



# Enterprise Adoption and Management of **ARTIFICIAL INTELLIGENCE**

**Thomas H. Davenport,**  
Babson College

*Artificial intelligence is the most important new technology of the age, but it comes in many varieties, and businesses face a range of challenges in effectively deploying it throughout their organizations. Tom Davenport takes a pragmatic but positive approach to AI's long-term potential, describing effective approaches to creating and implementing a strategy for this transformative technology.*

**A**rtificial intelligence (AI), often defined as technology that performs tasks which previously could only be done by the human brain, is currently viewed as the new technology most important and disruptive to large organizations. Many companies around the world are now exploring or implementing it, with varying degrees of enthusiasm. There is now sufficient evidence through which to evaluate its initial impact on organizational strategies, business models, product and service offerings, and employment. Overall, its effects are incremental and in keeping with previous analytical efforts. However, there is every reason to expect more dramatic effects over the long term,

particularly on companies that decide to make extensive and ambitious use of AI.

Results from various surveys, most of them conducted by consulting firms, suggest that between 20 and 37 percent of large companies globally are either adopting or experimenting with AI. Market researchers have found that a larger percentage—perhaps as many as 60 percent of large firms—now employ robotic process automation, the easiest form of AI to assimilate. In many of these firms, AI, particularly in the form of machine learning, is used to extend business analytics. However, some forms of AI have different capabilities.

### Technologies Adopted in the Enterprise

Artificial intelligence is not one technology, but several that are increasingly being used for specific applications. Three main resources underlie most AI, and they are employed in several different types of technology. These resources are statistical analysis, semantic or linguistic analysis, and logic, typically in the form of rules. Those tools support a variety of AI methods, which in turn drive AI applications (Figure 1).

### AI Methods

*Neural Networks, Deep Learning, Machine Learning*—Machine learning is used to create computer systems that improve themselves through experience, which is often codified in the form of training data. The machine learns by devising the best ways of fitting models to the training data. A diverse array of learning algorithms and statistical techniques has been developed to cover the wide variety of data and problems used to teach machines. As noted by the eminent researchers and pioneers of machine learning Professors Jordan and Mitchell, “Many developers of AI systems now recognize that, for many applications, it can be far easier to train a system by showing it examples of desired input-output behavior than to program it manually by anticipating the desired response for all possible inputs.”

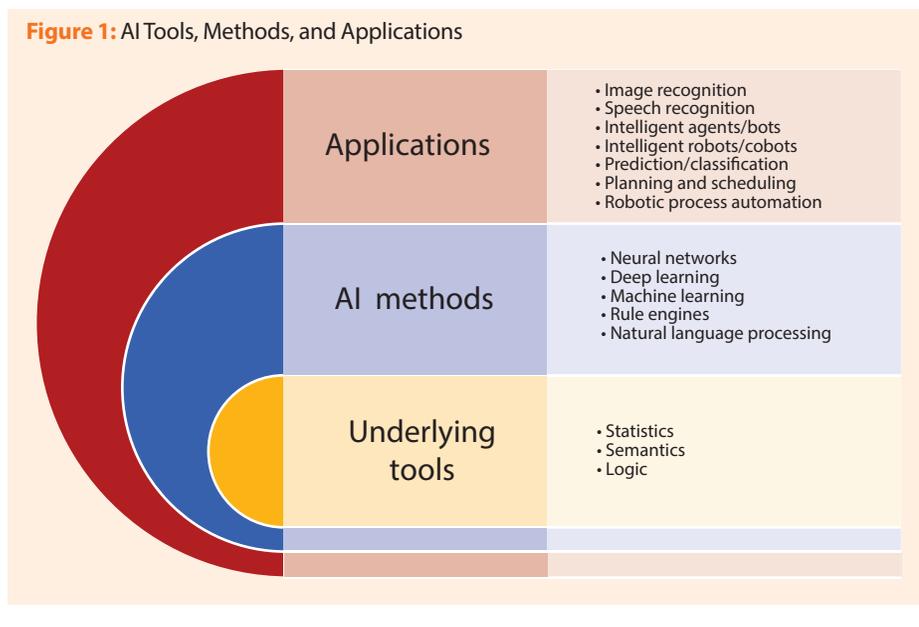
Over the past decade, machine learning has become one of the most commonly used forms of AI. A 2018 Deloitte survey of 1,100 U.S. managers whose organizations were pursuing AI found that 63 percent used machine learning. Indeed, the many forms of machine learning lie at the core of numerous approaches

In its most basic form, machine learning is synonymous with predictive analytics. Computers use models from data for which the results are known to predict the results from new data.

to artificial intelligence. In its most basic form, machine learning is synonymous with predictive analytics. Computers use models from data for which the results are known to predict the results from new data. AI systems have even been referred to as prediction machines. This method is called supervised learning, and it comprises the great majority of business applications for machine learning. Over the past several years, reinforcement learning, in which machine learning models are designed to maximize performance in support of a specified goal, has also become popular. Reinforcement learning is particularly common in games, but also has some business applications.

