

MCKINSEY GLOBAL INSTITUTE

ARTIFICIAL INTELLIGENCE THE NEXT DIGITAL FRONTIER?

DISCUSSION PAPER
JUNE 2017

Jacques Bughin | Brussels
Eric Hazan | Paris
Sree Ramaswamy | Washington, DC
Michael Chui | San Francisco
Tera Allas | London
Peter Dahlström | London
Nicolaus Henke | London
Monica Trench | London

MCKINSEY GLOBAL INSTITUTE

Since its founding in 1990, the McKinsey Global Institute (MGI) has sought to develop a deeper understanding of the evolving global economy. As the business and economics research arm of McKinsey & Company, MGI aims to provide leaders in the commercial, public, and social sectors with the facts and insights on which to base management and policy decisions. The Lauder Institute at the University of Pennsylvania ranked MGI the world's number-one private-sector think tank in its 2016 Global Think Tank Index for the second consecutive year.

MGI research combines the disciplines of economics and management, employing the analytical tools of economics with the insights of business leaders. Our “micro-to-macro” methodology examines microeconomic industry trends to better understand the broad macroeconomic forces affecting business strategy and public policy. MGI's in-depth reports have covered more than 20 countries and 30 industries. Current research focuses on six themes: productivity and growth, natural resources, labor markets, the evolution of global financial markets, the economic impact of technology and innovation, and urbanization.

Recent reports have assessed the economic benefits of tackling gender inequality, a new era of global competition, Chinese innovation, and digital globalization. MGI is led by four McKinsey and Company senior partners: Jacques Bughin, James Manyika, Jonathan Woetzel, and Frank Mattern, MGI's chairman. Michael Chui, Susan Lund, Anu Madgavkar, Sree Ramaswamy, and Jaana Remes serve as MGI partners. Project teams are led by the MGI partners and a group of senior fellows and include consultants from McKinsey offices around the world. These teams draw on McKinsey's global network of partners and industry and management experts. Input is provided by the MGI Council, which coleads projects and provides guidance; members are Andres Cadena, Sandrine Devillard, Richard Dobbs, Katy George, Rajat Gupta, Eric Hazan, Eric Labaye, Acha Leke, Scott Nyquist, Gary Pinkus, Oliver Tonby, and Eckart Windhagen. In addition, leading economists, including Nobel laureates, act as research advisers.

McKinsey & Company is a member of the Partnership on AI, a collection of companies and non-profits that have committed to sharing best practices and communicating openly about the benefits and risks of artificial intelligence research. The partners of McKinsey fund MGI's research; it is not commissioned by any business, government, or other institution.

For further information about MGI and to download reports, please visit

www.mckinsey.com/mgi.

CONTENTS

Preface

In brief

1. Artificial intelligence is getting ready for business, but are businesses ready for AI?

Page 6

2. Artificial intelligence promises to boost profits and transform industries

Page 20

3. Businesses, developers, and governments need to act now to realize AI's full potential

Page 31

Appendix A: Five case studies

Page 42

Retail

Page 42

Electric utility

Page 47

Manufacturing

Page 53

Health care

Page 58

Education

Page 65

Appendix B: Technical appendix

Page 70

Bibliography

Page 73

PREFACE

In this independent discussion paper, we examine investment in artificial intelligence (AI), describe how it is being deployed by companies that have started to use these technologies across sectors, and aim to explore its potential to become a major business disrupter. To do this, we looked at AI through several lenses. We analyzed the total investment landscape bringing together both investment of large corporations and funding from venture capital and private equity funds. We also reviewed the portfolio plays of major internet companies, the dynamics in AI ecosystems from Shenzhen to New York, and a wide range of case studies. As part of our primary research, we surveyed more than 3,000 senior executives on the use of AI technologies, their companies' prospects for further deployment, and AI's impact on markets, governments, and individuals. This report also leverages the resources of McKinsey Analytics, a global practice that helps clients achieve better performance through data. The research was conducted jointly with Digital McKinsey, a global practice that designs and implements digital transformations.

In addition to identifying a gap between AI investment and commercial application, which is typical of early technology development curves, we found that the new generation of AI applications is based on the foundation of digitization. Leading sectors in digital tend to be leading sectors in AI, and these are predicted to drive growth. We also found that AI has the potential to accelerate shifts in market share, revenue, and profit pools—all hallmarks of digitally disrupted sectors. This report leverages two MGI analyses of digitization, *Digital America: A tale of the haves and have-mores*, published in December 2015, and *Digital Europe: Pushing the frontier, capturing the benefits*, published in June 2016. These reports introduced the McKinsey Global Institute (MGI) Industry Digitization Index, which combines dozens of indicators to provide a comprehensive picture of where and how companies are building digital assets, expanding digital usage, and creating a more digital workforce. This report also builds on MGI's work on advanced analytics, *The age of analytics: Competing in a data-driven world*, published in December 2016, and on automation, *A future that works: Automation, employment, and productivity*, published in January 2017, as well as *Artificial intelligence: Implications for China*, published in April 2017; and an April 2017 Digital McKinsey report, *Smartening up with artificial intelligence (AI): What's in it for Germany and its industrial sector?*

This latest research has been led by Jacques Bughin, an MGI senior partner based in Brussels; Eric Hazan, a member of the MGI Council and a McKinsey senior partner based in Paris; Sree Ramaswamy, an MGI partner based in Washington, DC; Michael Chui, an MGI partner based in San Francisco; Tera Allas, an MGI visiting fellow in London; Peter Dahlström, a McKinsey senior partner in London; Nicolaus Henke, a McKinsey senior partner based in London; and Monica Trench, a McKinsey consultant based in London. The project team comprised Mathilde Castet, François Allain des Beauvais, Lindsay Macdonald, Oleg Pynda, Dariusz Smolen, and Jordan Ward. Sincere thanks go to Timothy Beacom, April Cheng, Paul-Louis Caylar, and Hugo Weber. We would also like to thank MGI senior editor Mark A. Stein; Matt Cooke, MGI director of external communications; MGI visual graphics specialist Marisa Carder, designer Margo Shimasaki, and infographic designers Richard Johnson and Jason Leder; MGI editorial production manager Julie Philpot; and Deadra Henderson, MGI manager of personnel and administration.

This report builds on a considerable body of expertise within MGI and McKinsey. We particularly want to acknowledge Tamim Saleh and Brian McCarthy from McKinsey Analytics, Louise Herring and Casper Louw for contributing to the retail sector case study, Arnout de Pee and Mike Munroe for contributing to the electric utilities case study, Richard Kelly for contributing to the manufacturing case study, Martin Markus and Sri Velamoor for contributing to the health care case study, and Jake Bryant, Mike Munroe, and Jimmy Sarakatsannis for contributing to the education case study. We would also like to thank all previous MGI teams that produced reports on digitization, automation, big data and analytics, the internet of things, and online talent platforms.

Our research was also enriched by insights from Eric Goubault and Jesse Read from Ecole Polytechnique.

This report contributes to MGI's mission to help business and policy leaders understand the forces transforming the global economy, identify strategic imperatives, and prepare for the next wave of growth. As with all MGI research, this work is independent and has not been commissioned or sponsored in any way by any business, government, or other institution. We welcome your comments on the research at MGI@mckinsey.com.

Jacques Bughin

Director, McKinsey Global Institute
Brussels

Eric Hazan

Co-leader, Digital McKinsey
Senior Partner, McKinsey and Company
Member, McKinsey Global Institute Council
Paris

James Manyika

Director, McKinsey Global Institute
San Francisco

Jonathan Woetzel

Director, McKinsey Global Institute
Shanghai

June 2017

IN BRIEF

ARTIFICIAL INTELLIGENCE: THE NEXT DIGITAL FRONTIER?

Artificial intelligence is poised to unleash the next wave of digital disruption, and companies should prepare for it now. We already see real-life benefits for a few early-adopting firms, making it more urgent than ever for others to accelerate their digital transformations. Our findings focus on five AI technology systems: robotics and autonomous vehicles, computer vision, language, virtual agents, and machine learning, which includes deep learning and underpins many recent advances in the other AI technologies.

- AI investment is growing fast, dominated by digital giants such as Google and Baidu. Globally, we estimate tech giants spent \$20 billion to \$30 billion on AI in 2016, with 90 percent of this spent on R&D and deployment, and 10 percent on AI acquisitions. VC and PE financing, grants, and seed investments also grew rapidly, albeit from a small base, to a combined total of \$6 billion to \$9 billion. Machine learning, as an enabling technology, received the largest share of both internal and external investment.
- AI adoption outside of the tech sector is at an early, often experimental stage. Few firms have deployed it at scale. In our survey of 3,000 AI-aware C-level executives, across 10 countries and 14 sectors, only 20 percent said they currently use any AI-related technology at scale or in a core part of their businesses. Many firms say they are uncertain of the business case or return on investment. A review of more than 160 use cases shows that AI was deployed commercially in only 12 percent of cases.
- Adoption patterns illustrate a growing gap between digitized early AI adopters and others. Sectors at the top of MGI's Industry Digitization Index, such as high tech and telecom or financial services, are also leading adopters of AI. They also have the most aggressive AI investment intentions. Leaders' adoption is both broad and deep: using multiple technologies across multiple functions, with deployment at the core of their business. Automakers use AI to develop self-driving vehicles and improve operations, for example, while financial services firms are more likely to use it in customer experience-related functions.
- Early evidence suggests that AI can deliver real value to serious adopters and can be a powerful force for disruption. In our survey, early AI adopters that combine strong digital capability with proactive strategies have higher profit margins and expect the performance gap with other firms to widen in the future. Our case studies in retail, electric utilities, manufacturing, health care, and education highlight AI's potential to improve forecasting and sourcing, optimize and automate operations, develop targeted marketing and pricing, and enhance the user experience.
- AI's dependence on a digital foundation and the fact that it often must be trained on unique data mean that there are no shortcuts for firms. Companies cannot delay advancing their digital journeys, including AI. Early adopters are already creating competitive advantages, and the gap with the laggards looks set to grow. A successful program requires firms to address many elements of a digital and analytics transformation: identify the business case, set up the right data ecosystem, build or buy appropriate AI tools, and adapt workflow processes, capabilities, and culture. In particular, our survey shows that leadership from the top, management and technical capabilities, and seamless data access are key enablers.
- AI promises benefits, but also poses urgent challenges that cut across firms, developers, government, and workers. The workforce needs to be reskilled to exploit AI rather than compete with it; cities and countries serious about establishing themselves as a global hub for AI development will need to join the global competition to attract AI talent and investment; and progress will need to be made on the ethical, legal and regulatory challenges that could otherwise hold back AI.



ARTIFICIAL INTELLIGENCE

The next digital frontier?

The current AI wave is poised to finally break through

Investment in AI is growing at a high rate, but adoption in 2017 remains low

In 2016, companies invested **\$26B to \$39B** in artificial intelligence

TECH GIANTS
\$20B to \$30B


STARTUPS
\$6B to \$9B

3x External investment growth since 2013

20% of AI-aware firms say they are **adopters**


- 3+ technologies
- 2 technologies
- 1 technology

41% of firms say they are **uncertain** about the benefits of AI

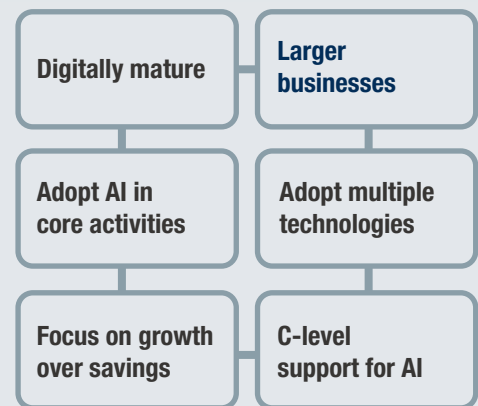


How companies are adopting AI

AI adoption is greatest in sectors that are already **strong digital adopters**



Six characteristics of early AI adopters



Four areas across the value chain where AI can create value



Five elements of successful AI transformations



1. ARTIFICIAL INTELLIGENCE IS GETTING READY FOR BUSINESS, BUT ARE BUSINESSES READY FOR AI?

Claims about the promise and peril of artificial intelligence are abundant, and growing. AI, which enables machines to exhibit human-like cognition, can drive our cars or steal our privacy, stoke corporate productivity or empower corporate spies. It can relieve workers of repetitive or dangerous tasks or strip them of their livelihoods. Twice as many articles mentioned AI in 2016 as in 2015, and nearly four times as many as in 2014.¹ Expectations are high.

AI has been here before. Its history abounds with booms and busts, extravagant promises and frustrating disappointments. Is it different this time? New analysis suggests yes: AI is finally starting to deliver real-life business benefits. The ingredients for a breakthrough are in place. Computer power is growing significantly, algorithms are becoming more sophisticated, and, perhaps most important of all, the world is generating vast quantities of the fuel that powers AI—data. Billions of gigabytes of it every day.

Companies at the digital frontier—online firms and digital natives such as Google and Baidu—are betting vast amounts of money on AI. We estimate between \$20 billion and \$30 billion in 2016, including significant M&A activity. Private investors are jumping in, too. We estimate that venture capitalists invested \$4 billion to \$5 billion in AI in 2016, and private equity firms invested \$1 billion to \$3 billion. That is more than three times as much as in 2013. An additional \$1 billion of investment came from grants and seed funding.

For now, though, most of the news is coming from the suppliers of AI technologies. And many new uses are only in the experimental phase. Few products are on the market or are likely to arrive there soon to drive immediate and widespread adoption. As a result, analysts remain divided as to the potential of AI: some have formed a rosy consensus about AI's potential while others remain cautious about its true economic benefit. This lack of agreement is visible in the large variance of current market forecasts, which range from \$644 million to \$126 billion by 2025.² Given the size of investment being poured into AI, the low estimate would indicate that we are witnessing another phase in a boom-and-bust cycle.

Our business experience with AI suggests that this bust scenario is unlikely. In order to provide a more informed view, we decided to perform our own research into how users are adopting AI technologies. Our research offers a snapshot of the current state of the rapidly changing AI industry, looking through the lenses of both suppliers and users to come up with a more robust view of the economic potential of AI and how it will unfold. To begin, we examine the investment landscape, including firms' internal investment in R&D and deployment, large corporate M&A, and funding from venture capital (VC) and private equity (PE) firms. We then look at the demand side, combining use case analyses and our AI adoption and use survey of C-level executives at more than 3,000 companies to understand how companies use AI technologies today, what is driving their adoption of AI, the barriers to further deployment, and the market, financial, and organizational impacts of AI. For further details on sources of our insights, see Box 1, "A multi-lens approach to understanding the AI story."

¹ Factiva.

² Tractica; Transparency Market Research.

AI generally refers to the ability of machines to exhibit human-like intelligence—for example, solving a problem without the use of hand-coded software containing detailed instructions. There are several ways to categorize AI technologies, but it is difficult to draft a list that is mutually exclusive and collectively exhaustive because people often mix and match several technologies to create solutions for individual problems. These creations sometimes are treated as independent technologies, sometimes as subgroups of other technologies, and sometimes as applications. Some frameworks group AI technologies by basic functionality, such as text, speech, or image recognition, and some group them by business applications such as commerce or cybersecurity.³

³ Gil Press, “Top 10 hot artificial intelligence (AI) technologies,” *Forbes.com*, January 23, 2017; “AI100: The artificial intelligence startups redefining industries,” *CBInsights.com*, January 11, 2017.

Box 1. A multi-lens approach to understanding the AI story

For the findings presented in this report, we drew on both primary and secondary research. We used six distinct sources of insight to assess the state of AI and its future potential. See Appendix B for further details.

AI adoption and use survey. We surveyed AI-aware C-level executives at 3,073 companies about how they are using digital technology and AI today, the drivers and barriers to further deployment, and the market, financial, and organizational impacts of AI. Our stratified sample covered 14 sectors of the economy, 10 countries across Europe, North America, and Asia, and companies with workforces ranging from fewer than 10 to more than 10,000. Responses were verified using algorithmic data cleansing techniques.

Use case review. We collated and reviewed over 160 use cases from both public sources and databases assembled for related MGI and Digital McKinsey research. Use cases were individually assessed to determine the extent to which they had achieved commercialization, and were classified according to their primary sector and business function of use.

Investment flows. To measure investment in the development and deployment of new AI technology, products, and services, we conducted an extensive review of publicly available information on both external funding sources (VC, PE, and M&A), as well as R&D and deployment activities internal to large corporations. Our assessment of external investment is based on global deal databases, scrutinizing deals for AI content, and classifying them by type of technology and geography.

Internal investment flows are based on an analysis of the top 35 companies investing in AI globally.

Investment hubs. We built a global picture of AI innovation activity, covering 10 countries and over 75 cities, and evaluating their local AI ecosystems. This included mapping AI investment flows from VC, PE, and M&A databases by city, scanning large AI-investing companies’ activities by geography, assessing the research outputs of and talent flows from universities, evaluating the broader business environment for entrepreneurs, and interviewing local investors.

Sector case studies. In five sectors, we conducted industry-expert interviews to understand the specific use cases today and promising applications for the future. External interviews were complemented by unique insights from across McKinsey sector practices and functions, including Digital McKinsey and McKinsey Analytics, as well as advanced-analytics firms that McKinsey has acquired or partnered with.

Previous MGI and McKinsey research. This report leverages other recent major research efforts by MGI and McKinsey Analytics, including databases of use cases. In particular, the findings in this report complement and build on the following research: *A future that works: Automation, employment, and productivity* (2017), *The age of analytics: Competing in a data-driven world* (2016), *Digital Europe: Pushing the frontier, capturing the benefits* (2016), and *Digital America: A tale of the haves and have-mores* (2015).

Trying to pin down the term more precisely is fraught for several reasons: AI covers a broad range of technologies and applications, some of which are merely extensions of earlier techniques and others that are wholly new. Also, there is no generally accepted theory of “intelligence,” and the definition of machine “intelligence” changes as people become accustomed to previous advances.⁴ Tesler’s theorem, attributed to the computer scientist Larry Tesler, asserts that “AI is whatever hasn’t been done yet.”⁵

The AI technologies we consider in this paper are what is called “narrow” AI, which performs one narrow task, as opposed to artificial general intelligence, or AGI, which seeks to be able to perform any intellectual task that a human can do. We focus on narrow AI because it has near-term business potential, while AGI has yet to arrive.⁶

In this report, we focus on the set of AI technology systems that solve business problems. We have categorized these into five technology systems that are key areas of AI development: robotics and autonomous vehicles, computer vision, language, virtual agents, and machine learning, which is based on algorithms that learn from data without relying on rules-based programming in order to draw conclusions or direct an action. Some are related to processing information from the external world, such as computer vision and language (including natural language processing, text analytics, speech recognition, and semantics technology); some are about learning from information, such as machine learning; and others are related to acting on information, such as robotics, autonomous vehicles, and virtual agents, which are computer programs that can converse with humans. Machine learning and a subfield called deep learning are at the heart of many recent advances in artificial intelligence applications and have attracted a lot of attention and a significant share of the financing that has been pouring into the AI universe—almost 60 percent of all investment from outside the industry in 2016.

ARTIFICIAL INTELLIGENCE’S ROLLER-COASTER RIDE TO TODAY

Artificial intelligence, as an idea, first appeared soon after humans developed the electronic digital computing that makes it possible. And, like digital technology, artificial intelligence, or AI, has ridden waves of hype and gloom—with one exception: AI has not yet experienced wide-scale commercial deployment (see Box 2, “Fits and starts: A history of artificial intelligence”).

That may be changing. Machines powered by AI can today perform many tasks—such as recognizing complex patterns, synthesizing information, drawing conclusions, and forecasting—that not long ago were assumed to require human cognition. And as AI’s capabilities have dramatically expanded, so has its utility in a growing number of fields. At the same time, it is worth remembering that machine learning has limitations. For example, because the systems are trained on specific data sets, they can be susceptible to bias; to avoid this, users must be sure to train them with comprehensive data sets. Nevertheless, we are seeing significant progress.

⁴ Marvin Minsky, “Steps toward artificial intelligence,” *Proceedings of the IRE*, volume 49, number 1, January 1961; Edward A. Feigenbaum, *The art of artificial intelligence: Themes and case studies of knowledge engineering*, Stanford University Computer Science Department report number STAN-CS-77-621, August 1977; Allen Newell, “Intellectual issues in the history of artificial intelligence,” in *The Study of Information: Interdisciplinary messages*, Fritz Machlup and Una Mansfield, eds., John Wiley and Sons, 1983.

⁵ Douglas R. Hofstadter, *Gödel, Escher, Bach: An eternal golden braid*, Basic Books, 1979. Hofstadter writes that he gave the theorem its name after Tesler expressed the idea to him firsthand. However, Tesler writes in his online CV that he actually said, “Intelligence is whatever machines haven’t done yet.”

⁶ William Vorhies, “Artificial general intelligence—the Holy Grail of AI,” DataScienceCentral.com, February 23, 2016.

Box 2. Fits and starts: A history of artificial intelligence

The idea of computer-based artificial intelligence dates to 1950, when Alan Turing proposed what has come to be called the Turing test: can a computer communicate well enough to persuade a human that it, too, is human?¹ A few months later, Princeton students built the first artificial neural network, using 300 vacuum tubes and a war-surplus gyropilot.²

The term “artificial intelligence” was coined in 1955, to describe the first academic conference on the subject, at Dartmouth College. That same year, researchers at the Carnegie Institute of Technology (now Carnegie Mellon University) produced the first AI program, Logic Theorist.³ Advances followed often through the 1950s: Marvin Lee Minsky founded the Artificial Intelligence Laboratory at MIT, while others worked on semantic networks for machine translation at Cambridge and self-learning software at IBM.⁴

Funding slumped in the 1970s as research backers, primarily the US government, tired of waiting for practical AI applications and cut appropriations for further work.⁵ The field was fallow for the better part of a decade.

University researchers’ development of “expert systems”—software programs that assess a set of facts using a database of expert knowledge and then offer solutions to problems—revived AI in the 1980s.⁶ Around this time, the first computer-controlled autonomous vehicles began to appear.⁷ But this burst of interest preceded another AI “winter.”

Interest in AI boomed again in the 21st century as advances in fields such as deep learning, underpinned by faster computers and more data, convinced investors and researchers that it was practical—and profitable—to put AI to work.⁸

¹ A. M. Turing, “Computing machinery and intelligence,” *Mind*, volume 49, number 236, October 1950.

² Jeremy Bernstein, “A.I.,” *The New Yorker*, December 14, 1981.

³ Leo Gugerty, “Newell and Simon’s Logic Theorist: Historical background and impact on cognitive modeling,” *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 50, issue 9, October 2006.

⁴ “The IBM 700 Series: Computing comes to business,” *IBM Icons of Progress*, March 24, 2011.

⁵ Michael Negnevitsky, *Artificial intelligence: A guide to intelligent systems*, Addison-Wesley, 2002.

⁶ Edward A. Feigenbaum, “Expert systems in the 1980s,” working paper, 1980.

⁷ Hans P. Moravec, “The Stanford Cart and the CMU Rover,” *Proceedings of the IEEE*, volume 71, issue 7, July 1983; Tom Vanderbilt, “Autonomous cars through the ages,” *Wired.com*, February 6, 2012.

⁸ Buchanan, Bruce G., “A (very) brief history of artificial intelligence,” *AI Magazine*, volume 26, number 4, Winter 2005.

