Insights You Need from Harvard Business Review

Business is changing. Will you adapt or be left behind?

Get up to speed and deepen your understanding of the topics that are shaping your company's future with the **Insights You Need from Harvard Business Review** series. Featuring HBR's smartest thinking on fast-moving issues—blockchain, cybersecurity, AI, and more—each book provides the foundational introduction and practical case studies your organization needs to compete today and collects the best research, interviews, and analysis to get it ready for tomorrow.

You can't afford to ignore how these issues will transform the landscape of business and society. The Insights You Need series will help you grasp these critical ideas—and prepare you and your company for the future.

Books in the series include:

Agile

Artificial Intelligence

Blockchain

Cybersecurity

Monopolies and Tech Giants

Strategic Analytics

Library of Congress Cataloging-in-Publication Data

Title: Artificial intelligence : the insights you need from Harvard Business Review.
Other titles: Artificial intelligence (Harvard Business Review Press) | Insights you need from Harvard Business Review.
Description: Boston, Massachusetts : Harvard Business Review Press, [2019] | Series: The insights you need from Harvard Business Review
Identifiers: LCCN 2019011886 | ISBN 9781633697898 (pbk.)
Subjects: LCSH: Artificial intelligence. | Artificial intelligence—Industrial applications. | Technological innovations—Economic aspects.
Classification: LCC Q335 .A778 2019 | DDC 658/.0563—dc23
LC record available at https://lccn.loc.gov/2019011886

ISBN: 978-1-63369-789-8
eISBN: 978-1-63369-790-4

Contents

Introduction

THE STATE OF AI IN BUSINESS

by Thomas H. Davenport

The most important general-purpose technology of our era is artificial intelligence. So Eric Brynjolfsson and Andrew McAfee describe AI in the first article of this book. But even as the significance of AI becomes irrefutable, it remains misunderstood. Executives view AI as a key disruptive technology, employees fear it as a job destroyer, consultants pitch it as a cure-all, and the media hype and deride it endlessly.

This book will help you tune out all this noise and understand AI's implications for you and your business. No matter your industry, level, or the size of your company, this collection of some of HBR's best recent articles on AI will show you where the technology is going.

Let's start with an overview of the state of AI in business today and its near-term implications.

AI *is* undoubtedly booming in business—at least in certain segments of it. In my work, I've helped design and analyze surveys suggesting that 25% to 30% of large U.S. companies are pursuing AI, many quite aggressively. Some have hundreds or even thousands of projects underway. The firms using AI most aggressively are large businesses with the most data—online platforms, financial services, telecommunications, and retail. Small- to medium-sized enterprises, business-to-business firms, and those in basic manufacturing industries are less likely to use AI. They typically lack not only the data to succeed with AI, but the expertise and awareness to pursue it effectively. Firms outside the United States are also pursuing AI at a slower pace, although there are aggressive adopters in China, the U.K., Canada, and Singapore.

A variety of different AI technologies are in use. You need to be aware of which ones do what. As Emma Martinho-Truswell explains in her article, machine learning is perhaps the most important component of AI, but it has multiple variations—ordinary statistical machine learning, neural networks, deep learning neural networks, and so on. Versions of AI also use semantic approaches to understanding language and logic-based rule engines for making simple decisions. Each technology performs a particular set of tasks; deep learning, for example, excels at recognizing images and speech.

AI is being applied for various business purposes. The most common uses enable us to make better decisions, improve operational processes, and enhance products and services. The first two are an extension of business analytics and typically employ machine learning; product-oriented objectives are common in high-tech firms, automobiles, and advanced manufacturing.

Many large companies are creating infrastructures and processes to manage AI. More than a third of large U.S. firms report in multiple surveys that they have a strategy in place for AI, have created a center of excellence to facilitate its use, and have identified its champions in the management team. As Vikram Mahidhar and I suggest in our article, late adopters may have difficulty catching up.

Companies are finding success by focusing their AI efforts in certain areas of the organization. Given the combination of short-term incremental value and long-term opportunity, many companies are tempering expectations about AI while still providing motivation to move forward aggressively with the technology. This is perhaps best accomplished by undertaking several projects focused in a particular area rather than spreading AI projects throughout the organization. Transforming customer service, for example, might include projects involving chatbots, intelligent agents, recommendation engines, and so forth.

AI hasn't transformed business—yet. While surveys suggest high expectations for transformation and high percentages of respondents say they have achieved economic returns, there are few examples of sweeping business reinvention thus far, for several reasons:

- It's still early in the life cycle of AI activity.
- Not every company has data that's suited for AI, as Ajay Agrawal, Joshua Gans, and Avi Goldfarb explain in their article (although H. James Wilson, Paul Daugherty, and my son Chase Davenport suggest in their article that data requirements for effective AI may lessen in the future).
- Companies are undertaking pilots with AI rather than production deployments, as Andrew Ng recommends in his piece.
- AI tends to be a narrow technology that supports particular tasks, not entire jobs or processes.
- Highly ambitious moon-shot projects, such as treating cancer, enabling autonomous vehicles, and powering drone deliveries, have been unsuccessful or slow to arrive.

Even at data powerhouses like Amazon, most AI activity has involved projects that "quietly but meaningfully improve core operations," according to CEO Jeff Bezos in his 2017 letter to shareholders. It's an evolutionary set of improvements that will eventually add up to revolution.

