

Why the United States Needs a National Artificial Intelligence Strategy and What It Should Look Like

By Joshua New | December 4, 2018

The United States should develop a national AI strategy to bolster U.S. competitiveness and national security and maximize the social and economic value of AI.

The United States is the global leader in developing and using artificial intelligence (AI), but it may not be for long.

Succeeding in AI requires more than just having leading companies make investments. It requires a healthy ecosystem of AI companies, robust AI inputs—including skills, research, and data—and organizations that are motivated and free to use AI. And that requires the federal government to support the development and adoption of AI. Many other countries, including China, France, and the United Kingdom, are developing significant initiatives to gain global market share in AI. While the U.S. government has taken some steps, it lacks a comprehensive strategy to proactively spur the development and adoption of AI. This report explains why a national AI strategy is necessary to bolster U.S.

competitiveness, strengthen national security, and maximize the societal benefits that the country could derive from AI. It then lays out six overarching goals and 40 specific recommendations for Congress and the administration to support AI development and adoption.

INTRODUCTION

Computer scientists have worked since the 1950s to develop artificial intelligence—computer systems that perform tasks characteristic of human intelligence, such as learning and decision-making. But it is only in the last

decade that they have had all the technological building blocks necessary to achieve their vision. Advances in hardware, including faster processors and more abundant storage, plus larger data sets and more capable algorithms, especially machine learning, have unlocked many more opportunities to use AI throughout the economy, such that AI is now poised to become a major driver of innovation, growth, and social welfare.

Today's economy is a data economy, with organizations using data and analytics to drive productivity and innovation; but this is transitioning into the algorithmic economy, in which many more organizations will invest in AI to automate processes, develop new products and services, improve quality, and increase efficiency. Though the industries that are developing AI applications will surely grow, the major beneficiaries of AI will be all the organizations in the public, private and nonprofit sectors that integrate AI into their operations. And while many factors will influence U.S. economic competitiveness, the extent to which key traded sectors adopt AI will be among the most significant, from agriculture to manufacturing and financial services.

Many factors allowed U.S. firms to lead the digital economy, including early adoption of information technologies such as e-commerce and cloud computing. But the ability of U.S. firms to hold their lead in the AI era is anything but assured. Other nations recognize the potential of the algorithmic economy and are taking actions to ensure their industries are well-positioned to leverage AI. China is developing a particularly sophisticated national AI strategy and is allocating massive funding to capture global market share. Moreover, as with many emerging technologies, there are a host of market failures that limit AI development and adoption. Absent an AI strategy tailored to the United States' political economy, U.S. firms developing AI will lose their advantage in global markets and U.S. organizations will adopt AI at a less-than-optimal pace.

It is therefore critical for Congress and the White House to craft and fund a comprehensive national AI strategy that addresses the significant challenges inhibiting the development and adoption of the technology. This strategy should be focused on six goals:

1. **Support key AI organizational inputs.** To effectively develop or use AI, organizations need abundant access to three key resources: high-value data, AI skills, and publicly funded research and development.
2. **Accelerate public-sector adoption of AI, including for national security.** One of the most straightforward and effective steps government can take spur AI progress is to rapidly adopt AI in support of its own missions. Government can help prove the value of deploying AI, as well as provide markets and increase

economies of scale for AI firms. To do this, a strategy needs to address challenges related to acquisition, funding, and oversight. One key area for public-sector adoption is defense. Defense agencies recognize the benefits AI can offer, but they face unique challenges in putting the technology to use to enhance national security.

3. **Spur AI development and adoption in industry, including through sector-specific AI strategies.** The federal government has significant influence and involvement in sectors such as health care, transportation, and education through funding, procurement, and regulation. Federal agencies should be charged with developing sector-specific AI strategies to shape their policies affecting these industries in ways that support AI transformation. Additionally, many nations look to AI as an important industry for future competitiveness. They are putting in place a host of development policies designed to grow their domestic AI industries. The United States needs to do the same.
4. **Support digital free trade policies.** Data is at the core of AI, and many nations are enacting policies that restrict cross-border data flows. The U.S. government needs to accelerate its efforts to establish free trade in data and fight other protectionist efforts that inhibit AI, such as source code disclosure requirements.
5. **Foster innovation-friendly regulation.** If poorly implemented, AI can produce undesirable outcomes. In response, some have called for strong regulations on AI, including through tougher enforcement of antitrust and regulation of algorithms. Most proposals thus far would harm AI innovation and use, often without providing meaningful protections. Policymakers should instead pursue a more innovation-friendly framework built around the principle of “algorithmic accountability,” in which the operators of algorithms are held accountable for explicit and serious harms. Additionally, antitrust regulators should resist viewing the possession of large amounts of data as a threat to competition.
6. **Provide workers with better tools to manage AI-driven workforce transitions.** AI-enabled automation will increase productivity and per-capita incomes but also will likely modestly increase the rate of worker displacement, which can lead to support for policies that restrict how firms can use AI. To help workers more effectively make transitions, policymakers need to modernize workforce training and worker dislocation policies and programs.

To implement these goals, the Center for Data Innovation offers the following 40 recommendations:

Ensuring Data Availability

1. Relevant federal agencies should support the development of shared pools of high quality, application-specific training and validation data in key areas of public interest, such as agriculture, education, health care, public safety and law enforcement, and transportation. For example, the National Institute of Standards and Technology (NIST) should work with law enforcement agencies, civil society, and other stakeholders to develop shared, representative datasets of faces that can serve as an unbiased resource for organizations developing facial recognition technology.
2. Relevant federal agencies, including the Department of Commerce and the Department of Health and Human Services, should develop and pilot data trusts to facilitate data sharing in specific application areas among academia, businesses, and government agencies. Congress should also fund and task the Department of Commerce to pursue additional innovative models for increasing the availability of data. This should include facilitating the creation of industry-led data councils to identify barriers to data sharing and developing strategies to overcome these barriers.
3. Congress and the administration should accelerate efforts to digitize all sectors of the economy, including health care, education, and municipal issues such as utilities and city management. Federal agencies, such as the Department of Housing and Urban Development (HUD), Department of Health and Human Services (HHS), Department of Transportation (DOT), and Federal Energy Regulatory Commission (FERC), should identify and implement policies that can drive digital transformation in relevant sectors.
4. Congress and the administration should encourage the private sector to share data for public benefit. There are many examples of firms voluntarily making data available for AI research, however this is not the norm. Policymakers should consider a variety of different approaches to encourage this. For example, France's AI strategy proposes requiring the private sector to share certain data sets in select circumstances, when it does not threaten a firms' business and relates to key public interests such as health and safety. The United States government should not appropriate data from the private sector, but instead take a more collaborative approach

to identifying such datasets and working with the private sector to increase their availability.

5. Congress should pass legislation codifying the federal government's responsibility to publish open data. Since 2013, the federal government has had a policy of treating its data as open and machine readable by default, opening vast troves of government data to the public. This data is a crucial public resource for businesses, academics, and public-sector employees alike, however the government has no legal obligation to continue publishing this data. This means that the government may decide to stop making certain data available at any time, such as with a new administration or agency priorities. Congress should pass legislation to codify open data requirements to ensure that government data remains available as a valuable source of data for AI systems. Congress should also allocate additional funding for the federal government's open data efforts to improve the timeliness, quality, and accessibility of its data.
6. Relevant federal agencies should ensure data collection efforts emphasize reducing the "data divide" and combatting data poverty. For the public sector, this means supporting and expanding data collection programs that focus on hard-to-reach and underrepresented communities. Additionally, this would mean ensuring that federal programs devoted to closing the digital divide also consider data poverty concerns. More broadly, an agency such as NIST could develop educational materials about how to improve data collection in the private sector to combat the data divide.
7. Congress should ensure that any national legislation addressing privacy considers the importance of data for the development and use of AI and does not impose undue restrictions on the collection, sharing, and use of data that come at the direct expense of AI innovation, such as an opt-in requirement for data sharing.

Developing AI Talent

8. Congress should invest in cultivating AI talent. Countries such as Canada and the United Kingdom have launched initiatives to do just that, and the United States should adapt these approaches. For example, Canada's AI strategy funds the creation of AI research institutes, programs to attract and retain AI talent in Canadian universities. Similarly, the United Kingdom's AI Sector Deal describes how the government will fund at least 1,000 AI PhD students by 2025. Congress should

fund and direct the National Science Foundation (NSF) to create a competitive AI fellowship program for at least 1,000 computer science students annually.

9. Congress should fund and authorize a program at the National Science Foundation to provide competitive awards for up to 1,000 academic AI researchers for a period of five years. Awards should be conditional on remaining in academia for five years. Though individual businesses may benefit from attracting the best AI talent away from universities, the overall AI innovation ecosystem in the United States suffers. These awards would incentivize more AI researchers to stay in academia and help U.S. universities meet the demand for AI skills.
10. Congress should enable more foreign AI talent to work in the United States by increasing the cap on H-1B visas to ensure U.S. firms can hire as much AI talent as they need. To the extent global AI talent works in the United States, they are not working for competitors in other nations.
11. Federal agencies should address barriers that limit the number of students able to take computer science courses at the university level. NSF should provide grants to colleges and universities that have increased or are implementing programs to increase enrollment and retention in computer science. The federal government should also require increased transparency as a prerequisite for receiving NSF awards. For instance, schools should be required to monitor and disclose the number of computer science applicants, prospective majors, and their retention rates in computer science subjects. Ideally this would be done for all STEM disciplines.

Conducting AI R&D

12. Congress should substantially increase R&D funding for AI, with an emphasis on basic and applied research. Exact dollar amounts can be debated, but at the very least Congress should provide funding to fully support R&D efforts where the potential supply of high-quality research is greater than the supply of funds. For example, Congress should appropriate at least an additional \$200 million annually to NSF for AI research.
13. Federal agencies should support R&D for all kinds of AI applications. It is a mistake for federal agencies to support only certain kinds of AI and related technologies. For example, NSF's National Robotics Initiative only supports research that augments, and does not replace workers. While augmenting human labor is a valuable application of AI, AI that can replace

a human worker does more to boost productivity and should at minimum be on equal footing for support of AI that complements workers.

14. Congress should increase the R&D tax credit to keep pace with competing countries. A healthy AI ecosystem requires both government and business funding of AI research. Companies will do more AI research in the United States if the R&D tax credit is more generous. However, as of 2017, the United States ranked 32nd of OECD nations in terms of R&D tax credit generosity, behind countries such as Canada, China, Germany, and Japan. And in 2018, a number of nations created or expanded their R&D tax incentives. Meanwhile, the 2017 tax legislation passed by Congress in 2017 actually increased the after-tax cost of research spending. As such, Congress should increase the Alternative Simplified Credit from 14 percent to 20 percent.

Transforming Government With AI

15. Congress and the administration should support efforts to foster communities of practice and raise awareness about AI within the public sector. For example, the General Services Administration's (GSA's) Emerging Citizen Technology Office (ECTO) coordinates government-wide deployments of AI applications and helps foster relationships between government employees interested in AI and firms working on public-sector applications of AI.
16. Congress should provide agencies with venture capital funds to pilot AI initiatives.
17. Federal agencies should establish domain-specific programs to spur AI adoption. For example, the Department of Defense (DoD) recently established the Joint Artificial Intelligence Center (JAIC) to help teams "to swiftly deliver new AI-enabled capabilities and effectively experiment with new operating concepts in support of DoD's military missions and business functions." Other departments, such as HHS and DOT, should consider developing similar programs.
18. The White House should establish a strategic initiative devoted to AI in the CIO Council. There is currently an initiative devoted to "data analytics and big data" which would likely cover certain aspects affecting AI adoption, however there should be a more explicit focus on AI.
19. GSA should work with state government CIOs to share best practices for AI implementation and develop shared resources

that make it easier for state officials to learn about and procure AI technologies.

20. Defense agencies should prioritize the use of AI to support their missions to protect national security.
21. DoD should create a body with both government and industry stakeholders to accelerate the adoption of dual-use AI technologies by the military. This could include publishing performance and safety standards for various key military AI applications so industry could more readily develop those solutions, or creating guidelines for modifying commercial AI applications for military use.
22. DoD should establish a cross-agency task force to identify opportunities to simplify the acquisition process for AI.
23. DoD should pursue and expand the use of alternative acquisition mechanisms as a workaround for cumbersome procurement policies. For example, the 2016 National Defense Authorization Act granted DoD the permanent Other Transaction Authority (OTA), which allowed DoD to circumvent the traditional acquisition process in certain circumstances.
24. DoD should foster better relationships between the defense community and the U.S. technology industry, such as by expanding industry outreach efforts like DoD's Defense Innovation Unit Experimental (DIUx) designed to make it faster for DoD to take advantage of emerging commercial technologies. There should be specific emphasis on creating greater incentives for technology firms to work with DoD.
25. DoD should establish a new Program Element (PE) specifically for AI to increase the visibility of AI appropriations.
26. Congress should prioritize the development and adoption of AI in defense spending. This could entail either focusing greater attention on AI projects as compared to less-important work or increasing overall spending.
27. Congress and the administration should support productive conversations about the appropriate way to oversee the use of AI for national security. This will include rejecting bans on lethal autonomous weapons (LAWS) and differentiating between concerns about LAWS specifically and broader concerns about different military activities, which are often the underlying concern in these discussions.
28. Congress and the administration should recognize that the benefits of AI to national security are too important to let concerns about LAWS oversight or other defense activities

involving AI limit national security AI support and adoption. This is particularly important because foreign adversaries will pay little heed to such oversight concerns and gain a competitive advantage in certain areas of AI.

Spurring AI Development and Adoption in Industry

29. Federal agencies should work with industry to create strategies for supporting AI adoption in relevant sectors of the economy. These strategies should provide guidance about how best to leverage AI to advance agency missions, as well as identify opportunities to encourage AI adoption in relevant industries, such as by proactively providing guidance on policy questions, ensuring that procurement supports AI, ensuring regulations do not limit AI usage, and creating incentives for firms to invest in AI. These strategies should be updated regularly as agencies become more familiar with the technology and as AI matures, creating new challenges and opportunities to address.
30. The Department of Commerce should establish organizations designed to advance the development of innovative AI applications in various sectors. For example, Manufacturing USA, overseen by federal agencies including the Department of Commerce and the Department of Energy, is a network of research institutes focused on fostering innovation and collaboration in the manufacturing sector. Among them is the Advanced Robotics for Manufacturing, a public-private partnership in Pittsburgh which focuses on AI and automation. Using this model, agencies should support similar institutes that include industry, academia, and government agency resources to advance AI in other sectors such as city management and precision medicine.
31. Congress should direct the Economic Development Administration to enable state governments to foster AI industry development. Congress should appropriate funds for the Economic Development Administration to create a state economic development competition in which states would compete for funds to establish their own state development plans and policies for supporting AI development, especially through new startups.

Ensuring Trade Policy Supports AI

32. The United States Trade Representative (USTR) should continue to advocate for cross-border data flow protections in all future trade negotiations.

-
33. USTR should continue to fight source code disclosure requirements other nations may enact to unfairly disadvantage U.S. firms or exploit their intellectual property.

Ensuring Any Regulation of AI Is Innovation Friendly

34. Regulators should encourage adherence to the principle of algorithmic accountability. Importantly, policymakers need to recognize that the goal of algorithmic accountability is not to achieve perfect, error-free algorithms, but to minimize risk—just as vehicle safety standards do not require cars to be 100 percent safe, but as reasonably safe as can be expected. The most important step is for regulators to formally recognize this framework for algorithmic accountability and integrate it into their oversight. This applies to both domain-specific and consumer-protection regulators.
35. Congress should reject blanket mandates for algorithms, such as algorithmic transparency requirements, or the creation of new regulatory bodies focused only on regulating algorithms.
36. Congress and the administration should support increasing the technical expertise of regulators and policymakers. Regulators should foster relationships with communities of developers, academics, civil society groups, and private-sector organizations invested in algorithmic decisionmaking to stay abreast of technical developments and concerns about algorithmic harms that could influence how algorithmic accountability is achieved or enforced. This requires ensuring regulators have the resources to hire staff with the necessary technical expertise to scrutinize algorithms.
37. Congress and the administration should caution regulators against viewing the mere act of collecting or possessing large amounts of data (which is necessary for certain uses of AI) as potentially anticompetitive behavior.
38. Congress should reject overly stringent rules in privacy legislation, such as broad opt-in requirement, purpose specification, data erasure, and data minimization, as well as other policies modeled on the EU's General Data Protection regulation.
39. Congress and the administration should emphasize that data is a crucial business input for the development of AI, and that companies should be encouraged to invest in collecting data, not punished for it. If policymakers are concerned that startups and small businesses cannot access the data necessary to develop AI and compete with larger firms, they should focus on making it easier to collect data and ensure data is readily

available, as described earlier in this report, rather than penalize a company that has succeeded in collecting and making data available.

Providing Workers with Better Tools to Manage AI-driven Workforce Transitions

40. Congress and the administration should implement comprehensive reforms to the nation's workforce training and adjustment policies, as the Information Technology and Innovation Foundation (ITIF) outlined in its February 2018 report, "How to Reform Worker-Training and Adjustment Policies for an Era of Technological Change."