These advances have allowed machine learning to be scaled up since 2000 and used to drive deep learning algorithms, among other things. The advances have been facilitated by the availability of large and diverse data sets, improved algorithms that find patterns in mountains of data, increased R&D financing, and powerful graphics processing units (GPUs), which have brought new levels of mathematical computing power. GPUs, which are specialized integrated circuits originally developed for video games, can process images 40 to 80 times faster than the fastest versions available in 2013. Advances in the speed of GPUs have enabled the training speed of deep learning systems to improve five- or sixfold in each of the last two years. More data—the world creates about 2.2 exabytes, or 2.2 billion gigabytes, of it every day—translates into more insights and higher accuracy because it exposes algorithms to more examples they can use to identify correct and reject incorrect answers. Machine learning systems enabled by these torrents of data have reduced computer error rates in some applications—for example, in image identification—to about the same as the rate for humans.

LED BY TECH GIANTS, AI INVESTMENT IS GROWING RAPIDLY, BUT COMMERCIAL ADOPTION IS LAGGING BEHIND


Tech giants and digital native companies such as Amazon, Apple, Baidu, and Google are investing billions of dollars in the various technologies known collectively as artificial intelligence. They see that the inputs needed to enable AI to finally live up to expectations—powerful computer hardware, increasingly sophisticated algorithmic models, and a vast and fast-growing inventory of data—are in place. Indeed, internal investment by large corporations dominates: we estimate that this amounted to \$18 billion to \$27 billion in

2016; external investment (from VCs, PE firms, M&A, grants, and seed funding) was around \$8 billion to \$12 billion (Exhibit 1).⁷

Exhibit 1

Technology giants dominate investment in AI

Investment in AI, 2016¹
\$ billion



1 Estimate of 2016 spend by corporations to develop and deploy AI-based products. Calculated for top 35 high tech and advanced manufacturing companies investing in AI. Estimate is based on the ratio of AI spend to total revenue calculated for a subset of the 35 companies.
 2 VC value is an estimate of VC investment in companies primarily focused on AI. PE value is an estimate of PE investment in AI-related companies. M&A value is an estimate of AI deals done by corporations. “Other” refers to grants and seed fund investments. Includes only disclosed data available in databases, and assumes that all registered deals were completed within the year of transaction. Compound annual growth rate values rounded.
 3 M&A and PE deals expressed by volume; VC deals expressed by value.

SOURCE: Capital IQ; Pitchbook; Dealogic; S&P; McKinsey Global Institute analysis

But for all the recent investment, the scope of AI deployment has been limited so far. That is partly due to the fact that one beneficiary of that investment, internal R&D, is largely focused on improving the firms’ own performance. But it is also true that there is only tepid demand for artificial intelligence applications for businesses, partly due to the relatively slow pace of digital and analytics transformation of the economy. Our survey of more than 3,000 businesses around the world found that many business leaders are uncertain about what exactly AI can do for them, where to obtain AI-powered applications, how to integrate them into their companies, and how to assess the return on an investment in the technology.

⁷ Internal investment includes research and development, talent acquisition, cooperation with scientific institutions, and joint ventures with other companies done by corporations. External investment includes mergers and acquisitions, private equity funding, venture capital financing, and seed funds and other early-stage investing. The estimates of external investment are based on data available in the Capital IQ, PitchBook, and Dealogic databases. Provided values are estimates of annual investment in AI, assuming that all registered deals were completed within the year of transaction. Internal investment is estimated based on the ratio of AI spend to revenue for the top 35 high tech and advanced manufacturing companies focused on AI technologies.

Most of the investment in AI has consisted of internal spending—R&D and deployment—by large, cash-rich digital native companies. What is the large corporate investment in AI focused on? Bigger companies, such as Apple, Baidu, and Google, are working on suites of technologies internally, but vary in the breadth and focus of their AI investment. Amazon is working on robotics and speech recognition, Salesforce on virtual agents and machine learning. BMW, Tesla, and Toyota are among the manufacturers making sizable commitments in robotics and machine learning for use in driverless cars. Toyota, for example, set aside \$1 billion to establish a new research institute devoted to AI for robotics and driverless vehicles.⁸ Industrial giants such as ABB, Bosch, GE, and Siemens also are investing internally, often in machine learning and robotics, seeking to develop specific technologies related to their core businesses. IBM has pledged to invest \$3 billion to make its Watson cognitive computing service a force in the internet of things.⁹ Baidu has invested \$1.5 billion in AI research over the last 2½ years. This is in addition to \$200 million it committed to a new in-house venture capital fund, Baidu Venture.¹⁰

At the same time, big tech companies have been actively buying AI startups, not just to acquire technology or clients but to secure qualified talent. The pool of true experts in the field is small, and Alibaba, Amazon, Facebook, Google, and other tech giants have hired many of them. Companies have adopted M&A as a way to sign up top talent, a practice known as “acqui-hiring,” for sums that typically work out to \$5 million to \$10 million per person. The shortage of talent and cost of acquiring it are underlined by a recent report that companies are seeking to fill 10,000 AI-related jobs and have budgeted more than \$650 million for salaries.¹¹

Overall, corporate M&A is the fastest-growing external source of funding for AI companies, increasing in terms of value at a compound annual growth rate of over 80 percent from 2013 to 2016, based on our estimates. Leading high tech companies and advanced manufacturers have closed more than 100 M&A deals since 2010. Google completed 24 transactions in that time, including eight in computer vision and seven in language processing. Apple, the second-most-active acquirer, has closed nine, split evenly among computer vision, machine learning, and language processing.

Companies are also expanding their search for talent abroad. Facebook, for instance, is opening an AI lab in Paris that will supplement similar facilities in New York and Silicon Valley—and make it easier for the company to recruit top researchers in Europe.¹² Google recently invested \$4.5 million in the Montreal Institute for Learning Algorithms, a research lab at the University of Montreal; Intel donated \$1.5 million to establish a machine learning and cybersecurity research center at Georgia Tech; and NVIDIA is working with the National Taiwan University to establish an AI laboratory in Taipei.¹³

⁸ Craig Trudell and Yuki Hagiwara, “Toyota starts \$1 billion center to develop cars that don’t crash,” Bloomberg.com, November 6, 2015.

⁹ “IBM invests to lead global internet of things market—shows accelerated client adoption,” IBM press release, October 3, 2006.

¹⁰ Phoenix Kwong, “Baidu launches \$200m venture capital unit focused on artificial intelligence,” *South China Morning Post*, September 13, 2016.

¹¹ “U.S. companies raising \$1 billion or more to fuel artificial intelligence (AI) development: Looking to staff 10,000+ openings, cites new Paysa research,” Paysa press release, April 18, 2017.

¹² Cade Metz, “Facebook opens a Paris lab as AI research goes global,” *Wired.com*, June 2, 2015.

¹³ Cade Metz, “Google opens Montreal AI lab to snag scarce global talent,” *Wired.com*, November 12, 2015; “Georgia Tech launches new research on the security of machine-learning systems,” Georgia Institute of Technology press release, October 31, 2016; “NVIDIA collaborates with Taipei Tech to establish Embedded GPU Joint Lab,” National Taipei University of Technology press release, September 4, 2014.

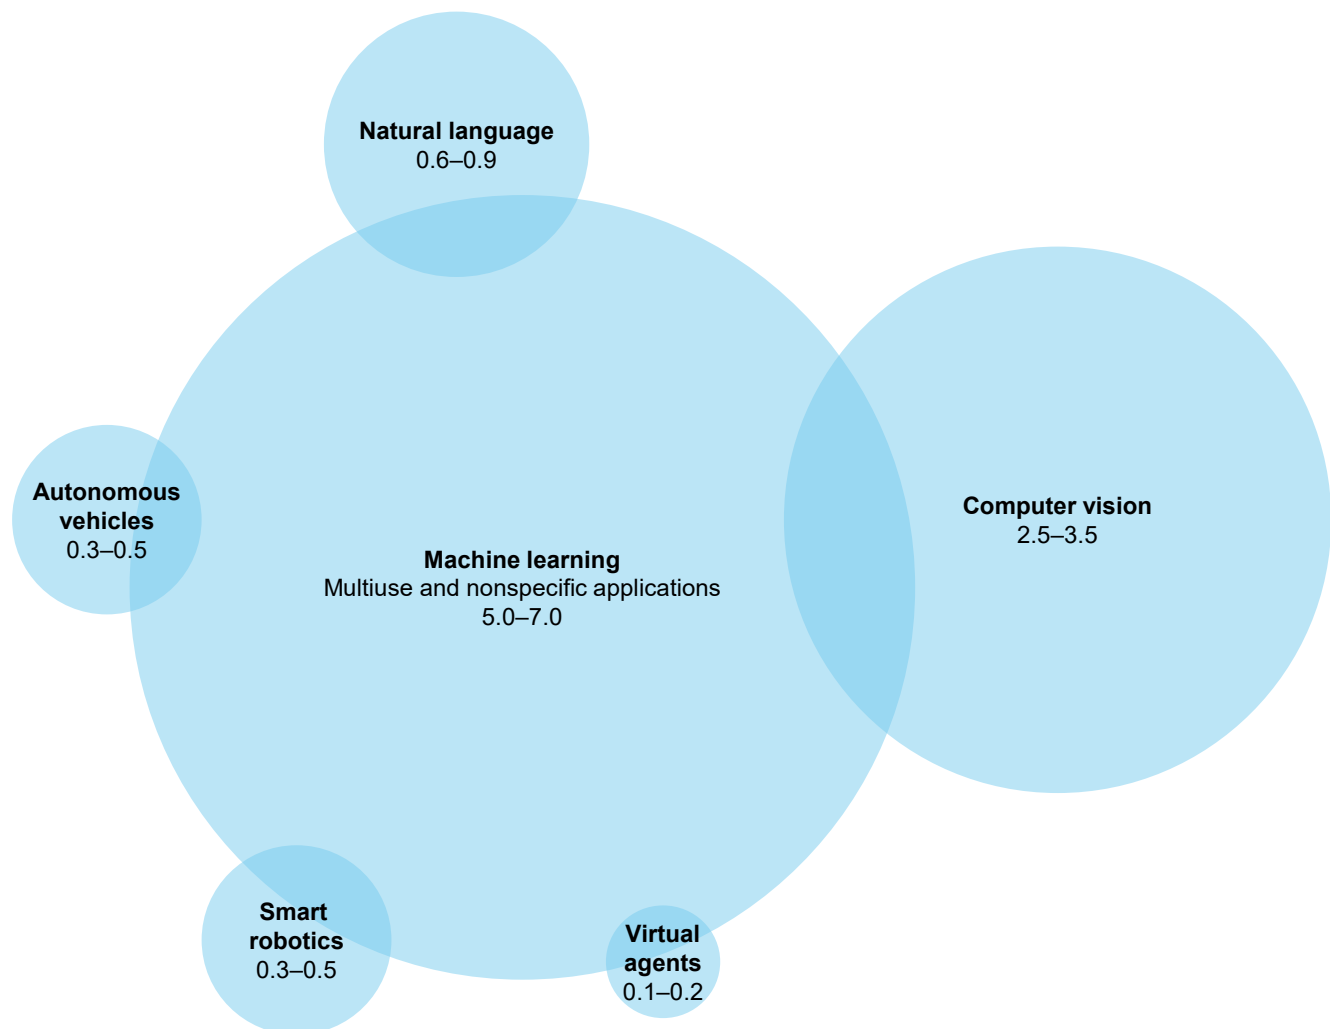
The buzz over AI has grown loud enough to encourage venture capital and private equity firms to step up their investment in AI. Other external investors, such as angel funds and seed incubators, also are active. We estimate total annual external investment was \$8 billion to \$12 billion in 2016.¹⁴

Machine learning attracted almost 60 percent of that investment, most likely because it is an enabler for so many other technologies and applications, such as robotics and speech recognition (Exhibit 2). In addition, investors are drawn to machine learning because, as has long been the case, it is quicker and easier to install new code than to rebuild a robot or other machine that runs the software. Corporate M&A in this area is also growing fast, with a compound annual growth rate of around 80 percent from 2013 through 2016.

Exhibit 2

Machine learning received the most investment, although boundaries between technologies are not clear-cut

External investment in AI-focused companies by technology category, 2016¹
\$ billion



¹ Estimates consist of annual VC investment in AI-focused companies, PE investment in AI-related companies, and M&A by corporations. Includes only disclosed data available in databases, and assumes that all registered deals were completed within the year of transaction.

SOURCE: Capital IQ; Pitchbook; Dealogic; McKinsey Global Institute analysis

¹⁴ Estimates of external investment in AI vary widely because measurement standards vary. For example, Venture Scanner puts total funding of AI-related startups in 2016 at \$2.5 billion, while Goldman Sachs estimates that the venture capital sector alone made \$13.7 billion of AI-related investment that year.

Investment in AI is still in the early stages and relatively small compared with the investment in the digital revolution. Artificial intelligence, for example, attracted 2 to 3 percent of all VC funding by value in 2016, while information technology in general soaked up 60 percent. AI also was a small fraction—1 to 3 percent—of all investment by PE firms in 2016.¹⁵ But AI investment is growing fast.

From 2013 through 2016, external investment in AI technologies had a compound annual growth rate of almost 40 percent. That compares with 30 percent from 2010 through 2013. Not only are deals getting bigger and more numerous, but they require fewer participants to complete the financing. This suggests that investors are growing more confident in the sector and may have a better understanding of the technology and its potential.

However, for the most part, investors are still waiting for their investments to pay off. Only 10 percent of startup companies that consider machine learning to be a core business say they generate revenue, according to PitchBook. Of those, only half report more than \$50 million in revenue. Moreover, external investment remains highly concentrated geographically, dominated by a few technology hubs in the United States and China, with Europe lagging far behind. We explore these issues further in Chapter 3.

FIRMS AND INDUSTRIES ALREADY ON THE DIGITAL FRONTIER ARE ADOPTING AI, BUT OTHERS ARE HESITANT TO ACT

Investors are pouring billions of dollars into AI companies based on the hope that a market of AI adopters will develop fairly quickly and will be willing to pay for AI infrastructure, platforms, and services. Clearly, Amazon, Google, and other digital natives are investing for their own applications, such as optimizing searches and personalizing marketing. But getting a sense of how much traditional companies in health care, retail, and telecom are spending on AI is not easy. For this reason, we conducted a survey to understand this situation in more depth.

In general, few companies have incorporated AI into their value chains at scale; a majority of companies that had some awareness of AI technologies are still in experimental or pilot phases. In fact, out of the 3,073 respondents, only 20 percent said they had adopted one or more AI-related technology at scale or in a core part of their business.¹⁶ Ten percent reported adopting more than two technologies, and only 9 percent reported adopting machine learning.¹⁷

Even this may overstate the commercial demand for AI at this point. Our review of more than 160 global use cases across a variety of industries found that only 12 percent had progressed beyond the experimental stage. Commercial considerations can explain why some companies may be reluctant to act. In our survey, poor or uncertain returns were the primary reason for not adopting reported by firms, especially smaller firms. Regulatory concerns, explored further in Chapter 3, also have become much more important.

As with every new wave of technology, we expect to see a pattern of early and late adopters among sectors and firms. We uncover six features of the early pattern of AI adoption, which is broadly in line with the ways companies have been adopting and using the recent cohort of digital technologies. Not coincidentally, the same players who were leaders in that earlier wave of digitization are leading in AI—the next wave.

¹⁵ It is worth noting that VC funds were focusing on AI technology when choosing investments, while PE funds were investing in AI-related companies.

¹⁶ Survey results throughout this discussion paper are weighted for firm size; “20 percent of firms” indicates firms representing 20 percent of the workforce. See Appendix B for an explanation of the weighting methodology.

¹⁷ The eight technologies are natural language processing, natural language generation, speech recognition, machine learning, decision management, virtual agents, robotics process automation, and computer vision. The five technology systems are robotics and autonomous vehicles, computer vision, language, virtual agents, and machine learning.

The first feature is that early AI adopters are from sectors already investing at scale in related technologies, such as cloud services and big data. Those sectors are also at the frontier of digital assets and usage.¹⁸ This is a crucial finding, as it suggests that there is limited evidence of sectors and firms catching up when it comes to digitization, as each new generation of tech builds on the previous one.

Second, independently of sectors, large companies tend to invest in AI faster at scale. This again is typical of digital adoption, in which, for instance, small and midsized businesses have typically lagged behind in their decision to invest in new technologies.

Third, early adopters are not specializing in one type of technology. They go broader as they adopt multiple AI tools addressing a number of different use cases at the same time.

Fourth, companies investing at scale do it close to their core business.

Fifth, early adopters that adopt at scale tend to be motivated as much by the upside growth potential of AI as they are by cutting costs. AI is not only about process automation, but is also used by companies as part of major product and service innovation. This has been the case for early adopters of digital technologies and suggests that AI-driven innovation will be a new source of productivity and may further expand the growing productivity and income gap between high-performing firms and those left behind.¹⁹

Finally, strong executive leadership goes hand in hand with stronger AI adoption. Respondents from firms that have successfully deployed an AI technology at scale tended to rate C-suite support nearly twice as high as those from companies that had not adopted any AI technology.

Early-adopting sectors are closer to the digital frontier

Sector-by-sector adoption of AI is highly uneven right now, reflecting many features of digital adoption more broadly. Our survey found that larger companies and industries that adopted digital technologies in the past are more likely to adopt AI. For them, AI is the next wave of digitization.

This pattern in the adoption of technology is not new—we saw similar behavior in firms adopting enterprise social technologies.²⁰ But this implies that, at least in the near future, AI deployment is likely to accelerate at the digital frontier, expanding the gap between adopters and laggards across companies, industries, and geographic regions.

The leading sectors include some that MGI's Industry Digitization Index found at the digital frontier, namely high tech and telecom and financial services.²¹ These are industries with long histories of digital investment. They have been leaders in developing or adopting digital tools, both for their core product offerings and for optimizing their operations. However, even these sectors are far behind in AI adoption when compared with overall digitization (Exhibit 3).

¹⁸ *Digital Europe: Pushing the frontier, capturing the benefits*, McKinsey Global Institute, June 2016; *Digital America: A tale of the haves and have-mores*, McKinsey Global Institute, December 2015.

¹⁹ Rosina Moreno and Jordi Suriñach, "Innovation adoption and productivity growth: Evidence for Europe," working paper, 2014; Jacques Bughin and Nicolas van Zeebroeck, "The right response to digital disruption," *MIT Sloan Management Review*, April 2017.

²⁰ Jacques Bughin and James Manyika, "How businesses are using web 2.0: A McKinsey global survey," *McKinsey Quarterly*, December 2007; Jacques Bughin and James Manyika, "Bubble or paradigm change? Assessing the global diffusion of enterprise 2.0," in Alex Koohang, Johannes Britz, and Keith Harman, eds., *Knowledge management: Research and applications*, Informing Science, 2008.

²¹ *Digital America: A tale of the haves and have-mores*, McKinsey Global Institute, December 2015.

Exhibit 3

AI adoption is occurring faster in more digitized sectors and across the value chain

AI Index

Relatively low  Relatively high

	Overall AI index	MGI Digitization Index ¹	Assets			Usage						Labor	
			Depth of AI technologies	AI spend	Supporting digital assets	Product development	Operations	Supply chain and distribution	Customer experience	Financial and general management	Workforce management	Exposure to AI in workforce	AI resources per worker
High tech and telecommunications													
Automotive and assembly													
Financial services													
Resources and utilities													
Media and entertainment													
Consumer packaged goods													
Transportation and logistics													
Retail													
Education													
Professional services													
Health care													
Building materials and construction													
Travel and tourism													

1 The MGI Digitization Index is GDP weighted average of Europe and United States. See Appendix B for full list of metrics and explanation of methodology.

SOURCE: McKinsey Global Institute AI adoption and use survey; *Digital Europe: Pushing the frontier, capturing the benefits*, McKinsey Global Institute, June 2016; *Digital America: A tale of the haves and have-mores*, McKinsey Global Institute, December 2015; McKinsey Global Institute analysis

Automotive and assembly is also highly ranked. It was one of the first sectors that implemented advanced robotics at scale for manufacturing, and today is also using AI technologies to develop self-driving cars.

In the middle are less digitized industries, including resources and utilities, personal and professional services, and building materials and construction. A combination of factors may account for this. These sectors have been slow to employ digital tools generally, except for some parts of the professional services industry and large construction companies. They are also industries in which innovation and productivity growth has lagged, potentially in part due to their domestic focus. Some of these sectors have a particularly high number of small firms—an important predictor for AI adoption, as explored below.

Toward the bottom of the pack for now are traditionally less digital fields such as education and health care. Despite ample publicity about cutting-edge AI applications in these industries, the reality is that uptake appears to be low so far. Weaker adoption reflects the particular challenges faced in these sectors. In health care, for example, practitioners and administrators acknowledge the potential for AI to reduce costs but quickly add that they believe that regulatory concerns and customer acceptance will inhibit adoption.

When it comes to adopting AI, the bigger, the bolder

A stylized fact in IT literature is that large firms usually are early adopters of innovative technology, while smaller firms are more reluctant to be first movers.²² We find the same digital divide when we look at AI: large firms have much higher rates of adoption and awareness. Across all sectors, larger firms—which we define as those with more than 500 employees—are at least 10 percent more likely than smaller firms to have adopted at least one AI technology at scale or in a core part of their business. In sectors with lower rates of AI uptake, the adoption rate of bigger companies was as much as 300 percent that of smaller companies.

Other digitization indicators reflect this fact, as highlighted in MGI's digitization work. Larger firms typically have access to more and better-structured data, and are more likely to have employees with the technical skills needed to understand the business case for AI investment and to successfully engage suppliers. Bigger firms also have an advantage because the kind of fixed-cost investment required for AI tends to generate higher returns when applied to a bigger base of costs and revenue.

Nonetheless, we find success stories among some smaller firms, too. Relative to larger companies, they can benefit from fewer issues with legacy IT systems and lower levels of organizational resistance to change. Smaller firms can also benefit from AI tools provided as a service.

Early AI adopters tend to become serial adopters

We looked at how firms deploy AI across eight different application areas and five technology systems.²³ Our results suggest that early-adopting firms are looking across multiple AI tools when they begin to adopt, rather than focusing on a particular technology. This is consistent with adoption patterns in other digital technologies.²⁴

The phenomenon of multitechnology application is persistent at a sector level. Industries with high rates of adopting one technology have higher rates in adopting others. High tech and telecom, for example, report the highest rates of adoption across all five technology groups, while construction is among the lowest among all five.

However, there are anomalies. Education and health care are notable for being slow to adopt AI technology. In frontier sectors—those with a relatively high percentage of early adopters—two-thirds of firms that had already adopted one of the eight AI technologies had adopted at least two others as well. In health care, only one-third had, with language technologies the most likely to be deployed at scale or in a core part of the business.

²² Kevin Zhu, Kenneth L. Kraemer, and Sean Xu, "The process of innovation assimilation by firms in different countries: A technology diffusion perspective on e-business," *Management Science*, volume 52, number 10, October 2006; Chris Forman, Avi Goldfarb, and Shane Greenstein, "The geographic dispersion of commercial Internet use," in *Rethinking rights and regulations: Institutional responses to new communication technologies*, Lorrie Faith Cranor and Steven S. Wildman, eds., MIT Press, 2003.

²³ The eight technologies are: natural language processing, natural language generation, speech recognition, machine learning, decision management, virtual agents, robotics process automation, and computer vision. The five technology systems are: robotics and autonomous vehicles, computer vision, language, virtual agents, and machine learning.