In commercial organizations, conventional machine learning is most commonly applied to marketing and sales (building models to predict consumer responses to marketing and determine optimal pricing), to human resources (HR) management (monitoring employee performance or predicting attrition), and to managing supply chains (predicting demand, required inventory, and more). In any of these applications, learning machines must start with at least a moderate amount of data to build a model. Almost all supervised machine learning is episodic and requires human intervention; the computers have to be retrained with new data if they are to improve their ability to predict.

Figure 1: AI Tools, Methods, and Applications



---

Some leading-edge vendor and user firms, though, are experimenting with models that learn continuously when new training data becomes available.

The *neural network* is more complex, a form of automated machine learning that has its roots in the 1940s and has steadily progressed ever since. Its developers made particular progress in the 1980's (with the invention of backpropagation, which streamlined the calculation of errors through multilayered neural nets and thus removed a major barrier to complex neural algorithms) and in the year 2012 (with the University of Toronto team's victory at Stanford's ImageNet competition for visual image recognition).

The most complex forms of machine learning involve deep learning, neural networks with many layers and many nodes per layer. These many levels permit the complexity of features or variables necessary to make predictions. Deep learning systems view problems in terms of inputs and outputs, and of the weighted variables or features that connect them. Thousands of hidden features may make up such models and are calculated by the extremely rapid processing of today's GPU (graphics processing unit) and cloud architectures. Deep learning is often used in image recognition, although it can also be used to increase the precision of applications like fraud detection, which would otherwise use conventional machine learning. Although deep learning has been likened to the processing of signals by neurons, the analogy holds only at a high conceptual level, not in terms of the actual details of computation or execution.

Deep learning is also increasingly used for speech recognition and other language-based analysis. It has therefore become an important tool for language processing. Unlike earlier forms of statistical analysis, the individual features of a deep

learning model have little meaning to a human observer, sometimes to the point of being impossible to interpret. As a result, highly regulated industries like banking, insurance, and healthcare find deep learning models problematic, though researchers are attempting to make them more transparent.

*Rule Engines and Rule-Based Systems*—Systems based on collections of if/then rules or those based on knowledge engineering, which tried to emulate human expertise, were the dominant technology for AI in the 1980s. They have since seen wide commercial use. Although these rule-based expert systems are no longer state of the art, one recent survey found that about half of large US firms still use them. They are quite reliable for tasks like insurance underwriting and providing clinical decision support in health care. These expert systems require human experts and knowledge engineers to construct a set of rules for each particular knowledge domain. The systems then perform their tasks well and are easy to understand. But when the number of rules grows, into the several thousands, and they begin to conflict with each other, these systems tend to break down. Moreover changing the rules as the field changes can be both difficult and time-consuming.

*Natural Language Processing (NLP)*—Making sense of human language has been a goal of AI researchers since the 1950s. This field, called natural language processing, includes speech recognition, text analysis, translation, and other tasks involving language. There are two basic approaches to NLP: semantic and statistical. Semantic NLP requires a graph or network of relationships between terms and phrases, which can be arduous to construct. It also requires the system to assess conversational intent. Semantic NLP is widely used for the intelligent

automated agents in customer service and call centers.

Statistical NLP is rooted in machine learning (particularly in increasingly deep learning neural networks) and is responsible for the recent increase in speech recognition accuracy. It requires a large body of language from which to learn. Statistical NLP is widely used for applications like machine translation, where there is a great deal of available data with which to train models.

## Applications of AI

*Image and Speech Recognition*—AI image and speech recognition have existed for many years, but new methods, particularly deep learning, have rapidly improved them. The accuracy of speech recognition in some applications is nearly 97 percent, while that of image recognition is well over the 95 percent that humans generally achieve. Both applications have improved considerably through the availability of large databases of speech and images.