THE IMPORTANCE OF AI

As an emerging general-purpose technology, like electricity and information technology, AI has the potential to drive innovation, competitiveness and productivity. According to the consulting firm PwC, AI will increase global GDP by up to 14 percent by 2030 due to its ability to drive productivity gains by automating business processes and augmenting human labor.¹ For the United States specifically, consulting firm Accenture estimates that AI could increase labor productivity by 35 percent and increase the annual growth rate of gross value added to the U.S. economy from 2.6 percent to 4.6 percent by 2035 through automation and improving labor and capital management.²

AI has the potential to generate a range of societal benefits, such as by helping develop new medical treatments, improving public safety, and fighting human trafficking.³ For example, Facebook has developed computer-vision algorithms that can describe images to blind users, making the Internet more accessible for people with visual impairments.⁴ San Francisco-based analytics company Kanjoya has developed machine learning software that can analyze workforce communications and flag signs of implicit gender bias so companies can treat employees more fairly.⁵ And conservation technology start-up Conservation Metrics uses a system of acoustic sensors and machine learning to improve conservation efforts for threatened species in California.⁶

WHY THE U.S. NEEDS A NATIONAL AI STRATEGY

There are three major reasons why the United States needs a national AI strategy: 1) to boost U.S. economic competitiveness; 2) to support U.S. defense capabilities; and 3) to overcome market failures, including the provisioning of public goods, that would otherwise slow AI development and adoption.

Competitiveness

The United States has unique advantages over other countries that have made it an early leader in AI, particularly its large technology sector and its innovation-friendly regulatory environment. For example, in 2016, 66 percent of global investment in AI went to the United States, with Silicon Valley and the San Francisco Bay Area in particular attracting 40 percent of global investment in AI.⁷ By contrast, China, which attracted the second largest share of global investment in AI, received just 17 percent.⁸ And while the United States has traditionally embraced a regulatory philosophy based on the innovation principle—the idea that the majority of innovations overwhelmingly benefit society, and the government’s role should be to pave the way for widespread innovation while building guardrails, where necessary, to ensure public safety—others have taken a much more precautionary approach, at the expense of innovation.⁹ For example, the European Union’s recently enacted General Data Protection Regulation (GDPR) contains a variety of provisions that harm the ability of European firms to take advantage of AI while doing little to protect European consumers.¹⁰

However, there are signs that the United States’ early lead is slipping. In 2017, China secured 48 percent of global investment in AI startups, while the United States received just 38 percent.¹¹ According to a study by *The Economist* about national readiness for automation, based on national innovation environments, labor market policies, and education policies, the United States ranked 9th out of 25 countries studied, while South Korea, Germany, Singapore, Japan, and Canada ranked 1st through 5th place, respectively.¹² And since 2015, China began publishing more patents related to AI and deep learning than the United States and has grown that lead considerably, publishing six times as many AI and deep learning patents than the United States in 2017.¹³ And while other governments are aggressively increasing their research funding for AI, U.S. government research has remained relatively flat.¹⁴ NSF and the National Science Board noted in February 2018 that if current trends continue, China will surpass the United States in all R&D investments by the end of 2018.¹⁵

National Security

The ability of the United States to harness AI for defense purposes will also have crucial implications for national security. The opportunities to use AI for national security are vast, including improving logistics, analyzing surveillance footage and satellite imagery, and improving training exercises.¹⁶ AI can support the use of autonomous and semi-autonomous weapons systems and provide analysis and threat assessments to soldiers on the battlefield. However, the bulk of the benefits AI can offer national security are similar to the benefits it can offer the private sector, such as automating routine processes, improving data analysis, identifying and fighting cyber threats, and processing large amounts of sensor data.

Beyond improving existing defense operations, ensuring the United States can use AI effectively for national security will be a necessity to maintain technological superiority over adversaries that are investing heavily in AI for military use. For example, China is developing a range of autonomous aerial, ground, surface, and underwater military vehicles, and Russia is using AI to augment its information warfare activities.¹⁷

To be sure, if the Defense Department needs cutting-edge AI capabilities it can attempt to procure military-specific AI technologies. But while DoD drove IT innovation in the 1950s and even the 60's, today the center of IT innovation, and AI innovation, is in the private sector. And DoD will be increasingly reliant on partnerships with and purchases from the commercial IT industry for cutting-edge AI capabilities. However, it would be extremely risky to national security to rely on those capabilities from military adversaries.

Market Failures

There are a number of market failures that will slow AI development and adoption in the absence of supportive policies. First is the fact that effective AI development requires access AI skills and a broad base of technical knowledge. These inputs—skilled workers and public R&D—are public goods. Though the private sector invests in worker training and R&D of its own, it does not capture all of the benefits of this investment. For example, a company could provide years of technical training for a new hire, only to have him or her leave to work for a competitor. Thus, firms often underinvest in these inputs relative to societally optimal levels. Smart policies to increase the availability of these public goods can help correct for this.

A second failure relates to risk and uncertainty. Because AI is an emerging technology, many potential users, including companies and government agencies, will minimize the benefits it promises and delay adoption until the technology is proven. Economists refer to this challenge as excess inertia or, more commonly, “the penguin effect”—in a group of hungry penguins, no individual penguin is willing to be the first to enter the water to search for food due to the risk of encountering a predator. Yet if no penguin is willing to test the waters, then the whole group risks starvation.¹⁸

A third failure relates to externalities. The widespread use of AI can generate significant social and economic benefits. However individual firms investing in AI are unable to capture some of these benefits, leading to underinvestment. For example, the ubiquitous use autonomous vehicles could drastically reduce traffic deaths, property damage, and congestion (and the resulting loss of economic activity). Yet no individual company

developing autonomous vehicles could expect to capture enough value from these benefits to recoup the costs of developing the technology.

Finally, there are collective action problems. One of the key drivers of AI is widespread data availability. However, organizations often have strong incentives not to share data. Even though all parties would be better off if each shared data for mutual benefit, without the proper incentives for participation and organizations willing and able to coordinate these efforts, there will be less data available overall for AI development.

Despite these compelling rationales, U.S. policymakers have not yet committed to developing such a strategy. For example, a spokesman for the Senate Commerce Committee stated that because the private sector is already using and investing heavily in AI, “the horse has already left the barn... and I think that any attempts by government to try and intervene could be constraining on the development of this technology.”¹⁹ To be sure, government action to regulate could be constraining. But that doesn’t mean that policies to proactively support AI development and adoption would not be helpful or needed.

Policy Approaches to AI

There are three main approaches a government can take regarding AI development and adoption: abdication, regulation, and facilitation. To be sure, most nations are incorporating all three approaches, but for most, one approach dominates.

Abdication: This approach is based on the belief that the private sector can, with little support, largely effectively drive AI development and adoption. Holders of this view generally see AI market failures as minimal and worry more about government failures, particularly the risk of using regulation, either motivated by anti-technology forces or incumbents seeking protection from competition, to limit or shape progress in AI.

To date, the United States has embraced this approach regarding AI, preferring to leave AI development and use largely up to the market. While this approach avoids the pitfalls that come with an approach centered on regulation, as restrictive regulations can make AI development and adoption more difficult, relying on this approach will lead to suboptimal AI use and weakened competitiveness. This is because, as described above, there are serious market failures that impede the development and adoption of AI that industry alone cannot overcome in a timely manner.

Regulation: The current generation of AI is new and for some, a risky technology. As a result, some advocate sweeping regulations to shape and constrain AI. Like those who support an abdication approach, the advocates of regulation pay little heed to market failures regarding the development and use AI. Instead, these advocates assume companies will

pursue their own self-interest and sufficiently advance AI development and adoption. Therefore, they believe that the main focus of AI policy should be to constrain and shape the self-developing technology. A recent Brookings Institution report on AI embodies this approach, stating, “If the private and public sectors can work together, each making its own contribution to an ethically aware system of regulation for AI, we have an opportunity to avoid past mistakes and build a better future.”²⁰ In other words, these advocates are not concerned with whether the United States will be the global leader in AI, only with whether AI will develop in what they believe to be the right way.

While no nation with an AI policy focuses exclusively on shaping AI, some, particularly in Europe, put relatively more emphasis on this. Seeing AI as a threat, rather than an opportunity, many policymakers in some nations place their focus on minimizing the technology’s potential harms. For example, the United Kingdom’s AI strategy focuses on developing “ethical AI” and France’s AI strategy seeks to “guarantee the principle of human responsibility” in AI.²¹ Though modernizing regulations may be beneficial in some cases, if not done smartly it will limit AI development and adoption. In addition, as with an abdication approach, a regulation approach does not address market failures that impede the technology.

Facilitation: Most nations that have developed AI policies have recognized that active government support can accelerate AI development and adoption. As such, most nations, particularly Southeast Asian nations, are focused primarily on facilitating AI development and adoption. To be sure, an overzealous approach to facilitation, such as China’s approach in particular, can introduce government failures such as supporting the wrong firms or limiting foreign competition. Done correctly however, a facilitation approach combines the best of enterprise-driven market forces with societal support and facilitation.

An ideal national AI strategy will incorporate elements of all three approaches, recognizing that the private sector and market forces will play decisive roles in the advancement of AI, modernizing regulations for the AI economy in ways that enable AI innovation, and emphasizing the right government role in facilitating widespread development and adoption of AI.

ASSESSING THE COMPETITION

Unlike the United States, a number of countries have translated recognition of the importance of AI into significant action. It is important to recognize that this response is unique. During the last major IT revolution of the Internet, virtually no U.S. competitor took it seriously, at least initially, and few developed robust policies to compete globally. Things are different this time. Many nations now understand that AI is a central general-

purpose technology of the next era, and they are loath to make the same mistakes they made 30 years ago.

To be sure not every country intends to compete with the United States in every aspect of AI development and adoption. Instead, most countries' strategies focus on capitalizing on their respective comparative advantages by prioritizing their efforts to secure the benefits of AI where it can have the most impact. However, this should not be interpreted as a chance for the United States to rest on its laurels; rather, because as the largest economy globally it needs to be competitive in every aspect of AI.

Several leading countries, including China, France, and the United Kingdom, have developed comprehensive national AI strategies and are in the process of implementing them. Additionally, others have made their intentions to support AI clear, though have yet to formally introduce a strategy. For example, German Chancellor Angela Merkel has announced a proposal to make €3 billion (US \$3.34 billion) available to Germany's private sector to support AI R&D through 2025, and Germany is in the process of drafting a broader AI strategy.²² And the European Commission published its Communication on Artificial intelligence in Spring 2018 explaining its plans to boost technical capacity and spur public and private-sector AI adoption, and has outlined its planned activities to support AI, such as by increasing investments and improving research centers, through 2020.²³ Many other countries have also signaled that they will be developing policies to support AI, though their plans are less clear.²⁴

Canada

In March 2017, Canada launched the Pan-Canadian Artificial Intelligence Strategy, to be led by the Canadian Institute for Advanced Research (CIFAR), a nonprofit research institute that receives government support.²⁵ Backed by a one-time CAD \$125 million (US \$98.7 million) government investment, the strategy has four goals: "increase the number of outstanding artificial intelligence researchers and skilled graduates in Canada; establish interconnected nodes of scientific excellence in Canada's three major centres for artificial intelligence in Edmonton, Montreal, and Toronto; develop global thought leadership on the economic, ethical, policy and legal implications of advances in artificial intelligence; and support a national research community on artificial intelligence."²⁶ CIFAR will oversee several programs over the next five years to advance the strategy that focuses on expanding Canada's human capital, raising Canada's international profile in the field of AI research, and translating AI research into public and private-sector applications. Canada has also allocated CAD \$950 million (US \$718 million) to fund the creation of five technology "superclusters" designed to foster collaboration and accelerate growth and job creation around different technology issues.²⁷ Several of these, including the Scale AI supercluster in Quebec, the Advanced

Manufacturing supercluster in Ontario, and the Digital Technology supercluster in British Columbia, have an explicit focus on AI.

China

China's State Council issued a development plan for AI in July 2017 with the goal of making China a leader in the field by 2030.²⁸ The document is primarily a statement of intent, but it details some of China's key objectives in advancing AI and creating a domestic AI industry worth ¥1 trillion (US \$147.8 billion).²⁹ The plan's goal is for China to be equal to countries leading in AI by 2020. Then over the subsequent five years, China will focus on developing breakthroughs in areas of AI that will be a "a key impetus for economic transformation." Finally, by 2030, China intends to be the world's "premier artificial intelligence innovation center." To support the development plan, China is also preparing a multibillion-dollar investment initiative to promote AI startups, academic research, and ambitious moonshot projects.³⁰

China, like the United States, has its own natural advantages when it comes to AI. First, the Chinese private sector often moves in lock-step with the government to secure favorable treatment. At the end of 2016, 67.9 percent of the 2.73 million non-state-owned companies in China had Communist Party cells within their organizations, and authorities often appoint business leaders with memberships in the National People's Congress, China's national legislature. Kenji Kawase, chief business news correspondent at the *Nikkei Asian Review*, describes that "These positions are largely ceremonial, but membership symbolizes a political recognition by the party and the state, and also provides a certain amount of protection. For the authorities, selecting these business leaders is a way to ensure that private companies cooperate with policy goals, including that of gaining a tighter grip over cyberspace."³¹ Thus, the Chinese government has a de facto authority to direct private-sector investment in AI as it sees fit to advance its goals rather than allow businesses to act pragmatically in their own best interests. This close alignment is also reflected in the actions of Chinese military, financial institutions, and subnational governments, ensuring that in China, "commercial companies, university research laboratories, the military, and the central government routinely work together closely. As a result, the Chinese government has a direct means of guiding AI development priorities and principles."³²

Second, whether as the result of different values or China's more authoritarian government, Chinese organizations using AI simply do not have to grapple with the same regulatory and consumer protection considerations that U.S. firms do. For example, China's plans to deploy ubiquitous facial recognition technology, surveil citizens with drones, monitor messaging apps, and implement a "social credit system" would all likely be met with widespread consumer backlash in the United States and

other western democracies, but are being implemented in China without notable domestic resistance (though it should be noted that vocal dissent is likely censored or otherwise discouraged).³³ The Chinese government's control over private technology firms and its ability to deploy largescale data collection systems with little regard to consumer protection or privacy give it an advantage over the United States because it gives it access to vast troves of consumer and personal data that U.S. firms have difficulty accessing.³⁴ As this paper will describe, data is a crucial input for AI development, and this gives China a large competitive edge in that regard.

Third, China has a long history of stealing intellectual property from foreign firms, particularly U.S. firms, and despite its claims that it respects intellectual property, there is every reason to believe it will continue this practice as international competition regarding AI escalates. Both the Office of the United States Trade Representative and the European Commission published reports in 2018 detailing numerous problems with how China protects the intellectual property of foreign firms.³⁵ These include forced technology transfers, requiring the disclosure of business information, and theft of trade secrets.³⁶ Thus, as China seeks to promote its domestic AI industry, it is safe to assume that it will simply steal intellectual property that U.S. firms invested large sums to develop.

Thus, when comparing national efforts to support AI, it is crucial for U.S. policymakers to recognize that when it comes to competing with China, this is not an even playing field. This makes a U.S. national AI strategy all the more important to ensure U.S. firms can compete effectively.

France

France published a report detailing its national AI strategy titled “For a Meaningful Artificial Intelligence,” in March 2018. This strategy includes investing €1.5 billion (\$1.85 billion) over five years to support R&D, promote AI startups, and foster the creation of valuable datasets that can support AI development.³⁷ Overall, the strategy fails to describe how the French government will promote widespread AI adoption, however it is noteworthy for its emphasis on making data available for AI.³⁸ The French strategy, drafted by mathematician and Member of Parliament Cédric Villani, calls for legislation to mandate repurposing both public and private-sector data, including personal data, to enable public-interest uses of AI by government or others, depending on the sensitivity of the data. For example, public health services could use data generated by Internet of Things (IoT) devices to help doctors better treat and diagnose patients. Researchers could use data captured by motorway CCTV to train driverless cars. Energy distributors could manage peaks and troughs in demand using data from smart meters. Repurposed data held by private companies could be made publicly available, shared with other companies, or processed securely by the public sector, depending on the extent to which

sharing the data presents privacy risks or undermines competition. The report suggests that the government would not require companies to share data publicly when doing so would impact legitimate business interests, nor would it require that any personal data be made public. Instead, Villani argues, if wider data sharing would do unreasonable damage to a company's commercial interests, it may be appropriate to give only public authorities access to the data. But where the stakes are lower, companies could be required to share the data more widely, to maximize reuse. Villani rightly argues that it is virtually impossible to come up with generalizable rules for how data should be shared that would work across all sectors. Instead, he argues for a sector-specific approach to determining how and when data should be shared.