AI's overall impact on employment isn't certain, but jobs will clearly change. Some observers have predicted dire levels of AI-driven unemployment. Thus far—as Wilson and Daugherty discuss in their article on "collaborative intelligence"—augmentation of human work by smart machines has been far more common than large-scale automation. Therefore, according to Mark Knickrehm, organizations need to begin preparing employees to work alongside smart machines and add value to their efforts.

Implementing AI raises ethical questions. Other articles in the book, including one by Roman Yampolskiy, suggest that it's not too early to consider the ethical concerns around AI. Algorithmic bias and lack of transparency are two critical issues that AI exacerbates. These powerful technologies have powerful implications for the workplace and the broader society.

With these fundamentals covered, it's time to dive into the articles. To best understand how AI will impact your company's situation, consider these questions as you read:

- Which particular AI technologies have the greatest potential benefit to your organization?
- How might those technologies enable new strategies, business models, or business process designs?
- What data resources do you have—or might you obtain—in order to power AI projects?
- How do you anticipate that AI will impact your workforce, and how can you begin to prepare employees to augment AI capabilities?

If you and your organization haven't already confronted these questions, let this book spark conversations. Think about how the right AI initiative could help your division perform better or make you more efficient at your own job. Simply asking the questions may be the first step in starting your company down the path of transformation.

As a professor and a consultant on information technology and business, I've spent the past several decades watching AI alternate from spring blooms to winter doldrums. This time is different. AI is deeply ensconced in business and is starting to bring about exciting change. Now, it appears that winter will not return.

Section 1

UNDERSTANDING AI AND MACHINE LEARNING

1

THE BUSINESS OF ARTIFICIAL INTELLIGENCE

by Erik Brynjolfsson and Andrew McAfee

For more than 250 years the fundamental drivers of economic growth have been technological innovations. The most important of these are what economists call general-purpose technologies—a category that includes the steam engine, electricity, and the internal combustion engine. Each one catalyzed waves of complementary innovations and opportunities. The internal combustion engine, for example, gave rise to cars, trucks, airplanes, chain saws, and lawnmowers, along with big-box retailers, shopping centers, cross-docking warehouses, new supply chains, and, when you think about it, suburbs. Companies as diverse as Walmart, UPS, and Uber found ways to leverage the technology to create profitable new business models.

The most important general-purpose technology of our era is artificial intelligence, particularly machine learning (ML)—that is, the machine's ability to keep improving its performance without humans having to explain exactly how to accomplish all the tasks it's given. Within just the past few years machine learning has become far more effective and widely available. We can now build systems that learn how to perform tasks on their own.

Why is this such a big deal? Two reasons. First, we humans know more than we can tell: We can't explain exactly how we're able to do a lot of things—from recognizing a face to making a smart move in the ancient Asian strategy game of Go. Prior to ML, this inability to articulate our own knowledge meant that we couldn't automate many tasks. Now we can.

Second, ML systems are often excellent learners. They can achieve superhuman performance in a wide range of activities, including detecting fraud and diagnosing disease. Excellent digital learners are being deployed across the economy, and their impact will be profound.

In the sphere of business, AI is poised to have a transformational impact, on the scale of earlier general-purpose technologies. Although it is already in use in thousands of companies around the world, most big opportunities have not yet been tapped. The effects of AI will be magnified in the coming decade, as manufacturing, retailing, transportation, finance, health care, law, advertising, insurance, entertainment, education, and virtually every other industry transform their core processes and business models to take advantage of machine learning. The bottleneck now is in management, implementation, and business imagination.

Like so many other new technologies, however, AI has generated lots of unrealistic expectations. We see business plans liberally sprinkled with references to machine learning, neural nets, and other forms of the technology, with little connection to its real capabilities. Simply calling a dating site "AI-powered," for example, doesn't make it any more effective, but it might help with fund-raising. This article will cut through the noise to describe the real potential of AI, its practical implications, and the barriers to its adoption.

What Can AI Do Today?

The term *artificial intelligence* was coined in 1955 by John McCarthy, a math professor at Dartmouth who organized the seminal conference on the topic the following year. Ever since, perhaps in part because of its evocative name, the field has given rise to more than its share of fantastic claims and promises. In 1957 the economist Herbert Simon predicted that computers would beat humans at chess within 10 years. (It took 40.) In 1967 the cognitive scientist Marvin Minsky said, "Within a generation the problem of creating 'artificial intelligence' will be substantially solved." Simon and Minsky were both intellectual giants, but they erred badly. Thus it's understandable that dramatic claims about future breakthroughs meet with a certain amount of skepticism.

Let's start by exploring what AI is already doing and how quickly it is improving. The biggest advances have been in two broad areas: perception and cognition. In the former category some of the most practical advances have been made in relation to speech. Voice recognition is still far from perfect, but millions of people are now using it—think Siri, Alexa, and Google Assistant. The text you are now reading was originally dictated to a computer and transcribed with sufficient accuracy to make it faster than typing. A study by the Stanford computer scientist James Landay and colleagues found that speech recognition is now about three times as fast, on average, as typing on a cell phone. The error rate, once 8.5%, has dropped to 4.9%. What's striking is that this substantial improvement has come not over the past 10 years but just since the summer of 2016.