*Intelligent Agents and Chatbots*—Hundreds of chatbots, intelligent agents, and intelligent assistants are now commercially available. They power smart speakers, smartphone assistants, and the basic interactions customers have with firms in a broad range of industries. Some employ standard rules and simple, natural language generators while others use sophisticated deep learning models. For standard transactions, they often augment human interactions, but they have only rarely replaced human agents in call centers.

*Prediction and Classification Systems*—Perhaps the oldest and most common form of AI is statistically based prediction or classification systems. Supervised machine learning, by far the most popular form in business, relies on data sets in which outcomes are used to train a model to predict or classify accurately.

---

These systems can then be used to score unknown results. While machine learning uses a variety of algorithms to predict or classify, in its simplest forms it is essentially indistinguishable from predictive analytical models, like credit scores, which have been used commercially since the 1950s.

*Planning and Scheduling*—Another enduring application of AI uses automated systems to achieve specific objectives in the physical or virtual world. Using methods that range from machine learning and reinforcement learning to rules or optimization techniques, these systems get autonomous vehicles to their destinations, manage workflow, or guide the actions of robots.

*Intelligent Robots and Cobots*—With almost 400,000 industrial robots installed globally each year, physical robots are well-known. In factories and warehouses, they perform defined tasks like lifting, repositioning, welding, or assembling objects. They distribute supplies in hospitals. Robotic mechanical arms, in concert with camera systems and other sensors, have been used in surgery for almost twenty years. Functioning under the guidance and control of a human surgeon, they are not, technically speaking, autonomous robots, but rather remotely operated systems. Nonetheless, these machine appendages expand the powers of surgeons dramatically, improving their ability to see, create precise incisions, sew tiny stitches, and more. A new generation of robots, dubbed cobots, are even better at collaborating with humans and are easily trained simply by moving them through a particular task. As broader AI capabilities are added to their brains, that is, their operating systems, robots are becoming ever more intelligent.

*Robotic Process Automation (RPA)*—This technology inhabits the gap between AI methods and the use of other technologies, some AI-based, and some not. RPA digitally

performs structured administrative tasks, using information systems, as if it were a human user following a script or set of rules. RPA is comparatively inexpensive, easy to program, and transparent in its actions. Its use is therefore growing very rapidly. The revenues of RPA vendors grew by 63 percent in 2018 and reached \$1.3 billion in 2019. Robotic process automation is actually a misnomer, since it doesn't involve any robots, only computer programs on ordinary servers. It relies on integrating workflow, business rules, and standardized data with information systems to act like a semi-intelligent user of the systems. In business it is generally applied to structured back-office processes including billing and customer service. When combined with other technologies, like image recognition, RPA can be used to extract data from, for example, faxed images and input them into transactional systems.<sup>7</sup>

While all these technologies are individual and separate, they can also be combined and integrated. Although they do not function like biological brains, AI is providing robots with brains in the sense that both cognitive automation systems and physical automation systems are becoming more capable, more adaptable and, in a limited sense, smarter. Machine learning-based image and text recognition are now being integrated with RPA. Perhaps in the future they will be so seamlessly integrated that artificial intelligence can be accurately discussed as a single technology. Today, however, the capabilities and applications of different AI technologies are sufficiently distinct that it is important for business leaders to understand the variations.

### How Aggressively Are Large Firms Adopting AI?

If a company wants to use AI to create a competitive advantage, it must adopt the technology broadly and

aggressively. Too often, companies create pilots or proofs of concept without planning how the technology will be deployed or fully understanding the depth and complexity of the process changes necessary to integrating it and realizing large scale benefits. Full production implementation of AI technology is relatively scarce outside of the largest and most capable firms, and for good reasons. One of these is the relative immaturity of the technology. Chatbots and intelligent agents, for example, are improving all the time, but they can still be an ordeal for customers and many companies hesitate to force them on their clients. Instead, companies ask their employees to use these applications in HR and IT, or make it easy for customers to opt out.

And if the AI requires changes to an existing process or new employee skills, that's another barrier, since the company must devise a plan to manage those changes. Most AI systems still need to interact with human workers, and teaching those workers new tasks and skills can be time-consuming and expensive. Surveyed workers often feel they are not being effectively trained to work with AI.

