

Getting AI Right: Introductory Notes on AI & Society

James Manyika

NATHAN: Do you know what the Turing Test is?

CALEB: . . . Yeah. I know what the Turing Test is. It's when a human interacts with a computer. And if the human doesn't know they're interacting with a computer, the test is passed.

NATHAN: And what does a pass tell us?

CALEB: That the computer has artificial intelligence. . . .

NATHAN: You got it. Because if that test is passed, you are dead center of the single greatest scientific event in the history of man.

CALEB: If you've created a conscious machine, it's not the history of man. It's the history of gods.

This dialogue is from an early scene in the 2014 film *Ex Machina*, in which Nathan has invited Caleb to determine whether Nathan has succeeded in creating artificial intelligence.¹ The achievement of powerful artificial general intelligence has long held a grip on our imagination not only for its exciting as well as worrisome possibilities, but also for its suggestion of a new, uncharted era for humanity. In opening his 2021 BBC Reith Lectures, titled “Living with Artificial Intelligence,” Stuart Russell states that “the eventual emergence of general-purpose artificial intelligence [will be] the biggest event in human history.”²

Over the last decade, a rapid succession of impressive results has brought wider public attention to the possibilities of powerful artificial intelligence. In machine vision, researchers demonstrated systems that could recognize objects as well as, if not better than, humans in some situations. Then came the games. Complex games of strategy have long been associated with superior intelligence, and so when AI systems beat the best human players at chess, Atari games, Go, shogi, StarCraft, and Dota, the world took notice. It was not just that AIs beat humans (although that was astounding when it first happened), but the escalating progression of how they did it: initially by learning from expert human play, then from self-play, then by teaching themselves the principles of the games from the ground up, eventually yielding single systems that could learn, play, and win at

several structurally different games, hinting at the possibility of generally intelligent systems.³

Speech recognition and natural language processing have also seen rapid and headline-grabbing advances. Most impressive has been the emergence recently of large language models capable of generating human-like outputs. Progress in language is of particular significance given the role language has always played in human notions of intelligence, reasoning, and understanding. While the advances mentioned thus far may seem abstract, those in driverless cars and robots have been more tangible given their embodied and often biomorphic forms. Demonstrations of such embodied systems exhibiting increasingly complex and autonomous behaviors in our physical world have captured public attention.

Also in the headlines have been results in various branches of science in which AI and its related techniques have been used as tools to advance research from materials and environmental sciences to high energy physics and astronomy.⁴ A few highlights, such as the spectacular results on the fifty-year-old protein-folding problem by AlphaFold, suggest the possibility that AI could soon help tackle science's hardest problems, such as in health and the life sciences.⁵

While the headlines tend to feature results and demonstrations of a future to come, AI and its associated technologies are already here and pervade our daily lives more than many realize. Examples include recommendation systems, search, language translators—now covering more than one hundred languages—facial recognition, speech to text (and back), digital assistants, chatbots for customer service, fraud detection, decision support systems, energy management systems, and tools for scientific research, to name a few. In all these examples and others, AI-related techniques have become components of other software and hardware systems as methods for learning from and incorporating messy real-world inputs into inferences, predictions, and, in some cases, actions. As director of the Future of Humanity Institute at the University of Oxford, Nick Bostrom noted back in 2006, “A lot of cutting-edge AI has filtered into general applications, often without being called AI because once something becomes useful enough and common enough it’s not labeled AI anymore.”⁶

As the scope, use, and usefulness of these systems have grown for individual users, researchers in various fields, companies and other types of organizations, and governments, so too have concerns when the systems have not worked well (such as bias in facial recognition systems), or have been misused (as in deepfakes), or have resulted in harms to some (in predicting crime, for example), or have been associated with accidents (such as fatalities from self-driving cars).⁷

Dædalus last devoted a volume to the topic of artificial intelligence in 1988, with contributions from several of the founders of the field, among others. Much of that issue was concerned with questions of whether research in AI was making progress, of whether AI was at a turning point, and of its foundations, mathemati-

cal, technical, and philosophical – with much disagreement. However, in that volume there was also a recognition, or perhaps a rediscovery, of an alternative path toward AI – the connectionist learning approach and the notion of neural nets – and a burgeoning optimism for this approach’s potential. Since the 1960s, the learning approach had been relegated to the fringes in favor of the symbolic formalism for representing the world, our knowledge of it, and how machines can reason about it. Yet no essay captured some of the mood at the time better than Hilary Putnam’s “Much Ado About Not Very Much.” Putnam questioned the *Dædalus* issue itself: “Why a whole issue of *Dædalus*? Why don’t we wait until AI achieves something and then have an issue?” He concluded:

Perhaps the optimistic view is right, but I do not see anyone on the scene, in either artificial intelligence or inductive logic, who has any interesting ideas about how the topic-neutral [general] learning strategy works. When someone does appear with such an idea, that will be time for *Dædalus* to publish an issue on AI.⁸

This volume of *Dædalus* is indeed the first since 1988 to be devoted to artificial intelligence. This volume does not rehash the same debates; much else has happened since, mostly as a result of the success of the machine learning approach that was being rediscovered and reimagined, as discussed in the 1988 volume. This issue aims to capture where we are in AI’s development and how its growing uses impact society. The themes and concerns herein are colored by my own involvement with AI. Besides the television, films, and books that I grew up with, my interest in AI began in earnest in 1989 when, as an undergraduate at the University of Zimbabwe, I undertook a research project to model and train a neural network.⁹ I went on to do research on AI and robotics at Oxford. Over the years, I have been involved with researchers in academia and labs developing AI systems, studying AI’s impact on the economy, tracking AI’s progress, and working with others in business, policy, and labor grappling with its opportunities and challenges for society.¹⁰

The authors of the twenty-five essays in this volume range from AI scientists and technologists at the frontier of many of AI’s developments to social scientists at the forefront of analyzing AI’s impacts on society. The volume is organized into ten sections. Half of the sections are focused on AI’s development, the other half on its intersections with various aspects of society. In addition to the diversity in their topics, expertise, and vantage points, the authors bring a range of views on the possibilities, benefits, and concerns for society. I am grateful to the authors for accepting my invitation to write these essays.

Before proceeding further, it may be useful to say what we mean by artificial intelligence. The headlines and increasing pervasiveness of AI and its associated technologies have led to some conflation and confusion about

what exactly counts as AI. This has not been helped by the current trend – among researchers in science and the humanities, startups, established companies, and even governments – to associate anything involving not only machine learning, but data science, algorithms, robots, and automation of all sorts with AI. This could simply reflect the hype now associated with AI, but it could also be an acknowledgment of the success of the current wave of AI and its related techniques and their wide-ranging use and usefulness. I think both are true; but it has not always been like this. In the period now referred to as the AI winter, during which progress in AI did not live up to expectations, there was a reticence to associate most of what we now call AI with AI.

Two types of definitions are typically given for AI. The first are those that suggest that it is the ability to artificially do what intelligent beings, usually human, can do. For example, artificial intelligence is:

the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings.¹¹

The human abilities invoked in such definitions include visual perception, speech recognition, the capacity to reason, solve problems, discover meaning, generalize, and learn from experience. Definitions of this type are considered by some to be limiting in their human-centricity as to what counts as intelligence and in the benchmarks for success they set for the development of AI (more on this later). The second type of definitions try to be free of human-centricity and define an intelligent agent or system, whatever its origin, makeup, or method, as:

Any system that perceives its environment and takes actions that maximize its chance of achieving its goals.¹²

This type of definition also suggests the pursuit of goals, which could be given to the system, self-generated, or learned.¹³ That both types of definitions are employed throughout this volume yields insights of its own.