After making the case for state-mandated repurposing of data, the report goes on to highlight four key sectors as priorities: health, transport, the environment, and defense. Since these all have clear implications for the public interest, France can create national laws authorizing extensive repurposing of personal data without violating the General Data Protection Regulation (GDPR), which permits the repurposing of personal data where it serves the public interest. The French strategy is the first clear effort by an EU member state to proactively use this clause in aid of national efforts to bolster AI.³⁹ However, the strategy primarily focuses on repurposing data in the public interest because the GDPR limits repurposing data for commercial purposes. But many important uses of AI—from banking to agriculture—are commercial. France is making the best of the situation created by the GDPR, but its limited strategy highlights the need for EU-level reform of data protection law in order to enable more ambitious AI strategies in the member states. Ultimately, unless the EU reforms the GDPR to enable greater collection, use, and sharing of personal data, any European country's AI strategy will be constrained by the GDPR's limitations.⁴⁰

India

In June 2008, India published a discussion draft of its national AI strategy, which is focused on overcoming barriers limiting the development and adoption of AI at scale, such as a lack of technical expertise and restrictions on data access.⁴¹ The strategy includes policy recommendations designed to address these barriers and ensure India can capture AI's economic and social value, including establishing research centers that focus on advancing AI applications in key sectors such as health care, education, and agriculture and creating annotated "foundational" datasets that could serve as a public resource to spur AI development.⁴² India's strategy also sets the goal of becoming the AI "garage" for the 40 percent of the world whose economies are developing. This means it would create an environment such that if a firm can successfully deploy an AI application in

India, it can be confident that it could deploy it in the rest of the developing world.⁴³

Japan

Japan launched its Artificial Intelligence Technology Strategy Council in April 2016 to develop a roadmap for the development and commercialization of AI.⁴⁴ It was published in May 2017.⁴⁵ The strategy outlines priority areas for research and development (R&D), focusing on the themes of productivity, health, medical care, and mobility. The strategy also encourages collaboration among industry, government, and academia to advance AI research, as well as stresses the need for Japan to develop the necessary human capital to work with AI. Japan also launched its Japan Revitalization Strategy 2017, which details how the government will work to support growth in certain areas of the economy. The 2017 strategy includes a push to promote the development of AI telemedicine as well as the development of self-driving vehicles to help address the shortage of workers in Japan's logistics sector.⁴⁶

South Korea

South Korea has made several large commitments to support the development and adoption of AI, starting with a March 2016 announcement of a ₩1 trillion (US \$886.7 million) investment in AI R&D over five years, a 55 percent increase in its annual spending on AI.⁴⁷ In May 2018, the government announced a ₩2.2 trillion (US \$1.95 billion) investment in AI over five years with the goal of establishing six AI graduate schools to train 5,000 AI specialists, advance the development of AI applications in areas including national defense and public safety, and foster AI startups and small businesses.⁴⁸

Taiwan

Taiwan's Premier William Lai announced the Taiwan AI Action Plan, which details strategies to grow Taiwan's AI industry, in January 2018.⁴⁹ The four-year plan states the government will allocate between NT\$9 billion and NT\$10 billion (US \$304.4 million and US \$338.3 million, respectively) annually to cultivate AI talent, including by recruiting international talent and making it easier for foreign workers to work in Taiwan, developing AI pilot projects, fostering AI startups, and increasing the availability of data, such as by creating open data platforms and developing flexible regulations about data use.⁵⁰

United Kingdom

The United Kingdom has taken several steps to better understand AI and identify ways the government could help secure its benefits. In October 2016, the House of Commons Science and Technology Committee published a report on robotics and AI detailing many of the potential benefits and challenges AI could offer.⁵¹ One of the report's main

conclusions was that the United Kingdom should place a greater focus on improving its education and worker training systems to ensure that the national workforce has the necessary skills to be successful as AI transforms the economy. The report also stressed the need of increased government leadership around robotics and autonomous systems, citing a lack of a government strategy to coordinate policymaking and guide investment. In November 2016, the Government Office for Science published a report detailing the potential implications AI poses for society and government and stressed the need for smart, flexible governance to promote the responsible development of AI.⁵² The UK Digital Strategy, published in March 2017, recognizes AI as a key field that can help grow the United Kingdom's digital economy, and includes £17.3 million (US \$22.3 million) in funding for UK universities to develop AI technologies.⁵³ And in April 2018, the government launched its AI Sector Deal, an extensive industrial strategy that combines and builds on its prior efforts to support AI designed to “boost the U.K.’s global position as a leader in developing AI technologies.”⁵⁴

The AI Sector Deal also lays out an interesting approach to increasing the amount of data available for AI called “data trusts,” which are “mechanisms where parties have defined rights and responsibilities with respect to share data.”⁵⁵ This is a promising approach as it could encourage businesses, government agencies, and researchers to share sensitive or valuable proprietary data with one another to advance AI research by reducing concerns that this data could be misused.

United States

While the United States has taken some steps to support the technology, it does not have a cohesive national strategy for AI. In 2016, the White House Office of Science and Technology (OSTP) hosted five workshops with academic leaders on different social, ethical, economic, and technological aspects of AI, and in June 2016, solicited public feedback about AI.⁵⁶ In October 2016, OSTP published a report titled “Preparing for the Future of Artificial Intelligence” detailing its findings and recommending that the government pursue policies that can help maximize the economic and social benefits of AI.⁵⁷ That same month, the Networking and Information Technology Research and Development Subcommittee (NITRD) published its National Artificial Intelligence Research and Development Strategic Plan detailing seven strategies to help guide AI R&D efforts, including “develop effective methods for human-AI collaboration,” “develop shared public datasets and environments for AI training and testing,” and “better understand the national AI R&D workforce needs.”⁵⁸ Finally, in December 2016, the White House published a report titled “Artificial Intelligence, Automation, and the Economy” reaffirming many of the recommendations from its prior efforts, particularly that the government should ensure the workforce is equipped with the skills to thrive in the transition to an AI-

driven economy.⁵⁹ However, most of the efforts from the Obama administration were foundational to support more explicit policies later.

The Trump administration has taken some additional steps to support AI. In May 2018, the White House hosted a summit titled “Artificial Intelligence for American Industry,” convening leading technology companies to discuss methods for fostering the advancement of AI. Though the summit did not result in the formulation of any new policies or lead to any additional public action, it did establish a Select Committee on Artificial Intelligence under the National Science and Technology Council to advise the White House on AI issues, improve coordination of federal AI R&D efforts, and identify opportunities to leverage federal data and computing resources to support AI R&D.⁶⁰ And in September 2018, NITRD issued a request for information to update its AI R&D Strategic Plan.⁶¹ In the absence of a national strategy, NITRD’s AI R&D Strategic Plan amounts to the most substantive and comprehensive effort to maximize the benefits of AI for the United States, however this document does not direct policy, funding, or regulation.

Congress has taken several actions to better understand the challenges and opportunities posed by AI and has introduced several pieces of legislation to support the technology. For example, the House Oversight Subcommittee on Information Technology held three hearings about AI, exploring how to improve the government’s use of the technology and how best to address policy questions about ethics and competitiveness in AI.⁶² Bipartisan members of the House and Senate have also introduced the FUTURE of Artificial Intelligence Act of 2017 (S. 2217) and the AI in Government Act of 2018 (S. 3502), which would direct the Department of Commerce to create an advisory committee to guide activities to support AI development, and spur the federal government’s adoption of AI, respectively.⁶³

CRAFTING A NATIONAL AI STRATEGY

The ultimate goal of a national AI strategy should be to make the United States a global leader in the development and use of AI. The United States does not need to be the only global leader in AI to benefit from the technology—indeed international collaboration on AI can be highly beneficial. For example, international collaboration can greatly benefit academic research, data sharing, and the use of AI to solve pressing global challenges. Furthermore, so long as other countries are competing fairly, competition in AI would stimulate the development and adoption of AI both in the United States and globally by putting pressure on firms to innovate. There are several mechanisms policymakers could use to create a national AI strategy, such as Congress passing legislation directing the administration to develop a strategy, the administration developing a strategy on its own, or an agency such as the Department of Commerce

leading the development of a strategy. The mechanism used matters little, so long as a strategy incorporates the provisions presented below.

Support Key AI Organizational Inputs

For a firm to leverage AI successfully, it needs to have access to key inputs necessary for the effective development and use of AI: data, AI skills, and public research and development (R&D). Policy can and should do more in all three areas.

Data

Data is a key enabler of AI, and public policy can both remove unwarranted barriers to, and proactively support, data collection, sharing, and use. However, organizations are restricted in their ability to access the data they need to develop and use AI effectively for a host of reasons. In many cases, large swaths of the U.S. economy are still not fully digitized, making the collection, sharing and analysis of data virtually impossible. In other cases, some individuals and communities are not included in data collection efforts, limiting the benefits the data can provide. Regulatory obstacles can unnecessarily limit data collection, access, and use. And in some cases, perverse incentives, or a lack of incentives, limit organizations' willingness to share data. All of these factors make obtaining value from AI "theoretically possible but pragmatically difficult."⁶⁴

The barriers in the health care sector typify the challenges firms face in developing and using AI effectively. The potential for AI to deliver benefits in this sector, such as by discovering new drugs, reducing costs, and improving patient care, is significant.⁶⁵ Yet organizations face a wide variety of factors limiting their ability to access the data necessary to take advantage of AI effectively. The Health Insurance Portability and Accountability Act (HIPAA), in particular, presents numerous unnecessary regulatory obstacles to accessing and sharing useful data. For example, without affirmative consent, HIPAA restricts the use of personally identifiable health information for 50 years after a patient's death, causing large troves of valuable data to be inaccessible to researchers despite negligible privacy risks.⁶⁶ In fact, many health care providers find the complex privacy provisions of HIPAA so daunting that they heavily limit how they share data beyond what the law requires in order to reduce their risk of non-compliance. The U.S. Government Accountability Office found that providers were uncertain about what their specific legal responsibilities regarding patient data actually were, "and often responded with an overly guarded approach to disclosing information."⁶⁷ This problem is exacerbated by the fact that, in addition to HIPAA, states can have their own varying regulations about health data; this creates additional confusion and complexity about how to manage health data.⁶⁸

Considerable nonregulatory barriers exist in the health care sector that also limit the availability of data. For example, despite the federal government investing to boost the use of information technology in the healthcare sector, the United States woefully lags behind global leaders in the extent of modernizing health care with IT, particularly through the use of electronic health records.⁶⁹ When data is not in electronic, interoperable formats the ability to use that data for AI is severely constrained. Many sectors similarly lag in digitization, and organizations in these sectors are limited in their ability to use AI as a consequence. Over half of all U.S. electricity customers do not yet have smart meters monitoring electricity usage.⁷⁰ No U.S. city yet can truly be considered “smart” because even leading cities have only a handful of efforts to deploy sensor networks and digitize municipal operations. And despite the promise of smart manufacturing, adoption is sluggish and it remains more a hopeful slogan than reality.⁷¹

Organizations can have incentives to not share data. For example, many healthcare providers and intermediaries actively limit the exchange and sharing of patient and other related research data.⁷² There is a strong market incentive for providers and healthcare technology companies to restrict the ability of patients to take their data to competing healthcare systems and for customers to take their data to other vendors, and though this practice, called “data blocking,” is illegal, it can be difficult, if not impossible, to prevent and regulators are still determining how to enforce the law effectively.⁷³ Similarly, health care researchers have strong incentives to keep the data they produce proprietary, such as the desire to maximize publications or to ensure that only they can extract value from their data, even if sharing the data broadly could advance research or improve care.⁷⁴

In some cases, there may be strong incentives to share data, but effective models to coordinate involved parties do not exist. In many cases, such as with health data, academia, government agencies, and the private sector could all benefit substantially from sharing data to develop new drugs, reduce costs, and gain insights, but stakeholders lack the mechanisms to do so while ensuring that this proprietary and sensitive data is protected. The United Kingdom is working to address this challenge by developing a model for data trusts, which it defines as “not a legal entity or institution, but rather a set of relationships underpinned by a repeatable framework, compliant with parties’ obligations to share data in a fair, safe, and equitable way.”⁷⁵ Without a coordinating body like a government agency specifically devoted to developing and supporting these models, it is unlikely that organizations will develop them on their own.

Finally, another key challenge in ensuring the availability of data to support the development and adoption of AI is that in many cases, datasets and

data collection efforts are flawed as a result of unequal digital inclusion. As every sector of society and the economy become more reliant on AI, being excluded or underrepresented in the data used to develop these AI systems means that certain individuals or demographics will be unable to fully enjoy the benefits that AI can offer. This concept is known as the “data divide”—the social and economic inequalities that may result from a lack of collection or use of data about an individual or community.⁷⁶ This occurs because data collection can be expensive, and organizations may be inclined to make do with the data they have, rather than invest in more representative data, or simply because developers are unaware that their data is flawed. This is problematic because it can exacerbate or create new social and economic inequalities by limiting the utility of AI applications—if a community is not represented in a dataset used to train an AI system, that system will not consider that community’s wants or needs, and the benefits of products and services that rely on that system will flow only to some, rather than to all. Additionally, because unrepresentative or inaccurate data produces less useful AI, demand for AI-driven products and services will be lessened and U.S. firms will reduce its investments in AI accordingly. Thus, failing to address the data divide is undesirable for both social and economic reasons.

These barriers exist in many sectors of the economy and can cause a less-than-optimal amount of data to be available to U.S. companies developing AI. France’s national AI strategy describes this challenge perfectly, noting that “The person who collects the data is frequently not the only one to benefit from it, or the best placed to capitalize on it; hence the need to promote its circulation so as to maximize its economic and social utility.”⁷⁷

As part of a national strategy:

- Relevant federal agencies should support the development of shared pools of high quality, application-specific training and validation data in key areas of public interest, such as agriculture, education, health care, public safety and law enforcement, and transportation. For example, the NIST should work with law enforcement agencies, civil society, and other stakeholders to develop shared, representative datasets of faces that can serve as an unbiased resource for organizations developing facial recognition technology.
- Relevant federal agencies, including the Department of Commerce and HHS, should develop and pilot data trusts to facilitate data sharing in specific application areas among academia, businesses, and government agencies. Congress should also fund and task the Department of Commerce to pursue additional innovative models for increasing the availability of data. This should include facilitating the creation of industry-led data councils to identify

barriers to data sharing and developing strategies to overcome these barriers.

- Congress and the administration should accelerate efforts to digitize all sectors of the economy, including health care, education, and municipal issues such as utilities and city management. Federal agencies, such as HUD, HHS, DOT, and FERC, should identify and implement policies that can drive digital transformation in relevant sectors.
- Congress and the administration should encourage the private sector to share data for public benefit. There are many examples of firms voluntarily making data available for AI research, however this is not the norm. Policymakers should consider a variety of different approaches to encourage this. For example, France's AI strategy proposes requiring the private sector to share certain data sets in select circumstances, when it does not threaten a firms' business and relates to key public interests such as health and safety. The United States government should not appropriate data from the private sector, but instead take a more collaborative approach to identifying such datasets and working with the private sector to increase their availability.
- Congress should pass legislation codifying the federal government's responsibility to publish open data. Since 2013, the federal government has had a policy of treating its data as open and machine readable by default, opening vast troves of government data to the public. This data is a crucial public resource for businesses, academics, and public-sector employees alike, however the government has no legal obligation to continue publishing this data. This means that the government may decide to stop making certain data available at any time, such as with a new administration or agency priorities. Congress should pass legislation to codify open data requirements to ensure that government data remains available as a valuable source of data for AI systems. Congress should also allocate additional funding for the federal government's open data efforts to improve the timeliness, quality, and accessibility of its data.
- Relevant federal agencies should ensure data collection efforts emphasize reducing the "data divide" and combatting data poverty. For the public sector, this means supporting and expanding data collection programs that focus on hard-to-reach and underrepresented communities. Additionally, this would mean ensuring that federal programs devoted to closing the digital divide also consider data poverty concerns. More broadly, an agency such as NIST could develop educational materials about how to improve data collection in the private sector to combat the data divide.

-
- Congress should ensure that any national legislation addressing privacy considers the importance of data for the development and use of AI and does not impose undue restrictions on the collection, sharing, and use of data that come at the direct expense of AI innovation, such as an opt-in requirement for data sharing.

