MIT Sloan Management Review

SPECIAL COLLECTION

FROM THE LEADERSHIP ARCHIVE

Strategies for AI and Cognitive Tools

How can cognitive technologies aid business leaders? Learn what AI can — and can't — do.

CONTENTS

SPECIAL COLLECTION

Strategies for AI and Cognitive Tools

When People Don't Trust Algorithms

Berkeley J. Dietvorst, interviewed by Paul Michelman

University of Chicago professor Berkeley Dietvorst explains why we can't let go of human judgment — to our own detriment.

Berkeley Dietvorst, assistant professor of marketing at the University of Chicago Booth School of Business

Even when faced with evidence that an algorithm will deliver better results than human judgment, we consistently choose to follow our own minds.

Why?

MIT Sloan Management Review editor in chief Paul Michelman sat down with Berkeley Dietvorst, assistant professor of marketing at the University of Chicago Booth School of Business, to discuss a phenomenon Dietvorst has studied in great detail. (See "Related Research.") What follows is an edited and condensed version of their conversation.

MIT Sloan Management Review: What prompted you to investigate people's acceptance or lack thereof of algorithms in decision-making?

Dietvorst: When I was a Ph.D. student, some of my favorite papers were old works by [the late psychology scholar and behavioral decision research expert] Robyn Dawes showing that algorithms outperform human experts at making certain types of predictions. The algorithms that Dawes was using were very simple and oftentimes not even calibrated properly.

A lot of others followed up Dawes's work and showed that algorithms beat humans in many domains — in fact, in most of the domains that have been tested. There's all this empirical work showing algorithms are the best alternative, but people still aren't using them.

So we have this disconnect between what the evidence says people should do and what people are doing, and no one was researching why.

What's an example of these simple algorithms that were already proving to be superior?

Dietvorst: One of the areas was predicting student performance during an admission review. Dawes built a simple model: Take four or five variables — GPA, test scores, etc. — assign them equal weight, average them on a numerical scale, and use that result as your prediction of how students will rank against each other in actual performance. That model — which doesn't even try to determine the relative value of the different variables — significantly outperforms admissions experts in predicting a student's performance.

What were the experiments you conducted to try to get at the reasons we resist algorithms?

Dietvorst: We ran three sets of experiments.

For the first paper, we ran experiments where the participants' job was to complete a forecasting task, and they were incentivized to perform well. The better they performed, the more money they would earn in each experiment. There were two stages: first a practice round — for both humans and algorithms — and then a stage where participants were paid based on the quality of their performance.

In the practice round, we manipulated what forecasts participants were exposed to. Some made their own forecasts and saw those of the algorithm. Some made only their own forecasts. Some saw only the algorithm's results. Some saw neither. So each group had different information about how well each forecasting option had performed during the practice round.

For the second stage, participants could choose to forecast the results themselves or rely on the algorithm. The majority of participants who had *not* seen the algorithm's results from the first round chose to use it in the second round. However, those people who had seen the algorithm's results were significantly *less* likely to use it, even if it beat their own performance.

Once people had seen the algorithm perform and learned that it was imperfect, that it makes mistakes, they didn't want to use it. But there wasn't a similar effect for them. Once I made a forecast and learned that I was imperfect, I wasn't less likely to use my own forecast. We saw that effect only for the algorithm.

And for the second experiment?

Dietvorst: In the second paper, we tried to address the problem: How can we get people to use algorithms once they know that they're imperfect?

We began with the same basic question for participants: human or algorithm? In these experiments, however, there was an additional twist. Some participants were given the choice between using the algorithm as it existed or not at all. Other participants, if they chose to use the algorithm, could make some adjustments to it.

We found that people were substantially more willing to use algorithms when they could tweak them, even if just a tiny amount. People may be unwilling to use imperfect algorithms as they exist — even when the algorithm's performance has been demonstrated superior to their own — but if you give the person any freedom to apply their own judgment through small adjustments, they're much more willing.

So those are the key findings from the first two papers I wrote with my coauthors Joe Simmons and Cade Massey. Following on those, I have a solo paper where I'm investigating more about why people weren't willing to use algorithms once they learned that they're imperfect.

Most people in my experiment used human forecast by default, which positions the algorithm as an alternative. And the way they make the decision about whether or not to use the algorithm is by asking, "Will this algorithm meet my performance goal?" even if that goal is unrealistic for human forecasts, too. They don't choose the algorithm if it won't meet some lofty goal.

What they should more reasonably ask is, "Is this algorithm better than me?" — which it usually is. So people fail to ask the right question and end up holding the two options to different standards.

And to what do you attribute that?

Dietvorst: That's an interesting question. I'm not sure how this decision process came about or why people are making the decision this way. And I've found it's not actually unique to algorithms.

When choosing between two human forecasters, people do the same thing. If you assign them to have one forecaster as their default and you ask them how well would the other forecaster have to perform in order for you to switch, people say the other forecaster would have to meet my performance goals, just as with the algorithm.

It seems like people are naturally making what I would call the wrong comparison.

So it's kind of a switching cost?

Dietvorst: Not necessarily. The way I would think about a switching cost would be I'm used to using human judgment, so an algorithm has to perform X percent better or X points better than me, or a human, for me to switch to it, right?

But that's not really how it works. People are comparing the alternative to their performance goal, rather than comparing the two options. So, the higher the performance goal I give you, the better you need the algorithm to perform in order to switch to it, even though your own performance is staying constant.

So it doesn't seem like a switching cost, at least as we tend to think of the term.

What I find so interesting is that it's not limited to comparing human and algorithmic judgment; it's my current method versus a new method, irrelevant of whether that new method is human or technology.

Dietvorst: Yes, absolutely. That's exactly what I've been finding.

I think one of the questions that's going to come up is, "Well, what do I do about this? Is simple recognition of the bias enough to counter it?"

Dietvorst: If I can convince someone that the right question to ask is, "Does this algorithm outperform what you're currently using?" instead of, "Does this algorithm meet some lofty performance goal?" and that person buys in and says, "Yes, you're right, I should use algorithms that outperform what I'm currently doing," then, yes, that would work. I don't know how easy or hard it would be to get people to buy into that, though.

And in a larger organization, thousands of decisions are being made every day. Without this bias being known, there really isn't an obvious corrective measure, is there?

Dietvorst: The studies I've done suggest a couple restrictions that could reduce the bias.

People are deciding whether or not to use the algorithm by comparing it to the performance goal that they have. If you incentivize people to attempt to deliver performance much better than an algorithm has shown it's capable of, it's not so surprising that they ditch the algorithm to chase down that incentive with human judgment — even if it's unrealistic they will achieve it.

If you lower their performance goal, the algorithm will be compared more favorably and people may be more likely to use it.

So the problem exists in situations where the goal itself is unreasonable.

Dietvorst: Yes, if you have some forecasting goal that is very hard to achieve and an algorithm hasn't achieved it in the past, then you could see how it would make sense, in a certain way, for people not to use the algorithm. They're pretty sure it's not going to achieve the goal. So they use human judgment and end up performing even worse than the algorithm.

Presumably, we're in an age now where the quality of algorithms is increasing — perhaps dramatically. I'm wondering whether this phenomenon will make our biases more or less pronounced. On the one hand, you could see the quality of algorithms catching up to people's reference points. But the inverse of that is the reference point will continue to move at a speed as high if not higher than the ability of the algorithm.

Dietvorst: I agree: That could go either way. But I would like to push back a little bit on this idea that algorithms are really great. The literature shows that on average, when predicting human behavior, algorithms are about 10% to 15% better than humans. But humans are very bad at it. Algorithms are significantly better but nowhere near perfection. In many domains, I don't see any way that they're going to get close to perfection very soon.

There is a lot of uncertainty in the world that can't be resolved or reduced — that is unknowable. Like when you roll a die you don't know what number is going to come up until it happens. A lot of that type of aleatory uncertainty is determining outcomes in the real world. Algorithms can't explain that.

Suppose Google Maps is telling you the fastest route to a new place. It can't predict if there's going to be a giant accident right in front of you when you're halfway there. And so, as long as there's random error and there's aleatory uncertainty that factors into a lot of these outcomes — which it does to a larger extent than people recognize — algorithms aren't going to be perfect, and they aren't really even going to be close to perfect. They'll just be better than humans.

So what's next? Is this an ongoing field of study for you?

Dietvorst: Absolutely. There's a lot more to understand about how people think algorithms operate; what they think are the differences between algorithms and humans; and how that affects their use of algorithms. There's still really interesting research to be done.

About The Author

Paul Michelman is editor in chief of *MIT Sloan Management Review*. He tweets @pmichelman on Twitter.

Related Research

• B.J. Dietvorst, J.P. Simmons, and C. Massey, "Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err," Journal of Experimental Psychology: General 144, no. 1 (February 2015): 114-126.

• B.J. Dietvorst, J.P. Simmons, and C. Massey, "Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them," Management Science, published online Nov. 4, 2016.

• B.J. Dietvorst, "People Reject (Superior) Algorithms Because They Compare Them to Counter-Normative Reference Points," working paper, Dec. 6, 2016, http://papers.ssrn.com.

Reshaping Business With Artificial Intelligence

Executive Summary

Expectations for artificial intelligence (AI) are sky-high, but what are businesses actually doing now? The goal of this report is to present a realistic baseline that allows companies to compare their AI ambitions and efforts. Building on data rather than conjecture, the research is based on a global survey of more than 3,000 executives, managers, and analysts across industries and in-depth interviews with more than 30 technology experts and executives. (See “About the Research,” page 2.)

The gap between ambition and execution is large at most companies. Three-quarters of executives believe AI will enable their companies to move into new businesses. Almost 85% believe AI will allow their companies to obtain or sustain a competitive advantage. But only about one in five companies has incorporated AI in *some* offerings or processes. Only one in 20 companies has *extensively* incorporated AI in offerings or processes. Less than 39% of all companies have an AI strategy in place. The largest companies — those with at least 100,000 employees — are the most likely to have an AI strategy, but only half have one.

Our research reveals large gaps between today’s leaders — companies that already understand and have adopted AI — and laggards. One sizeable difference is their approach to data. AI algorithms are not natively “intelligent.” They learn inductively by analyzing data. While most leaders are investing in AI talent and have built robust information infrastructures, other companies lack analytics expertise and easy access to their data. Our research surfaced several misunderstandings about the resources needed to train AI. The leaders not only have a much deeper appreciation about what’s required to produce AI than laggards, they are also more likely to have senior leadership support and have developed a business case for AI initiatives.

ABOUT THE RESEARCH

To understand the challenges and opportunities associated with the use of artificial intelligence, *MIT Sloan Management Review*, in collaboration with The Boston Consulting Group, conducted its inaugural annual survey of more than 3,000 business executives, managers, and analysts from organizations around the world.

The survey, conducted in the spring of 2017, captured insights from individuals in 112 countries and 21 industries, from organizations of various sizes. More than two-thirds of the respondents were from outside of the United States. The sample was drawn from a number of sources, including *MIT Sloan Management Review* readers, and other interested parties.

In addition to our survey results, we interviewed business executives from a number of industries and academia to understand the practical issues facing organizations today. Their insights contributed to a richer understanding of the data.

For the purpose of our survey, we used the definition of artificial intelligence from the Oxford Dictionary: "AI is the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages." However, AI is evolving rapidly, as is the understanding and definition of the term.

AI has implications for management and organizational practices. While there are already multiple models for organizing for AI, organizational flexibility is a centerpiece of all of them. For large companies, the culture change required to implement AI will be daunting, according to several executives with whom we spoke.

Our survey respondents and interviewees are more sanguine than conventional wisdom on job loss. Most managers we surveyed do not expect that AI will lead to staff reductions at their organization within the next five years. Rather, they hope that AI will take over some of their more boring and unpleasant current tasks.

AI at Work

As Airbus started to ramp up production of its new A350 aircraft, the company faced a multibillion-euro challenge. In the words of Matthew Evans, vice president of digital transformation at the Toulouse, France-based company, "Our plan was to increase the production rate of that aircraft faster than ever before. To do that, we needed to address issues like responding quickly to disruptions in the factory. Because they will happen."

Airbus turned to artificial intelligence. It combined data from past production programs, continuing input from the A350 program, fuzzy matching, and a self-learning algorithm to identify patterns in production problems. In some areas, the system matches about 70% of the production disruptions to solutions used previously — in near real time. Evans describes how AI enables the entire Airbus production line to learn quickly and meet its business challenge:

> *What the system does is essentially look at a problem description, taking in all of the contextual information, and then it matches that with the description of the issue itself and gives the person on the floor an immediate recommendation. The problem might be new to them, but in fact, we've seen something very similar in the production line the weekend before, or on a different shift, or on a different section of the line. This has allowed us to shorten the amount of time it takes us to deal with disruptions by more than a third.*

AI empowered Airbus to solve a business problem more quickly and efficiently than prior approaches (such as root-cause analysis based on manual analysis of hundreds or thousands of cases).

Just as it is enabling speed and efficiency at Airbus, AI capabilities are leading directly to new, better processes and results at other pioneering organizations. Other large companies, such as BP, Infosys, Wells Fargo, and Ping An Insurance, are already solving important business problems with AI. Many others, however, have yet to get started.

High Expectations Amid Diverse Applications

Expectations for AI run high across industries, company sizes, and geography. While most executives have not yet seen substantial effects from AI, they clearly expect to in the next five years. Across all organizations, only 14% of respondents believe that AI is currently having a large effect (a lot or to a great extent) on their organization's offerings. However, 63% expect to see these effects within just five years.

Expectations for Change Across Industries and Within Organizations

Expectations for AI's effects on companies' offerings are consistently high across industry sectors. (See Figure 1.) Within the technology, media, and telecommunications industry, 72% of respondents expect large effects from AI in five years, a 52-percentage-point increase from the number of respondents currently reporting large effects. However, even in the public sector — the industry with the lowest overall expectations for AI's effects — 41% of respondents expect large effects from AI within five years, an increase of 30 percentage points from current levels. This bullishness is apparent regardless of the size or geography of the organization.

Within organizations, respondents report similarly high expectations for the large effects of AI on processes. While 15% of respondents reported a large effect of AI on current processes, over 59% expect to see large effects within five years. (See Figure 2.) Most organizations foresee sizable effects on information technology, operations and manufacturing, supply chain management, and customer-facing activities. (See Figure 3, page 4.) For example:

Information technology: Business process outsourcing providers serve as an example of the potential of AI. "IT services, where Infosys plays a big role, has seen tremendous growth in the last 20 or so years," says Infosys Ltd. CEO and managing director Vishal Sikka.[1] "Many jobs that moved to low labor-cost countries were the ones that were more mechanical: system administration, IT administration, business operations, verification. With AI techniques, we now have systems that can do more and more of those kinds of jobs. We are still in the early stages and portions of these activities can be automated, but we will get to the point in the next few years where the majority if not all of these jobs will be automated. However, just as AI technologies automate existing, well-defined activities, they

FIGURE 1: Expectations for AI's effect on businesses' offerings in five years are consistently high across industries.

Expectations for AI adoption across industries: impact on offerings

To what extent will the adoption of AI affect your organization's offerings today and five years from today?

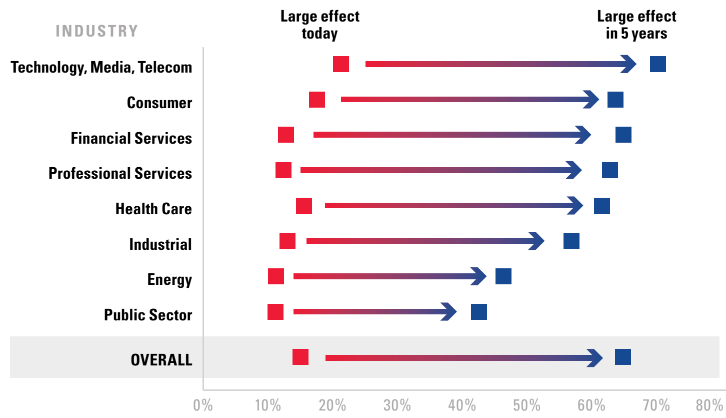

Percentage of respondents who expect a large ("a lot" or "great") effect on a five-point scale

FIGURE 2: As with offerings, organizations expect AI to have a great impact on processes within the next five years.

Expectations for AI adoption across industries: impact on processes

To what extent will the adoption of AI affect your organization's processes today and five years from today?