These definitional distinctions notwithstanding, the term AI, much to the chagrin of some in the field, has come to be what cognitive and computer scientist Marvin Minsky called a “suitcase word.”¹⁴ It is packed variously, depending on who you ask, with approaches for achieving intelligence, including those based on logic, probability, information and control theory, neural networks, and various other learning, inference, and planning methods, as well as their instantiations in software, hardware, and, in the case of embodied intelligence, systems that can perceive, move, and manipulate objects.

Three questions cut through the discussions in this volume: 1) Where are we in AI’s development? 2) What opportunities and challenges does AI pose for society? 3) How much about AI is really about us?

Where are we in AI's development?

Notions of intelligent machines date all the way back to antiquity.¹⁵ Philosophers, too, among them Hobbes, Leibnitz, and Descartes, have been dreaming about AI for a long time; Daniel Dennett suggests that Descartes may have even anticipated the Turing Test.¹⁶ The idea of computation-based machine intelligence traces to Alan Turing's invention of the universal Turing machine in the 1930s, and to the ideas of several of his contemporaries in the mid-twentieth century. But the birth of artificial intelligence as we know it and the use of the term is generally attributed to the now famed Dartmouth summer workshop of 1956. The workshop was the result of a proposal for a two-month summer project by John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon whereby "An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves."¹⁷

In their respective contributions to this volume, "From So Simple a Beginning: Species of Artificial Intelligence" and "If We Succeed," and in different but complementary ways, Nigel Shadbolt and Stuart Russell chart the key ideas and developments in AI, its periods of excitement as well as the aforementioned AI winters. The current AI spring has been underway since the 1990s, with headline-grabbing breakthroughs appearing in rapid succession over the last ten years or so: a period that Jeffrey Dean describes in the title of his essay as a "golden decade," not only for the pace of AI development but also its use in a wide range of sectors of society, as well as areas of scientific research.¹⁸ This period is best characterized by the approach to achieve artificial intelligence through learning from experience, and by the success of neural networks, deep learning, and reinforcement learning, together with methods from probability theory, as ways for machines to learn.¹⁹

A brief history may be useful here: In the 1950s, there were two dominant visions of how to achieve machine intelligence. One vision was to use computers to create a logic and symbolic representation of the world and our knowledge of it and, from there, create systems that could reason about the world, thus exhibiting intelligence akin to the mind. This vision was most espoused by Allen Newell and Hebert Simon, along with Marvin Minsky and others. Closely associated with it was the "heuristic search" approach that supposed intelligence was essentially a problem of exploring a space of possibilities for answers. The second vision was inspired by the brain, rather than the mind, and sought to achieve intelligence by learning. In what became known as the connectionist approach, units called perceptrons were connected in ways inspired by the connection of neurons in the brain. At the time, this approach was most associated with Frank Rosenblatt. While there was initial excitement about both visions, the first came to dominate, and did so for decades, with some successes, including so-called expert systems.

Not only did this approach benefit from championing by its advocates and plentiful funding, it came with the suggested weight of a long intellectual tradition – exemplified by Descartes, Boole, Frege, Russell, and Church, among others – that sought to manipulate symbols and to formalize and axiomatize knowledge and reasoning. It was only in the late 1980s that interest began to grow again in the second vision, largely through the work of David Rumelhart, Geoffrey Hinton, James McClelland, and others. The history of these two visions and the associated philosophical ideas are discussed in Hubert Dreyfus and Stuart Dreyfus’s 1988 *Dædalus* essay “Making a Mind Versus Modeling the Brain: Artificial Intelligence Back at a Branchpoint.”²⁰ Since then, the approach to intelligence based on learning, the use of statistical methods, back-propagation, and training (supervised and unsupervised) has come to characterize the current dominant approach.

Kevin Scott, in his essay “I Do Not Think It Means What You Think It Means: Artificial Intelligence, Cognitive Work & Scale,” reminds us of the work of Ray Solomonoff and others linking information and probability theory with the idea of machines that can not only learn, but compress and potentially generalize what they learn, and the emerging realization of this in the systems now being built and those to come. The success of the machine learning approach has benefited from the boon in the availability of data to train the algorithms thanks to the growth in the use of the Internet and other applications and services. In research, the data explosion has been the result of new scientific instruments and observation platforms and data-generating breakthroughs, for example, in astronomy and in genomics. Equally important has been the co-evolution of the software and hardware used, especially chip architectures better suited to the parallel computations involved in data- and compute-intensive neural networks and other machine learning approaches, as Dean discusses.

Several authors delve into progress in key subfields of AI.²¹ In their essay, “Searching for Computer Vision North Stars,” Fei-Fei Li and Ranjay Krishna chart developments in machine vision and the creation of standard data sets such as ImageNet that could be used for benchmarking performance. In their respective essays “Human Language Understanding & Reasoning” and “The Curious Case of Commonsense Intelligence,” Chris Manning and Yejin Choi discuss different eras and ideas in natural language processing, including the recent emergence of large language models comprising hundreds of billions of parameters and that use transformer architectures and self-supervised learning on vast amounts of data.²² The resulting pretrained models are impressive in their capacity to take natural language prompts for which they have not been trained specifically and generate human-like outputs, not only in natural language, but also images, software code, and more, as Mira Murati discusses and illustrates in “Language & Coding Creativity.” Some have started to refer to these large language models as foundational models in that once they are trained, they are adaptable to a wide range of tasks and outputs.²³ But

despite their unexpected performance, these large language models are still early in their development and have many shortcomings and limitations that are highlighted in this volume and elsewhere, including by some of their developers.²⁴

In “The Machines from Our Future,” Daniela Rus discusses the progress in robotic systems, including advances in the underlying technologies, as well as in their integrated design that enables them to operate in the physical world. She highlights the limitations in the “industrial” approaches used thus far and suggests new ways of conceptualizing robots that draw on insights from biological systems. In robotics, as in AI more generally, there has always been a tension as to whether to copy or simply draw inspiration from how humans and other biological organisms achieve intelligent behavior. Elsewhere, AI researcher Demis Hassabis and colleagues have explored how neuroscience and AI learn from and inspire each other, although so far more in one direction than the other, as Alexis Baria and Keith Cross have suggested.²⁵

Despite the success of the current approaches to AI, there are still many shortcomings and limitations, as well as conceptually hard problems in AI.²⁶ It is useful to distinguish on one hand problematic shortcomings, such as when AI does not perform as intended or safely, or produces biased or toxic outputs that can lead to harm, or when it impinges on privacy, or generates false information about the world, or when it has characteristics such as lack of explainability, all of which can lead to a loss of public trust. These shortcomings have rightly captured the attention of the wider public and regulatory bodies, as well as researchers, among whom there is an increased focus on technical AI and ethics issues.²⁷ In recent years, there has been a flurry of efforts to develop principles and approaches to responsible AI, as well as bodies involving industry and academia, such as the Partnership on AI, that aim to share best practices.²⁸ Another important shortcoming has been the significant lack of diversity—especially with respect to gender and race—in the people researching and developing AI in both industry and academia, as has been well documented in recent years.²⁹ This is an important gap in its own right, but also with respect to the characteristics of the resulting AI and, consequently, in its intersections with society more broadly.

