Darrell M. West John R. Allen

TURNING POINT

Policymaking in the Era of Artificial Intelligence

THE BROOKINGS INSTITUTION
1775 Massachusetts Avenue, N.W.
Washington, D.C. 20036
www.brookings.edu

Library of Congress Control Number: 2020936148
ISBN 9780815738596 (hc)
ISBN 9780815738602 (ebook)

9 8 7 6 5 4 3 2 1

Typeset in Garamond Premier Pro

Composition by Elliott Beard

CONTENTS

PREFACE

A group of technologists was asked by the Pew Research Center whether artificial intelligence (AI) would empower individuals and enhance human capacities. Sixty-three percent of these experts answered in the affirmative and indicated most people would be better off.[1] As shown below, many of these visionaries foresaw a bright future of better health, wider prosperity, and improved access to information:

ERIK BRYNJOLFSSON, *director of the Massachusetts Institute of Technology (MIT) Initiative on the Digital Economy:* "We can virtually eliminate global poverty, massively reduce disease and provide better education to almost everyone on the planet."

LEONARD KLEINROCK, *Internet Hall of Fame member:* "As AI and machine learning improve, we will see highly customized interactions between humans and their healthcare needs."

MICAH ALTMAN, *head scientist in the program on information science at MIT Libraries:* "These technologies will help to adapt learning to the needs of each individual by translating language, aiding memory and

providing us feedback on our own emotional and cognitive state and on the environment."

STEVE CROCKER, *Internet Hall of Fame member:* "Encyclopedic knowledge [will be] available at our fingertips."

KEN BIRMAN, *professor of computer science at Cornell University:* "I believe that our homes and offices will have evolved to support app-like functionality. . . . People will customize their living and working spaces."

JOHN HAVENS, *executive director of the Institute of Electrical and Electronics Engineers Global Initiative on Ethics of Autonomous and Intelligent Systems:* "By utilizing blockchain or similar technologies and adopting progressive ideals toward citizens and their data . . . we can usher in genuine digital democracy in the age of the algorithm."[2]

Yet beneath that general optimism, a number of prominent thinkers also expressed doubts regarding these rosy scenarios. They worried about wealth concentration, algorithmic bias, and political authoritarianism. Witness these ominous predictions:

ERIK BRYNJOLFSSON, *MIT:* "AI and ML [machine learning] can also be used to increasingly concentrate wealth and power, leaving many people behind."

ANDREW MCLAUGHLIN, *executive director of the Yale University Center for Innovative Thinking:* "[It will] enable hidden discrimination and arbitrary penalization of individuals in areas like insurance, job seeking and performance assessment."

DANAH BOYD, *principal researcher for Microsoft and president of the Data & Society Research Institute:* "There will be abuses of power that involve AI."

NATHANIEL BORENSTEIN, *chief scientist at Mimecast:* "I foresee a world in which IT and so-called AI produce an ever-increasing set of minor benefits, while simultaneously eroding human agency and privacy and supporting authoritarian forms of governance."

SONIA KATYAL, *co-director of the Berkeley Center for Law and Technology:* "Questions about privacy, speech, the right of assembly and technological construction of personhood will all reemerge in this new AI context."

MARINA GORBIS, *executive director of the Institute for the Future:* "[AI] is likely to create greater economic inequalities."

AMY WEBB, *founder of the Future Today Institute:* "The social safety net structures currently in place in the U.S. and in many other countries around the world weren't designed for our transition to AI."

THAD HALL, *coauthor of* Politics for a Connected American Public*:* "Fake videos, audio and similar media are likely to explode and create a world where 'reality' is hard to discern."[3]

This litany of concerns suggests we are at a watershed moment in human history. Based on these expert assessments, there are vigorous debates regarding whether the emerging digital world will usher in a utopia of good health and widespread prosperity or a dystopian world of unfettered surveillance, income inequality, and pervasive bias. An online survey of over 34,000 adults from around the world undertaken by the 2020 Edelman Trust Barometer echoed experts' fears when it found that 61 percent of adults believe "the pace of change in technology is too fast" and 66 percent "worry technology will make it impossible to know if what people are seeing or hearing is real."[4]

The anxiety among experts and the general public is pertinent because each group understands we are on the cusp of numerous AI deployments, and humankind will face choices that no previous wave of technology has demanded. Self-driving vehicles powered by AI already are being pilot tested; students are being assigned to charter schools via AI algorithms; companies are using AI on their e-commerce sites to make product recommendations; healthcare is being assessed through data analytics and AI; and warfare is incorporating autonomous weapons in national defense.