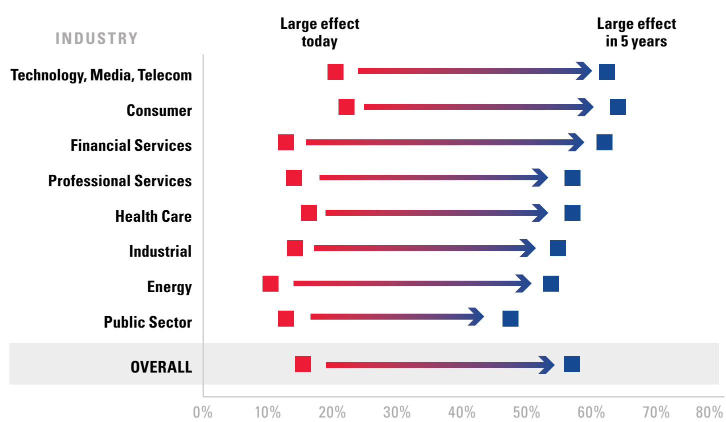

Percentage of respondents who expect a large ("a lot" or "great") effect on a five-point scale

FIGURE 3: Most organizations foresee a sizable effect on IT, operations, and customer-facing activities.

Most affected functional areas across industries

What areas within your organization do you anticipate AI will affect the most? Select three.

INDUSTRY	FUNCTIONAL AREA: RANK 1	RANK 2	RANK 3
Technology, Media, Telecom	Information Technology	Customer Service	Marketing
Consumer	Supply Chain Management	Sales	Marketing
Financial Services	Customer Service	Finance/ Accounting	Information Technology
Professional Services	Strategy	Information Technology	Marketing
Health Care	R&D	Operations/ Manufacturing	Information Technology
Industrial	Operations/ Manufacturing	Supply Chain Management	Information Technology
Energy	Operations/ Manufacturing	Information Technology	Supply Chain Management
Public Sector	Information Technology	Customer Service	Operations/ Manufacturing
OVERALL	Information Technology	Customer Service	Operations/ Manufacturing

Operations, including Manufacturing, SCM, and R&D · Information Technology · Customer-facing Functions · Corporate Center Functions

Functional areas that were not in the top three of any industry: communications, human resources, legal or compliance, procurement

FIGURE 4: More than 80% of organizations see AI as a strategic opportunity, while almost 40% also see strategic risks.

AI as strategic opportunity and risk

Do you perceive AI as a strategic opportunity or risk to your organization?

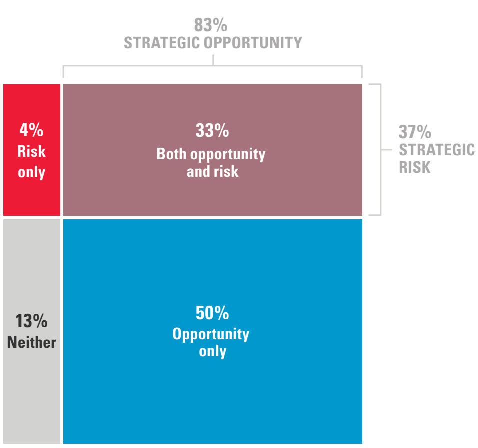

also create opportunities for new, breakthrough kinds of activities that did not exist."

Operations and manufacturing: Executives at industrial companies expect the largest effect in operations and manufacturing. BP plc, for example, augments human skills with AI in order to improve operations in the field. "We have something called the BP well advisor," says Ahmed Hashmi, global head of upstream technology, "that takes all of the data that's coming off of the drilling systems and creates advice for the engineers to adjust their drilling parameters to remain in the optimum zone and alerts them to potential operational upsets and risks down the road. We are also trying to automate root-cause failure analysis to where the system trains itself over time and it has the intelligence to rapidly assess and move from description to prediction to prescription."

Customer-facing activities: Ping An Insurance Co. of China Ltd., the second-largest insurer in China, with a market capitalization of $120 billion, is improving customer service across its insurance and financial services portfolio with AI. For example, it now offers an online loan in three minutes, thanks in part to a customer scoring tool that uses an internally developed AI-based face-recognition capability that is more accurate than humans. The tool has verified more than 300 million faces in various uses and now complements Ping An's cognitive AI capabilities including voice and imaging recognition.

Adoption as Opportunity and Risk

While expectations for AI run high, executives recognize its potential risks. Sikka is optimistic but cautions against hyping AI's imminent triumph: "If you look at the history of AI since its origin in 1956, it has been a story of peaks and valleys, and right now we are in a particularly exuberant time where everything looks like there is one magnificent peak in front of us." More than 80% of the executives surveyed are eyeing the peaks and view AI as a strategic opportunity. (See Figure 4.) In fact, the largest group of respondents, 50%, consider AI to be only an opportunity. Some see risks and the potential for increased competition from AI as well as benefits. Almost 40% of managers see AI as a strategic risk as well. A much smaller group (13%) does not view AI as either an opportunity or risk.

What is behind these high expectations and business interest in AI? There is no single explanation. (See Figure 5.) Most respondents believe that AI will benefit their organization, such as through new business or reduced costs; 84% believe AI will allow their organization to obtain or sustain a competitive advantage. Three in four managers think AI will allow them to move into new businesses.

Executives simultaneously recognize that their organization will not likely be the sole beneficiary of AI in their markets. Respondents expect that both new entrants and incumbents would similarly see the potential for benefits. Three-quarters of respondents foresee new competitors using AI to enter their markets while 69% expect current competitors to adopt AI in their businesses. Furthermore, they realize that suppliers and customers in their business ecosystem will increasingly expect them to use AI.

FIGURE 5: Organizations expect to create competitive advantage from AI — but also anticipate increased competition.

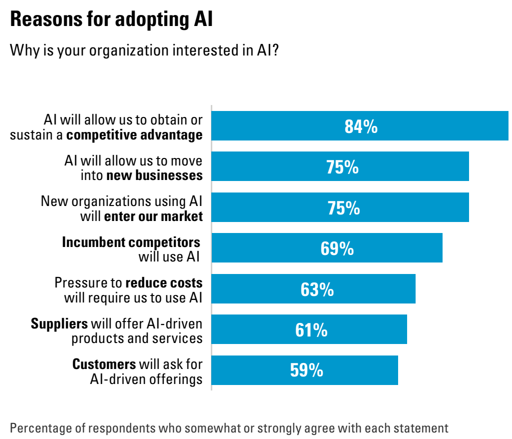

FIGURE 6: Only about a quarter of all organizations have adopted AI so far.

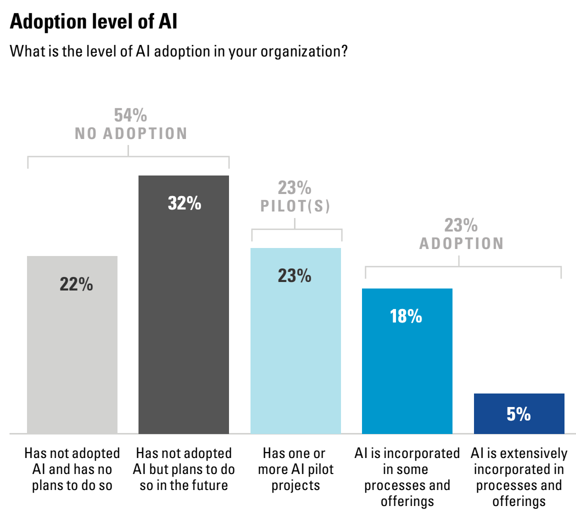

Disparity in Adoption and Understanding

Despite high expectations, business adoption of AI is at a very early stage: There is a disparity between expectation and action. Although four in five executives agree that AI is a strategic opportunity for their organization, only about one in five has incorporated AI in *some* offerings or processes. Only one in 20 has *extensively* incorporated AI in their offerings or processes. (See Figure 6.)

The differences in adoption can be striking, particularly within the same industry. For example, Ping An, which employs 110 data scientists, has launched about 30 CEO-sponsored AI initiatives that support, in part, its vision "that technology will be the key driver to deliver top-line growth for the company in the years to come," says the company's chief innovation officer, Jonathan Larsen. Yet in sharp contrast, elsewhere in the insurance industry, other large companies' AI initiatives are limited to "experimenting with chatbots," as a senior executive at a large Western insurer describes his company's AI program.

Organizations also report significant differences in their overall understanding of AI. For example, 16% of respondents strongly agreed that their organization understands the costs of developing AI-based products and services. And almost the same percentage (17%) strongly disagreed that their organization understands these costs. Similarly, while 19% of respondents strongly agreed that their organization understands the data required to train AI algorithms, 16% strongly disagreed that their organization has that understanding.

Combining survey responses to questions around AI understanding and adoption, four distinct organizational maturity clusters emerged: Pioneers, Investigators, Experimenters, and Passives.[2]

- **Pioneers (19%):** Organizations that both understand and have adopted AI. These organizations are on the leading edge of incorporating AI into both their organization's offerings and internal processes.
- **Investigators (32%):** Organizations that understand AI but are not deploying it beyond the pilot stage. Their investigation into what AI may offer emphasizes looking before leaping.
- **Experimenters (13%):** Organizations that are piloting or adopting AI without deep understanding. These organizations are learning by doing.
- **Passives (36%):** Organizations with no adoption or much understanding of AI.

If expectations and sense of opportunity are so high, what prevents organizations from adopting AI? Even in industries with extensive histories of integrating new technologies and managing data, barriers to AI adoption can be difficult to overcome. In financial services, for example, Simon Smiles, chief investment officer, ultra high net worth at UBS, puts it this way: "The potential for larger-scale financial institutions to leverage technology more actively, including artificial intelligence, within their business, and to harness their data to deliver a better client experience to the end user, is huge. The question there is whether these traditional institutions will actually grab the opportunity." Taking advantage of AI opportunities requires organizational commitment to get past the inevitable difficulties that accompany many AI initiatives.

These differences are less about technological limitations and much more about business. In aggregate, respondents ranked competing investment priorities and unclear business cases as more significant barriers to AI implementation than technology capabilities. Evans of Airbus makes the critical distinction: "Well, strictly speaking, we don't invest in AI. We don't invest in natural language processing. We don't invest in image analytics. We're always investing in a business problem." Airbus turned to AI because it solved a business problem; it made business sense to invest in AI instead of other approaches.

Smiles at UBS notes that organizations do not all face the same challenges. With respect to incumbents and fintech startups, he says: "There is a bifurcation between the groups that have the scale needed to develop incredibly valuable platforms and those unencumbered by legacy business models and systems to arguably have the better model going forward, but don't have the

FIGURE 7: While AI talent limits Pioneers, Passives don't yet discern a business case for AI.

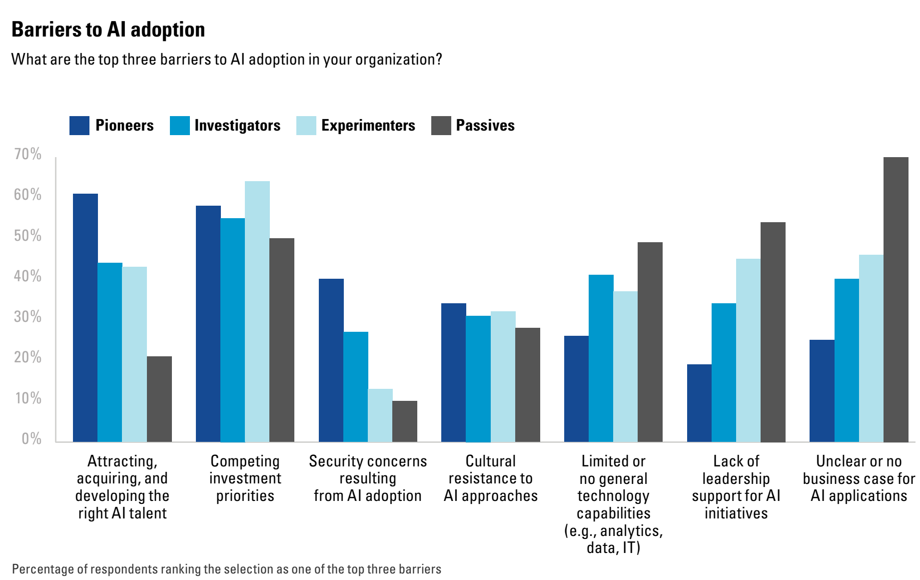

FIGURE 8: Organizations have different levels of understanding for AI-related technology and business context.

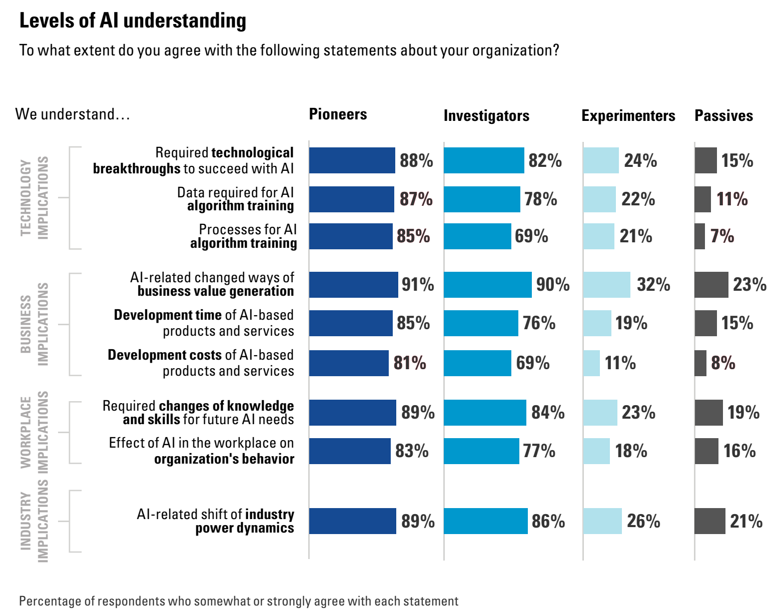

clients and accompanying data to capitalize fully on the opportunity." Differences like these lead to differences in rates of AI adoption.

Barriers to Adoption

The clusters of organizations demonstrate how barriers to AI differ and affect rates of adoption. (See Figure 7, page 6.) Pioneers have overcome issues related to understanding: three-quarters of these companies have identified business cases for AI. Senior executives are leading organizational AI initiatives. Their biggest hurdles are grappling with the practicalities of developing or acquiring the requisite AI talent and addressing competing priorities for AI investment. They are also much more likely to be attuned to the security concerns resulting from AI adoption. Passives, by contrast, have yet to come to grips with what AI can do for them. They have not identified solid business cases that meet their investment criteria. Leadership may not be on board. Technology is a hurdle. Many are not yet even aware of the difficulties in sourcing and deploying talent with AI expertise.

Our clustering also reveals nuanced differences in understanding among the clusters.

- **Business potential:** AI may change how organizations create business value. Pioneers (91%) and Investigators (90%) are much more likely to report that their organization recognizes how AI affects business value than Experimenters (32%) and Passives (23%). Evans at Airbus reports that "there was no question of value; it was trying to address an in-service issue on one of our aircraft."
- **Workplace implications:** Integrating the capabilities of humans and machines is a looming issue. AI stands to change much of the daily work environment. Pioneers and Investigators better appreciate that the presence of machines in the workplace will change behavior within the organization. Julie Shah, an associate professor of aeronautics at MIT, says, "What people don't talk about is the integration problem. Even if you can develop the system to do very focused, individual tasks for what people are doing today, as long as you can't entirely remove the person from the process, you have a new problem that arises — which is coordinating the work of, or even communication between, people and these AI systems. And that interaction problem is still a very difficult problem for us, and it's currently unsolved."
- **Industry context:** Organizations operate in regulatory and industry contexts; respondents from Experimenter and Passive organizations do not feel that their organization appreciates how AI may affect industry power dynamics.

The Need for Data, Training, and Algorithms

Perhaps the most telling difference among the four maturity clusters is in their understanding of the critical interdependence between data and AI algorithms. Compared to Passives, Pioneers are 12 times more likely to understand the process for training algorithms, 10 times more likely to understand the development costs of AI-based products and services, and 8 times more likely to understand the data that's needed for training AI algorithms. (See Figure 8.)

Most organizations represented in the survey have little understanding of the need to train AI algorithms on their data so they can recognize the sort of problem patterns that Airbus's AI application revealed. Less than half of respondents said their organization understands the processes required to train algorithms or the data needs of algorithms.

Generating business value from AI is directly connected to effective training of AI algorithms. Many current AI applications start with one or more "naked" algorithms that become intelligent only upon being trained (predominantly on company-specific data). Successful training depends on having well-developed information systems that can pull together relevant training data. Many Pioneers already have robust data and analytics infrastructures along with a broad understanding of what it takes to develop the data for training AI algorithms. Investigators and Experimenters, by contrast, struggle because they have little analytics expertise and keep their data largely in silos, where it is difficult to integrate. While over half of Pioneer organizations invest significantly in data and training, organizations from the other maturity clusters invest substantially less. For example, only one-quarter of Investigators have made significant investments in AI technology, the data required to train AI algorithms, and processes to support that training.