On the other hand, there are limitations and hard problems associated with the things that AI is not yet capable of that, if solved, could lead to more powerful, more capable, or more general AI. In their Turing Lecture, deep learning pioneers Yoshua Bengio, Yann LeCun, and Geoffrey Hinton took stock of where deep learning stands and highlighted its current limitations, such as the difficulties with out-of-distribution generalization.³⁰ In the case of natural language processing, Manning and Choi highlight the hard challenges in reasoning and common-sense understanding, despite the surprising performance of large language models. Elsewhere, computational linguists Emily Bender and Alexander Koller have challenged the notion that large language models do anything resembling under-

standing, learning, or meaning.³¹ In “Multi-Agent Systems: Technical & Ethical Challenges of Functioning in a Mixed Group,” Kobi Gal and Barbara Grosz discuss the hard problems in multi-agent systems, highlighting the conceptual difficulties – such as how to reason about other agents, their belief systems, and intentionality – as well as ethical challenges in both cooperative and competitive settings, especially when the agents include both humans and machines. Elsewhere, Allan Dafoe and others provide a useful overview of the open problems in cooperative AI.³² Indeed, there is a growing sense among many that we do not have adequate theories for the sociotechnical embedding of AI systems, especially as they become more capable and the scope of societal use expands.

And although AI and its related techniques are proving to be powerful tools for research in science, as examples in this volume and elsewhere illustrate – including recent examples in which embedded AI capabilities not only help evaluate results but also steer experiments by going beyond heuristics-based experimental design and become what some have termed “self-driving laboratories”³³ – getting AI to understand science and mathematics and to theorize and develop novel concepts remain grand challenges for AI.³⁴ Indeed the possibility that more powerful AI could lead to new discoveries in science, as well as enable game-changing progress in some of humanity’s greatest challenges and opportunities, has long been a key motivation for many at the frontier of AI research to build more capable systems.

Beyond the particulars of each subfield of AI, the list of more general hard problems that continue to limit the possibility of more capable AI includes one-shot learning, cross-domain generalizations, causal reasoning, grounding, complexities of timescales and memory, and meta-cognition.³⁵ Consideration of these and other hard problems that could lead to more capable systems raises the question of whether current approaches – mostly characterized by deep learning, the building of larger and larger and more foundational and multimodal models, and reinforcement learning – are sufficient, or whether entirely different conceptual approaches are needed in addition, such as neuroscience-inspired cognitive agent approaches or semantic representations or reasoning based on logic and probability theory, to name a few. On whether and what kind of additional approaches might be needed, the AI community is divided, but many believe the current approaches³⁶ along with further evolution of compute and learning architectures have yet to reach their limits.³⁷

The debate about the sufficiency of the current approaches is closely associated with the question of whether artificial general intelligence can be achieved, and if so, how and when. *Artificial general intelligence* (AGI) is defined in distinction to what is sometimes called *narrow AI*: that is, AI developed and fine-tuned for specific tasks and goals, such as playing chess. The development of AGI, on the other hand, aims for more powerful AI – at least as powerful as humans – that is generally applicable to any problem or situation and, in some conceptions, includes the capacity to evolve and improve itself, as well as set and evolve its own goals and

preferences. Though the question of whether, how, and when AGI will be achieved is a matter for debate, most agree that its achievement would have profound implications – beneficial and worrisome – for humanity, as is often depicted in popular books³⁸ and films such as *2001: A Space Odyssey* through *Terminator* and *The Matrix* to *Ex Machina* and *Her*. Whether it is imminent or not, there is growing agreement among many at the frontier of AI research that we should prepare for the possibility of powerful AGI with respect to safety and control, alignment and compatibility with humans, its governance and use, and the possibility that multiple varieties of AGI could emerge, and that we should factor these considerations into how we approach the development of AGI.

Most of the investment, research and development, and commercial activity in AI today is of the narrow AI variety and in its numerous forms: what Nigel Shadbolt terms the *speciation* of AI. This is hardly surprising given the scope for useful and commercial applications and the potential for economic gains in multiple sectors of the economy.³⁹ However, a few organizations have made the development of AGI their primary goal. Among the most well-known of these are DeepMind and OpenAI, each of which has demonstrated results of increasing generality, though still a long way from AGI.

What opportunities and challenges does AI pose for society?

Perhaps the most widely discussed societal impact of AI and automation is on jobs and the future of work. This is not new. In 1964, in the wake of the era's excitement about AI and automation, and concerns about their impact on jobs, President Lyndon Johnson empaneled a National Commission on Technology, Automation, and Economic Progress.⁴⁰ Among the commission's conclusions was that such technologies were important for economic growth and prosperity and "the basic fact that technology destroys jobs, but not work." Most recent studies of this effect, including those I have been involved in, have reached similar conclusions and that over time, more jobs are gained than are lost. These studies highlight that it is the sectoral and occupational transitions, the skill and wage effects – not the existence of jobs broadly – that will present the greatest challenges.⁴¹ In their essay "Automation, AI & Work," Laura Tyson and John Zysman discuss these implications for work and workers. Michael Spence goes further, in "Automation, Augmentation, Value Creation & the Distribution of Income & Wealth," to discuss the distributional issues with respect to income and wealth within and between countries, as well as the societal opportunities that are created, especially in developing countries. In "The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence," Erik Brynjolfsson discusses how the use of human benchmarks in the development of AI runs the risk of AI that substitutes for, rather than complements, human labor. He concludes that the direction AI's development

will take in this regard, and resulting outcomes for work, will depend on the incentives for researchers, companies, and governments.⁴²

Still, a concern remains that the conclusion that more jobs will be created than lost draws too much from patterns of the past and does not look far enough into the future and at what AI will be capable of. The arguments for why AI could break from past patterns of technology-driven change include: first, that unlike in the past, technological change is happening faster and labor markets (including workers) and societal systems' ability to adapt are slow and mismatched; and second, that, until now, automation has mostly mechanized physical and routine tasks, but that going forward, AI will be taking on more cognitive and nonroutine tasks, creative tasks, tasks based on tacit knowledge, and, if early examples are any indication, even socioempathic tasks are not out of the question.⁴³ In other words, "There are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until – in a visible future – the range of problems they can handle will be coextensive with the range to which the human mind has been applied." This was Herbert Simon and Allen Newell in 1957.⁴⁴

Acknowledging that this time could be different usually elicits two responses: First, that new labor markets will emerge in which people will value things done by other humans for their own sake, even when machines may be capable of doing these things as well as or even better than humans. The other response is that AI will create so much wealth and material abundance, all without the need for human labor, and the scale of abundance will be sufficient to provide for everyone's needs. And when that happens, humanity will face the challenge that Keynes once framed: "For the first time since his creation man will be faced with his real, his permanent problem – how to use his freedom from pressing economic cares, how to occupy the leisure, which science and compound interest will have won for him, to live wisely and agreeably and well."⁴⁵ However, most researchers believe that we are not close to a future in which the majority of humanity will face Keynes's challenge, and that until then, there are other AI- and automation-related effects that must be addressed in the labor markets now and in the near future, such as inequality and other wage effects, education, skilling, and how humans work alongside increasingly capable machines – issues that Laura Tyson and John Zysman, Michael Spence, and Erik Brynjolfsson discuss in this volume.