²⁴ Sanjeev Dewan, Dale Ganley, and Kenneth L. Kraemer, "Complementarities in the diffusion of personal computers and the internet: Implications for the global digital divide," *Information Systems Research*, volume 21, number 5, December 2010.

Users are keeping artificial intelligence close to their core

Functionally, AI technologies are finding applications across the value chain, but with some parts of the value chain getting more attention than others. For example, customer service functions such as sales and marketing, as well as operations and product development, all tend to use the most commonly cited AI applications. General and financial management, by contrast, lag well behind. A similar pattern is found in big data. The literature shows that the most frequent big data applications originate in sales and marketing functions.²⁵

In general, firms queried in our survey say they tend to adopt AI technologies affecting the part of their value chain closest to the core. Operations are an important area of adoption in the automotive and assembly, and consumer packaged goods sectors, as well as utilities and resources. Operations and customer service are the most important areas for financial services. This is new. Previously, new digital technology tended to remain on the margins, away from the core of the business.

However, in line with trends in technology, we also see sectors going deeper and broader as they increase their degree of AI adoption. Leading sectors are not only more extensively deploying AI in the core parts of their value chain, they are also deploying it in more parts of their value chain.

Early adopters see AI increasing revenue while companies experimenting with AI expect lower costs

As companies become more familiar with AI, their perceptions about its benefits change. The results of survey analysis show that early AI adopters are driven to employ AI technologies in order to grow revenue and market share, and the potential for cost reduction is a secondary idea. Firms that we consider more advanced AI adopters were 27 percent more likely to report using AI to grow their market than companies only experimenting with or partially adopting AI, and 52 percent more likely to report using it to increase their market share. Experimenters, by contrast, were more focused on costs. They were 23 percent more likely than advanced AI adopters to point to labor cost reductions, and 38 percent more likely to mention non-labor cost reductions.

In other words, the more companies use and become familiar with AI, the more potential for growth they see in it. Companies with less experience tend to focus more narrowly on reducing costs. The employment implications are further discussed in Chapter 3.

AI is not only about technical adoption, it is about enterprise acceptance

To be successful, AI adoption requires buy-in by the executive suite to generate the momentum needed to overwhelm organizational inertia.

Successful AI adopters, according to our survey, have strong executive leadership support for the new technology. Representatives of firms that have successfully deployed an AI technology at scale tended to rate C-suite support nearly twice as high as those of companies that had not adopted any AI technology. They added that strong support came not only from the CEO and IT executives—that is, chief information officer, chief digital officer, and chief technology officer—but from all other C-level officers and the board of directors as well.

Successful adopters also adjusted their firm-wide strategy to become proactive toward AI. See more details in Chapter 2.

²⁵ Jacques Bughin, “Ten big lessons learned from big data analytics,” *Applied Marketing Analytics*, volume 2, number 4, 2017.

AI'S NEXT CHALLENGE: GET USERS TO ADAPT AND ADOPT

IT industry analysts concur that the market size for AI technology will experience strong growth over the next three years. Most of the firms we surveyed expected to increase spending on AI in the coming three years, a finding echoed in other recent surveys. For example, 75 percent of the 203 executives queried in an Economist Intelligence Unit survey said AI would be “actively implemented” in their firms within three years (3 percent said it had already happened).

Expectations of how large this growth will be vary widely. Our survey documented relatively modest growth projections—only one-fifth of firms expected to increase expenditure by more than 10 percent. Industry analysts’ forecasts of the compound annual growth rate ranged from just under 20 percent to nearly 63 percent, including both adoption by additional companies and increased spending within companies.²⁶ The actual growth rate may need to be toward the upper end of that range to meet the expectations of investors piling into the industry.

Growth will hinge on the ability of sectors and firms to overcome technical, commercial, and regulatory challenges. Our survey respondents and outside forecasters expect financial services, retail, health care, and advanced manufacturing to be in the AI vanguard. These are the industries where technical feasibility is relatively high (reflected in the case studies on the market today) and the business case for AI is most compelling. They are also the sectors with the highest degree of digital adoption to date—a key foundation for AI (Exhibit 4).

Technical challenges are an important differentiating factor between industries. While big tech and academia are pushing advances in the performance of the underlying technology, engineering solutions need to be worked out for specific use cases, requiring both data and talent. Industries such as financial services, and high tech and telecom have generated and stored large volumes of structured data, but others, including construction and travel, lag far behind.²⁷

Commercial drivers also differ between sectors. Industries most likely to lead the adoption of AI technologies at scale are those with complex businesses in terms of both operations and geography, whose performance is driven by forecasting, fast and accurate decision making, or personalized customer connections. In financial services, there are clear benefits from improved accuracy and speed in AI-optimized fraud-detection systems, forecast to be a \$3 billion market in 2020. In Chapter 2 (and the supporting appendix) we explore how these commercial drivers play out in other industries. For example, in retail, there are compelling benefits from improved inventory forecasts, automated customer operations, and highly personalized marketing campaigns. Similarly, in health care, AI-powered diagnosis and treatment systems can both save costs and deliver better outcomes for patients.

Even where compelling commercial use cases have been engineered and are demanded by firms, regulatory and social barriers can raise the cost and slow the rate of adoption. Product liability is one such concern; it is especially troublesome for automakers and other manufacturers. Privacy considerations restrict access to data and often require it to be anonymized before it can be used in research. Ethical issues such as trained biases and algorithmic transparency remain unresolved. (For further discussion, see Box 4 in Chapter 3, “An overview of ongoing challenges.”) Preferences for a human relationship in settings such as health care and education will need to be navigated. Job security concerns could also limit market growth—there are already serious calls for taxes on robots.

²⁶ The full range of forecasts: BCC Research, 19.7 percent; Transparency Market Research, 36.1 percent, Tractica, 57.6 percent; IDC, 58 percent; and Markets and Markets, 62.9 percent.


²⁷ *A future that works: Automation, employment, and productivity*, McKinsey Global Institute, January 2017.

Exhibit 4

Sectors leading in AI adoption today also intend to grow their investment the most

Future AI demand trajectory¹

Average estimated % change in AI spending, next 3 years, weighted by firm size²



1 Based on the midpoint of the range selected by the survey respondent.
2 Results are weighted by firm size. See Appendix B for an explanation of the weighting methodology.

SOURCE: McKinsey Global Institute AI adoption and use survey; McKinsey Global Institute analysis

These forces will help determine the industries that AI is likely to transform the most. However, if current trends hold, variation of adoption within industries will be even larger than between industries. We expect that large companies with the most digital experience will be the first movers because they can leverage their technical skills, digital expertise, and data resources to develop and smoothly integrate the most appropriate AI solutions.



After decades of false starts, artificial intelligence is on the verge of a breakthrough, with the latest progress propelled by machine learning. Tech giants and digital natives are investing in and deploying the technology at scale, but widespread adoption among less digitally mature sectors and companies is lagging. However, the current mismatch between AI investment and adoption has not stopped people from imagining a future where AI transforms businesses and entire industries. In the next chapter, we explore the four major ways in which AI can create value across the value chain in different sectors.

2. ARTIFICIAL INTELLIGENCE PROMISES TO BOOST PROFITS AND TRANSFORM INDUSTRIES

AI technologies have advanced significantly in recent years. Adoption, however, remains in its infancy. This makes it challenging to assess the true potential impact of AI on firms and sectors. We do know that many non-adopters report that they have trouble making the business case for AI investment, but what about the firms that have adopted? Looking at case studies of digital natives and responses from our survey, we find early evidence that AI implemented at scale delivers attractive returns. Through a review of a large number of case studies in five sectors, we also show how AI can transform some business activities and has the potential to fundamentally change others. The cases demonstrate how AI can shape different functions across the whole value chain and in diverse sectors. The cases have wide-ranging implications for many stakeholders, including multinational corporations, startups, governments, and social institutions.

FIRMS THAT COMBINE STRONG DIGITAL CAPABILITY, ROBUST AI ADOPTION, AND A PROACTIVE AI STRATEGY SEE OUTSIZE FINANCIAL PERFORMANCE

Digital native companies made some of the most significant and earliest investment in AI, providing test cases for potential return on investments in AI. Amazon has achieved impressive results from its \$775 million acquisition of Kiva, a robotics company that automates picking and packing. “Click to ship” cycle time, which ranged from 60 to 75 minutes with humans, fell to 15 minutes with Kiva, while inventory capacity increased by 50 percent. Operating costs fell an estimated 20 percent, giving a return of close to 40 percent on the original investment.²⁸

Netflix has also achieved impressive results from the algorithm it uses to personalize recommendations to its 100 million subscribers worldwide. Helping customers quickly find desirable content is critical—customers tend to give up if it takes longer than 90 seconds to find a movie or TV show they want to watch. Through better search results, Netflix estimates that it is avoiding cancelled subscriptions that would reduce its revenue by \$1 billion annually.²⁹

Firms adopting AI at scale or in a core part of their business already see the technology’s potential, and those implementing proactive AI strategies are anticipating even greater benefits. Using our survey results, we compared the current self-reported profit margins of firms with differing levels of AI adoption, digital maturity (as reflected in their use of big data and cloud services), and critically, their strategic posture (Exhibit 5). We found that serious adopters have significantly higher projected margins than others.


This suggests AI can deliver significant competitive advantages, but only for firms that are fully committed to it. Take any ingredient away—a strong digital starting point, serious adoption of AI, or a proactive strategic posture—and profit margins are much less impressive. This is consistent with our findings in the broader digital space. Technology is a tool and in itself does not deliver competitiveness improvements. In Chapter 3, we explore what firms need to do to successfully adopt AI, including identifying value-adding use cases, building capabilities, and embedding a collaborative culture.

²⁸ Eugene Kim, “Amazon’s \$775 million deal for robotics company Kiva is starting to look really smart,” *Business Insider*, June 15, 2016.

²⁹ Nathan McAlone, “Why Netflix thinks its personalized recommendation engine is worth \$1 billion per year,” *Business Insider*, June 14, 2016.

Exhibit 5

AI adopters with a proactive strategy have significantly higher profit margins



1 Firms that are big data and cloud services users and report their strategic posture toward AI to be: “Disrupting our industry using AI technology is at the core of our strategy,” “We have changed our longer-term corporate strategy to address the AI threat or opportunity disruption,” or “We have developed a coordinated plan to respond to the AI threat or opportunity but have not changed our longer-term corporate strategy.”

2 Operating profit margin for selected sectors as a share of turnover, for continuing operations and before exceptional items.

SOURCE: McKinsey Global Institute AI adoption and use survey; McKinsey Global Institute analysis

The same pattern is apparent when we analyze expectations of future profit margins. Not only do serious AI adopters with proactive strategies report current profit margins that are three to 15 percentage points higher than the industry average in most sectors, but they also expect this advantage to grow in the future, when they could expect their AI investment to mature and start paying substantial dividends. In the next three years, these AI leaders expect their margins to increase by up to five percentage points more than the industry average.

How can AI tangibly provide new sources of value for businesses and institutions? We have looked at use cases across selected sectors to understand how AI could impact organizations and their value chains.

INDUSTRY CASE STUDIES DEMONSTRATE THE DISRUPTIVE POTENTIAL OF AI

To develop a vision of AI as it could be applied across business domains, we created five case studies in different sectors that suggest how AI in its many guises could affect specific activities (see Appendix A). The sectors we chose—retail, electric utilities, manufacturing, health care, and education—involve a mixture of private, public, and social enterprises, as well as covering the range from highly labor-intensive customer-oriented sectors to more asset-heavy business-to-business operations.

To fulfil the expectations being heaped upon it, AI will need to deliver economic applications that significantly reduce costs, increase revenue, and enhance asset utilization. We categorized the ways in which AI can create value in four areas: enabling companies to better **project** and forecast to anticipate demand, optimize R&D, and improve sourcing; increasing companies' ability to **produce** goods and services at lower cost and higher quality; helping **promote** offerings at the right price, with the right message, and to the right target customers; and allowing them to **provide** rich, personal, and convenient user experiences.

These four areas of value creation are based on specific use cases that are being explored or have been deployed in businesses today. The list, which may not be exhaustive, is based on our current knowledge of narrow AI technologies. Also, they will have different degrees of relevance for sectors and industries, with the **project** and **produce** levers being particularly rich with opportunities to leverage AI. In addition, while machine learning can bring highly valuable benefits to all sectors, some technologies are particularly suited for business application in specific sectors, such as robotics for retail and manufacturing, computer vision for health care, and natural language processing and generation for education (see Exhibit 6).

Project: Accurately forecast demand, optimize supply, and shape future offerings for success


The first area in which AI can create value is projection and forecasting. Organizations need to constantly anticipate the future to gain competitive advantage. AI allows businesses to provide better forecasts for their supply chain and design better offerings. Reliably forecasting demand is a way to use AI's ability to digest disparate data and automatically adjust to new information. It can discern trends and patterns that can be acted on. Businesses use this tool in a number of ways, such as forecasting demand to stock only the specific quantities of specific products they will sell and thus minimize waste, and anticipating sales trends so they can order more soon-to-be-popular items. Indeed, the benefits of projecting demand go beyond traditional business sectors: for example, by using sophisticated algorithms, health systems can increasingly predict—and prevent—major epidemics.

The benefits of AI-enabled demand forecasting are impressive in retail, for instance. In some settings, AI-based approaches to demand forecasting are expected to reduce forecasting errors by 30 to 50 percent from conventional approaches (Exhibit 7). Lost sales due to product unavailability can be reduced by up to 65 percent. Costs related to transportation and warehousing and supply chain administration are expected to decrease by 5 to 10 percent and 25 to 40 percent, respectively. With AI, overall inventory reductions of 20 to 50 percent are feasible.³⁰

³⁰ *Smartening up with artificial intelligence (AI): What's in it for Germany and its industrial sector?* McKinsey & Company, April 2017.

Exhibit 6

Artificial intelligence can create value across the value chain in four ways



	Applicable technologies	Project	Produce	Promote	Provide
		Enlightened R&D, real-time forecasting, and smart sourcing	Operations with higher productivity, lower cost, and better efficiency	Products and services at the right price, with the right message, and to the right targets	Enriched, tailored, and convenient user experience
Retail		Anticipate demand trends, while optimizing and automating supplier negotiation and contracting	Automate warehouse and store operations; optimize merchandising, product assortment, and microspace	Optimize pricing, personalize promotions, and tailor website displays in real time	Personalize tips and suggestions, offer immediate assistance with virtual agents, automate in-store checkout, and complete last-mile delivery by drones
Electric utilities		Enhance demand and supply prediction, assess reliability of integrated generation assets, and automate demand-side response	Optimize preventive maintenance, improve electricity production yield, reduce energy waste, and prevent electricity theft	Optimize pricing with time-of-day and dynamic tariffing; match producers and consumers in real time	Automate supplier selection, provide consumption insights, automate customer service with virtual agents, and tailor usage to consumer's preferences
Manufacturing		Improve product design yield and efficiency, automate supplier assessment, and anticipate parts requirements	Improve processes by the task, automate assembly lines, reduce errors, limit product rework, and reduce material delivery time	Predict sales of maintenance services, optimize pricing, and refine sales-leads prioritization	Optimize flight planning and route and fleet allocation; enhance maintenance engineer and pilot training
Health care		Predict disease, identify high-risk patient groups, and launch prevention therapies	Automate and optimize hospital operations; automate diagnostic tests and make them faster and more accurate	Predict cost more accurately, focus on patients' risk reduction	Adapt therapies and drug formulations to patients, use virtual agents to help patients navigate their hospital journey
Education		Anticipate job market demand, identify new drivers of performance to assess students, and help graduates highlight their strengths	Automate teachers' routine tasks, identify early disengagement signs, and optimize group formation for learning objectives		Personalize learning, shift from stop-and-test model to continuous learning cadenced by virtual coaches and tutors, and build student self-awareness

1 Machine learning for multiuse and nonspecific applications.

SOURCE: McKinsey Global Institute analysis

Exhibit 7

AI can help capture significant gains across the value chain

Examples of AI-related business impact from current use cases

	Project	Produce	Promote	Provide
	Accurate demand forecasting, smart sourcing, and enlightened R&D	Higher productivity and minimized maintenance and repairs	Products and services at the right price, with the right message, to the right targets	Enriched, tailored, and convenient user experience
Retail	<ul style="list-style-type: none"> ▪ 1–2% EBIT¹ improvement using machine learning to anticipate fruit and vegetable sales ▪ 20% stock reduction using deep learning to predict e-commerce purchases ▪ 2 million fewer product returns per year 	<ul style="list-style-type: none"> ▪ 30% reduction of stocking time using autonomous vehicles in warehouses 	<ul style="list-style-type: none"> ▪ 50% improvement of assortment efficiency ▪ 4–6% sales increase using geospatial modeling to improve micromarket attractiveness ▪ 30% online sales increase by using dynamic pricing and personalization 	
Electric utilities	<ul style="list-style-type: none"> ▪ Objective to cut 10% in national electricity usage by using deep learning to predict power demand and supply 	<ul style="list-style-type: none"> ▪ 20% energy production increase using machine learning and smart sensors to optimize assets' yield ▪ 10–20% EBIT improvement by using machine learning to enhance predictive maintenance, automate fault prediction, and increase capital productivity 		<ul style="list-style-type: none"> ▪ \$10–\$30 savings on monthly bills by using machine learning to automatically switch electricity supply deals
Manufacturing	<ul style="list-style-type: none"> ▪ 10% yield improvement for integrated-circuit products using AI to improve R&D process ▪ 39% IT staff reduction by using AI to fully automate procurement processes 	<ul style="list-style-type: none"> ▪ 30% increase of material delivery time using machine learning to determine timing of goods' transfer ▪ 3–5% production yield improvement 	<ul style="list-style-type: none"> ▪ 13% EBIT improvement by using machine learning to predict sources of servicing revenues and optimize sales efforts 	<ul style="list-style-type: none"> ▪ 12% fuel savings for manufacturers' customers, airlines, by using machine learning to optimize flight routes
Health care	<ul style="list-style-type: none"> ▪ \$300 billion possible savings in the United States using machine learning tools for population health forecasting ▪ £3.3 billion possible savings in the United Kingdom using AI to provide preventive care and reduce nonelective hospital admissions 	<ul style="list-style-type: none"> ▪ 30–50% productivity improvement for nurses supported by AI tools ▪ Up to 2% GDP savings for operational efficiencies in developed countries 	<ul style="list-style-type: none"> ▪ 5–9% health expenditure reduction by using machine learning to tailor treatments and keep patients engaged 	<ul style="list-style-type: none"> ▪ \$2 trillion–\$10 trillion savings globally by tailoring drugs and treatments ▪ 0.2–1.3 additional years of average life expectancy
Education		<ul style="list-style-type: none"> ▪ Virtual teaching assistants can answer 40% of students' routine questions 	<ul style="list-style-type: none"> ▪ 1% increase in enrollment by using a virtual assistant to follow up with applicants 	<ul style="list-style-type: none"> ▪ 85% match with human grading, using machine learning and predictive modelling

1 Earnings before interest and taxes.

SOURCE: McKinsey Global Institute analysis

When it comes to matching supply and demand, electric utilities are a special case: they need to do it, literally, in real time. Making short-term load forecasts more accurate in order to adjust supply to meet anticipated demand can deliver enormous savings, reduce waste and emissions, and add to system resilience. National Grid in the United Kingdom is collaborating with DeepMind, an AI startup bought by Google in 2014, to predict supply and demand variations based on weather-related variables and smart meters as exogenous inputs. The goal is to cut national energy use by 10 percent and maximize the use of renewable power. AI also is used to briefly switch off air conditioning at participating businesses as it forecasts the approach of peak consumption, easing the load for all and postponing or even forgoing the need to fire up peak generating capacity.

Forecasting is not a new idea. A well-functioning supply chain is the backbone of virtually every industry. Accurate projections to ensure just the right amount of inventory are critical to achieving a competitive advantage. Factors such as product introductions, distribution network expansion, weather forecasts, extreme seasonality, and changes in customer perception or media coverage can severely affect the performance of the supply chain. Traditional systems for forecasting and replenishment might not take advantage of the amount of data associated with the internet of things and the sheer number of influencing factors. Supply-chain leaders are starting to realize the ability of machine learning–based methods to increase forecasting accuracy and optimize replenishment. The objective is to reduce swings in inventory levels and increase flexibility. Machine learning approaches not only incorporate historical sales data and the setup of the supply chains, but also rely on near-real-time data regarding variables such as advertising campaigns, prices, and local weather forecasts.

Using AI to forecast demand also allows businesses to optimize their sourcing more broadly, including fully automating purchases and order processing. The German online retailer Otto uses an AI application that is 90 percent accurate in forecasting what the company will sell over the next 30 days. The forecasts are so reliable that Otto now builds inventory in anticipation of the orders AI has forecast, enabling the retailer to speed deliveries to customers and reduce returns. Otto is confident enough in the technology to let it order 200,000 items a month from vendors with no human intervention.³¹

But AI is helpful not just for forecasting demand for current products; it can also be used by R&D departments, partly to help researchers quickly assess whether a prototype would be likely to succeed or fail in the market—and why. Engineers and researchers today face difficult challenges, from the sharp growth in demand in emerging countries to market fragmentation driven by consumers' taste for customization. At the same time, budget constraints require engineering teams to improve their productivity and efficiency, even as limits on the number of designs that can be tested restrict the predictability of product performance. These features may be particularly attractive to the health and pharmaceuticals sectors, which spend \$160 billion annually on R&D in the United States alone.

AI-powered technologies can help deliver more efficient designs than were previously achievable by eliminating waste in the design process. Innovations can be brought to market faster as AI facilitates faster process cycle times and an increased focus on real-time negotiations and other interactions. AI-based approaches to increasing R&D project performance can result in productivity gains of 10 to 15 percent. Time to market could be reduced by 10 percent or more.³² Motivo, an artificial intelligence startup, compresses

³¹ "How Germany's Otto uses artificial intelligence," *The Economist*, April 12, 2017.

³² *Smartening up with artificial intelligence (AI): What's in it for Germany and its industrial sector?* McKinsey & Company, April 2017.

design processes that used to take months—sometimes a year or more—to roughly four weeks, saving chip makers the cost of iterations and testing.

Produce: Get more out of machines while minimizing maintenance and repairs

The second area where AI can help create value is production, or the transformation of inputs into outputs, whether products (semiconductors, aircraft engines) or services (teaching, health care, or power or consumer goods distribution). AI can help businesses produce by continually optimizing assets and processes, assembling the best teams of people and robots, improving quality and reliability, and preventing downtime for maintenance—all of which increase productivity.

The obvious role for AI is to replace humans through automation. However, in some situations, AI is complementing teams of people. Ocado, the UK online supermarket, illustrates how this happens when a company embeds AI and robotics at the core of its operations. In the retailer's warehouse, robots steer thousands of product-filled bins over a maze of conveyor belts and deliver them to human packers just in time to fill shopping bags. Other robots whisk the bags to delivery vans whose drivers are guided to customers' homes by an AI application that picks the best route based on traffic conditions and weather.

Advances in computer vision are behind many developments in collaborative and context-aware robots. Enhanced vision is enabled by more powerful computers, new algorithmic models, and large training data sets. Within the field of computer vision, object recognition and semantic segmentation—that is, the ability to categorize object type, such as distinguishing a tool from a component—have recently advanced significantly in their performance. They allow robots to behave appropriately for the context in which they operate, for example by recognizing the properties of the materials and objects they interact with. They are flexible and autonomous systems that are capable of safely interacting with the real world and humans.