Misunderstandings About Data for AI

Our research revealed several data-related misconceptions. One misunderstanding is that sophisticated AI algorithms alone can provide valuable business solutions without sufficient data. Jacob Spoelstra, director of data science at Microsoft, observes:

> *I think there's still a pretty low maturity level in terms of people's understanding of what can be done through machine learning. A mistake we often see is that organizations don't have the historical data required for the algorithms to extract patterns for robust predictions. For example, they'll bring us in to build a predictive maintenance solution for them, and then we'll find out that there are very few, if any, recorded failures. They expect AI to predict when there* will *be a failure, even though there are no examples to learn from.*

No amount of algorithmic sophistication will overcome a lack of data. This is particularly relevant as organizations work to use AI to advance the frontiers of their performance.

Some forms of data scarcity go unrecognized: Positive results alone may not be enough for training AI. Citrine Informatics, a materials-aware AI platform helping to accelerate product development, uses data from both published experiments (which are biased toward successful experiments) and unpublished experiments (which include failed experiments) through a large network of relationships with research institutions. "Negative data is almost never published, but the corpus of negative results is critical for building an unbiased database," says Bryce Meredig, Citrine's cofounder and chief science officer. This approach has allowed Citrine to cut R&D time in half for specific applications. W.L. Gore & Associates, Inc., developer of Gore-Tex waterproof fabric, similarly records both successful and unsuccessful results in its push to innovate; knowing what does *not* work helps it to know where to explore next.[3]

Sophisticated algorithms can sometimes overcome limited data if its quality is high, but bad data is simply paralyzing. Data collection and preparation are typically the most time-consuming activities in developing an AI-based application, much more so than selecting and tuning a model. As Airbus' Evans says:

> *For every new project that we build, there's an investment in combining the data. There's an investment sometimes in bringing in new sources to the data platform. But we're also able to reuse all of the work that we've done in the past, because we can manage those business objects effectively. Each and every project becomes faster. The upfront costs, the nonrecurring costs, of development are lower. And we're able to, with each project, add more value and more business content to that data lake.*

Pioneer organizations understand the value of their data infrastructure to fuel AI algorithms.

Additionally, companies sometimes erroneously believe that they already have access to the data they need to exploit AI. Data ownership is a vexing problem for managers across all industries. Some data is proprietary, and the organizations that own it may have little incentive to make it available to others. Other data is fragmented across data sources, requiring consolidation and agreements with multiple other organizations in order to get more complete information for training AI systems. In other cases, ownership of important data may be uncertain or contested. Getting business value from AI may be theoretically possible but pragmatically difficult.

Even if the organization owns the data it needs, fragmentation across multiple systems can hinder the process of training AI algorithms. Agus Sudjianto, executive vice president of corporate model risk at Wells Fargo & Co., puts it this way:

> *A big component of what we do is dealing with unstructured data, such as text mining, and analyzing enormous quantities of transaction data, looking at patterns. We work on continuously improving our customer experience as well as decision-making in terms of customer prospecting, credit approval, and financial crime detection. In all these fields, there are significant opportunities to apply AI, but in a very large organization, data is often fragmented. This is the core issue of the large corporation — dealing with data strategically.*

Make Versus Buy

The need to train AI algorithms with appropriate data has wide-ranging implications for the traditional make-versus-buy decision that companies typically face with new technology investments. Generating value from AI is more complex than simply making or buying AI for a business process. Training AI algorithms involves a variety of skills, including understanding how to build algorithms, how to collect and integrate the relevant data for training purposes, and how to supervise the training of the algorithm. "We have to bring in people from different disciplines. And then, of course, we need the machine learning and AI people," says Sudjianto. "Somebody who can lead that type of team holistically is very important."

Pioneers rely heavily on developing internal skills through training or hiring. Organizations with less experience and understanding of AI put more emphasis on gaining access to outsourced AI-related skills, but this triggers some problems. (See Figure 9.)

The chief information officer of a large pharma company describes the products and services that AI vendors provide as "very young children." The AI tech suppliers "require us to give them tons of information to allow them to learn," he says, reflecting his frustration. "The amount of effort it takes to get the AI-based service to age 17, or 18, or 21 does not appear worth it yet. We believe the juice is not worth the squeeze."

To be sure, for some support functions, such as IT management and payroll support, companies might choose to outsource the entire process (and pass

FIGURE 9: Pioneers build AI-related skills through training and hiring, while Passives more heavily rely on external resources.

Approach to building AI-related skills

How does your organization build AI-related skills?

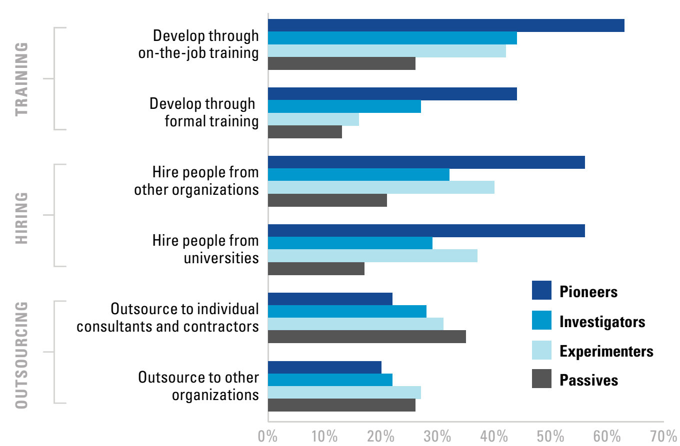

Percentage of respondents indicating they acquire skills in each way. Respondents could choose more than one option.

along all of their data). Even if companies expect to rely largely on external support, they need their own people who know how to structure the problem, handle the data, and stay aware of evolving opportunities. "Five years ago, we would have leveraged labor arbitrage arrangements with large outsourcers to access lower cost human labor to do that work," the pharma company CIO says. "What the vendors have done in the meantime is begin to automate those processes, oftentimes on our systems using our infrastructure, but using their technology. And I would not want it to be characterized as just rule-based. They actually have quite a bit more sophisticated heuristics to automate those things." But such an approach is clearly not suited for companies' distinctive offerings or core processes.

Eric Horvitz, director of Microsoft Research, argues that the tech sector is quickly catching up with the new model of offering technology tools to use with proprietary data, or "providing industry with toolsets, computation, and storage that helps to democratize AI." Many AI algorithms and tools are already in the public domain, including Google's TensorFlow, GitHub, and application programming interfaces from tech vendors. According to Horvitz:

> *Because this is a competitive space now in itself, the tools are getting easier to use and people that are there to help sell, market, and use these tools are becoming more efficacious in their abilities. That doesn't mean that people don't need to have their own in-house expertise and experts. While the tools and services are out there and that will make things easier, it is still going to be important for organizations to have their own experts in machine learning and AI more generally.*

FIGURE 10: Pioneers rate their companies higher across general management and leadership dimensions.

Link between AI and general organizational capabilities

To what extent do you agree with the following statements about your organization?

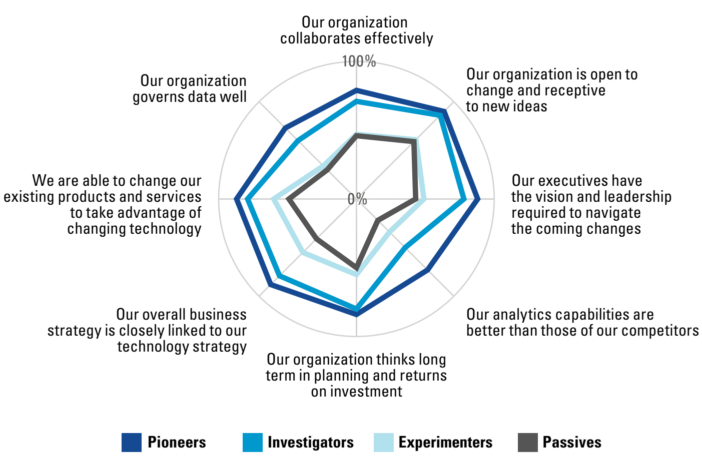

Percentage of respondents who somewhat or strongly agree with each statement

Privacy and Regulation

The data and the algorithms constituting AI cannot simply be accurate and high performing; they also need to satisfy privacy concerns and meet regulatory requirements. Yet only half the respondents in our survey agree that their industries have established data privacy rules.

Ensuring data privacy depends on having strong data governance practices. Pioneers (73%) are far more likely to have good data governance practices than the Experimenters (34%) and Passives (30%). (See Figure 10.) This large gap represents another barrier for companies that are behind in developing their AI capabilities.

The data issues can be pronounced in heavily regulated industries such as insurance, which is shifting from a historic model based on risk pooling toward an approach that incorporates elements that predict specific risks. But some attributes are off limits. For example, while sex and religion factors could be used to predict some risks, they are unacceptable to regulators in some applications and jurisdictions.

Regulators in other financial markets also have stringent transparency requirements. As Wells Fargo's Sudjianto says: "Models have to be very, very transparent and checked by the regulators all the time. When we choose not to use machine learning as the final model, it's because regulatory requirements oftentimes demand solutions be less 'black box' and something the regulator can see very clearly. But we use machine learning algorithms to assess the model's non-linear construction, variables and features entered, and as a benchmark for how well the traditional model performs."

As technology races ahead of consumer expectations and preferences, companies and the public sector tread an increasingly thin line between their AI initiatives, privacy protections, and customer service. Some financial services providers are using voice-recognition technology to identify customers on the phone to save time verifying identity. Customers welcome rather than balk at this experience, in part because they value the service and trust the company not to misuse the capability or the data that enables it. Likewise, a technology vendor offers an AI-based service to help call center operators recognize when customers are getting frustrated, using real-time sentiment analysis of voice data. Less welcome applications may be on the horizon, however. In a few years, any of the 170 million installed cameras in China or the 50 million cameras in the U.S. will be able to recognize faces. In fact, jaywalkers in Shanghai can already be fined (or shamed) based on such images.[4]

Beyond Technology: Management Challenges

AI requires more than data mastery. Companies also face many managerial challenges in introducing AI into their organizations.

Unsurprisingly, respondents at Pioneer organizations rate their companies higher in several general management and leadership areas: vision and leadership, openness and ability to change, long-term thinking, close alignment between business and technology strategy, and effective collaboration. As with other technology-driven transformations, these are essential general capabilities for high-performing companies.

However, there are also some specific challenges: Executives may still need to (1) learn more about AI; (2) deepen their perspective on how to organize their business around AI; and (3) develop a more expansive view of the competitive landscape in which their business operates.

Challenge 1: Develop an Intuitive Understanding of AI

The notion that executives and other managers need at least a basic understanding of AI is echoed by executives and academics. J.D. Elliott, director of enterprise data management at TIAA, a Fortune 100 financial services organization with nearly $1 trillion in assets under management, adds, "I don't think that every frontline manager needs to understand the difference between deep and shallow learning within a neural network. But I think a basic understanding that — through the use of analytics and by leveraging data — we do have techniques that will produce better and more accurate results and decisions than gut instinct is important." Avi Goldfarb, professor of marketing at the University of Toronto's Rotman School of Management, notes, "You worry that the unsophisticated manager might see one prediction work once and think that it's always good, or see one prediction that was bad and think it's always bad." Joi Ito, head of the MIT Media Lab, contends that "every manager has to develop an intuitive understanding of AI."[5]

To develop their understanding of digital, many executives have taken trips to Silicon Valley to experience digital natives, design-thinking approaches, fail-fast cultures, and more. While these are all core to building digital businesses, such trips are not particularly rewarding to learn about AI. For those who have already been exposed to the marvels of robots, self-driving vehicles, or poker-playing machines, there is little new to experience at AI companies. Instead, managers should take some time to learn the basics, possibly starting with simple online courses or online tools. They should understand how programs learn from data, maybe the most important facet of understanding how AI can benefit a particular business.

FIGURE 11: While the majority of organizations see developing an AI strategy as urgent, only half already have one.

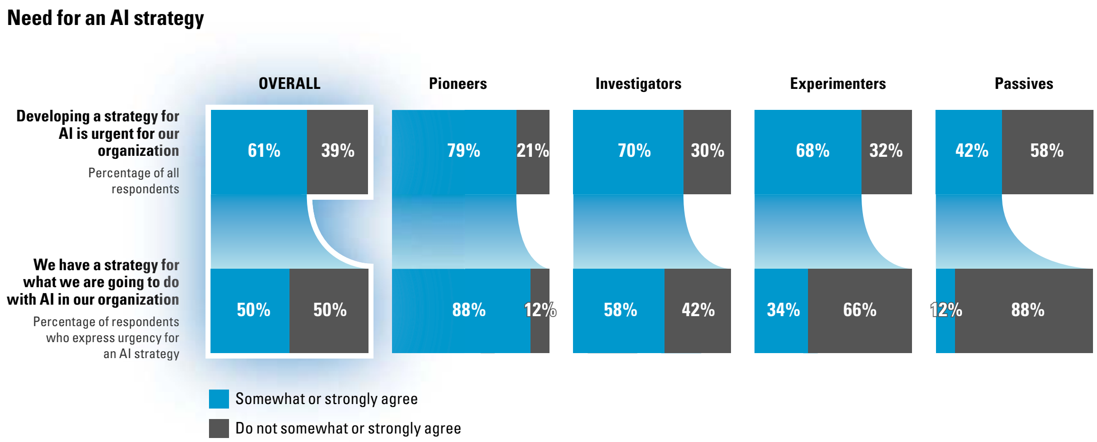

Challenge 2: Organize for AI

Adopting AI broadly across the enterprise will likely place a premium on soft skills and organizational flexibility that enable new forms of collaboration, including project teams composed of humans and machines.

Our survey finds companies exploring many approaches to developing AI capabilities. Pioneers are relatively evenly split among centralized, distributed, and hybrid organizational models. Investigators and Experimenters also pursue a mix of approaches, but almost 30% of both clusters have not yet set clear responsibility for AI in their organization. Some 70% of Passives also have not even started to lay out clear responsibilities for AI initiatives, perhaps (in part) because fewer than 50% of Passives see AI having a large effect on their processes and offerings in the next five years.

Ultimately, a hybrid model may make the most sense since many companies need AI resources both centrally and locally. TIAA, for example, has an analytics center of excellence and a number of decentralized groups. "The center of excellence is not intended to be the group that will provide all analytics for the entire organization. It provides expertise, guidance, and direction to other internal teams that are working to deploy AI and analytics," says TIAA's Elliott.

While companies in all four clusters rate cultural resistance to AI approaches relatively low on the list of barriers, only about half said that their company understands the required changes of knowledge and skills for future AI needs. Jessica Tan, group executive vice president, group chief operating officer, and chief information officer of Ping An, says the biggest challenges at her company have been getting units to work together; acknowledging the fact that "humans don't want to train algorithms"; establishing centralized and decentralized technology teams; and finding the right people. It's looking in particular for three types of people: technical people who have the means to try different ways of working, technical people who understand specific business domains, and people with consulting or project management skills who are able to network and bring them all together.

Challenge 3: Re-think the Competitive Landscape

More than 60% of respondents say that a strategy for AI is urgent for their organizations, but only half of those say their organizations have a strategy in place. (See Figure 11.) In terms of company size, the largest companies (those with more than 100,000 employees) are the most likely to have an AI strategy, but only half (56%) have one.

Amy Hoe, chief technology and operations officer of insurer FWD Group, says that she sees access to data as key for competitive advantage for her company. FWD aims to secure a wide range of data sources, including partnerships with other companies, such as telecommunications companies and ride-hailing services, its customer base, agencies, social media, the public domain, and external data analysis providers. As the volume of data doubles every few years, gaining privileged access to data is nonstop work.

What to Do Next

Is AI just an element of a company's overall digital transformation — or does AI require new approaches? On the one hand, AI presents many of the same issues and challenges as other digital technologies, and companies can build in many ways on their digital and analytics programs. However, AI also has distinctive features.

Ensure customer trust. AI capabilities are similar to many digital initiatives that depend on both customer data and customers' trust that the company will respect and safeguard their personal data. Ensuring that AI is trustworthy is different from other data-dependent digital initiatives, however, in several ways. First, managers may not be able to explain exactly how a customer's personal data is being used to produce a certain outcome from an AI product. The inner workings of some machine-learning programs are opaque. Second, a growing number of AI systems are able to mimic human agents, putting the onus on managers to clearly communicate to customers whether they are engaging with machines or human agents in a given setting. Third, some AI systems are able to assess emotions and discern quite personal details — at a distance. This capability creates new information management issues, including which employees have access to such information and under what circumstances.

