MITSloan
Management Review

What to Expect From Artificial Intelligence

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

Ajay Agrawal
Joshua S. Gans
Avi Goldfarb

Reprint #58311 | http://mitsmr.com/2jZdf1Y

What to Expect From Artificial Intelligence

AJAY AGRAWAL, JOSHUA S. GANS, AND AVI GOLDFARB

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

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 both in color and shape. We could continue encoding our knowledge of 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.

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.

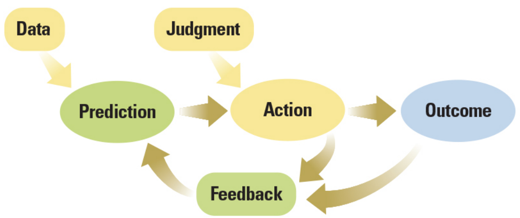

Seen in this way, it's useful to distinguish between the cost versus the value of prediction. As we have noted, advances in AI have reduced the *cost* of prediction. Just as important is what has happened to the *value*. Prediction becomes more valuable when data is more widely available and more accessible. 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 itself millions of times. It then incorporated the feedback on actions and outcomes to develop more accurate predictions and new strategies.

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 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, much as returns to numeracy and literacy increased with the diffusion of computing. 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.

About the Authors

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.

Acknowledgments

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

MITSloan
Management Review

PDFs ■ Reprints ■ Permission to Copy ■ Back Issues

Articles published in MIT Sloan Management Review are copyrighted by the Massachusetts Institute of Technology unless otherwise specified at the end of an article.

MIT Sloan Management Review articles, permissions, and back issues can be purchased on our Web site: sloanreview.mit.edu or you may order through our Business Service Center (9 a.m.-5 p.m. ET) at the phone numbers listed below. Paper reprints are available in quantities of 250 or more.

To reproduce or transmit one or more MIT Sloan Management Review articles by electronic or mechanical means (including photocopying or archiving in any information storage or retrieval system) **requires written permission.**

To request permission, use our Web site: sloanreview.mit.edu
or
E-mail: smr-help@mit.edu
Call (US and International):617-253-7170 Fax: 617-258-9739

Posting of full-text SMR articles on publicly accessible Internet sites is prohibited. To obtain permission to post articles on secure and/or password-protected intranet sites, e-mail your request to smr-help@mit.edu.