To be used to full advantage, AI must also be interfaced with production information systems and architectures. A 2017 Deloitte survey found that the number one obstacle to the successful deployment of AI was that it was "difficult to integrate cognitive projects with existing processes and systems." Even RPA systems, which are quite easy to set up for small volumes of robots, can become an architectural challenge when the volume increases. Moreover, because RPA applications act as users of production systems, they are sensitive to changes in those systems and may have to be reprogrammed.

The number of AI projects underway in a company serves as a

---

rough measure of the aggressiveness with which it is adopting AI. Of course, AI projects vary in scale and scope, and some projects call for several underlying AI applications. As AI matures and becomes more integrated into firms, the combined extent to which businesses have been successfully transformed may become a better measure of its use.

There are several ways for firms to count AI projects, and many do not have any accurate count, but in interviews with company executives I have been told that:

- Vanguard has only a few AI projects, in the single digits, but one with some AI is the very successful “robo-advisor” embedded in the company’s “Personal Advisor Services;”
- Intermountain Healthcare, a provider which has been quite aggressive with AI, has several dozen projects;
- Pfizer has over 150 projects in its pharmaceutical and life sciences business;
- Capital One, a pioneer in applying analytics to banking, has about 1000 AI projects; and
- Google/Alphabet had over 2700 machine learning projects in 2016 and has now stopped counting.

Clearly tech firms like Google are often the quickest and most successful at adopting AI. Facebook and Amazon, likewise, are known to have widely embraced it. Some banks, like Capital One, Bank of America, and JPMorgan Chase, have many AI projects underway as well. Yet other industries are only dipping their toes in the AI water or staying on the beach. Using the technology requires a fairly high degree of technical skill as well as a substantial amount of data, so business to business firms and small or medium companies are least likely to use AI.

## How Ambitious Should AI Projects Be?

Surveys suggest that many executives feel that AI will have a transformative impact on their businesses and industries. In a 2018 global survey, Deloitte found that 57 percent of executives believed that AI technology would substantially transform their companies within three years, and 38 percent believed that their industries would also be transformed. While these numbers are lower (and perhaps more realistic) than those in Deloitte’s 2017 survey, they still suggest high expectations.

---

### Even at highly tech-oriented firms like Amazon, ambitious AI projects often fail or dramatically exceed their budgets and timelines.

---

In order to fulfill those expectations, companies will have to enact large and highly ambitious projects. Early results, however, suggest that, even at tech-oriented firms like Amazon, ambitious AI projects often fail or dramatically exceed their budgets and timelines. Still, even companies whose grander projects have failed have succeeded with less ambitious ones.

How should organizations reconcile, then, their desire and expectation for AI to be transformative with the knowledge that simpler AI projects have a much higher success rate? One approach is to recognize the overall nature of AI, which tends to automate individual tasks rather than entire jobs or processes. Most successful AI projects will therefore be relatively small in scope. However, small projects can be combined to achieve greater impact. A company interested in transforming its customer experience could combine intelligent assistance and

recommendations with customer feedback analysis. Together they might produce a noticeable improvement in customer experience.

Many companies say that their primary objectives with AI concern improvement or innovation in products and services. In 2017 and 2018 Deloitte’s surveys found that the most common anticipated benefit from AI, at about 50 percent, was “enhancing the performance of products and services.” “Creating new products” was chosen by about 30 percent of respondents and “pursuing new markets” by about 25 percent. Products and services, then, offer extensive opportunities for companies to add AI. These innovation-oriented objectives may diminish somewhat compared to operational improvements or cost reduction in a more challenging economic climate.

Given the desire and expectation of executives that their business will be transformed, companies should take a strategic approach to determining how AI fits in their firm. An approach that allows them to think big, but start small, is likely to produce a higher success rate than plunging in at the deep end. Senior executives should discuss all possible effects of AI on their business, and create a high-level model of how their AI strategy should evolve, as well as whether it will involve internal processes, external offerings, or both.

The most aggressive AI strategies can also change existing business models. The world’s most valuable companies today (Microsoft, Alphabet/Google, Facebook, Amazon, Alibaba, Tencent, Ping An) all have business models which are to some extent based on a digital platform. These platforms, which connect buyers and sellers, work effectively only because of machine learning. This technology allows massive platforms to connect the right buyer with the right seller and

---

recommend suitable products and services. Their success may entice other firms with an interest in AI to move toward platform business models powered by machine learning. While success would almost certainly raise their equity valuations, changing a company's entire business model is complex and difficult, and adopting AI is a relatively small piece of the puzzle. Nevertheless, the employee benefits firm Benefit-Focus, has recently announced that it is enacting a "AI-driven platform pivot," a multi-year, multi-faceted change program.