In this book, we examine the use of artificial intelligence and emerging technologies in healthcare, education, transportation, e-commerce,

and national defense. We study these topics because each area represents a major sector where AI is poised to alter how businesses, governments, and consumers operate. Using a number of different illustrations, we ask how emerging technologies are being utilized, what opportunities are being created, and which risks are arising.

It is crucial to analyze AI implementation in specific domains to decipher its full implications for governance, society, and ethics. One cannot address big questions about digital health, personalized education, autonomous vehicles, e-commerce, and national defense without a thorough discussion of how AI is being deployed. We believe regulations and policies will be crucial in determining AI's future, as will corporate decisions, legal liability rules, and consumer sentiments.

In addition, we analyze how AI affects broader geopolitical considerations in a world bedeviled by pandemics, income inequality, and climate change, among other challenges. International relations is changing rapidly, and a number of nations are making tremendous progress in their technological capabilities. China, for example, is expected to become the world's largest economy by 2050 and a global leader in AI. It is putting considerable money into AI-related commerce and defense. There also are important AI developments in France, Germany, the United Kingdom, Russia, India, Israel, and elsewhere that will affect how AI expands.

To preview our argument, we do not think technologies in and of themselves will preordain utopia or dystopia, but people will control the future through the choices they make in the coming period. Humans will interact in complex ways with digital technology, and analysts will not be able to resolve broader concerns about human safety, individual privacy, national security, and societal well-being without assessing AI's particular characteristics. However AI unfolds, the resulting tapestry will be a product of policy and operational decisions made in the next several years. It is important to limit the possibility of an inadvertent or accidental dystopia that could emerge from poorly-thought-out choices.

Chapter 1 defines AI and provides examples of how it is being used. We argue AI is a multifaceted technology with the ability to analyze, learn, and adapt to changing circumstances. As such, it has enormous

power to alter how people make decisions and the manner in which organizations operate. The reason AI is becoming so prevalent today is advances in computing power over the past couple of decades have improved the ability to analyze large data sets and develop sophisticated algorithms that act on those insights. Combined with machine learning and data analytics, AI has the means to interpret information in real time and help humans make more informed decisions. In thinking about AI's deployment, our goal should be to use it to advance humanity while minimizing detrimental governance, ethical, legal, and geopolitical consequences. Since AI is not value-neutral, technologists must be careful data analytics do not lead to unfair outcomes for segments of the population.

Chapter 2 explores the growing reliance on AI in healthcare. Algorithms are helping providers diagnose illnesses, treat diseases, improve clinical trials, identify new drugs, and detect financial fraud. These are especially important virtues at a time when the coronavirus has inflicted serious harm on health and the economy. In combination with wearable devices, remote monitoring tools, data analytics, and machine learning, AI is bringing new tools into the medical arena. Through natural language processing, it is possible to analyze thousands of scientific articles for insights that can lead to new drugs, medical treatments, or clinical processes. That could expedite drug discovery and improve the treatment of large numbers of patients. Yet medical algorithms also introduce a number of possible biases into healthcare based on race, gender, age, income, and geography. For these reasons, it is important that policies and practices limit biases based on incomplete or misleading data.

Chapter 3 suggests ways in which AI can help with school enrollment decisions, personalizing and assessing learning, assisting teachers, dealing with at-risk students, and protecting against school violence. To accomplish these tasks, though, requires improving access to electronic resources, promoting greater technical literacy, making sure algorithms do not reinforce societal inequities, and making better use of digital learning platforms. Both international competitiveness and national security depend on having an equitable, inclusive, and forward-looking ed-

ucational system. Without widespread access to high-speed broadband, advanced technologies could widen inequality and create information "haves" and "have-nots." Thus it is crucial that everyone has access to digital resources so that important segments of the society are not left behind.