Perform an AI health check. This has some similarities with digital health checks, from applications across processes to enabling infrastructure, technical skills, agile processes, and a fail-fast atmosphere. As with many digital initiatives, success with AI depends on access to data sources, be they existing internal or external data or investments in data infrastructure. Big companies may well have the data they need, but if it is fragmented and siloed, this significantly constrains strategy development and progress. Unlike other digital initiatives, an AI health check involves an assessment of the skills necessary to properly execute the *training* of AI, from first nurturing the system to become intelligent all the way to continuing to learn after deployment. This is both new and decisive — and a capability most companies need to build themselves. Off-the-shelf AI programs are likely to be limited in their capability and effect.

Brace for uncertainty. The adage "No idea is born good; you have to nurture it over time" applies to AI as well as to digital technologies — only more so. Companies often prioritize their initiatives by estimating the value of, and time required for, establishing a process or offering. But hard estimates are particularly difficult with AI. As a consequence, experimentation and learning with AI can take much longer than other digital initiatives, with a higher variability of success and failure. Managers need to brace themselves for more uncertainty, and this affects how effective they are at prioritizing projects and investments.

Adopt scenario-based planning. Like digital, AI has the potential to shift the ways in which businesses generate value — in multiple markets, processes, and functions. AI requires even more radical thinking, as it affects knowledge- and judgment-based professions, and the new entrants in markets could be machines. Thus, companies need to think even more expansively about their businesses, build cohesive future scenarios, and test the resilience of their directional plans against such scenarios.[6] This kind of scenario-based planning can also sharpen the ability to recognize events that could trigger large effects on their business.

Add a workforce focus. AI stands to create significant unease, since even the most knowledgeable expert has difficulties specifying how programs will play out, which functions or processes should be off limits, or where AI should stop. The threat to jobs and careers in their current form is real and can easily lead to employee anxiety and unrest. Establishing a clear focus and work plan for AI initiatives — where they will be pursued and how, including regular communication, education, and training — should be a component of an AI program. Attracting and developing people who combine both business and technical skills will be critical, as will the ability to deploy cross-functional teams, requiring both individual and organizational flexibility.

The Way Forward: Implications for the Future

The adoption of AI may have profound effects on the workplace, value creation, and competitive advantage. Beyond the near term, how should companies prepare for these changes?

FIGURE 12: Organizations suggest cautious optimism about AI's effect on the workforce in the next five years.

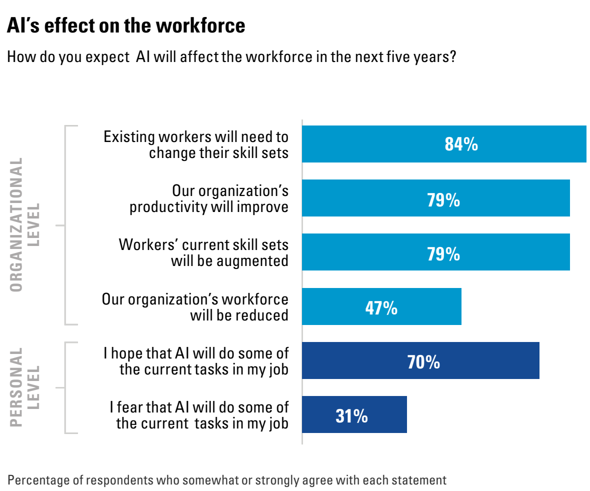

The Future of Work

As AI is increasingly applied to knowledge work, a significant shift will likely take place in the workplace, affecting many jobs in the Western middle class. Contrary to recent dire predictions about AI's effect on employment, our survey suggests cautious optimism. Most respondents, for example, do not expect that AI will lead to a reduction of jobs at their organization within the next five years. Nearly 70% also said they are not fearful that AI will automate their own jobs. By a similar margin, respondents hope that AI will take over some of their presumably boring and unpleasant current tasks. However, respondents overwhelmingly agree that AI will both require employees to learn new skills within the next five years and augment their existing skills. (See Figure 12.) Taken together, these portend adjustment, not annihilation. "Even with rapid advances," says Erik Brynjolfsson, Schussel Family Professor at the MIT Sloan School of Management, "AI won't be able to replace most jobs anytime soon. But in almost every industry, people using AI are starting to replace people who don't use AI, and that trend will only accelerate."[7]

Shifting Value Creation

Where will AI create, destroy, or shift economic value?

Consider the health care industry, one of the largest and most resilient sources of economic activity in the world. Health care spending makes up one-sixth of the U.S. economy, and on average, about one-tenth of the economies of Organisation for Economic Co-operation and Development (OECD) member nations. AI is already altering the health care value chain: Machines read diagnostic images, surgeons rely on robots, and an ever-increasing number of real-time medical devices contribute and communicate data to improve preventive and chronic care.

While AI may create value within an industry, it is far from clear exactly which organizations will see their fortunes rise and which will see decline. When IT vendors, medtech companies, radiologist networks, hospitals, specialized startups, and even insurance companies all strive to take advantage of AI to improve and lower the costs of diagnostics, the effects of AI will likely be uneven.

It's too early to tell which types of organizations may benefit from AI in health care. But if regulatory concerns can be worked out, the industry has numerous sources of detailed data. And as Marcus Winter, head of reinsurance development at Munich Re Group, remarks, "In today's world, with the proliferation of Big Data, there are precious few exclusive data sets. Most of the time, we can triangulate what we need to know via other sources." In other words, the combination of data and AI algorithms create the possibility of new and more effective workarounds. For example, when diagnostic imaging is unavailable, an evermore accurately analyzed sample of blood or other body fluids might help with diagnosis. As a result, shifts in value creation are difficult to predict.

Building Competitive Advantage

Managers expect significant improvement in performance of current processes or products from AI. Many companies are focused on addressing those. However, mere improvement does not create a sustainable competitive advantage — when everyone finds the same efficiencies, only the baseline shifts. For AI to become a prominent feature in future strategies, companies must figure out how humans and computers can build off each other's strengths to create competitive advantage. This is not easy: Companies need *privileged* access to data — which, as we've seen, many do not now have. They must learn how to make people and machines work effectively together — a capability relatively few Pioneers have at present. And they need to put in place flexible organizational structures, which means cultural changes for both company and employee.

Just about any company today needs a plan with respect to AI. Most do not have one, and those that have been slower to move have some catching up to do. Those that continue to fall behind may find the playing field tilted evermore steeply against them.

REFERENCES

1. At the time of his interview, Vishal Sikka was serving as CEO and managing director of Infosys. He has since resigned from that position to become executive vice chairman prior to the publication of this report.

2. We built a composite index of organizational understanding of AI based on the responses to nine survey questions related to AI understanding. This index, combined with the level of organizational adoption of AI, determined classification into the four clusters of organizations.

3. S. Ransbotham and D. Kiron, "Analytics as a Source of Business Innovation," Feb. 28, 2017, www.sloanreview.mit.edu.

4. Y. Wang, "China Is Quickly Embracing Facial Recognition Tech, for Better and Worse," July 11, 2017, www.forbes.com.

5. J. Ito and D. Kirkpatrick, "Davos — An Insight, an Idea with Joi Ito," World Economic Forum interview, Jan. 20, 2017, www.youtube.com.

6. R. Ramirez, S. Churchhouse, A. Palermo, and J. Hoffman, "Using Scenario Planning to Reshape Strategy," MIT Sloan Management Review 58 no. 4 (summer 2017): 31-37.

7. We did not ask respondents to look beyond five years, a horizon that is reasonably foreseeable. For some thoughts on what is possible in a 10- or 20-year time frame, see the Appendix.

8. World Economic Forum, "The Future of Jobs: Employment, Skills, and Workforce Strategy for the Fourth Industrial Revolution" (January 2016), 13.

9. Ibid., 8.

Appendix: Work in the Longer Term

Survey respondents and most interviewees both expect big changes from AI in the next five years. But the more dramatic effects of AI may occur within 10 to 20 years. What can we expect in that time frame?

Automation of Tasks ≠ Automation of Jobs. History shows that jobs evolve as tasks shift. BP's Ahmed Hashmi says the company's engineers used to spend a lot of time hunting for data to put together their reports, but "now that's all automated. We've got a data lake, which gives engineers ready access to all the data. We employ the same number of engineers, but they're improving the business rather than searching for data to get ready to improve the business." In other words, extrapolation from the automation of repetitive tasks to the automation of jobs in a high tech industry is risky business.

AI as Job Creator. Increased organizational reliance on AI will create new needs as it meets current needs. The job of an insurance underwriter, for example, tops many "most endangered species" lists. However, AI simultaneously expands the universe of insurable events. And, as James Platt, chief operating officer of Aon Risk Solutions, has said, "Many things that people would like to insure themselves against, such as brand and reputational risks or wider cybersecurity coverage, are 'uninsurable' today. There is simply no one offering an insurance option." As new methods of assessing risks become available, underwriters can start offering such new services. Missy Cummings, director of the Humans and Autonomy Laboratory at Duke University, puts it this way: "What we often don't think of are the jobs that are created as other new businesses come up around a technology."

If it's hard to imagine AI as doing anything other than *eliminating* jobs, step back and consider the scope of the problem. The 2016 World Economic Forum report, "The Future of Jobs," noted that "upcoming disruptions to the employment landscape are going to be a lot more complex and multifaceted than conveyed by a narrow focus only on automation"[8] — saying, in a nutshell, that digital technologies and AI are not the only forces transforming the nature of work. It has been clear for some time that technological change — not just AI — obliges employees to become lifelong learners and embrace career flexibility, but as the WEF report observes, it's far from alone: "technological, socioeconomic, geopolitical, and demographic developments and the interactions between them will generate new categories of jobs and occupations while partly or wholly displacing others. They will change the skill sets required in both old and new occupations in most industries and transform how and where people work."[9] Yet we have also seen digital technologies be used to address this problem. Accompanying the expansion of AI are many new learning options for humans: Augmented reality, new training tools, and digitally accessible forms of education (such as massive open online courses [MOOCs] and "nanodegrees") are proliferating.

Against a canvas of even broader social, demographic, environmental, and global political developments, predictions of aggregate employment levels based on AI alone are difficult; there are too many countervailing forces to discuss any one of them in isolation. But it is not unreasonable to imagine an opportunity for AI to cushion some of its own impacts, and perhaps the impacts of other factors, by helping to anticipate the coming changes in the job market and identify (and meet) workforce training needs as they arise.

Even So, Inertia Is Not an Option. Big global uncertainties should not deter corporations from acting today, when action is required. Infosys, for example, has trained more than 120,000 employees in design thinking. This new capability will enable its employees both to shape a world of new AI-based service offerings and automate historic business processing services.

Reprint 59181.

ACKNOWLEDGMENTS

We thank each of the following individuals, who were interviewed for this report:

Fabien Beckers, cofounder and chief executive officer, Arterys

Erik Brynjolfsson, director, MIT Initiative on the Digital Economy; Schussel Family Professor, MIT Sloan School; research associate, NBER

Eric Colson, chief algorithms officer, Stitch Fix

Missy Cummings, director, Humans and Autonomy Lab, Duke University

Steve Derbis, director, innovation development, Anthem

Steve Eglash, executive director, strategic research initiatives, computer science, Stanford University

J.D. Elliott, director, enterprise data management, TIAA

Eldad Elnekave, MD, chief medical officer, Zebra Medical Vision Ltd.

Matthew Evans, vice president, digital transformation, Airbus

Avi Goldfarb, professor of marketing, Rotman School of Management, University of Toronto

Mirsad Hadzikadic, director, data science and business analytics, UNC Charlotte

Ahmed Hashmi, global head of upstream technology, BP plc

Amy Hoe, chief technology and operations officer, FWD Group

Eric Horvitz, director, Microsoft Research, Microsoft

Michael Jordan, professor, computer science, University of California, Berkeley

Jonathan Larsen, chief innovation officer, Ping An Insurance Co. of China Ltd.

Bryce Meredig, cofounder and chief science officer, Citrine Informatics

James Platt, chief operating officer, Aon Risk Solutions

Julie Shah, associate professor, aeronautics, MIT

Vishal Sikka, chief executive officer and managing director, Infosys Ltd.

Simon Smiles, chief investment officer, ultra high net worth, UBS

Beth Smith, general manager, IBM Watson Platform

Alfred Spector, chief technology officer, Two Sigma

Jacob Spoelstra, director, data science, Microsoft

Agus Sudjianto, executive vice president, corporate model risk, Wells Fargo & Co.

Jessica Tan, group executive vice president, group chief operating officer, and chief information officer, Ping An Insurance Co. of China Ltd.

Marcus Winter, head of reinsurance development, Munich Re Group

AUTHORS

SAM RANSBOTHAM is an associate professor in the information systems department at the Carroll School of Business at Boston College, as well as guest editor for *MIT Sloan Management Review*'s Artificial Intelligence Big Ideas Initiative. He can be reached on Twitter @ransbotham.

DAVID KIRON is the executive editor of *MIT Sloan Management Review*, which brings ideas from the world of thinkers to the executives and managers who use them.

PHILIPP GERBERT is a senior partner and managing director at The Boston Consulting Group's Munich, Germany, office. He is BCG's global topic leader for digital strategy and a BCG Henderson Institute Fellow for the Impact of Artificial Intelligence on Business. He can be reached at gerbert.philipp@bcg.com.

MARTIN REEVES is a senior partner and managing director at The Boston Consulting Group and the director of the BCG Henderson Institute, which brings ideas and inspiration to forward-looking leaders.

CONTRIBUTORS

Sebastian Steinhäuser, principal and member of the AI core team, BCG
Patrick Ruwolt, consultant and member of the AI core team, BCG
Allison Ryder, senior project editor, *MIT Sloan Management Review*

The research and analysis for this report was conducted under the direction of the authors as part of an *MIT Sloan Management Review* research initiative in collaboration with and sponsored by The Boston Consulting Group.

To cite this report, please use:
S. Ransbotham, D. Kiron, P. Gerbert, and M. Reeves, "Reshaping Business With Artificial Intelligence," *MIT Sloan Management Review* and The Boston Consulting Group, September 2017.

[ARTIFICIAL INTELLIGENCE]

The Fundamental Flaw in AI Implementation

Many executives are enthusiastic about the business potential of machine learning applications. But business leaders often overlook a key issue: To fully unlock the benefits of artificial intelligence, you'll need to upgrade your people's skills — and build an empowered, AI-savvy workforce.

BY JEANNE ROSS

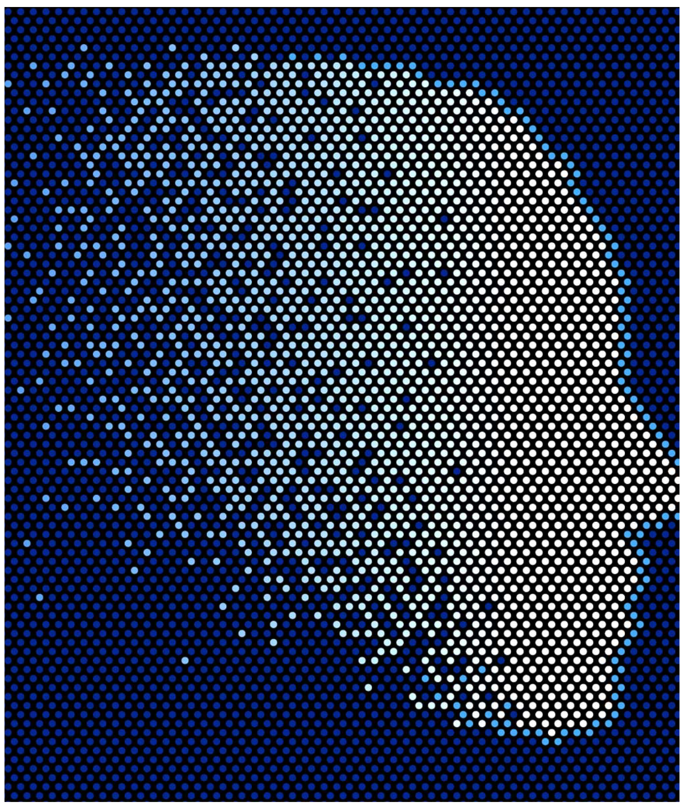

There is no question that artificial intelligence (AI) is presenting huge opportunities for companies to automate business processes. However, as you prepare to insert machine learning applications into your business processes, I recommend that you not fantasize about how a computer that can win at Go or poker can surely help you win in the marketplace. A better reference point will be your experience implementing your enterprise resource planning (ERP) system or another enterprise system. Yes, effective ERP implementations enhanced the competitiveness of many companies, but many other companies found the experience more of a nightmare. The promised opportunity never came to fruition.