### The Impact on the Workforce

One of the most common concerns about AI is that it will eliminate jobs. Yet so far almost none of the organizations in which I have conducted interviews have reported significant job cuts. Almost all say they are "freeing up human workers to do more creative or complex tasks" or something similar. Two technologies, though, have contributed to exceptions: industrial robots and robotic process automation. Two economists studied the impact of industrial robots on jobs. They found that, per thousand US workers, each robot replaced six humans and decreased wages by less than one percent.

Job losses from robotic process automation have yet to be accurately reported, but some companies do say they are planning or hoping to substantially reduce jobs. The Swiss bank UBS, for example, plans to use thousands of RPA robots to do work formerly performed by human office and information workers. Their stated goal is to reduce employment and thereby labor costs. A report created by the bank declared that the company is seeking cost savings through automation, and noted, "We make no secret that a certain portion of cost savings will come from reducing staff numbers." Likewise, some offshore outsourcing firms

have attributed their reduction of human jobs in offshore back office processes to RPA use by onshore clients.

Despite a few such admissions, though, there may be a conspiracy of silence around the topic of job loss due to automation. While many companies claim to be planning to augment their workforce with AI, rather than replacing it, there are signs that they too may instead be hoping that large-scale automation will allow them to cut human employment and save on labor costs. One report of unofficial conversations at the 2019 World Economic Forum in Davos suggests that executives privately hope and plan to cut jobs on a large scale. Similarly, a 2018 Deloitte survey of US executives familiar with their companies' AI initiatives found that 63 percent agreed that, "to cut costs, my company wants to automate as many jobs as possible with AI." Several vendors of AI technology have told me that, while they don't talk about it publicly, their customers are intent upon using AI to eliminate jobs. It seems likely that any economic recession would lead to more substantial job and cost reductions from AI in its various forms.

This objective is not surprising in a capitalist society, and companies have used technology to automate jobs for centuries. However, if AI does eliminate a substantial percentage of jobs or exacerbate inequality, there may be a major backlash against the technology.

### Key Trends in Enterprise AI

There are four current trends which are beginning to reshape the use of AI in large companies: embedding AI into transactional systems, democratization through automation, creation of AI centers of excellence and other management structures, and sparse data technologies. These innovations are easing what are now inherent necessities of using AI: integration with existing systems and processes,

highly skilled technicians, special abilities in every business unit and function, and large quantities of data, respectively.

*Embedding AI into Transactional Systems*—Many of the AI systems on the market or developed from scratch by organizations solve relatively isolated problems and are therefore stand-alone solutions. In order to be effectively deployed in large organizations, they need to be integrated with existing systems and processes. To help meet that need, many software vendors are beginning to offer systems with embedded AI capabilities.

For example, if a company wants to qualify and rank its sales leads by predicted likelihood of purchase (typically through some combination of machine learning and natural language processing), it has two choices; it can develop its own AI application to predict and score leads, and then try to integrate it with its customer relationship management (CRM) system, or it can buy the same capabilities from an established CRM vendor. Salesforce.com, for example, includes some AI capabilities, under the title of Einstein, that are already integrated with transaction software. And the company continues to invest substantially in embedding more AI capabilities into its products. Several other vendors, including SAP and Oracle, offer similar products. By using an existing vendor, a company can use all the data collected by its enterprise software, and users don't have to learn a new interface or system.

Companies are increasingly adopting their vendors' embedded capabilities. In a 2018 survey of U.S. executives from companies already using AI, 59 percent said they were employing enterprise software with AI capabilities, making it the most popular means of moving into AI. And more companies will surely use this approach as more enterprise

---

software vendors develop AI capabilities. Stand-alone AI software vendors will therefore be under increasing pressure to provide at least an interface to existing enterprise software.

*Democratization of Machine Learning Through Automation*—The scarcity of trained personnel, including data scientists and AI engineers, has always been a substantial constraint on the use of AI. Now, however, several vendors offer automated machine learning, which makes it possible for less skilled analysts—sometimes called citizen data scientists—to do more sophisticated work. Automation software is increasingly able to perform key tasks required for machine learning. These include some aspects of data preparation, feature engineering or variable transformation, exploring different algorithms, selecting the best model, writing program code or APIs for models, and explaining what factors are particularly important to a model. These systems not only allow less skilled users to direct machine learning, they can also make machine learning which *is* guided by sophisticated data scientists more productive.