Chapter 4 examines AI in transportation through the use of advanced data analytics in autonomous vehicles. Advances in artificial intelligence and machine learning allow onboard computers connected to cloud platforms to integrate data instantly and independently power cars, trucks, buses, planes, and trains. Over coming years, driverless vehicles guided by AI will move onto roads and affect transportation in major cities around the world. Connected vehicles are likely to improve highway safety, alleviate traffic congestion, and reduce air pollution and carbon emissions, but designers must overcome poor infrastructure, policy and regulatory barriers, cybersecurity challenges, privacy concerns, and public doubts regarding the safety of autonomous vehicles. As society navigates this transition to a world of autonomous vehicles in future decades, there will be consequences for drivers, auto manufacturers, insurance companies, and transportation policymakers. Drivers and insurers, in particular, will face a complex situation because there likely will be fewer accidents and, therefore, less expensive insurance policies, but much greater information regarding vehicular operations from sensors and cameras in cars and trucks. In that situation, the analysis of accidents will rely much less on witness testimony and more on digital data regarding vehicles, software, and operators.

Chapter 5 reviews advances in AI, machine learning, and data analytics in e-commerce. Drawing on several cases, we study how e-commerce is accelerating, especially due to the coronavirus. We look at e-commerce growth and ways to overcome current barriers. There needs to be progress on digital infrastructure, delivery mechanisms, zoning rules, tax policy, dynamic pricing, labor laws, and the handling of data breaches to build an e-commerce system that works for employers, employees, and communities.

Chapter 6 calls for technology innovation in national defense. For the

United States, security and defense capabilities need to get faster, smarter, and better integrated through AI. Because of the accelerated pace of conflict, countries will conduct simultaneous operations in many warfare domains, and the cyber domain increasingly will become the decisive realm of conflict. There will be major opportunities to deploy AI-based tools to improve decisionmaking, speed, and scalability as well as to address a variety of rising threats. As with many AI applications, policy and operational shifts will be necessary in order to facilitate the proper integration of emerging technologies and make sure they strengthen leadership capacity and military performance.

Having surveyed AI applications in many areas, Chapter 7 looks at public reactions to artificial intelligence, robotics, and facial recognition software. Using national public opinion surveys, we measure how people are reacting to technological deployments, the opportunities they see, and the fears they hold. At a time of digital change, there is a discernible "techlash"—that is, a backlash against emerging technologies and the sector as a whole. Technology companies and other proponents must determine how to assuage human fears about technology, manage privacy issues, and deal effectively with people's concerns.

Chapter 8 examines the ethical aspects of artificial intelligence and presents several recommendations for incorporating ethics in decisionmaking. As society moves forward into an AI-based world, there are legitimate worries about a lack of transparency, poor accountability, unfairness, and possible biases in automated and autonomous tools. With millions of lines of code in each application, it is difficult to know what values are inculcated in software and how algorithms actually reach decisions. Corporate coders increasingly are digital sovereigns who set the rules of the road and terms of service for consumers. What they decide, directly or indirectly, has far-reaching consequences for the people who are using their products. In this situation, it is important to analyze AI ethics in "sensitive-use" cases such as criminal justice, border enforcement, and government surveillance. To protect human values, we recommend developing responsible AI principles based on fairness, transparency, and human safety; hiring ethicists, designing codes of ethics, and developing

AI ethics review boards; and having annotated software and AI audit trails, mandating AI training programs, and providing a means of remediation in cases of discernible consumer harm.

Chapter 9 makes a number of proposals for building responsible AI and emerging technologies. We argue that advancing technology requires many actions on the part of business executives and political leaders:

- addressing AI's challenges in a time of pandemics and megachange;
- improving governance through distributed collaboration;
- establishing guiding ethical principles;
- adopting horizontal and vertical rules;
- strengthening oversight through AI impact assessments;
- restoring the Office of Technology Assessment;
- creating advisory boards of relevant stakeholders for federal agencies;
- defining corporate culpability;
- enforcing current laws;
- improving digital access;
- reducing algorithmic biases through third-party audits;
- moving beyond notice and consent requirements for personal privacy;
- using insurance to mitigate exposure to AI risks;
- diversifying decisionmaking;
- penalizing malicious abuses of new technologies;
- establishing a national research cloud;
- developing a national data strategy;
- addressing geographic disparities and workforce training; and
- ensuring human oversight of technology.

