CARNEGIE INDIA

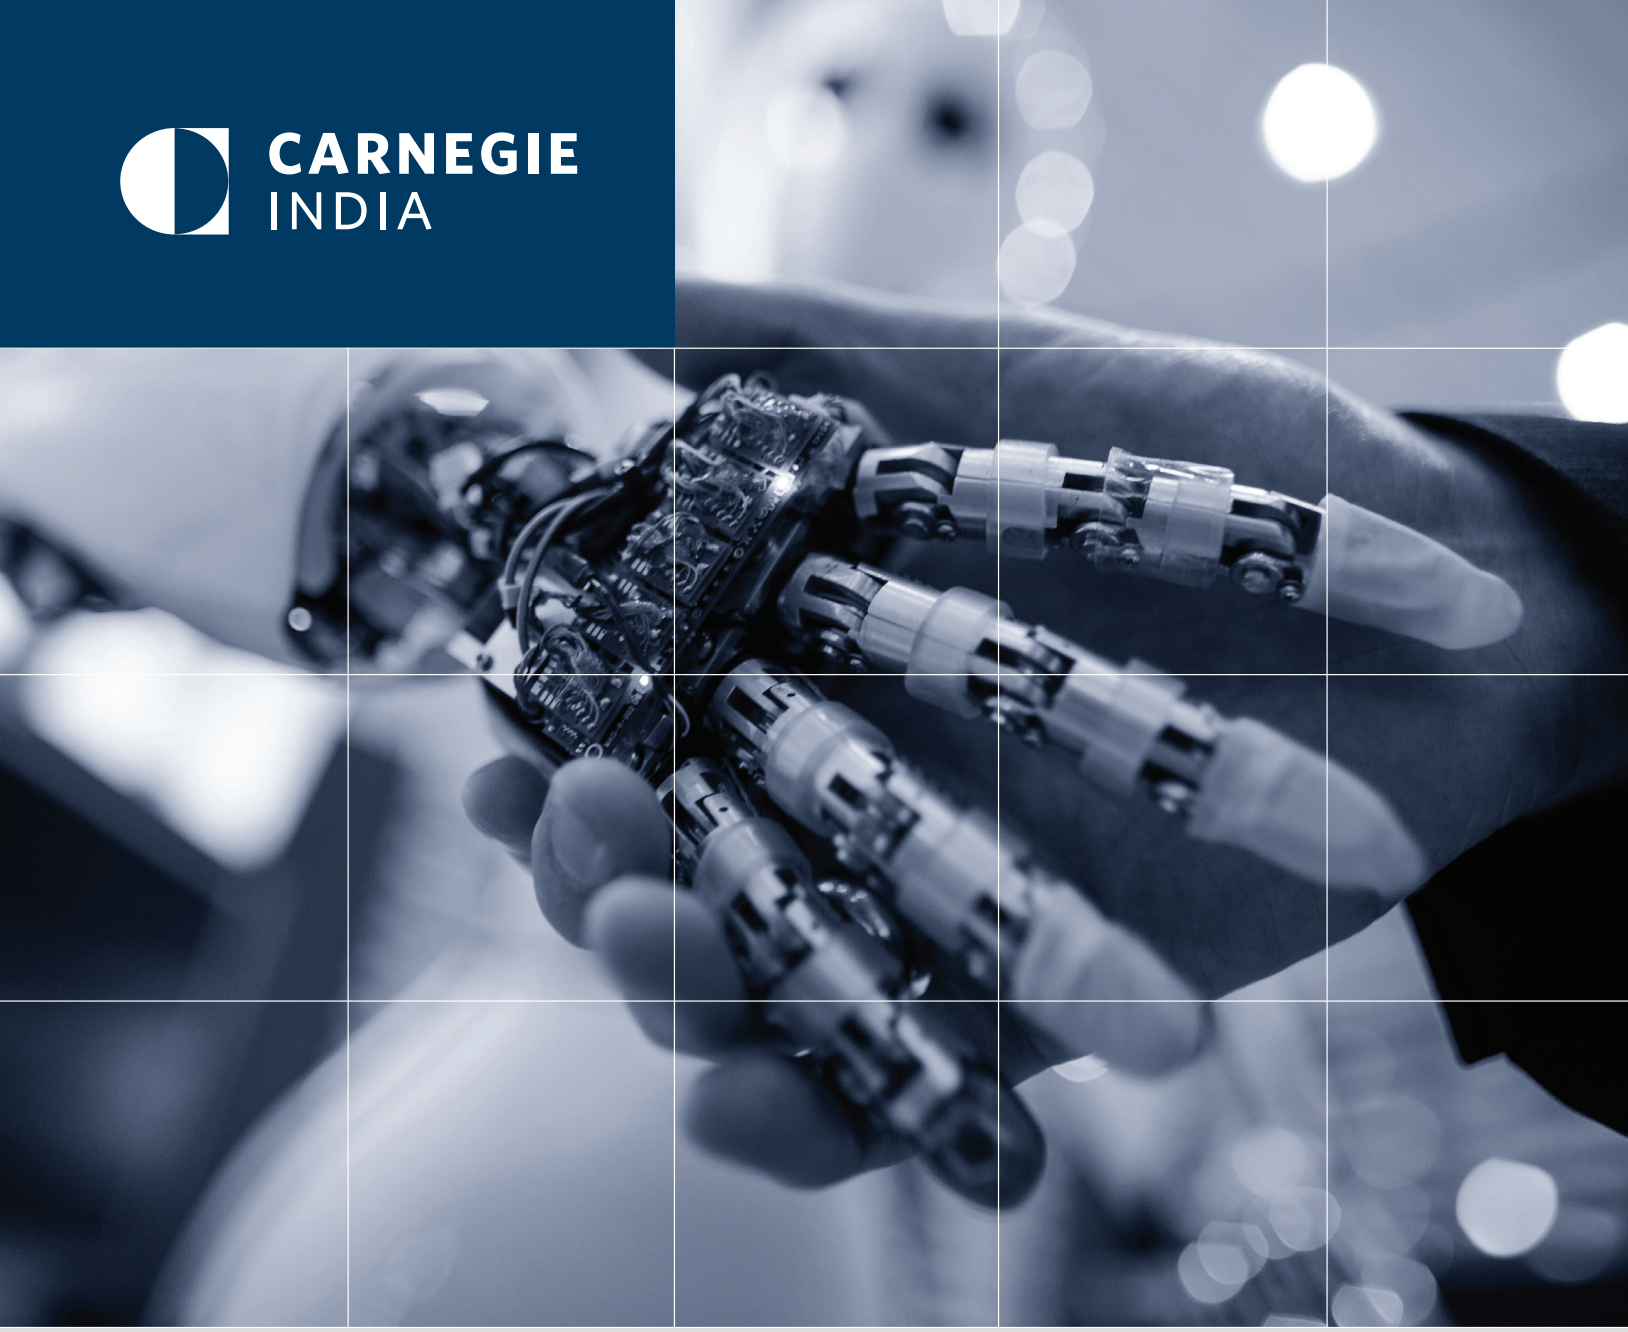

INDIA AND THE ARTIFICIAL INTELLIGENCE REVOLUTION

Shashi Shekhar Vempati

AUGUST 2016

CARNEGIE
INDIA

INDIA AND THE ARTIFICIAL INTELLIGENCE REVOLUTION

Shashi Shekhar Vempati

Carnegie Endowment for International Peace
Publications Department
1779 Massachusetts Avenue NW
Washington, DC 20036
P: +1 202 483 7600
F: +1 202 483 1840
CarnegieEndowment.org

This publication can be downloaded at no cost at CarnegieIndia.org

CP 283

Contents

About the Author

Shashi Shekhar Vempati is a digital strategist and a commentator on politics and public policy in India. An alumnus of the Indian Institute of Technology, Bombay, Vempati serves on the board of India's public broadcaster Prasar Bharati. Prior to that position, Shashi was CEO at Niti Digital, which played a pivotal role in the digital campaign of Prime Minister Narendra Modi during the 2014 Indian election. Before joining Niti Digital, Shashi was a principal architect at Infosys Technologies for over a decade.