AI Skills

If the AI economy is to expand it will require considerably more workers with AI skills. However, the United States is already struggling to meet the demand for workers with such skills and this limitation, more so than any technical limitation, holds back AI progress.⁷⁸

The demand for AI talent is global, compounding the competition for skilled workers. A 2017 report from Chinese technology company Tencent speculates that there are approximately 300,000 “AI researchers and practitioners” employed worldwide, but that the demand for such workers is in the millions.⁷⁹ More concretely, according to Montreal-based research lab Element AI, there are only approximately 22,000 people who have received PhDs related to AI since 2015 working in the field, with up to 90,000 people in the world who have the skills necessary to work in the field but who have not received their degrees as recently.⁸⁰

The United States does have a sizeable talent pool to draw from, as nearly half of recent AI Ph.D. holders live and work in the United States.⁸¹ AI talent is not just limited to Ph.D. holders however, and fortunately the United States is home to many strong universities that attract large numbers of students interested in pursuing undergraduate and graduate degrees related to AI around the world. For example in 2015, the United States had slightly over 58,000 graduate students in computer science, 45,800 of whom were international students.⁸² However, only a small percentage of these students will be able to stay and work at U.S. firms as the fixed cap of H-1B visas, which allow U.S. employers to hire foreign workers with specialized skills, is too low to accommodate them and has been exhausted for the past 16 years due to high demand.⁸³ In 2015, 348,669 people applied for an H-1B visa, and only 275,317 were approved.⁸⁴ Of course not all of these applicants were recent graduates with AI skills, but this stands in stark contrast to China, which is actively courting young AI talent internationally.⁸⁵ In 2002, just over 1 in 10 Chinese students returned home to work after studying abroad, and in 2017, that number jumped to just under 8 in 10.⁸⁶

Regardless of the supply shortage, the demand for AI talent is growing at a dramatic clip.⁸⁷ According to LinkedIn, job postings seeking “machine learning engineers” increased by 980 percent from 2012 to 2017, a far larger increase than any other (emerging) job.⁸⁸ Element AI also found that while about 3,000 of the 22,000 recent AI PhD holders are looking for

work, there are at least 10,000 open positions that require their skills in the United States.⁸⁹ According to a 2017 MIT Sloan Management Review and Boston Consulting Group survey, difficulties in acquiring workers skilled in AI was one of the largest reported barriers to AI adoption for surveyed organizations.⁹⁰ And a May 2018 survey of 122 business leaders conducted by Ernst and Young found that 80 percent of respondents identified a lack of talent as the biggest barrier to AI adoption.⁹¹

The demand for people with AI-specific skillsets alone does not completely describe the problem the United States faces. People developing machine learning algorithms, for example, do not work in a vacuum and rely heavily on a broad ecosystem of workers with more general data science skills, such as data governance and security experts, data engineers, and data-literate managers. A 2017 report from Burning Glass Technologies, which analyzes labor market trends, forecasts that the number of positions for workers and managers with data science and analytics skills in the United States will increase by 364,000 openings by 2020, to 2,720,000.⁹² Though not all of such positions would involve work with AI directly, all of these skills would be necessary to support a firm's ability to use AI effectively, and there are already signs that the United States cannot produce enough workers with the skills necessary to meet this demand. The same report found that jobs requiring data science and analytics skills go unfilled for 45 days, five days longer than average, and more sophisticated analytics roles take considerably longer to fill.⁹³

The shortage of skilled AI workers also creates negative externalities for the education pipeline. Because there is such a shortage, firms are willing to pay a premium for workers with much-sought-after AI skills, often attracting leading academics away from universities, which limits the pool of AI experts available to teach these skills. For example, in May 2015, Uber hired 40 researchers and scientists from Carnegie Mellon University's National Robotics Engineering Center to support its autonomous vehicle development.⁹⁴ Similarly, as of late 2017, six of twenty AI professors at the University of Washington are on full or partial leave to work for the private sector, and four of the most renowned AI researchers in academia have taken leave or fully left their professorships at Stanford University.⁹⁵ Data from the National Science Foundation suggests this trend is endemic, with 58 percent of new computer science PhDs taking jobs in the private sector in 2015, rather than staying in academia, up from 38 percent in 2005.⁹⁶ This is due in no small part to the fact that in 2014, the median annual salary for postdocs in computer science was \$55,000 at universities and \$110,000 in the private sector, but also because of the severe shortage of graduates with doctoral degrees in AI.⁹⁷

This shortage of workers with AI and complementary data science skills will only grow worse without policymakers' intervention.⁹⁸

As part of a national strategy:

- Congress should invest in cultivating AI talent. Countries such as Canada and the United Kingdom have launched initiatives to do just that, and the United States should adapt these approaches. For example, Canada's AI strategy funds programs to attract and retain AI talent in Canadian universities.⁹⁹ Similarly, the United Kingdom's AI Sector Deal describes how the government will fund at least 1,000 AI PhD students by 2025.¹⁰⁰ NSF should create a competitive computer science fellowship program that pays for a stipend to at least 1,000 U.S. citizens or permanent residents to receive their PhD in computer science.
- Congress should fund and authorize a program at NSF to provide competitive awards for up to 1,000 academic AI researchers for a period of five years. Awards should be conditional on remaining in academia for five years. Though individual businesses may benefit from attracting the best AI talent away from universities, the overall AI innovation ecosystem in the United States suffers. These awards would incentivize more AI researchers to stay in academia and help U.S. universities meet the demand for AI skills.
- Congress should enable more foreign AI talent to work in the United States. Increase the cap on H-1B visas to ensure U.S. firms can hire as much AI talent as they need. To the extent global AI talent works in the United States, they are not profiting competitors in other nations.
- Federal agencies should address barriers that limit the number of students able to take computer science courses at the university level. NSF should provide grants to colleges and universities that have increased or are implementing programs to increase enrollment and retention in computer science. The federal government should also require increased transparency as a prerequisite for receiving NSF awards. For instance, schools should be required to monitor and disclose the number of computer science applicants, prospective majors, and their retention rates in computer science subjects. Ideally this would be done for all STEM disciplines.

Research and Development

Public and private R&D investment is a crucial driver of innovation and economic growth. Government investments in R&D have played a key role in many major technology advances over the past several decades, including the Internet, search engines, GPS, supercomputing, speech recognition, assistive robotics, and the underlying technologies in the smartphone.¹⁰¹

However, federal investment in R&D as a share of GDP has fallen from 1.86 percent in 1964 to just 0.66 percent in 2015.¹⁰² And private-sector R&D as a share of GDP has largely stagnated with almost no increase since 2000, and with less going to basic and applied research.¹⁰³ This is why public and private R&D investment in applied research as a share of GDP has fallen over the last two decades.¹⁰⁴

Notwithstanding the fact that major IT companies invest in AI research, federal R&D investment is needed to continue advance AI. This is because companies tend to invest in later stage applied research and development, rather than more foundational basic and early stage research. As stated in NITRD's 2016 National Artificial Intelligence Research and Development Strategic Plan, "some important areas of [AI] research are unlikely to receive sufficient investment by industry, as they are subject to the typical underinvestment problem surrounding public goods."¹⁰⁵

In 2015, the federal government invested approximately \$1.1 billion in unclassified AI R&D.¹⁰⁶ And the Defense Advanced Research Projects Agency's (DARPA) "AI Next" campaign has committed to investing \$2 billion in new and existing AI research initiatives.¹⁰⁷ Efforts to increase R&D coordination are likely to be beneficial. This is why NITRD's AI R&D Strategic Plan includes steps for better coordination of federal R&D investment in AI. But more coordination is not enough; more money is needed as it appears that much more AI research could be productively funded. For example, while NSF funded \$122 million in core AI research in 2017 it received \$174 million in additional proposals for AI research that it deemed either competitive or highly competitive, but that it did not have the budget to support.¹⁰⁸

Other nations are increasing AI R&D at a more rapid pace. As Amir Khosrowshahi, vice president and chief technology officer of Intel's artificial intelligence group, writes, "federal funding levels are not keeping pace with the rest of the industrialized world."¹⁰⁹ For example, China's AI development plan, released in 2017, details the country's plans to grow its AI industries by \$22.15 billion by 2020 and \$59.07 billion by 2025 in part through significant R&D investments.¹¹⁰ The European Union has acknowledged the importance of increasing both public and private R&D in AI, with the goal of increasing it from €4-5 billion (\$4.6-5.8 billion) in 2017 to €20 billion (\$23.2 billion) by 2020, and continuing to increase it beyond that.¹¹¹ To achieve this, the European Commission has stated its intention to increase funding of AI R&D in its Horizon 2020 research program to €1.5 billion by the end of 2020, and has encouraged member states to make similar investments.¹¹²

In summary, failure to boost funding for AI research will put the United States at a disadvantage in the global AI race.

As part of a national strategy:

- Congress should substantially increase R&D funding for AI, with a particular emphasis on basic and applied research. Exact dollar amounts can be debated, but at the very least Congress should provide funding to fully support R&D efforts where the potential supply of high-quality research is greater than the supply of funds. For example, Congress should appropriate at least an additional \$200 million annually to NSF for AI research.
- Federal agencies should support R&D for all kind of AI applications. It is a mistake for federal agencies to support only certain kinds of AI and related technologies. For example, NSF's National Robotics Initiative supports only research that augments, not replaces workers.¹¹³ While augmenting human labor is a valuable application of AI, AI that can replace a human worker does more to boost productivity and should at minimum be on as equal a footing for support as AI that complements workers.
- Congress should increase the R&D tax credit to keep pace with competing countries. A healthy AI ecosystem requires both government and business funding of AI research. Companies will do more AI research in the United States if the R&D tax credit is more generous. However as of 2017, the United States ranked 32nd of OECD in terms of R&D tax credit generosity, behind countries such as Canada, China, Germany, and Japan.¹¹⁴ And in 2018, a number of nations created or expanded their R&D tax incentives.¹¹⁵ Meanwhile, the 2017 tax legislation passed by Congress in 2017 actually increased the after-tax cost of research spending, reducing the incentives of companies to invest in research.¹¹⁶ As such, Congress should increase the Alternative Simplified Credit from 14 percent to 20 percent.¹¹⁷

Accelerate Public Sector Adoption of AI, Including for National Security

One of the most important things government can do to spur AI is to be a robust adopter of AI technologies. This helps drive AI innovation and cost reduction.

However, while the public and private sectors face many of the same challenges in adopting AI, such as a need for skilled workers and robust data sets, the public sector faces unique challenges. These include outdated IT infrastructures, limited funding for capital expenditures, and risk aversion, among others.¹¹⁸ One key challenge is a lack of awareness about the technology and its benefits. As Niall Brennan, former chief data officer at the Centers for Medicare and Medicaid Services (CMS), lamented: “While we did a ton of data work and re-centered and reengineered CMS as a more data-driven organization, I’m afraid AI is so

far off its radar screen that if you said AI to somebody at CMS they might think you were talking about Allen Iverson.”¹¹⁹

To its credit, the Trump administration has directed federal agencies to pursue AI, particularly in the contexts of R&D and national security.¹²⁰ In addition, while the President’s Management Agenda, released in March 2018, does not specifically reference AI, it does emphasize IT modernization and improved use of data as key drivers of government transformation.¹²¹ The Office of Management and Budget (OMB) is also working in conjunction with the Office of Science and Technology Policy (OSTP) and other agencies to develop a Federal Data Strategy, which will emphasize principles such as enterprise data governance, decisionmaking, and innovation.¹²² This could involve AI directly or support government AI adoption indirectly.

Compounding these challenges is the shortage of government workers equipped to work with AI. While this challenge is not unique to government, the public sector struggles to compete with the private sector in attracting and retaining AI talent. As the demand for workers with AI skills increases, the public sector will have an even harder time recruiting this talent, as the private sector has greater flexibility to offer more attractive salaries and benefits than the government. With fewer employees familiar with AI, government agencies will be less aware of the ways in which AI could benefit their missions. And without this expertise, procurement managers will be less able to effectively facilitate AI adoption.

Congress has introduced the bipartisan Artificial Intelligence in Government Act (S. 3502) explicitly to direct agencies to consider AI in their planning and would address many of these strategic shortcomings, including by establishing an AI advisory board for government and directing OMB to study how to foster the necessary workforce skills for effective AI adoption within government. However the bill has yet to be passed.¹²³

Removing barriers to public-sector adoption of AI to enable the government to be a lead adopter of the technology would have important secondary benefits beyond improving agency mission delivery. Federal agencies implementing AI on an enterprise scale would help reduce the perceived risk of the technology and spur additional adoption in the private sector. This would also spur state and local governments, which often look to the federal government as a model, to also support and adopt. For example, the launch of the Obama administration’s open data policy for federal data in May 2013 prompted many state, city, and county government to follow suit.¹²⁴

AI adoption should be of particular concern for the defense community. U.S. national security agencies are well aware of the benefits AI can offer for national security and intelligence, recognizing that maintaining the lead

in AI would help the United States maintain information superiority and enable faster and more accurate decisions both on and off the battlefield.¹²⁵ DoD is investing in AI in fields including intelligence collection and analysis, logistics, cyberspace operations, command and control, and autonomous military vehicles.¹²⁶

Though the DoD and the intelligence community have been at the forefront of AI development and adoption in government, they are still “waiting for AI,” as they face the same obstacles in taking advantage of AI as the rest of the public sector as well as unique challenges that make doing so difficult.¹²⁷ Addressing these obstacles should be a key component of a national AI strategy because, as an April 2018 report from the Congressional Research Service assessing the implications of AI for national security notes, “International rivals in the AI market are creating pressure for the United States to compete for innovative military AI applications.”¹²⁸ Concerns about other countries overtaking the United States in the use of AI for defense has been a focus of the Senate Select Intelligence Committee’s annual hearing on worldwide threat assessment for the past two years. At the 2017 hearing, Director of National Intelligence Daniel Coates stated that ‘The implications of our adversaries’ abilities to use AI are potentially profound and broad. They include an increased vulnerability to cyber-attack, difficulty in ascertaining attribution, [and] facilitation of advances in foreign weapon and intelligence systems...’¹²⁹

The obstacles limiting the use of AI to enhance national security include acquisition challenges, insufficient funding, and oversight questions.

Acquisition challenges are likely the largest barrier to the use of AI for national security purposes. The commercial sector drives the development of AI. However, defense agencies have different priorities, expectations, and limitations than the private sector, making purchasing AI systems or contracting their development from the private sector problematic. First, many dual-use AI technologies—commercial technologies that also have military applications—will likely require significant modification to meet the needs of defense agencies. For example, the private sector is investing heavily in developing autonomous vehicles, but for civilian use—they will operate on public roads which have common rules and known obstacles, large amounts of navigation and mapping data, and a set range of likely driving conditions. By contrast, an autonomous military vehicle could need to drive off-road, in poorly mapped or unmapped areas, in extreme conditions, and in close proximity to humans—operating conditions that would likely make the use of commercial autonomous vehicles highly specialized for civilian use infeasible.¹³⁰

Simply put, the performance and safety standards for civilian AI applications will likely differ significantly from military needs and

standards. However, there is no commercial or government body responsible for the testing and validation of AI performance or safety standards, making determining whether a commercial AI system is fit for military use a cumbersome process, likely requiring tailored versions of commercial technologies and experts within agencies to determine whether individual products meet their standards.¹³¹

DoD's acquisition process itself poses significant obstacles to purchasing AI technologies due to its sluggishness. DoD Instruction 5000.02, which governs this acquisition process, requires a lengthy, multi-phase process for acquiring new technologies that an internal analysis found to take an average of 91 months.¹³² This process is off-putting to technology firms that might otherwise be interested in doing business with DoD. For example, the Government Accountability Office found that 12 companies that decided to not do business with DoD all cited the laboriousness of this process as their reason for not doing business.¹³³

A study from the Center for New American Security also found that leading U.S. technology companies lack the incentives to collaborate with the Department of Defense.¹³⁴ Conversely, DoD leadership found working with the technology sector to be challenging due to technology firms' insistence on preserving their intellectual property. This poses challenges for AI because firms typically sell or license their applications to DoD, rather than the code itself, whereas DoD has a strong interest in scrutinizing this code to ensure it can meet its safety and performance standards.¹³⁵ Overall, the study found that 80 percent of leadership in top Silicon Valley technology firms rated their collaboration with DoD as "poor" or "very poor."¹³⁶

Thus, even though its strong technology sector is one of the United States' main advantages over other countries when it comes to AI, DoD is limited in its ability to work with U.S. technology firms.

Even if procurement challenges were to be addressed, defense agencies may simply lack the funding to pursue AI effectively. Lieutenant General John Shanahan, who leads the Pentagon's algorithmic warfare team Project Maven, has singled out funding as a barrier to the increased use of AI, and the Army Science Board said in a 2017 report that it lacks sufficient funding to pursue AI and other disruptive technologies.¹³⁷ While more funding may be needed, DoD also needs an effective method for tracking defense spending on AI. DoD appropriations are coded with different program elements (PE) based on their area, such as "computer and software technology" or "ground robotics," however there is no PE specifically for AI.¹³⁸ Organizations within DoD may categorize AI funding in many different ways because of this, making it difficult to identify exactly how much DoD is spending on AI compared to how much it is spending on larger projects for which AI is just a component.

Finally, questions about how to oversee DoD's use of AI, if unanswered or if answered poorly, can stymie the ability of the United States to capture the benefits of AI for national security. Most pressing are questions about the development and application of Lethal Autonomous Weapons Systems (LAWS). The debate about if, when, and how a country should deploy LAWS, whether fully or only partially autonomous, is contentious and has drawn the attention of numerous national and international bodies, as well as many activist groups. As of an April 2018 meeting of the United Nations Convention on Conventional Weapons (UN CCW), 26 nations have endorsed banning fully autonomous weapons, while France, Israel, Russia, the United Kingdom, and the United States opposed this call.¹³⁹

Even if these questions are answered, some will continue to oppose defense development and use of AI. Indeed, some leading AI scientists have publicly stated that they will not do work that helps the military develop any AI capabilities and they encourage other AI scientists to follow suit.¹⁴⁰ And groups such as the Campaign to Stop Killer Robots have argued against military funding for AI despite the fact that the U.S. Department of Defense has a policy on autonomous weapons systems.¹⁴¹

While debating the ethical and safety implications of new kinds of military technologies is appropriate, there should be no mistake: opposition to military use of AI will slow the development of AI technology and harm both U.S. national security and competitiveness.¹⁴² For example, a fully autonomous tank will likely rely on large portions of the same algorithms and data used to develop a fully autonomous military transport vehicle.

As part of a national strategy:

- Congress and the administration should support efforts to foster communities of practice and raise awareness about AI within the public sector. For example, the General Services Administration's Emerging Citizen Technology Office (ECTO) coordinates government-wide deployments of AI applications and helps foster relationships between government employees interested in AI and firms working on public-sector applications of AI.¹⁴³
- Congress should provide agencies with venture capital funds to pilot AI initiatives.
- Federal agencies should establish domain-specific programs to spur AI adoption. For example, the Department of Defense's recently established Joint Artificial Intelligence Center (JAIC) will help teams "to swiftly deliver new AI-enabled capabilities and effectively experiment with new operating concepts in support of DoD's military missions and business functions."¹⁴⁴ Other departments, such as HHS and DOT, should consider developing similar programs.