A preliminary version of these ideas was presented in our 2018 Brookings Institution paper entitled "How Artificial Intelligence Is Transform-

ing the World." In it, we looked at AI applications in several areas and the policy, legal, and ethical issues associated with them. This book also draws on our AI work analyzing ethics, public opinion, and transportation, among other areas.

We would like to thank the organizations that invited us to make AI presentations in the United States and around the world. Those engaging conversations with leading experts helped us refine our ideas and test our recommendations. Kevin Desouza, Dipayan Ghosh, Frank Rose, Tom Stefanick, and Tom Wheeler provided valuable feedback and made a number of suggestions that improved the book. Alex Engler, Chris Meserole, and John Villasenor reviewed our glossary of technical terms and helped clarify our definitions. We also wish to thank a number of individuals for their research help. Quinn Bornstein, Miku Fujita, Raj Gambhir, Grace Gilberg Gedye, Gabrielle Jackson, Amritha Jayanti, Jack Karsten, Jacob Lineberry, Lia Newman, Hillary Schaub, Julia Slisz, and Kristjan Tomasson provided valuable assistance on this project. Julia Slisz helped to write the healthcare chapter. Corey Broschak and Jeesue Lee contributed substantially to the defense chapter. Thanks also go to William Finan and Yelba Quinn of the Brookings Institution Press for their sound advice and assistance with this project. Carla Huelsenbeck did an excellent job copy editing the manuscript and Cecilia González ably shepherded the book through the production process. None of the people cited here is responsible for the analysis or interpretations presented in this book.

TURNING POINT

ONE | WHAT IS AI?

Artificial intelligence (AI) is a difficult concept to grasp. As an illustration, only 17 percent of 1,500 U.S. business leaders in 2017 said they were familiar with how AI would affect their companies.[1] These executives understood there was considerable potential for revolutionizing business processes but were not clear how AI would be deployed within their own organizations or how it would alter their industries.

Hollywood offers little help in improving people's understanding of advanced technologies. Many movies conflate AI with malevolent robots or hyperintelligent beings such as the Terminator or the evil HAL in Arthur C. Clarke's *2001: A Space Odyssey*. In film depictions, superpowered entities inevitably gain humanlike intelligence, go rogue, and inflict tremendous harm on the human population. The ominous message in these cinematic renditions is that AI is dangerous because it will eventually enslave humans.[2]

One of the difficulties in understanding AI is the lack of a uniform definition. People often intertwine many things in their conceptions and then imagine the worst possible outcomes. They assume advanced technologies will have omniscient capabilities, destructive motivations, and lim-

ited human oversight and will be impossible to control. For those reasons, it is important to clarify what we mean by artificial intelligence, provide understandable examples of how it is being used, and outline its major risks.

AI ORIGINS

Alan Turing generally is credited with conceptualizing the idea of AI in 1950, when he speculated about "thinking machines" that could reason at the level of a human being. His well-known "Turing Test" specifies that computers need to complete "reasoning puzzles" as effectively as humans to be considered autonomous.[3]

Turing was followed a few years later by visionary John McCarthy. In 1956, he was the first scientist to coin the term "artificial intelligence" to denote machines that could think and act independently. He defined AI as "getting a computer to do things which, when done by people, are said to involve intelligence."

Since that time, scientists have argued over what constitutes "thinking," "intelligence," and "autonomy" when it comes to hardware and software. Advanced computers such as IBM's Watson already have beaten humans at chess and poker, and are capable of instantly processing enormous amounts of information.[4] Deep Mind's AlphaGo and AlphaGo Master have advanced even further and used AI to defeat expert players of the board game Go.[5] Learning from a series of competitions, these AI-based systems figured out how to make complex moves and develop strategies that humans never had considered.