Jobs are not the only aspect of the economy impacted by AI. Russell provides a directional estimate of the potentially huge economic bounty from artificial general intelligence, once fully realized: a global GDP of \$750 trillion, or ten times today's global GDP. But even before we get to fully realized general-purpose AI, the commercial opportunities for companies and, for countries, the potential productivity gains and economic growth as well as economic competitiveness from narrow AI and its related technologies are more than sufficient to ensure intense

pursuit and competition by companies and countries in the development, deployment, and use of AI. At the national level, while many believe the United States is ahead, it is generally acknowledged that China is fast becoming a major player in AI, as evidenced by its growth in AI research, infrastructure, and ecosystems, as highlighted in several reports.⁴⁶ Such competition will likely have market structure effects for companies and countries, given the characteristics of such technologies as discussed by Eric Schmidt, Spence, and others elsewhere.⁴⁷ Moreover, the competitive dynamics may get in the way of responsible approaches to AI and issues requiring collective action (such as safety) between competitors, whether they are companies or countries, as Amanda Askill, Miles Brundage, and Gillian Hadfield have highlighted.⁴⁸

Nations have reasons beyond the economic to want to lead in AI. The role of AI in national security – in surveillance, signals intelligence, cyber operations, defense systems, battle-space superiority, autonomous weapons, even disinformation and other forms of sociopolitical warfare – is increasingly clear. In “AI, Great Power Competition & National Security,” Eric Schmidt, who cochaired the U.S. National Security Commission on Artificial Intelligence, paints a stark picture of current and future risks that AI technologies pose to international security and stability. Schmidt calls for the exploration of shared limits and treaties on AI, even among rivals. Short of that, he points to confidence-building measures to limit risks and increase trust.⁴⁹ At the same time, Russell and Shadbolt, respectively, spotlight concerns regarding autonomous weapons and weaponized AI.

In “The Moral Dimension of AI-Assisted Decision-Making: Some Practical Perspectives from the Front Lines,” former Secretary of Defense Ash Carter identifies lessons for AI drawn from other national security-related technologies, such as nuclear weapons, while focusing on the ethics of automated decision-making. However, there are important differences between AI and nuclear technologies: for example, AI’s development has been led by a private sector in pursuit of global opportunities. And, as Schmidt points out, AI technologies in their development and use have network effects and tend to consolidate around those who lead in their development, whether they are companies or countries. This pits commercial and economic interests for companies and countries on one hand, and the national security interests of countries on the other.⁵⁰ Not fully explored in this volume are the implications for companies (as well as other types of organizations) and countries not at the forefront of AI’s development but that could benefit from its use. This is of particular significance given that many have highlighted the potential for AI and its related technologies to contribute, along with other social and developmental efforts, to tackling many current and future global and societal challenges.⁵¹ The COVID-19 pandemic has given us a live example of the human cost when countries at the forefront of a globally valuable discovery, such as a vaccine, do not or are slow to share it with poorer parts of the world.

As the use of AI has grown to encompass not only consumer applications and services, but also those in health care, financial services, public services, and commerce generally, it has in many instances improved effectiveness and decision quality and enabled much-needed cost and performance optimization. At the same time, in some cases, the use of algorithms has led to issues of bias and fairness, often the result of bias in the training data and the societal systems through which such data are collected.⁵² Sonia Katyal uses examples from facial recognition, policing, and sentencing to argue in “Democracy & Distrust in an Era of Artificial Intelligence” that, when there is an absence of representation and participation, AI-powered systems carry the same risks and potential for distrust as political systems. In “Distrust of Artificial Intelligence: Sources & Responses from Computer Science & Law,” Cynthia Dwork and Martha Minow highlight the absence of ground truth and what happens when utility for users and commercial interests are at odds with considerations of privacy and the risks of societal harms.⁵³ In light of these concerns, as well as the beneficial possibilities of AI, Mariano-Florentino Cuéllar, a former California Supreme Court Justice, and Aziz Huq frame how we might achieve the title of their essay: artificially intelligent regulation.

It is easy to see how governments and organizations in their desire to observe, analyze, and optimize everything would be tempted to use AI to create increasingly powerful “seeing rooms.” In “Socializing Data,” Diane Coyle discusses the history and perils of seeing rooms, even when well intentioned, and the problems that arise when markets are the primary mechanism for how AI uses social data. For governments, the opportunity to use AI to improve the delivery and effectiveness of public services is also hard to ignore. In her essay “Rethinking AI for Good Governance,” Helen Margetts asks what a public sector AI would look like. She draws on public sector examples from different countries to highlight key challenges, notably those related to issues like resource allocation, that are more “normatively loaded” in the public sector than they are for firms. She concludes by exploring how and in which areas governments can make the most ambitious and societally beneficial use of AI.

How much about AI is really about us?

At the end of her essay, Katyal quotes J. David Bolter from his 1984 *Dædalus* essay: “I think artificial intelligence will grow in importance as a way of looking at the human mind, regardless of the success of the programs themselves in imitating various aspects of human thought.” Taking this suggestion, one can ask various kinds of questions about us using the mirror AI provides, especially as it becomes more capable: What does it mean to be intelligent, creative, or, more generally, cognitively human when many of the ways we have defined these characteristics of ourselves increasingly can be imitated or even, in the future, done better or

better done by machines? How much of being human needs the mystery of not knowing how it works, or relies on our inability to mimic it or replicate it artificially? What happens when this changes? To what extent do our human ability-bounded conceptions of *X* (where *X* could be intelligence, creativity, empathy, relations, and so on) limit the possibility of other forms of *X* that may complement or serve humanity better? To what extent must we reexamine our socioeconomic systems and institutions, our social infrastructure, what lies at the heart of our social policies, at our notions of justice, representation, and inclusion, and face up to what they really are (and have been) and what they will need to be in the age of AI?

Their shortcomings notwithstanding, the emergence of large language models and their ability to generate human-like outputs provides a “laboratory” of sorts, as Tobias Rees calls it, to explore questions about us in an era of increasingly capable machines. We may have finally arrived at what Dennett suggests at the end of his 1988 essay, that “AI has not yet solved any of our ancient riddles . . . but it has provided us with new ways of disciplining and extending philosophical imagination that we have only just begun to exploit.”⁵⁴ Murati explores how humans could relate to and work alongside machines when machines can generate outputs approaching human-like creativity. She illustrates this with examples generated by GPT-3, OpenAI’s large language model. The possibilities she describes echo what Scott suggests: that we humans may have to rethink our relation to work and other creative activities.

Blaise Agüera y Arcas explores the titular question of his essay “Do Large Language Models Understand Us?” through a series of provocations interspersed with outputs from LaMDA, Google’s large language model. He asks whether we are gatekeeping or constantly moving the goalposts when it comes to notions such as intelligence or understanding, even consciousness, in order to retain these for ourselves. Pamela McCorduck, in her still-relevant history of the field, *Machines Who Think*, first published in 1979, put it thus: “It’s part of the history of the field of artificial intelligence that every time somebody figured out how to make a computer do something – play good checkers, solve simple but relatively informal problems – there was a chorus of critics to say, ‘that’s not thinking.’”⁵⁵ As to what machines are actually doing or not actually doing when they appear to be thinking, one could ask whether whatever they are doing is different from what humans do in any way other than how it is being done. In “Non-Human Words: On GPT-3 as a Philosophical Laboratory,” while engaging in current debates about the nature of these models, Rees also discusses how conceptions of the human have been intertwined with language in different historical eras and considers the possibility of a new era in which language is separated from humans.