-
- The White House should establish a strategic initiative devoted to AI in the CIO Council. There is currently an initiative devoted to “data analytics and big data” that would likely cover certain aspects affecting AI adoption, however there should be a more explicit focus on AI.¹⁴⁵
 - The General Services Administration (GSA) should work with state government CIOs to share best practices for AI implementation and develop shared resources to make it easier for state officials to learn about and procure AI technologies.
 - Defense agencies should prioritize the use of AI to support their missions to protect national security.
 - DoD should create a body with both government and industry stakeholders to accelerate the adoption of dual-use AI technologies by the military. This could include publishing performance and safety standards for various key military AI applications so industry could more readily develop those solutions or creating guidelines for modifying commercial AI applications for military use.
 - DoD should establish a cross-agency task force to identify opportunities to simplify the acquisition process for AI.
 - DoD should pursue and expand the use of alternative acquisition mechanisms as a workaround for cumbersome procurement policies. For example, the 2016 National Defense Authorization Act granted DoD the permanent Other Transaction Authority (OTA), which allowed DoD to circumvent the traditional acquisition process in certain circumstances.¹⁴⁶
 - DoD should foster better relationships between DoD and the U.S. technology industry, such as by expanding industry outreach efforts like DoD’s Defense Innovation Unit Experimental (DIUx) designed to make it faster for DoD to take advantage of emerging commercial technologies.¹⁴⁷ There should be a specific emphasis on creating greater incentives for technology firms to work with DoD.
 - DoD should establish a new Program Element (PE) specifically for AI to increase the visibility of AI appropriations.¹⁴⁸
 - Congress should prioritize the development and adoption of AI in defense spending. This could entail either focusing greater attention on AI projects relative to less-important work or increasing overall spending.
 - Congress and the administration should support productive conversations about the appropriate way to oversee the use of AI for national security. This will include rejecting bans on LAWS and differentiating between concerns about LAWS specifically and

broader concerns about different military activities, which are often the underlying concern in these discussions.

- Congress and the administration should recognize that the benefits of AI to national security are too important to let concerns about LAWS oversight or other defense activities involving AI limit national security AI support and adoption, particularly because foreign adversaries will pay little heed to such oversight concerns and gain a competitive advantage in certain areas of AI.

Spur AI Development and Adoption in Industry, Including Through Sector-Specific AI Strategies

Many nations look to AI as an important driver of growth and as such they are putting in place a host of development policies designed to grow their domestic AI industry. The United States needs to do the same to ensure its AI industry is robust and globally competitive.

The first place to start is with an understanding that sectors face different challenges in AI adoption and use. As such, a federal AI strategy needs to have a sector-based focus where agencies that have a significant impact on particular sectors identify barriers and opportunities for AI transformation. In some cases, that influence is regulation. For example, the financial sector is influenced by a number of federal regulators. In other cases, the influence is through direct funding, such as for education. For example, state departments of education often rely on the federal Department of Education for guidance, coordination, and funding. In addition to finance and education, these key sectors include, but are not limited to, agriculture, energy and utilities, health care, and transportation.

Agencies have broad authority to influence the ability and willingness of firms and organizations under their purview to develop and use AI. For example, DOT has published several versions of its Federal Automated Vehicles Policy since its original introduction in September 2016.¹⁴⁹ The current draft version, “AV 3.0,” focuses on reducing policy uncertainty that could impede the development and adoption of autonomous vehicles and outlines key considerations and best practices for state transportation authorities and regulators, local governments, and the private sector.¹⁵⁰ Likewise, the Department of Education could support AI tools to help improve pedagogy in K-12 schools and the National Institutes of Health could expand its efforts to develop large health data sets for training AI-powered diagnostic support systems, which would encourage their adoption throughout the healthcare sector. In addition, within health care the federal government is a major purchaser of services, with the Veterans Administration and the DoD both providing health care services.

In addition, local, state and national government economic development programs can help firms develop and implement AI.

As part of a national strategy:

- Federal agencies should work with industry to create strategies for supporting AI adoption in relevant sectors of the economy. These strategies should provide guidance about how best to leverage AI to advance agency missions as well as identify opportunities to encourage AI adoption in relevant industries, such as by proactively providing guidance on policy questions, ensuring that procurement supports AI, ensuring regulations do not limit AI usage, and creating incentives for firms to invest in AI. These strategies should be updated regularly as agencies become more familiar with the technology and as AI matures, creating new challenges and opportunities to address.
- The Department of Commerce should establish organizations designed to advance the development of innovative AI applications in various sectors. For example, Manufacturing USA, overseen by federal agencies including the Department of Commerce and the Department of Energy, is a network of research institutes focused on fostering innovation and collaboration in the manufacturing sector, including the Advanced Robotics for Manufacturing public-private partnership in Pittsburgh, which focuses on AI and automation.¹⁵¹ Using this model, agencies should support similar institutes that include industry, academia, and government agency resources to advance AI in other sectors such as city management and precision medicine.
- Congress should direct the Economic Development Administration to enable state governments to foster AI industry development. Congress should appropriate funds for the Economic Development Administration to create a state economic development competition in which states would compete for funds to establish their own state development plans and policies for supporting AI development, especially through new startups.

Shape Trade Policies to Allow AI to Flourish

The practices of other countries can have a significant impact on how effectively U.S. firms can develop and deploy AI. In particular, efforts to restrict how data can move across borders limits the amount of data at the disposal of U.S. businesses. As of May 2017, 34 countries have proposed or enacted restrictions to the free flow of certain kinds of valuable data, including certain kinds of financial data, personal data, and data from emerging digital services such as online publishing.¹⁵² Countries often attempt to justify these barriers on the grounds of preserving privacy and security, despite the fact that data localization does not in any way guarantee either. In reality these approaches are mercantilist in nature, designed to prop up domestic industries at the expense of productivity.¹⁵³ It is difficult to quantify exactly the harm these barriers cause to the United

States' ability to develop AI, but it is no doubt significant. A 2014 study from the International Trade Commission (ITC) found that removing these barriers to data flows and digital trade would increase U.S. GDP by \$16.7 to \$41.4 billion annually.¹⁵⁴

The United States has sought to protect cross border data flows in previous trade agreements but with limited success. The Trans-Pacific Partnership (TPP), now the Comprehensive and Progressive Agreement for Trans-Pacific Partnership (CPTPP), was the first international trade agreement with explicit language governing the flow of data across borders.¹⁵⁵ However, since the United States withdrew from the TPP, it has less leverage to lobby for these protections abroad.

Fortunately, the United States secured protections for cross border data flows in the United States-Mexico-Canada Agreement (USMCA). Article 19.11 of USMCA prevents parties from restricting “the cross-border transfer of information, including personal information, by electronic means if this activity is for the conduct of the business of a covered person.”¹⁵⁶ Additionally, Article 19.12 prevents parties from requiring people or firms “to use or locate computing facilities in that Party’s territory as a condition for conducting business in that territory.”¹⁵⁷ In effect, these provisions prevent parties from enacting protectionist data localization requirements that inhibit the flow of data across borders.

USMCA includes several other data-related provisions important for AI that can serve as a model for future trade negotiations. For example, USMCA is the first trade agreement in the world to promote the publication of open government data, and though the deal does not require parties to publish open government data, it nonetheless supports the availability of valuable open data as a public resource that can spur AI development.¹⁵⁸ Additionally, USMCA provides much needed protections for source code. Article 19.16 prevents parties from requiring “the transfer of, or access to, source code of software owned by a person of another Party, or to an algorithm expressed in that source code, as a condition for the import, distribution, sale or use of that software, or of products containing that software, in its territory.”¹⁵⁹ This is important for data innovation because it reduces the risk of parties imposing mandates for algorithmic transparency on AI systems developed in other countries thereby exposing them to considerable intellectual property risks. It is easy to imagine how some countries could use algorithmic transparency requirements to force foreign firms to reveal intellectual property that would aid domestic firms. While the agreement would still allow parties to enact algorithmic transparency mandates for all firms, both foreign and domestic, this provision prohibits them using algorithmic transparency as a protectionist measure.

As part of a national strategy:

- The United States Trade Representative (USTR) should continue to advocate for cross-border data flow protections in all future trade negotiations.
- USTR should continue to fight source code disclosure requirements other nations may enact to unfairly disadvantage U.S. firms or exploit their intellectual property.

Foster Innovation-Friendly Regulation

As the public and private sectors increasingly rely on AI to make important decisions with substantial social and economic impacts, there are concerns that AI will exacerbate systemic bias, reduce privacy, or cause other harms. Concerns about the potential harms of AI can have a chilling effect on AI acceptance. Consumers' distrust of AI will drive down demand for AI products and services, and businesses in turn will be less likely to develop or offer these products or services. For example, according to a 2017 Pew survey, 76 percent of Americans say they would not want to apply for jobs where the employer uses a computer program to make hiring decisions.¹⁶⁰ Few, if any, companies, would invest in AI hiring software if it made their hiring process that much less competitive, regardless of the benefits to efficiency or productivity the AI could offer.

Thus, there is a need for regulatory action that addresses potential harms without inhibiting AI innovation. Unfortunately, the most popular proposals to address these concerns would do little to protect consumers but would substantially harm AI innovation. To develop innovation-friendly regulation, it is necessary to understand how consumer harms could arise from AI and why most current proposals would be both ineffective and detrimental to U.S. competitiveness in AI. In most cases, the challenges AI poses to traditional governance stems from its complexity and scalability.¹⁶¹ AI can involve extraordinarily complex decision models involving millions of data points and thousands of lines of code, and these models can change over time as they encounter new data. Further complicating things, in many cases, developers lack the ability to precisely explain how their algorithms make decisions, and instead can only express the degree of confidence they have in the accuracy of the algorithms' decisions.¹⁶² The difficulty arises from the fact that while developers or operators can control what data goes into their systems, and instruct algorithms how to weigh different variables, it can be challenging, if not impossible, to program their systems to explain or justify their decisions.¹⁶³ As a result, many have labeled these algorithms as impenetrable black boxes that defy scrutiny.¹⁶⁴

It is important to recognize that in most cases flawed AI will hurt the organization using it. Therefore, firms do have some incentives not to use biased or otherwise flawed algorithmic decision-making. For example,

banks making loans would be motivated to ensure their AI is not biased because, by definition, errors such as granting a loan to someone who should not receive one, or not granting a loan to someone who is qualified, costs banks money. In some cases, the cost of using flawed AI may be reputational, and a company will lose market share to a competitor that ensured their AI did not cause harm, creating similar financial incentives to avoid using biased or otherwise flawed AI.

However, there are cases where the cost of the harm falls largely on the subject of an AI algorithm's decision, and in these cases incentives to avoid harmful AI may be diminished or may not exist at all. Biased algorithms in parole decision systems, for instance, hurt individuals who are unfairly denied parole, but impose little cost on a court system. A court system does not have an explicit financial incentive to be fair, and a person standing trial cannot simply take their case to a competing court system. In such cases, existing legal frameworks may not be sufficiently equipped to respond quickly or effectively to mitigate this risk. These situations are more likely to arise in the public sector, where market forces play a diminished role in influencing how AI is used compared to the private sector.

Thus, there is some validity to the concerns that existing legal and regulatory frameworks may be insufficient to effectively protect against these kinds of potential harms AI can produce. However, the most popular ideas about how to address potential harms are flawed. In particular, many have expressed support for mandating algorithmic transparency, which requires organizations to expose their algorithms and information about their data to some degree of public scrutiny.¹⁶⁵ Similarly, many have expressed support for algorithmic explainability, which requires the parties responsible for deploying an algorithm, or an “operator,” to make their algorithms interpretable to end users, such as by having operators describe how their algorithms work or by using algorithms capable of articulating rationales for their decisions.¹⁶⁶

While transparency and explainability are fundamentally different concepts, they share many of the same flaws as a solution for regulating algorithms. First, they hold algorithmic decisions to a standard that simply does not exist for human decisions. As the Electronic Privacy Information Center (EPIC) describes, “Without knowledge of the factors that provide the basis for decisions, it is impossible to know whether government and companies engage in practices that are deceptive, discriminatory, or unethical. Therefore, algorithmic transparency is crucial to defending human rights and democracy online.”¹⁶⁷ This argument fails to recognize that algorithms are simply a recipe for decision-making. If proponents of algorithmic transparency and explainability are concerned that these decisions are harmful, then it is counterproductive to call for algorithmic

decisions alone to be transparent or explainable, rather than for all aspects of all decisionmaking to be made public or explained. If blanket mandates for transparency and explainability are appropriate for algorithmic decisionmaking, but not human decisionmaking (which itself is often supported by computers), logic would dictate that human decisions are already transparent, fair, and free from unconscious and overt biases. In reality, bias permeates every aspect of human decisionmaking, so to hold algorithms to a higher standard than for humans is simply unreasonable. For example, research shows taxicabs frequently do not pick up passengers based on their race, and employers may eliminate job applicants with African-American sounding names despite their sufficient qualifications.¹⁶⁸ Yet, understandably, taxi drivers are not required to publicly report their reasons for not picking up every passenger they pass by, and employers do not have to publish a review of every resume they receive, with detailed notes explaining why they choose not to offer a particular candidate a job, because laws and regulations for these sectors focus on outcomes, not unconscious bias. If EPIC and other proponents of algorithmic transparency and explainability worry that such broad categories of decisions have the potential to be harmful due to the influence of bias, then they should advocate for transparency and explainability in all significant decisionmaking, as an algorithm's involvement in those decisions is irrelevant. By targeting algorithms specifically however, such proposals risk imposing significant regulatory costs on the use of AI for a wide array of benign or beneficial applications, disincentivizing its adoption at the direct expense of the productivity of U.S. firms.

Importantly, calls for algorithmic transparency and, sometimes, for algorithmic explainability discount the value of proprietary software. Requirements to publicly disclose source code or information about the inner workings of software would reduce incentives for a company to invest in developing algorithms, as competitors could simply copy them. While copyright laws could reduce this risk in countries with strong intellectual property protections like the United States, this would make it significantly easier for bad actors in countries that routinely flout intellectual property protections, such as China, to steal source code.¹⁶⁹

The ideal method to minimize the potential harms of AI would be to pursue regulatory frameworks based on algorithmic accountability: the principle that an algorithmic system should employ a variety of controls to ensure the operator (i.e., the party responsible for deploying the algorithm) can verify it acts in accordance with its intentions, as well as identify and rectify harmful outcomes.¹⁷⁰ Algorithmic accountability promotes desirable outcomes, protects against harmful ones, and ensures algorithmic decisions are subject to the same requirements as human decisions. This approach is technology neutral, granting operators flexibility to employ a

variety of different technical and procedural mechanisms to achieve algorithmic accountability. Importantly, algorithmic accountability is relevant only when an application of algorithmic decisionmaking poses potential harms significant enough to warrant regulatory scrutiny, and not, for example, applications that only pose the risk of minor inconveniences should the algorithms involved be flawed.

The first step in achieving algorithmic accountability is determining whether algorithms are working the way their operators intended. If the answer is yes, and it is causing harm, then it is important to recognize that regulations already exist in various industries that prohibit racial discrimination, require due process, and so on. When an operator intends to cause harm, whether or not they use an algorithm to do so should be irrelevant. There are a variety of different technical and procedural mechanisms that can be employed, when contextually relevant, to make the determination of whether a harm is intentional. These include: transparency, explainability, confidence measures, and procedural regularity. In most cases, operators would likely have to employ a combination of several of these mechanisms in order to be confident an algorithm is acting as they intended. This is not meant to be a comprehensive list of all the ways an operator can verify an algorithm is acting as intended, as there may be methods that are only useful in niche circumstances or that have yet to be developed.

Simply taking steps to verify an algorithm is acting as intended is not enough to ensure it is not also producing harmful outcomes. Thus, an accountable algorithmic system must also allow operators to identify and minimize harmful outcomes. This is an important capability because it allows for organizations to responsibly deploy algorithms despite not being able to predict or control for every possible harmful outcome that could arise from an algorithm's decisions—which would likely be impossible and could severely limit the utility of algorithms. There are a variety of methods to accomplish this that allow operators to take meaningful steps to minimize harms. These include, but are not limited to, impact assessment, error analysis, and bias testing. Importantly, these are not simply just post hoc controls—operators can and should apply these steps throughout the entire process of developing and deploying an algorithm, and continuously employ them throughout the time an algorithm is in use.

Using an algorithmic accountability framework, regulators can take a straightforward approach to evaluating and punishing operators whose algorithms violate existing laws or regulations and produce significant harms worthy of regulatory scrutiny. Importantly, this standard is open to interpretation and will change over time as market forces, social norms, new technologies, and other factors shape the use of algorithmic decision-making.

When operators violate existing laws or regulations using an algorithm, regulators should first examine whether and how effectively the operator can demonstrate they had controls to ensure the algorithm was acting as intended. Operators could be subject to higher levels of punishment if a significant harm occurred and no such controls were present, or if operators were careless or superficial in their approach to meeting this standard. If these controls were thorough and implemented appropriately, a regulator could likely determine the operator was not acting with negligence or with intent to harm. At this point, regulators could conduct a similar analysis of whether and how effectively the operator could identify and rectify harmful outcomes. If operators fail to meet this standard, then a regulator could conclude they were irresponsible in their efforts to minimize the potential harms of their algorithms and again be subject to more punishment.

In select cases where market forces are muted, and significant harm is possible, it may be appropriate for policymakers to dictate specific requirements for algorithmic accountability. This is particularly relevant in the criminal justice system. Caleb Watney, a technology policy fellow at the R Street Institute, argues that because the concept of transparency is central to the goals of the justice system, as indicated by countless court precedents and statutory obligations, such as the Freedom of Information Act and other “sunshine” laws, it would be appropriate to mandate all algorithms that influence judicial decisionmaking be open-source.¹⁷¹ Though this transparency may not shed much light on how more advanced machine learning systems work, there is likely a compelling public interest in ensuring these algorithms are nonetheless exposed to the highest degree of scrutiny possible. Similarly, it would likely be appropriate for policymakers to mandate that public agencies conduct thorough impact assessments for algorithms they intend to use in decisions with high social or economic consequences, such as the administration of entitlement programs.¹⁷² However, any such rules should be narrow and targeted to identifiable harms that algorithmic decisionmaking could cause in a specific context.