Summary

Recent advances in artificial intelligence (AI) are a wake-up call to policymakers in India, with every one of Prime Minister Narendra Modi's flagship programs likely to be directly affected within the next few years. With China making rapid progress in AI-based research, it is imperative that India view AI as a critical element of national security strategy. Spurring AI-based innovation and establishing AI-ready infrastructure are thus necessary to prepare India's jobs and skills markets for an AI-based future and to secure its strategic interests.

The Challenges Facing India's AI Development

- AI-based applications to date have been driven largely by the private sector and have been focused primarily in consumer goods. The emergent scale and implications of the technology make it imperative for policymakers in government to take notice.
- Early lessons of AI success in the United States, China, South Korea, and elsewhere offer public and private funding models for AI research that India should consider.
- The sequential system of education and work is outdated in today's economic environment as the nature of jobs shifts rapidly and skills become valuable and obsolete in a matter of years.

Recommendations

For India to maximally benefit from the AI revolution, it must adopt a deliberate policy to drive AI innovation, adaptation, and proliferation in sectors beyond consumer goods and information technology services.

Policymakers should make AI a critical component of the prime minister's flagship Make in India, Skill India, and Digital India programs by offering incentives for manufacturers, creating regional innovation clusters for manufacturing automation and robotics in partnership with universities and start-ups, incorporating market-based mechanisms for identifying the kind of skills that employers will value in the future, and promoting cloud infrastructure capacity building inside India.

The National Education Policy must make radical recommendations on alternative models of education that would be better suited to an AI-powered economy of the future.

The government should identify public sector applications like detecting tax fraud, preventing subsidy leakage, and targeting beneficiaries, where current advances in AI could make a significant impact.

India must view machine intelligence as a critical element of its national security strategy and evaluate models of defense research in collaboration with the private sector and universities.

Introduction

Recent advances in artificial intelligence (AI) have stimulated fervent interest from both the private sector and governments across the globe, as the possibility of mass-produced consumer product machinery with humanlike intelligence inches closer to reality.

The big breakthrough for artificial intelligence in recent months was the victory of machine over human in the ancient board game Go. AlphaGo, an AI-based computer developed by London-based Google DeepMind,[1] challenged the world champion of the Chinese board game, Lee Sedol of South Korea, to a series of five games in which the machine defeated the human four to one. While AlphaGo deservedly captured headlines across the globe, the real breakthrough in artificial intelligence is not this singular event but the impressive advances artificial intelligence–based computer programs have made as a technology, to the point that they can learn and intelligently respond across a wide range of problem domains.

AI-based applications today have already touched people's lives in ways that are often not fully perceived.

AI-based applications today have already touched people's lives in ways that are often not fully perceived or fathomed. Until now, this subtle proliferation of AI technology has been driven largely by the private sector and has been focused primarily on consumer goods. The technology, however, is of such great potential and importance that its development and implementation cannot be left solely to a few Silicon Valley corporations and their distributors: the emergent scale and implications of AI's applications make it imperative for policymakers in government to take notice.[2]

To better appreciate how AI's propagation has affected people's lives, consider the following examples. Every time someone visits a website, a network of algorithms running in the background processes the data of the user's online behavior: analyzing the browsing history and collapsing thousands of data points into an intelligent guess on which products would entice the user, all in order to decide which advertisements to display. From the intelligent keyboards on smartphones, which reduce a user's typing burden by anticipating the next words, to the voice-activated assistants in tablets and desktops, capable of following voice commands and completing follow-on tasks, the machines in users' immediate personal space have become far more intelligent than is commonly realized.

While there is a substantive body of literature on recent advances in AI and the resulting implications for jobs, skills, and society at large, few analyses have

examined the specific impact of AI on India's emerging economy. The Indian government is aggressively trying to increase human capital on a national scale, with a specific emphasis on its younger population through the Skill India initiative, while seeking to attract global manufacturing to India via its Make in India program. The other part of this modernizing triad is the Digital India initiative: a determined push to expand digital access nationwide. AI will have a direct impact on each of these flagship initiatives of Prime Minister Narendra Modi in the immediate future, making it all the more urgent for policymakers in India to take seriously AI's potential for national strategies and to be on the front line in developing AI technology.

AI technology has enormous potential to shape India's economic and national security future.

For India to maximally benefit from the AI revolution, it must adopt a deliberate policy to drive AI innovation, adaptation, and proliferation in sectors beyond merely consumer goods and information technology (IT) services. AI's rapid diffusion begets unique opportunities and challenges for India: establishing policies specifically designed for the short term and tailoring the Modi policies to incorporate and emphasize AI, as well as for the medium and long terms, will allow India to appreciate the technology's full potential. While India has undoubtedly been a benefactor of AI's rapid ascent, AI has yet to capture the imagination of the country's policymakers. In foregoing the opportunity to inaugurate national strategies concerning AI, India runs the risk of falling behind the United States and China. AI technology has enormous potential to shape India's economic and national security future; in the absence of a specific policy regime, however, India will find it difficult to realize the full power of AI while potentially falling prey to the detrimental effects of AI proliferation.

Moore's Law Spurs Recent Advances in AI

Before addressing the challenges and opportunities for India, it is critical to understand what has changed in the field of AI in recent times. The answer lies in what experts of the industry refer to as Moore's law and its broader impact on the computing space. This maxim states that the number of transistors in computer chips doubles every two years, meaning that computer power increases and the relative cost of computer hardware decreases at an exponential pace; the fruit of this relation—faster and cheaper computing power coupled with the diminishing size of computer hardware—has had a decisive impact on AI. The faster, smaller, and less expensive hardware available today enables the computer algorithms of the 1990s to solve far more complex problems than they were previously capable of handling. In this, the profit of Moore's law is obvious: by merely saving time and space, without even accounting for

significant improvement to software, computers can perform far more complex operations, creating a multitude of possibilities for advancements in AI.[3]

This efficiency is compounded by the proliferation of cloud technology: with the massive computing power available on demand via cloud-based computing platforms, again at a much lower cost, it is possible today to store and analyze expansive sets of data, potentially processing billions of data points in a matter of seconds. This allows for several layers of computers working in tandem to interpret information in these large data sets, recognizing patterns, discerning behavior, and making intelligent decisions. It has also allowed these computers to learn from their own past performances so that they can improve their capabilities in recognition and deduction.

The combined impact of superior hardware, cloud-based on-demand computing, and the increasingly ubiquitous analysis of big data has spurred significant improvement in the performance of machine learning—the ability of computers to learn without being explicitly programmed.[4] A variety of techniques such as supervised learning, unsupervised learning, and reinforcement learning have become increasingly common in commercial business, leading to consumer applications like Apple's Siri and Microsoft's Cortana.[5] It has also resulted in the viability of complex architectures known as deep learning or deep neural networks, paving the way for recent innovations like the AlphaGo computer.[6]

Internet Boosts AI Proliferation in India, but Steep Barriers Persist

Given the technology's rapid propagation in the commercial space, it would not be an exaggeration to suggest that the Indian consumer is far ahead of the Indian state in the use of AI. Thanks to the increasingly digital economy, fueled by improving education and globalization, the Indian consumer is unknowingly the country's biggest beneficiary of recent advances in AI. From utilizing various applications powered by artificial intelligence to using a range of online services such as Amazon Marketplace and Netflix that learn from consumers' online behavior to make intelligent product and service recommendations, consumers are readily engaged with the proliferation of AI in India, whether they appreciate it or not. Policymakers, meanwhile, lag behind, not exploiting AI for national security, public services, or other priorities.

While technologies such as genetically modified foods have prompted sharp public debates in India and drawn the attention of bureaucrats and political parties, AI-powered applications and services have quietly become widely available to Indian consumers thanks to the global digital economy and relatively few barriers to their ensconcing without much political consideration. Indian academics, public researchers, labs, and entrepreneurs face a different challenge than the corporations that dominate the space, however, as the infrastructure necessary for an AI revolution in India has been neglected by policymakers.