More recently, futurist Ray Kurzweil raised the AI notion to a new level when he predicted a "singularity," that is, a "machine-based superintelligence [that is] greater than human intelligence."[6] With advances in AI, data analytics, and machine learning, it no longer seems far-fetched to foresee sentient machines that can take on advanced functions.[7] As computing power grows, Kurzweil believes there will be a convergence between humans and computers.

The wide scope in how people define artificial intelligence lends itself to a broad range of interpretations regarding its societal impact. Those who

are worried about a loss of humanity can point to plausible scenarios where robots make everything and Terminator-like war machines operate independently of human control. Conversely, people who are more optimistic can point to ways that automated machines could spread convenience, relieve humans of boring or dangerous tasks, and improve the overall quality of life. Being specific about what we mean by AI is important for how observers evaluate the possible array of opportunities and risks.

DEFINING AI

Engineers Shukla Shubhendu and Jaiswal Vijay define AI as "machines that respond to stimulation consistent with traditional responses from humans, given the human capacity for contemplation, judgment, and intention."[8] This interpretation is helpful because it suggests several qualities that separate AI from mechanical devices or traditional computer software: intentionality, intelligence, and adaptability. It is these features that enable AI algorithms to find patterns or associations through "neural networks" that group data based on common similarities (see Glossary of Key Terms). In addition, AI-based systems learn from insights gathered via "deep learning" techniques that use statistics to spot underlying trends or patterns in data and apply that knowledge to changing circumstances. And it is not just numbers that can be mined, as AI can employ "natural language processing" that analyzes textual information in order to make sense of its meaning. As we outline below, AI employs a number of sophisticated traits to bring powerful computational strengths to a wide range of endeavors.[9]

Intentionality

Artificial intelligence algorithms are designed to undertake actions or make decisions using real-time data. They are unlike passive machines that are only capable of static or mechanical responses. Using sensors, digital data, or remote inputs, advanced algorithms combine information from a variety of sources, analyze the material instantly, and act on the insights derived from those data.[10]

With massive improvements in storage systems, processing speeds, and analytic techniques, AI algorithms are capable of tremendous sophistication in analysis and decisionmaking.[11] For example, financial algorithms can spot minute differentials in stock valuations and undertake market transactions based on that information. Similarly, artificially intelligent sensors on dynamic systems such as electric generators or aircraft turbine engines can sense operating anomalies far earlier than previous instrumentation and thereby save millions of dollars by signaling the need for preventive maintenance. The same logic applies in environmental sustainability systems that use sensors to determine whether someone is in a room and automatically adjust heating, cooling, and lighting utilizing that sensory input.

As long as these systems conform to human-based values, there is little risk of AI endangering human beings. Sophisticated algorithms can be intentional while analyzing information in ways that augment human performance and understanding. However, if the software is poorly designed or based on incomplete or biased information, it can make decisions that are unfair or discriminatory and therefore create either intentional or unintentional harm.

Intelligence

AI is adept at finding statistical associations through deep learning techniques. Using sophisticated data analytics, software designers can develop algorithms that identify data patterns, and use that knowledge to make specific decisions.

As we discuss in chapter 3, there are AI systems for managing school enrollments. These systems compile information on neighborhood location, desired schools, and enrollment preferences, and assign pupils based on that compiled information. As long as there is little disagreement regarding basic criteria, these systems can work intelligently and effectively.[12]

Of course, that standard often is not the case in many policy areas. Reflecting the importance of education for life outcomes, parents, teachers, and school administrators fight over the weighting of different factors.

Should students be assigned to their neighborhood school or should other criteria, such as a desire to overcome residential segregation and diversify schools, override locational considerations? Those criteria are important because in a city with widespread racial segregation and economic inequalities, emphasizing neighborhood school assignments invariably exacerbates inequality and heightens racial segregation.

This is an important point because making those kinds of decisions increasingly falls to computer programmers, many of whom are inexperienced when it comes to resolving ethical disputes or values-based disagreements. Programmers are adept at building intelligent algorithms that compile information based on specified considerations, but figuring out how to reconcile conflicting values is not something for which most of them are trained. Computer programmers' expertise is in analyzing data and writing code, not resolving societal disagreements about fundamental objectives.