Image recognition, too, has improved dramatically. You may have noticed that Facebook and other apps now recognize many of your friends' faces in posted photos and prompt you to tag them with their names. An app running on your smartphone will recognize virtually any bird in the wild. Image recognition is even replacing ID cards at corporate headquarters. Vision systems, such as those used in self-driving cars, formerly made a mistake when identifying a pedestrian as often as once in 30 frames (the cameras in these systems record about 30 frames a second); now they err less often than once in 30 million frames. The error rate for recognizing images from a large database called ImageNet, with several million photographs of common, obscure, or downright weird images, fell from higher than 30% in 2010 to about 4% in 2016 for the best systems. (See figure 1-1.)

FIGURE 1-1

Machines have made real strides in distinguishing among similar-looking categories of images

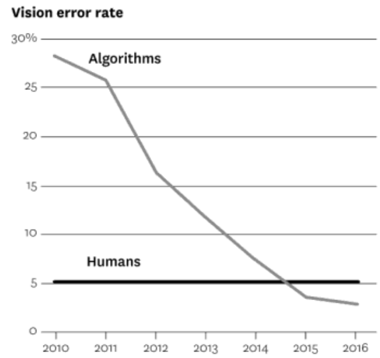

Source: Electronic Frontier Foundation

The speed of improvement has accelerated rapidly in recent years as a new approach, based on very large or "deep" neural nets, was adopted. The ML approach for vision systems is still far from flawless—but even people have trouble quickly recognizing puppies' faces or, more embarrassingly, see their cute faces where none exist.

The second type of major improvement has been in cognition and problem solving. Machines have already beaten the finest (human) players of poker and Go—achievements that experts had predicted would take at least another decade. Google's DeepMind team has used ML systems to improve the cooling efficiency at data centers by more than 15%, even after they were optimized by human experts. Intelligent agents are being used by the cybersecurity company Deep Instinct to detect malware, and by PayPal to prevent money laundering. A system using IBM technology automates the claims process at an insurance company in Singapore, and a system from Lumidatum, a data science platform firm, offers timely advice to improve customer support. Dozens of companies are using ML to decide which trades to execute on Wall Street, and more and more credit decisions are made with its help. Amazon employs ML to optimize inventory and improve product recommendations to customers. Infinite Analytics developed one ML system to predict whether a user would click on a particular ad, improving online ad placement for a global consumer packaged goods company, and another to improve customers' search and discovery process at a Brazilian online retailer. The first system increased advertising ROI threefold, and the second resulted in a $125 million increase in annual revenue.

Machine learning systems are not only replacing older algorithms in many applications, but are now superior at many tasks that were once done best by humans. Although the systems are far from perfect, their error rate—about 5%—on the ImageNet database is at or better than human-level performance. Voice recognition, too, even in noisy environments, is now nearly equal to human performance. Reaching this threshold opens up vast new possibilities for transforming the workplace and the economy. Once AI-based systems surpass human performance at a given task, they are much likelier to spread quickly. For instance, Aptonomy and Sanbot, makers respectively of drones and robots, are using improved vision systems to automate much of the work of security guards. The software company Affectiva, among others, is using them to recognize emotions such as joy, surprise, and anger in focus groups. And Enlitic is one of several deep-learning startups that use them to scan medical images to help diagnose cancer.

These are impressive achievements, but the applicability of AI-based systems is still quite narrow. For instance, their remarkable performance on the ImageNet database, even with its millions of images, doesn't always translate into similar success "in the wild," where lighting conditions, angles, image resolution, and context may be very different. More fundamentally, we can marvel at a system that understands Chinese speech and translates it into English, but we don't expect such a system to know what a particular Chinese character means—let alone where to eat in Beijing. If someone performs a task well, it's natural to assume that the person has some competence in related tasks. But ML systems are trained to do specific tasks, and typically their knowledge does not generalize. The fallacy that a computer's narrow understanding implies broader understanding is perhaps the biggest source of confusion, and exaggerated claims, about AI's progress. We are far from machines that exhibit general intelligence across diverse domains.

Understanding Machine Learning

The most important thing to understand about ML is that it represents a fundamentally different approach to creating software: The machine learns from examples, rather than being explicitly programmed for a particular outcome. This is an important break from previous practice. For most of the past 50 years, advances in information technology and its applications have focused on codifying existing knowledge and procedures and embedding them in machines. Indeed, the term "coding" denotes the painstaking process of transferring knowledge from developers' heads into a form that machines can understand and execute. This approach has a fundamental weakness: Much of the knowledge we all have is tacit, meaning that we can't fully explain it. It's nearly impossible for us to write down instructions that would enable another person to learn how to ride a bike or to recognize a friend's face.

In other words, we all know more than we can tell. This fact is so important that it has a name: Polanyi's Paradox, for the philosopher and polymath Michael Polanyi, who described it in 1964. Polanyi's Paradox not only limits what we can tell one another but has historically placed a fundamental restriction on our ability to endow machines with intelligence. For a long time that limited the activities that machines could productively perform in the economy.

Machine learning is overcoming those limits. In this second wave of the second machine age, machines built by humans are learning from examples and using structured feedback to solve on their own problems such as Polanyi's classic one of recognizing a face.