In “Signs Taken for Wonders: AI, Art & the Matter of Race,” Michele Elam illustrates how, throughout history, socially transformative technologies have played a formalizing and codifying role in our conceptions of what constitutes

humanity and who the “us” is. In how they are developed, used, and monetized, and by whom, she argues that technologies like AI have the effect of universalizing particular conceptions of what it is to be human and to progress, often at the exclusion of other ways of being human and of progressing and knowing, especially those associated with Black, Latinx, and Indigenous communities and with feminist, queer, disability, and decolonial perspectives; further highlighting the need for diversity among those involved in AI’s development. Elsewhere, Timnit Gebru has clearly illustrated how, like other technologies with the potential to benefit society, AI can also worsen systematic discrimination of already marginalized groups.⁵⁶ In another example of AI as formalizer to ill-effect, Blaise Agüera y Arcas, Margaret Mitchell, and Alexander Todorov examine the use of machine learning to correlate physical characteristics with nonphysical traits, not unlike nineteenth- and twentieth-century physiognomy, and point out the harmful circular logic of essentialism that can result when AI is used as a detector of traits.⁵⁷

Progress in AI not only raises the stakes on ethical issues associated with its application, it also helps bring to light issues already extant in society. Many have shown how algorithms and automated decision-making can not only perpetuate but also formalize and amplify existing societal inequalities, as well as create new inequalities.⁵⁸ In addition, the challenge to remove bias or code for fairness may also create the opportunity for society to examine in a new light what it means by “fair.”⁵⁹ Here it is worth recalling Dennett being unimpressed by Putnam’s indictment of AI, that “AI has utterly failed, over a quarter century, to solve problems that philosophy has utterly failed to solve over two millennia.”⁶⁰ Furthermore, examining the role of algorithms and automated decision-making and the data needed to inform algorithms may shed light on what actually underlies society’s goals and policies in the first place, issues that have begun to receive attention in the literature of algorithms, fairness, and social welfare.⁶¹ In “Toward a Theory of Justice for Artificial Intelligence,” Jason Gabriel, drawing on Rawls’s theory of justice, explores the intersection of AI and distributive justice by considering the role that sociotechnical systems play. He examines issues including basic liberties and equality of opportunity to suggest that considerations of distributive justice may now need to grapple with the particularities of AI as a technological system and that could lead to some novel consequences.

And as AI becomes more powerful, a looming question becomes how to align AI with humans with respect to safety and control, goals and preferences, even values. The question of AI and control is as old as the field itself; Turing himself raised it, as Russell reminds us. Some researchers believe that concerns about these sorts of risks are overblown given the nature of AI, while others believe we are a long way away from existential control risks but that research must begin to consider approaches to the control issue and factor it into how we develop more powerful AI systems.⁶² Russell proposes an approach to alignment and human compatibility

that capitalizes on uncertainty in goals and human preferences, and makes use of inverse reinforcement learning as a way for machines to learn human preferences. Elsewhere, Gabriel has discussed the range of possibilities as to what we mean by alignment with AI, with each possibility presenting its own complexities.⁶³ But in Gabriel, as in Russell, there are considerable normative challenges involved, along with complications due to the plasticity of human preferences.

In “Artificial Intelligence, Humanistic Ethics,” John Tasioulas argues that designing AI that aligns with human preferences is one thing, but it does not obviate the need to determine what those human preferences should be in the first place. He challenges the tendency to default to preference utilitarianism and its maximization by AI developers, as well as by economic and governmental actors (who often use wealth maximization and GDP as proxies), which leads to market mechanisms dominating solutions at the expense of nonmarket values and mechanisms, echoing some of Coyle’s concerns. Here again it seems that the mirror provided by more capable AI highlights, and with higher stakes, the unfinished (perhaps never to be finished) business of humanistic ethics, not unlike how AI may be pushing us to clarify fairness and serving notice that trolley problems are no longer just the stuff of thought experiments, since we are building autonomous systems that may have to make such choices.

Throughout the history of AI, we have asked: how good is it now? This question has been asked about every application from playing chess or Go, to knowing things, performing surgery, driving a car, writing a novel, creating art, independently making mathematical conjectures or scientific discoveries, or simply having a good bedside manner. In asking the question, it may be useful also to ask: compared to what? With an eye toward implications for society, one might compare AI with the humans best at the respective activity. There remain plenty of activities in which the “best” humans perform better than AI – as they likely will for the foreseeable future – and society is well served by these humans performing these activities. One might also compare with other samplings of humanity, such as the average person employed in or permitted to conduct that activity, or a randomly selected human. And here, as AI becomes more capable, is where the societal implications get more complicated. For example, do we raise permission standards for humans performing safety-critical activities to keep up with machine capabilities? Similarly, what determines when AI is good enough? A third comparison might be with respect to how co-extensive the range of AI capabilities become with those of humans – what Simon and Newell, as mentioned earlier, thought would eventually come to pass. How good AI systems become in this respect would likely herald the beginning of a new era for us and for society of the sort discussed previously. But perhaps the most important comparison is with respect to what we choose to use AI for and what we need AI to be capable of in order to benefit society. It would seem that in any such comparisons, along with how we

design, develop, and deploy AI, the societal implications are not foregone conclusions, but choices that are up to us.

Is all this worth it? If not, a logical response might be to stop everything, stop further development and deployment of AI, put the curses back in Pandora's box. This hardly seems realistic, given the huge economic and strategic stakes and the intense competition that has been unleashed between countries and between companies, not to mention the usefulness of AI to its users and the tantalizing beneficial possibilities, some already here, for society. My response to the question is a conditional yes.

At an AI conference a few years ago, I participated on a panel to which the host, Stuart Russell, posed a thought experiment. I forget the exact formulation, or even how I responded, but I have come to express it as follows:

It's the year 2050, AI has turned out to be hugely beneficial to society and generally acknowledged as such. What happened?

This thought experiment aims to elicit the most worthwhile possibilities we achieved, the most beneficial opportunities we realized, the hard problems we solved, the risks we averted, the unintended consequences, misuses, and abuses we avoided, and the downsides we mitigated all in order to achieve the positive outcome in a not-too-distant future. In other words, it is a way of asking what we need to get right if AI is to be a net benefit to society.

The essays in this volume of *Dædalus* highlight many of the things we must get right. Drawing from these and other discussions, and a growing literature,⁶⁴ one can compile a long working list⁶⁵ whose items can be grouped as follows: The first group is related to the challenges of building AI powerful and capable enough to achieve the exciting beneficial possibilities for humanity, but also safe and without causing or worsening individual or group harms, and able to earn public trust, especially where societal stakes are high. A second set of challenges concerns focusing AI's development and use where it can make the greatest contributions to humanity – such as in health and the life sciences, climate change, overall well-being, and in the foundational sciences and in scientific discoveries – and to deliver net positive socioeconomic outcomes for all people. The *all* is all-important, given the likelihood that without purposeful attention to it, the characteristics of the resulting AI and its benefits could accrue to a few individuals, organizations, and countries, likely those leading in its development and use. The third group of challenges centers on the responsible development, deployment, use, and governance of AI. This is especially critical given the huge economic and geopolitical stakes and the intense competition for leadership in AI that has been unleashed between companies and between countries as a result. Not prioritizing responsible approaches to AI could lead to harmful and unsafe deployment and uses, out-

right misuses, many more unintended consequences, and destabilizing race conditions among the various competitors. A fourth set of challenges concerns us: how we co-evolve our societal systems and institutions and negotiate the complexities of how to be human in an age of increasingly powerful AI.

Readers of this volume will undoubtedly develop their own perspectives on what we collectively must get right if AI is to be a net positive for humanity. While such lists will necessarily evolve as our uses and societal experience with AI grow and as AI itself becomes more powerful, the work on them must not wait.

Returning to the question, is this worth it? My affirmative answer is conditioned on confronting and getting right these hard issues. At present, it seems that the majority of human ingenuity, effort, and financial and other resources are disproportionately focused on commercial applications and the economic potential of AI, and not enough on the other issues that are also critical for AI to be a net benefit to humanity given the stakes. We can change that.

AUTHOR'S NOTE

I am grateful to the American Academy for the opportunity to conceive this *Dædalus* volume on AI & Society and to bring together diverse perspectives on AI across a range of topics. On a theme as broad as this, there are without doubt many more topics and views that are missing; for that I take responsibility.

