M. Taddeo, The limits of deterrence theory in cyberspace. *Philos. Technol.* **2017**, 1–17 (2017).
 103. J. Aguilar, T. Zhang, F. Qian, M. Kingsbury, B. McInroe, N. Mazouchova, C. Li, R. Maladen, C. Gong, M. Travers, R. L. Hatton, A review on locomotion robophysics: The study of movement at the intersection of robotics, soft matter and dynamical systems. *Rep. Prog. Phys.* **79**, 110001 (2016).
 104. House of Commons Rebuilding, <http://hansard.millbanksystems.com/commons/1943/oct/28/house-of-commons-rebuilding>.
 105. S.-J. Park, M. Gazzola, K. S. Park, S. Park, V. Di Santo, E. L. Blevins, J. U. Lind, P. H. Campbell, S. Dauth, A. K. Capulli, F. S. Pasqualini, S. Ahn, A. Cho, H. Yuan, B. M. Maoz, R. Vijaykumar, J.-W. Choi, K. Deisseroth, G. V. Lauder, L. Mahadevan, K. K. Parker, Phototactic guidance of a tissue-engineered soft-robotic ray. *Science* **353**, 158–162 (2016).
 106. H. Singh, T. Maksym, J. Wilkinson, G. Williams, Inexpensive, small AUVs for studying ice-covered polar environments. *Sci. Robot.* **2**, eaan4809 (2017).

Acknowledgments: We thank all those who have responded to our online survey. In particular, we thank H. Choset, P. Dario, P. Fiorini, T. Fukuda, D. Leff, and S. Russell, who provided detailed comments, review, and original materials for this article.

10.1126/scirobotics.aar7650

Citation: G.-Z. Yang, J. Bellingham, P. E. Dupont, P. Fischer, L. Floridi, R. Full, N. Jacobstein, V. Kumar, M. McNutt, R. Merrifield, B. J. Nelson, B. Scassellati, M. Taddeo, R. Taylor, M. Veloso, Z. L. Wang, R. Wood, The grand challenges of *Science Robotics*. *Sci. Robot.* **3**, eaar7650 (2018).

The grand challenges of *Science Robotics*

Guang-Zhong Yang, Jim Bellingham, Pierre E. Dupont, Peer Fischer, Luciano Floridi, Robert Full, Neil Jacobstein, Vijay Kumar, Marcia McNutt, Robert Merrifield, Bradley J. Nelson, Brian Scassellati, Mariarosaria Taddeo, Russell Taylor, Manuela Veloso, Zhong Lin Wang and Robert Wood

Sci. Robotics **3**, eaar7650.
DOI: 10.1126/scirobotics.aa7650

ARTICLE TOOLS <http://robotics.sciencemag.org/content/3/14/eaar7650>

REFERENCES This article cites 89 articles, 14 of which you can access for free
<http://robotics.sciencemag.org/content/3/14/eaar7650#BIBL>

PERMISSIONS <http://www.sciencemag.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of Service](#)

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. 2017 © The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. The title *Science Robotics* is a registered trademark of AAAS.