Different Flavors of Machine Learning

Artificial intelligence and machine learning come in many flavors, but most of the successes in recent years have been in one category: supervised learning systems, in which the machine is given lots of examples of the correct answer to a particular problem. This process almost always involves mapping from a set of inputs, X, to a set of outputs, Y. For instance, the inputs might be pictures of various animals, and the correct outputs might be labels for those animals: dog, cat, horse. The inputs could also be waveforms from a sound recording and the outputs could be words: "yes," "no," "hello," "good-bye." (See table 1-1.)

TABLE 1-1

Supervised learning systems

As two pioneers in the field, Tom Mitchell and Michael I. Jordan, have noted, most of the recent progress in machine learning involves mapping from a set of inputs to a set of outputs. Some examples:

Input X	Output Y	Application
Voice recording	Transcript	Speech recognition
Historical market data	Future market data	Trading bots
Photograph	Caption	Image tagging
Drug chemical properties	Treatment efficacy	Pharma R&D
Store transaction details	Is the transaction fraudulent?	Fraud detection
Recipe ingredients	Customer reviews	Food recommendations
Purchase histories	Future purchase behavior	Customer retention
Car locations and speed	Traffic flow	Traffic lights
Faces	Names	Face recognition

Successful systems often use a training set of data with thousands or even millions of examples, each of which has been labeled with the correct answer. The system can then be let loose to look at new examples. If the training has gone well, the system will predict answers with a high rate of accuracy.

The algorithms that have driven much of this success depend on an approach called *deep learning*, which uses neural networks. Deep learning algorithms have a significant advantage over earlier generations of ML algorithms: They can make better use of much larger data sets. The old systems would improve as the number of examples in the training data grew, but only up to a point, after which additional data didn't lead to better predictions. According to Andrew Ng, one of the giants of the field, deep neural nets don't seem to level off in this way: More data leads to better and better predictions. Some very large systems are trained by using 36 million examples or more. Of course, working with extremely large data sets requires more and more processing power, which is one reason the very big systems are often run on supercomputers or specialized computer architectures.

Any situation in which you have a lot of data on behavior and are trying to predict an outcome is a potential application for supervised learning systems. Jeff Wilke, who leads Amazon's consumer business, says that supervised learning systems have largely replaced the memory-based filtering algorithms that were used to make personalized recommendations to customers. In other cases, classic algorithms for setting inventory levels and optimizing supply chains have been replaced by more efficient and robust systems based on machine learning. JPMorgan Chase introduced a system for reviewing commercial loan contracts; work that used to take loan officers 360,000 hours can now be done in a few seconds. And supervised learning systems are now being used to diagnose skin cancer. These are just a few examples.

It's comparatively straightforward to label a body of data and use it to train a supervised learner; that's why supervised ML systems are more common than *un*supervised ones, at least for now. Unsupervised learning systems seek to learn on their own. We humans are excellent unsupervised learners: We pick up most of our knowledge of the world (such as how to recognize a tree) with little or no labeled data. But it is exceedingly difficult to develop a successful machine learning system that works this way.

If and when we learn to build robust unsupervised learners, exciting possibilities will open up. These machines could look at complex problems in fresh ways to help us discover patterns—in the spread of diseases, in price moves across securities in a market, in customers' purchase behaviors, and so on—that we are currently unaware of. Such possibilities lead Yann LeCun, the head of AI research at Facebook and a professor at NYU, to compare supervised learning systems to the frosting on the cake and unsupervised learning to the cake itself.

Another small but growing area within the field is *reinforcement learning*. This approach is embedded in systems that have mastered Atari video games and board games like Go. It is also helping to optimize data center power usage and to develop trading strategies for the stock market. Robots created by Kindred use machine learning to identify and sort objects they've never encountered before, speeding up the "pick and place" process in distribution centers for consumer goods. In reinforcement learning systems the programmer specifies the current state of the system and the goal, lists allowable actions, and describes the elements of the environment that constrain the outcomes for each of those actions. Using the allowable actions, the system has to figure out how to get as close to the goal as possible. These systems work well when humans can specify the goal but not necessarily how to get there. For instance, Microsoft used reinforcement learning to select headlines for MSN.com news stories by "rewarding" the system with a higher score when more visitors clicked on the link. The system tried to maximize its score on the basis of the rules its designers gave it. Of course, this means that a reinforcement learning system will optimize for the goal you explicitly reward, not necessarily the goal you really care about (such as lifetime customer value), so specifying the goal correctly and clearly is critical.

Putting Machine Learning to Work

There are three pieces of good news for organizations looking to put ML to use today. First, AI skills are spreading quickly. The world still has not nearly enough data scientists and machine learning experts, but the demand for them is being met by online educational resources as well as by universities. The best of these, including Udacity, Coursera, and fast.ai, do much more than teach introductory concepts; they can actually get smart, motivated students to the point of being able to create industrial-grade ML deployments. In addition to training their own people, interested companies can use online talent platforms such as Upwork, Topcoder, and Kaggle to find ML experts with verifiable expertise.

The second welcome development is that the necessary algorithms and hardware for modern AI can be bought or rented as needed. Google, Amazon, Microsoft, Salesforce, and other companies are making powerful ML infrastructure available via the cloud. The cutthroat competition among these rivals means that companies that want to experiment with or deploy ML will see more and more capabilities available at ever-lower prices over time.