This AutoML software, as it is often known, is available both from specialized startups like DataRobot and H2O.ai, and from more established cloud providers like Google and Microsoft. Some versions of AutoML focus on traditional statistical machine learning while others emphasize deep learning. Some are specifically aimed at citizen data scientists, and others at highly skilled professionals. Regardless of these variations in technology, AutoML is likely to hugely expand the use of machine learning by organizations and take the creation of statistical models into a post-algorithmic age. It may also free data scientists and quantitative analysts to devote themselves to managing change and designing processes for AI models.

*Creation of AI Centers of Excellence and Other Management Structures*—Ma-

ny executives are establishing dedicated organizational units to weave AI firmly into their businesses. As with other new technologies (e-commerce, analytics, and even blockchain), establishing a principal support group or center of excellence (CoE) provides workers with coordination and leadership. Companies are devoting considerable financial resources to AI, and the necessary skills and experience are too rare to leave scattered around the organization with little coordination or collaboration.

In the 2018 Deloitte survey of large firms using AI, 37 percent said they had already established such a unit. Deutsche Bank, JPMorgan Chase, Nielsen, Pfizer, Procter & Gamble, Anthem, and Farmers Insurance are among the non-tech firms that have centralized AI oversight groups. They create a vision and strategy for AI in the company, establish a prioritized list of uses, delineate the data resources needed, and manage relationships with external partners. In some cases, CoEs help to develop AI applications for specific units and functions. They also may assist with management structures and processes related to AI. The 2018 Deloitte survey also found that:

- 54 percent have created a process for moving AI prototypes into production;
- 52 percent have created a road map for using AI;
- 45 percent have appointed senior executives as AI champions; and
- 37 percent have put a comprehensive strategy for AI into effect.

There is no perfect home for an AI center of excellence, but the most common approach is to house it within a broader data and analytics department. The role of chief data and analytics officer, which then becomes responsible for AI, is well established in a plethora of companies including General Motors, JPMorgan Chase, Travelers, Wells Fargo, MetLife, Partners Healthcare, Marsh & McLennan, Walmart, CVS

Health, and many more. That office reports to a variety of executives, from chief of operations to marketing and financial officers, and even to the CEO, but not, generally, to the chief information officer.

*Sparser Data AI Models*—One of the key requirements for many AI systems, particularly deep learning models, is voluminous data. These data, their outcomes labeled, are used in supervised learning to train models. Less commonly, unlabeled data are used to discover patterns using unsupervised learning methods. But many companies can't assemble enough data, and labeling it is almost always labor-intensive. AI vendors are therefore working on a variety of technologies that make it possible to build high quality models and systems with smaller amounts of real world data.

Some of these systems, such as generative adversarial networks, create their own synthetic data. Other projects strive to give machines common sense. It's not clear which of these approaches will prevail, since they are currently the purview of research labs rather than commercial enterprises, but it seems likely that the data constraint on AI development and use will be eased before long.

## Whither AI in the Enterprise?

Artificial intelligence has waned and waxed in the past. It has lagged during times of disenchantment, only to burst forth in growth and hype. Now it seems unlikely that there are any more cold, fallow winters in its future. There are thousands of AI startups, many organizations which believe strongly and invest in AI, and enormous technological progress along multiple fronts. AI groups are increasingly well established in universities, research institutions, and companies. Even some governments, including those of China, Canada, the United Kingdom, and Singapore, have launched programs for developing AI and accelerating its use in industry.

However, we should not underestimate the time it will take for AI to become ubiquitous in business and society. Unlike microprocessors, AI is not an exponential technology, doubling in capacity every year or two over a sustained period of time. It is improving only at a linear rate. Even if AI technology improves quickly in some cases, regulation or organization must still be changed before it becomes truly useful. Autonomous vehicles driven by AI are a prime example. Many such vehicles are already being tested on public roads and some experts suggest that 90 percent of the requisite

technologies have already been developed. However, these experts also argue that the last 10 percent will require just as much effort and time as what has already been done.