Why am I raining on the AI parade? Because, as with enterprise systems, AI inserted into businesses drives value by improving processes through automation. But eventually, the outputs of most automated processes require people to do something. As most managers have learned the hard way, computers can process data just fine, but that processing isn't worth much if people are feeding them bad data in the first place or don't know what to do with information or analysis once it's provided.

With my fellow researchers, Cynthia Beath, Monideepa Tarafdar, and Kate Moloney, I've been studying how companies insert value-adding AI algorithms into their processes. As other researchers and managers have also observed, we are finding that most machine learning applications augment, rather than replace, human efforts. In doing so, they demand changes in what people are doing. And in the case of AI — even more than was true with ERP systems — those changes eliminate many nonspecialized tasks and create skilled tasks that require good judgment and domain expertise.

For example, fraud detection applications may reduce the time that people spend looking for anomalies but increase requirements for deciding what to do about those anomalies. An AI application might allow financial analysts to spend less time extracting data on financial performance, but it adds value only if someone spends more time considering the implications of that performance. With the help of AI applications, customer service staff can spend fewer hours resolving routine problems, but they are more likely to improve operations if at least some of that saved time is reallocated to better understanding the problems customers are experiencing with the company's most recent offerings.

Many leaders think that they will generate value from AI by recruiting more data scientists. Of course, there's a shortage of data scientists — and some of them are more attracted to the challenge of building an application that wins at poker than solving a business need. Others will be inspired to find a cure for cancer or to mitigate global warming. So financial services and insurance companies attempting to uncover fraud and technology companies hoping to improve customer satisfaction will be fighting over the remaining talent.

But recruiting data scientists is not your biggest challenge. Data scientists can develop useful algorithms, but domain experts are needed to help train the machine to recognize important patterns and understand new data. Domain experts include top analysts, contract managers, salespeople, recruiters, and other specialists who are not only experts at their jobs but are also acutely aware of how they deliver excellence. That may involve just a few key people for a given application, but they'd better be good. And we still haven't gotten to the really hard part!

Ultimately, you need people who can use probabilistic output to guide actions that make your company more effective. Probabilistic outputs are no problem when, say, an application such as Salesforce.com Inc.'s AI tool, Einstein, indicates that one lead has a 95% chance of converting into a sale while another has a 60% chance. The salesperson knows what to do with that information. But what's the next step when a recruiter learns from an AI application that a job candidate has a 50% likelihood of being a good fit for a particular opening?

When a machine learning application is helping a lawyer identify potentially relevant legal precedents, helping a vendor management team ensure compliance with a contract, or helping a banker decide whether a customer qualifies for a loan, the machine is taking over mundane tasks. Machines can surely learn to develop spreadsheets and search large databases for relevant information. But to generate competitive advantage from machine learning applications, you'll need to upgrade your employees' skills. You'll also need to redesign their accountabilities, so that they are empowered and motivated to deploy machines when they believe that doing so will enhance outcomes. In short, you will need to build an entire workforce of intelligence-consuming, action-oriented superstars.

There are, of course, examples of AI algorithms fully automating a process rather than augmenting human efforts. Google DeepMind might automatically adjust temperature settings in a data center. Similarly, IBM Watson can trigger automated alerts to insurance customers in an area likely to be hit by a hailstorm. But these are exceptions. More often, machine learning applications are helping people accomplish something. Like people, machines have natural limits, which tend to leave parts of the tasks — the parts that don't fit the algorithms well — to people. When a machine detects fraud or predicts customer or employee churn with 90% accuracy, people must address the other 10% — and that will be the toughest 10%. The machine will assuredly take care of the easy cases.

Addressing the toughest instances is particularly challenging because AI algorithms can produce indecipherable results. When a machine learning algorithm decides who gets a loan and who doesn't, forget about trying to advise a client about how to qualify. Machine intelligence is not a substitute for human intelligence, because, as organizations, we need to be able to understand why we're doing what we're doing.

> Machine intelligence is not a substitute for human intelligence, because, as organizations, we need to be able to understand why we're doing what we're doing.

None of the issues associated with using AI to augment your employees' skills are insurmountable. Great companies are already empowering their people with better information produced by smart machines. Those machines sift through far more data, and do it much faster, than people can. They also discover complex relationships that can be exposed only with massive amounts of data and a large pool of contrasting outcomes. Companies are succeeding with AI by partnering smart machines with smart people who are learning how to take advantage of what those machines can do. In short, AI implementation success depends on your ability to hire and develop problem-solvers, equip them with data (and potentially AI), and then empower them to actually solve problems. Note that addressing skill requirements this way may well require major changes to your existing hiring and development practices.

Companies that view smart machines purely as a cost-cutting opportunity are likely to insert them in all the wrong places and in all the wrong ways. These companies will automate existing processes rather than imagine new ones. They will cut jobs rather than upgrade roles. These are the companies who will find that implementing AI is little more than a reprise of the ERP nightmare.

***Jeanne Ross** is a principal research scientist at the MIT Center for Information Systems Research (@mit_cisr) in Cambridge, Massachusetts. Comment on this article at http://sloanreview.mit.edu/x/59212.*

What You Need to Know Before Starting a Platform Business

Richard Schmalensee and David S. Evans, interviewed by Martha E. Mangelsdorf

There's probably never been a better time for platform businesses. But, warn two experienced economists, that doesn't make them easy to launch successfully.

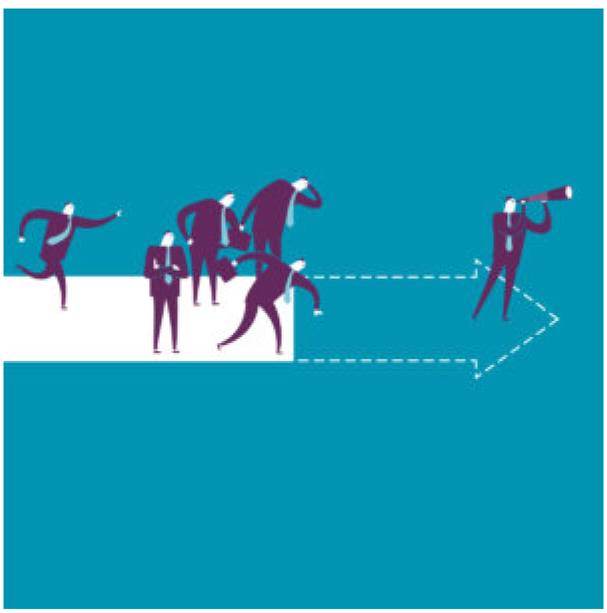

There's a great deal of enthusiasm about platform strategies these days. Entrepreneurs pitch their startups as the next Uber, the next Facebook, or the next Airbnb, while executives in established companies are retooling their strategies around platforms to drive growth and compete digitally.

But creating a successful platform business is not easy, as economists Richard Schmalensee and David S. Evans will tell you. Schmalensee and Evans have studied multisided platforms — in other words, businesses that create value by connecting two or more sets of participants via a physical or virtual platform — for more than a decade. Their most recent book for executives, *Matchmakers: The New Economics of Multisided Platforms*, was published by Harvard Business Review Press in 2016.

Schmalensee, a former dean and professor emeritus of management and of economics at the MIT Sloan School of Management, and Evans, chairman of the consulting firm Global Economics Group and a lecturer at the University of Chicago Law School, don't sugarcoat the difficulties associated with pursuing a successful platform strategy. As they wrote in *Matchmakers*, "A multisided platform is one of the toughest businesses to get right. The entrepreneur has to solve a tough puzzle and use counterintuitive strategies to make a go of it."

MIT Sloan Management Review editorial director Martha E. Mangelsdorf spoke with Schmalensee and Evans to learn more about just what it takes to get a platform strategy right. What follows is an edited and condensed version of their conversation.

In your most recent book, you point out that what you call the "matchmaking" business model — businesses that create value by connecting different groups — have been around for a long time, and have included everything from a trading area in ancient Athens to credit card companies. Is it safe to say that digital technologies make these kinds of businesses more common — and more important?

Schmalensee: Yes, the "matchmaker" model of making it easier for two sides to connect and create value — and then capturing some of that value — goes back at least to the ancient Greeks. But it's the new digital technologies that have really turbocharged this business model. It's easier to connect now.

Evans: Think about Uber. As a business, it could not have existed in 2008.

Schmalensee: Right. Without a lot of smartphone penetration, Uber's business model wouldn't work.

Evans: What has created massive opportunities for these platform businesses is the fact that, in effect, the internet is pervasively available in physical space. Between mobile phones that people carry around, and then the internet of things connected to wireless networks, you basically have plugged the physical world into the internet. More and more points globally are connected. That provides an immense opportunity to develop platform businesses that connect different groups.

So we're in an era that, thanks to digital technologies, is highly conducive to platform businesses — and to their existing on a greater scale than in the past. Yet one of the points you make in *Matchmakers* is that the multisided platform business model — the matchmaker model — is not a very easy one to succeed with. Why is that?

Schmalensee: Our editor said what he liked about our book *Matchmakers* was that we were both curmudgeons. Our perspective is that there's too much hype about this business model. It's really easy to look at companies like Uber and say, "Look at these wildly successful businesses!" but forget all the people who tried to do similar things and failed.

In particular, there are two big difficulties associated with starting a matchmaker business. First, you have to solve a real problem for both groups you seek to connect. You have to be enabling something that wasn't there before — and it's got to be a sufficiently big thing that the outcome is good for Group A, and it's good for Group B, and there's a little of the value created left over for you.

The second difficulty is that you have to get started somehow. It's easy to imagine matchmaker business models that work when you have 10 million people on one side of the platform and 10 million businesses or people on the other side. That's easy, but getting there — figuring out how to get off the ground and attract enough of both groups to make it interesting to both — is not. You're giving one side of the platform access to the other side for the purpose of creating value — and if there's nobody on the other side, you don't have a product. We refer to it as the "chicken-and-egg" problem.

Let's talk about the first issue: solving a big enough problem. Can you give an example that illustrates that?

Schmalensee: Sure. The mobile payment system M-Pesa solved a big problem in Kenya: transferring money from family members who work in the city to the family members left behind in the village.

Could you make that system work in the U.S. with feature phones and convenience stores? No, because we don't have a big problem transferring money in this country.

Ask yourself: What's the economic friction that your platform will reduce? What problem are you solving? It doesn't have to be something that people have been trying to do for years. YouTube is an interesting example. They thought video sharing might be neat, but nobody had done it before. It turned out they were right. So the platform could involve something completely novel that technology enables, but that's a harder gamble than figuring out how to solve a big problem that already exists.

Let's talk about the second problem you mentioned: the "chicken-and-egg" issue.

Schmalensee: Once you get lots of people participating on a platform, it can be really attractive to new participants, because there are a lot of people already there who they want to deal with. But if you can't get momentum, you don't ever have an attractive product, because you're not selling access to an attractive group of potential partners.

Think of the restaurant-booking service OpenTable. If there aren't enough restaurants on the system, then it's not of much interest to consumers, and if there aren't enough consumers, it's not of much interest to restaurants to sign up. So in its early days, OpenTable had to solve a hard problem: How do you get both restaurants and consumers?

OpenTable figured out the trick was to offer software to restaurants that would enable them to manage tables and reservations and also get online. And OpenTable's first idea was to start working with restaurant chains across the country to get a lot of restaurants onto its service. But the fact that there are lots of restaurants in Omaha on OpenTable is not of much interest to me if I live in Boston. I want there to be a dense set of restaurants where I live.

So OpenTable backed off its nationwide strategy for a while, focused on San Francisco and then Chicago, and signed up a bunch of leading restaurants in those markets. Then the company could go out and get consumers to sign up because it had a product. But the trick was to go first for restaurants and then to consumers and to give restaurants a reason independent of the existence of consumers to sign up with OpenTable — in this case, table-management software. That's one way around the chicken-and-egg problem: You sell them something that is of value to one side of the platform and gives them access.

What are some of the other ways to address the chicken-and-egg problem?

Evans: Sometimes, like OpenTable, it's a ratcheting or zig-zag effect where you get more restaurants, then you get more diners, and you keep moving the numbers up that way. There are other cases where you really need to go out and sign up — in effect, anchor tenants for one side of the platform — in order to really increase the number of customers on the other side. Take game consoles, where you need to sign the game developers up a year or two in advance in order to make sure games are available for your platform when you release it. In the game console case, it's essential that you figure out a way to sign up developers, so that console customers have something to interact with.

Another solution to the chicken-and-egg problem is to make your own chickens. We all think of the iPhone as a multisided platform with an app store, developers, and users. But when Apple introduced the iPhone in June 2007, there was no app store. It wasn't open to developers at all.

Steve Jobs and Apple did two things. One is that they created a bunch of applications themselves, so out of the box, the iPhone was a great experience for the consumer. And then they went out and did specific deals with a small handful of developers. The iPhone in June 2007 wasn't what we would call a matchmaker or a platform at all. It wasn't until about a year later that Apple opened it up and made it that way.

These different strategies — gradually ratcheting up the numbers on both sides of the platform through a zig-zag effect, signing up key participants such as anchor tenants, and creating your own chickens — aren't mutually exclusive.

And do most successful platforms end up employing a range of them?

Schmalensee: I think they try all of the above, unless there is some obvious best approach.

Evans: In some cases, it is possible to go out and just develop one side of the platform, and then you can be pretty confident you can then sign up the other side. That's how ad-supported platforms work: You come up with a bunch of great content. You get a bunch of people to come look at the content. Once you have enough eyeballs looking at the content, then you go sell ads.

You mention eyeballs. One of the points you make in your book *Matchmakers* is that, in the early days of e-commerce, one of the fallacies was that if you just got enough eyeballs, then you would necessarily win. Can you talk a little bit about the problems with that strategy?

Evans: Sometimes that strategy will work, but sometimes it won't. The fallacy is in believing it's just a numbers game. It's not. It's a game of coming up with the right number of participants on each side that actually want to do business with or interact with the participants on the other side.

If you go back to the OpenTable example, if you had a choice between having 50,000 restaurants and 100,000 consumers spread all across the U.S. or having 1,000 restaurants and 10,000 consumers in San Francisco, you'd probably rather have the latter. In the former case, you're not going to have enough restaurants and consumers that actually want to do business with each other because the groups are just spread over too large an area. That was the flaw in the OpenTable strategy initially: They tried to go national when what they needed to do was to go local first.

Another point you make in *Matchmakers* is that pricing — how you allocate the pie — is more complicated in multisided platform businesses than if you're simply dealing with one set of customers. Can you talk a little bit about that?

Schmalensee: One of the first things that strikes you when you study multisided platform businesses is that very often, one of the sides rides for free. For example, OpenTable gives me bonus points for using its service. I don't pay a thing, and it's a valuable service. And you can go on down a list of matchmaker businesses and see that very often, one set of participants goes free or doesn't pay enough to cover its share of costs.

Sometimes it's obvious which side should be subsidized. Sometimes it's not — and it's not always the case that one side goes free. The balance may shift over time as the market changes. So getting the pricing balance right initially doesn't mean you get it right forever. You have to keep thinking about it. If you're going to do some completely novel platform business, figuring out how to price it for both sides may be hard; you may have to experiment, and you may have to change pricing strategies.

Businesses that aren't platforms don't have the same kind of balancing act. You may have one group of customers over here, and another group of customers over there, and you price to each group, but in a normal business, those customers don't interact. In a platform business, the interaction is key — and pricing to make sure you get that interaction right requires a balancing act.

Evans: If you see a successful platform business, the pricing solution looks obvious — after the fact. But it isn't. If you're actually starting one of these businesses, it's not obvious at all. These entrepreneurs typically struggle, and it's the ones who latch on to the right model who become successful. But it's very, very complicated to figure out the right things to do to launch these businesses — both in terms of how to price them and how to get them off the ground. With these businesses, there are more dimensions of things you need to get right — and more uncertainty.

What are some of the other common mistakes that people make when launching platform businesses?

Schmalensee: Well, one that we talk about is getting the governance structure wrong. Again, it depends on the platform, but a lot of these businesses involve interactions that need to be governed in ways that are platform-specific.

Take the social networks. The actress Lindsay Lohan at one point got thrown off of Facebook for posting under an assumed name — a violation of Facebook's rules. OpenTable will throw you off the platform if you don't show up for dinner reservations, and if you want to sell through the Amazon Marketplace, there are a whole set of rules you have to comply with. Securities exchanges have rules governing market makers, liquidity providers, traders, and liquidity takers.

If you get the rules wrong, even if your pricing is right, you can get bad behavior that drives people away from the platform.

Interesting. On a different topic, Peter Weill and Stephanie Woerner of the MIT Center for Information Systems Research have written in *MIT Sloan Management Review* about how some established companies are creating digital ecosystems that may include their competitors. It seems to me that that's a different kind of platform problem — when you have an established business and you're thinking about making it into a platform that connects not just you to your customers, but also maybe involves other companies that may be competitors of yours. Is that a question you've looked at, or does that just get into even more complicated economic questions?