This is a notable shift since current robotic systems are often limited in their ability to identify a specific object and react in a flexible manner to changes in their environment, and therefore are programmed to follow predefined steps. However, new AI-enhanced, camera-equipped logistics robots can be trained to recognize empty shelf space. This leads to a dramatic speed advantage over conventional methods in picking objects. Deep learning can also be used to correctly identify an object and its position. This enables robots to handle objects without requiring the objects to be in fixed, predefined positions. AI-enhanced logistics robots are also able to integrate disturbances in their movement routines via an unsupervised learning engine for dynamics. This capability leads to more precise makeovers and an overall improved robustness of processes. Rethink Robotics is one company designing collaborative robots. Its learning algorithms allow for joint human-robot collaborative work spaces. AI enables the “programming” of a robot by simply showing it the desired movements. A human “robot instructor” can take the machine's arm and guide it through the desired movement, including gripping and releasing objects. Robot movement is then the result of the robot's replication and further improvement of the freshly learned movement combined with a computer-vision-based assessment of an object's position in space.

Collaborative robots are particularly relevant with respect to tasks that are not fully automatable. In such settings they hold the potential to increase productivity by up to 20 percent.³³ Using AI to improve humans' efficiency is a critical enabler of productivity. In education, AI tools can help teachers accelerate administrative tasks and daily operations. AI-enabled applications to help teachers quickly grade students' work are already available. They decipher students' handwriting, learn how the teacher grades the first few tests, and

³³ Ibid.

apply the same standards to subsequent tests. This works well with objectively correct answers, such as math problems. The true step change will come with the cracking of natural language understanding, which would allow the automatic grading of more creative work, such as essays and presentations. While this seems to be feasible, it is still in development. In a multivendor evaluation at the University of Akron in 2012, grading software gave marks to 16,000 essays that instructors had already reviewed; computers matched the teachers' grades about 85 percent of the time.³⁴ Still unresolved are questions around creativity and pattern disruption: how can machine learning be taught not to encourage uniform thinking or the upholding of biases?

AI tools can also improve productivity by better aligning team formation with the team's objective. Collaboration.ai uses artificial intelligence models to assist teachers in forming optimized classes. It analyzes data from students' education profiles, social media accounts, and surveys to populate a class or study group with students that have compatible skills and personalities. The emergence of collaborative robots is further improving human-machine compatibility.

In addition to speeding up processes, reducing costs, and increasing output, AI has huge potential to improve quality by reducing errors. Yield losses—losses incurred when products have to be disposed of or reworked due to defects—play an important role in complex manufacturing environments.

The multistep semiconductor chip-production process is a good example because cycle times from the first processing of the wafer to the final chip are typically several weeks to months and include various intermediate quality-testing processes. Testing costs and yield losses in semiconductor production can constitute up to 30 percent of the total production cost. Ample data are usually available in semiconductor fabrication plants, or fabs, due to their high degree of automation and advanced production equipment. They often have archives that allow insights into detailed production information dating back months or years. However, systematic analysis and linkage of data sources across multiple tool groups is not always performed. Semiconductor manufacturers are starting to use AI engines to identify root causes of yield losses that can be avoided by changing production processes. Enhanced applications are designed to monitor and adjust subprocesses in real time. AI techniques can help determine the optimized product operating conditions or process conditions to significantly reduce defects in manufacturing.

In asset-heavy businesses, keeping complex systems running with minimal downtime is another key opportunity for AI. Utility companies can shift from regularly scheduled maintenance of their extensive electrical grids to condition-based maintenance run by AI. Using data from sensors, drones, and other hardware, machine learning applications can help grid operators avoid decommissioning assets before their useful lives have ended, while simultaneously enabling them to perform more frequent remote inspections and maintenance to keep assets working well. One European power distribution company was able to reduce its cash costs by 30 percent over five years by changing its maintenance patterns based on remote analysis of 20 variables to determine the overall health of power transformers.

AI is enabling the “preventive maintenance” of people, too, and will do even more in the future. Applications powered by artificial intelligence will enable health-care providers to dramatically accelerate the shift toward personalized preventive medicine. Clinicians will focus on managing patients' health remotely via wearable wireless sensors, aiming to keep them healthy, fit, and out of hospitals. To do this, AI tools will take into account not only

³⁴ “Man and machine: Better writers, better grades,” University of Akron press release, April 12, 2012.

patients' medical histories and genetic makeup but also environmental factors that can influence health, such as pollution and noise where they live and work.

Promote: Charge the right price and deliver the right message to the right target

The third area where AI can create value is promotion, or marketing offerings at the right price, with the right message, and to the right target. Armed with enough of the right kind of data, companies can use artificial intelligence to price goods and services dynamically, raising prices when demand rises or a consumer appears willing to pay more, and lowering them when the opposite happens. Yield management programs have been dynamically pricing airline seats, hotel rooms, and other perishables for years, but AI will allow sellers to extend dynamic pricing to the rest of the marketplace.

Today, the requirements of intelligent price management are high: customers expect a good price, and price transparency for brand-name products is close to 100 percent. The basic question to ask for each item is: what price is the customer willing to pay? Hyperconnected consumers continuously redefine value by comparing prices online, even when browsing in a brick-and-mortar store. The optimal price for a product depends on many factors: the day of the week, season, time of day, weather, channel and device, competitors' prices, and much more. The challenge is to set the optimal price in relation to time. The right price at the right time increases customer satisfaction and leads to more sales and higher profit. AI can determine the price elasticity for every item and automatically adjust prices according to the chosen product strategy.

Similarly, energy retailers can use AI to create custom benefits such as low rates or extra service in order to hold on to their most valuable clients. While price sensitivity is a key consideration in attracting new customers and reducing churn, machine learning can also help address another critical component for utilities' marketing strategies—determining which customers are the most profitable. They are typically a small proportion (less than 20 percent) of a utility's customer base.

Retailers have accepted that smartphone penetration necessitates an omni-channel sales strategy, and AI can help optimize, update, and tailor it to each shopper in real time. Retailers can send mobile coupons—usually text messages with a discount code—to shoppers' smartphones as they approach a store, adding others based on how long the shopper stays. The size of the discount, the merchandise on offer, and other variables can be determined by an AI program that has looked for clues about what the shopper will like based on previous purchases, age, home address, web browsing habits, and mounds of other data. This kind of insights-based selling, including personalized promotions, optimized assortment, and tailored displays, can increase sales by 1 to 5 percent. Online, a focus on the most valuable customers, combined with dynamic pricing, can lead to a 30 percent growth in sales.

Aerospace companies, too, are using AI technologies to prioritize sales targets and optimize the price of services. For years, they prioritized maintenance, repair, and overhaul (MRO) sales leads manually, a cumbersome, resource-heavy, and not always efficient process. Using AI to improve the accuracy of forecasting MRO work and focusing the firm's sales efforts on the most promising leads can have a significant effect on profitability. One firm said its profit improved by over \$300 million by using machine learning to forecast 10 years of repair events for its transportation fleet and by developing a deal-scoring tool to advise on “what good looks like” when pricing after-market services.

Provide: Give customers a rich, personalized, and convenient experience

The fourth area where AI can create value is in enhancing the user experience and creating new sources of value to make it richer, more tailored, and more convenient. Making your best customers feel special and welcome is one way to foster loyalty and increase revenue, but it is difficult and expensive to do and so is often reserved for only the most lucrative clients. AI technologies like computer vision and machine learning can open a scaled-down version of the experience to many more people.

If, for example, a regular supermarket shopper puts a bunch of bananas in his cart, cameras or sensors could relay the information to an AI application that would have a good idea of what the shopper likes based on previous purchases. The app could then, via a video screen in the cart, suggest that bananas would be delicious with a chocolate fondue, which the purchase history suggests the shopper likes, and remind the shopper of where to find the right ingredients. Or a runner could download an app from an athletic shoe company, which would monitor her exercise regimen and recommend footwear tailored to her routine and running paths she may like.

Some of the same AI applications are also being employed in an experimental supermarket that may give new meaning to the idea of a “convenience store.” Amazon has built a retail outlet in Seattle that allows shoppers to take food off the shelves and walk directly out of the store without stopping at a checkout kiosk to pay. The store, called Amazon Go, relies on computer vision to track shoppers after they swipe into the store and associate them with products taken from shelves. When shoppers leave, Amazon debits their accounts for the cost of the items in their bag and emails them a receipt. In the future, the shopping experience will be completed by delivery drones for full convenience. Most of today’s efforts—by players as big as Amazon and as small as the Reno, Nevada, startup Flirtey—focus on unpiloted aerial drones. Flirtey made its first delivery, a box of snacks from a local convenience store, to a private residence in July 2016. Delivery drones will significantly benefit from breakthroughs in deep learning, which will help them categorize and handle anomalous situations, such as when no one is home to accept a delivery.

Personalizing user experiences has huge advantages in health care and education. In health care, treatment decisions based on AI analysis of existing science, data from tests, and patient monitoring with remote diagnostic devices carry the promise of significantly increased efficacy. Researchers are moving in this direction because standardized treatments do not work for everyone, given the complexity of each person’s history and genetic makeup. For a cancer patient, the technology models cell biology on the molecular level and seeks to identify the best drug to use for specific tumors. It can also identify complex biomarkers and search for combination therapies by performing millions of simulated experiments each day.

Several companies already use AI technologies to tailor treatments to individuals. Mindmaze uses AI to optimize rehabilitation for stroke patients. Ginger.io uses it to recommend the best time to take medication based each patient’s metabolism and other factors. A startup called Turbine uses AI to design personalized cancer treatment regimens.

This kind of tailored treatments may reduce health expenditures by 5 to 9 percent, add 0.2 to 1.3 years to average life expectancy, and increase productivity by \$200 per person per year. Globally, the economic impact of such advances could range from \$2 trillion to \$10 trillion.³⁵

³⁵ *The age of analytics: Competing in a data-driven world*, McKinsey Global Institute, December 2016.

In education, adaptive learning has been a growing trend, with some 40 companies, such as Knewton and DreamBox Learning, already marketing adaptive learning systems to schools in North America, Europe, and Asia.³⁶ Adaptive learning attempts to address the limitations of conventional classroom teaching by capturing information about what each student knows and crafting custom lesson plans based on individuals' knowledge and progress. Instead of marching the class through the same lesson (and often covering the same material repeatedly until most of the students get it), adaptive learning claims to deliver the right content, at the right time, in the best way for each student. Providing a personalized experience to students can have significant impact. At Arizona State University, an adaptive learning program helps students who are struggling with remedial math. Student pass rates have improved from 66 percent to 75 percent, and dropout rates are down by 7 percent.³⁷

Deep learning algorithms could further unleash the ability of AI techniques to tailor the teaching experience to students' specific needs and progress. Indicators such as facial expressions, digital interactions, group interactions, and eye tracking can be captured through computer vision to gauge students' engagement and coach and assess students in real time. A European Union project called iTalk2Learn is currently developing an open-source intelligent-tutoring platform to help primary school students learn mathematics. A combination of machine learning, computer vision, user modeling, and natural language processing enables the platform to interact with and respond to a student's speech throughout a tutoring session.

Common to all sectors is the need for employees to adapt their skills to support and complement the AI-powered user experience. As teaching becomes more adaptive, teachers and professors might have to adjust to doing less lecturing and more tutoring and coaching. Similarly, as sales and call centers become more virtual, human representatives must hone their emotional intelligence skills to provide a service experience beyond the capability of machines.



Early adopters and early case studies demonstrate AI's potential to transform business processes, shake up entire sectors, increase profits, and create new sources of value. AI applications are starting to reach maturity, and companies with serious, proactive adoption strategies stand to gain significant competitive advantages. There are many industry- and sector-specific use cases to inform companies when they define a focused strategy. Also, while machine learning and deep learning underpin most opportunities, industries will need to identify the AI technologies that will bring the most benefits to them, and then start to develop their infrastructure, talent, and knowledge as early as possible to catch up on the learning and adoption curves. AI is more than the sum of its parts: for truly impressive gains, companies are building their AI capability across the value chain, integrating it into core processes, and using it to enable their employees to be more productive.

In this chapter, we presented four areas where AI can create value. Such gains, however, do not come overnight. In the next chapter, we will look at what is required from businesses, AI vendors, and governments to seize this opportunity.

³⁶ *Learning to adapt: Understanding the adaptive learning supplier landscape*, Education Growth Advisors and Bill and Melinda Gates Foundation, 2013.

³⁷ Arizona State University.

3. BUSINESSES, DEVELOPERS, AND GOVERNMENTS NEED TO ACT NOW TO REALIZE AI'S FULL POTENTIAL

While AI has the potential to fundamentally reshape society, significant uncertainty remains about how the technology will develop. For firms, governments and workers, this might suggest a “wait and see” approach. However, we think there is a need for urgent but clear-headed action to respond to the opportunities and risks that are already apparent.

For many firms, this will mean accelerating their digital journey to ensure that they can effectively deploy AI tools. AI becomes impactful when it has access to large amounts of high-quality data and is integrated into automated work processes. AI is not a shortcut to these digital foundations. Rather, it is a powerful extension of them.

Developers have a crucial role to play in helping businesses realize the potential of the technology. AI products need to address real-world business problems, not just provide interesting solutions, and they must work at scale.

Finally, governments and workers should prepare for the wide-reaching changes ahead. Public education systems and workforce training programs will have to be rethought to ensure that workers have the skills to complement rather than compete with machines. Cities and countries hoping to establish local AI ecosystems must enter the global competition for AI talent and investment. And society as a whole will need to navigate the unresolved legal and ethical issues that could become major barriers to realizing the benefits of AI.


FIRMS NEED TO FOCUS ON AI USE CASES THAT ADD THE MOST VALUE AND ENSURE THEY HAVE THE RIGHT DIGITAL ENABLERS

While the overall impact of AI to date may be relatively small in many industries, its potential for disruption is high. We are already seeing examples of how AI-powered new entrants are able to take on incumbents and win, as Uber and Lyft have done to the taxi industry. AI can go beyond changing business processes to changing entire business models with winner-take-all dynamics, and firms that wait for the AI dust to settle risk being left behind.

This underscores the need for action now. Firms must conduct sober analysis of what the most valuable AI use cases are. They should also build out the supporting digital assets and capabilities. Indeed, the core elements of a successful AI transformation are the same as those for data and analytics generally (Exhibit 8). This includes building the data ecosystem, adopting the right techniques and tools, integrating technology into workplace processes, and adopting an open, collaborative culture while reskilling the workforce.

Exhibit 8

Successful AI transformations require elements similar to those found in successful digital and analytics transformations



SOURCE: *The age of analytics: Competing in a data-driven world*, McKinsey Global Institute, December 2016; McKinsey Global Institute analysis

Find the true source of value, and build a business case for it

The first step is to establish a solid AI business case and connect it to a firm’s strategy. This requires separating the hype and buzz around AI from its actual capabilities in a specific, real-world context. It includes a realistic view of AI’s capabilities and an honest accounting of its limitations, which requires at least a high-level grasp of how AI works and how it differs from conventional technological approaches.

A portfolio-based approach, looking at use cases over a one- to five-year horizon, can be helpful. In the immediate future, focus on use cases where there are proven technology solutions today that can be adopted at scale, such as robotic process automation and some applications of machine learning. Further out, identify use cases where a technology is emerging but not yet proven at scale. Over the longer term, pick one or two high-impact but unproven use cases and partner with academia or other third parties to innovate, gaining a potential first-mover advantage in the future. Examples include commercial underwriting in insurance and geo-modeling of subsurface deposits. Across all horizons, a “test and learn” approach can help validate the business case, conducting time-limited experiments to see

what really works and then scaling up successes. The importance of fast, agile approaches is discussed further later in this chapter.

To ensure a focus on the most valuable use cases, AI initiatives should be assessed and co-led by both business and technical leaders. Given the significant advancements in AI technologies in recent years, there is a tendency to compartmentalize accountability for AI with functional leaders in IT, digital, or innovation. This can result in a “hammer in search of a nail” outcome, or technologies being rolled out without compelling use cases. The orientation should be the opposite: business led and value focused. This business-led approach follows successful adoption approaches in other digital waves such as mobile, social, and analytics.³⁸

Build the data ecosystem

Data is at the heart of the disruptions occurring across economies and is recognized as an increasingly critical corporate asset. Without data, getting the AI engine started is impossible. Because of this, business leaders should know what data they already have access to and where they can obtain additional data relevant for their company’s future success. Google and Facebook are well-known examples of companies that obtain most of their revenue through insights they extract from the enormous quantities of data their customers generate on a daily basis by using their services.

One important capability will be making data usable that is not available in a relational format or that cannot be analyzed with traditional methodologies. Much data being produced in industry today is “flat data,” without relational structure—in manufacturing, an estimated 90 percent of data is flat.³⁹ Making this data usable requires new approaches that can efficiently handle large volumes of different types of data—for example, NoSQL and Hadoop technologies.

Firms also need to recognize potentially distinctive types of data that can create a competitive edge in AI-enabled product offerings. Customer sentiment and geo-locational real-time event data are examples of differentiating data for which competition for exclusive access is likely to intensify. Certain data may become valuable only if combined with other data sources in a larger ecosystem.

Given the rapidly increasing data output from sensors, machinery, and social networks, organizations face challenges in how to handle such massive streams of information. While some use cases for such data will be very concrete, with clear requirements, other potential uses will be fuzzy or not yet fully defined. Some use cases will require significant time series of data, while others may require real-time data. Companies will need to decide which data to store in their original granularity and which to aggregate or pre-analyze. With increasing data storage capacities in the cloud as well as more powerful “edge” computing capabilities close to sensors, flexibility increases rapidly.

Know what—and whom—you need

To really capture the performance-boosting value of AI, companies need to build internal capabilities and partner with or acquire additional know-how from AI startups or leading AI firms. Roles that companies often have to fill are “translators” and data scientists (“quants”). Translators bridge the gap between the techniques available to data scientists and the real-world problems of management. Quants design, develop, deploy, and train AI technologies.

³⁸ *Raising your Digital Quotient*, McKinsey Quarterly, June 2015.

³⁹ *Smartening up with artificial intelligence: How AI will transform Germany’s industrial sector*, Digital McKinsey, April 2017.

Finding the talent to fill these roles is challenging. Out of approximately 150 million workers in the United States, for example, only 235,000 are data scientists.⁴⁰ Given the scarcity of talent, it can be useful to partner to get the AI engine up and running while simultaneously developing and recruiting AI talent. One option for partnership is the “build, operate, transfer” model. Regardless of the approach, a key challenge is finding tools that are fit for the purpose. Off-the-shelf solutions can play an important role, but only if they are well aligned to company-specific needs and properly account for the costs and time of integration.

As with other digital technologies, a fast, agile, “test and learn” approach is important. Small, fast steps ensure the right focus, for example, through simulation-based pilots that allow companies to quickly test the impact estimates in the business case. Best-practice companies set up cross-functional AI task forces that are able to prototype a solution within weeks (where data is available), test it with the business units, and decide how to proceed. Two-speed approaches can help, pushing ahead on newer, more flexible IT architecture while gradually migrating legacy systems.

Internal and external collaboration are also important for agile companies. Internally, teamwork and collaboration are especially important for digital technologies like AI, which often cuts across traditionally siloed parts of organizations, from customer service to fulfillment to supply-chain management to financial reporting. Agile companies also collaborate beyond the boundaries of the firm, tapping into broader networks of learning and innovation, and bringing suppliers and other partners along for the ride. This requires digital leaders to recognize what they’re good at themselves and what others might do better, and to improve their ability to work collaboratively with people and institutions.

Integrate to capture the benefits

Once companies’ capabilities are producing AI-powered insights, they must be integrated to capture the benefits promised in the business case. Typically, this involves process redesign to incorporate AI insights into the workflow. In some cases, this will involve automation. In others, it requires getting the right data insights into the hands of the right personnel within the organization. In either case, optimizing the human-machine interface is critical. Careful thought needs to be given to determine what tasks the machine automates and how the hand-off with humans will work (for example, between virtual agents and advisers). In many cases, the change-management challenges of transforming what many people do within an organization are greater than many of the technical challenges of implementing AI.

Adopt an open culture and reskill the workforce

To get the most out of AI in the long run, an organizational culture open to the collaboration of humans and machines is required. Trust is a key enabler here. Due to the interplay of training and inference in AI, the relationship between a machine’s inner workings and the results it produces can become rather obscure. Instead of an algorithm’s predetermined steps, it is, in a sense, the data used to train it that leads to a certain outcome. Humans will need some time to adjust to this paradigm shift. That means the creation of an AI-ready culture should be a priority early on. It may also require investment to build the capabilities of workers, especially mid-level managers, to understand how to use data-driven AI insights—and to trust them as the basis for making decisions.

Companies should be forthright in addressing employee concerns if they expect them to help develop and implement AI tools that will change business processes and potentially automate some activities. Some workers will have to be retrained to work alongside AI-directed machines, while others will have to be redeployed within the company or elsewhere in the economy; businesses have a vital role to play in aiding these transitions. This will

⁴⁰ *The age of analytics: Competing in a data-driven world*, McKinsey Global Institute, December 2016.

require changes in skills, mindsets, and culture as we transition into a world where “co-workers” include machines as well as other people.

As with all cultural and organizational changes, leadership is critical, a point reinforced by our survey findings. Strong leadership support goes hand in hand with stronger AI adoption. Respondents from firms that have successfully deployed an AI technology at scale tend to report C-suite support that is nearly twice as high as those from companies that have not adopted any AI technology.

AI PROVIDER COMPANIES NEED TO MATCH TECHNOLOGICAL SOLUTIONS TO REAL-WORLD BUSINESS PROBLEMS

AI providers must go beyond offering interesting AI products to offering AI solutions for real-world business problems. Our survey highlights the scale of this challenge in the market today. Of the more than 3,000 firms we surveyed, 41 percent reported that one of the biggest barriers preventing them from further adoption of AI was the uncertain return on investment, while 26 percent reported a lack of relevant AI products on the market.

To close that gap, AI providers need to develop sector-specific expertise to understand where major business opportunities are and what potential sector-specific solutions would look like. This requires AI developers and engineers to become more familiar with potential customers’ value chains and data systems. It also requires more emphasis on demonstrating the technologies and applications, and the benefits they bring, at scale.

Meeting the market will also require AI developers to address discomfort with human-machine interactions. In our survey, one-fifth of respondents said their customers were not interested in AI-enabled products and services, and the figure was significantly higher in health care and other sectors where the interactions are particularly intimate. Machines can frustrate people if they cannot understand what the people are saying, deliver inaccurate or inappropriate answers, or simply do not live up to promises made about them.

There are no easy answers to this evolving problem. AI providers must take the lead in overcoming this hurdle, as when Amazon reduced the error rate for its intelligent personal assistant, Alexa, by a factor of two between 2014 and 2016.⁴¹ Google took a different approach by managing customer expectations for what to expect from its voice search feature, making clear that it was not an intelligent personal assistant, in order to avoid some of the customer disappointment expressed about Apple’s Siri.⁴²

GOVERNMENTS SHOULD TAKE THE LEAD IN TACKLING CROSS-CUTTING AI CHALLENGES

AI’s impacts extend well beyond firms. There are serious implications for governments, the workforce, and broader society. Indeed, the importance is highlighted by the growing number of governments crafting national AI development plans— in the past 18 months, the United States, the United Kingdom, China, and South Korea have all issued national strategic plans with significant AI dimensions, in some cases backed up by billions of dollars of AI-specific funding initiatives. MGI has also looked at country-specific implications of AI (see Box 3, “China’s path forward”).