Adaptability

The last quality that marks AI systems is the ability to learn and adapt as they compile information and make decisions. Effective artificial intelligence makes adjustments as circumstances or conditions shift. This may involve alterations in financial situations, road conditions, environmental considerations, or military circumstances pertaining to the AI system. In each of these areas, AI has to integrate new information quickly in its algorithms and make shifts based on relevant data.

Adaptability represents a valuable characteristic of AI because conditions change quickly in many areas and organizations need the capacity to assess needs and requirements for adaptation and make relevant decisions that can be significantly assisted and even accelerated by trained algorithms. That agility is an important benefit of autonomous algorithms over earlier systems that merely project past trends into the future in a linear manner.

Transportation illustrates how this adaptability and agility can happen. Autonomous vehicles use machine-to-machine communications to alert other cars on the road about upcoming congestion, potholes,

highway construction, or traffic impediments. Vehicles can take advantage of other vehicles' road experience, without human involvement. Everything that is learned can be immediately and fully transferable to other connected vehicles. Their advanced algorithms, sensors, and cameras integrate real-time experiences and use dashboards and visual displays to help people make sense of ongoing traffic and vehicular conditions. This is made possible by advances in computing power.

Similar capabilities are incorporated into software for scheduling appointments. Digital personal assistants can ascertain a person's preferences and respond to email requests for appointments in a dynamic manner. Without any human intervention, the digital assistant can make appointments, adjust schedules, and communicate preferences to other individuals.

Increasingly, human interactions with computers are taking place through verbal conversation. Rather than the conventional technique of point and click based on mouse navigation, people can talk to computers and provide oral instructions of what they want to do.[13] Through "chatbots" built into devices such as Amazon Alexa and Apple Siri, people can order pizza, pay bills, or request music through conversational interfaces.

The simplicity of these tools dramatically expands the possibilities for AI and human-machine interactions. Voice-activated features can turn what used to be complex interfaces into easy-to-use consumer devices. People no longer have to be computer scientists or engineers to utilize complex systems or operate sophisticated products. They can employ AI without even knowing they are interacting with an advanced algorithm. The software code may have millions of programming lines built into it but only require verbal commands to operate.

AI DEPLOYMENT

With its growing capabilities, AI is being deployed in a number of sectors. Fields ranging from finance and shopping to smart cities, energy management, and criminal justice are incorporating AI, data analytics, and machine learning. Motivated by efforts to build efficiency and improve

effectiveness, algorithms have altered the way many leaders decide and organizations operate. Below are brief descriptions of some ways in which algorithms are being utilized.

Finance

The financial services industry employs many AI applications.[14] According to industry observers, "decisions about loans are now being made by software that can take into account a variety of finely parsed data about a borrower, rather than just a credit score and a background check."[15] In addition, there are so-called robo-advisers that "create personalized investment portfolios, obviating the need for stockbrokers and financial advisers."[16] These advances take the emotion out of investing and make decisions based on analytical considerations.[17]

A prominent illustration of this is stock exchanges, where automated, high-frequency trading has replaced much of human decisionmaking. Systems make buy and sell orders based on an analysis of trading inefficiencies or financial differentials.[18] Powered in some cases by advanced computing, these tools have much greater capacities for storing information and processing data.[19]

Fraud detection is another area where AI is helpful in financial systems. It sometimes is difficult to discern fraudulent activities in large organizations, but AI can identify early or telltale abnormalities, outliers, or deviant cases. These detection systems help human managers find problems before they reach dangerous levels.[20]

For some individuals, however, AI creates consumer financial protection problems. Digital tools alter the retail experience and can expose consumers to deceptive marketing strategies or fraudulent practices. High-tech appeals may target people with limited financial literacy and encourage them to buy products for which they are not well suited. According to Brookings Institution scholar Makada Henry-Nickie, "AI can lead to a surge of wicked, legacy problems: product steering, discriminatory pricing, unfair credit rationing, exclusionary filtering, and digital redlining."[21] All of those practices harm consumer well-being and exacerbate financial inclusion problems.