Schmalensee: I think it is more complicated, but some of the same principles hold. The question is, how are you adding value, and how are you capturing value? If you're going to become a platform, what's the proposition you offer to folks on one side, and what's the proposition on the other side? What problem are you solving?

It may be a problem for all participants or a problem for the end customer, but you have to be solving a significant problem. You have to give value to all sides. You've got to solve the chicken-and-egg question. You've got to get the incentives right. So, the basic economics are the same. The context may just be much more complicated.

Who's done that successfully? Maybe Amazon with Marketplace?

Evans: It's worth noting that Amazon tried to take on eBay early on with a Marketplace-like initiative and failed miserably. Initially, Amazon was not able to crack the chicken-and-egg problem, and it took them several tries before they were able to get Marketplace off the ground. It was initially an abysmal failure; they had trouble attracting merchants.

Dick described us earlier as curmudgeons. That's probably an unfair term, but we're realists; we're thoughtful skeptics. So one thing I'd like to point out is that lots of companies shouldn't become platforms. There are probably lots of areas of commerce where it probably doesn't make sense to have a platform. Companies need to be mindful of not getting caught up in the hype — the idea that since Uber and so many other companies are doing one of these platform things that therefore that's the right solution for you.

It may be that you're in a part of commerce that is just not suitable for a platform business. Or it may be that you're better off joining someone else's platform, like a retailer joining Amazon Marketplace, rather than trying to create your own platform. Because you may not be successful at creating your own platform, whereas you may be very successful as a participant on someone else's platform.

Good point. You've described the increased potential for platform businesses that exists because of the ability to connect groups, wherever they are, more seamlessly through internet-enabled devices. At the same time, you're making it very clear that a platform business model is absolutely not a slam dunk. When

considering a plan for a platform business, how can an executive gauge whether it represents a good opportunity?

Evans: One way to figure out whether there's an opportunity is to see if there is an existing platform business that's operated by traditional methods where, using modern technology, it is possible to provide the intermediation in a much more efficient and scalable way. It's identifying platforms that already exist, but where there's a lot of economic friction and inefficiency. And a great example of that is Uber.

The taxi industry was a two-sided platform. Taxi companies were basically intermediaries that connected taxi drivers and riders through dispatch systems. It was just a very inefficient platform, and Uber came along and figured out a way to operate much more efficiently within a city, but then also in a way that could be scaled geographically.

Schmalensee: I would emphasize the economic friction part of that point more than the technology, because a common mistake made by platform entrepreneurs is to think, "Oh, gee whiz, I could do a platform and use the internet and it would be wonderful." But their platform turns out not be enough of an improvement over existing solutions. You have to ask: Does the proposed platform actually solve a real problem — as opposed to being just "gee-whizzy"?

And the second question to ask is: How are you going to get this thing off the ground? What's the launch strategy? What's the strategy for solving the chicken-and-egg problem? How do you get to this wonderful place where you're making a lot of money because you've got the half the world on one side of your platform and half the world on the other side? What's the plan?

Not every successful platform business had a perfect plan at the start, but they sure had to think about it — and often readjust on the fly when the first idea didn't work.

Evans: Dick and I, I think, are resistant to making things sound too simple. But we have a lot of experience now at studying these platforms at an early stage and tracking them and finding out which ones succeed and which ones don't. And while there's all sorts of complicated stuff we can talk about, in my experience, the mistake that these businesses make time after time after time is not focusing enough on the point that Dick raised: Making sure the friction they eliminate — the problem they solve for users on both sides of the platform — is really significant. Many would-be platforms fail because they're trying to use technology to solve something that might be a friction, but just really isn't a big enough problem.

Schmalensee: That said, the combination of the matchmaker platform business model with all of today's digital technologies means people are likely to come up with some new matchmaking businesses where, in five years, you're going to say, "Oh, what a great idea! I wish I'd thought of that." We don't know what those new platform business ideas will be yet, but there will be some.

About The Author

Martha E. Mangelsdorf is the editorial director of *MIT Sloan Management Review*. She tweets at @memangelsdorf.

[DATA & ANALYTICS]

The Subtle Sources of Sampling Bias Hiding in Your Data

Plummeting data acquisition costs have contributed to a surge in business analytics. But more data doesn't inherently remove sampling bias — and in some cases, it could make it worse.

BY SAM RANSBOTHAM

When a group of Boston College students started an analytics project using data about UFO sightings, they thought they'd learn something about visits from spaceships and alien creatures — such as how weather and movie releases influence sightings. *The Economist* had done something similar, finding that most UFO reports are made during what it called "drinking hours" (5 to 11 p.m.), when people could be "nursing their fourth beer" — a possible connection that the publication dubbed "close encounters of the slurred kind."

Instead, the students learned about sampling biases.

UFO sighting reports in the United States have increased substantially since the National UFO Reporting Center, a private organization based in Davenport, Washington, started tracking them in 1974. But this might not mean that we are getting more visitors from outer space.

When the reporting center first opened, communicating a sighting required making a telephone call to file a report. Once the internet became publicly available and people could make reports using an online form, the number of sightings began to rise. This easier and cheaper collection system provided more data about sightings. But the increase in the availability of data fundamentally changed the sample set — and any change in data affects the conclusions we can draw from that data.

Looking beyond the world of UFOs, lower costs of data collection provide value in many ways: We have much more data to work with and learn from than ever before. But managers must be careful to understand *how* the data was generated and how that might influence its value. The sources of bias in data sets can be far subtler than the ones that could be at play in the UFO data. What's more, the task of interpreting data is falling on the shoulders of more people in organizations. What biases should managers be on the lookout for as they work to gain insight from increasing amounts of available data? And how can managers help their employees become better at spotting such biases?

Here are four practices that can help:

Understand the history behind your data. New data can be fundamentally different from older data in ways that managers must understand. In the infamous 1948 *Chicago Daily Tribune* "Dewey Defeats Truman" example, when the newspaper prematurely printed an incorrect headline about the winner of the U.S. presidential election, the paper had based its conclusion on telephone polling, rather than door-to-door polling, which had been used in the past. That turned out to be a mistake because the underlying demographics of telephone owners at that time differed substantially from those of the national electorate as a whole.

In a similar vein, businesses today must be savvy about how they interpret the rich, low-cost data from online forums, keeping in mind that what they learn about customers using social media may show different trends than data from prior sources, such as phone or written surveys. Social media is an amazing new source of detailed data about consumer activity; it gives businesses access to unprecedented amounts of information about individuals. But not every customer uses social media. And not everyone is honest online: People intentionally shape their images on social media. Customers may not necessarily be behaving differently, but they could be *responding differently* based on the medium or the visibility of their response.

Organizations seeking to become more data driven must make it easy for managers to understand data's lineage — its origins, the systems involved in collecting it, and intermediate processing steps. Managers must understand the need to ask: What do we know about where the data we are using comes from? What might have changed the data since its origin?

Acknowledge that more data may not mean better data. With increased amounts of data, statistically significant results are easier to find but distract from the larger problem of sampling errors — that the sample may be internally consistent but not reflect the desired population. Data volume can give false comfort: Managers may fall into the trap of thinking they have "better data" when they just have heavier weighting of the prior data. In a worse case, the increased data volume hides sampling errors deeper in a haystack of information.

It is here that the "big" of Big Data can fail us. Managers risk falling into thinking that an enormous number of data points just can't be wrong, that the sample is too big to fail. What's more, ingesting and processing new information may require substantial processing to transform unstructured data into structured data. Or the new data may require that business processes change to incorporate real-time feeds. Each of these tasks takes resources, time, and effort. Failure comes from generating more data that doesn't add to what the organization already knows.

Before embarking on projects to acquire more data, managers need to assess what new information the additional data will bring — or do a pilot to find out. They should ask: What insights are we looking for? And if we don't know, how can we find out without large investments?

Recognize that old data sources were imperfect, too. While it may be tempting to benchmark new data sources against old sources, old sources had sampling biases, too. In the UFO example, most reported UFO sightings were in Washington state. But in the days before widespread use of the internet, reports of sightings had to be made by phone, and anyone outside of the local calling area in Washington state would have had to make a long-distance call to reach the National UFO Reporting Center — a factor that may have made cost-conscious people think twice before picking up the phone.

Managers risk falling into thinking that an enormous number of data points just can't be wrong, that the sample is too big to fail.

So how do you know which data source is best? It is likely that managers understand their older data sources better than their new data sources, since managerial experience grows over time. Just like it took time to get to know the old data, it will take time to learn the new. There are trade-offs in the biases of one source versus another, and understanding those biases will take experience. Using both old and new data sources can sometimes provide more insight than either alone, since each can help illuminate the sampling bias in the other.

To acknowledge this challenge, managers should ask: How is our existing data limited? Can a new data source get around that limitation?

Remember that intuition remains important. With an increasing volume of data from an ever-expanding variety of sources, blending information with intuition and understanding potential sampling bias may become more vital than ever.

Understanding sampling bias is an inherently human task. It requires knowing what is *not* in the data — and the data itself cannot tell you what it's missing. Despite the rise of artificial intelligence and machine learning, human domain expertise is still needed to look at the big picture and understand which portion of that picture a particular data source shows, as well as what it doesn't show. Human domain expertise is still needed to understand ongoing trends that likely began before the new data source existed. And human domain expertise is still needed to know what might happen in the business context that affects the sampling.

As data becomes increasingly ubiquitous, expertise in applying data to specific business problems will become a key resource — and a potential source of differentiation for both individuals and the organizations that employ them. Managers should ask: What do I know about my business that the data does not?

Combining Two Types of Knowledge

Unfortunately for organizations, the burden of understanding sampling bias falls squarely on whichever staff members try to use and interpret analytical results. Just because an organization can produce analytics, that does not mean everyone (or

The Subtle Sources of Sampling Bias Hiding in Your Data
(Continued from page 21)

anyone!) in the organization will be able to apply those results well. Getting value from data requires deep knowledge of the data-generating processes. But in many organizations, this information is not attached to the analytical results. That presents a crucial, growing challenge: conveying information about data lineage through inherently distributed processes.

As use of analytics becomes more pervasive throughout organizations, an increasing number of people will need to become savvy consumers of analytical results. To succeed as data consumers, managers must combine two vastly different types of knowledge. First, they must know the details behind the data-generating processes to understand what the data can and cannot say. Second, they must have a broad understanding and general knowledge of their business.

The burden of understanding sampling bias cannot be handled centrally. It is a challenge that will be felt throughout organizations, requiring many individuals to learn more to compensate for it. To address this challenge, managers need to ask: How can I develop employees who combine a broad knowledge of the business with an ability to interpret detailed data accurately?

***Sam Ransbotham** is an associate professor of information systems at the Carroll School of Management at Boston College in Chestnut Hill, Massachusetts. Comment on this article at http://sloanreview.mit.edu/x/59126, or contact the author at smrfeedback@mit.edu.*

ACKNOWLEDGMENTS

The author thanks Boston College students Matthew Frederick, Puneet Nayyar, Amanda Valdes, Alexa Villalobos, and Valeria Yanes for insights from their analytics project about UFO sightings.

Reprint 59126.

[TECHNOLOGY]

What to Expect From Artificial Intelligence

To understand how advances in artificial intelligence are likely to change the workplace — and the work of managers — you need to know where AI delivers the most value.

BY AJAY AGRAWAL, JOSHUA S. GANS, AND AVI GOLDFARB

Major technology companies such as Apple, Google, and Amazon are prominently featuring artificial intelligence (AI) in their product launches and acquiring AI-based startups. The flurry of interest in AI is triggering a variety of reactions — everything from excitement about how the capabilities will augment human labor to trepidation about how they will eliminate jobs. In our view, the best way to assess the impact of radical technological change is to ask a fundamental question: How does the technology reduce costs? Only then can we really figure out how things might change.

To appreciate how useful this framing can be, let's review the rise of computer technology through the same lens. Moore's law, the long-held view that the number of transistors on an integrated circuit doubles approximately every two years, dominated information technology until just a few years ago. What did the semiconductor revolution reduce the cost of? In a word: *arithmetic.*

This answer may seem surprising since computers have become so widespread. We use them to communicate, play games and music, design buildings, and even produce art. But deep down, computers are souped-up calculators. That they appear to do more is testament to the power of arithmetic. The link between computers and arithmetic was clear in the early days, when computers were primarily used for censuses and various military applications. Before semiconductors, "computers" were humans who were employed to do arithmetic problems. Digital computers made arithmetic inexpensive, which eventually resulted in thousands of new applications for everything from data storage to word processing to photography.

AI presents a similar opportunity: to make something that has been comparatively expensive abundant and cheap. The task that AI makes abundant and inexpensive is *prediction* — in other words, the ability to take information you have and generate information you didn't previously have. In this article, we will demonstrate how improvement in AI is linked to advances in prediction. We will explore how AI can help us solve problems that were not previously prediction oriented, how the value of some human skills will rise while others fall, and what the implications are for managers. Our speculations are informed by how technological change has affected the cost of previous tasks, allowing us to anticipate how AI may affect what workers and managers do.

Machine Learning and Prediction

The recent advances in AI come under the rubric of what's known as "machine learning," which involves programming computers to learn from example data or past experience. Consider, for example, what it takes to identify objects in a basket of groceries. If we could describe how an apple looks, then we could program a computer to recognize apples based on their color and shape. However, there are other objects that are apple-like in both color and shape. We could continue encoding our knowledge of

What to Expect From Artificial Intelligence (Continued from page 23)

apples in finer detail, but in the real world, the amount of complexity increases exponentially.

Environments with a high degree of complexity are where machine learning is most useful. In one type of training, the machine is shown a set of pictures with names attached. It is then shown millions of pictures that each contain named objects, only some of which are apples. As a result, the machine notices correlations — for example, apples are often red. Using correlates such as color, shape, texture, and, most important, context, the machine references information from past images of apples to predict whether an unidentified new image it's viewing contains an apple.

When we talk about prediction, we usually mean anticipating what will happen in the future. For example, machine learning can be used to predict whether a bank customer will default on a loan. But we can also apply it to the present by, for instance, using symptoms to develop a medical diagnosis (in effect, *predicting* the presence of a disease). Using data this way is not new. The mathematical ideas behind machine learning are decades old. Many of the algorithms are even older. So what has changed?

Recent advances in computational speed, data storage, data retrieval, sensors, and algorithms have combined to dramatically reduce the cost of machine learning-based predictions. And the results can be seen in the speed of image recognition and language translation, which have gone from clunky to nearly perfect. All this progress has resulted in a dramatic decrease in the cost of prediction.

The Value of Prediction

So how will improvements in machine learning impact what happens in the workplace? How will they affect one's ability to complete a task, which might be anything from driving a car to establishing the price for a new product? Once actions are taken, they generate outcomes. (See "The Anatomy of a Task.") But actions don't occur in a vacuum. Rather, they are shaped by underlying conditions. For example, a driver's decision to turn right or left is influenced by predictions about what other drivers will do and what the best course of action may be in light of those predictions.

Seen in this way, it's useful to distinguish between the cost versus the value of prediction. As we have noted, advances in AI have reduced the cost of prediction. Just as important is what has happened to the *value*. Prediction becomes more valuable when data is more widely available and more accessible. The computer revolution has enabled huge increases in both the amount and variety of data. As data availability expands, prediction becomes increasingly possible in a wider variety of tasks.

Autonomous driving offers a good example. The technology required for a car to accelerate, turn, and brake without a driver is decades old. Engineers initially focused on directing the car with what computer scientists call "if then else" algorithms, such as "If an object is in front of the car, then brake." But progress was slow; there were too many possibilities to codify everything. Then, in the early 2000s, several research groups pursued a useful insight: A vehicle could drive autonomously by predicting what a human driver would do in response to a set of inputs (inputs that, in the vehicle's case, could come from camera images, information using the laser-based measurement method known as LIDAR, and mapping data). The recognition that autonomous driving was a prediction problem solvable with machine learning meant that autonomous vehicles could start to become a reality in the marketplace years earlier than had been anticipated.

Who Judges?

Judgment is the ability to make considered decisions — to understand the impact different actions will have on outcomes in light of predictions. Tasks where the desired outcome can be easily described and there is limited need for human judgment are generally easier to automate. For other tasks, describing a precise outcome can be more difficult, particularly when the desired outcome resides in the minds of humans and cannot be translated into something a machine can understand.