Enforcing algorithmic accountability in this way would have important benefits. If operators know this framework exists, they can take proactive steps to ensure they embrace algorithmic accountability, such as by modifying existing systems to increase their transparency, or by discontinuing the use of algorithms that fail to meet these standards. Similarly, this would send a market signal to developers about what customers will expect of an algorithmic system, thus encouraging them to provide algorithms with the necessary capabilities or risk losing market share to competitors that do so.¹⁷³

Finally, as noted above, one of the most important inputs for the AI economy is data, and as such any national strategy should enable organizations to collect, manage and use data effectively. However, even if the United States were to develop a national AI strategy, new privacy rules could easily undermine this strategy. For example, many in the United States have advocated for the adoption of privacy laws similar to those of the EU's General Data Protection Regulation (GDPR), which would limit AI development and use.¹⁷⁴ Some provisions of GDPR, including the right to erasure and the prohibition on repurposing data, will significantly limit the ability of organizations to access the data necessary to develop and use AI effectively. For example, AI systems that operate using unsupervised machine learning—those that improve themselves, without outside help, by learning from the data they process—will be required to “remember” all the data they used to train themselves in order to sustain rules derived from that data.¹⁷⁵ However, erasing data that underpins key rules in an AI system's behavior can both make it less accurate and limit its benefit to other data subjects—or even break it entirely. Additionally, the GDPR imposes a general prohibition on using data for any purposes other than that for which it was first collected, thus making it difficult for firms to innovate using data. Moreover, the GDPR's broad requirements for obtaining express consent before collecting data from users serves as a general deterrent to organizing collecting and using data. If the United States were to adopt similar data protection legislation, these types of restriction will limit the ability of domestic companies developing or using AI to experiment with new functions that could improve their services. As a result, U.S. consumers and businesses will be slow to receive the benefits of the latest innovations in AI.¹⁷⁶

Acknowledging the impact of restrictive regulation should be just as important for antitrust regulators as it is for consumer protection regulators. Having a competitive marketplace in which firms have strong incentives to invest in AI is crucial to advancing development and adoption of the technology. However, overly zealous antitrust efforts can reduce these incentives, leading to decreased investment in AI.

Developing and using AI effectively requires large amounts of data. However some commentators have begun to argue that competition policy should be extended to incorporate concerns about the collection and use of data beyond clear examples of anticompetitive behavior.¹⁷⁷ Their argument is that the mere act of collecting large amounts of data, such as the personal data collected by social-networking platforms, search engines, and e-commerce sites, gives companies an unfair competitive advantage and that competition policy needs to incorporate this analysis.

To date, U.S. regulators have not adopted this line of reasoning, nor should they. While it is true that data can be used in anticompetitive ways,

competition policy is capable of dealing with such abuses.¹⁷⁸ In fact, when analyzing allegations of such behavior, it is often helpful to imagine whether agencies would object if the activity complained about involved some input of critical importance, such as machinery, other than data. This helps clarify whether the threat to competition is truly due to control of an important resource or to ungrounded fears about the uniqueness of data.

The collection of large amounts of data does not by itself represent a threat to competition. Although use of data might in specific circumstances justify regulatory intervention, in most cases the acquisition and use of data does not reduce competition, and the existing legal framework, including traditional interpretations of existing statutes, gives competition regulators all the flexibility they need to protect markets. On the contrary, large amounts of data, including personal information, are increasingly a vital input for some of the most valuable applications of AI, such as medical diagnostics, digital assistants, language translation, and autonomous vehicles.

Economists Anja Lambrecht and Catherine Tucker examined big data using a resource-based view of the firm, which holds that for a resource such as data to provide a company with a competitive advantage, it must be inimitable, rare, valuable, and non-substitutable.¹⁷⁹ They conclude that:

The unstable history of digital business offers little evidence that the mere possession of big data is a sufficient protection for an incumbent against a superior product offering. To build a sustainable competitive advantage, the focus of a digital strategy should therefore be on how to use digital technologies to provide value to customers in ways that were previously impossible.¹⁸⁰

While developing and using AI does indeed often require large amounts of data, the possession of such data does not necessarily constitute an unfair competitive advantage over new AI firms. Many industries have high start-up costs. Few would argue that Ford and GM have an unfair advantage just because companies must first build an expensive factory before they sell a single car. Nor does amassing a large number of workers represent a barrier to competition, even though these same workers are not available to competitors. In contrast, collecting data can be relatively cheap, and the data remains available to others.

Although barriers to entry are an element of antitrust analysis, these barriers can be less imposing than they look. Companies have often been able to overcome high upfront costs, provided they have a compelling business plan for eventually earning enough profits to deliver an appropriate risk-adjusted rate of return. An entire ecosystem of angel investors, incubators, and hedge funds exists to invest in promising young

companies capable of growing rapidly. Although funding is often a challenge, the larger bottleneck remains a lack of innovative and workable ideas.

Finally, it is important to recognize that different companies have different strategies and business models around data. For some companies their competitive advantage is the algorithm; for others, including some who are making their algorithms open source, it is the data. For the former case, IBM is training its cognitive computing system, Watson, to help analyze medical information, including the discovery of new drugs for immuno-oncology.¹⁸¹ To do this, it needs lots of data. But the data would be much less valuable without Watson's sophisticated artificial intelligence capabilities. Sometimes these algorithms are protected as intellectual property, but that does not prohibit competitors from trying to write better ones. And sometimes these algorithms are made public.¹⁸² For example, Google published the source code for its artificial intelligence engine, TensorFlow, to encourage others to find uses for it, and ways of improving it, which Google might not have considered.¹⁸³ But even the best algorithms can be defeated by poor business strategy. As an example, one former executive attributes the fall of MySpace largely to poor business decisions.¹⁸⁴

As part of a national strategy:

- Regulators should encourage adherence to the principle of algorithmic accountability. Importantly, policymakers need to recognize that the goal of algorithmic accountability is not to achieve perfect, error-free algorithms, but to minimize risk—just as vehicle safety standards do not require cars to be 100 percent safe, but only as safe as can reasonably be expected. The most important step is for regulators to formally recognize this framework for algorithmic accountability and integrate it into their oversight. This applies to both domain-specific and consumer-protection regulators.
- Congress should reject blanket mandates for algorithms, such as algorithmic transparency requirements, or the creation of new regulatory bodies focused only on regulating algorithms.
- Congress and the administration should increase the technical expertise of regulators and policymakers. Regulators should foster relationships with communities of developers, academics, civil society groups, and private sector organizations invested in algorithmic decisionmaking to stay abreast of technical developments and concerns about algorithmic harms that could influence how algorithmic accountability is achieved or enforced. This requires ensuring regulators have the resources to hire staff

with the necessary technical expertise to scrutinize algorithms. Similarly, a national strategy should ensure that policymakers in Congress and the administration have access to the technical expertise necessary to understand the potential harms of AI and the implications regulations might have on national competitiveness. This could include expanding the roster of staffers in the White House Office of Science and Technology Policy or reviving the Office of Technology Assessment to provide technical expertise to Congressional offices.¹⁸⁵

- Congress and the administration should caution regulators against viewing the mere act of collecting or possessing large amounts of data, which is necessary for certain uses of AI, as potentially anticompetitive behavior.
- Congress should reject overly stringent privacy legislation rules, such as broad opt-in requirement, purpose specification, data erasure, and data minimization, as well as other policies modeled on the EU's General Data Protection regulation.
- Congress and the administration should emphasize to policymakers and regulators that data is a crucial business input for the development of AI, and that companies should be encouraged to invest in collecting data, not punished for it. If policymakers are concerned that startups and small businesses cannot access the data necessary to develop AI and compete with larger firms, they should focus on making it easier to collect data and ensure data is readily available, as described earlier in this report, rather than penalize a company that has succeeded in doing so.

Provide Workers with Better Tools to Manage AI-Driven Workforce Transitions

Perhaps one of the biggest challenges to ubiquitous AI adoption is the growing opposition to AI based on the view that it is a job killer and will lead to high rates of unemployment. Many who hold this view believe the government should slow the deployment of AI. A 2017 Pew survey found that 72 percent of adults are worried about a future in which computers can perform many jobs done by humans, 76 percent believe economic inequality will grow worse if computers can perform these jobs, and 75 percent do not believe the economy will create enough new, higher-paying jobs for humans should computers be able to perform these jobs.¹⁸⁶ This anxiety has undoubtedly negatively colored Americans' attitudes about policies governing the use of AI. Pew also found that 87 percent of Americans would support a law requiring a human to be in the driver seat of an autonomous vehicle to take over in an emergency, and should AI be able to perform most of the jobs done by humans today, 85 percent would

support laws limiting the use of automation to only jobs that are hazardous to human health or safety.¹⁸⁷ Additionally, 58 percent would support establishing a national service program that would pay humans for tasks that machines could complete faster or for less money.¹⁸⁸ Perceptions about this skills gap have already engendered support for policies that would be bad for innovation and growth, such as a universal basic income or a tax on automation equipment and software.

This is not to imply no workers will be displaced and required to find new jobs. A number of studies have tried to estimate this impact, with perhaps the most widely cited among them, by Oxford's Osborne and Frey, estimating 47 percent of U.S. jobs could be eliminated by technology over the next 20 years.¹⁸⁹ But their study appears to significantly overstate the real number by including many jobs that have little chance of automation, such as fashion models, school bus drivers, and barbers.

The OECD estimates around 15 percent of U.S. jobs will be lost to automation over the next 15 years. The Information Technology and Innovation Foundation (ITIF) estimates 20 percent of U.S. jobs are likely to be automated over the next 15 years.¹⁹⁰ And the McKinsey Global Institute estimates the percent of jobs displaced by technology in the G7 to range from 20 percent of jobs in the United Kingdom to 26 percent of jobs in Japan.¹⁹¹

One reason actual job-loss rates are unlikely to reach the higher-end estimates of nearly 50 percent or above is that automation affects jobs as a whole to a lesser extent than certain specific tasks that comprise those jobs. As McKinsey concludes, "Very few occupations will be automated in their entirety in the near or medium term. Rather, certain activities are more likely to be automated, requiring entire business processes to be transformed, and jobs performed by people to be redefined."¹⁹² In other words, technology will lead to more jobs being redefined and opportunities created in order to add more value, rather than to outright job destruction.

ITIF has written extensively on the impact of AI and automation on the labor market. For a detailed list of proposals to help ease labor-market transitions, see ITIF's report, "How to Reform Worker-Training and Adjustment Policies for an Era of Technological Change."¹⁹³

CONCLUSION

Today's global competition for advanced technology industries is unique. Not only is the competition fiercer than ever as many more nations are competing for the same emerging advanced technologies, including AI, but for the first time since before WWII, the United States faces a true global competitor for military preeminence, raising the stakes on the competition. In addition, AI is likely to be a key technology to spur productivity and

economic growth as well generate as a wide range of social benefits. These factors make the development of a proactive national AI strategy to support the development and adoption of the technology a necessity for the United States. The benefits of AI—to the competitiveness of firms in the United States, to economic growth, to government operations, and to social welfare—and the risks of falling behind are too vast for policymakers to either sit on the sidelines hoping private-sector action is enough, or to believe that the government's main role should be shaping and constraining AI through regulation without concerning themselves with the challenges the private sector faces. It is time for a national AI development and adoption strategy.

REFERENCES

1. “Sizing the Prize” (PWC, June 2017), <https://www.pwc.com/gx/en/issues/analytics/assets/pwc-ai-analysis-sizing-the-prize-report.pdf>.
2. “Country Spotlights: Why Artificial Intelligence Is the Future of Growth” (Accenture, 2016), https://www.accenture.com/t20170202T122451Z__w__/us-en/_acnmedia/PDF-33/Accenture-Why-AI-is-the-Future-of-Growth-Country-Spotlights.pdf?la=en.
3. Daniel Castro and Joshua New, “The Promise of Artificial Intelligence” (Center for Data Innovation, October 2016), <http://www2.datainnovation.org/2016-promise-of-ai.pdf>.
4. Casey Newton, “Facebook Begins Using Artificial Intelligence to Describe Photos to Blind Users,” *The Verge*, April 5, 2016, <http://www.theverge.com/2016/4/5/11364914/facebook-automatic-alt-tags-blind-visually-impaired>.
5. Vivian Giang, “The Growing Business of Detecting Unconscious Bias,” *Fast Company*, May 5, 2015, <https://www.fastcompany.com/3045899/hit-the-ground-running/the-growing-business-of-detecting-unconscious-bias>.
6. Jody Avirgan, “Big Data Is Saving This Little Bird,” *FiveThirtyEight*, October 16, 2015, <http://fivethirtyeight.com/datalab/big-data-is-saving-this-little-bird/>.
7. Jacques Bughin, “Artificial Intelligence: The Next Digital Frontier?,” (McKinsey Global Institute, June 2017), <https://www.mckinsey.com/~media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx>.
8. Ibid.
9. Joshua New and Daniel Castro, “How Policymakers Can Foster Algorithmic Accountability” (Center for Data Innovation, May 2017), <http://www2.datainnovation.org/2018-algorithmic-accountability.pdf>.
10. Nick Wallace and Daniel Castro, “The Impact of the EU’s New Data Protection Regulation on AI,” (Center for Data Innovation, March 2018), <http://www2.datainnovation.org/2018-impact-gdpr-ai.pdf>.
11. “Top AI Trends to Watch in 2018,” (CB Insights, February 2018), https://www.cbinsights.com/reports/CB-Insights_State-of-Artificial-Intelligence-2018.pdf.
12. Keith Shaw, “New Index Ranks Robot, AI and Automation Readiness for 25 Countries,” *Robotics Business Review*, April 23, 2018, <https://www.roboticsbusinessreview.com/research/new-index-ranks-countries-that-are-ready-for-robots-ai/>.
13. “Top AI Trends to Watch in 2018,” (CB Insights, February 2018), https://www.cbinsights.com/reports/CB-Insights_State-of-Artificial-Intelligence-2018.pdf.

-
14. *Rise of the Machines: Artificial intelligence and its Growing Impact on U.S. Policy* (Washington, DC: House of Representatives Subcommittee on Information Technology, September 2018), <https://oversight.house.gov/wp-content/uploads/2018/09/AI-White-Paper-.pdf>.
 15. “National Science Board Statement on Global Research and Development (R&D) Investments NSB-2018-9,” February 7, 2018, https://www.nsf.gov/nsb/news/news_summ.jsp?cntn_id=244465.
 16. Robert Button, “Artificial Intelligence and the Military,” Rand Corporation, September 7, 2017, <https://www.rand.org/blog/2017/09/artificial-intelligence-and-the-military.html>.
 17. Tom Upchurch, “How China Could Beat the West in the Deadly Race for AI Weapons,” *Wired*, August 8, 2018, <https://www.wired.co.uk/article/artificial-intelligence-weapons-warfare-project-maven-google-china>.
 18. Tim Weitzel, *Economics of Standards in Information Networks*, (PhysicaVerlag Heidelberg, 2004), <https://books.google.com/books?id=Ae37CAAAQBAJ&printsec=frontcover#v=onepage&q&f=false>
 19. Brendan Bordelon, “Why Won’t Congress Craft a National AI Strategy?,” *National Journal*, August 15, 2018, <https://www.nationaljournal.com/s/671537/why-wont-congress-craft-a-national-ai-strategy>.
 20. William Galston, “Why the Government Must Help Shape the Future of AI,” Brookings Institution, September 13, 2018, <https://www.brookings.edu/research/why-the-government-must-help-shape-the-future-of-ai/>.
 21. “France’s AI Strategy,” France in the UK, March 29, 2018, <https://uk.ambafrance.org/France-s-AI-strategy>.
 22. “Artificial Intelligence Industry in Germany to Receive €3 Billion,” *Government Europa*, November 16, 2018, <https://www.governmenteuropa.eu/artificial-intelligence-industry-germany/91164/>; *Key Points for a Federal Government Strategy on Artificial Intelligence* (Germany: Federal Cabinet, 2018), <https://www.bmwi.de/Redaktion/EN/Downloads/E/key-points-for-federal-government-strategy-on-artificial-intelligence.pdf>.
 23. Tim Dutton, “A Timeline for Europe’s AI Strategy,” Medium, May 7, 2018, <https://medium.com/swlh/a-timeline-for-europes-ai-strategy-d2fc9f7bbcf1>.
 24. Tim Dutton, “An Overview of National AI Strategies,” Medium, June 28, 2018, <https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd>.
 25. “Pan-Canadian Artificial Intelligence Strategy,” CIFAR, Accessed November 26, 2018, <https://www.cifar.ca/ai/pan-canadian-artificial-intelligence-strategy>.
 26. “*Building a Strong Middle Class* (Canada: House of Commons, March 2017), <https://www.budget.gc.ca/2017/docs/plan/budget-2017-en.pdf>.

