

Article

McKinsey Quarterly

June 2015

An executive's guide to machine learning

By Dorian Pyle and Cristina San Jose

It's no longer the preserve of artificial-intelligence researchers and born-digital companies like Amazon, Google, and Netflix.

Machine learning is based on algorithms that can learn from data without relying on rules-based programming. It came into its own as a scientific discipline in the late 1990s as steady advances in digitization and cheap computing power enabled data scientists to stop building finished models and instead train computers to do so. The unmanageable volume and complexity of the big data that the world is now swimming in have increased the potential of machine learning—and the need for it.

In 2007 Fei-Fei Li, the head of Stanford's Artificial Intelligence Lab, gave up trying to program computers to recognize objects and began labeling the millions of raw images that a child might encounter by age three and feeding them to computers. By being shown thousands and thousands of labeled data sets with instances of, say, a cat, the machine could shape its own rules for deciding whether a particular set of digital pixels was, in fact, a cat.¹ Last November, Li's team unveiled a program that identifies the visual elements of any picture with a high degree of accuracy. IBM's Watson machine relied on a similar self-generated scoring system among hundreds of potential answers to crush the world's best *Jeopardy!* players in 2011.

Dazzling as such feats are, machine learning is nothing like learning in the human sense (yet). But what it already does extraordinarily well—and will get better at—is relentlessly chewing through any amount of data and every combination of variables. Because machine learning's emergence as a mainstream management tool is relatively recent, it often raises questions. In this article, we've posed some that we often hear and answered them in a way we hope will be useful for any executive. Now is the time to grapple with these issues, because the competitive significance of business models turbocharged by machine learning is poised to surge. Indeed, management author Ram Charan suggests that “any organization that is not a math house now or is unable to become one soon is already a legacy company.”²

1. How are traditional industries using machine learning to gather fresh business insights?

Well, let's start with sports. This past spring, contenders for the US National Basketball Association championship relied on the analytics of Second Spectrum, a California machine-learning start-up. By digitizing the past few seasons' games, it has created predictive models that allow a coach to distinguish between, as CEO Rajiv Maheswaran puts it, “a bad shooter who takes good shots and a good shooter who takes bad shots”—and to adjust his decisions accordingly.

You can't get more venerable or traditional than General Electric, the only member of the original Dow Jones Industrial Average still around after 119 years. GE already makes hundreds of millions of dollars by crunching the data it collects from deep-sea oil wells or jet engines to optimize performance, anticipate breakdowns, and streamline maintenance. But Colin Parris, who joined GE Software from IBM late last year as vice president of software research, believes that continued advances in data-processing power, sensors, and predictive algorithms will soon give his company the same sharpness of insight into the individual vagaries of a jet engine that Google has into the online behavior of a 24-year-old netizen from West Hollywood.

2. What about outside North America?

In Europe, more than a dozen banks have replaced older statistical-modeling approaches with machine-learning techniques and, in some cases, experienced 10 percent increases in sales of new products, 20 percent savings in capital expenditures, 20 percent increases in cash collections, and 20 percent declines in churn. The banks have achieved these gains by devising new recommendation engines for clients in retailing and in small and medium-sized companies. They have also built microtargeted models that more accurately forecast who will cancel service or default on their loans, and how best to intervene.

Closer to home, as a recent article in *McKinsey Quarterly* notes,³ our colleagues have been applying hard analytics to the soft stuff of talent management. Last fall, they tested the ability of three algorithms developed by external vendors and one built internally to forecast, solely by examining scanned résumés, which of more than 10,000 potential recruits the firm would have accepted. The predictions strongly correlated with the real-world results. Interestingly, the machines accepted a slightly higher percentage of female candidates, which holds promise for using analytics to unlock a more diverse range of profiles and counter hidden human bias.

As ever more of the analog world gets digitized, our ability to learn from data by developing and testing algorithms will only become more important for what are now seen as traditional businesses. Google chief economist Hal Varian calls this “computer *kaizen*.” For “just as mass production changed the way products were assembled and continuous improvement changed how manufacturing was done,” he says, “so continuous [and often automatic] experimentation will improve the way we optimize business processes in our organizations.”⁴

3. What were the early foundations of machine learning?

Machine learning is based on a number of earlier building blocks, starting with classical statistics. Statistical inference does form an important foundation for the current implementations of artificial intelligence. But it's important to recognize that classical statistical techniques were developed between the 18th and early 20th centuries for much smaller data sets than the ones we now have at our disposal. Machine learning is unconstrained by the preset assumptions of statistics. As a result, it can yield insights that human analysts do not see on their own and make predictions with ever-higher degrees of accuracy.