For example, cloud-computing infrastructure, capable of storing the vast amounts of data and possessing the massive amount of computing power required by AI, largely resides in servers beyond India's borders. From the Amazon Web Services (AWS) elastic cloud to the Google machine learning infrastructure, almost all of the online tools that have made AI accessible to the entrepreneurial community rely on infrastructure that exists outside of India. A 2016 study on the entire data center market in India estimates its value at $2 billion—a stark contrast to the $17 billion of Amazon's last-quarter global revenue from AWS.[7] It is only recently that cloud infrastructure providers have made efforts to invest in this technology in India: Microsoft has pledged to invest in three data centers for its Azure cloud infrastructure, and Amazon has promised to locate some of its AWS infrastructure in India in 2016.[8] This delay in investment has had an injurious effect on the Indian business community, as a number of Indian start-ups have incorporated themselves outside of India in no small part due to easier access to cutting-edge technology and infrastructure with the simplicity of doing business abroad offering a further incentive for corporate flight.[9]

Furthermore, the absence of a large native-install base of on-demand cloud-computing infrastructure in India puts most recent advances in AI out of the reach of government-funded research labs. Perhaps even more damagingly, many industries cannot risk storing their data outside of India and being accessed by algorithms over which the Indian government has little direct control. This constraint makes adoption of AI technology all the more expensive for the private sector. Without this critical infrastructure, India has struggled, and will continue to struggle, to reach the inflexion point in AI efficiency and productivity, all while consumer adoption marches ahead.

AI Research—China Steals a March

While lack of physical infrastructure is certainly a major impediment, India's AI development also suffers from the paucity of the necessary cultural infrastructure, which is key for recent advances from lab to marketplace in AI.

Fostering a culture of innovation and a commitment to research and, most important, nurturing an ecosystem beyond the four walls of the organization are all common to Google's DeepMind, IBM's Watson, and Baidu's Institute of Deep Learning, the most successful AI projects of the past half decade. While it must come as no surprise that Google, IBM, Microsoft, Facebook, and other global technology giants have invested significantly over the decades in machine intelligence, it is the story of Baidu that holds the most pertinent lessons for India.[10] This Chinese search engine has built formidable AI capabilities that could serve as both an inspiration and a model for India's own domestic companies.[11]

The story of AI at Baidu is the story of Andrew Ng. Ng is an associate professor at Stanford University who teaches a popular course on machine learning

(also available via Coursera, an online learning platform that he co-founded).[12] More relevantly, Ng was hired away from the Google Brain project in 2014, in no small part to install similar capabilities for Baidu. With Ng on board, Baidu invested heavily in physical infrastructure and in a nimble enabling environment in which multifaceted AI research could prosper.

Baidu is investing in deep speech for voice-based searches that leverage speech recognition. This intelligence is being built for understanding and interpreting queries in Mandarin rather than English; as such, Baidu is constructing a uniquely Chinese platform, independent of that which is used ubiquitously in the English-speaking West. This version of AI thus offers the possibility for the development of AI in India, as an example of AI technology successfully developed in and employed by a non-Western nation.

Baidu's advancements in AI are the result of ambitious investments in homegrown infrastructure, an institutionalized approach to research, and an innovation ecosystem that extends from Beijing to Silicon Valley.

Baidu's advancements in AI are the result of ambitious investments in homegrown infrastructure, an institutionalized approach to research, and an innovation ecosystem that extends from Beijing to Silicon Valley—all strategies that to some extent are lacking in India. Baidu's innovation will repay and further this investment in infrastructure, as it has released code of its deep-learning algorithm for speech recognition to the general public. As other organizations improve the technology and appreciate the benefit of public access, China's homegrown infrastructure and research and development ecosystem will advance further.

Baidu's investment in AI research exposes the relative backwardness of India's technological infrastructure: China has recognized the importance of bridging the gap between the lab and the market while nurturing a research and innovation ecosystem unbounded by national borders and corporate firewalls. India, in contrast, boasts neither the material nor the cultural institutions required for such innovation.

Prime Minister Narendra Modi often challenges Indian IT entrepreneurs, asking when India will give birth to the next Google or Microsoft. But until India attains the infrastructure omnipresent in the United States, and increasingly existent in China, the deep-learning capabilities necessary to address the vast linguistic diversity across India using machine intelligence may prove elusive. Consequently, the policy focus in India should be not only on the creation of an Indian Google but also on the organizational structure, facilities, and customs required for these technological leviathans.

Funding AI Research—Global Lessons for India

Research in AI across the globe has a long history of public funding with periodic cycles of hope and despair. In the past two decades, however, with the advent of the Internet economy, there has been a significant shift toward private-sector funding. An interesting aspect of this shift has been the attrition

from academia to industry as major Internet companies have drawn talent from universities to put private-sector AI efforts on a fast track. Indian policymakers intent on spurring AI research will have to look closely at global lessons from both models.

Much of the early support for and investment in AI in the United States came from the government, which was seeking to use machine intelligence to autotranslate Russian documents during the Cold War. These efforts hit a wall in the 1960s after achieving only meager progress; as a result, AI underwent its first winter—reduced funding—after $20 million in public spending dried up, halting research and stalling scientific impetus.[13] Meanwhile, in the United Kingdom, AI research was virtually abandoned in the 1970s after the Lighthill report, which was commissioned by the British Science Research Council to evaluate AI progress, cited its utter failure to achieve its grandiose objectives.[14] Similar funding cuts by the Defense Advanced Research Projects Agency (DARPA) in the United States prolonged AI's long winter. The 1970s failure of a speech recognition project at Carnegie Mellon University prompted DARPA to cancel a multimillion-dollar research grant, a particularly prominent setback. Similar debacles occurred in the 1980s and 1990s as LISP, an early family of computer programming languages, failed to take off and expert systems gained little commercial traction.

Yet AI research ultimately not only survived but also thrived, as AI applications found their way into embedded systems within other applications as well as standalone technologies. Much of the public funding for AI in the past decade has come from DARPA's Cyber Grand Challenge, a competition with prize money, and the European Union's EU-FP7 technology funding program.[15] The BRAIN Initiative, created in 2013, is a ten-year, multibillion-dollar fund for AI research in the United States, while the EU's Human Brain Project envisages spending 1 billion euros on AI over the next decade.[16]

The commercial advancements in AI in recent years have come primarily from the private sector, with the labs of Google, IBM, Microsoft, and Facebook proving to be especially fruitful. While the United States and Western Europe enjoy innovation clusters in Silicon Valley, Boston, and London, as well as a network of elite research universities driving private research, India, bestowed with none of these resources, must turn to the public funding approaches employed elsewhere in the world.

For example, South Korea recently announced an $840 million public-private partnership spanning six corporations to drive AI research and innovation,[17] all in the name of building a "platform for intelligence information society," as reported by KBS World Radio. This follows South Korean investment in projects such as Exobrain, a brainstorming tool, and DeepView, a molecular graphics program, alongside AI initiatives by Samsung, LG, and AI-based news provider KakaoTalk.[18] Even Australia, a Western nation in all but its geography, has directly invested in AI rather than relied on its considerably

advanced private sector to spur growth. The country's intellectual property rights agency pioneered the integration and innovation of AI in the public sector in its trial of and collaboration with IBM's Watson.[19] Public funding for AI has even transcended borders: the Malaysian government's strategic investment fund recently invested in an augmented reality/visual recognition start-up founded by an Indian-origin entrepreneur with offices in Silicon Valley.[20]

Privately funded research in AI in India is relatively novel in comparison.[21] Infosys, for example, recently announced its decision to support AI research efforts at the Indraprastha Institute of Information Technology, Delhi, apart from its commitment to open-source AI research efforts led by SpaceX founder and Tesla Motors CEO Elon Musk (and others) under the OpenAI Project.[22] With Arya.ai and just a handful of other Indian AI start-ups beginning to make their mark domestically,[23] India has yet to see its private sector make a significant global impact in AI.[24]

India has yet to see its private sector make a significant global impact in AI.

In recent years, the Indian Institutes of Technology (IIT) have largely carried out publicly funded research in artificial intelligence. The most comprehensive insight into the state of AI research in India comes from a 2012 paper for *AI Magazine* by Deepak Khemani, a professor of computer science at IIT Madras.[25] According to Khemani, AI research in India has been limited to a handful of passionate researchers and thus has significantly lagged behind the United States and Europe. Listing the few labs dedicated to AI at various public institutions of higher education and research, Khemani underscores the limited focus of AI research in India on what he calls societal needs. He argues that because of India's vast linguistic diversity, most research effort has been applied toward machine translation, natural language, and text- and speech-related applications.