The final piece of good news, and probably the most underappreciated, is that you may not need all that much data to start making productive use of ML. The performance of most machine learning systems improves as they're given more data to work with, so it seems logical to conclude that the company with the most data will win. That might be the case if "win" means "dominate the global market for a single application such as ad targeting or speech recognition." But if success is defined instead as significantly improving performance, then sufficient data is often surprisingly easy to obtain.

For example, Udacity cofounder Sebastian Thrun noticed that some of his salespeople were much more effective than others when replying to inbound queries in a chat room. Thrun and his graduate student Zayd Enam realized that their chat room logs were essentially a set of labeled training data—exactly what a supervised learning system needs. Interactions that led to a sale were labeled successes, and all others were labeled failures. Zayd used the data to predict what answers successful salespeople were likely to give in response to certain very common inquiries and then shared those predictions with the other salespeople to nudge them toward better performance. After 1,000 training cycles, the salespeople had increased their effectiveness by 54% and were able to serve twice as many customers at a time.

The AI startup WorkFusion takes a similar approach. It works with companies to bring higher levels of automation to back-office processes such as paying international invoices and settling large trades between financial institutions. The reason these processes haven't been automated yet is that they're complicated; relevant information isn't always presented the same way every time ("How do we know what currency they're talking about?"), and some interpretation and judgment are necessary. WorkFusion's software watches in the background as people do their work and uses their actions as training data for the cognitive task of classification ("This invoice is in dollars. This one is in yen. This one is in euros …"). Once the system is confident enough in its classifications, it takes over the process.

Machine learning is driving changes at three levels: tasks and occupations, business processes, and business models. An example of task-and-occupation redesign is the use of machine vision systems to identify potential cancer cells—freeing up radiologists to focus on truly critical cases, to communicate with patients, and to coordinate with other physicians. An example of process redesign is the reinvention of the workflow and layout of Amazon fulfillment centers after the introduction of robots and optimization algorithms based on machine learning. Similarly, business models need to be rethought to take advantage of ML systems that can intelligently recommend music or movies in a personalized way. Instead of selling songs à la carte on the basis of consumer choices, a better model might offer a subscription to a personalized station that predicted and played music a particular customer would like, even if the person had never heard it before.

Note that machine learning systems hardly ever replace the entire job, process, or business model. Most often they complement human activities, which can make their work ever more valuable. The most effective rule for the new division of labor is rarely, if ever, "give all tasks to the machine." Instead, if the successful completion of a process requires 10 steps, one or two of them may become automated while the rest become more valuable for humans to do. For instance, the chat room sales support system at Udacity didn't try to build a bot that could take over all the conversations; rather, it advised human salespeople about how to improve their performance. The humans remained in charge but became vastly more effective and efficient. This approach is usually much more feasible than trying to design machines that can do everything humans can do. It often leads to better, more satisfying work for the people involved and ultimately to a better outcome for customers.

Designing and implementing new combinations of technologies, human skills, and capital assets to meet customers' needs requires large-scale creativity and planning. It is a task that machines are not very good at. That makes being an entrepreneur or a business manager one of society's most rewarding jobs in the age of ML.

Risks and Limits

The second wave of the second machine age brings with it new risks. In particular, machine learning systems often have low "interpretability," meaning that humans have difficulty figuring out how the systems reached their decisions. Deep neural networks may have hundreds of millions of connections, each of which contributes a small amount to the ultimate decision. As a result, these systems' predictions tend to resist simple, clear explanation. Unlike humans, machines are not (yet!) good storytellers. They can't always give a rationale for why a particular applicant was accepted or rejected for a job, or a particular medicine was recommended. Ironically, even as we have begun to overcome Polanyi's Paradox, we're facing a kind of reverse version: Machines know more than they can tell us.

This creates three risks. First, the machines may have hidden biases, derived not from any intent of the designer but from the data provided to train the system. For instance, if a system learns which job applicants to accept for an interview by using a data set of decisions made by human recruiters in the past, it may inadvertently learn to perpetuate their racial, gender, ethnic, or other biases. Moreover, these biases may not appear as an explicit rule but, rather, be embedded in subtle interactions among the thousands of factors considered.

A second risk is that, unlike traditional systems built on explicit logic rules, neural network systems deal with statistical truths rather than literal truths. That can make it difficult, if not impossible, to prove with complete certainty that the system will work in all cases—especially in situations that weren't represented in the training data. Lack of verifiability can be a concern in mission-critical applications, such as controlling a nuclear power plant, or when life-or-death decisions are involved.

Third, when the ML system does make errors, as it almost inevitably will, diagnosing and correcting exactly what's going wrong can be difficult. The underlying structure that led to the solution can be unimaginably complex, and the solution may be far from optimal if the conditions under which the system was trained change.

While all these risks are very real, the appropriate benchmark is not perfection but the best available alternative. After all, we humans, too, have biases, make mistakes, and have trouble explaining truthfully how we arrived at a particular decision. The advantage of machine-based systems is that they can be improved over time and will give consistent answers when presented with the same data.