**AI will no doubt become a revolutionary force in the fullness of time, but right now it is largely evolutionary.**

The greatest mistake that enterprises can make with AI is to expect too

much too soon. It will no doubt become a revolutionary force in the fullness of time, but right now it is largely evolutionary. As Amara's Law suggests, we are likely to overestimate AI in the short run and underestimate it in the long run. ■



**Thomas H. Davenport** is the President's Distinguished Professor of Information Technology and Management at Babson College, a Fellow at the MIT Initiative on the Digital Economy, and Senior Advisor to Deloitte Analytics. His most recent book is *The AI Advantage: How to Put the Artificial Intelligence Revolution to Work*.

## End Notes

1. See, for example, NewVantage Partners Big Data and AI Executive Survey 2019, Executive Summary of Findings, <https://newvantage.com/wp-content/uploads/2018/12/Big-Data-Executive-Survey-2019-Findings-Updated-010219-1.pdf>
2. "Gartner Says Worldwide Spending on Robotic Process Automation Software to Reach \$680 Million in 2018," Gartner Inc. press release, Nov. 13, 2018.
3. M. I. Jordan and T. M. Mitchell, "Machine Learning: Trends, Perspectives and Prospects," *Science*, 17 July 2015, Volume 349, Issue 6245, pages 255-260.
4. *Ibid*, page 255.
5. Eric Siegel, *Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die* (Wiley, 2013).
6. Ajay Agarwal, Joshua Gans, and Avi Goldfarb, *Prediction Machines: The Simple Economics of Artificial Intelligence*, Harvard Business Review Press, 2018.
7. Karen Hao, "We Analyzed 16,625 Papers to Figure Out Where AI Is Headed Next," *MIT Technology Review*, Jan. 25, 2019, <https://www.technologyreview.com/s/612768/we-analyzed-16625-papers-to-figure-out-where-ai-is-headed-next/>
8. A brief history of Neural Networks from the 1940's through the 1980's can be found at this Stanford Computer Science website: <https://cs.stanford.edu/people/eroberts/courses/soco/projects/neural-networks/History/history1.html> A more complete history, from the 1940's all the way to 2012 when multi-level neural networks, graphics processing units, and large training data sets were combined to win the Stanford ImageNet competition, is given at this Carnegie Mellon School of Computing website: [http://www.cs.cmu.edu/~10701/slides/Perceptron\\_Reading\\_Material.pdf](http://www.cs.cmu.edu/~10701/slides/Perceptron_Reading_Material.pdf)
9. See Terrence J. Sejnowski, *The Deep Learning Revolution*, MIT Press, 2018, for an overview of the evolution of deep learning, with an emphasis on how perspectives and insights from neuroscience and from computer science have developed in parallel in recent decades and have influenced one another. The book also explains some of the important differences between the function of the human brain and that of multi-layer neural networks.
10. Thomas H. Davenport, Jeff Loucks, and David Schatsky, "Bullish on the Business Value of Cognitive: The 2017 Deloitte State of Cognitive Survey," <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/deloitte-analytics/us-da-2017-deloitte-state-of-cognitive-survey.pdf>
11. For speech recognition, see Zhong Qiu Lin, Audrey G. Chung, and Alexander Wong, "EdgeSpeechNets: Highly Efficient Deep Neural Networks for Speech Recognition on the Edge," arXiv:1810.08559 [eess.AS], Nov. 13, 2018; for image recognition, see "Image Classification on ImageNet," <https://paperswithcode.com/sota/image-classification-on-imagenet>
12. "Executive Summary World Robotics 2018: Industrial Robots," International Federation of Robotics, [https://ifr.org/downloads/press2018/Executive\\_Summary\\_WR\\_2018\\_Industrial\\_Robots.pdf](https://ifr.org/downloads/press2018/Executive_Summary_WR_2018_Industrial_Robots.pdf)
13. "Gartner Says Worldwide Robotic Process Automation Software Market Grew 63% in 2018," Gartner Inc. press release, June 24, 2019, <https://www.gartner.com/en/newsroom/press-releases/2019-06-24-gartner-says-worldwide-robotic-process-automation-sof>