Indian AI research in defense is housed under the Defense Research and Development Organization, specifically within the Center for Artificial Intelligence and Robotics (CAIR).[26] CAIR lists artificial neural networks, computer vision, and situational awareness as its areas of primary focus, mentioning two other products that are under development for network-centric operations and decisionmaking using a vast knowledge base of battlefield tactics data. CAIR also lists civilian applications such as experimental robots, including a machine that can play chess by leveraging an AI-based decision engine.

The debate in China over public funding for AI should be of great interest to India, for the funding's audacity and the scale of its ambition, and even more so for its model of publicly funded AI research. In what appears to be a Hollywood script on "pre-crime"—a term coined by a science fiction writer to describe a police focus on yet-to-be-committed crimes—the Chinese are leveraging predictive policing to profile individual behavior and to predict crimes and terrorist incidents (or as the Chinese prefer to call them, security events).[27] The most definitive insight into the scale of Chinese ambition comes

from recent comments by Pan Jianwei, a professor of quantum physics at the University of Science and Technology of China.[28] Speaking in Beijing during the National People's Congress, he called for China's adoption of a DARPA-style defense research funding model that prioritized engagement with the private sector and universities. The joint research by Alibaba and Pan's team in quantum computing and artificial intelligence, a program designed to replicate and surpass Google's DeepMind, is but one example motivating this drive toward commercial and university engagement.[29]

The antiquated model of defense research employed by India has repeatedly struggled to bring research ideas from the lab into the real world, with several critical projects that have been pending for decades. Other developing nations have utilized public funding to advance their domestic AI regimes. But India's defense and other discretionary research apparatuses have yet to take a quantum leap to match China's progress, fueled by the lessons of the DARPA experience.

Understanding AI's Impact on Indian Jobs

One unintended consequence of AlphaGo's victory over Lee Sedol is the fear taking root within the popular discourse over job losses to technology.[30] Reported numbers in early 2016 from the Indian information technology industry hint as to how AI-influenced automation is beginning to make an impact on Indian jobs: a recent interview with Tata Consultancy Services' CEO on the projected decline in hiring by major Indian IT companies attributes much of this anticipated downturn to automation, with software replacing employees even as enterprises see greater use of bots and robots.[31]

The IT services sector is not the first to see a deleterious impact from AI on jobs. Manufacturing was arguably the first to bear the consequences of what has been dubbed the second machine age or fourth industrial revolution. While India dreams of its own manufacturing revolution through Modi's Make in India program, it is important for policymakers to closely examine how the advent of industrial robots and their impact on manufacturing transformed companies in other developing nations. The case of Foxconn, one of the world's largest contract manufacturers for electronics, is thus instructive.

It is important for policymakers to closely examine how the advent of industrial robots and their impact on manufacturing transformed companies in other developing nations.

The rise of Foxconn is a reflection of China's precipitous ascent as the manufacturing locus for much of the world's consumer electronics, as the company has twelve factories in nine Chinese cities. In 2015, Foxconn made news when its CEO predicted that 70 percent of all manufacturing in Foxconn's assembly lines would be automated with robots displacing humans.[32] Some backpedaling later, his estimate was scaled down to 30 percent.[33] The numbers, however, were irrelevant: this radical transformation in manufacturing, automation through the use of industrial robots, signaled an epochal revolution.

Foxconn is among the top owners of robotics patents filed with the United States Patent and Trademark Office and produces thousands of industrial robots a year that in aggregate are capable of performing more than ten types of manufacturing tasks.[34] This is expected to have a significant impact on the workforce: as many as 60,000 workers have been displaced by robots in one Foxconn factory alone in the Kunshan region of China.[35] China is projected to have more installed industrial robots by the end of 2016 than any other country, with more than 30 robots for every 10,000 industrial workers. If China were to increase that density, employment would be further damaged.

China may well be the last major world economy in which mass job creation from manufacturing was possible, with robots for manufacturing automation becoming increasingly ubiquitous.[36] While the dire predictions of a robot takeover of manufacturing have not come to pass, the reality of automation is that manufacturing is unlikely to create jobs at the scale that it did in the past. Quoting the U.S. Bureau of Labor Statistics in 2013 on future employment projections, Darrell M. West of the Brookings Center for Technology Innovation highlights how jobs will decline over the next decade in manufacturing and information technology among other sectors. "One business leader I know had 500 workers for his $100 million business and now has the same size workforce even though the company has grown to $250 million in revenues," West wrote. "He did this by automating certain functions and using robots and advanced manufacturing techniques to operate the firm."[37]

Despite the perils to manufacturing, the nascent debate on the skills and jobs crisis of the AI revolution until now has focused disproportionately on the importance of Indian IT companies building AI-related skills and capabilities and on the need to train Indian IT workers with AI skills to service the next wave of enterprise IT automation in the West. In a repeat of Y2K e-commerce and Web 2.0 waves, the needs and demands of the global IT services market have mistakenly driven much of the focus of Indian IT on AI and machine learning.[38] Writing in *Pacific Standard* magazine, Frank Levy, professor emeritus at the Massachusetts Institute of Technology, places jobs dislocation on account of artificial intelligence in perspective when he dismisses the dire projections of a robot takeover.[39] According to Levy, the greatest area of concern for policymakers ought to be the impact of AI on jobs in the middle-skill category—assembly line workers, clerical workers, and the like. Levy also warns of significant dislocation caused by automation leading to a reduction in an individual's potential for upward mobility. Expanding further on what it will take to survive this mass dislocation, he emphasizes that a good education will be critical to acquire the necessary skills and to be competitive in this evolved labor market.

In their book *The Second Machine Age*, Erik Brynjolfsson and Andrew McAfee make several specific policy recommendations for coping with the job crisis likely to be spurred by AI.[40] One such key recommendation of particular importance for India, given Modi's Startup India initiative, is the need to

"restart startups." Brynjolfsson and McAfee view creative destruction inherent within the start-up economy as the best bet for experimenting with the new jobs and industries that can thrive in an AI-driven economy. Expanding on the concept of a "peer economy," in which algorithms and machines collaborate with humans to create economic value, the authors provide examples of start-ups such as TaskRabbit and Airbnb that contrive previously nonexistent economic opportunities for ordinary people with spare time and assets, thus creating economically productive work.

Though this volatility in the middle-skill labor market represents perhaps the most profound danger to the new economy, there exists a more subtle danger within the AI community itself. According to the *Economist*, Uber recently hired 40 of the 140 staff engineers at the National Robotics Engineering Center at Carnegie Mellon University to work on a self-driving car.[41] Examples of flight to the private sector are rife in the United States, as universities continue to struggle to retain talent, especially for researchers and academics studying the in-demand machine learning. The result, according to the *Economist*, is that AI talent is increasingly and disproportionately concentrated in a few private corporations that have the ability to pay the most.[42] If all AI talent is thus concentrated, then AI research may not be as diversified, and research priorities could be narrowly focused on a few commercial ideas while many areas of social and national importance could suffer for want of talent. This could slow the proliferation and innovation potential of AI, which could displace any number of middle-skill jobs but then fail to create new jobs to replace them. This stagnation would represent the worst possible outcome, as market forces could allow a wealthy AI elite to prosper while stifling competition from start-ups, nonprofits, and institutions focused on the progression, rather than profitability, of AI technology. Were this to happen, the panacea offered by Brynjolfsson and McAfee would be rendered irrelevant.

Though India does not suffer from a brain drain of top-quality AI talent from university research labs to the industry, it must be wary to avoid this concentration of intellectual energy.

However, this crisis is hardly inevitable: one effort to open up AI for the common public good is the nonprofit OpenAI, a group led by Tesla's Musk. India's Infosys is among many major IT companies that have joined the OpenAI effort. Though India does not suffer from a brain drain of top-quality AI talent from university research labs to the industry, it must be wary to avoid this concentration of intellectual energy.