Does that mean there is no limit to what artificial intelligence and machine learning can do? Perception and cognition cover a great deal of territory—from driving a car to forecasting sales to deciding whom to hire or promote. We believe the chances are excellent that AI will soon reach superhuman levels of performance in most or all of these areas. So what *won't* AI and ML be able to do?

We sometimes hear "Artificial intelligence will never be good at assessing emotional, crafty, sly, inconsistent human beings—it's too rigid and impersonal for that." We don't agree. ML systems like those at Affectiva are already at or beyond human-level performance in discerning a person's emotional state on the basis of tone of voice or facial expression. Other systems can infer when even the world's best poker players are bluffing well enough to beat them at the amazingly complex game Heads-Up No-Limit Texas Hold'em. Reading people accurately is subtle work, but it's not magic. It requires perception and cognition—exactly the areas in which ML is currently strong and getting stronger all the time.

A great place to start a discussion of the limits of AI is with Pablo Picasso's observation about computers: "But they are useless. They can only give you answers." They're actually far from useless, as ML's recent triumphs show, but Picasso's observation still provides insight. Computers are devices for answering questions, not for posing them. That means entrepreneurs, innovators, scientists, creators, and other kinds of people who figure out what problem or opportunity to tackle next, or what new territory to explore, will continue to be essential.

Similarly, there's a huge difference between passively assessing someone's mental state or morale and actively working to change it. ML systems are getting quite good at the former but remain well behind us at the latter. We humans are a deeply social species; other humans, not machines, are best at tapping into social drives such as compassion, pride, solidarity, and shame in order to persuade, motivate, and inspire. In 2014 the TED Conference and the XPRIZE Foundation announced an award for "the first artificial intelligence to come to this stage and give a TED Talk compelling enough to win a standing ovation from the audience." We doubt the award will be claimed anytime soon.

We think the biggest and most important opportunities for human smarts in this new age of superpowerful ML lie at the intersection of two areas: figuring out what problems to work on next, and persuading a lot of people to tackle them and go along with the solutions. This is a decent definition of leadership, which is becoming much more important in the second machine age.

The status quo of dividing up work between minds and machines is falling apart very quickly. Companies that stick with it are going to find themselves at an ever-greater competitive disadvantage compared with rivals who are willing and able to put ML to use in all the places where it is appropriate and who can figure out how to effectively integrate its capabilities with humanity's.

A time of tectonic change in the business world has begun, brought on by technological progress. As was the case with steam power and electricity, it's not access to the new technologies themselves, or even to the best technologists, that separates winners from losers. Instead, it's innovators who are open-minded enough to see past the status quo and envision very different approaches, and savvy enough to put them into place. One of machine learning's greatest legacies may well be the creation of a new generation of business leaders.

In our view, artificial intelligence, especially machine learning, is the most important general-purpose technology of our era. The impact of these innovations on business and the economy will be reflected not only in their direct contributions but also in their ability to enable and inspire complementary innovations. New products and processes are being made possible by better vision systems, speech recognition, intelligent problem solving, and many other capabilities that machine learning delivers.

Some experts have gone even further. Gil Pratt, who now heads the Toyota Research Institute, has compared the current wave of AI technology to the Cambrian explosion 500 million years ago that birthed a tremendous variety of new life forms. Then as now, one of the key new capabilities was vision. When animals first gained this capability, it allowed them to explore the environment far more effectively; that catalyzed an enormous increase in the number of species, both predators and prey, and in the range of ecological niches that were filled. Today as well we expect to see a variety of new products, services, processes, and organizational forms and also numerous extinctions. There will certainly be some weird failures along with unexpected successes.

Although it is hard to predict exactly which companies will dominate in the new environment, a general principle is clear: The most nimble and adaptable companies and executives will thrive. Organizations that can rapidly sense and respond to opportunities will seize the advantage in the AI-enabled landscape. So the successful strategy is to be willing to experiment and learn quickly. If managers aren't ramping up experiments in the area of machine learning, they aren't doing their job. Over the next decade, AI won't replace managers, but managers who use AI will replace those who don't.

TAKEAWAYS

The most important new general-purpose technology is artificial intelligence, particularly machine learning. ML systems are replacing older algorithms in many applications and are now superior at many tasks previously done best by humans.

- ✓ Machine learning is fundamentally different from the software that preceded it: The machine learns from examples, rather than being explicitly programmed for a particular outcome.
- ✓ Organizations looking to put ML to use should be aware that AI skills are spreading quickly; the necessary algorithms and hardware for modern AI can be bought or rented as needed; and they may not need much data to start using ML productively.
- ✓ ML systems have low "interpretability," meaning that humans have difficulty figuring out how the systems reach their decisions. This creates three risks: The machines may have hidden biases; it is often impossible to prove that an ML system will work in all mission-critical situations; and when the ML system does make errors, diagnosing the problem and correcting it can be difficult.
- ✓ Companies that continue to divide up work between minds and machines will increasingly lose their competitive advantage to rivals that effectively integrate AI's capabilities with human capabilities.

Adapted from content posted on hbr.org, August 7, 2017 (product #BG1704).

2

INSIDE FACEBOOK'S AI WORKSHOP

An interview with Joaquin Candela by Scott Berinato

Within Facebook's cavernous Building 20, about halfway between the lobby (panoramic views of the Ravenswood Slough) and the kitchen (hot breakfast, smoothies, gourmet coffee), in a small conference room called Lollapalooza, Joaquin Candela is trying to explain artificial intelligence to a layperson.