I would like to thank the Fellows of All Souls College, Oxford, where I have been a Visiting Fellow during the editing of this *Dædalus* volume. I would also like to thank my colleagues at the McKinsey Global Institute, the AI Index, and the 100-Year Study of AI at Stanford, as well as my fellow members on the National Academies of Sciences, Engineering, and Medicine Committee on Responsible Computing Research and Its Applications, for our many discussions as well as our work together that informed the shape of this volume. I am grateful for the conversations with the authors in this volume and with others, including Hans-Peter Brondmo, Gillian Hadfield, Demis Hassabis, Mary Kay Henry, Reid Hoffman, Eric Horvitz, Margaret Levi, Eric Salobir, Myron Scholes, Julie Su, Paul Tighe, and Ngaire Woods. I am grateful for valuable comments and suggestions on this introduction from Jack Clark, Erik Brynjolfsson, Blaise Agüera y Arcas, Julian Manyika, Sarah Manyika, Maithra Raghu, and Stuart Russell, but they should not be held responsible for any errors or opinions herein.

This volume could not have come together without the generous collaboration of the Academy's editorial team of Phyllis Bendell, Director of Publications and Managing Editor of *Dædalus*, who brought her experience as guide and editor, and enthusiasm from the very beginning to the completion of this effort, and Heather Struntz and Peter Walton, who were collaborative and expert copyeditors for all the essays in this volume.

ABOUT THE AUTHOR

James Manyika, a Fellow of the American Academy since 2019, is Chairman and Director Emeritus of the McKinsey Global Institute and Senior Partner Emeritus of McKinsey & Company, where he spent twenty-six years. He was appointed by President Obama as Vice Chair of the Global Development Council at the White House (2012–2017), and by two U.S. Commerce Secretaries to the Digital Economy Board and the National Innovation Board. He is a Distinguished Fellow of Stanford’s Human-Centered AI Institute, a Distinguished Research Fellow in Ethics & AI at Oxford, and a Research Fellow of DeepMind. He is a Visiting Professor at Oxford University’s Blavatnik School of Government. In early 2022, he joined Google as Senior Vice President for Technology and Society.

ENDNOTES

- ¹ The Turing Test was conceived by Alan Turing in 1950 as a way of testing whether a computer’s responses are indistinguishable from those of a human. Though it is often discussed in popular culture as a test for artificial intelligence, many researchers do not consider it a test of artificial intelligence; Turing himself called it “the imitation game.” Alan M. Turing, “Computing Machinery and Intelligence,” *Mind*, October 1950.
- ² “The Reith Lectures: Living with Artificial Intelligence,” BBC, <https://www.bbc.co.uk/programmes/m001216k>.
- ³ Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, et al., “MuZero: Mastering Go, Chess, Shogi and Atari without Rules,” DeepMind, December 23, 2020; and Christopher Berner, Greg Brockman, Brooke Chan, et al., “Dota 2 with Large Scale Deep Reinforcement Learning,” arXiv (2019), <https://arxiv.org/abs/1912.06680>.
- ⁴ The Department of Energy’s report on AI for science provides an extensive review of both the current state-of-the-art uses of AI in various branches of science as well as the grand challenges for AI in each. See Rick Stevens, Valerie Taylor, Jeff Nichols, et al., *AI for Science: Report on the Department of Energy (DOE) on Artificial Intelligence (AI) for Science* (Oak Ridge, Tenn.: U.S. Department of Energy Office of Scientific and Technical Information, 2020), <https://doi.org/10.2172/1604756>. See also the Royal Society and the Alan Turing Institute, “The AI Revolution in Scientific Research” (London: The Royal Society, 2019).
- ⁵ See Kathryn Tunyasuvunakool, Jonas Adler, Zachary Wu, et al., “Highly Accurate Protein Structure Prediction for the Human Proteome,” *Nature* 596 (7873) (2021); Janet Thornton and colleagues discuss the contributions of AlphaFold to the life sciences, including its use in predicting the structure of some of the proteins associated with SARS-CoV-2, the virus that causes COVID-19. See Janet Thornton, Roman A. Laskowski, and Neera Borkakoti, “AlphaFold Heralds a Data-Driven Revolution in Biology and Medicine,” *Nature Medicine* 27 (10) (2021).
- ⁶ “AI Set to Exceed Human Brain Power,” CNN, August 9, 2006, <http://edition.cnn.com/2006/TECH/science/07/24/ai.bostrom/>.
- ⁷ See Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner, “Machine Bias,” ProPublica, May 23, 2016, <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>; and Joy Buolamwini and Timnit Gebru, “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification,” in *Proceedings*

of the 1st Conference on Fairness, Accountability and Transparency (New York: Association for Computing Machinery, 2018).

- ⁸ See Hilary Putnam, “Much Ado About Not Very Much,” *Daedalus* 117 (1) (Winter 1988): 279, https://www.amacad.org/sites/default/files/daedalus/downloads/Daedalus_Wi98_Artificial-Intelligence.pdf. In the same volume, see also Daniel Dennett’s essay “When Philosophers Encounter Artificial Intelligence,” in which he provides a robust response to Putman while also making observations about AI and philosophy that, with the benefit of hindsight, remain insightful today, even as the field has progressed.
- ⁹ Robert K. Appiah, Jean H. Daigle, James M. Manyika, and Themuso Makhurane, “Modeling and Training of Artificial Neural Networks,” *African Journal of Science and Technology Series B, Science* 6 (1) (1992).
- ¹⁰ Founded by Eric Horvitz, the 100-Year Study of AI that I have been involved in publishes a report every five years; its most recent report takes stock of progress in AI as well as concerns as it is more widely deployed in society. See Michael L. Littman, Ifeoma Ajunwa, Guy Berger, et al., *Gathering Strength, Gathering Storms: The One Hundred Year Study on Artificial Intelligence (AI100) 2021 Study Panel Report* (Stanford, Calif.: Stanford University, 2021). Separately, at the AI Index, we provide an annual view of developments in AI. See Artificial Intelligence Index, Stanford University Human-Centered Artificial Intelligence, <https://aiindex.stanford.edu/>.
- ¹¹ B. J. Copeland, “Artificial Intelligence,” *Britannica*, <https://www.britannica.com/technology/artificial-intelligence> (last edited December 14, 2021).
- ¹² David Poole, Alan Mackworth, and Randy Goebel, *Computational Intelligence: A Logical Approach* (New York: Oxford University Press, 1998).
- ¹³ The goal-orientation in this second type of definition is considered by some also as limiting, hence variations such as Stuart Russell and Peter Norvig’s, that focus on perceiving and acting. See Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. (Hoboken, N.J.: Pearson, 2021). See also Shane Legg and Marcus Hutter, “A Collection of Definitions of Intelligence,” arXiv (2007), <https://arxiv.org/abs/0706.3639>.
- ¹⁴ See Marvin Minsky, *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind* (New York: Simon & Schuster, 2007).
- ¹⁵ See Stephen Cave and Kanta Dihal, “Ancient Dreams of Intelligent Machines: 3,000 Years of Robots,” *Nature* 559 (7715) (2018).
- ¹⁶ Dennett, “When Philosophers Encounter Artificial Intelligence.”
- ¹⁷ John McCarthy, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon, “A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence,” August 31, 1955, <http://raysolomonoff.com/dartmouth/boxa/dart564props.pdf>.
- ¹⁸ Many of the pioneers of the current AI spring and their views are featured in Martin Ford, *Architects of Intelligence: The Truth about AI from the People Building It* (Birmingham, United Kingdom: Packt Publishing, 2018).
- ¹⁹ Yoshua Bengio, Yann Lecun, and Geoffrey Hinton, “Deep Learning for AI,” *Communications of the ACM* 64 (7) (2021). Reinforcement learning adds the notion of learning through sequential experiences that involve state transitions and making use of reinforcing rewards. See Richard S. Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction* (Cambridge, Mass.: MIT Press, 2018).