⁴¹ Jordan Novet, “Amazon has reduced Alexa’s mistakes in completing tasks by a factor of 2,” *Venturebeat.com*, July 13, 2016.

⁴² Farhad Manjoo, “Siri is a gimmick and a tease but Google voice search is getting close to fulfilling Apple’s broken promise,” *Slate.com*, November 15, 2012.

Box 3. China's path forward

While many of the challenges of AI are global in nature, the implications for specific governments vary across countries. In April 2017, McKinsey Global Institute published a discussion paper assessing the specific priorities for the Chinese government.¹

- **Strategic priority 1: Build a robust data ecosystem.** China can move to set and implement data standards, open public-sector data for private exploration, and encourage international exchange of data streams.
- **Strategic priority 2: Broaden adoption of AI within traditional industries.** Adoption of AI by Chinese firms is being held back by a lack of strategic awareness, the costs of adoption, and a shortage of technical know-how. The government should look to address these barriers using tax credit and subsidy tools, as well as pioneering adoption within the government.
- **Strategic priority 3: Strengthen the pipeline of specialized AI talent.** To address China's AI talent gap, the government needs to invest in AI-related education and research programs, reorient the education system for a greater focus on innovation and digital skills, and devise an immigration policy to attract the best global talent.
- **Strategic priority 4: Ensure the education and training systems are prepared to develop technology skills and retain large segments of the workforce.** The government must proactively identify the jobs that are most likely to be automated and ensure that retraining programs are made available to the segments of the labor force whose livelihoods are at risk. These efforts could involve collaborating closely with vocational training schools and providing educational vouchers to workers.
- **Strategic priority 5: Establish an ethical and legal consensus among Chinese citizens and in the global community.** The Chinese legislature will need to provide a framework to deal with the current legal uncertainties around AI. The government should also look to establish a regulatory body to monitor and regulate AI activities. China could also take the lead in forming an international governing body to promote peaceful, inclusive, and sustainable development of AI technology.

¹ *Artificial intelligence: Implications for China*, McKinsey Global Institute, April 2017.

The range of issues confronting governments is vast, ranging from ethical and equity concerns to the setting of data standards (see Box 4, “An overview of ongoing challenges”). In this report, we explore two particularly pressing issues—reskilling of the workforce and supporting local AI industries. These are not issues for government to solve alone. The Partnership on AI is one effort at helping to ensure progress, bringing together civil society organizations such as the American Civil Liberties Union, tech leaders such as Amazon, Apple, Facebook, Google, and Microsoft, and other partners from the private sector, including McKinsey & Company.

Box 4. An overview of ongoing challenges

AI presents a wide range of issues for government and society. We outline some of these issues below, along with some ideas about how to address them. Progress on these issues will be critical both to realizing the potential benefits of AI and to safeguarding people from the risks.

- **Encourage broader uptake of AI.** Adoption of AI tends to be concentrated in relatively digitized industries and, within those industries, in the firms that are already at the digital frontier. Broader adoption of AI and supporting digital technologies, especially in smaller firms, could be important to supporting productivity growth across the economy and to healthy, competitive markets. Broader adoption could also support more equal wage growth. AI can drive improvements in labor productivity, which in turn drives higher wages. Broadening adoption will ensure that the productivity benefits of AI are spread across more firms and workers, rather than pushing up the wages of only those in frontier firms, who are already located near the top of the income distribution.
- **Address employment and income-distribution concerns.** AI-powered automation could have a profound impact on jobs and wages. In our survey, the majority of firms did not expect AI to significantly reduce the size of their workforce. However, there might still be people whose skills and capabilities are mismatched to the work that needs doing. Governments may have to rethink models of social support. Various ideas are under consideration, including work sharing, negative income taxes, and universal basic income.¹
- **Resolve ethical, legal, and regulatory issues.** AI presents a range of ethical, legal, and regulatory issues. Real-world biases risk being embedded into training data. Since the real world is racist, sexist, and biased in many other ways, real-world data that feeds algorithms will also have these features—and when AI algorithms learn from biased training data, they internalize the biases, exacerbating those problems.² There are also concerns about the algorithms themselves—whose ethical guidelines will be encoded into them, what rights should people have to understand the decision-making process, and who will be responsible for their conclusions? This has led to calls for algorithmic transparency and accountability.³ Privacy is likewise a concern—who should have ownership of data, and what safeguards are needed to protect highly sensitive data, such as health-care data, without destroying its usefulness? Organizations leading efforts to tackle these questions include the Partnership on AI, OpenAI, the Foundation for Responsible Robotics, and the Ethics and Governance of Artificial Intelligence Fund.
- **Ensure the availability of training data.** An abundance of data is critical for AI training systems. Opening up public-sector data can spur private-sector innovation. Setting common data standards can also help. In the United States, for example, the Securities and Exchange Commission mandated in 2009 that all public companies must disclose their financial statements in XBRL (extensible business reporting language) format, thereby ensuring that public data is machine readable.
- **Deploy AI within government.** AI has tremendous potential benefits for the public sector as well as the private sector. Improved planning, targeting, and personalization could deliver a much-needed step change in both the efficiency and effectiveness of government services.⁴ In our sector deep dives (see Appendix A), we explore what an AI-powered future could look like for two particular areas of government—health care and education.

¹ *A future that works: Automation, employment, and productivity*, McKinsey Global Institute, January 2017.

² *The age of analytics: Competing in a data-driven world*, McKinsey Global Institute, December 2016.

³ See, for example, *Statement on algorithmic transparency and accountability*, Association for Computing Machinery US Public Policy Council, January 2017.

⁴ For further discussion of digital technology and analytics on government productivity, see *Government productivity: Unlocking the \$3.5 trillion opportunity*, McKinsey Center for Government, April 2017.

Prepare the workforce for a lifetime of reskilling


Artificial intelligence tools have the promise to change our lives as fundamentally as personal computers did a generation ago. But, as we saw then, gains will be accompanied by losses.

In some cases, full or partial automation from AI will displace labor. In *A future that works*, we show that almost half of all work activities have the potential to be automated by adapting currently proven technology. We estimate that 60 percent of all occupations have at least 30 percent technically automatable activities (Exhibit 9). However, automation will change far more occupations—by, for example, partially automating them—than it will replace.⁴³

Exhibit 9

While few occupations are fully automatable, 60 percent of all occupations have at least 30 percent technically automatable activities

Automation potential based on demonstrated technology of occupation titles in the United States (cumulative)¹



¹ We define automation potential according to the work activities that can be automated by adapting currently demonstrated technology.

SOURCE: US Bureau of Labor Statistics; *A future that works: Automation, employment and productivity*, McKinsey Global Institute, January 2017; McKinsey Global Institute analysis

Not all AI innovations will displace labor. Many AI applications target non-labor cost savings, as when an AI algorithm reduces a factory's energy use or when AI is used in predictive maintenance. Indeed, less than a fifth of AI adopters in our survey said their primary driver for adoption was to reduce labor costs. Improving capital efficiency was cited more often, as were revenue-focused drivers such as enhancing their product offering.

Although one of the near-term benefits of automation is labor cost reduction, our survey shows 24 percent of firms that have adopted AI at scale expect to increase the size of their workforce in response to AI, anticipating growth in their business activities. Moreover, 82 percent of AI-aware firms do not expect to significantly reduce the size of their workforce,

⁴³ *A future that works: Automation, employment, and productivity*, McKinsey Global Institute, January 2017.

owing in part to increased demand for new skills, such as those from data scientists and data translators.

The implication of these changes is clear: companies need to update the skills in their workforces, and individuals need to acquire skills that work with, not compete against, machines.⁴⁴ For people already in the workforce, reskilling will be essential. Much of this reskilling can occur on the job through stronger professional development programs. For people transitioning between jobs, vocational and adult education programs must be strengthened. These work best when they are short, affordable, and closely linked to the job market. Nanodegree programs are one recent innovation in this space. Governments and training institutions also need to do better to ensure alignment of training programs with the jobs of the future. The World Economic Forum argues that one bet governments should look to make is on the care economy. AI-powered predictive analysis could also be put to use to anticipate future pockets of skill shortages and oversupply.

For students, a future-focused curriculum is a necessity. The World Economic Forum identified 16 skills that are needed in the 21st century—including creativity, collaboration, initiative, and adaptability—but are not included in standard curricula.⁴⁵ People are also questioning the “learn then earn” model, wondering if lengthy degree programs still make sense in a world of fast-changing jobs. Instead, there are calls for a new deal on lifelong learning.

Building local industries

Governments serious about establishing their cities or countries as hubs for AI development will need to join the global competition to attract AI talent and investment.

Our global review finds that AI investment is highly concentrated geographically: in 2016, the United States absorbed around 66 percent of external investment (VC, PE, and M&A activity). China was a distant second, at 17 percent, but it is growing fast. Unsurprisingly, the San Francisco Bay Area and Silicon Valley is the biggest and most important ecosystem, attracting around 40 percent of global external investment in 2016 (Exhibit 10). Other US cities, such as Boston and New York, have also drawn AI investment, and are followed by two strong China AI tech ecosystems: Beijing and Shenzhen. In Europe the only clear and strong ecosystem is currently London. Other vibrant European AI ecosystems may appear within the next few years, given exceptional startup activity and VC investment in 2016, especially in Germany, France, and the Nordic region.⁴⁶

Many of the most important hubs for external investment are also key locations for tech giants’ AI research and development work. However, we also see a pattern of big tech investing in new areas to scout and attract local talent, including Google at the University of Montreal, Intel at Georgia Tech, NVIDIA at the National Taiwan University, and Facebook in Paris.

These hubs benefit not only from the creation of highly skilled, highly paid jobs, but also, critically, knowledge and innovation spillovers. Employees become AI entrepreneurs, AI-savvy workers switch between companies, and innovative AI products can be developed for and deployed in the local market. Indeed, our survey suggests that countries that lead the way in AI investment and innovation—the United States and China—also lead the way in AI adoption.

⁴⁴ *Digital America: A tale of the haves and have-mores*, McKinsey Global Institute, December 2015.

⁴⁵ Jenny Soffel, “What are the 21st-century skills every student needs?” World Economic Forum, March 10, 2016.

⁴⁶ Galina Degtyareva, “European AI startups landscape,” Medium.com, March 21, 2017.

The United States and China dominate the AI landscape, with Europe falling behind

The most vibrant AI hubs ...

Silicon Valley


- Top global hub for startups
 - 12,700–15,600 active startups
 - 2 million tech workers
- Global leader for VC investment
- Headquarters of many top high-tech companies

New York

- Leading hub for financial and media industries
- AI talent pipeline from universities such as Cornell
- Strong funding ecosystem—second in the world after Silicon Valley for the absolute number of early-stage investments

Beijing

- Leading in volume of academic research output in AI coming from Tsinghua, Beihang, and Peking universities
- Extensive involvement of tech leaders, especially Baidu
- AI identified as a strategically important technology by the Chinese government



Boston

- Long history of cooperation between science and industry
- World-class universities such as MIT developing advanced technologies and providing a talent pipeline

London


- Global finance center, supporting both investment and fin-tech applications
- European leader of VC startup investment
- Presence of top high-tech companies
- Talent pipeline and research expertise from universities such as University of Cambridge, Imperial College, and Oxford

Shenzhen

- Hub for electronics manufacturing firms such as Huawei and ZTE
- Strong expertise in hardware
- AI identified as a strategically important technology by the Chinese government

... and the external investment behind their growth¹

\$ billion (estimate)



¹ Estimates consist of annual VC investment in AI-focused companies, PE investment in AI-related companies, and M&A done by corporations. Includes only disclosed data available in databases, and assumes that all registered deals were completed within the year the transactions were announced.

SOURCE: Capital IQ; Pitchbook; Dealogic; S&P; McKinsey Global Institute analysis

There are opportunities for governments to act on multiple fronts. Thoughtful incentives to attract both investment and talent are helpful—tax breaks for AI entrepreneur immigrants and special tech visa quotas are two approaches governments have adopted. The United Kingdom, for example, has established the Tech Nation Visa Scheme, which awards up to 200 visas annually, without work sponsorship requirements, for applicants with exceptional talent or promise in the digital space.⁴⁷

Funding for leading-edge science programs is also important, be it through grants to universities, creation of government laboratories, or joint research initiatives with the private sector. US federal government investment into unclassified AI R&D topped \$1 billion in 2015.⁴⁸ The South Korean government has announced it is investing 1 trillion won (\$900 million) to build a public-private AI research center jointly with leading Korean conglomerates.⁴⁹ Meanwhile, China's National Development and Research Commission has launched a national engineering laboratory for deep learning, led by Baidu.⁵⁰

Governments also have other important levers to foster and consolidate AI hubs. They can act as lead customers, providing an important boost to startups by giving them a credible reference case. They can ensure that local regulations are AI-friendly, providing space for new technologies to be safely piloted and for greater clarity around legal liability and data-ownership issues. As mentioned above, they can also improve the availability of training data, both by opening up their own data banks and by setting data standards. These actions can help secure an environment that is uniquely attractive to AI investors and entrepreneurs alike.



After decades of hopes and disappointment, AI is back and could be set to drive profound changes in the global economy. Investment has been growing fast since 2013, with tech giants making huge plays on AI technology development and deploying it across their businesses. We are already seeing examples of real-life business benefits for early-adopting firms.

However, adoption of AI technologies remains largely at an experimental stage. Indeed, the gap between early adopters and the rest is set to grow. While many companies have yet to be convinced of AI's benefits, frontier firms are charging ahead. Early AI adopters tend to be larger firms with mature digital strategies that display both deep and broad adoption patterns. Their focus on AI is driven by the desire to increase revenue and reduce cost, and they have momentum and support for AI from the top of the company.

Significant gains are there for the taking. For many companies, this means accelerating the digital transformation journey. They will have to put the right digital assets and skills in place to be able to effectively deploy AI.

For governments, there is an urgent need to support both firms and citizens to ensure that they benefit from the AI-powered digital economy. This means preparing workers for rapidly changing job requirements, facilitating AI investment and adoption, and solving for broad-based, productivity-driven economic growth.

⁴⁷ "Tech Nation Visa Scheme: High demand as scheme enters fourth year," Tech City UK blog, April 13, 2017.

⁴⁸ National Science and Technology Council, "Preparing for the future of artificial intelligence," Executive Office of the President, October 2016.

⁴⁹ Mark Zastrow, "South Korea trumpets \$860-million AI fund after AlphaGo 'shock,'" *Nature*, March 18, 2016.

⁵⁰ Weining Hu, "How China is becoming a world leader in artificial intelligence," China Briefing blog, March 14, 2017.

APPENDIX A: FIVE CASE STUDIES

RETAIL

AI CAN HELP PREDICT DEMAND, AUTOMATE OPERATIONS, AND DELIVER A BETTER SHOPPING EXPERIENCE

Retailers are already beginning to apply AI, machine learning, and robotics in major parts of the value chain. Most importantly, AI technologies could eliminate many levels of manual activities in areas such as promotions, assortments, and supply chain. The three areas of greatest opportunity in the short to medium term are promotions, assortment, and replenishment. Major retailers are experimenting with AI in all of these areas.

“Digital native” e-commerce companies are leading the way, using AI to predict trends, optimize warehousing and logistics, set prices, and personalize promotions. Some even aim at the full anticipation of customers’ orders, shipping goods without waiting for a purchase confirmation.

The use of AI in retail can generate several benefits. First, it helps people make smarter decisions, with more accurate and real-time forecasting. Good forecasts help improve supply management, define impactful thematic promotions, and optimize assortment and pricing. Second, AI can make operations more efficient, thanks to a combination of robotics and process optimization that enhances productivity and reduces manual labor costs. AI will enable retailers to increase both the number of customers and the average amount they spend by creating personal and convenient shopping experiences.

Can artificial intelligence help traditional, non-digital retailers catch up, or will it further widen the divide between the agile and data-driven internet pure players and historical brands that are lagging behind? Success will hinge on the capacity of retailers to get on board and secure access to strategic data while reinventing the shopping experience. But before we delve into the conditions of full achievement, let’s explore what the future could look like in 2030.

Retailers can know more about what shoppers want—sometimes before shoppers themselves

In the future, artificial intelligence could help forecast and automate retailers’ decision making in real time. By identifying and learning from patterns in large volumes of data, spanning many disparate sources—previous transactions, weather forecasts, social media trends, shopping patterns, online viewing history, facial expression analysis, seasonal shopping patterns—AI can help companies adjust to and master an increasingly dynamic market environment. By improving forecasting accuracy, machine learning and computer vision can help better anticipate consumer expectations while optimizing and automating supplier negotiations.

The impact of AI-enabled forecasting is already being demonstrated. For instance, a European retailer was able to improve its earnings before interest and taxes (EBIT) by 1 to 2 percent by using a machine learning algorithm to anticipate fruit and vegetable sales. The company automatically orders more produce based on this forecast to maximize turnover and minimize waste. Similarly, German e-commerce merchant Otto has cut surplus stock by 20 percent and reduced product returns by more than two million items a year, using deep learning to analyze billions of transactions and predict what customers will buy before they place an order. The system is 90 percent accurate in forecasting what the firm will sell over

the next 30 days, so Otto allows it to order 200,000 items a month from vendors with no human intervention.

AI technologies can also help retailers predict future store performance when expanding their physical footprints. As more sales migrate online, non-digital store sales per square meter are declining. In the United Kingdom, retailers would need to shave space by more than 20 percent to return to 2010 sales densities in real terms if all other factors were held equal. It has become crucial for retailers to optimize store space and location. One Japanese retailer applied machine learning to understand profitability drivers when picking the location of a new concept store.

Operations are ripe for automation

Warehousing and store operations present a rich set of optimization opportunities for AI application. For some non-digital retailers, notably supermarkets, automation of operations would actually make an existential difference. Many supermarkets offer online sales and home delivery to match online grocers but still carry the full cost of physical stores, so the costs of online service—in the United Kingdom, roughly £5 to pick items off the shelves for an £80 order and £8 to deliver them—wipe out the industry's 2 percent profit margin many times over.

Autonomous robots can work alongside people to increase productivity and reduce injuries. Swisslog has reduced stocking time by 30 percent since it began using autonomous guided vehicles in its warehouses. DHL unleashed a pair of fully automated trolleys last year that follow pickers through the warehouse and relieve them of physical work.

In store, machine learning can help optimize merchandising, with opportunities to improve assortment efficiency by 50 percent. A retailer was able to generate a sales uplift of 4 to 6 percent by using geo-spatial modeling to determine micromarket attractiveness and leveraging statistical modeling to predict and minimize running out of stock. With machine learning, these efficiencies would be realized in real time and would gain in accuracy as they learn from new data.

Ocado, a UK online supermarket, is one company that has embedded AI at the core of its operations. In the retailer's warehouse, machine learning algorithms steer thousands of products over a maze of conveyor belts and deliver them to humans just in time to fill shopping bags. Other robots whisk bags to delivery vans whose drivers are guided by an AI application that picks the best route based on weather and traffic conditions.

Retailers are getting personal

Empowered by the ease, economy, and immediacy of online shopping, many consumers already expect personalized, immediate, pitch-perfect help. In the future, AI will be invaluable to marketers trying to reach hyperconnected consumers who continuously redefine value by comparing prices online—even, and particularly, when browsing in a non-digital store. Smartphone penetration necessitates an omni-channel strategy, and AI can help optimize, update, and tailor it to each shopper in real time. Insights-based selling, including personalized promotions, optimized assortment, and tailored displays, could increase sales by 1 to 5 percent. Online, this kind of personalization, combined with dynamic pricing, can lead to a 30 percent growth in sales.

The pure internet players are several steps ahead in targeted marketing, thanks to data gathered online. Traditional retailers need to start gaining access to data assets to compete. Carrefour, the global retailer based in France, and Target in the United States have both deployed electronic beacons in stores to collect data about customer behaviors and purchasing patterns, and they use machine learning algorithms to determine which

personalized promotions to send customers as they shop. Carrefour reported a 600 percent increase in app users after it deployed beacons in just 28 stores.⁵¹

As natural language processing develops, AI-enabled personalization could go far beyond the realm of targeted promotions. In store, virtual assistants could identify repeat customers using facial recognition, analyze their shopping history to make suggestions, and communicate in a conversational way using natural-language processing and generation. In the meantime, online retailers are trying to give more of a human touch to the Web and make personal recommendations to shoppers. Stitch Fix, an online personal-shopper service, has an algorithm that reviews the images clients display on Pinterest to better understand their styles, even if they have a difficult time articulating it online. Online retailers also use smart agents to understand shoppers' needs. One example is floral retailer 1-800-Flowers' digital gift concierge, powered by machine learning and language recognition, which proposes a selection of products based on a chat with the shopper.

Bringing it all back home

Enhanced user experience is the area that offers probably the most futuristic perspectives for AI in retail. Deep learning and computer vision technologies also will help store owners compete with the one-click convenience of online retailers by eliminating checkout altogether. Amazon Go, an experimental grocery in Seattle, allows shoppers to take goods off shelves and leave without seeing a cashier or stopping at a self-checkout kiosk. Computer vision identifies them as they enter the store, then links them with products taken from shelves. When shoppers leave, the system deducts the cost of the items in their bag from their Amazon accounts and sends an email receipt.


At home, virtual assistants further push the convenience boundaries. In the future, they could alert users that they're about to run out of a product and suggest buying more. Google's smart speaker service, Google Home, allows shoppers to complete orders with 50 participating Google Express retailers, such as Costco, Whole Foods, and PetSmart, while Amazon's Alexa has partnerships with more than 100 third-party services. Recent developments in smart home assistants pave the way for a significant shopping disruption, where computer vision helps identify desired goods by taking a picture or the assistant identifies preference patterns from images and videos liked online by consumers. Amazon's new Echo Look device, for instance, which was unveiled in April 2017, incorporates a camera into Alexa's virtual assistant function and recommends styles based on the user's wardrobe and body shape, combining machine learning and computer vision.

In the future, artificial intelligence technologies could also be used at scale to deliver goods minutes after purchase. Most of today's efforts—by players as big as Amazon and as small as the Reno, Nevada, startup Flirtey—focus on unpiloted aerial drones. Flirtey made its first delivery to a private residence, a box of snacks from a local convenience store, in July 2016. In Europe, an Estonian startup, Starship Technologies, has taken a different path: six-wheeled delivery robots that pitter along city pavements at 4 miles per hour.⁵² The ramp-up of drones and robots hinges on the application of deep learning technology to allow innovative problem solving and decision making, and on airspace regulation, as civil aviation authorities question the operating of drones over populated areas and close to the flight paths of piloted aircraft.


⁵¹ Shubhi Mittal, "25 retailers nailing it with their proximity marketing campaigns," Beaconstac.com, February 11, 2016.

⁵² "Starship Technologies launches testing program for self-driving delivery robots with major industry partners," Starship Technologies press release, July 6, 2016.