-
27. “Supercluster Tech Groups Share in Up to \$950M in Federal Funds,” CTV News, February 15, 2018, <https://www.ctvnews.ca/business/supercluster-tech-groups-to-share-in-up-to-950m-in-federal-funds-1.3804876>.
 28. “China Issues Guideline on Artificial Intelligence Development,” The State Council, July 20, 2017, http://english.gov.cn/policies/latest_releases/2017/07/20/content_281475742458322.htm.
 29. Graham Webster et al., “Full Translation: China’s ‘New Generation Artificial Intelligence Development Plan’ (2017),” New America, August 1, 2017, <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>.
 30. Paul Mozur, “Beijing Wants A.I. to Be Made in China by 2030,” *The New York Times*, July 20, 2017, <https://www.nytimes.com/2017/07/20/business/china-artificial-intelligence.html>.
 31. Kenji Kawase, “In China, Private Companies Walk a Fine Line,” *Nikkei Asian Review*, May 23, 2018, <https://asia.nikkei.com/Spotlight/Cover-Story/In-China-private-companies-walk-a-fine-line2>.
 32. Lindsey Sheppard et al., Artificial Intelligence and National Security (Center for Strategic and International Studies, November 2018), https://csis-prod.s3.amazonaws.com/s3fs-public/publication/181102_AI_interior.pdf?6jofgIIr0rJ2qFc3.TCg8jQ8p.Mpc81X.
 33. Zhou Jiaquan, “Drones, Facial Recognition and a social credit system: 10 Ways China Watches its Citizens,” *South China Morning Post*, August 4, 2018, <https://www.scmp.com/news/china/society/article/2157883/drones-facial-recognition-and-social-credit-system-10-ways-china>.
 34. Kai-Fu Lee, *AI Superpowers: China, Silicon Valley, and the New World Order* (Boston: Houghton Mifflin Harcourt, 2018).
 35. *Findings of the Investigation into China’s Acts, Policies, and Practices Related to Technology, Intellectual Property, and Innovation Under Section 301 of the Trade Act of 1974* (Washington, D.C.: Office of the United States Trade Representative, 2018), <https://ustr.gov/sites/default/files/Section%20301%20FINAL.PDF>; *Report on the Protection and Enforcement of Intellectual Property Rights in Third Countries* (Brussels: European Commission, 2018), http://trade.ec.europa.eu/doclib/docs/2018/march/tradoc_156634.pdf; Kelly Olson, “China Respects Others’ Trade Secrets – Unless it Wants Something, Experts, Say,” *CNBC*, October 4, 2018, <https://www.cnbc.com/2018/10/04/trade-secret-protection-remains-a-challenge-in-china-even-experts.html>.
 36. Ibid.
 37. Cédric Villani, *For a Meaningful Artificial Intelligence* (France: Parliament, March 2018), https://www.aiforhumanity.fr/pdfs/MissionVillani_Report_ENG-VF.pdf; Laura Stevens, “Emmanuel Macron Talks to Wired About France’s AI Strategy,” *Wired*, March 31, 2018,

-
- <https://www.wired.com/story/emmanuel-macron-talks-to-wired-about-frances-ai-strategy/>.
38. Nick Wallace, "Countries Can Learn from France's Plan for Public Interest Data and AI," Center for Data Innovation, August 14, 2018, <https://www.datainnovation.org/2018/08/countries-can-learn-from-frances-plan-for-public-interest-data-and-ai/>.
 39. Ibid.
 40. Ibid; Nick Wallace and Daniel Castro, "The Impact of the EU's New Data Protection Regulation on AI," (Center for Data Innovation, March 2018), <http://www2.datainnovation.org/2018-impact-gdpr-ai.pdf>.
 41. *National Strategy for Artificial Intelligence* (India: NTIA Aayog, June 2018), http://niti.gov.in/writereaddata/files/document_publication/NationalStrategy-for-AI-Discussion-Paper.pdf.
 42. Ibid.
 43. Ibid.
 44. *Japan Revitalization Strategy 2016* (Japan, 2016), http://www.kantei.go.jp/jp/singi/keizaisaisei/pdf/zentaihombun_160602_en.pdf.
 45. *Artificial Intelligence Technology Strategy* (Japan: Strategic Council for AI Technology, March 2017), <http://www.nedo.go.jp/content/100865202.pdf>.
 46. "Japan Eyes Self-Driving Trucks, AI in New Growth Strategy," *The Mainichi*, June 4, 2017, <https://mainichi.jp/english/articles/20170604/p2a/00m/0na/006000c>.
 47. Mark Zastrow, "South Korea Trumpets \$860-million AI fund after AlphaGo 'Shock'," *Nature*, March 18, 2016, <https://www.nature.com/news/south-korea-trumpets-860-million-ai-fund-after-alphago-shock-1.19595>.
 48. "Korean AI, 1.8 Years Behind US, China Overtake...Government Gets 2.2 Trillion Won," *JoongAng Daily*, May 15, 2018, <https://news.joins.com/article/22625271>; Tim Dutton, "An Overview of National AI Strategies," Medium, June 28, 2018, <https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd>.
 49. Tim Dutton, "An Overview of National AI Strategies," Medium, June 28, 2018, <https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd>.
 50. Ibid.
 51. *Robotics and Artificial intelligence* (London: House of Commons, September 2016), <https://publications.parliament.uk/pa/cm201617/cmselect/cmsctech/145/145.pdf>.
 52. *Artificial Intelligence: Opportunities and Implications for the Future of Decision Making* (London: Government Office for Science, 2016), https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/566075/gs-16-19-artificial-intelligence-ai-report.pdf.
 53. "UK Digital Strategy," Department for Digital, Culture, Media, and Sport, March 2017, <https://www.gov.uk/government/publications/uk-digital->
-

-
- strategy; Department for Digital, Culture, Media, and Sport, “£17 Million Boost for the UK's Booming Artificial Intelligence Sector,” news release, February 26, 2017, <https://www.gov.uk/government/news/17-million-boost-for-the-uks-booming-artificial-intelligence-sector>.
54. *AI Sector Deal* (London: Department for Digital, Culture, Media, and Sport and Department for Business, Energy, and Industrial Strategy, April 2016), <https://www.gov.uk/government/publications/artificial-intelligence-sector-deal>.
55. Ibid.
56. Ed Felten and Terah Lyons, “Public Input and Next Steps on the Future of Artificial Intelligence,” Medium, September 6, 2016, <https://medium.com/@USCTO/public-input-and-next-steps-on-the-future-of-artificial-intelligence-458b82059fc3>; “Request for Information on Artificial Intelligence,” Federal Register, June 27, 2016, <https://www.federalregister.gov/documents/2016/06/27/2016-15082/request-for-information-on-artificial-intelligence>.
57. *Preparing for the Future of Artificial Intelligence* (Washington, D.C.: National Science and Technology Council, October 2016), https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf.
58. *The National Artificial Intelligence Research and Development Strategic Plan* (Washington, D.C.: Networking and Information Technology Research and Development Subcommittee, October 2016), https://www.nitrd.gov/PUBS/national_ai_rd_strategic_plan.pdf.
59. *Artificial Intelligence, Automation, and the Economy* (Washington, D.C.: Executive Office of the President, December 2016), <https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.PDF>.
60. *Summary of the 2017 White House Summit on Artificial Intelligence for American Industry* (Washington, D.C.: Office of Science and Technology policy, May 2015), <https://www.whitehouse.gov/wp-content/uploads/2018/05/Summary-Report-of-White-House-AI-Summit.pdf>.
61. “Request for Information on Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan,” Federal Register, September 26, 2018, <https://www.federalregister.gov/documents/2018/09/26/2018-20914/request-for-information-on-update-to-the-2016-national-artificial-intelligence-research-and>.
62. “Game Changers: Artificial Intelligence Part I,” Subcommittee on Information Technology, February 14, 2018, <https://oversight.house.gov/hearing/game-changers-artificial-intelligence-part/>.
63. FUTURE of Artificial Intelligence Act of 2017, S. 2217, 115th Congress (2017); AI in Government Act of 2018, S. 3502, 115th Congress (2018).
64. Sam Ransbotham et al., “Reshaping Business with Artificial Intelligence,” MIT Sloan Management Review, September 6, 2017, <https://sloanreview.mit.edu/projects/reshaping-business-with-artificial-intelligence/>.

-
65. Megan Molteni, "Health Care is Hemorrhaging Data. AI is Here to Help," *Wired*, December 30, 2017, <https://www.wired.com/story/health-care-is-hemorrhaging-data-ai-is-here-to-help/>.
 66. Eleni Manis, "Americans Want to Share Their Medical Data. So Why Can't They?," *RealClear Health*, July 26, 2018, https://www.realclearhealth.com/articles/2018/07/26/americans_want_to_share_their_medical_data_so_why_cant_they_110807.html.
 67. "Health Insurance Portability and Accountability Act Privacy Rule Causes Ongoing Concerns among Clinicians and Researchers," *Annals of Internal Medicine* 145 (2016) <http://annals.org/aim/article-abstract/727746/health-insurance-portability-accountability-act-privacy-rule-causes-ongoing-concerns?doi=10.7326%2f0003-4819-145-4-200608150-00019>.
 68. "Legal Issues Relating to Patient Privacy," in *Sharing Clinical Research Data: Workshop Summary*, (Institute of Medicine, 2013), <https://www.nap.edu/read/18267/chapter/4#20>.
 69. "Top 10 Countries for EHR Adoption," *Becker's Hospital Review*, June 27, 2013, <https://www.beckershospitalreview.com/healthcare-information-technology/top-10-countries-for-ehr-adoption.html>.
 70. "Nearly Half of All U.S. Electricity Customers Have Smart Meters," U.S. Energy Information Administration, December 6, 2017, <https://www.eia.gov/todayinenergy/detail.php?id=34012>.
 71. Stephen Ezell et al., "Manufacturing Digitalization: Extent of Adoption and Recommendations for Increasing Penetration in Korea and the U.S." (ITIF, August 2018), <http://www2.itif.org/2018-korean-manufacturing-digitalization.pdf>.
 72. Joshua New, "Congress Has a New Cure to Stop Companies from Blocking Patient Data," Center for Data Innovation, November 10, 2015, <https://www.datainnovation.org/2015/11/congress-has-a-new-cure-to-stop-companies-from-blocking-patient-data/>.
 73. Ibid; Jennifer Bresnick, "FHIR Takes off as ONC Teases Upcoming Data Blocking Definitions," *Health IT Analytics*, October 1, 2018, <https://healthitanalytics.com/news/fhir-takes-off-as-onc-teases-upcoming-data-blocking-definitions>.
 74. Benedikt Fecher et al., "What Drives Academic Data Sharing," *PLoS One* 10, no:2 (February 2015), <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4340811/>; "Cultural Barriers," in *Sharing Clinical Research Data: Workshop Summary*, (Institute of Medicine, 2013), <https://www.nap.edu/read/18267/chapter/4#22>.
 75. Dame Wendy Hall and Jérôme Pesenti, "Growing the Artificial Intelligence Industry in the UK," (London: Independent Report, October 2017), https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/652097/Growing_the_artificial_intelligence_industry_in_the_UK.pdf.
 76. Daniel Castro, "The Rise of Data Poverty in America," (Center for Data Innovation, September 2014), <http://www2.datainnovation.org/2014-data-poverty.pdf>.

-
77. Cédric Villani, *For a Meaningful Artificial Intelligence* (France: Parliament, March 2018), https://www.aiforhumanity.fr/pdfs/MissionVillani_Report_ENG-VF.pdf.
 78. Hal Varian, "Artificial Intelligence, Economics, and Industrial Organization," National Bureau of Economic Research, November 2017, <https://www.nber.org/chapters/c14017.pdf>.
 79. James Vincent, "Tencent Says There Are Only 300,000 AI Engineers Worldwide, but Millions are Needed," *The Verge*, December 5, 2017, <https://www.theverge.com/2017/12/5/16737224/global-ai-talent-shortfall-tencent-report>
 80. "Global AI Talent Report," JG Gagné, Accessed November 26, 2018, <http://www.jfgagne.ai/talent/>.
 81. Ibid.
 82. Caleb Watney, "To Win the AI Race, We Need More Humans," *The Hill*, August 14, 2018, <https://thehill.com/opinion/cybersecurity/401773-to-win-the-ai-race-we-need-more-humans>; "The Importance of International Students to American Science and Engineering," National Foundation for American Policy, October 2017, <http://nfap.com/wp-content/uploads/2017/10/The-Importance-of-International-Students.NFAP-Policy-Brief.October-20171.pdf>.
 83. Ibid.
 84. "Report on H-1B Petitions," (Washington, D.C.: U.S. Citizenship and Immigration Services, February 2016), <https://www.uscis.gov/sites/default/files/USCIS/Resources/Reports%20and%20Studies/H-1B/H-1B-FY-2015-Petitions.pdf>.
 85. Caleb Watney, "To Win the AI Race, We Need More Humans," *The Hill*, August 14, 2018, <https://thehill.com/opinion/cybersecurity/401773-to-win-the-ai-race-we-need-more-humans>.
 86. Youyou Zhou, "Chinese Students Increasingly Return Home After Studying Abroad," *Quartz*, July 29, 2018, <https://qz.com/1342525/chinese-students-increasingly-return-home-after-studying-abroad/>.
 87. "Global AI Talent Report," JG Gagné, Accessed November 26, 2018, <http://www.jfgagne.ai/talent/>.
 88. "LinkedIn's 2017 U.S. Emerging Jobs Report," LinkedIn, December 7, 2017, <https://economicgraph.linkedin.com/research/LinkedIns-2017-US-Emerging-Jobs-Report>.
 89. Jeremy Kahn, "Ski-High Salaries Are the Weapons in the AI Talent War," *Bloomberg*, February 14, 2018, <https://www.bloomberg.com/news/articles/2018-02-13/in-the-war-for-ai-talent-sky-high-salaries-are-the-weapons>.
 90. Sam Ransbotham et al., "Reshaping Business with Artificial Intelligence," MIT Sloan Management Review, September 6, 2017, <https://sloanreview.mit.edu/projects/reshaping-business-with-artificial-intelligence/>.
 91. "The Growing Impact of AI on Business," *MIT Technology Review*, April 30, 2018, <https://www.technologyreview.com/s/611013/the-growing-impact-of-ai-on-business/>.
 92. Will Markow et al., "The Quant Crunch," (Burning Glass Technologies, 2017),

-
- <https://public.dhe.ibm.com/common/ssi/ecm/im/en/iml14576usen/analytics-analytics-platform-im-analyst-paper-or-report-iml14576usen-20171229.pdf>.
93. Ibid.
94. Mike Ramsey and Douglas MacMillan, "Carnegie Mellon Reels After Uber Lures Away Researchers," *The Wall Street Journal*, May 31, 2015, <https://www.wsj.com/articles/is-uber-a-friend-or-foe-of-carnegie-mellon-in-robotics-1433084582>.
95. Cade Metz, "Tech Giants Are Paying Huge Salaries for Scarce A.I. Talent," *The New York Times*, October 22, 2017, <https://www.nytimes.com/2017/10/22/technology/artificial-intelligence-experts-salaries.html>.
96. Daniela Hernandez and Rachael King, "Universities' AI Talent Poached by Tech Giants," *The Wall Street Journal*, November 24, 2016, <https://www.wsj.com/articles/universities-ai-talent-poached-by-tech-giants-1479999601>.
97. Ibid.
98. Daron Acemoglu and Pascual Restrepo, "Artificial Intelligence, Automation and Work," NBER Working Paper No. 24196, January 4, 2018, <https://economics.mit.edu/files/14641>.
99. "Pan-Canadian Artificial Intelligence Strategy," CIFAR, Accessed November 26, 2018, <https://www.cifar.ca/ai/pan-canadian-artificial-intelligence-strategy>.
100. Conner Forrest, "UK Government Funds 1000 AI PhDs to Push Country's Tech Innovation," *TechRepublic*, April 26, 2018, <https://www.techrepublic.com/article/uk-government-funds-1000-ai-phds-to-push-countrys-tech-innovation/>.
101. Peter Singer, "Investing in 'Innovation Infrastructure' to Restore Growth" (ITIF, January 2017), <http://www2.itif.org/2017-innovation-infrastructure.pdf>
102. "U.S. gross domestic product, R&D, and ratio of R&D to gross domestic product (and components): 1953–2016," National Science Foundation, Accessed October 25, 2018, <https://www.nsf.gov/statistics/2018/nsf18309/pdf/np16-dst-tab001.pdf>.
103. John Wu, "Why U.S. Business R&D Is Not as Strong as It Appears" (ITIF, June 2018), <http://www2.itif.org/2018-us-business-rd.pdf>.
104. "Federal R&D Budget Dashboard," AAAS, Accessed November 26, 2018, <https://www.aaas.org/programs/r-d-budget-and-policy/federal-rd-budget-dashboard>.
105. *The National Artificial Intelligence Research and Development Strategic Plan* (Washington, D.C.: Networking and Information Technology Research and Development Subcommittee, October 2016), https://www.nitrd.gov/PUBS/national_ai_rd_strategic_plan.pdf.
106. Ibid.
107. "AI Next Campaign," DARPA, Accessed November 26, 2018, <https://www.darpa.mil/work-with-us/ai-next-campaign>.