Encouraging Skill Development for Future Jobs

Not everyone is as sanguine as Erik Brynjolfsson and Andrew McAfee are regarding the coming AI revolution. The recent victory of AlphaGo over the world champion in Go has prompted fears of the threat posed by intelligent machines that are capable of superhuman tasks.[43] The direst warning comes from noted physicist Stephen Hawking, who apocalyptically predicts the end of the human

race with the development of "full artificial intelligence."[44] While this doomsday scenario is many decades away, in a recent *Guardian* profile about the DARPA Cyber Grand Challenge on autonomous hacking to detect and fix cybersecurity vulnerabilities, Konstantinos Karagiannis, a technology executive with a telecom firm, raised the red flag of the likelihood of criminals acquiring AI capabilities to mount sophisticated cyberattacks.[45] In another recent column in the *Guardian*, Jason Millar, a robotics ethics expert, while calling for an interdisciplinary effort to frame ethics standards for robots and AI, raised the possibility of the "loss of human control" likely to result from advances in AI.[46]

Anxiety over AI mirrors fears over advances in cloning, genetically modified foods, nuclear weaponry, and other areas of technology.

Anxiety over AI mirrors fears over advances in cloning, genetically modified foods, nuclear weaponry, and other areas of technology. The politics over technologies such as AI in India could see a repeat of the kind of protests against the use of computers in machine tools in the 1960s and banking in the 1980s. Indeed, India would have lost out in connecting 1 billion people through cellular telephony if fears over radiation had delayed the rollout of mobile services in the 1990s. Today, through the Aadhaar universal identification program, India's poor have an identity and, overcoming barriers and bypassing middlemen, are empowered to claim their entitlements. This would not have been possible if fears over privacy had stalled the rollout of Aadhaar over the past seven years. Likewise, the AI opportunity in India today should be seen in a similar light as a means for India to not merely leapfrog into the future but also as a means to empower its population.

In the provocative *Humans Need Not Apply*, Jerry Kaplan, a U.S. computer scientist and futurist, explores this and several other questions while attempting to paint a picture of an apocalyptic future and what it might look like if and when machines take over.[47] (Fortunately, that futuristic scenario need not be a concern yet because Kaplan rules out this happening anytime in the immediate future.) Of greater interest is his devotion of specific attention to the immediate challenges and opportunities at hand with the evolution of machine intelligence.

Specifically, Kaplan raises two issues that should be of interest to policymakers in India. The first is the education system and the second involves skills and jobs. While discussing the likely impact of AI on labor markets, Kaplan poses the radical question: Is the current system of sequential education and work outdated, and does it require an overhaul? According to Kaplan, the sequential system—as he put it, "First you go to school, then you get a job"[48]—made sense in an economic era in which an individual was expected to hold a single job throughout a career. In an economic environment that is rapidly changing, however, the nature of available jobs constantly shifts, with skills becoming valuable and obsolete in a matter of years. This is not a reality of a distant future—it is already happening, including in India. The number of engineering

graduates who perform jobs distinctly unrelated to the engineering they specialized in and the number of arts and commerce graduates in business processing outsourcing jobs is proof of this phenomenon. A few generations back, a mid-career change in profession was unthinkable; today, several career changes within a span of two decades are routine as individuals reinvent themselves to adapt and stay relevant in a fast-changing economic environment.

Jerry Kaplan's question on the sequential system of education becomes all the more important in an AI-powered economy in which the rate of obsolescence of skills would be faster than ever before. As India looks to formulate a new education policy, it is imperative that the government thinks ahead and asks hard questions about the continued relevance of a sequential education system, both now and in the future. Going beyond the debate on formative education, Kaplan takes on the issue of a mismatch between the supply of and demand for skills in the job market in order to highlight how inapt the current system of apprenticeship has become. One solution Kaplan proposes in his book is particularly promising: what he calls a "job mortgage" as a new type of financial instrument through which employers, vocational schools, and colleges would have an incentive to collaborate in a new way. In this proposed job mortgage market, Kaplan attempts to use free market mechanisms to match current skills acquisition to future job opportunities. He proposes to accomplish this by compelling employers to commit to an intent to employ an individual in the future if that person commits to acquire a specific set of skills over a certain time frame. This market mechanism, in other words, would ensure that training institutions tailor their programs around specific skills that employers are willing to commit to and students are willing to develop. According to Kaplan, skills acquisition would then occur in an environment in which an individual has the assurance that these skills are valued by the market, backed by an "intent to employ" guarantee.

The challenges and opportunities of AI's advances call for urgent response.

India will have to experiment with the kind of innovative instruments that Kaplan proposes if it is to prepare itself for the challenges from a machine intelligence–driven economy in the near future. Although this is hardly the only means by which human capital can be efficiently developed, this formalization of the economic search model is a promising, if admittedly rough, approach to the challenges AI poses to the future Indian economy.

AI Policy Road Map for India

Although this regime of programs may appear to be a projection ill-suited for the present, the challenges and opportunities of AI's advances call for urgent responses, because every one of Prime Minister Narendra Modi's flagship programs is likely to be directly affected by AI's inexorable advance. Automation

in manufacturing will likely lead Make in India to fall short of its promise of job creation; the pace of obsolescence in skills will likely cause Skill India to miss its goal by a wide margin; overblown fears of so-called digital colonization in recent months may become realized if consumers' data and online behavior in a Digital India are at the mercy of AI-powered platforms and services located entirely outside India. The AI-powered global economy could wipe away India's demographic dividend if its new education policies fail to anticipate and adequately adapt for the future. Furthermore, India could face a strategic and near permanent disadvantage in the balance of power against China, given Beijing's considerable capability and ambition in its defense-driven artificial intelligence research agenda.

Short-Term Actions

In the immediate term, policymakers in India should make AI a critical component of the prime minister's flagship programs. As an example, within the Make in India program, India must create special incentives for manufacturers, such as relaxing regulations and lowering trade barriers, so they:

- Invest in automation research in India by building research labs and design studios in India
- Create regional innovation clusters, districts, and corridors by building strong linkages around manufacturing automation and robotics between universities and start-ups in India
- Make India a global hub for machine intelligence–based innovation in manufacturing

Similarly, the Skill India initiative should be reworked with the twin objectives of being resilient to skills obsolescence through market-based instruments that tie together the employers, the training institutes, and the students as well as paying special attention to new skills needed to survive in an AI-led economy in the future.

Digital India must be reconfigured to establish cloud infrastructure inside India on a fast-track basis: the limited capacity as it stands today is a critical infrastructure gap and a national security risk. As a part of the Digital India initiative, New Delhi must create specific incentives for building large-scale data centers in India, ideally in partnership with the state governments. The government must identify specific regions in India that are geographically suited for building massive data centers, with an assured supply of power and other critical public infrastructure required for such facilities, and promote these as preferred destinations for investment in cloud infrastructure in India under the Digital India scheme.

The spirit of Startup India, that of creative destruction rather than protectionism, must be allowed to prevail. Recent regulatory decisions in India across

cities and states bearing down on taxi aggregators Uber and Ola are regressive and counterproductive; these well-intentioned regulations must be eliminated if AI is to achieve its full potential. Unless the government and domestic industry allow the marketplace to experiment with untested business models enabled by the so-called peer economy, it is unlikely that the economy will create jobs resilient to and in an AI-driven economy. It is imperative to recognize that improving start-ups' ease of doing business is not merely a regulatory measure to incubate or liquidate a business, but also a free market initiative to allow new and efficient business models to thrive.

For the first time in the history of democratic India, the formulation of a new education policy has been undertaken with a nationwide process of consultation and crowdsourcing. The massive task of analyzing the inputs received and devising the new policy is challenged both by the volume of the inputs and the complexity of myriad issues across India. Formulating the new education policy must not, however, be based solely on inputs received that are affected by current challenges, constraints, and limitations. The National Education Policy must take a long-term view of the skills economy, evaluate the continued relevance of the current system of sequential education, and make radical recommendations on alternative models of education that would be better suited to the economy of the future. Piloting and experimenting in such new models of education must commence in the immediate future before the rapid obsolescence of the current system begins.