- ²⁰ Hubert L. Dreyfus and Stuart E. Dreyfus, “Making a Mind Versus Modeling the Brain : Artificial Intelligence Back at a Branchpoint,” *Daedalus* 117 (1) (Winter 1988) : 15–44, https://www.amacad.org/sites/default/files/daedalus/downloads/Daedalus_Wi98_Artificial-Intelligence.pdf.
- ²¹ For a view on trends in performance versus benchmarks in various AI subfields, see Chapter 2 in Human-Centered Artificial Intelligence, *Artificial Intelligence Index Report 2022* (Stanford, Calif. : Stanford University, 2022), https://aiindex.stanford.edu/wp-content/uploads/2022/03/2022-AI-Index-Report_Master.pdf.
- ²² At the time of developing this volume (2020–2021), the most well-known large language models included OpenAI’s GPT-3, Google’s LaMDA, Microsoft’s MT-NLG, and DeepMind’s Gopher. These models use transformer architectures first described in Ashish Vaswani, Noam Shazeer, Niki Parmar, et al., “Attention Is All You Need,” arXiv (2017), <https://arxiv.org/abs/1706.03762>.
- ²³ Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, et al., “On the Opportunities and Risks of Foundation Models,” arXiv (2021), <https://arxiv.org/abs/2108.07258>.
- ²⁴ Ibid. See also Laura Weidinger, John Mellor, Maribeth Rauh, et al., “Ethical and Social Risks of Harm from Language Models,” arXiv (2021), <https://arxiv.org/abs/2112.04359>. On toxicity, see Samuel Gehman, Suchin Gururangan, Maarten Sap, et al., “RealToxicity-Prompts : Evaluating Neural Toxic Degeneration in Language Models,” in *Findings of the Association for Computational Linguistics : EMNLP 2020* (Stroudsburg, Pa. : Association for Computational Linguistics, 2020), 3356–3369 ; and Albert Xu, Eshaan Pathak, Eric Wallace, et al., “Detoxifying Language Models Risks Marginalizing Minority Voices,” arXiv (2014), <https://arxiv.org/abs/2104.06390>.
- ²⁵ Demis Hassabis, Dhharshan Kumaran, Christopher Summerfield, and Matthew Botvinick, “Neuroscience-Inspired Artificial Intelligence,” *Neuron* 95 (2) (2017); and Alexis T. Baria and Keith Cross, “The Brain Is a Computer Is a Brain : Neuroscience’s Internal Debate and the Social Significance of the Computational Metaphor,” arXiv (2021), <https://arxiv.org/abs/2107.14042>.
- ²⁶ See Littman, *Gathering Strength, Gathering Storms*.
- ²⁷ For an overview of trends in AI technical and ethics issues as well as AI regulation and policy, see Chapters 3 and 6, respectively, in Human-Centered Artificial Intelligence, *Artificial Intelligence Index Report 2022*. See also Mateusz Szczepański, Michał Choraś, Marek Pawlicki, and Aleksandra Pawlicka, “The Methods and Approaches of Explainable Artificial Intelligence,” in *Computational Science – ICCS 2021*, ed. Maciej Paszynski, Dieter Kranzlmüller, Valeria V. Krzhizhanovskaya, et al. (Cham, Switzerland : Springer, 2021). See also Cynthia Dwork and Aaron Roth, “The Algorithmic Foundations of Differential Privacy,” *Foundations and Trends in Theoretical Computer Science* 9 (3–4) (2014).
- ²⁸ For an overview of the types of efforts as well as three case studies (Microsoft, OpenAI, and OECD’s observatory), see Jessica Cussins Newman, *Decision Points in AI Governance : Three Case Studies Explore Efforts to Operationalize AI Principles* (Berkeley : Center for Long-Term Cybersecurity, UC Berkeley, 2020).
- ²⁹ See Chapter 6, “Diversity in AI,” in Human-Centered Artificial Intelligence, *Artificial Intelligence Index Report 2021* (Stanford, Calif. : Stanford University, 2021), https://aiindex.stanford.edu/wp-content/uploads/2021/11/2021-AI-Index-Report_Master.pdf. See also Sarah Myers West, Meredith Whittaker, and Kate Crawford, “Discriminating Systems : Gender, Race and Power in AI” (New York : AI Now Institute, 2019).

- ³⁰ Bengio et al., “Deep Learning for AI.”
- ³¹ Emily B. Bender and Alexander Koller, “Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (New York: Association for Computing Machinery, 2020).
- ³² See Allan Dafoe, Edward Hughes, Yoram Bachrach, et al., “Open Problems in Cooperative AI,” arXiv (2020), <http://arxiv.org/abs/2012.08630>; and Allan Dafoe, Yoram Barach, Gillian Hadfield, Eric Horvitz, et al., “Cooperative AI: Machines Must Learn to Find Common Ground,” *Nature* 593 (2021).
- ³³ See the Department of Energy’s report *AI for Science* for examples in several scientific fields. See also Alex Davies, Petar Veličković, Lars Buesing, et al., “Advancing Mathematics by Guiding Human Intuition with AI,” *Nature* 600 (7887) (2021); and Anil Ananthaswamy, “AI Designs Quantum Physics Experiments Beyond What Any Human Has Conceived,” *Scientific American*, July 2021.
- ³⁴ On AI’s grand challenges, Raj Reddy posed probably the first list in his 1988 AAAI Presidential Address, “Foundations and Grand Challenges of Artificial Intelligence,” *AI Magazine*, 1988. Ganesh Manni provides a useful history of AI grand challenges in “Artificial Intelligence’s Grand Challenges: Past, Present, and Future,” *AI Magazine*, Spring 2021.
- ³⁵ On the challenges and progress in causal reasoning, see Judea Pearl and Dana Mackenzie, *The Book of Why: The New Science of Cause and Effect* (New York: Basic Books, 2018).
- ³⁶ For example, see David Silver, Satinder Singh, Doina Precup, and Richard S. Sutton, “Reward is Enough,” *Artificial Intelligence* 299 (4) (2021).
- ³⁷ In a recent paper, Chinese researchers describe an approach that has the potential to train models of up to 174 trillion parameters, a size that rivals the number of synapses in the brain (hence the claim of “brain-scale” models), on high performance supercomputers. See Zixuan Ma, Jiaao He, Jiezhong Qiu, et al., “BaGuaLu: Targeting Brain Scale Pretrained Models with over 37 Million Cores,” March 2022, <https://keg.cs.tsinghua.edu.cn/jietang/publications/PPOPP22-Ma%20et%20al.-BaGuaLu%20Targeting%20Brain%20Scale%20Pretrained%20Models%20ow.pdf>.
- ³⁸ See Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford: Oxford University Press, 2014); Max Tegmark, *Life 3.0: Being Human in the Age of AI* (New York: Knopf, 2017); and Martin Ford, *Rule of the Robots: How Artificial Intelligence Will Transform Everything* (New York: Basic Books, 2021).
- ³⁹ See Michael Chui, James Manyika, Mehdi Miremadi, et al., “Notes from the AI Frontier: Applications and Value of Deep Learning” (New York: McKinsey Global Institute, 2018); and Jacques Bughin, Jeongmin Seong, James Manyika, et al., “Notes from the AI Frontier: Modeling the Impact of AI on the World Economy” (New York: McKinsey Global Institute, 2018). And for trends on adoption of AI in business and the economy as well as AI labor markets, see Chapter 4 in *Human-Centered Artificial Intelligence, Artificial Intelligence Index Report 2022*.
- ⁴⁰ See National Commission on Technology, Automation and Economic Progress, *Technology and the American Economy*, vol. 1 (Washington, D.C.: U.S. Government Printing Office, 1966), <https://files.eric.ed.gov/fulltext/ED023803.pdf>.
- ⁴¹ James Manyika, Susan Lund, Michael Chui, et al., *Jobs Lost, Jobs Gained: What the Future of Work Will Mean for Jobs, Skills, and Wages* (New York: McKinsey Global Institute, 2017); Daron Acemoglu and Pascual Restrepo, “Artificial Intelligence, Automation and Work,”