More recently, in the 1930s and 1940s, the pioneers of computing (such as Alan Turing, who had a deep and abiding interest in artificial intelligence) began formulating and tinkering with the basic techniques such as neural networks that make today's machine learning possible. But those techniques stayed in the laboratory longer than many technologies did and, for the most part, had to await the development and infrastructure of powerful computers, in the late 1970s and early 1980s. That's probably the starting point for the machine-learning adoption curve. New technologies introduced into modern economies—the steam engine, electricity, the electric motor, and computers, for example—seem to take about 80 years to transition from the laboratory to what you might call cultural invisibility. The computer hasn't faded from sight just yet, but it's likely to by 2040. And it probably won't take much longer for machine learning to recede into the background.

4. What does it take to get started?

C-level executives will best exploit machine learning if they see it as a tool to craft and implement a strategic vision. But that means putting strategy first. Without strategy as a starting point, machine learning risks becoming a tool buried inside a company's routine operations: it will provide a useful service, but its long-term value will probably be limited to an endless repetition of “cookie cutter” applications such as models for acquiring, stimulating, and retaining customers.

We find the parallels with M&A instructive. That, after all, is a means to a well-defined end. No sensible business rushes into a flurry of acquisitions or mergers and then just sits back to see what happens. Companies embarking on machine learning should make the same three commitments companies make before embracing M&A. Those commitments are, first, to investigate all feasible alternatives; second, to pursue the strategy wholeheartedly at the C-suite level; and, third, to use (or if necessary acquire) existing expertise and knowledge in the C-suite to guide the application of that strategy.

The people charged with creating the strategic vision may well be (or have been) data scientists. But as they define the problem and the desired outcome of the strategy, they will need guidance from C-level colleagues overseeing other crucial strategic initiatives. More broadly, companies must have two types of people to unleash the potential of machine learning. “Quants” are schooled in its language and methods. “Translators” can bridge the disciplines of data, machine learning, and decision making by reframing the quants’ complex results as actionable insights that generalist managers can execute.

Access to troves of useful and reliable data is required for effective machine learning, such as Watson’s ability, in tests, to predict oncological outcomes better than physicians or Facebook’s recent success teaching computers to identify specific human faces nearly as accurately as humans do. A true data strategy starts with identifying gaps in the data, determining the time and money required to fill those gaps, and breaking down silos. Too often, departments hoard information and politicize access to it—one reason some companies have created the new role of chief data officer to pull together what’s required. Other elements include putting responsibility for generating data in the hands of frontline managers.

Start small—look for low-hanging fruit and trumpet any early success. This will help recruit grassroots support and reinforce the changes in individual behavior and the employee buy-in that ultimately determine whether an organization can apply machine learning effectively. Finally, evaluate the results in the light of clearly identified criteria for success.

5. What's the role of top management?

Behavioral change will be critical, and one of top management's key roles will be to influence and encourage it. Traditional managers, for example, will have to get comfortable with their own variations on A/B testing, the technique digital companies use to see what will and will not appeal to online consumers. Frontline managers, armed with insights from increasingly powerful computers, must learn to make more decisions on their own, with top management setting the overall direction and zeroing in only when exceptions surface. Democratizing the use of analytics—providing the front line with the necessary skills and setting appropriate incentives to encourage data sharing—will require time.