Medium-Term Applications of AI in the Public Sector

The government should identify public-sector applications in India where current advances in AI could make a significant impact toward building skills and capabilities domestic applications of AI. For example, New Delhi could:

- Apply AI-based techniques to recognize patterns and learn about tax evasion behavior, in an effort to mine public databases to detect tax fraud and money obtained illegally with the goal of minimizing tax evasion and maximizing tax revenue
- Use AI to scan records, recognize patterns of fraud, and correlate subsidy claims with other consumer data to help detect leakages of subsidies and to learn and better target direct benefits to citizens through interventions most relevant to them
- Develop natural language–processing capabilities to automate multilingual communication and interactivity across a whole range of government services and interfaces, for example, crowdsourcing via MyGov.in, voice calls, automated helplines, and chatbots for the most routine citizen-government interactions

- Use AI-based training and teaching software in various skilling and educational applications

In each of these areas, the government should collaborate with the private sector and university research labs to leverage existing technologies effectively and to rapidly create new technologies to address specific and well-defined problems.

Long-Term Strategy

India must view machine intelligence as a critical element of its national security strategy. At a time when AI is being viewed as a key component of foreign policy between the United States and Japan, with similar proposals of treatment being floated in India, the Indian government must formulate a national strategy on emerging technology trends with long-term strategic consequences.[49] India must seriously evaluate the DARPA model of defense research in conjunction with private sector and university collaboration in order to create dual-purpose technologies with a scope large enough to allow for development of civilian technology applications. Specifically, the Cyber Grand Challenge model of DARPA needs to be examined for its successful incentivization of academia and the private sector.

India must view machine intelligence as a critical element of its national security strategy.

The proposed National Intelligence Grid (NATGRID) platform, which would link citizen databases, might be a good pilot candidate for creating a machine intelligence–based platform with both national security and civilian benefits and should thus be taken up on a mission mode. Another possibility is Aadhaar, a platforms-based approach to governance founded on massive data sets, which builds on the possibility articulated in *Rebooting India*.[50] Authors Nandan Nilekani and Viral Shah enumerate five data-based platforms that could address a broad range of governance challenges. Expanding this list to ten public databases encompassing state and local governments would enable a robust machine intelligence architecture capable of plugging leakages in subsidies, better targeting benefits, and expanding the tax base.

Conclusion

From NATGRID to Aadhaar, machine intelligence–powered platforms can become a strategic instrument of governance in India across a wide range of public services. These platforms are not without their challenges: a machine intelligence–powered approach to governance will require robust digital privacy laws and a code of ethics on limits to using AI. However, the range of AI's possibilities is so vast that the full spectrum of its opportunities is difficult to fully comprehend. While India may be late to wake up to the AI revolution,

Indians of many hues—consumers, technocrats, researchers, and entrepreneurs—are already participants in this revolution with many of Indian origin driving and influencing research in the United States and elsewhere. A clarion call from the prime minister to all of them, to come together and help build an AI ecosystem in India, will go a long way for India to not merely catch up to but to take a quantum leap into the AI-driven future.

Recommended Reading on AI for Indian Policymakers

Indian policymakers looking to understand the recent advances in artificial intelligence will likely find themselves challenged by the technical complexity in most available literature. The *MIT Technology Review*'s "AI Takes Off" is a good starting point to get a quick overview of the latest developments.[51] For a more popular overview of these developments, *Rolling Stone*'s two-part series makes for an interesting read.[52] For those looking for an in-depth understanding of why AI is at an inflexion point, the recent lecture on "The Future of AI" by Professor Rich Sutton from the University of Alberta, Canada, is an excellent resource.[53] For the more technically inclined, the O'Reilly Media publication on "The Future of Machine Intelligence" gives a comprehensive view of ongoing research in AI from a practitioner's perspective.[54] Aspiring students of AI who are looking for a solid grounding in machine learning should make use of the entire lecture series by Andrew Ng on machine learning at Stanford University.[55] Finally, policymakers ought to read the books by Jerry Kaplan and by Erik Brynjolfsson and Andrew McAfee to obtain a comprehensive perspective on all of the policy issues in the global debate on AI.[56]

Glossary

Algorithm – A set or sequence of step-by-step operations that need to be carried out to perform a calculation, to process a set of data, or to test a logical statement.

AlphaGo – A computer designed to play the ancient Chinese board game Go, developed by Google DeepMind.

Analytics – A discipline of mathematics and computer science that concerns itself with statistical analysis of data to discover and interpret patterns.

Artificial intelligence (AI, or machine intelligence) – The ability of machines or software to behave in a humanlike, intelligent manner.

AI winter – A period of reduced funding to artificial intelligence research during the 1970s and 1980s on account of AI progress failing to match expectations.

Autonomous vehicles – Vehicles that are capable of sensing within their environment and navigating terrains without human input. Driverless (self-driving) cars from Google and Tesla are the most recent popular examples.

Autonomous weapons – Weapons that can engage and perform without human input.

Big Data – Large and complex sets of data that cannot be efficiently processed using traditional techniques. These typically entail millions or billions of records, information that does not have a predefined structure, and complex formats. Examples include website, image, and video banks on the Internet.

Chatbots – Computer programs that can engage in humanlike conversation. Facebook's announcement auguring AI-based chatbots available as a service via its Messenger Platform is a recent example.

Cloud computing – Internet-based remote computing infrastructure that can be made available on demand via sharing of computing resources and data storage.

Cognitive computing – Technology that can mimic the human brain in how it senses and responds to stimuli.

Data center – A facility where computing infrastructure and data storage infrastructure are housed.

Deep Blue – A chess-playing computer developed by IBM that made history in the 1990s by defeating the then world champion of chess, Garry Kasparov.

Deep learning – A branch of machine learning that involves algorithms that analyze data through multiple layers of complex processing. Each layer's output becomes the input to the next layer to carry out pattern analysis and classification and to establish hierarchical relationships for both supervised and unsupervised learning.

Deep neural networks – A kind of deep-learning architecture based on artificial neural networks that uses multiple layers of processing units that can model complex nonlinear relationships.

Defense Advanced Research Projects Agency (DARPA) – An organ of the United States Department of Defense. It is credited with having funded much of the early stage research in artificial intelligence, most notably through its Cyber Grand Challenge projects.

Expert systems – An artificial intelligence system that can make decisions like a human expert in a given subject matter area or domain. Developed in 1970s and 1980s, these systems lost their name in the 1990s as technology evolved beyond standalone AI systems.

Google DeepMind – A British artificial intelligence company based in London that was acquired by Google in 2014. It recently made news for its computer AlphaGo, which made history in defeating the world champion in Go.

Image recognition – Algorithms used in computer vision and image processing in order to determine whether an image contains a specific object, feature, or activity.

Industry 4.0 or Manufacturing 4.0 – A collection of manufacturing technologies powered by data analytics and automation based on AI. This is occasionally referred to as the "fourth industrial revolution."

LISP – A family of computer programming languages that were favored by artificial intelligence programmers during the technology's early development in the 1950s. LISP-based machines fell out of favor due to performance issues during the 1970s as AI research expanded.

Machine learning – A subfield of artificial intelligence. It refers to the ability of computers to learn without being explicitly programmed. It entails algorithms that can recognize patterns in data, learn from these patterns, and subsequently make predictions based on these data.

Moore's law – Named for Gordon E. Moore, co-founder of Intel, the rule observes that computing power on a chip (integrated circuit) doubles every two years. The recent success of AI techniques such as deep learning and deep neural networks is attributed to the improved performance of algorithms thanks to the increase in computing power, a phenomenon most easily understood through Moore's law.

Natural language processing – A field of study that deals with human-computer interactions through natural languages or human languages by interpreting or understanding natural language inputs to a computer and by synthesizing outputs from a computer in natural languages.

Neural networks – Artificial neural networks are an architecture of computing used in machine learning. Inspired by the organization and processing mechanisms of biological neural networks, artificial neural networks have been used in speech recognition, image recognition, and other problem areas involving machine learning.

OpenAI – Open-source nonprofit AI company founded by Elon Musk and others with the intent of developing AI in a way that is beneficial, rather than detrimental, to humanity. OpenAI has received support from Indian IT company Infosys.[57]

Reinforcement learning – A subcategory of machine learning in which a computer learns from its past performance to make improvements to its future behavior. It has been applied to robot control and elevator scheduling among other areas. Using deep neural networks, an application of reinforced learning, AlphaGo beat world champion Lee Sedol in the board game Go.