This is not to say that our understanding of human judgment won't improve and therefore become subject to automation. New modes of machine learning may find ways to examine the relationships between actions and outcomes, and then use the information to improve predictions. We saw an example of this in 2016, when AlphaGo, Google's DeepMind artificial intelligence program, succeeded in beating one of the top players in the world in the game of Go. AlphaGo honed its capability by analyzing thousands of human-to-human Go games and playing against

THE ANATOMY OF A TASK

Actions are shaped by the underlying conditions and the resolution of uncertainty to lead to outcomes. Drivers, for example, need to observe the immediate environment and make adjustments to minimize the risk of accidents and avoid bottlenecks. In doing so, they use judgment in combination with prediction.

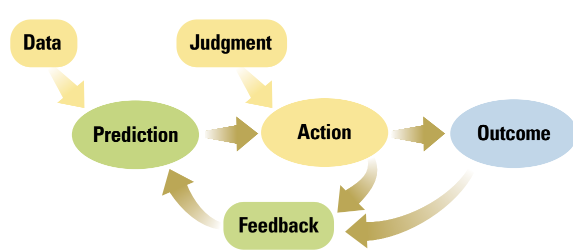

itself millions of times. It then incorporated the feedback on actions and outcomes to develop more accurate predictions and new strategies.

In cases where whole decisions can be clearly defined with an algorithm, we expect to see computers replace humans.

Examples of machine learning are beginning to appear more in everyday contexts. For instance, x.ai, a New York City-based artificial intelligence startup, provides a virtual personal assistant for scheduling appointments over email and managing calendars. To train the virtual assistants, development team members had the virtual assistants study the email interactions between people as they schedule meetings with one another so that the technology could learn to anticipate the human responses and see the choices humans make. Although this training didn't produce a formal catalog of outcomes, the idea is to help virtual assistants mimic human judgment so that over time, the feedback can turn some aspects of judgment into prediction problems.

By breaking down tasks into their constituent components, we can begin to see ways AI will affect the workplace. Although the discussion about AI is usually framed in terms of machines versus humans, we see it more in terms of understanding the level of judgment necessary to pursue actions. In cases where whole decisions can be clearly defined with an algorithm (for example, image recognition and autonomous driving), we expect to see computers replace humans. This will take longer in areas where judgment can't be easily described, although as the cost of prediction falls, the number of such tasks will decline.

Employing Prediction Machines

Major advances in prediction may facilitate the automation of entire tasks. This will require machines that can both generate reliable predictions and rely on those predictions to determine what to do next. For example, for many business-related language translation tasks, the role of human judgment will become limited as prediction-driven translation improves (though judgment might still be important when translations are part of complex negotiations). However, in other contexts, cheaper and more readily available predictions could lead to increased value for human-led judgment tasks. For instance, Google's Inbox by Gmail can process incoming email messages and propose several short responses, but it asks the human judge which automated response is the most appropriate. Selecting from a list of choices is faster than typing a reply, enabling the user to respond to more emails in less time.

Medicine is an area where AI will likely play a larger role — but humans will still have an important role, too. Although artificial intelligence can improve diagnosis, which is likely to lead to more effective treatments and better patient care, treatment and care will still rely on human judgment. Different patients have different needs, which humans are better able to respond to than machines. There are many situations where machines may never be able to weigh the relevant pros and cons of doing things one way as opposed to another way in a manner that is acceptable to humans.

The Managerial Challenge

As artificial intelligence technology improves, predictions by machines will increasingly take the place of predictions by humans. As this scenario unfolds, what roles will humans play that emphasize their strengths in judgment while recognizing their limitations in prediction? Preparing for such a future requires considering three interrelated insights:

1. Prediction is not the same as automation. Prediction is an input in automation, but successful automation requires a variety of other activities. Tasks are made up of data, prediction, judgment, and action. Machine learning involves just one component: prediction. Automation also requires that machines be involved with data collection, judgment, and action. For example, autonomous driving involves vision (data); scenarios — given sensory inputs, what action would a human take? (prediction); assessment of consequences (judgment); and acceleration, braking, and steering (action). Medical care can involve information about the patient's condition (data); diagnostics (prediction); treatment choices (judgment); bedside manner (judgment and action); and physical intervention (action). Prediction is the aspect of automation in which the technology is currently improving especially rapidly, although sensor technology (data) and robotics (action) are also advancing quickly.

2. The most valuable workforce skills involve judgment. In many work activities, prediction has been the bottleneck to automation. In some activities, such as driving, this bottleneck has meant that human workers have remained involved in prediction tasks. Going forward, such human involvement is all but certain to diminish. Instead, employers will want workers to augment the value of

What to Expect From Artificial Intelligence (Continued from page 25)

The future's most valuable skills will be those that are complementary to prediction — in other words, those related to judgment.

prediction; the future's most valuable skills will be those that are complementary to prediction — in other words, those related to judgment. Consider this analogy: The demand for golf balls rises if the price of golf clubs falls, because golf clubs and golf balls are what economists call complementary goods. Similarly, judgment skills are complementary to prediction and will be in greater demand if the price of prediction falls due to advances in AI. For now, we can only speculate on which aspects of judgment are apt to be most vital: ethical judgment, emotional intelligence, artistic taste, the ability to define tasks well, or some other forms of judgment. However, it seems likely that organizations will have continuing demand for people who can make responsible decisions (requiring ethical judgment), engage customers and employees (requiring emotional intelligence), and identify new opportunities (requiring creativity).

Judgment-related skills will be increasingly valuable in a variety of settings. For example, if prediction leads to cheaper, faster, and earlier diagnosis of diseases, nursing skills related to physical intervention and emotional comfort may become more important. Similarly, as AI becomes better at predicting shopping behavior, skilled human greeters at stores may help differentiate retailers from their competitors. And as AI becomes better at anticipating crimes, private security guards who combine ethical judgment with policing skills may be in greater demand. The part of a task that requires human judgment may change over time, as AI learns to predict human judgment in a particular context. Thus, the judgment aspect of a task will be a moving target, requiring humans to adapt to new situations where judgment is required.

3. Managing may require a new set of talents and expertise. Today, many managerial tasks are predictive. Hiring and promoting decisions, for example, are predicated on prediction: Which job applicant is most likely to succeed in a particular role? As machines become better at prediction, managers' prediction skills will become less valuable while their judgment skills (which include the ability to mentor, provide emotional support, and maintain ethical standards) become more valuable.

Increasingly, the role of the manager will involve determining how best to apply artificial intelligence, by asking questions such as: What are the opportunities for prediction? What should be predicted? How should the AI agent learn in order to improve predictions over time? Managing in this context will require judgment both in identifying and applying the most useful predictions, and in being able to weigh the relative costs of different types of errors. Sometimes there will be well-acknowledged objectives (for example, identifying people from their faces). Other times, the objective will be less clear and therefore require judgment to specify the desired outcome. In such cases, managers' judgment will become a particularly valuable complement to prediction technology.

Looking Ahead

At the dawn of the 21st century, the most common prediction problems in business were classic statistical questions such as inventory management and demand forecasting. However, over the last 10 years, researchers have learned that image recognition, driving, and translation may also be framed as prediction problems. As the range of tasks that are recast as prediction problems continues to grow, we believe the scope of new applications will be extraordinary. The key challenges for executives will be (1) shifting the training of employees from a focus on prediction-related skills to judgment-related ones; (2) assessing the rate and direction of the adoption of AI technologies in order to properly time the shifting of workforce training (not too early, yet not too late); and (3) developing management processes that build the most effective teams of judgment-focused humans and prediction-focused AI agents.

Ajay Agrawal *is the Peter Munk Professor of Entrepreneurship at the University of Toronto's Rotman School of Management in Toronto, Ontario, as well as founder of the Creative Destruction Lab, a Canada-based training and mentoring program for founders of technology startups.* ***Joshua S. Gans*** *holds the Jeffrey S. Skoll Chair in Technical Innovation and Entrepreneurship at the Rotman School and is chief economist of the Creative Destruction Lab.* ***Avi Goldfarb*** *is the Ellison Professor of Marketing at the Rotman School and chief data scientist of the Creative Destruction Lab. Comment on this article at http://sloanreview.mit.edu/x/58311, or contact the authors at smrfeedback@mit.edu.*

ACKNOWLEDGMENTS

The authors wish to thank James Bergstra, Tim Bresnahan, and Graham Taylor for helpful discussions. All views remain our own.

Reprint 58311.

Digital Today, Cognitive Tomorrow

Ginni Rometty

Digital is not the destination. Rather, it is laying the foundation for a more profound transformation to come.

Editor's Note: This article is one of a special series of 14 commissioned essays MIT Sloan Management Review *is publishing to celebrate the launch of our new Frontiers initiative. Each essay gives the author's response to this question:*

"Within the next five years, how will technology change the practice of management in a way we have not yet witnessed?"

In today's economy, we are seeing companies, business models, products, and processes undergoing major transformation. Enterprises and governments are rapidly "becoming digital" as they seek to capture the cost savings, agility, and collaboration enabled by cloud, analytics, mobile, and social technologies.

However, digital is not the destination. Rather, it is laying the foundation for a much more profound transformation to come. Within five years, I believe all major business decisions will be enhanced by *cognitive* technologies.

I sensed the magnitude of the transition for the first time in 2011, when I watched IBM's Watson system win on "Jeopardy!" At the time, I felt that I was watching history in the making: The technology known as artificial intelligence (AI) was finally moving from the lab into the world.

Why are we seeing this now?

First, the technologies required for cognitive systems — not just AI, but a broad spectrum of capabilities that include natural language processing, human-computer interaction, deep learning, neural nets, and more — have made exponential advances in recent years.

Second, the abundance of data being generated throughout the world today requires cognitive technology. Much of this data is "unstructured": video, audio, sensor outputs, and everything we encode in language, from medical journals to tweets. However, such unstructured data are "dark" to traditional computer systems. Computers can capture, move, and store the data, but they cannot understand what the data mean (which is why cognitive systems are so vital).

Finally, and most important, we will see systems that learn. We *need* systems that learn. Think of the challenges and issues we face today: predicting risk in financial markets, anticipating consumer behavior, ensuring public safety, managing traffic, optimizing global supply chains, personalizing medicine, treating chronic diseases, and preventing pandemics.

The challenges today go beyond information overload. In many ways, we live in an era of *cognitive overload*, characterized by an exponential increase in the complexity of decision making. It's impossible to create protocols, algorithms, or software code to successfully anticipate all the potential permutations, trajectories, and interactions. But cognitive systems are not simply programmed. They actually improve with use, as they receive expert training, interact with clients and customers, and ingest data from their own experiences, successes, and failures.

Some people think of cognitive systems as supercomputers, and there is no question that the computational power behind systems like Watson is considerable. But thanks to the increasing prevalence of application program interfaces (APIs) — which can be encoded into digital services and easily accessed or combined in new ways in the cloud — it's possible to build a kind of thinking into virtually every digital application, product, and system.

And because we can, we will. If it's digital today, it will be cognitive tomorrow — and not a distant tomorrow. IDC Research Inc. has estimated that by 2018, more than half of the teams developing apps will embed some kind of cognitive services in them, up from 1% in 2015.

Cognitive systems are already transforming everything from the world-changing to the everyday. For example, cognitive oncology is a reality thanks to technology developed in partnership with Memorial Sloan Kettering Cancer Center in New York City that helps oncologists identify personalized, evidence-based treatment options based on massive volumes of data. This breakthrough technology is now helping scale access to knowledge at Bumrungrad International Hospital in Thailand, Manipal Hospitals in India, and more than 20 hospitals in China. Cognitive assistants are at work helping build more intimate, personalized relationships at the Brazilian bank Banco Bradesco, the insurance company GEICO, and the retailer The North Face. Dublin-based Medtronic plc, a global health care solutions company, is creating a cognitive app for people with diabetes to predict a hypoglycemic event hours in advance. These are just a few examples of organizations that are using cognitive systems today.

It's important to note that we are not talking about the AI we see in movies. This isn't about creating a synthetic brain or an artificial human. Rather, this is about augmenting human intelligence. Indeed, there is nothing in either cognitive science or its application that implies either sentience or autonomy.

Of course, anyone familiar with the history of technology knows that technological breakthroughs often have major effects on work and jobs. Some jobs are eliminated, while others are created. With cognitive systems, we are already beginning to see the emergence of new disciplines — from data curation to system training, as well as new fields of scientific knowledge and new kinds of work — quite possibly more than in any prior technology revolution.

Data can be seen as the world's great new resource. What steam power, electricity, and fossil fuels did for earlier eras, data promises to do for the 21st century — if we can mine, refine, and apply it. Thanks to the new generation of cognitive technologies, we can. Intelligence *augmentation* — IA as opposed to AI — will change how humans work together, make decisions, and manage organizations.

About The Author

Ginni Rometty is chairman, president, and CEO of IBM Corp.

Predicting a Future Where the Future Is Routinely Predicted

Andrew W. Moore

Artificial intelligence systems will be able to give managers real-time insights about their business operations — as well as detect early warnings of problems before they occur.

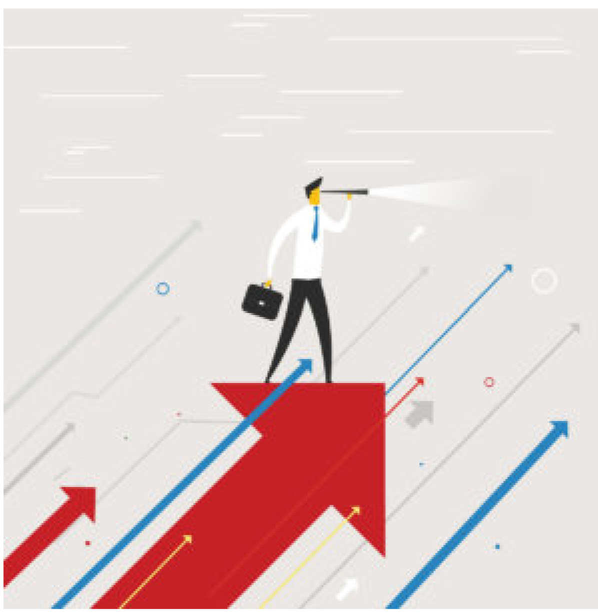

Editor's Note: This article is one of a special series of 14 commissioned essays MIT Sloan Management Review *is publishing to celebrate the launch of our new Frontiers initiative. Each essay gives the author's response to this question:*

"Within the next five years, how will technology change the practice of management in a way we have not yet witnessed?"

Workers on the factory floor have suddenly gathered at a point along the production line. Some are scratching their heads. Others are gesticulating wildly. Most stand with their hands in their pockets. Something is wrong, and no one has thought to call management.

In the near future, scenes like this one will be obsolete. Thanks to advances in artificial intelligence (AI), managers will be alerted to workplace anomalies as soon they occur. Unusual behaviors will be identified in real time by cameras and image-processing software that continuously analyze and *comprehend* scenes across the enterprise.

The hunch-based bets of the past already are giving way to far more reliable data-informed decisions. But AI will take this further. By analyzing new types of data, including real-time video and a range of other inputs, AI systems will be able to provide managers with insights about what is happening in their businesses at any moment in time and, even more significantly, detect early warnings of bigger problems that have yet to materialize.

As a researcher, I learned to appreciate the value of early warnings some years ago, while developing algorithms for analyzing data from hospital emergency rooms and drugstores. We discovered that we could alert public health officials to potential epidemics and even the possibility of biological warfare attacks, giving them time to take countermeasures to slow the spread of disease.

Similar analytic techniques are being deployed to detect early signs of problems in aircraft. The detailed maintenance and flight logs for the U.S. Air Force's aging fleet of F-16 fighter jets are analyzed automatically to identify patterns of equipment failures that may affect only a handful of aircraft at present, but have the potential to become widespread. This has enabled officials to confirm and diagnose problems and take corrective action before the problems spread.

With AI, we can have machines look for millions of worrying patterns in the time it would take a human to consider just one. But that capability includes a terrible dilemma: the multiple hypotheses problem. If you sound an alarm whenever something is anomalous at a 99% confidence level, and you check millions of things an hour, then you will receive hundreds of alarms every minute.

Statisticians and AI researchers are working together to identify situations and conditions that tend to sound false alarms, like a truckload of potassium-rich bananas that can set off a radiation detector meant to identify nuclear materials. By reducing the risk of false alarms, it will be possible to set sensor thresholds even lower, enhancing sensitivity.

The predictive benefits of AI will stretch well beyond equipment and process analysis. For instance, researchers are having great success with algorithms that closely monitor subtle facial movements to assess the emotional and psychological states of individuals. Some of the most interesting applications now are in the mental health sphere, but imagine if the same tools could be deployed on checkout lines in stores, lines at theme parks, or security queues at airports. Are your customers happy or agitated? Executives wouldn't need to wait weeks or even days for a survey to be completed; these systems could tell you the emotional state of your customers right now.