In addition, AI creates the means by which financial institutions can bypass conventional nondiscrimination rules. For example, Brookings Institution scholar Aaron Klein notes that banks are prohibited from incorporating marital status into their lending decisions. Yet, algorithms can gauge the potential for marriage problems through proxy measures that examine travel, hotel, gift, and restaurant bills and therefore estimate the odds of impending divorces that could affect the ability to repay loans. When AI has the ability to assess relationship problems through indirect means, Klein argues it should be limited in the same way decisions based on marital status would be in lending decisions.[22]

Shopping

Many retail establishments have pioneered technologies designed to improve the consumer experience. Amazon, for example, has opened convenience stores with no salesclerks. A consumer downloads the company's mobile app, enters the store through turnstiles, picks out desired items, and exits through turnstiles that charge the person's credit card or mobile payment system.[23]

Giant Food Stores has introduced robots powered by AI and equipped with cameras and voice-paging capabilities that can spot food spills and empty shelves. These devices page human employees to let them know of the problem so that debris can be cleaned up and items restocked. The robots can also do price checks and answer customer questions through audio interfaces.[24] Company executives hope emerging technologies will make it easier for shoppers to find what they want and speed the checkout process. Such systems furthermore offer the virtue of instantaneous updates on inventory control and reorder points, thereby facilitating shelf restocking.

Walmart is adding 3,900 advanced robots to its retail stores. These devices look for shelves that need restocking, trucks that can be unloaded, or floors that need to be cleaned, with the goal of improving business operations. The robots are being spread across the company's 4,700 stores to enhance customer shopping and make sure there are fewer barriers to people's purchases.[25]

Yet there are concerns these robots will take human jobs and invade personal privacy. For example, some companies use smartphones to track consumer movements and shopping preferences while inside a store. They can use this information to tailor ads for specific people or learn what products are of greatest interest to the shoppers. Those practices raise a number of questions regarding the ethics of ad targeting and shopper profiling.

Smart Cities

Metropolitan governments are using AI to improve urban service delivery. According to Brookings Institution writers Kevin Desouza, Rashmi Krishnamurthy, and Gregory Dawson, "the Cincinnati Fire Department is using data analytics to optimize medical emergency responses. The new analytics system recommends to the dispatcher an appropriate response to a medical emergency call—whether a patient can be treated on-site or needs to be taken to the hospital—by taking into account several factors, such as the type of call, location, weather, and similar calls."[26]

Since the Cincinnati Fire Department fields 80,000 requests each year, city officials are deploying this technology to prioritize responses and determine the best ways to handle emergencies. They see AI as a way to deal with large volumes of data and figure out efficient ways of responding to public requests. Rather than address service issues in an ad hoc manner, authorities are trying to be proactive and systematic in how they provide urban services. So far, city officials have been pleased with the deployment and feel it has put them in a better position to manage emergency services.

A number of metropolitan areas are adopting smart city applications that use AI to improve environmental planning, resource management, energy utilization, and crime prevention, among other things. For its smart cities index, the magazine *Fast Company* ranked American locales and found Seattle, Boston, San Francisco, Washington, D.C., and New York City as the top adopters of smart city solutions. Seattle, for example, has embraced sustainability and is using AI to manage energy usage and resource management. Boston has launched "City Hall to Go," a mobile

app that makes sure underserved communities receive needed public services. It also has deployed "cameras and inductive loops to manage traffic and acoustic sensors to identify gunshots."[27]

Through these means, metropolitan areas are leading the country in the deployment of AI solutions. According to a National League of Cities report, 66 percent of American cities are investing in smart city technology. Among the top applications are "smart meters for utilities, intelligent traffic signals, e-governance applications, Wi-Fi kiosks, and radio frequency identification sensors in pavement."[28]

But in some places, these innovations have raised concerns regarding the collection of personal information and how city officials are deploying technology. Even worse are fears that technology contributes to "urban fragility." As noted by Kevin Desouza and David Selby, digital innovations sometimes can disrupt financial safety nets, weaken societal connections, and thereby undermine the ability of communities to deal with problems.[29]

Energy and Climate Change

One of the important ways AI helps is through improved analysis of energy resources and climate change. Estimating consumer and business demands for energy has always been challenging, and that limitation creates major challenges for the sector. Advanced analytics make it possible to model supply and demand in a granular manner and thereby project needs and consequences more clearly.