Retailers can know more about what shoppers want—sometimes before shoppers themselves




Machine learning personalizes promotions to match shoppers' profiles; in-store beacons send offers to their smartphones as they browse through the store





Autonomous drones using deep learning technology complete last-mile delivery, and are able to handle obstacles or absent recipients


Facial recognition software, machine learning, and natural language enable virtual agents to greet you personally, anticipate orders, and provide directions





Computer vision with deep learning identifies articles bagged by shoppers; adding data from sensors, AI allows non-stop checkout and automatic payment




Interactive screens and tablespots enabled with computer vision and deep learning can identify articles and recommend complementary products and uses that fit shoppers' lifestyle profile



Stores update and optimize prices in real time, with machine learning leveraging data on competitors' prices, weather, and inventory levels to maximize revenues



An autonomous shopping cart follows you in the store, and can find its way to your vehicle or to a robot or drone for home delivery



AI-enhanced robots continuously track inventory, recognize empty shelves, and replenish them; other robots fill bags in the warehouse

Moving forward

Keeping up with the competition in this new realm of retailing will be as hard as it is important. Retailers would be wise to invest in data gathering up and down their supply chain to seek a competitive advantage. This implies shifting to a collaborative mindset to allow insights to flow through the entire value chain.

On the one hand, partnerships between retailers and their suppliers will become important to improve supply chains and marketing, optimize pricing, and achieve efficient merchandising. Major retailers are already experimenting with data integration. For instance, Walmart has started sharing real-time data with major consumer packaged goods players using a data lake.

On the other hand, cross-industry partnerships between retailers and other players will evolve to create better customer insights. Ecosystems will emerge with third-party legal entities that bring together shareholders such as retailers, loyalty card providers, and banks. The trend is already under way in Brazil, Turkey, Thailand, Indonesia, the United Kingdom, and other countries.

Computer vision technology in physical retail outlets can give tremendous insight through customer footfall data, customer reactions, and promotions (for example, how long shoppers stand in front of a promotion and whether they take the product), and can enable automatic checkout. Natural language processing and deep learning will facilitate the rise of virtual assistants by making dialogue more intuitive and improving it faster.

However, the rise of virtual assistants that take orders directly from shoppers may someday lead to the disintermediation of retailers. With automation, retailers also will need to rethink their sales teams' skills, as human interaction will be increasingly reserved for higher engagement types, involving emotional intelligence, excellent product knowledge, and brand championing. The retailing revolution has room to run.

ELECTRIC UTILITY

AI CAN MAKE THE SMART GRID SMARTER AND REDUCE THE NEED FOR UTILITIES TO ADD POWER PLANTS

The electric utilities sector has great potential to embrace artificial intelligence in the coming years. At every step of the value chain, from power generation to end consumers, opportunities for machine learning, robotics, and decision-making automation exist that could help electric utilities better predict supply and demand, balance the grid in real time, reduce downtime, maximize yield, and improve end-users' experience.

The increasing deployment of renewable energy sources has introduced significant volatility in energy supply, with variation of up to 60 percent. Demand also fluctuates dramatically and frequently by time and by region, with weather and events such as the Super Bowl creating demand peaks sometimes for period of less than an hour. Following an investment wave during the 1970s and 1980s, transmission and distribution companies, under financial constraints and regulatory requirements to hold down costs and rates, put far less money into improving their networks over the following decades. As a result, current energy grids are ill-equipped to smooth out these spikes, and excess power is regularly lost at top dollar.

The increasingly complex network of stakeholders and assets, combined with aging critical capital assets, highly uncertain demand and supply, non-linear power loads, cost pressures, and price deregulation, is engineering real momentum for artificial intelligence and robotics.

Ultra-accurate forecasts can make possible the integration of both additional renewable energy and excess power into the grid

One of the most acute cases for AI in electricity is around demand and supply prediction. An inaccurate load forecast in the power grid can affect many stakeholders: power generation determines which power sources should be allocated for the next 24 hours based on one- to two-day forecasts; transmission grids assign resources based on power transmission requirements; and electricity retailers calculate energy prices based on estimated demand.

Electric utilities are starting to explore artificial intelligence technologies to produce more accurate short-term load forecasts. DeepMind, the AI startup bought by Google in 2014, is currently working with National Grid to predict supply and demand peaks in the United Kingdom by using weather-related variables and smart meters as exogenous inputs, hoping to cut national energy usage by 10 percent and maximize the use of renewable power despite its intermittence.

In the future, machine and deep learning technologies could forecast demand and supply in real time and optimize load dispatch, thereby saving energy and cost. For a network that experiences demand ranges between 10 and 18 gigawatts, savings could reach 100 megawatts over periods of one to four hours per day.⁵³ More reliable forecasts would allow utilities to delay or even avoid ramping up a fossil-fuel-powered station. It would also offer cost-effective alternatives to operators, who currently consider building new plants to absorb seemingly impossible variability.

Grid modernization and deployment of smart meters are already under way in most countries, aiming at creating a more dynamic matching of supply and demand; the use of artificial intelligence enables suppliers to better predict and optimize load dispatch. Smart grid initiatives allow small, private energy producers (even individual homeowners) to sell excess capacity back to the grid. The technology is developing quickly: the United

⁵³ Jaime Buitrago and Shihab Asfour, "Short-term forecasting of electric loads using nonlinear autoregressive artificial neural networks with exogenous vector inputs," *Energies*, volume 10, number 1, 2017.

States alone has committed more than \$9 billion in public and private funds since 2010.⁵⁴ In Europe, Sweden and Italy have replaced nearly all meters with smart meters; other European countries are on track to finish the conversion within 10 years. In 2011, China's State Grid Corporation announced plans to invest \$45 billion in smart grid technologies and another \$45 billion between 2016 and 2020.⁵⁵

AI could also help utilities assess the reliability of new small supply players, such as households, by predicting the lifetime of their storage units and their suitability for integration in a power storage scheme. Demand-side response aggregators currently focus only on larger firms due to the cost of scanning hundreds of small participants. In the future, the grid could become a marketplace where grid operators bid on power offered for sale by a large number of small players with electricity from sources such as electric-car batteries and rooftop solar cells. Machine learning could help automatically assess this large volume of minisuppliers, who, collectively, could be instrumental in addressing demand peaks.

Yield optimization, predictive outage, and preventive maintenance can help better plan the grid

The other lever where AI and robotics could help reduce costs is operations, from power generation to transmission and distribution.

With AI, power providers could maximize their generation efficiency with real-time adjustments across assets. For instance, machine learning can help optimize wind turbines' yield based on their own past performance, real-time communication with other wind farms, the grid status, and changes in wind speed and direction. GE Renewables recently introduced a "digital wind farms" concept, which optimizes yields with machine learning applied to turbine sensors data, and modular turbines that can be customized to conditions at each installation site.⁵⁶ GE says the technology could boost a wind farm's energy production by as much as 20 percent and create \$100 million in extra value over the lifetime of a 100-megawatt farm.

Power generation yield can also be bolstered by reducing downtime and improving preventive maintenance. To date, preventive maintenance efforts have had a limited impact because firms can be overwhelmed by the sheer volume of sensor data and inaccurate alerts. This is an opportunity for AI technologies, which thrive on mountains of information. Advanced analytics already demonstrate the benefit of intelligent maintenance. Some coal power plants, for instance, were able predict the timing of failures within one week six to nine months in advance, with 74 percent accuracy. Overall, we estimate that optimizing preventive maintenance, automating fault prediction, and increasing capital productivity through AI applications could increase power generation earnings before interest, taxes, depreciation, and amortization (EBITDA) by 10 to 20 percent.

Machine learning can be instrumental in reducing energy losses in transmission and distribution. Looking across multiple assets on a customer site, AI tools can make intelligent, real-time decisions to balance the grid, such as rescheduling the consumption of power assets within a building in synchrony and leveraging the inherent flexibility of air conditioning and refrigeration systems to maximize the benefit of solar panels and batteries installed on site. Upside Energy, a startup, received a UK government grant in 2017 to use AI to manage a portfolio of batteries and other storage assets and provide real-time energy reserves to the grid. The company estimates that machine learning could be used to help unlock up to

⁵⁴ *2014 smart grid system report*, US Department of Energy, August 2014.

⁵⁵ Melanie Hart, "China pours money into smart grid technology," Center for American Progress, October 24, 2011.

⁵⁶ "GE launches the next evolution of wind energy making renewables more efficient, economic: The digital wind farm," General Electric press release, May 19, 2015.

six gigawatts of demand-side flexibility that can be shifted during the evening peak without affecting end users.⁵⁷

By shaving off demand peaks during the day, utilities could postpone or even forgo the need to add generating capacity that would be required—and would generate revenue—only for short periods.

Transmission and distribution firms also could shift from time-based maintenance to condition-based maintenance. Leveraging data from sensors, communications devices, and other hardware that tracks and controls objects remotely, machine learning applications could liberate grid operators from decommissioning assets before their useful lives have ended, while enabling them to perform more frequent inspections and maintenance to keep assets working well. A European power distribution company was able to reduce its cash costs by 30 percent over five years by analyzing 20 variables to determine the overall health of power transformers and diagnose the condition of individual components. The company could predict parts failures more accurately and prioritize repairs based on which faults cause the most disruption; it can even identify faults that do not yet have enough data to be spotted by today's systems.

In the future, one can imagine that operational trade-offs among several power stations or within the distribution network—such as outage planning refinement, replacement vs. run-to-failure decisions, and cutoff for spare delivery—will be automatically made by advanced analytics and machine learning algorithms. When necessary, inspections could be automatically scheduled, with machine learning algorithms making a judgment call on whether the deployment of drones and smart robots could be sufficient at first, or whether human intervention is needed right away. Drone-based surveillance can replace the time-intensive and risky manual inspection of turbines, for example. This reduces turbine downtime because inspections by drones or robots can happen while turbines remain running, helping keep costs down.

At last, AI tools can help tackle non-technical energy losses, such as electricity theft, which is a significant problem in some developing countries. In Brazil, electricity theft accounts for up to 40 percent of electricity distributed.⁵⁸ In Hungary, Eon, after realizing that consumption in one town was three times the value of the electricity invoiced, found that 50 percent of businesses inspected and more than 80 percent of households in the town were stealing electricity.⁵⁹ The utility started to use machine learning to narrow down the list of suspicious users for its private detectives to investigate, and it was able to reduce theft by 30 percent. Machine learning can analyze customer data, including usage patterns and payment history, and compare it to known irregular behavior.

By our reckoning, network operators that adopt the full suite of AI apps—including inspection automation, preventive maintenance, demand management, and theft detection—could raise their EBITDA by 20 to 30 percent.

⁵⁷ Michael Bironneau, "How AI is shaping the future of energy," *Clean Energy News*, February 21, 2017.

⁵⁸ Patrick Glauner, "Large-scale detection of non-technical losses in imbalanced data sets," PhD dissertation, University of Luxembourg, March 2016, and expert interviews.

⁵⁹ Kester Eddy, "Hungary's power thieves," *Financial Times*, March 25, 2011.

AI can unleash market forces on retail electricity markets

Machine learning applications have the power to tailor electricity prices based on the large volume of data currently being provided by the growing number of smart meters and other devices and sensors. In the future, if regulators allow dynamic tariffs, utilities could adopt machine learning–based dynamic pricing that would permit them to protect their margins and reduce customer churn while maximizing their assets. For instance, they could use time-of-day pricing to encourage customers to shift nonessential consumption to earlier in the morning or later in the evening, when demand is lower.

Energy retailers also would be able to use AI to create custom benefits such as low rates or extra service in order to hold on to their most valuable, high-volume clients. While price sensitivity is a key consideration in attracting new customers and reducing churn, machine learning can also help address another critical component for utilities' marketing strategies, that is, identifying which customers are the most profitable.

The emergence of smart grids around the world also creates an opportunity for AI to support energy trading, not only for utilities but also for “prosumers,” consumers who will be able to sell excess power back to grid operators. Data and analytics are transforming the way markets connect sellers and buyers for many products and services, and this holds particularly true for grid operators. “Hyperscale” digital platforms can have a notable impact as electricity demand and supply fluctuate frequently. AI can help produce better and faster matches. AI-powered hyperscale platforms could transform energy markets by enabling smart grids to deliver distributed energy from many small producers. In the Netherlands, some startups are using the peer-to-peer model to match individual households directly with small providers (such as farmers) who produce excess energy. Vandebrom, for instance, charges a fixed subscription fee to connect consumers with renewable energy providers; in 2016, this service provided electricity to about 80,000 Dutch households. Utilities could also inform their purchases and sales for trading activities on liquid, unpredictable over-the-counter markets or through more reliable power-purchasing agreements.

AI can transform the user experience with consumption tailoring and automation for more convenience

Energy consumers also can benefit from AI. Since the liberalization of energy retailing, new entrants have piled into the market. In Europe, customers can choose from more than 20 suppliers, many competing on price alone. AI can help understand consumption patterns, tailor the value proposition as well as consumption to the users' preferences, and limit the hurdles for switchers.

Machine learning can help consumers deal with the complex task of selecting their electricity supplier based on users' preferences in terms of pricing and energy generation type, as well as metering measurements. Lumator has developed software with Carnegie Mellon University in Pittsburgh, Pennsylvania, that scans the market for the most suitable electricity supply deal. Lumator claims it can save people between \$10 and \$30 a month on their bills. In the future, AI could automatically switch energy plans, without consulting consumers or interrupting service, as the best deals become available for that specific user's profile.


Consumers could also benefit from detailed real-time insights on their energy consumption. Machine learning applied to metering data could extract energy profiles of the home's largest appliances and see how much each device contributes to the electricity bill. Startup Bidgeley mines home-meter data and disaggregation as well as machine learning algorithms to isolate the consumption of energy-hungry appliances such as ovens and clothes dryers. The company claims that its customers who use this information save between 4 and 12 percent on their electricity bills.

AI can make the smart grid smarter and reduce the need for utilities to add power plants





Sensors and machine learning allow for by-the-minute adjustments to maximize generation efficiency by adjusting to changes in wind conditions, for example

Machine learning-enabled forecasting anticipates supply and demand peaks, and maximizes the use of intermittent renewable power




Smart wires combine with machine learning to enable real-time power dispatching, and optimize it to current grid load and to buildings' asset portfolios


Drones and insect-size robots identify defects, predict failures, and inspect assets without interrupting production



Few technicians remain, but they spend more time on problem solving; in place of logging inspection status by hand, documents are automatically logged and routed



Virtual agents automate call centers, and automatically segment consumers based on service history; machine learning offers early warning of bad debts



Field workforce receives real-time updates to decrease response times and reduce the impact of outages



Smart-meter data and machine learning enable utilities to offer services based on usage, weather and other factors

Users could also benefit from tailored consumption management. AI-enabled devices can identify heating usage patterns and learn about inhabitants' preferences to tailor heating regulation. For instance, Google-owned Nest's Wi-Fi-enabled thermostat creates a personalized "heating schedule" by monitoring users' habits with a motion sensor to detect when the home is empty and doesn't need to be heated. The hope is that eventually users will never have to change the temperature by hand.

Another area of user experience enhancement that AI can help with is customer service. Virtual agents offer the opportunity to reduce customer service costs. Some retail utilities are already testing virtual agents to answer consumers' queries and provide instant help, using natural language to understand consumers' comments and machine learning to help customers deploy and use their power. The development of natural language technologies will eventually unlock the capacity to fully automate customer service.

Toward an entirely independently run grid?

Employing AI opens a vista of possibilities for the electricity sector. It may lead to a world where power generation, distribution, and transmission operations are automatically optimized, where the grid is balanced independently of any human interventions, where trading and arbitrage decisions are made in nanoseconds at a scale that only machines could tackle, and where end-users never have to worry about searching for a better supplier or changing the temperature manually.

As the price of solar cells and battery storage falls and their popularity rises, it is not inconceivable to imagine a day when distributed generation—power generation at the point of consumption—becomes the primary source of electricity and century-old grids, with their tens of thousands of generating stations and hundreds of thousands of miles of transmission and distribution lines, become the backups.

A study by the Lawrence Berkeley National Laboratory, for example, suggests that a 2.5 percent penetration rate of home-installed generation sources would reduce the earnings of a utility by 4 percent.⁶⁰ Utilities will have to work closely with policy makers to balance their own interests with the lowest economic cost and highest efficiency of resource use. To address this issue, a number of interventions could be explored. These include redefining rate structures for distributed power, creating a more flexible rate structure such as time-of-use pricing to drive efficiency and demand management, and expanding storage to balance the load and avoid massive spikes in demand. Not all of these changes will be easy, of course. Nor will they be without some discomfort for some industries and individuals. Regulators will be cautious in reviewing time-of-day pricing and other fundamental changes, considering the fraud, mis-selling, and consumer ignorance encountered when the United States deregulated energy markets in the 1980s.

⁶⁰ Andrew Satchwell et al., *Financial impacts of net-metered PV on utilities and ratepayers: A scoping study of two prototypical U.S. utilities*, Lawrence Berkeley National Laboratory, September 2014.

MANUFACTURING

SMARTER, MORE NIMBLE, AND LESS PRONE TO ERROR

Manufacturing is on the verge of a revolution in which artificial intelligence applications, from virtual assistants to advanced robotics, will disrupt end-to-end value chains amid radical shifts in demand. The scope of change will compel many manufacturers to adopt new plant designs, reshape their manufacturing footprints, and devise new supply chain models.

Advances in AI technologies will enable the industry to leverage rapid growth in the volume of data to optimize processes in real time. They can shorten development cycles, improve engineering efficiency, prevent faults, increase safety by automating risky activities, reduce inventory costs with better supply and demand planning, and increase revenue with better sales lead identification and price optimization.

However, many companies are unprepared to face the future: they lack visibility on the nature of changes such as new plant models, the implications for the way they do business, and what they require to manage the transition toward more a collaborative working model.

Intelligent manufacturing is a “smart” approach to production where machines linked through the internet assemble parts and adapt to new processes with minimal guidance from human operators. It has become a priority for governments and corporations around the world as they prepare for a restructuring of global industrial value chains driven by artificial intelligence.

A glimpse of the future of intelligent manufacturing may be had at Siemens' Electronic Works Amberg. People manage and control the production of programmable logic circuits through a virtual factory that replicates the factory floor. Via bar codes, products communicate with the machines that make them, and the machines communicate among themselves to replenish parts and identify problems. Nearly 75 percent of the production process is fully automated, and 99.99988 percent of the logic circuits are defect-free.⁶¹

In this section, we detail how AI could transform manufacturing in two industries, aerospace and semiconductors.

For engineering and R&D, artificial intelligence tools can mean quicker turnaround and fewer iterations

Engineers and researchers today face difficult challenges, from the sharp growth in demand in emerging countries to market fragmentation driven by consumers' taste for customization. At the same time, budget constraints require engineering teams to improve their productivity and efficiency, even as limits on the number of designs that can be considered and optimized against process capabilities restrict the predictability of product performance. AI-powered technologies can help deliver more efficient designs than previously achievable by eliminating waste in the design process. Innovation can be brought to the market faster as AI facilitates lower process cycle times and an increased focus on real-time negotiations and other interactions.

⁶¹ Siemens, “Welcome to Electronics Works Amberg (EWA),” presentation to analysts, September 29, 2015; Massimo Barbato, “Inside Amberg: Industry 4.0 in action,” Chartered Management Institute, November 4, 2015.

Intel deployed a team of data scientists in its R&D department to speed up data integration and the use of advanced analytics. Afterward, the company achieved 10 percent higher yield for integrated-circuit products compared with other players pursuing similar designs at a similar, pre-production development stage.

Predictive analytics using machine learning is indeed a powerful tool to reduce the time required to solve design problems for semiconductor manufacturers. Motivo, an artificial intelligence startup, managed to compress semiconductor design processes from years to a few weeks, saving chip makers the cost of iterations and testing.

Without machine learning, an aerospace manufacturer used advanced analytics to develop productivity tools for engineering teams, such as team travel norms, team composition, and supplier communication. The firm reduced its development costs by nearly €200 million. Machine learning could multiply this performance, unleashing the speed, accuracy, and relevance of advanced analytics.

We expect that manufacturers will apply machine learning to consolidated comprehensive data, enhancing the “design for manufacturability” process from beginning to end. AI also will permit more explicit enforcement of value-based decisions when making trade-offs, as when balancing safety and cost weighting. AI allows manufacturers to integrate production and client feedback in real time to refine the product design. With suppliers, AI-based tools could provide better accountability throughout the supply chain, which helps aerospace manufacturers, for instance, conform to safety regulations. Meanwhile, deep learning and network theory will help engineering development teams to better optimize their composition, ways of working, and key performance indicators, all in real time.

Untangling the procurement process to get a better grip on costs and supplies around the world

Keeping manufacturers stocked with parts is a complex challenge. Thousands of different parts must be sourced from tens of thousands of suppliers all over the world. When manufacturers are digitally linked with their suppliers’ systems, AI technologies can provide transparency on supplier machine availability, performance, and downtime. They can also help balance the supply chain and optimize inventories in real time.

One aerospace manufacturer attacked the problem by applying advanced algorithms to spending data to search for discrepancies between what it paid suppliers and what it actually owed them. The company also sought data that would help it understand differences across suppliers to develop effective procurement levers and reduce administrative costs.

In the future, supplier research and analysis will be automated and optimized with machine learning algorithms; e-auctions will be supplemented with virtual agents and, when needed, will automatically program in-person interactions. We estimate that full automation could reduce IT staff numbers by 39 percent.

Program managers can use advanced analytics to improve the effectiveness of review processes

Program reviews sometimes fail to detect emerging problems or prioritize critical decisions and tasks, resulting in delays that cost millions of dollars. Faulty communications between the marketing and manufacturing functions, which are still mostly manual, hinder the process.

One aerospace manufacturer sought to improve the effectiveness of its program reviews by using advanced analytics to predict key performance indicators and identify “traffic lights,” such as heavy email traffic, that could be a bellwether of problems that would appear later on. Doing this, the manufacturer generated a pre-tax run-rate cash flow impact of approximately €40 million.

In the future, manufacturers will be able to use deep learning technology to optimize the key performance indicators of program reviews in real time. Deep learning networks are already applied to make predictions in real time but are trained with batches of historic data (that is, not updated with real-time data streams). The tailoring of a model in real time implies performing deep learning on real-time information. This will help to better predict, identify, and prevent material and staffing bottlenecks and optimize energy consumption. Virtual agents could alert program leadership and team members when problems arise and recommend solutions. The step change will come with the development of natural language, because it will allow those virtual agents to engage with team members and help them solve problems.

AI can rethink manufacturing processes and assembly line practices to cut costs, reduce waste, and speed time to market

Manufacturing and assembly inefficiencies cost manufacturers billions of dollars every year. Existing fault detection and classification tools can be wildly inaccurate and cause expensive and unnecessary interruptions on the assembly line. An aerospace manufacturer applied AI technologies to these problems and reaped €350 million in savings. Nearly 60 percent came from using advanced analytics to review data from every step in the assembly process and then rewriting standard operating procedures based on the results. The rest came from using machine learning algorithms, collaborative robots, and self-driving vehicles to improve warehouse costs and reduce inventory levels.

A semiconductor maker reduced its material-delivery time by 30 percent by using machine learning to propose the best time to leave the office or warehouse. It also improved its production yield by 3 to 5 percent.

Factory managers can apply deep learning to real-time information flows in order to update and increase the accuracy of standard-operating-procedure predictions, especially during ramp-up, with visibility on component availability and risk management. Asset reliability can also be enhanced with AI tools, notably thanks to machine learning improving the predictive accuracy of defaults or production interruptions.