Other researchers are deploying AI in the classroom. When I taught, I couldn't tell whether the lecture I was giving was any good — at least not when it would still benefit me or my students. But simple sensors like microphones and cameras can be used by AI programs to detect when active learning is taking place. Just the sounds alone — Who's talking? Who isn't? Is anyone laughing? — can provide a lot of clues about teaching effectiveness and when adjustments should be made.

Such a tool could also be used to gauge whether your employees are buying in to what you're sharing with them in a meeting, or if potential customers are engaged during focus groups. A "managerial Siri" might take this even further. If you asked your digital assistant, "Do the folks in my staff meeting seem to be more engaged since we had the retreat?" you might receive an answer such as, "Yes, there is an increase in eye contact between team members and a slight but significant increase in laughter."

As a manager, I absolutely detest being surprised. And like everyone else, despite the petabytes of data at my fingertips, I too often am. But AI doesn't get overwhelmed by the size and complexity of information the way we humans do. Thus, its promise to keep managers more in the know about what's really happening across their enterprise is truly profound.

About The Author

Andrew W. Moore is dean of the School of Computer Science at Carnegie Mellon University in Pittsburgh, Pennsylvania.

[AUTOMATION]

Just How Smart Are Smart Machines?

The number of sophisticated cognitive technologies that might be capable of cutting into the need for human labor is expanding rapidly. But linking these offerings to an organization's business needs requires a deep understanding of their capabilities.

BY THOMAS H. DAVENPORT AND JULIA KIRBY

If popular culture is an accurate gauge of what's on the public's mind, it seems everyone has suddenly awakened to the threat of smart machines. Several recent films have featured robots with scary abilities to outthink and manipulate humans. In the economics literature, too, there has been a surge of concern about the potential for soaring unemployment as software becomes increasingly capable of decision making. Yet managers we talk to don't expect to see machines displacing knowledge workers anytime soon — they expect computing technology to augment rather than replace the work of humans. In the face of a sprawling and fast-evolving set of opportunities, their challenge is figuring out what forms the augmentation should take. Given the kinds of work managers oversee, what cognitive technologies should they be applying now, monitoring closely, or helping to build?

To help, we have developed a simple framework that plots cognitive technologies along two dimensions. (See "What Today's Cognitive Technologies Can — and Can't — Do," p. 23.) First, it recognizes that these tools differ according to how autonomously they can apply their intelligence. On the low end, they simply respond to human queries and instructions; at the (still theoretical) high end, they formulate their own objectives. Second, it reflects the type of tasks smart machines are being used to perform, moving from conventional numerical analysis to performance of digital and physical tasks in the real world. The breadth of inputs and data

types in real-world tasks makes them more complex for machines to accomplish.

By putting those two dimensions together, we create a matrix into which we can place all of the multitudinous technologies known as "smart machines." More important, this helps to clarify today's limits to machine intelligence and the challenges technology innovators are working to overcome next. Depending on the type of task a manager is targeting for redesigned performance, this framework reveals the various extents to which it might be performed autonomously and by what kinds of machines.

Four Levels of Intelligence

Clearly, the level of intelligence of smart machines is increasing. The general trend is toward greater autonomy in decision making — from machines that require a highly structured data and decision context to those capable of deciphering a more complex context.

Support for Humans For decades, the prevailing assumption has been that cognitive technologies would provide insight to human decision makers — what used to be known as "decision support." Even with IBM Corp.'s Watson and many of today's other cognitive systems, most people assume that the machine will offer a recommended decision or course of action but that a human will make the final decision.

Repetitive Task Automation It is a relatively small step to go from having machines support humans to having the machines make decisions, particularly in structured contexts. Automated decision making has been gaining ground in recent years in several domains, such as insurance underwriting and financial trading; it typically relies on a fixed set of rules or algorithms, so performance doesn't improve without human intervention. Typically, people monitor system performance and fine-tune the algorithms.

Context Awareness and Learning Sophisticated cognitive technologies today have some degree of real-time contextual awareness. As data flow more continuously and voluminously, we need technologies that can help us make sense of the data in real time — detecting anomalies, noticing patterns, and anticipating what will happen next. Relevant information might include location, time, and/or a user's identity, which might be used to make recommendations (for example, the best

Just How Smart Are Smart Machines
(Continued from page 21)

route to work based on the time of day, current traffic levels, and the driver's preference for highways versus back roads).

One of the hallmarks of today's cognitive computing is its ability to learn and improve performance. Much of the learning takes place through continuous analysis of real-time data, user feedback, and new content from text-based articles. In settings where results are measurable, learning-oriented systems will ultimately deliver benefits in the form of better stock trading decisions, more accurate driving time predictions, and more precise medical diagnoses.

Self-Awareness So far, machines with self-awareness and the ability to form independent objectives reside only in the realm of fiction. With substantial self-awareness, computers may eventually gain the ability to work beyond human levels of intelligence across multiple contexts, but even the most optimistic experts say that general intelligence in machines is three to four decades away.

Four Cognitive Task Types

A straightforward way to sort out tasks performed by machines is according to whether they process only numbers, text, or images — the building blocks of cognition — or whether they know enough to take informed actions in the digital or physical world.

Analyzing Numbers The root of all cognitive technologies is computing machines' superior performance at analyzing numbers in structured formats (typically, rows and columns). Classically, this numerical analysis was applied purely in support of human decision makers. People continued to perform the front-end cognitive tasks of creating hypotheses and framing problems, as well as the back-end interpretation of the numbers' implications for decisions. Even as analysts added more visual analytics displays and more predictive analytics in the past decade, people still did the interpretation.

Today, companies are increasingly embedding analytics into operational systems and processes to make repetitive automated decisions, which enables dramatic increases in both speed and scale. And whereas it used to take a human analyst to develop embedded models, "machine learning" methods can produce models in an automated or semiautomated fashion.

Analyzing Words and Images A key aspect of human cognition is the ability to read words and images and to determine their meaning and significance. But today, a wide variety of technological tools, such as machine learning, natural language processing, neural networks, and deep learning, can classify, interpret, and generate words. Some of them can also analyze and identify images.

The earliest intelligent applications involving words and images involved text, image, and speech recognition to allow humans to communicate with computers. Today, of course, smartphones "understand" human speech and text and can recognize images. These capabilities are hardly perfect, but they are widely used in many applications.

When words and images are analyzed on a large scale, this comprises a different category of capability. One such application involves translating large volumes of text across languages. Another is to answer questions as a human would. A third is to make sense of language in a way that can either summarize it or generate new passages.

IBM Watson was the first tool capable of ingesting, analyzing, and "understanding" text well enough to respond to detailed questions. However, it doesn't deal with structured numerical data, nor can it understand relationships between variables or make predictions. It's also not well suited for applying rules or analyzing options on decision trees. However, IBM is rapidly adding new capabilities included in our matrix, including image analysis.

There are other examples of word and image systems. Most were developed for particular applications and are slowly being modified to handle other types of cognitive situations. Digital Reasoning Systems Inc., for example, a company based in Franklin, Tennessee, that developed cognitive computing software for national security purposes, has begun to market intelligent software that analyzes employee communications in financial institutions to determine the likelihood of fraud. Another company, IPsoft Inc., based in New York City, processes spoken words with an intelligent customer agent programmed to interpret what customers want and, when possible, do it for them.

IPsoft, Digital Reasoning, and the original Watson all use similar components, including the ability to classify parts of speech, to identify key entities and facts in text, to show the relationships among entities and facts in a graphical diagram, and to relate entities and relationships with objectives. This category of application is best suited for situations with much more — and more rapidly changing — codified textual information than any human could possibly absorb and retain.

Image identification and classification are hardly new. "Machine vision" based on geometric pattern matching technology has been used for decades to locate parts in production lines and read bar codes. Today, many companies want to perform more sensitive vision tasks such as facial recognition, classification of photos on the Internet, or assessment of auto collision damage. Such tasks are based on machine learning and neural network analysis that can match particular patterns of pixels to recognizable images.

The most capable machine learning systems have the ability to "learn" — their decisions get better with more data, and they "remember" previously ingested information. For example, as Watson is introduced to new information, its reservoir of

WHAT TODAY'S COGNITIVE TECHNOLOGIES CAN — AND CAN'T — DO

Mapping cognitive technologies by how autonomously they work and the tasks they perform shows the current state of smart machines — and anticipates how future technologies might unfold.

LEVELS OF INTELLIGENCE

TASK TYPE	SUPPORT FOR HUMANS	REPETITIVE TASK AUTOMATION	CONTEXT AWARENESS AND LEARNING	SELF-AWARENESS
Analyze Numbers	Business intelligence, data visualization, hypothesis-driven analytics	Operational analytics, scoring, model management	Machine learning, neural networks	Not yet
Analyze Words and Images	Character and speech recognition	Image recognition, machine vision	IBM Watson, natural language processing	Not yet
Perform Digital Tasks	Business process management	Rules engines, robotic process automation	Not yet	Not yet
Perform Physical Tasks	Remote operation of equipment	Industrial robotics, collaborative robotics	Autonomous robots, vehicles	Not yet

THE GREAT CONVERGENCE

information expands. Other systems in this category get better at their cognitive task by having more data for training purposes. But as Mike Rhodin, senior vice president of business development for IBM Watson, noted, "Watson doesn't have the ability to think on its own," and neither does any other intelligent system thus far created.

Performing Digital Tasks One of the more pragmatic roles for cognitive technology in recent years has been to automate administrative tasks and decisions. In order to make automation possible, two technical capabilities are necessary. First, you need to be able to express the decision logic in terms of "business rules." Second, you need technologies that can move a case or task through the series of steps required to complete it. Over the past couple of decades, automated decision-making tools have been used to support a wide variety of administrative tasks, from insurance policy approvals to information technology operations to high-speed trading.

Lately, companies have begun using "robotic process automation," which uses work flow and business rules technology to interface with multiple information systems as if it were a human user. Robotic process technology has become popular in banking (for back-office customer service tasks, such as replacing a lost ATM card), insurance (for processing claims and payments), information technology (IT) (for monitoring system error messages and fixing simple problems), and supply chain management (for processing invoices and responding to routine requests from customers and suppliers).

The benefits of process automation can add up quickly. An April 2015 case study at Telefónica O2, the second-largest mobile carrier in the United Kingdom, found that the company had automated over 160 process areas using software "robots." The overall three-year return on investment was between 650% and 800%.

Performing Physical Tasks Physical task automation is, of course, the realm of robots. Though people love to call every form of automation technology a robot, one of Merriam-Webster's definitions of robot is "a machine that can do the work of a person and that works automatically or is controlled by a computer."

In 2014, companies installed about 225,000 industrial robots globally, more than one-third of them in the automotive industry. However, robots often fall well short of expectations. In 2011, the founder of Foxconn Technology Co., Ltd., a Taiwan-based multinational electronics contract manufacturing company, said he would install one million robots within three years, replacing one million workers. However, the company found that employing only robots to build smartphones was easier said than done. To assemble new iPhone models in 2015, Foxconn planned to hire more than 100,000 new workers and install about 10,000 new robots.

Historically, robots that replaced humans required a high level of programming to do repetitive tasks. For safety reasons, they had to be segregated from human workers. However, a new type of robots — often called "collaborative robots" — can work safely alongside humans. They can be programmed simply by having a human move their arms.

Robots have varying degrees of autonomy. Some, such as remotely piloted drone

Just How Smart Are Smart Machines
(Continued from page 23)

aircraft and robotic surgical instruments and mining equipment, are designed to be manipulated by humans. Others become at least semiautonomous once programmed but have limited ability to respond to unexpected conditions. As robots get more intelligence, better machine vision, and increased ability to make decisions, they will integrate other types of cognitive technologies while also having the ability to transform the physical environment. IBM Watson software, for example, has been installed in several different types of robots.

The Great Convergence

Slowly but surely, the worlds of artificially intelligent software and robots seem to be converging, and the boundaries between different cognitive technologies are blurring. In the future, robots will be able to learn and sense context, robotic process automation and other digital task tools will improve, and smart software will be able to analyze more intricate combinations of numbers, text, and images.

We anticipate that companies will develop cognitive solutions using the building blocks of application program interfaces (APIs). One API might handle language processing, another numerical machine learning, and a third question-and-answer dialogue. While these elements would interact with each other, determining which APIs are required will demand a sophisticated understanding of cognitive solution architectures.

This modular approach is the direction in which key vendors are moving. IBM, for example, has disaggregated Watson into a set of services — a "cognitive platform," if you will — available by subscription in the cloud. Watson's original question-and-answer services have been expanded to include more than 30 other types, including "personality insights" to gauge human behavior, "visual recognition" for image identification, and so forth. Other vendors of cognitive technologies, such as Cognitive Scale Inc., based in Austin, Texas, are also integrating multiple cognitive capabilities into a "cognitive cloud."

Despite the growing capabilities of cognitive technologies, most organizations that are exploring them are starting with small projects to explore the technology in a specific domain. But others have much bigger ambitions. For example, Memorial Sloan Kettering Cancer Center, in New York City, and the University of Texas MD Anderson Cancer Center, in Houston, Texas, are taking a "moon shot" approach, marshaling cognitive tools like Watson to develop better diagnostic and treatment approaches for cancer.

Designing a Cognitive Architecture

Hardware and software will continue to get better, but rather than waiting for next-generation options, managers should be introducing cognitive technologies to workplaces now and discovering their human-augmenting value. The most sophisticated managers will create IT architectures that support more than one application. Indeed, we expect to see organizations building "cognitive architectures" that interface with, but are distinct from, their regular IT architectures. What would that mean? We think a well-designed cognitive architecture would emphasize several attributes:

The Ability to Handle a Variety of Data Types Cognitive insights don't just come from a single data type (text, for example). In the future, they will come from combining text, numbers, images, speech, genomic data, and so forth to develop broad situational awareness.

The Ability to Learn Although this should be the essence of cognitive technologies, most systems today (such as rules engines and robotic process automation) don't improve themselves. If you have a choice between a system that learns and one that doesn't, go with the former.

Transparency Humans and cognitive technologies will be working together for the foreseeable future. Humans will always want to know how the cognitive technologies came up with their decision or recommendation. If people can't open the "black box," they won't trust it. This is a key aspect of augmentation, and one that will facilitate rapid adoption of these technologies.

A Variety of Human Roles Once programmed, some cognitive technologies, like most industrial robots, run their assigned process. By contrast, with surgical robots it's assumed that a human is in charge. In the future, we will probably need multiple control modes. As with self-driving vehicles, there needs to be a way for the human to take control. Having multiple means of control is another way to facilitate augmentation rather than automation.

Flexible Updating and Modification One of the reasons why rule-based systems have become successful in insurance and banking is that users can modify the rules. But modifying and updating most cognitive systems is currently a task only for experts. Future systems will need to be more flexible.

Robust Reporting Capabilities Cognitive technologies will need to be accountable to the rest of the organization, as well as to other stakeholders. We've spoken, for example, with representatives of several companies using automated systems to buy and place digital ads, and they say that customers insist on detailed reporting so that the data can be "sliced and diced" in many different ways.

State-of-the-Art IT Hygiene Cognitive technologies will need all the attributes of modern information systems, including an easy user interface, state-of-the-art data security, and the ability to handle multiple users at once. Companies won't want to compromise on any of these objectives in the cognitive space, and eventually they won't have to.

What's more, if the managerial goal is augmentation rather than automation, it's essential to understand how human capabilities fit into the picture. People will continue to have advantages over even the smartest machines. They are better able to interpret unstructured data — for example, the meaning of a poem or whether an image is of a good neighborhood or a bad one. They have the cognitive breadth to simultaneously do a lot of different things well. The judgment and flexibility that come with these basic advantages will continue to be the basis of any enterprise's ability to innovate, delight customers, and prevail in competitive markets — where, soon enough, cognitive technologies will be ubiquitous.

Clearly, smart machines are advancing at the things they do well at a much faster rate than we humans are. And granted, many workers will need to call on and cultivate different capabilities than the ones they have relied on in the past. But for the foreseeable future, there are still unlimited ways for humans to contribute tremendous value. To the extent that wise managers leverage their talents with advanced technology, we can all stop dreading the rise of smart machines.

Thomas H. Davenport *is the President's Distinguished Professor of Information Technology and Management at Babson College in Wellesley, Massachusetts, as well as a fellow of the MIT Initiative on the Digital Economy.* ***Julia Kirby*** *is a Boston-based editor and writer. They are the authors of the book* Only Humans Need Apply: Winners and Losers in the Age of Smart Machines *(forthcoming, HarperCollins). Comment on this article at http://sloanreview.mit.edu/x/57306, or contact the authors at smrfeedback@mit.edu.*

Reprint 57306.