- NBER Working Paper 24196 (Cambridge, Mass. : National Bureau of Economic Research, 2018); David Autor, David Mindell, and Elisabeth Reynolds, *The Work of the Future: Building Better Jobs in an Age of Intelligent Machines* (Cambridge, Mass. : MIT Work of the Future, 2020); and Erik Brynjolfsson, “The Problem Is Wages, Not Jobs,” in *Redesigning AI: Work, Democracy, and Justice in the Age of Automation*, ed. Daron Acemoglu (Cambridge, Mass. : MIT Press, 2021).
- ⁴² See Daron Acemoglu and Pascual Restrepo, “The Wrong Kind of AI? Artificial Intelligence and the Future of Labor Demand,” NBER Working Paper 25682 (Cambridge, Mass. : National Bureau of Economic Research, 2019); and Bryan Wilder, Eric Horvitz, and Ece Kamar, “Learning to Complement Humans,” arXiv (2020), <https://arxiv.org/abs/2005.00582>.
- ⁴³ Susskind provides a broad survey of many of the arguments that AI has changed everything with respect to jobs. See Daniel Susskind, *A World Without Work: Technology, Automation, and How We Should Respond* (New York: Metropolitan Books, 2020).
- ⁴⁴ From their 1957 lecture in Herbert A. Simon and Allen Newell, “Heuristic Problem Solving: The Next Advance Operations Research,” *Operations Research* 6 (1) (1958).
- ⁴⁵ John Maynard Keynes, “Economic Possibilities for Our Grandchildren,” in *Essays in Persuasion* (New York: Harcourt Brace, 1932), 358–373.
- ⁴⁶ See our most recent annual AI Index report, Human-Centered Artificial Intelligence, *Artificial Intelligence Index Report 2022*. See also Daniel Castro, Michael McLaughlin, and Eline Chivot, “Who Is Winning the AI Race: China, the EU or the United States?” Center for Data Innovation, August 19, 2019; and Daitian Li, Tony W. Tong, and Yangao Xiao, “Is China Emerging as the Global Leader in AI?” *Harvard Business Review*, February 18, 2021.
- ⁴⁷ Tania Babina, Anastassia Fedyk, Alex Xi He, and James Hodson, “Artificial Intelligence, Firm Growth, and Product Innovation” (2021), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3651052.
- ⁴⁸ See Amanda Askill, Miles Brundage, and Gillian Hadfield, “The Role of Cooperation in Responsible AI Development,” arXiv (2019), <https://arxiv.org/abs/1907.04534>.
- ⁴⁹ See Henry A. Kissinger, Eric Schmidt, and Daniel Huttenlocher, *The Age of AI: And Our Human Future* (Boston: Little, Brown and Company, 2021).
- ⁵⁰ Issues that we explored in a Council on Foreign Relations Taskforce on Innovation and National Security. See James Manyika and William H. McRaven, *Innovation and National Security: Keeping Our Edge* (New York: Council on Foreign Relations, 2019).
- ⁵¹ For an assessment of the potential contributions of AI to many of the global development challenges, as well as gaps and risks, see Michael Chui, Martin Harryson, James Manyika, et al., “Notes from the AI Frontier: Applying AI for Social Good” (New York: McKinsey Global Institute, 2018). See also Ricardo Vinuesa, Hossein Azizpour, Iolanda Leita, et al., “The Role of Artificial Intelligence in Achieving the Sustainable Development Goals,” *Nature Communications* 11 (1) (2022).
- ⁵² Timnit Gebru, Jamie Morgenstern, Briana Vecchione, et al., “Datasheets for Datasets,” arXiv (2021), <https://arxiv.org/abs/1803.09010>.
- ⁵³ See also Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: Public Affairs, 2019).
- ⁵⁴ See Dennett, “When Philosophers Encounter Artificial Intelligence.”

- ⁵⁵ Pamela McCorduck, *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*, 2nd ed. (Abingdon-on-Thames, United Kingdom: Routledge, 2004).
- ⁵⁶ Timnit Gebru, "Race and Gender," in *The Oxford Handbook of Ethics of AI*, ed. Markus D. Dubber, Frank Pasquale, and Sunit Das (Oxford: Oxford University Press, 2020).
- ⁵⁷ Blaise Agüera y Arcas, Margaret Mitchell, and Alexander Todorov, "Physiognomy in the Age of AI" (forthcoming).
- ⁵⁸ See Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner, "Machine Bias: There's Software Used across the Country to Predict Future Criminals. And It's Biased against Blacks," *ProPublica*, May 23, 2016; Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor* (New York: St. Martin's Press, 2018); and Maximilian Kasy and Rediet Abebe, "Fairness, Equality, and Power in Algorithmic Decision-Making," in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (New York: Association for Computing Machinery, 2021).
- ⁵⁹ See Sorelle A. Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian, "On the (im)Possibility of Fairness," arXiv (2016), <https://arxiv.org/abs/1609.07236>; and Arvind Narayanan, "Translation Tutorial: 21 Fairness Definitions and their Politics," in *Proceedings of the 2018 Conference on Fairness, Accountability, and Transparency* (New York: Association for Computing Machinery, 2018).
- ⁶⁰ Dennett, "When Philosophers Encounter Artificial Intelligence."
- ⁶¹ See Sendhil Mullainathan, "Algorithmic Fairness and the Social Welfare Function," in *Proceedings of the 2018 ACM Conference on Economics and Computation* (New York: Association for Computing Machinery, 2018). See also Lily Hu and Yiling Chen, "Fair Classification and Social Welfare," in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (New York: Association for Computing Machinery, 2020), 535–545; and Hoda Heidari, Claudio Ferrari, Krishna Gummadi, and Andreas Krause, "Fairness Behind a Veil of Ignorance: A Welfare Analysis for Automated Decision Making," *Advances in Neural Information Processing Systems* 31 (2018): 1265–1276.
- ⁶² See discussion in Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford: Oxford University Press, 2014); and Stuart Russell, *Human Compatible: Artificial Intelligence and the Problem of Control* (New York: Viking, 2019).
- ⁶³ Iason Gabriel, "Artificial Intelligence, Values, and Alignment," *Minds and Machines* 30 (3) (2020).
- ⁶⁴ At an AI conference organized by the Future of Life Institute, we generated a list of priorities for robust and beneficial AI. See Stuart Russell, Daniel Dewey, and Max Tegmark, "Research Priorities for Robust and Beneficial Artificial Intelligence," *AI Magazine*, Winter 2015. See also the issues raised in Littman, *Gathering Strength, Gathering Storms*.
- ⁶⁵ Such a working list in response to the 2050 thought experiment can be found at "AI2050's Hard Problems Working List," https://drive.google.com/file/d/1toSEnQszftuW9-RikM76oJSuP-Heauq_/view (accessed February 17, 2022).