Candela—bald, compact, thoughtful—runs Facebook's Applied Machine Learning (AML) group, the engine room of AI at Facebook, which increasingly makes it the engine room of Facebook in general. After some verbal searching, he finally says:

> *Look, a machine learning algorithm really is a lookup table, right? Where the key is the input, like an image, and the value is the label for the input, like "a horse." I have a bunch of examples of something. Pictures of horses. I give the algorithm as many as I can. "This is a horse. This is a horse. This isn't a horse. This is a horse." And the algorithm keeps those in a table. Then, if a new example comes along—or if I tell it to watch for new examples—well, the algorithm just goes and looks at all those examples we fed it. Which rows in the table look similar? And how similar? It's trying to decide, "Is this new thing a horse? I think so." If it's right, the image gets put in the "This is a horse" group, and if it's wrong, it gets put in the "This isn't a horse" group. Next time, it has more data to look up.*
>
> *One challenge is how do we decide how similar a new picture is to the ones stored in the table. One aspect of machine learning is to learn similarity functions. Another challenge is, What happens when your table grows really large? For every new image, you would need to make a zillion comparisons … So another aspect of machine learning is to approximate a large stored table with a function instead of going through every image. The function knows how to roughly estimate what the corresponding value should be. That's the essence of*

machine learning—to approximate a gigantic table with a function. This is what learning is about.

There's more to it than that, obviously, but it's a good starting point when talking about AI because it makes it sound real, almost boring. Mechanical. So much of the conversation around AI is awash in mystical descriptions of its power and in reverence for its near-magic capabilities. Candela doesn't like that and tries to use more-prosaic terms. It's powerful, yes, but not magical. It has limitations. During presentations, he's fond of showing a slide with a wizard and a factory, telling audiences that Facebook thinks of AI like the latter, because "wizards don't scale."

And that's what Facebook has done with AI and machine learning: scaled it at a breakneck pace. A few years ago the company's machine learning group numbered just a few and needed days to run an experiment. Now, Candela says, several hundred employees run thousands of experiments a day. AI is woven so intricately into the platform that it would be impossible to separate the products—your feed, your chat, your kid's finsta—from the algorithms. Nearly everything users see and do is informed by AI and machine learning.

Understanding how and why Facebook has so fully embraced AI can help any organization that's ready to invest in an algorithmic future. It would be easy to assume that Facebook, with all its resources, would simply get the best talent and write the best algorithms—game over. But Candela took a different approach. Certainly the talent is strong, and the algorithms are good. Some of them are designed to "see" images or automatically filter them. Some understand conversations and can respond to them. Some translate between languages. Some try to predict what you'll like and buy.

But in some ways the algorithms are not his main focus. Instead, he's been busy creating an AI workshop in which anyone in the company can use AI to achieve a goal. Basically, Candela built an AI platform for the platform. Whether you're a deeply knowledgeable programmer or a complete newbie, you can take advantage of his wares.

Here's how he did it and what you can learn from it.

Soyuz

Candela, a veteran of Microsoft Research, arrived at Facebook in 2012 to work in the company's ads business. He and a handful of staffers inherited a ranking algorithm for better targeting users with ads.

Candela describes the machine learning code he inherited as "robust but not the latest." More than once he compares it to Soyuz, the 1960s Soviet spacecraft. Basic but reliable. Gets the job done even if it's not the newest, best thing. "It'll get you up there and down. But it's not the latest covnet [convolutional neural net] of the month."

You might assume, then, that the first thing Candela set out to do was to upgrade the algorithm. Get rid of Soyuz in favor of a space plane. It wasn't. "To get more value, I can do three things," he says. "I can improve the algorithm itself, make it more sophisticated. I can throw more and better data at the algorithm so that the existing code produces better results. And I can change the speed of experimentation to get more results faster.

"We focused on data and speed, not on a better algorithm."

Candela describes this decision as "dramatic" and "hard." Computer scientists, especially academic-minded ones, are rewarded for inventing new algorithms or improving existing ones. A better statistical model is the goal. Getting cited in a journal is validation. Wowing your peers gives you cred.

It requires a shift in thinking to get those engineers to focus on business impact before optimal statistical model. He thinks many companies are making the mistake of structuring their efforts around building the best algorithms, or hiring developers who claim to have the best algorithms, because that's how many AI developers think.

But for a company, a good algorithm that improves the business is more valuable than vanguard statistical models. In truth, Candela says, real algorithmic breakthroughs are few and far between—two or three a year at best. If his team focused its energies there, it would take lots of effort to make marginal gains.

He hammers these points home constantly: Figure out the impact on the business first. Know what you're solving for. Know what business challenge you need to address. "You might look for the shiniest algorithm or the people who are telling you they have the most advanced algorithm. And you really should be looking for people who are most obsessed with getting any algorithm to do a job. That's kind of a profound thing that I think is lost in a lot of the conversation. I had a conversation with our resident machine learning geek at our office, and we were just talking about different people doing AI. He said, 'Nobody really thinks their algorithms are very good or whatever.' It makes me think, maybe that's fine.