Speech recognition – The ability of computers to recognize and interpret spoken language. Speech recognition is today the most ubiquitous application of artificial intelligence with applications ranging from automated customer care agents reachable via toll-free phone lines to voice-activated personal assistants on tablets and smartphones.

Supervised learning – A machine learning technique in which a computer is trained to recognize patterns and interpret data using a large training data set. Speech recognition is a one of the well-known applications of supervised learning: these algorithms are trained with a large data set of prerecorded audio samples to cover the vocabulary and spoken accents within a given language. Once trained, the system is able to recognize spoken words and phrases in that language in order to make interpretations based on the context.

Unsupervised learning – A machine learning technique that analyzes large data sets by determining patterns and relationships without any prior knowledge or instruction.

Virtual reality – Computer-simulated reality that mimics an environment by creating artificial sensor experiences such as vision/sight, hearing, touch, smell.

Watson – Technology created by IBM based on AI/cognitive-computing that acts like a human and can answer questions posed in a natural language. Watson made news in 2011 when it defeated humans on the quiz show *Jeopardy!* [58]

Notes

1 Google DeepMind, "AlphaGo," accessed August 1, 2016, https://deepmind.com/alpha-go.

2 Nanette Byrnes, "AI Hits the Mainstream," in *AI Takes Off*, ed. Nanette Byrnes (Cambridge, MA: MIT Technology Review, 2016), https://www.technologyreview.com/s/600986/ai-hits-the-mainstream/.

3 Rich Sutton, "The Future of AI," YouTube video, 1:22:40, presented as part of the University of British Columbia's Department of Computer Science's Distinguished Lecture Series, February 25, 2016, posted by "UBCCPSC," February 26, 2016, https://www.youtube.com/watch?v=pD-FWetbvN8.

4 David Beyer, *The Future of Machine Intelligence: Perspectives From Leading Practitioners* (Sebastopol, CA: O'Reilly Media, 2016), http://www.oreilly.com/data/free/files/future-of-machine-intelligence.pdf.

5 Tom Simonite, "The Missing Link of Artificial Intelligence," *MIT Technology Review*, February 18, 2016, https://www.technologyreview.com/s/600819/the-missing-link-of-artificial-intelligence/.

6 John Brownlee, "Neural Networks: What Are They, and Why Is the Tech Industry Obsessed With Them?," *Co.Design* (blog), *Fast Company*, April 7, 2016, http://www.fastcodesign.com/3058665/neural-networks-explained.

7 Mohd Ujaley, "Upswing Continues in Indian Data Center Market," *Express Computer*, February 27, 2016, http://computer.financialexpress.com/magazine/upswing-continues-in-indian-data-center-market/16157/.

8 Yevgeniy Sverdlik, "Microsoft Launches Three Cloud Data Centers in India," Data Center Knowledge, September 29, 2015, http://www.datacenterknowledge.com/archives/2015/09/29/microsoft-launches-three-cloud-data-centers-india/; and "Amazon to Open AWS Infrastructure Region in India in 2016," *India Today*, last updated on June 30, 2015, http://indiatoday.intoday.in/technology/story/amazon-to-open-aws-infrastructure-region-in-india-in-2016/1/448079.html.

9 Peerzada Abrar, "India's Tech Stars Shifting to Silicon Valley," *Hindu*, last updated June 30, 2015, http://www.thehindu.com/news/national/red-tape-forces-top-indian-entrepreneurs-to-shift-overseas/article7367731.ece; and "Stop Startups From Setting Base Abroad, iSPIRT Tells Govt," *Business Standard,* February 6, 2016, http://www.business-standard.com/article/companies/stop-startups-from-setting-base-abroad-ispirt-tells-govt-116020500414_1.html.

10 Will Knight, "How AI Is Feeding China's Internet Dragon," *MIT Technology Review,* March 2016, http://www.technologyreview.pk/how-ai-is-feeding-chinas-internet-dragon/.

11 Nanette Byrnes, ed., *AI Takes Off* (Cambridge, MA: MIT Technology Review, 2016), https://www.technologyreview.com/business-report/ai-takes-off/.

12 Andrew Ng, "Lecture 1 | Machine Learning (Stanford)," YouTube video, 1:08:39, lecture for the Stanford Computer Science Department, posted by "Stanford," July 22, 2008, http://www.youtube.com/view_play_list?p=A89DCFA6ADACE599/.

13 Daniel Faggella, "AI Is So Hot, We've Forgotten All About the AI Winter," *Singularity Weblog*, August 25, 2015, https://www.singularityweblog.com/ai-is-so-hot-weve-forgotten-all-about-the-ai-winter/.

14 The report was named for its author. See James Lighthill, "Artificial Intelligence: A General Survey," Computing at Chilton, accessed August 1, 2016, http://www.chilton-computing.org.uk/inf/literature/reports/lighthill_report/contents.htm.

15 Defense Advanced Research Projects Agency, "Cyber Grand Challenge (CGC) Final Event," accessed August 1, 2016, http://www.darpa.mil/news-events/cyber-grandchallenge; and European Union, "Understanding the Seventh Framework Programme," last modified August 25, 2015, https://ec.europa.eu/research/fp7/index_en.cfm?pg=understanding.

16 DARPA is one of the partners in the Brain Research Through Advancing Innovative Neurotechnologies Initiative, whose goal is to revolutionize understanding of the human brain. Defense Advanced Research Projects Agency, "DARPA and the Brain Initiative," accessed August 1, 2016, http://www.darpa.mil/program/our-research/darpa-and-the-brain-initiative; Human Brain Project, "Human Brain Project," accessed March 30, 2016, https://www.humanbrainproject.eu/.

17 Lee Chi-dong, "Gov't to Invest 1 Tln Won in Artificial Intelligence," Yonhap News Agency, March 17, 2016, http://english.yonhapnews.co.kr/news/2016/03/17/0200000000AEN20160317003751320.html?3be7ff58.

18 "The Status of Artificial Intelligence Projects in Korea," KBS World Radio, March 14, 2016, http://world.kbs.co.kr/english/program/program_economyplus_detail.htm?No=5420; Lee Ji-yoon, "Korea Considers Control Tower for AI Push," *Korea Herald*, March 14, 2016, http://www.koreaherald.com/view.php?ud=20160314000839.

19 Stephen Easton, "How and Why the Public Sector Must Make Friends With Artificial Intelligence," *Mandarin*, July 24, 2015, http://www.themandarin.com.au/45780-public-sector-must-adopt-artificial-intelligence-innovation-seminar/?pgnc=1.

20 Arjun Kharpal, "App That Recognizes Objects With AI Raises $54M," CNBC, March 2, 2016, http://www.cnbc.com/2016/03/02/blippar-app-that-recognizes-objects-with-artificial-intelligence-raises-54-million.html.

21 Pankaj Mishra, "Indian IT's AI: More Artificial Less Intelligent," *Economic Times*, July 13, 2015, http://tech.economictimes.indiatimes.com/news/corporate/indian-its-ai-more-artificial-less-intelligent/48048313.

22 Indraprastha Institute of Information Technology, Delhi, "Infosys Grants IIITD 24 Crore for Center for Artificial Intelligence," April 2016, https://www.iiitd.ac.in/webstory/cai; and "Artificial Intelligence Is for All: Sikka," *Business Standard*, December 15, 2015, http://www.business-standard.com/article/companies/artificial-intelligence-is-for-all-sikka-115121400757_1.html.

23 See the Arya.ai website, http://www.arya.ai/.

24 Krithika Krishnamurthy, "Startups Build on Artificial Intelligence, Save Costs and Drive Efficiency," *Economic Times*, December 1, 2015, http://articles.economictimes.indiatimes.com/2015-12-01/news/68688432_1_technology-startups-tools-google.

25 Deepak Khemani, "A Perspective on AI Research in India," *AI Magazine* 33, no. 1 (Spring 2012): 96–98.

26 Defense Research and Development Organization, "Center for Artificial Intelligence & Robotics (CAIR)," Ministry of Defense, last modified November 2, 2012, http://www.drdo.gov.in/drdo/labs/CAIR/English/index.jsp?pg=homebody.jsp.