Virtual agents will deliver instructions and information on tablets or other interactive personal-communications devices to reduce assembly errors and flatten the learning curve for new operators.

Assembly lines will be significantly automated and optimized in real time, with control and scheduling directly connected with real-time dispatching systems.

Manufacturers can deliver after-sales service just when it is needed

Aircraft maintenance and service is a large part of the value chain for aerospace firms. In the late 2000s, GE broke new ground in this space with its “power by the hour” service concept, which calls for operators to pay for aeroengines only when planes are flying. This business model is becoming increasingly dominant in aerospace manufacturing, notably thanks to the internet of things. The network of sensors and actuators enables companies to monitor the actual use of their products, provides customized pay-as-you-go services, and better prevents service disruption and downtime. AI technologies are ideally suited to leverage the vast amount of data collected from operating engines, while the manual prioritization of maintenance, repair, and overhaul sales leads is cumbersome, resource-heavy, and not always optimized. This results in unnecessarily keeping aircraft grounded—not generating revenue—and inefficiently allocating expensive engineering labor. With real-time feedback between plane and ground support facilities, machine learning algorithms will be able to make judgment calls on when to deploy drones and smart microrobots to conduct aircraft inspections and quality checks on intermediate and final products. Manufacturers with the best algorithms and data will be able to offer service contracts that promise better performance at a lower cost.

Improving the accuracy of forecasting MRO work and focusing sales efforts on the most promising leads can have a significant effect on manufacturers’ EBIT. One firm reported a profit improvement of approximately €300 million from using machine learning to forecast 10 years of repair events for a fleet of over 17,000 commercial aircraft, and to develop a deal-scoring tool to advise on “what good looks like” when pricing MRO work.

Advanced analytics and AI tools are also used to optimize processes around unplanned maintenance events, allow the manufacturer to respond more effectively to disruptions, and increase uptime. For instance, GE turned to Kaggle, a platform for predictive modeling and analytics competitions, and invited data scientists to design new routing and machine learning algorithms for flight planning that optimized fuel consumption by looking at variables such as weather patterns, wind, and airspace restraints. The winning routing algorithm showed a 12 percent improvement in efficiency over actual flight data.

In the future, AI tools will shift predictive analytics to cognitive assessments. The algorithms themselves will discover new rules, automatically optimizing sales and servicing for manufacturers. Preventive maintenance will be conducted with real-time feedback between plane and ground support facilities, using machine learning and virtual assistants to identify issues and using drones for inspections and quality checks on intermediate and final products. Robots the size of insects will be able to inspect airframes without removing panels and identify common defects with computer vision and machine learning. Smart flight systems will use predictive models based on deep learning technology to provide real-time feedback to flight crews, helping them to optimize energy consumption throughout flight.

We may see manufacturers shift from the traditional business model of providing spare parts and MRO services to an end-to-end approach amid an increase in joint ventures or industry consolidation.

In any case, virtual agents and deep learning technology will be able to improve the training delivered to maintenance operators and pilots, and machine learning could optimize global supply networks for spare parts to proactively stock critical and non-critical parts, thereby limiting aircraft downtime and reducing inventory costs.

People and factories must collaborate and communicate better to realize the full potential of AI

The key feature of the future manufacturing paradigm will be collaborative agility, which is the ability to adapt almost instantly to changes in demand and the evolution of regulation, input prices, technologies, and other parts of the industry landscape. From a human capital perspective, manufacturing will become more collaborative, based on increasingly complex activities and interdependent workers. From a technological perspective, manufacturing plants around the globe, supply chains, and value chains will be more interconnected and collaborative via a global digital backbone. This will combine highly automated plants that extensively use smart robotics for mass production of limited product ranges as well as a network of customer-centric plants close to higher-end market segments, and “e-plants in a box”—small-scale, low-capital-expenditure, mobile plants that can produce a limited range of products at a competitive cost. These e-plants will work well in niche markets where demand is temporarily strong and in remote markets where production must be local and low cost. Machine learning applications will help these integrated manufacturing systems self-adjust in real time.

The risk-conscious aerospace industry adopts new technology only as fast as it can ensure long-term safety. The US Federal Aviation Administrator and other regulators have change-management systems—for example, the parts manufacturer approval procedures in the United States—to monitor engineering alterations as they are made and then subject the changes to two to three years of testing before approving their use.

Manufacturers will need to significantly invest in training as technicians’ roles become more focused on knowledge and exception handling. Machine learning, advanced robotics, virtual agents, and other technologies may be new and difficult to master for mechanics, 96 percent of whom are more than 30 years old.

Aircraft maintenance automation can create approximately \$5 billion to \$15 billion of value for the industry annually, about 35 percent from performance gains and the rest from labor substitution. Given the high level of investment required in both technology and knowledge for full AI scale-up and automation, airlines may rely on partial measure for many years.

Real-time automation of the supplier selection process will blur where legal liability will lie if an accident occurs, because machine learning does not yet allow humans to fully understand how the algorithm got to the assessment.

Procurement agents’ roles will focus more on joint value creation than on negotiation prowess, requiring more cross-functional skills. Certain AI integration changes will require significant capital investment. However, after the technology demonstrates its value, automakers will be quick to invest because shortening time to market offers a significant competitive advantage.

HEALTH CARE

AI WILL MAKE QUICKER DIAGNOSES, CREATE BETTER TREATMENT PLANS, AND ENABLE NEW APPROACHES TO INSURANCE

Health care is a promising market for AI. There is enormous potential in its ability to draw inferences and recognize patterns in large volumes of patient histories, medical images, epidemiological statistics, and other data. AI has the potential to help doctors improve their diagnoses, forecast the spread of diseases, and customize treatments. Artificial intelligence combined with health care digitization can allow providers to monitor or diagnose patients remotely as well as transform the way we treat the chronic diseases that account for a large share of health-care budgets.

AI's impact on the health-care industry could be significant. In our survey, we asked industry executives to review case studies about AI adoption. Executives from health-care companies that were early adopters of AI said they expect the technologies will raise operating profit margins by five percentage points within the next three years.

Several AI technologies appear to be suitable for use in medical practices. Machine learning is starting to be applied in payments and claims management, but its further application in health care may arrive at scale soon. Machine learning is suited to analyzing the data in millions of medical histories to forecast health risks at the population level. This could be an early win for AI because it brings the potential for large savings and would not require the regulatory scrutiny to be expected when trying to anticipate individual health risks.

Using AI to diagnose illnesses may not happen so fast. While machine learning is able to use data to make a diagnosis, completely automated diagnosis is not likely to happen quickly, partly because of questions about whether patients will accept it, and partly because of the technical difficulty of integrating data from multiple sources and complying with strong regulatory requirements.

Hospitals also could improve their capacity utilization by employing AI solutions to optimize many ordinary business tasks. Virtual agents could automate routine patient interactions. Speech recognition software has been used in client services, where it has reduced the expense of processing patients by handling routine tasks such as scheduling appointments and registering people when they enter a hospital. Natural language processing can analyze journal articles and other documents and digest their contents for quick access by doctors. These kinds of applications can have a significant impact without needing to pass a regulatory review.

Before the medical profession can realize this potential, however, health care providers must adopt significant changes in the way they do business, commit to a substantial investment in computing power and technical expertise, and work to increase the availability of the fuel that will power progress: data, including medical records. (Specialized data brokers, such as Explorys, which IBM bought in 2015, already offer to aggregate health care data and sell it to potential AI solutions providers and users.)

In any case, the success of AI-based tools in medicine will hinge on whether public officials sign on and pitch in with financing, research support, and legislation that protects patients' privacy and gives medical professionals access to anonymized data on illnesses, treatments, and outcomes to teach computers how to identify and treat a wide range of maladies.

Making these changes will not be easy. But there are considerable rewards for success: AI is capable of improving care while reducing costs—no small matter when health-care spending globally reached 9.9 percent of GDP in 2014 (it was 11.5 percent in France and 17.1 percent in the United States), according to the World Health Organization.⁶²

Despite AI's potential, health care trails other industries in adopting the technology, according to our survey. AI use is concentrated in operations and customer service; the technologies adopted most often are speech recognition and computer vision, by 9 and 7 percent, respectively, of health care companies in our survey sample, which included organizations that already were aware of AI. In most hospitals, operations management functions such as appointment scheduling are still done manually.

We have found that if a sector was slow to adopt digital technologies, it tends to trail the pack in putting AI to use, too. Our report *Digital America* found that almost one-quarter of the nation's hospitals and more than 40 percent of its office-based physicians have not yet adopted electronic health record systems.⁶³ Even those that do have electronic record systems may not be sharing data seamlessly with the patient or with other providers; tests are repeated needlessly and patients are required to recount their medical histories over and over because these systems are not interoperable. Another MGI report, *The age of analytics*, found that the US health-care sector has realized only 10 to 20 percent of its opportunities to use advanced analytics and machine learning.⁶⁴

This slow progress does not stem from a lack of interest among medical professionals and executives. There is interest, but medicine faces some uniquely high hurdles to adoption. The sensitive nature of medical records and strict regulations to keep them private has stymied the collection of the high-quality aggregated data required by deep learning applications and other AI tools. Also slowing adoption are the complexity of both that data and the industry itself, the fragmentation of the health-care industry, and other regulatory barriers.

AI can identify public-health threats and the most at-risk patients

AI technology adoption rates are low. The most advanced application area currently is payment and claims management supported by machine learning algorithms. Some clinicians are using AI to forecast the spread of certain diseases and try to anticipate which patients would be most likely to succumb. Armed with this information, they offer preventive care. They also use the forecasts to help hospital administrators schedule staff members, negotiate reimbursement rates with insurers, set budgets, and optimize inventory levels.

⁶² World Health Organization, "Total expenditure on health as a percentage of gross domestic product (US\$)," Global Health Observatory, February 13, 2017.

⁶³ *Digital America: A tale of the haves and have-mores*, McKinsey Global Institute, December 2015.

⁶⁴ *The age of analytics: Competing in a data-driven world*, McKinsey Global Institute, December 2016.

This idea of leveraging medical and social data to better manage costs has made forecasting one of the few areas of active AI applications in health care, attracting top tech, pharma, and medical players as well as small startups. Johnson & Johnson, in partnership with SAP, has used machine learning to anticipate customer demand, inventory levels, and product mix. Careskore, a predictive analytics platform, uses machine learning to determine the likelihood of a patient's being readmitted to a hospital.

Indeed, in the future, AI tools will enable health care to dramatically accelerate its shift toward preventive medicine. Medical professionals will focus on managing patients' health remotely and keeping them out of hospitals. To do this, AI tools will analyze not only patients' medical histories but also environmental factors that can influence health, such as pollution and noise where they live and work. This can identify risk groups and inform local authorities' decisions about where to implement preventive-care programs.

Primary care providers will have information to engage patients about preventive actions that involve both medical services and lifestyle and environmental factors such as nutrition, exercise, and pollution avoidance. Hospital administrators will be better equipped to forecast spikes in admissions in ways not available today. Tracking the incidence of communicable diseases, combined with personal medical records, weather data, and other information will help an AI tool estimate how many people will need hospitalization.

In routine times, population health management analytics can identify service gaps. An AI application, for example, could use medical and demographic data to anticipate a rise in births and alert health care administrators if obstetric clinics require additional staff.

We estimate the full potential health care service cost savings of AI-enabled initiatives would be \$300 billion a year in the United States, or about 0.7 percent of GDP.⁶⁵ In the United Kingdom, using AI to target preventive care and reduce non-elective hospital admissions can save £3.3 billion annually.⁶⁶

AI can help medical professionals diagnose disease and improve operations

Machine learning has enormous potential to enhance diagnostic accuracy. The Sloan Kettering Institute estimates that doctors use only 20 percent of the available trial-based knowledge when diagnosing cancer patients and prescribing treatment. AI applications can sift through millions of pages of medical evidence to provide a diagnosis and treatment options in seconds.

AI-based image recognition and machine learning can see far more detail in MRI and X-ray images than human eyes can register.⁶⁷ For example, different types of glioblastomas have distinct genetic abnormalities, and doctors treat each one based on those abnormalities. But radiologists cannot identify genetic abnormalities of these brain cancers from images alone. The Mayo Clinic has a machine learning program that can quickly and reliably identify the abnormalities.⁶⁸

⁶⁵ Ibid.

⁶⁶ "Camden's CCG's analysis of the scale of opportunity for reducing non-elective admissions," in *Better Care Fund, Appendix B—Defining our Focus and Ambitions*, Camden Clinical Commissioning Group, September 2014; and Alicia O'Cathain et al., "A system-wide approach to explaining variation in potentially avoidable emergency admissions: National ecological study," *BMJ Quality and Safety*, volume 23, issue 1, January 2014.

⁶⁷ Bertalan Meskó, "Top artificial intelligence companies in healthcare to keep an eye on," *MedicalFuturist.com*, February 2, 2017.

⁶⁸ Marianne Matthews, "Machine learning can bring more intelligence to radiology," *Healthdatamanagement.com*, May 2, 2017.

Other institutions also have recognized the medical potential of machine learning and moved to harness it. Moorfields Eye Hospital in London has teamed with Google's DeepMind, and the Cleveland Clinic Lerner College of Medicine is collaborating with IBM Watson.⁶⁹ Software developers do not need to partner with a medical center to innovate. GE Healthcare and Alteryx have built ViosWorks, which uses an AI algorithm to improve and accelerate MRI scanning. ViosWorks can use a series of MRI scans to create a 3-D moving picture of a beating heart while also showing in real time how much blood is pumped with each contraction.⁷⁰

Innovation is not limited to imaging. Entrepreneurs are working to change each step in the patient care process. A startup called Enlitic is working on a deep learning app that could improve the accuracy of disease diagnostics. Oncora Medical has developed an AI tool that helps oncologists draft personalized radiation treatment plans for cancer patients.

AI-powered automation has the potential to increase health care productivity by relieving doctors and nurses of routine activities. Someday, chatbots equipped with deep learning algorithms could relieve emergency room personnel of tending to large numbers of walk-in patients with non-emergencies like sore throats and urinary tract infections.

Together, AI-enabled operational efficiencies could represent sizable savings in developed countries. Estimates for the United States range from 1 to 2 percent of GDP. Across other high-income countries, estimated savings would be 0.5 to 1 percent of GDP.⁷¹ Full AI adoption could raise the productivity of registered nurses by 40 to 50 percent.⁷² McKinsey research has found that this could allow hospitals to cut staffing costs in half while still significantly reducing patient waiting time.

Insurers can devise new ways to encourage preventive care and incentivize providers

The ability of machine learning technologies to predict patient behavior and calculate disease probabilities better than current methods will lift the profitability of life- and health-insurance providers.

New business models can use AI combined with behavioral health interventions to focus on prevention, disease management, and wellness—addressing unhealthy behaviors before people become patients. A South African insurer, Discovery Health, tracks the diet and fitness activity of people it insures and offers incentives for healthy behaviors.⁷³

AI also will encourage new partnerships among payers, providers, and pharma companies and will facilitate pay-for-performance models that will accelerate the shift toward preventive care. Payers may become more involved in care management or encourage their providers to do so by introducing contract models based on risks uncovered by machine learning or the potential for AI-based risk-management modeling.

⁶⁹ Bertalan Meskó, "Top artificial intelligence companies in healthcare to keep an eye on," MedicalFuturist.com, February 2, 2017.

⁷⁰ Drew Field, "See the heart in 7 dimensions: This team of researchers attacks world's biggest killer with software," GE Reports, December 3, 2015.

⁷¹ *The age of analytics: Competing in a data-driven world*, McKinsey Global Institute, December 2016.

⁷² *The age of analytics: Competing in a data-driven world*, McKinsey Global Institute, December 2016; *A future that works: Automation, employment, and productivity*, McKinsey Global Institute, January 2017.

⁷³ Discovery Holdings.

Episode-based payment plans, which reimburse doctors and hospitals based on the average cost of treatment across all providers in the group, will be significantly extended when more insurers use machine learning to analyze historical inpatient data. Based on McKinsey's client experience, we believe that this approach can have a clear impact on costs, reducing orthopedic surgeons' fees by 8 to 12 percent and the fees paid to diagnosing physicians by 4 to 5 percent.

Doctors will be able to tailor treatments—even drugs—to individual patients

Patients also can benefit directly from the rise of AI in health care. Standardized treatments do not work for every patient, given the complexity of each person's history and genetic makeup, so researchers are using advanced analytics to personalize regimens. Decisions can be based on data analysis and patient monitoring with use of remote diagnostic devices. A startup called Turbine uses AI to design personalized cancer-treatment regimens. The technology models cell biology on the molecular level and seeks to identify the best drug to use for specific tumors. It can also identify complex biomarkers and search for combination therapies by performing millions of simulated experiments each day.

AI's ability to use a universe of data to solve a narrow problem resonates with advocates of customized medical treatments. Knowing the health outcomes of millions of other people with similar symptoms, prognoses, and ages is invaluable to providers that promise one-of-a-kind pharmaceuticals, physiotherapies, and other actions designed to deliver the greatest benefit with the fewest side effects. Several companies are already using machine learning or other AI technologies to tailor their treatments to individual patients. Mindmaze uses machine learning to optimize rehabilitation activities for stroke patients. Ginger.io uses machine learning to recommend the best time to take medication based each patient's metabolism and other factors.

Tailored treatments may reduce health expenditures by 5 to 9 percent, add 0.2 to 1.3 years to average life expectancy, and increase productivity by \$200 per person annually. Globally, the economic impact could range from \$2 trillion to \$10 trillion.⁷⁴


Virtual agents can serve as primary touchpoints for patients

Medical practices have taken small steps toward incorporating AI into patient management, introducing speech recognition and other language AI technologies to automate steps in the process. In the future, virtual assistants equipped with speech recognition, image recognition, and machine learning tools will be able to conduct consultations, make diagnoses, and even prescribe drugs. If these systems lack enough information to reach a conclusion, a virtual agent could order additional tests and schedule them with the patient. In rural areas, virtual agents will be able to conduct remote consultations. However, this scenario would require patients, providers, and regulators to become comfortable with fully automated diagnosis and prescriptions.

Less controversially, in hospitals, virtual agents will be able to register patients and refer them to the appropriate doctor to address their issues. Virtual agents would be able to help patients navigate hospital bureaucracy, prepare them for tests, and make sure they are on time for appointments.

⁷⁴ *The age of analytics: Competing in a data-driven world*, McKinsey Global Institute, December 2016.

AI in health care: quicker diagnoses, better treatment plans, and improved health insurance




Machine learning program analyzes patients' health remotely via mobile device, compares it to medical records, and recommends a fitness routine or warns of possible disease

Autonomous diagnostic devices using machine learning and other AI technologies can conduct simple medical tests without human assistance, relieving doctors and nurses of routine activities




AI-powered diagnostic tools identify diseases faster and with greater accuracy, using historical medical data and patient records

AI algorithms optimize hospital operations, staffing schedules, and inventory by using medical and environmental factors to forecast patient behavior and disease probabilities




AI tools analyze patients' medical histories and environmental factors to identify people at risk of an illness and steer them to preventive care programs





Virtual agents in the form of interactive kiosks register patients and refer them to appropriate doctors, improving their experience and reducing waiting time

Personalized treatment plans designed by machine learning tools improve therapy efficiency by tailoring treatment to specific patients' needs and medical



AI insights from population health analyses give payers an opportunity to reduce hospitalization and treatment costs by encouraging care providers to manage patients' wellness



Several hurdles stand in the way, starting with data availability

One of the biggest potential bottlenecks that could inhibit or derail AI development and adoption in health care is the availability of sufficient quantities of high-quality data in standardized formats. As noted earlier, information today is highly fragmented and spread across the industry, residing in diverse, mostly uncoordinated repositories like electronic medical records, laboratory and imaging systems, physician notes, and health-insurance claims. Merging this information into large, integrated databases, which is required to empower AI to develop the deep understanding of diseases and their cures, is difficult. Cultural barriers, for example, stand in the way of partnerships among health-care data owners—hospitals, insurers, drug makers, and diagnostic companies. And, of course, there is the highly sensitive nature of the data itself. While people routinely allow access to certain kinds of personal data when they buy online or join loyalty programs, they may resist any attempt to give wider access to their more intimate medical histories, particularly if they do not see it as a necessity and the potential benefit is abstract. They may also be concerned that centralized collections of intimate health details would be a natural target for hackers and data thieves. Regulators will need to proactively develop clear rules defining who would be able to use the data, what they could be used for, how they must be stored, and how they would be anonymized.

Technological limitations are another hurdle. To work, AI technologies must know a great deal about the patient and other humans, but humans know little about how AI technologies actually make a diagnosis or choose a treatment plan. How much patients would trust AI tools and be willing to believe an AI diagnosis or follow an AI treatment plan remains unresolved. Regulators would not be eager to risk an incorrect computer decision harming a patient when no one would be able to explain how the computer made its decision—or how to prevent a repeat of the situation. This is a particular issue for the most powerful AI tools, such as deep neural networks, and could remain the case for some time, even though AI tools are, in theory, less likely to make mistakes than individual human clinicians.

AI companies will have to address market fragmentation. Hundreds of vendors offer thousands of different machine learning programs, each one designed for a specific clinical situation. But in daily practice, physicians need platforms that can address many different situations.

Health-care providers also have work to do if they want to take advantage of AI capabilities. To start, they must hire or develop people with the education and skills needed to deploy, maintain, and operate AI systems. In addition to data analysts and technical staff, this includes people with project management, team development, and problem-solving skills. At the same time, legacy staff—doctors, nurses, and other medical professionals—will have to become accustomed to working with the support of machines and AI tools. This may require them to overcome significant skepticism, although it will give them the opportunity to focus more on clinical cases and leave administration and low-risk work to AI and digital solutions.

EDUCATION

VIRTUAL TUTORS POWERED BY AI CAN PERSONALIZE LEARNING AND OPTIMIZE TEACHING

For decades people have discussed how to revolutionize education with technology, whether “gamifying” instructional materials or expanding access to knowledge via massive open online courses. EdTechXGlobal and Ibis capital estimated that schools spent nearly \$160 billion on education technology, or ed tech, in 2016, and forecast spending to grow 17 percent annually through 2020.⁷⁵ Private investment in educational technology, broadly defined as the use of computers or other technology to enhance teaching, grew 32 percent annually from 2011 through 2015, rising to \$4.5 billion globally.

AI’s share of these flows has not been tabulated, but it is likely to increase because artificial intelligence technologies are well suited to achieving crucial education objectives, such as enhancing teaching efficiency and effectiveness, providing education for all, and developing the skills that will be essential in the 21st century. So where will education be in 2030 in terms of artificial intelligence? Most probably, it will play an important part. However, success hinges not only on technical issues but on ethical issues, starting with who owns data on students, who can see it, who can use it, and for what purposes.

Bridging the skills gap

Many countries suffer from significant skills mismatches, brought about both by the inability of the education system to accurately reflect the demands of employers and by the frictions in the labor market that prevent optimal matching of individuals to jobs.⁷⁶ In a survey across 10 developed and developing countries, only half of students believed their post-secondary studies improved their employability, and more than a third of employers thought skills shortages are a leading reason for entry-level vacancies.⁷⁷ Not only does the resulting skills gap lead to economic underperformance, it also means that many individuals are not given the opportunity to reach their full potential.