"I'm not saying don't work on the algorithm at all. I'm saying that focusing on giving it more data and better data, and then experimenting faster, makes a lot more sense."

So rather than defining success as building the best natural language processing algorithm, he defines it as deploying one that will help users find a restaurant when they ask their friends, "Where can I get a good bite around here?" Instead of being thrilled that some computer vision algorithm is nearing pixel-perfect object recognition, he gets excited if that AI is good enough to notice that you post a lot of pictures of the beach and can help you buy a swimsuit.

The strategy worked when he started at Facebook. Ad revenues rose. Candela's profile rose. It was suggested that AML become a centralized function for all of Facebook. Candela said no. Twice. "I was concerned about the 'If you build it, they will come' phenomenon." Just creating bits of artificial intelligence in the hope that people would see the value and adopt it wouldn't work.

But he did pick his spots. He collaborated with the feeds team while saying no to many other groups. Then he worked with the Messenger team. His team grew and took on more projects with other teams.

By 2015 Candela could see that his group would need to centralize, so he turned his attention to how he'd build such an operation. He was still worried about the "build it and they will come" phenomenon, so he focused less on how his team would be structured and more on how the group would connect to the rest of Facebook.

“You build a factory that makes amazing widgets, and you forget to design the loading docks into your factory?” He laughs. “Well, enjoy your widgets.”

Only then, about three years in, did Candela think about upgrading some of his algorithms. (Incidentally, even today, the emergency escape spacecraft attached to the International Space Station is a Soyuz.)

H2

Candela goes to a whiteboard to describe how he built his AI factory inside Facebook. The key, he says, was figuring out where on the product development path AI fits. He draws something like the graph in figure 2-1.

FIGURE 2-1

Where AI fits in at Facebook

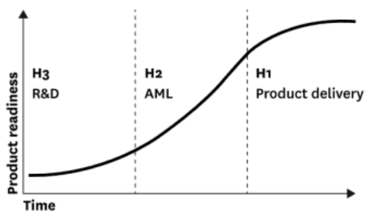

Source: Facebook

H3—Horizon 3, or three years out from product—is the realm of R&D and science. Often, data scientists who work on AI think of themselves as here, improving algorithms and looking for new ways to get machines to learn. Candela didn’t put his team here for the reasons already mentioned. It’s too far from impact on the business. H1, approaching product delivery, is where the product teams live—the feeds team, the Instagram team, the ads team. AI doesn’t go here either, because it would be difficult to retrofit products this deeply developed. It would be like building a car and then deciding that it should be self-driving after you started to put it together.

That leaves H2, between the science and the product, as the place AML lives at Facebook. AML is a conduit for transferring the science into the product. It does not do research for research’s sake, and it does not build and ship products. As the upward slope in the product’s readiness shows, it’s a dynamic space. Pointing to H2, Candela says, “This needs to feel uncomfortable all the time. The people you need to hire need to be okay with that, and they need to be incredibly selfless. Because if your work is successful, you spin it out. And you need to fail quite a bit. I’m comfortable with a 50% failure rate.”

If the team is failing less, Candela suspects its members are too risk averse, or they’re taking on challenges that are sliding them closer to H1’s product focus. “Maybe we do something like that and it works, but it’s still a failure, because the product teams should be taking that on, not us. If you own a piece of technology that the ads team should operate themselves to generate value, give it to them, and then increase your level of ambition in the machine learning space before something becomes product."

So Candela's team is neither earning the glory of inventing new statistical models nor putting products out into the world. It's a factory of specialists who translate others' science for others' products and fail half the time.

Push/Pull

All that being said, the lines between the three realms—H3, H2, and H1—still aren't crisp. In some cases Candela's team does look at the science of machine learning to solve specific problems. And sometimes it does help build the products.

That was especially true as AML got off the ground, because many people in the business hadn't yet been exposed to AI and what it could do for them. In one case AML built a translation algorithm. The team dipped into the research space to look at how existing translation algorithms worked and could be improved, because bad translations, which either don't make sense or create a misleading interpretation, are in some ways worse than no translation.

"Early on it was more push, more tenacity on our part," Candela says. "But it was gentle tenacity. We weren't going to throw something over the fence and tell the product team, 'This is great, use it.' " That meant that his team helped write some product code. Doing a little bit of the science and a little bit of the product in addition to its core function was meant to inspire the product team members to see what AML could do for them.

What the two teams built—a product that allowed community pages to instantly translate into several languages—worked. Other projects were similarly pushed out, and now the international team and other product groups at Facebook are pulling from AML, asking to use code in their products themselves.

"Look, it's nowhere near where I want it to be," Candela says. "I'd like to have all the product leaders in the company get together quarterly for AI reviews. That will certainly happen. But the conversation in the past two years has completely changed. Now if I walk from one end of this building to the other and I bump into, I don't know, the video team or the Messenger team, they'll stop me and say, 'Hey, we're excited to try this. We think we can build a product on this.' That didn't happen before."

AML's success, though, has created a new challenge for Candela. Now that everyone wants a piece of AML, the factory has to scale.

Layer Cake

Candela couldn't scale just by saying yes to every project and adding bodies to get the work done. So he organized in other ways. First he subdivided his team according to the type of AI its members would focus on (see figure 2-2).