The same is true in regard to climate change. There are many factors that contribute to climate shifts, and AI enables scientists to analyze these dynamics very effectively. For example, they can examine the root sources of emission levels and project possible impacts over the next several decades. It may be impossible to stop the warming of the planet, but AI puts humans in a better position to manage and to mitigate some of the negative ramifications of climate change.

According to David Victor of the University of California at San Diego, "AI helps make markets more efficient and easier for analysts and market participants to understand highly complex phenomena—from

the behavior of electrical power grids to climate change. But AI itself won't assure that outcome without clear policy incentives."[30] To take advantage of the possible benefits, it is necessary to provide consumers and businesses with relevant information and thereby enable them to make efficient energy decisions. Victor argues that relatively small policy "nudges" can generate significant cost and resource savings. Figuring out how to manage emissions is a key to mitigating their consequences.

Law Enforcement and Criminal Justice

In a number of places, AI is being deployed in law enforcement and criminal justice. For example, the city of Chicago has developed an AI-driven "Strategic Subject List" that analyzes people who have been arrested for their risk of becoming future perpetrators. It ranks 400,000 people on a scale of 0 to 500, using items such as age, criminal activity, victimization, drug arrest records, and gang affiliation. Analysts have found that youth is a strong predictor of violence and being a shooting victim is associated with becoming a future perpetrator.[31]

Judicial experts claim AI and data analytic programs reduce human bias in law enforcement and lead to a fairer sentencing system. For example, R Street Institute associate Caleb Watney writes that "empirically grounded questions of predictive risk analysis play to the strengths of machine learning, automated reasoning and other forms of AI. One machine-learning policy simulation concluded that such programs could be used to cut crime up to 24.8 percent with no change in jailing rates or reduce jail populations by up to 42 percent with no increase in crime rates."[32]

Others hope AI can improve courtroom outcomes through natural language processing of case material, evidence, legal opinions, and judicial verdicts. Judges could use the insights derived from big data analytics to determine flight risks, set reasonable bail, find legal precedents for similar crimes, assist in sentencing decisions, reduce racial bias, and provide broad guidelines for pronouncing verdicts.

However, critics charge that AI algorithms represent "a secret system to punish citizens for crimes they haven't yet committed. The risk scores

have been used numerous times to guide large-scale roundups."[33] Others worry about a "police industrial complex" that deploys tools that target people of color unfairly and do not help cities reduce crime rates.[34] Some analysts, such as Rashida Richardson, Jason Schultz, and Kate Crawford of the AI Now Institute, go even further and argue that unrepresentative crime data yield inaccurate predictions and ultimately lead to civil rights violations.[35]

GOVERNANCE, ETHICS, POLICY, LEGAL, AND GEOPOLITICAL ISSUES

AI's increasing penetration into many aspects of life creates tremendous opportunities for advancement and innovation. Yet, its algorithms also raise important governance, ethical, policy, legal, and geopolitical challenges. For example, who should make decisions about emerging technologies: software designers, corporate executives, policymakers, or consumers? What types of ethical problems are introduced through software programming, and how transparent should developers be about their coding choices? How can we guard against biased or unfair data used in algorithms? How can we protect personal privacy and security, and guard against offensive material? What about questions of legal liability in cases where algorithms cause harm? How does AI affect geopolitical relationships around the world and the ability to handle large-scale changes arising from pandemics, income inequality, and climate change? How do we maintain human control over advanced technologies?[36]

Governance

Emerging technologies raise governance questions about who should decide issues of innovation, deployment, and remediation in case of consumer harm.[37] In the old order, national governments were the relevant decisionmakers on basic questions about public policy and society. Leaders would pass laws and enact regulations designed to address how innovations from the telegraph and telephone to television and nuclear energy