Artificial intelligence will also play a key role in better connecting education systems and labor markets. Digital technologies are already making a difference by connecting talent with opportunities in the job market. A recent MGI study estimated that by 2025 online talent platforms could enable as many as 60 million people find work that more closely suits their skills or preferences and reduce the cost of human resources management, including recruitment, by as much as 7 percent.⁷⁸ With an increasing emphasis on lifelong learning, the opportunities for artificial intelligence in employment-to-education settings have already started to attract new players. In 2015, the employment-oriented social network LinkedIn acquired the educational website Lynda.com, hoping to leverage AI to offer a personalized online class selection for members considering a new job or career.

⁷⁵ “2016 Global EdTech industry report: A map for the future of education,” EdTechXGlobal and Ibis Capital.

⁷⁶ Müge Adalet McGowan and Dan Andrews, *Labor market mismatch and labor productivity: Evidence from PIAAC data*, OECD, April 28, 2015.

⁷⁷ *Education to employment: Designing a system that works*, McKinsey & Company, 2013.

⁷⁸ *A labor market that works: Connecting talent with opportunity in the digital age*, McKinsey Global Institute, June 2015.

In the future, improved pattern recognition enabled by machine learning, along with detailed data on potential employees, could improve recruitment results further. It could enable hiring companies to pinpoint the precise skill sets and personal traits that would enable someone to be successful in a job, and surface insights that have previously not been used in talent management. Artificial intelligence may also help detect promising candidates with less conventional credentials and free recruiters from using school reputation as a proxy to assess candidates' potential.

More fundamentally, artificial intelligence will enable education systems to better meet the needs of future employers. AI can be used by governments to forecast detailed job-market demand more accurately and steer educational institutions to adapt their curricula and approaches accordingly, making sure students have the skills required to fill those jobs. An example of this is Saudi Arabia's current exploration of machine learning as a tool to reduce unemployment. The administration hopes to leverage large amounts of past and forecasted economic and social information about the country to possibly guide students toward an education best matched to their abilities.

There is a risk, however: in an Orwellian scenario, AI could also be used to optimize labor markets without regard to nuanced social preferences, make education decisions on citizens' behalf, and sell valuable data on people's skills to private companies or political parties. Protecting individuals' data privacy is therefore critical for enabling AI to help bridge the skills gap.

Attracting students, and keeping them

Alongside governments, educators themselves will be able to use personal, academic, and professional data to ensure that students benefit from the courses they choose. The value stems not only from the ability of students to flourish academically but also from the institutions' ability to help them to find meaningful jobs. Machine learning could also be used to identify people who appear unsuitable based on traditional measures of academic success but have high potential based on other abilities and traits. Better targeting will benefit students and enable institutions attract the most appropriate mix of individuals, improve learning outcomes, and help schools and universities continually improve their offerings.

Universities are already exploring AI applications that can improve student retention. Some schools and colleges are testing advanced analytics and machine learning to identify students in trouble and offer them support before they drop out. Civitas Learning and Salesforce collaborate on a service for universities that identifies and engages with students at risk of quitting. The Salesforce tools use machine learning to recommend engagement strategies to optimize retention and graduation rates.

In the future, computer vision could identify signs of students' disengagement by monitoring them as they work, tracking their eye movements and observing their expressions to check whether they are engaged, confused, or bored. In the United Kingdom, some institutions are experimenting with computer vision, natural language, and deep learning algorithms to better understand students' learning difficulties and learning preferences, incorporating novel types of data, such as students' activities on social media.

Unleashing personalized learning

Attracting and retaining students is key, but the real step change for education will likely emerge from a radically different approach to learning itself, whether in the classroom or outside of it. For the last few decades, many efforts have been made to tailor learning to each student and to shift away from a standardized approach. Adaptive learning solutions aim to address the limitations of conventional classroom teaching by personalizing lesson plans to the student's existing knowledge, particular learning preferences, and individual progress. Instead of delivering one lesson to the entire class, which can leave behind struggling students or disengage fast learners, adaptive learning claims to deliver the right content, at the right time, in the best way to each student. As of 2015, some 80 companies, such as Knewton and DreamBox Learning, were marketing adaptive-learning systems to schools in North America, Europe, and Asia.⁷⁹

Artificial intelligence could improve adaptive learning and personalized teaching by identifying factors or indicators of successful learning for each student that were previously not possible to capture. In addition to monitoring such variables as the number of times a student pauses during a lesson, the amount of time needed to answer a question, and the number of times a question was attempted before getting it right, computer vision and deep learning could call in new information such as mouse movements, eye tracking, and sentiment analysis, delivering a deeper insights on the student's performance, confidence, mindset, and cognitive ability.

Implemented at scale, AI-enabled adaptive learning could restructure education. First, it could put an end to traditional testing systems and measure academic abilities and achievement in a more nuanced way. Class formats would give more room for students to learn according to their own preferences, with teachers focusing less on lecturing and more on coaching, aided by prescriptive analytics to choose the most effective methods.

Finally, artificial intelligence could empower students by providing them with control over how fast they learn, awareness of how they learn best, and the lifelong feedback of one's own cognitive and behavioral preferences. Ultimately, deep learning algorithms could not only predict outcomes and prescribe accurate solutions, but explain how the algorithm reached its conclusion and help retro-engineer drivers of educational success. This will allow students to reflect on their cognitive abilities and understand their own optimal learning setting, even after they leave school. Empowered by AI, students could build their own virtual robot teachers to help them navigate lifelong learning experiences.

The promise of AI is to use digital tools as a cognitive window into students' minds, and help tailor learning to maximize individuals' potential. However, the data it would leverage is by nature private and sensitive. Education institutions and education technology providers will need to put in place solid guarantees of data protection for students and families to embrace the use of AI-powered tools.

⁷⁹ "Learning to Adapt 2.0: The evolution of adaptive learning in higher education," Tyton Partners, 2016

Releasing teachers' true value add

In the future, the role of teachers could be stripped of time-consuming administrative tasks, such as supervising and answering routine questions. Teachers would have more time to mentor and coach students—value-adding tasks that are (at least currently) uniquely suited to human beings.

Natural language, computer vision, and deep learning could help replace teachers in answering students' routine questions or acting as tutorial supervisors. In 2014, a Georgia Tech professor and his team created a robot teaching assistant, which provided responses to students' online questions for five months without the students noticing.⁸⁰ The professor estimates that in 2017, the robot will be able to answer 40 percent of the students' questions.

A virtual supervisor could harness AI to track students' work and behavior and to support teachers by supplying statistically based insights on students' progress and constructive feedback. Coursera's online classes rely on machine learning to alert teachers when a large number of students make similar errors on an assignment, suggesting possible gaps in the teachers' lectures or course materials. In the future, AI solutions may also leverage voice and facial recognition to supervise an entire classroom and call out students individually.

Finally, AI could assist teachers in forming the most effective groups or classes by applying machine learning algorithms to data from students' education profile, social media, and surveys. Startups such as Collaboration.ai use artificial intelligence to process data on each student's experience, knowledge, and capabilities; to create instantaneous maps of connections and networks; to highlight each student's specific potential; to break down preferences and bias; and to recommend group formations best suited for the learning objective. Machine learning can identify complementary skills that will maximize critical thinking and test students' capacity to adapt and collaborate.

Toward virtual teachers

UNESCO estimates that the world will need to recruit and train 24.4 million primary school teachers in order to achieve universal primary education by 2030 and another 44.4 million teachers to fill openings at secondary schools. Many of these new hires—more than 85 percent of them, in the case of primary schools—will be required just to replace teachers who leave education. Artificial intelligence could be one part of a solution. With a much wider reach, AI-assisted teaching could have a significant impact in third world countries and remote locations by supporting two key enablers of teaching: coaching and assessing.

Coaching and assessing require specific skills, such as emotional intelligence, creativity, and communication, that are beyond machines' current capabilities. Yet with new indicators, such as facial expressions, digital interactions, group interactions, and attention tracking, deep learning algorithms could recognize patterns, attitude toward the learning situation, and affective states, and could support students in real time. A European Union project called iTalk2Learn is currently developing an open-source intelligent tutoring platform to help primary school students learn mathematics. Using a combination of machine learning, user modeling, and natural language processing, the tool is already able to interact with and respond to a student's speech throughout a tutoring session.

⁸⁰ Jason Maderer, "Artificial intelligence course creates AI teaching assistant," Georgia Tech press release, May 9, 2016.

AI-powered machines also are making headway in student assessment. Companies like GradeScope already use computer vision and machine learning to grade students' work quicker than a teacher, starting by deciphering handwriting and remembering the teacher's initial decisions on marks to automatically grade subsequent students. The technology today can assess only work with objectively correct answers, such as math problems, and rule-based learning, such as orthography, languages, and history events.

In the future, advances in natural language could expand AI's usefulness in automatically grading more creative work, such as essays and presentations. In a multivendor evaluation conducted by the University of Akron in 2012 and hosted by predictive modeling and analytics competitions platform Kaggle, some 16,000 essays were assessed by both grading software and teachers. Computers matched the grades from human teachers as much as 85 percent of the time.⁸¹ However, automatic grading immediately raises questions around creativity and pattern disruption: how can machine learning be tempered to prevent uniform thinking or the upholding of biases?

An important role for AI

Artificial intelligence clearly presents significant opportunities to raise the quality of education to a level that our current standardized-curriculum-and-testing systems have not been able to achieve and allow a shift of teachers' focus on higher-value creative and interpersonal tasks. Yet implementation at scale hinges on overcoming a structurally fragmented sector and will require support from key stakeholders—students, parents, teachers, administrators, and policy makers. Investment in technology and new capabilities in the classroom and in educational administration will also be crucial.

Most importantly, data privacy issues must be addressed. The use of AI in education raises legitimate concerns about how educational data, like other intimate personal data, are gathered and used. What are the risks that a student's longitudinal performance data, intended for teachers to improve instruction, become public, or that poorly performing students are denied educational and employment opportunities? To reap the benefits of AI in education, it will be essential that students and families take on roles as advocates and advisers in the design of AI solutions, that teachers and administrators familiarize themselves with AI technologies, and that policy makers and regulators create a safe and protected environment for AI-enabled education.

⁸¹ "Man and machine: Better writers, better grades," University of Akron press release, April 12, 2012.

APPENDIX B: TECHNICAL APPENDIX

ESTIMATE OF INVESTMENT IN AI

Investment flows into AI technologies were divided into two categories: external and internal investment. The external investment category captures annual investment in artificial intelligence companies by VC and PE funds, as well as M&A activities by corporations. The values provided are estimates of annual investment in AI based on data available in the Capital IQ, PitchBook, and Dealogic databases. The estimate assumes that all registered deals were completed within the year of transaction. For VC and M&A deals, only AI companies whose core technology is AI were included. For PE investment, target companies needed to be strongly related to AI. Deals announced but not completed by the end of 2016 were excluded.

The internal investment category includes expenditure on all corporate activities both for developing AI-based products or services and for building and leveraging the company's AI capabilities. It excludes M&A. The estimate of internal investment is based on a sector-specific AI spend-to-revenue ratio for 35 major companies (18 from North America, ten from Asia, and seven from Europe) making strong plays in AI. Companies were divided into the following sectors: advanced manufacturing, automotive, hardware/electronics, internet services, retail, and systems/software. For each sector, companies that disclosed information on internal AI expenditure were used to estimate a sector-wide average ratio. The estimate of total internal investment was arrived at by applying these calculated ratios to total revenues by sector.

SURVEY OF AI ADOPTION AND USE

As part of this research, we executed a survey designed to understand trends in AI investment, adoption, use, and business implications by sector, geography, type of company, technology, and application. The survey was conducted in March 2017. The final survey sample after quality checks consisted of C-level executives from 3,073 companies. The survey targeted respondents who were aware of at least one AI technology or application and its use in business from the following list: natural language processing, natural language generation, speech recognition, image recognition and video processing, machine learning and deep learning, virtual agents or artificial conversational entities, robotics, robotics process automation, and decision management (Exhibit 11).

The survey sample covered 14 sectors of the economy, ten countries (from Europe, North America, and Asia), and companies with workforces ranging from fewer than ten to more than 10,000 employees.

The survey consisted of four groups of questions, in addition to basic information about the company and the respondent. The first group of questions asked respondents about their awareness and adoption rates of AI technologies or applications. The second group asked about the current and future impact of AI in the respondent's sector, including the most important AI technology. It also asked about the parts of the business in which AI was being deployed, and drivers and barriers for AI deployment in the respondent's business. The third group of questions investigated the financial impact of AI, asking respondents to report current and future operating profit margins as well as spend on AI as a share of digital investment. The last group addressed the organizational impact of AI technologies, specifically on levels of employment and skills requirements.

Results were weighted by the share of employment represented by each firm size category, such that the overall results reflect the relative economic importance of different sizes of firms. Firm size weights were calculated separately for each sector, comparing the share of respondents in a particular firm size category with the share of the workforce in that category, using an average of the 10 countries in the survey.

Exhibit 11

MGI AI adoption and use survey sample overview

% of respondents (n = 3,073)

Geography		Company size		Sector	
Sweden	5	>10,000	7	Other	12
South Korea	9	5,000–10,000	6	Energy and resources	3
China	10	1,000–5,000	15	Travel and tourism	4
Germany	10	500–1,000	10	Automotive and assembly	4
Japan	10	250–500	10	Transportation and logistics	4
Italy	11	50–250	11	Telecommunications	5
Canada	11	10–50	14	Consumer packaged goods	5
France	11			Education	5
United States	11	<10	27	Media and entertainment	5
United Kingdom	12			Financial services	5
				Health-care systems and services	7
				Construction	8
				Retail	8
				High tech	10
				Professional services	14

NOTE: Numbers may not sum due to rounding.

SOURCE: McKinsey Global Institute analysis

AI INDEX

The AI index measures the extent of AI adoption and usage in 13 sectors across the 10 countries in the survey. (The telecommunications and high tech sectors were merged to align with the MGI Industry Digitization Index.) The index is based on 16 input metrics, divided into three categories: AI assets (three metrics), AI usage (11 metrics), and AI-enabled labor (two metrics). Using principal component analysis, the input metrics were combined into an overall AI adoption score. The data for these metrics were primarily obtained from the AI adoption and use survey, proprietary databases, and the MGI Industry Digitization Index (Exhibit 12).

Exhibit 12

Metrics included in the industry AI index

Metric		Description	
Assets	Depth of AI technologies	Average number of AI technologies adopted at scale or in core part of business per company	
	AI spend	Average AI spend as share of total annual investment ¹	
	Supporting digital assets	Percentage of firms using cloud and big data	
Usage	Product development	An entirely new product or service	Percentage of firms in sector using AI for entirely new product or service
		Research and development	Percentage of firms in sector using AI in R&D
	Operations		Percentage of firms in sector using AI in operations
	Supply chain and distribution	Supply chain management	Percentage of firms in sector using AI in supply chain management
		Distribution	Percentage of firms in sector using AI in distribution
	Customer experience	Customer services	Percentage of firms in sector using AI in customer services
		Sales and marketing	Percentage of firms in sector using AI in sales and marketing
	Financial and general management	Executive management	Percentage of firms in sector using AI in executive management
		Financial and risk management	Percentage of firms in sector using AI in financial and risk management
	Workforce management	Management of operational staff	Percentage of firms in sector using AI in operational staff management
HR		Percentage of firms in sector using AI in HR	
Labor	Exposure to AI in workforce	Percentage of workforce in firms adopting AI at scale or in core part of business	
	AI resources per worker	Average AI spend per employee (€ thousand)	

1 Results unweighted.

SOURCE: McKinsey Global Institute analysis

In order to calculate the overall AI score for an industry, a weight was assigned to each variable. Principal component analysis was used to determine the weights. The analysis is a mathematical transformation that converts a set of potentially correlated input variables into principal components, or new sets of values that explain the variance in the input variables. In this case, the resulting principal components aim to explain the variance across the 16 input variables in use. A principal component analysis yields multiple components, so the component that explains the most variance of the original 16 variables was used. Each component has corresponding variable loadings or weights, which were applied as the weights for each value in the index calculation. Since the 16 input variables are not in the same units, each value was standardized by subtracting from it the mean value for the variable and dividing the result by the standard deviation for the variable. The standardized scores for each variable were added to arrive at the overall AI index value for each sector. The analysis was performed within each of the three categories of variables: assets, usage, and labor.

BIBLIOGRAPHY

- Buchanan, Bruce G., "A (very) brief history of artificial intelligence," *AI Magazine*, volume 26, number 4, winter 2005.
- Bughin, Jacques, and James Manyika, "Bubble or paradigm change? Assessing the global diffusion of enterprise 2.0," in Alex Koohang, Johannes Britz, and Keith Harman, eds., *Knowledge management: Research and applications*, Informing Science, 2008.
- Bughin, Jacques, and James Manyika, "How businesses are using web 2.0: A McKinsey global survey," *McKinsey Quarterly*, December 2007.
- Bughin, Jacques, and Nicolas van Zeebroeck, "The right response to digital disruption," *MIT Sloan Management Review*, April 2017.
- Bughin, Jacques, "Ten big lessons learned from big data analytics," *Applied Marketing Analytics*, volume 2, number 4, 2017.
- Buitrago, Jaime, and Shihab Asfour, "Short-term forecasting of electric loads using nonlinear autoregressive artificial neural networks with exogenous vector inputs," *Energies*, volume 10, number 1, 2017.
- Catlin, Tanguy, Jay Scanlan, and Paul Willmott, "Raising your Digital Quotient," *McKinsey Quarterly*, June 2015.
- Dewan, Sanjeev, Dale Ganley, and Kenneth L. Kraemer, "Complementarities in the diffusion of personal computers and the internet: Implications for the global digital divide," *Information Systems Research*, volume 21, number 5, December 2010.
- Feigenbaum, Edward A., "Expert systems in the 1980s," working paper, 1980.
- Feigenbaum, Edward A., *The art of artificial intelligence: Themes and case studies of knowledge engineering*, Stanford University Computer Science Department report number STAN-CS-77-621, August 1977.
- Forman, Chris, Avi Goldfarb, and Shane Greenstein, "The geographic dispersion of commercial Internet use," in *Rethinking rights and regulations: Institutional responses to new communication technologies*, Lorrie Faith Cranor and Steven S. Wildman, eds., MIT Press, 2003.
- Glauner, Patrick, "Large-scale detection of non-technical losses in imbalanced data sets," PhD dissertation, University of Luxembourg, March 2016.
- Gugerty, Leo, "Newell and Simon's Logic Theorist: Historical background and impact on cognitive modeling," *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 50, issue 9, October 2006.
- Hofstadter, Douglas R., *Gödel, Escher, Bach: An eternal golden braid*, Basic Books, 1979.
- McGowan, Müge Adalet, and Dan Andrews, *Labor market mismatch and labor productivity: Evidence from PIAAC data*, OECD, April 28, 2015.
- McKinsey & Company, *Education to employment: Designing a system that works*, 2013.

McKinsey & Company, *Smartening up with artificial intelligence (AI): What's in it for Germany and its industrial sector?* April 2017.

McKinsey Center for Government, *Government productivity: Unlocking the \$3.5 trillion opportunity*, April 2017.

McKinsey Global Institute, *The age of analytics: Competing in a data-driven world*, December 2016.

McKinsey Global Institute, *Artificial intelligence: Implications for China*, April 2017.

McKinsey Global Institute, *Digital America: A tale of the haves and have-mores*, December 2015.

McKinsey Global Institute, *Digital Europe: Pushing the frontier, capturing the benefits*, June 2016.

McKinsey Global Institute, *A future that works: Automation, employment, and productivity*, January 2017.

McKinsey Global Institute, *Harnessing automation for a future that works*, 2017.

McKinsey Global Institute, *A labor market that works: Connecting talent with opportunity in the digital age*, June 2015.

McKinsey Global Institute, *Technology, jobs, and the future of work*, 2017.

Minsky, Marvin, "Steps toward artificial intelligence," *Proceedings of the IRE*, volume 49, number 1, January 1961.

MIT Technology Review, "Asia's AI agenda: How Asia is speeding up global artificial adoption," November 2016.

Moravec, Hans P., "The Stanford cart and the CMU Rover," *Proceedings of the IEEE*, volume 71, issue 7, July 1983.

Moreno, Rosina, and Jordi Suriñach, "Innovation adoption and productivity growth: Evidence for Europe," working paper, 2014.

Narrative Science, *Outlook on artificial intelligence in the enterprise*, 2016.

National Science and Technology Council, "Preparing for the future of artificial intelligence," Executive Office of the President, October 2016.

Negnevitsky, Michael, *Artificial Intelligence: A guide to intelligent systems*, Addison-Wesley, 2002.

Newell, Allen, "Intellectual issues in the history of artificial intelligence," in *The study of information: Interdisciplinary messages*, Fritz Machlup and Una Mansfield, eds., John Wiley and Sons, 1983.

O'Cathain, Alicia et al., "A system-wide approach to explaining variation in potentially avoidable emergency admissions: National ecological study," *BMJ Quality and Safety*, volume 23, issue 1, January 2014.

Satchwell, Andrew et al., *Financial impacts of net-metered PV on utilities and ratepayers: A scoping study of two prototypical U.S. utilities*, Lawrence Berkeley National Laboratory, September 2014.

Turing, A. M., "Computing machinery and intelligence," *Mind*, volume 49, number 236, October 1950.

Zastrow, Mark, "South Korea trumpets \$860-million AI fund after AlphaGo 'shock,'" *Nature*, March 18, 2016.

Zhu, Kevin, Kenneth L. Kraemer, and Sean Xu, "The process of innovation assimilation by firms in different countries: A technology diffusion perspective on e-business," *Management Science*, volume 52, number 10, October 2006.

RELATED MGI AND MCKINSEY RESEARCH

A future that works: Automation, employment, and productivity (January 2017)

Improvements in robotics, artificial intelligence, and machine learning have put us on the cusp of a new automation age. This will change work for everyone, from miners to bankers.

The age of analytics: Competing in a data-driven world (December 2016)

Data and analytics are advancing rapidly. Organizations that harness them will create significant value and differentiate themselves; those that don't will be at a disadvantage.

The rise of the machines: How Chinese executives think about developments in artificial intelligence (December 2016)

Software developers have approached the dream of teaching computers to think, aided by advances in computing power, the development of machine learning, and an explosion in data.

Digital globalization: The new era of global flows (March 2016)

Conventional wisdom says globalization has stalled because the global goods trade has flattened and cross-border capital flows have declined. But global flows of data and information are soaring.

The Internet of Things: Mapping the value beyond the hype (June 2015)

If policy makers and businesses get it right, linking the physical and digital worlds could generate up to \$11.1 trillion a year in economic value by 2025.

No ordinary disruption: The four global forces breaking all the trends (April 2015)

The global economy's operating system is being rewritten. Four great forces are colliding and transforming the global economy, and the world's leaders need to adjust to a new reality.

Smartening up with artificial intelligence: How AI will transform Germany's industrial sector (April 2017)

Artificial intelligence is finally bringing a multitude of capabilities to machines that were long thought to belong to humans. We explain how and where AI could affect the German industrial sector.

Artificial intelligence: Implications for China (April 2017)



Artificial intelligence has gone from speculation to reality. This paper explores AI's potential to fuel China's productivity growth—and to disrupt the nation's workforce.

Open data: Unlocking innovation and performance with liquid information (October 2013)

Open data—from both public and private sources—are adding a new dimension to analytics and giving rise to novel, data-driven innovations.



McKinsey Global Institute
June 2017
Copyright © McKinsey & Company
www.mckinsey.com/mgi

 @McKinsey_MGI
 McKinseyGlobalInstitute