-
108. “Game Changers: Artificial Intelligence Part II, Artificial Intelligence and the Federal Government,” YouTube, March 7, 2018, https://www.youtube.com/watch?time_continue=3182&v=pehKIVgPiSQ.
 109. *Rise of the Machines: Artificial intelligence and its Growing Impact on U.S. Policy* (Washington, DC: House of Representatives Subcommittee on Information Technology, September 2018), <https://oversight.house.gov/wp-content/uploads/2018/09/AI-White-Paper-.pdf>.
 110. Cate Cadell and Adam Jourdan, “China Aims to Become World Leader in AI, Challenges U.S. Dominance,” *Reuters*, July 20, 2017, <https://www.reuters.com/article/us-china-ai/china-aims-to-become-world-leader-in-ai-challenges-u-s-dominance-idUSKBN1A5103?feedType=RSS&feedName=technologyNews>; Graham Webster et al., “China’s Plan to ‘Lead’ in AI: Purpose, Prospects, and Problems,” August 1, 2017, <https://www.newamerica.org/cybersecurity-initiative/blog/chinas-plan-lead-ai-purpose-prospects-and-problems/>.
 111. “Communication Artificial Intelligence for Europe,” European Commission, April 25, 2018, <https://ec.europa.eu/digital-single-market/en/news/communication-artificial-intelligence-europe>
 112. Ibid.
 113. “National Robotics Initiative 2.0: Ubiquitous Collaborative Robots (NRI-2.0)” National Science Foundation, Accessed November 26, 2018, <https://nifa.usda.gov/sites/default/files/grant/FY%202017%20NRI%20nsf17518%20from%20published%20NSF%20web%20www%20nsf%20gov.pdf>.
 114. “Measuring Tax Support for R&D and Innovation: Indicators,” OECD, Accessed November 26, 2018, <http://www.oecd.org/sti/rd-tax-incentive-indicators.htm>.
 115. Joe Kennedy, “Countries Continue to Use Tax Incentives to Boost R&D,” Information Technology and Innovation Foundation, October 25, 2018, <https://itif.org/publications/2018/10/25/countries-continue-use-tax-incentives-boost-rd>
 116. Ibid.
 117. Joe Kennedy and Robert Atkinson, “Why Expanding the R&D Tax Credit Is Key to Successful Corporate Tax Reform” (Information Technology and Innovation Foundation, July 2017), <http://www2.itif.org/2017-rd-tax-credit.pdf>.
 118. Kevin Desouza, “Delivering Artificial Intelligence in Government: Challenges and Opportunities” (IBM Center for the Business of Government, 2018,) <http://www.businessofgovernment.org/sites/default/files/Delivering%20Artificial%20Intelligence%20in%20Government.pdf>.
 119. Greg Slabodkin, “Lack of Access to Health Data Said to Limit Potential of Machine Learning,” May 30, 2017, <https://www.information-management.com/news/lack-of-access-to-health-data-said-to-limit-potential-of-machine-learning>.
 120. “Artificial Intelligence for the American People,” The White House, May 20, 2018, <https://www.whitehouse.gov/briefings-statements/artificial-intelligence-american-people/>.

-
121. “President’s Management Agenda” (Washington, D.C.: Office of Management and Budget, March 2018), <https://www.whitehouse.gov/wp-content/uploads/2018/03/Presidents-Management-Agenda.pdf>.
 122. “What is the Federal Data Strategy?,” Federal Data Strategy, Accessed November 26, 2018, <https://strategy.data.gov/>.
 123. AI in Government Act of 2018, S. 3502, 115th Congress (2018).
 124. “Open Data Policy Collection,” Sunlight Foundation, Accessed November 26, 2018, <https://opendatapolicyhub.sunlightfoundation.com/collection/>.
 125. Lindsey Sheppard et al., Artificial Intelligence and National Security (Center for Strategic and International Studies, November 2018), https://csis-prod.s3.amazonaws.com/s3fs-public/publication/181102_AI_interior.pdf?6jofgIIIR0rJ2qFc3.TCg8jQ8p.Mpc81X.
 126. “Daniel Hoadley and Nathan Lucas, “Artificial Intelligence and National Security” (Washington, D.C.: Congressional Research Service, April 2016), <https://fas.org/sgp/crs/natsec/R45178.pdf>.
 127. Lindsey Sheppard et al., Artificial Intelligence and National Security (Center for Strategic and International Studies, November 2018), https://csis-prod.s3.amazonaws.com/s3fs-public/publication/181102_AI_interior.pdf?6jofgIIIR0rJ2qFc3.TCg8jQ8p.Mpc81X.
 128. “Daniel Hoadley and Nathan Lucas, “Artificial Intelligence and National Security” (Washington, D.C.: Congressional Research Service, April 2016), <https://fas.org/sgp/crs/natsec/R45178.pdf>.
 129. Daniel Coats, “Worldwide Threat Assessment of the US Intelligence Community” (Washington, D.C.: Office of the Director of National Intelligence, May 2017), <https://www.intelligence.senate.gov/sites/default/files/documents/os-coats-051117.pdf>.
 130. Daniel Hoadley and Nathan Lucas, “Artificial Intelligence and National Security” (Washington, D.C.: Congressional Research Service, April 2016), <https://fas.org/sgp/crs/natsec/R45178.pdf>.
 131. Ibid.
 132. Ibid.
 133. Ibid; “DOD Is Taking Steps to Address Challenges Faced by Certain Companies” (Washington, D.C.: U.S. Government Accountability Office, July 2017), <https://www.gao.gov/assets/690/686012.pdf>.
 134. Loren Schulman et al., “The Rocky Relationship Between Washington and Silicon Valley,” Center for a New American Security, July 19, 2017, <https://s3.amazonaws.com/files.cnas.org/documents/COPIA-CNAS-Rocky-Relationship-Between-Washington-And-Silicon-Valley.pdf>.
 135. Ibid.
 136. Ibid.
 137. Daniel Hoadley and Nathan Lucas, “Artificial Intelligence and National Security” (Washington, D.C.: Congressional Research Service, April 2016), <https://fas.org/sgp/crs/natsec/R45178.pdf>.

-
138. Daniel Hoadley and Nathan Lucas, "Artificial Intelligence and National Security" (Washington, D.C.: Congressional Research Service, April 2016), <https://fas.org/sgp/crs/natsec/R45178.pdf>.
 139. Tucker Davey, "Lethal Autonomous Weapons: An Update from the United Nations," Future of Life Institute, April 30, 2018, <https://futureoflife.org/2018/04/30/lethal-autonomous-weapons-an-update-from-the-united-nations/>.
 140. For example, one leading Canadian AI scientist, Yoshua Bengio, when asked if AI scientists should not work with the military stated "If they had the right moral values, fine. But I don't completely trust military organizations, because they tend to put duty before morality." Will Knight, "One of the Fathers of AI is Worried About its Future," *MIT Technology Review*, November 17, 2018, <https://www.technologyreview.com/s/612434/one-of-the-fathers-of-ai-is-worried-about-its-future/>.
 141. Robert Atkinson, "The 2015 ITIF Luddite Award Nominees: The Worst of the Year's Worst Innovation Killers" (Information Technology and Innovation Foundation, December 2015), <http://www2.itif.org/2015-itif-luddite-award.pdf>.
 142. Daniel Hoadley and Nathan Lucas, "Artificial Intelligence and National Security" (Washington, D.C.: Congressional Research Service, April 2016), <https://fas.org/sgp/crs/natsec/R45178.pdf>.
 143. "Game Changers: Artificial Intelligence Part II; Artificial Intelligence and the Federal Government," General Services Administration, March 7, 2018, <https://www.gsa.gov/about-us/newsroom/congressional-testimony/game-changers-artificial-intelligence-part-ii-artificial-intelligence-and-the-federal-government>.
 144. Aaron Mehta, "DoD Stands Up its Artificial Intelligence Hub," *C4ISRNET*, June 29, 2018, <https://www.c4isrnet.com/it-networks/2018/06/29/dod-stands-up-its-artificial-intelligence-hub/>.
 145. "Innovation Committee," CIO.gov, Accessed November 26, 2018, <https://www.cio.gov/about/innovation-committee/>.
 146. "Other transaction Authority (OTA)," AcqNotes, Accessed November 26, 2018, <http://acqnotes.com/acqnote/careerfields/other-transaction-authority-ota>.
 147. Patrick Tucker, "Pentagon Shakes Up Silicon Valley Outreach," *Defense One*, May 11, 2016, <https://www.defenseone.com/technology/2016/05/pentagon-shakes-silicon-valley-outreach/128198/>.
 148. Daniel Hoadley and Nathan Lucas, "Artificial Intelligence and National Security" (Washington, D.C.: Congressional Research Service, April 2016), <https://fas.org/sgp/crs/natsec/R45178.pdf>.
 149. "USDOT Automated Vehicles 3.0 Activities," U.S. Department of Transportation, October 4, 2018, <https://www.transportation.gov/av/3>.
 150. Ibid.
 151. "Advanced Robotics for Manufacturing," Arm Institute, Accessed November 26, 2018, <http://arminstitute.org/>.
 152. Nigel Cory, "Cross-Border Data Flows: Where Are the Barriers, and What Do They Cost?" (Information Technology and Innovation Foundation, May 2017), <http://www2.itif.org/2017-cross-border-data-flows.pdf>.

-
153. Daniel Castro, “The False Promise of Data Nationalism” (Information Technology and Innovation Foundation, December 2013), <http://www2.itif.org/2013-false-promise-data-nationalism.pdf>.
 154. Nigel Cory, “Cross-Border Data Flows: Where Are the Barriers, and What Do They Cost?” (Information Technology and Innovation Foundation, May 2017), <http://www2.itif.org/2017-cross-border-data-flows.pdf>.
 155. Susan Aaronson, “Data Minefield? How AI is Prodding Governments to Rethink Trade in Data,” Centre for International Governance Innovation, April 3, 2018, <https://www.cigionline.org/articles/data-minefield-how-ai-prodding-governments-rethink-trade-data>.
 156. “United States-Mexico-Canada Agreement Text,” Office of the United States Trade Representative, Accessed November 26, 2018, <https://ustr.gov/trade-agreements/free-trade-agreements/united-states-mexico-canada-agreement/united-states-mexico>.
 157. Ibid.
 158. Ibid.
 159. Ibid.
 160. Aaron Smith and Monica Anderson, “5. American’s Attitudes Towards Hiring Algorithms,” Pew Research Center, October 4, 2017, <http://www.pewinternet.org/2017/10/04/americans-attitudes-toward-hiring-algorithms/>.
 161. Joshua New and Daniel Castro, “How Policymakers Can Foster Algorithmic Accountability” (Center for Data Innovation, May 2017), <http://www2.datainnovation.org/2018-algorithmic-accountability.pdf>.
 162. Will Knight, “The Dark Secret at the Heart of AI,” MIT Technology Review, April 11, 2017, <https://www.technologyreview.com/s/604087/the-darksecret-at-the-heart-of-ai/>
 163. Cliff Kuang, “Can A.I. Be Taught to Explain Itself?” The New York Times, November 21, 2017, <https://www.nytimes.com/2017/11/21/magazine/can-ai-be-taught-toexplain-itself.html>.
 164. “Humans May Not Always Grasp Why AIs Act. Don’t Panic.” The Economist, February 15, 2018, <https://www.economist.com/news/leaders/21737033-humans-areinscrutable-too-existing-rules-and-regulations-can-applyartificial?frsc=dg%7Ce>.
 165. Joshua New and Daniel Castro, “How Policymakers Can Foster Algorithmic Accountability” (Center for Data Innovation, May 2017), <http://www2.datainnovation.org/2018-algorithmic-accountability.pdf>.
 166. Ibid.
 167. “Algorithmic Transparency: End Secret Profiling,” EPIC, Accessed May 8, 2018, <https://www.epic.org/algorithmic-transparency/>
 168. Cornell Belcher and Dee Brown, “Hailing While Black—Navigating the Discriminatory Landscape of Transportation” (key findings from the hailing while black survey of Chicago voters, Brilliant Corners, February 12, 2015), <http://www.brilliant-corners.com/post/hailing-while-black>; David R. Francis, “Employers’ Replies to Racial Names” (The National Bureau of

-
- Economic Research, accessed May 16, 2016), <http://www.nber.org/digest/sep03/w9873.html>
169. Patrick Gillespie, “China Broke Hacking Pact Before New Tariff Fight,” Axios, April 10, 2018, <https://www.axios.com/china-broke-hacking-pactbefore-new-tariff-tiff-d19f5604-f9ce-458a-a50a-2f906c8f12ab.html>.
170. Joshua New and Daniel Castro, “How Policymakers Can Foster Algorithmic Accountability” (Center for Data Innovation, May 2017), <http://www2.datainnovation.org/2018-algorithmic-accountability.pdf>.
171. Caleb Watney, “When it Comes to Criminal Justice AI, We Need Transparency and Accountability,” R Street Institute, December 1, 2017, <http://www.rstreet.org/2017/12/01/when-it-comes-to-criminal-justice-aiwe-need-transparency-and-accountability/>
172. The Center for Data Innovation’s report, “How Policymakers Can Foster Algorithmic Accountability,” provides an in-depth analysis of the ways AI could cause harm, a more thorough critique of the most popular proposals to regulate AI, and a more detailed explanation of how regulators should implement algorithmic accountability.
173. Ibid.
174. Gus Rossie, “Is the GDPR Right for the United States?” Public Knowledge, April 9, 2018, <https://www.publicknowledge.org/news-blog/blogs/is-the-gdpr-right-for-the-united-states>; “America Should Borrow from Europe’s Data-Privacy Law,” *The Economist*, April 5, 2018, <https://www.economist.com/leaders/2018/04/05/america-should-borrow-from-europes-data-privacy-law>.
175. Eduard Forsch, Peter Kiseberg, and Tiffany Li, “Humans Forget, Machines Remember: Artificial Intelligence and the Right to be Forgotten,” *Computer Security & Law Review* (forthcoming), August 15, 2017, accessed December 16, 2017, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3018186
176. Nick Wallace and Daniel Castro, “The Impact of the EU’s New Data Protection Regulation on AI,” (Center for Data Innovation, March 2018), <http://www2.datainnovation.org/2018-impact-gdpr-ai.pdf>.
177. For the most comprehensive argument see Maurice E. Stucke and Allen P. Grunes, *Big Data and Competition Policy* (New York: Oxford University Press, 2016). <http://www2.itif.org/2017-data-competition.pdf>
178. Joe Kennedy, “The Myth of Data Monopoly: Why Antitrust Concerns About Data Are Overblown” (Information Technology and Innovation Foundation, March 2017), <http://www2.itif.org/2017-data-competition.pdf>.
179. Anja Lambrecht and Catherine E. Tucker, “Can Big Data Protect a Firm From Competition?” *Antitrust Chronicle* 1, no. 12 (January 2017).
180. Ibid., 17.
181. “Will Artificial Intelligence Help to Crack Biology?” *The Economist*, January 7, 2017, <https://www.economist.com/science-and-technology/2017/01/07/will-artificial-intelligence-help-to-crack-biology>.
182. Cade Metz, “Google Just Open Sourced TensorFlow, Its Artificial Intelligence Engine” *Wired*, November 9, 2015,

-
- <https://www.wired.com/2015/11/google-open-sources-its-artificial-intelligence-engine/>.
183. Cade Metz “Google Just Open Sourced TensorFlow, Its Artificial Intelligence Engine,” *Wired*, September 9, 2015, <https://www.wired.com/2015/11/google-open-sources-its-artificial-intelligence-engine/>.
184. Stuart Dredge, “MySpace—What Went Wrong: ‘The Site Was a Massive Spaghetti-Ball Mess,’” *The Guardian*, March 6, 2015, <https://www.theguardian.com/technology/2015/mar/06/myspace-what-went-wrong-sean-percival-spotify>.
185. Zach Graves and Kevin Kosar, “Bring In the Nerds: Reviving the Office of Technology Assessment,” R Street Institute, January 2018, <https://2o9ub0417chl2lg6m43em6psi2i-wpengine.netdna-ssl.com/wp-content/uploads/2018/04/Final-128-1.pdf>.
186. Aaron Smith and Monica Anderson, “5. American’s Attitudes Towards Hiring Algorithms,” Pew Research Center, October 4, 2017, <http://www.pewinternet.org/2017/10/04/americans-attitudes-toward-hiring-algorithms/>.
187. Ibid.
188. Ibid.
189. Carl Benedikt Frey and Michael A. Osborne, “The Future of Employment: How Susceptible Are Jobs to Computerisation?” (Oxford Martin School, University of Oxford, Oxford, September 17, 2013), http://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf.
190. Ben Miller, “Automation Not So Automatic,” Information Technology and Innovation Foundation, September 20, 2013, <http://www.innovationfiles.org/automation-not-so-automatic/>.
191. James Manyika et al., “Jobs Lost, Jobs Gained: Workforce Transitions In A Time Of Automation,” (McKinsey Global Institute, December 2017), <https://www.mckinsey.com/global-themes/future-of-organizations-and-work/what-the-future-of-work-will-mean-for-jobs-skills-and-wages>.
192. Michael Chui, James Manyika, and Mehdi Miremadi, “Four Fundamentals of workplace automation,” (McKinsey & Company: November 2015), <http://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/four-fundamentals-of-workplace-automation>.
193. Robert Atkinson, “How to Reform Worker-Training and Adjustment Policies for an Era of Technological Change” (Information Technology and Innovation Foundation, February 2018), <http://www2.itif.org/2018-innovation-employment-workforce-policies.pdf>.

ABOUT THE AUTHOR

Joshua New is a senior policy analyst at the Center for Data Innovation. He has a background in government affairs, policy, and communication. New graduated from American University with degrees in C.L.E.G. (communication, legal institutions, economics, and government) and public communication.

ABOUT THE CENTER FOR DATA INNOVATION

The Center for Data Innovation is the leading global think tank studying the intersection of data, technology, and public policy. With staff in Washington, D.C. and Brussels, the center formulates and promotes pragmatic public policies designed to maximize the benefits of data-driven innovation in the public and private sectors. It educates policymakers and the public about the opportunities and challenges associated with data, as well as technology trends such as predictive analytics, open data, cloud computing, and the Internet of Things. The center is a nonprofit, nonpartisan research institute proudly affiliated with the Information Technology and Innovation Foundation.

contact: info@datainnovation.org

datainnovation.org