14. Thomas H. Davenport, "The Business Value of Digital Workflows," *Workflow Quarterly*, ServiceNow, Spring 2019, <https://workflow.servicenow.com/quarterly/issue/1/>
15. Davenport, Loucks, and Schatsky, op. cit.
16. Jeff Loucks et al, "Future in the balance: How countries are pursuing an AI advantage," *Deloitte*, May 1, 2019, <https://www2.deloitte.com/insights/us/en/focus/cognitive-technologies/ai-investment-by-country.html>
17. Thomas H. Davenport and Rajeev Ronanki, "Artificial Intelligence for the Real World," *Harvard Business Review*, Jan.-Feb. 2018, <https://hbr.org/2018/01/artificial-intelligence-for-the-real-world>
18. Barry Libert, Yoram Wind, and Megan Beck, "What Airbnb, Uber, and Alibaba Have in Common," *Harvard Business Review*, Nov. 20, 2014, <https://hbr.org/2014/11/what-airbnb-uber-and-alibaba-have-in-common>
19. Thomas H. Davenport, "AI and the 'Platform Pivot' at BenefitFocus," *Forbes*, Jan. 10, 2020, <https://www.forbes.com/sites/tomdavenport/2020/01/10/ai-and-the-platform-pivot-at-benefitfocus/#425039d06b1c>
20. Daron Acemoglu and Pascual Restrepo, "Robots and Jobs: Evidence from US Labor Markets," National Bureau of Economic Research working paper, March 2017, <https://www.nber.org/papers/w23285>
21. "Intelligent Automation," UBS Group Innovation White Paper, 2016, file:///Users/tdavenport/Downloads/ubs-intelligent-automation-wp.PDF
22. Venkatesh Ganesh, "Automation to Kill 70% of IT Jobs," *The Hindu BusinessLine*, January 9, 2018, <https://www.thehindubusinessline.com/info-tech/automation-to-kill-70-of-it-jobs/article9960555.ece>
23. Kevin Roose, "The Hidden Automation Agenda of the Davos Elite," *The New York Times*, Jan. 25, 2019, <https://www.nytimes.com/2019/01/25/technology/automation-davos-world-economic-forum.html>
24. Jeff Loucks, Tom Davenport, and David Schatsky, "State of AI in the Enterprise, 2nd Edition," *Deloitte*, 2018, [https://www2.deloitte.com/content/dam/insights/us/articles/4780\\_State-of-AI-in-the-enterprise/2nd-ed.pdf](https://www2.deloitte.com/content/dam/insights/us/articles/4780_State-of-AI-in-the-enterprise/2nd-ed.pdf)
25. Thomas H. Davenport and D.J. Patil, "Data Scientist: Sexiest Job of the 21st Century," *Harvard Business Review*, October 2012, <https://hbr.org/2012/10/data-scientist-the-sexiest-job-of-the-21st-century>
26. Clint Boulton, "The Age of the 'Citizen Data Scientist' Has Arrived," *CIO*, Nov. 26, 2018, <https://www.cio.com/article/3322940/the-age-of-the-citizen-data-scientist-has-arrived.html>
27. Bill Franks, "The Post-Algorithmic Era Has Arrived," *International Institute for Analytics*, Dec. 14, 2017, <https://www.iianalytics.com/blog/2017/12/14/the-post-algorithmic-era-has-arrived>
28. Thomas H. Davenport and Shivaji Dasgupta, "How to Set Up an AI Center of Excellence," *Harvard Business Review*, Jan. 19, 2019, <https://hbr.org/2019/01/how-to-set-up-an-ai-center-of-excellence>
29. H. James Wilson, Paul Daugherty, and Chase Davenport, "The Future of AI Will Be About Less Data, Not More," *Harvard Business Review*, Jan. 14, 2019, <https://hbr.org/2019/01/the-future-of-ai-will-be-about-less-data-not-more>
30. Connor Shorten, "Generative Models for Data Augmentation," *Medium: Towards Data Science*, Sept. 12, 2018, <https://towardsdatascience.com/generative-adversarial-networks-for-data-augmentation-experiment-design-2873d586eb59>
31. Thomas H. Davenport, "AI Is a Linear—Not Exponential—Technology," *International Institute for Analytics*, Dec. 20, 2018, <https://www.iianalytics.com/blog/2018/12/20/ai-is-a-linear-not-exponential-technology>
32. John D. Stoll, "Driverless Cars Are 90% Here. Another 90% Is Left to Go," *The Wall Street Journal*, May 24, 2019, <https://www.wsj.com/articles/driverless-cars-are-90-here-theres-another-90-left-to-go-11558717322>
33. Matt Ridley, "Don't Write Off the Next Big Thing too Soon," *The Times (London)*, Nov. 6, 2017, <https://www.thetimes.co.uk/article/dont-write-off-the-next-big-thing-too-soon-rbfzq9sck>