27 Sean Gallagher, "China Is Building a Big Data Platform for 'Precrime,'" *Ars Technica*, March 8, 2016, http://arstechnica.com/information-technology/2016/03/china-is-building-a-big-data-plaform-for-precrime/.

28 Stephen Chen, "China Should Copy America's DARPA to Create Cutting-Edge Technology – For Military and Civilian Use, Says Top Scientist," *South China Morning Post*, March 10, 2016, http://www.scmp.com/news/china/policies-politics/article/1922893/china-should-copy-americas-darpa-create-cutting-edge.

29 Cheng Yingqi, "Jack Ma Aims to Make Big Dream Come True," *China Daily*, September 1, 2015, http://www.chinadaily.com.cn/china/2015-09/01/content_21764766.htm.

30 Girish Shahane, "Getting Watsoned: Why Robots Could Spell Disaster for India's IT Sector," Scroll.in, March 16, 2016, http://scroll.in/article/805184/getting-watsoned-why-a-computer-chess-programme-could-spell-disaster-for-indias-it-sector.

31 Shivani Shinde Nadhe and Sheetal Agarwal, "IT Biggies Hit the Slow Lane on Hiring," *Business Standard*, April 22, 2016, http://www.business-standard.com/article/companies/it-biggies-hit-the-slow-lane-on-hiring-116042200069_1.html.

32 Enru Lin, "Foxconn Targets 70% Automation in 3 Years: Gou," *China Post*, February 27, 2015,http://www.chinapost.com.tw/taiwan/business/2015/02/27/429772/Foxconn-targets.htm.

33 Michael Kan, "Foxconn's CEO Backpedals on Robot Takeover at Factories," *Computerworld*, June 26, 2015, http://www.computerworld.com/article/2941272/emerging-technology/foxconns-ceo-backpedals-on-robot-takeover-at-factories.html.

34 Lorraine Luk, "Foxconn's Robot Army Yet to Prove Match for Humans," *Digits* (blog), *Wall Street Journal*, May 5, 2015, http://blogs.wsj.com/digits/2015/05/05/foxconns-robot-army-yet-to-prove-match-for-humans/.

35 Jane Wakefield, "Foxconn Replaces '60,000 Factory Workers With Robots,'" BBC News, May 25, 2016, http://www.bbc.com/news/technology-36376966.

36 Timothy Aeppel and Mark Magnier, "China's Hunger for Robots Marks a Significant Shift," *Wall Street Journal*, July 5, 2015, http://www.wsj.com/articles/chinas-hunger-for-robots-marks-significant-shift-1436118228.

37 Darrell M. West, "What Happens if Robots Take the Jobs? The Impact of Emerging Technologies on Employment and Public Policy," Brookings Institution, October 26, 2015, http://www.brookings.edu/research/papers/2015/10/26-robots-emerging-technologies-public-policy-west; and U.S. Bureau of Labor Statistics, "Employment Projections — 2012-2022," news release, December 19, 2013, http://www.bls.gov/news.release/archives/ecopro_12192013.pdf.

38 Mishra, "Indian IT's AI."

39 Frank Levy, "The Future of Work: The Three Dimensions of Artificial Intelligence," *Pacific Standard*, October 2, 2015, https://psmag.com/the-future-of-work-the-three-dimensions-of-artificial-intelligence-86b18f760ae7#.tbztwxz36.

40 Erik Brynjolfsson and Andrew McAfee, *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies* (New York: W. W. Norton and Company, 2014).

41 "Artificial Intelligence: Million-Dollar Babies," *Economist*, April, 2, 2016, http://www.economist.com/news/business/21695908-silicon-valley-fights-talent-universities-struggle-hold-their?zid=291&ah=906e69ad01d2ee51960100b7fa502595.

42 Ibid.

43 Jason Millar, "The Momentous Advance in Artificial Intelligence Demands a New Set of Ethics," *Guardian*, March 12, 2016, http://www.theguardian.com/commentisfree/2016/mar/13/artificial-intelligence-robots-ethics-human-control.

44 Rory Cellan-Jones, "Stephen Hawking Warns Artificial Intelligence Could End Mankind," BBC News, December 2, 2014, http://www.bbc.com/news/technology-30290540.

45 Olivia Solon, "These Engineers Are Developing Artificial Intelligence Hackers," *Guardian*, March 3, 2016, https://www.theguardian.com/technology/2016/mar/03/artificial-intelligence-hackers-security-autonomous-learning.

46 Millar, "The Momentous Advance."

47 Jerry Kaplan, *Humans Need Not Apply: A Guide to Wealth and Work in the Age of Artificial Intelligence* (New Haven, CT: Yale University Press, 2015).

48 Ibid., 13.

49 James Schoff, "The Japan-U.S. Alliance and Robotics Diplomacy," *Japan Times*, March 17, 2016, http://www.japantimes.co.jp/opinion/2016/03/17/commentary/japan-commentary/japan-u-s-alliance-robotics-diplomacy/#.V0s8NPkrLrd; and C. Raja Mohan, "Raja-Mandala: India, US and Artificial Intelligence," *Indian Express*, April 12, 2016, http://indianexpress.com/article/opinion/columns/raja-mandala-india-us-and-artificial-intelligence/.

50 Nandan Nilekani and Viral Shah, *Rebooting India: Realizing a Billion Aspirations* (New Delhi: Penguin Books India, 2015).

51 Byrnes, *AI Takes Off*.

52 Jeff Goodell, "Inside the Artificial Intelligence Revolution: A Special Report, Pt. 1," *Rolling Stone*, February 29, 2016, http://www.rollingstone.com/culture/features/inside-the-artificial-intelligence-revolution-a-special-report-pt-1-20160229. See also Goodell, "Inside the Artificial Intelligence Revolution: A Special Report, Pt. 2," *Rolling Stone*, March 9, 2016, http://www.rollingstone.com/culture/features/inside-the-artificial-intelligence-revolution-a-special-report-pt-2-20160309.

53 Sutton, "The Future of AI."

54 Beyer, *The Future of Machine Intelligence*.

55 Andrew Ng, "Lecture Collection | Machine Learning," YouTube video playlist, lecture series on machine learning presented at Stanford University's Computer Science Department, posted by "Stanford," last updated September 25, 2014, https://www.youtube.com/playlist?list=PLA89DCFA6ADACE599.

56 Kaplan, *Humans Need Not Apply*; and Brynjolfsson and McAfee, *The Second Machine Age*.

57 John Markoff, "Artificial-Intelligence Research Center Is Founded by Silicon Valley Investors," *New York Times*, December 11, 2015, http://www.nytimes.com/2015/12/12/science/artificial-intelligence-research-center-is-founded-by-silicon-valley-investors.html.

58 Steve Lohr, "The Promise of Artificial Intelligence Unfolds in Small Steps," *New York Times*, February 28, 2016, http://www.nytimes.com/2016/02/29/technology/the-promise-of-artificial-intelligence-unfolds-in-small-steps.html.

Carnegie India

Carnegie India, the sixth international center of the Carnegie Endowment for International Peace, opened in New Delhi in April 2016. As with Carnegie's centers in Beijing, Beirut, Brussels, Moscow, and Washington, Carnegie India is led by local experts who collaborate extensively with colleagues around the world. The center's research and programmatic focus includes the political economy of reform, foreign and security policy, and the role of innovation and technology in India's internal transformation and international relations. It will build on decades of regional scholarship from across Carnegie's programs while placing special emphasis on developing a cadre of up-and-coming Indian scholars.

The **Carnegie Endowment for International Peace** is a unique global network of policy research centers in Russia, China, Europe, the Middle East, India, and the United States. Our mission, dating back more than a century, is to advance the cause of peace through analysis and development of fresh policy ideas and direct engagement and collaboration with decisionmakers in government, business, and civil society. Working together, our centers bring the inestimable benefit of multiple national viewpoints to bilateral, regional, and global issues.

BEIJING BEIRUT BRUSSELS MOSCOW NEW DELHI WASHINGTON

THE GLOBAL THINK TANK

CARNEGIE INDIA

CarnegieIndia.org