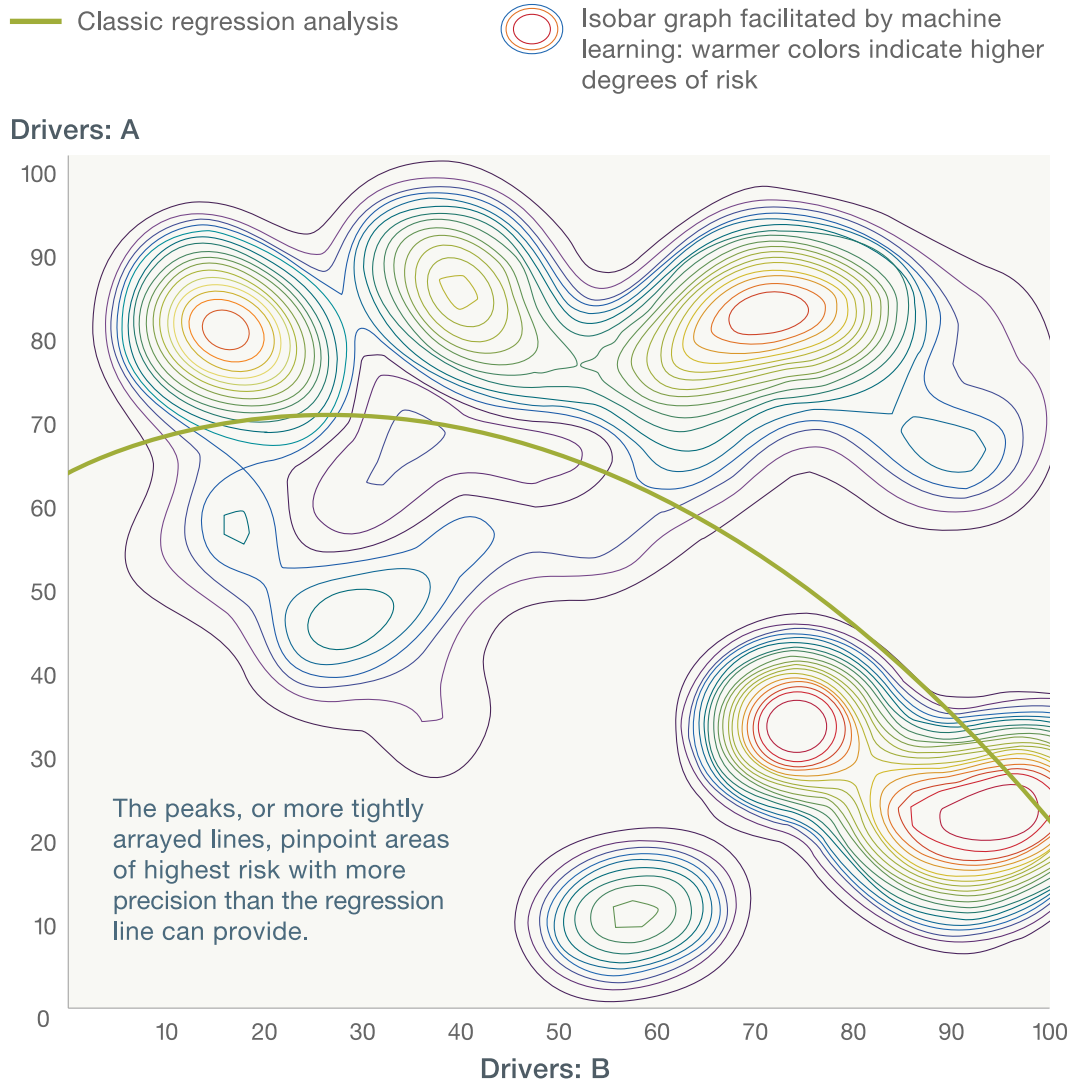
C-level officers should think about applied machine learning in three stages: machine learning 1.0, 2.0, and 3.0—or, as we prefer to say, description, prediction, and prescription. They probably don't need to worry much about the description stage, which most companies have already been through. That was all about collecting data in databases (which had to be invented for the purpose), a development that gave managers new insights into the past. OLAP—online analytical processing—is now pretty routine and well established in most large organizations.

There's a much more urgent need to embrace the prediction stage, which is happening right now. Today's cutting-edge technology already allows businesses not only to look at their historical data but also to predict behavior or outcomes in the future—for example, by helping credit-risk officers at banks to assess which customers are most likely to default or by enabling telcos to anticipate which customers are especially prone to “churn” in the near term (exhibit).

Exhibit

The contrast between routine statistical analysis and data generated by machine learning can be quite stark.

Value at risk from customer churn, telecom example



McKinsey&Company

A frequent concern for the C-suite when it embarks on the prediction stage is the quality of the data. That concern often paralyzes executives. In our experience, though, the last decade's IT investments have equipped most companies with sufficient information to obtain new insights even from incomplete, messy data sets, provided of course that those companies choose the right algorithm. Adding exotic new data sources may be of only marginal benefit compared with what can be mined from existing data warehouses. Confronting that challenge is the task of the “chief data scientist.”

Prescription—the third and most advanced stage of machine learning—is the opportunity of the future and must therefore command strong C-suite attention. It is, after all, not enough just to predict *what* customers are going to do; only by understanding *why* they are going to do it can companies encourage or deter that behavior in the future. Technically, today's machine-learning algorithms, aided by human translators, can already do this. For example, an international bank concerned about the scale of defaults in its retail business recently identified a group of customers who had suddenly switched from using credit cards during the day to using them in the middle of the night. That pattern was accompanied by a steep decrease in their savings rate. After consulting branch managers, the bank further discovered that the people behaving in this way were also coping with some recent stressful event. As a result, all customers tagged by the algorithm as members of that microsegment were automatically given a new limit on their credit cards and offered financial advice.

The prescription stage of machine learning, ushering in a new era of man-machine collaboration, will require the biggest change in the way we work. While the machine identifies patterns, the human translator's responsibility will be to interpret them for different microsegments and to recommend a course of action. Here the C-suite must be directly involved in the crafting and formulation of the objectives that such algorithms attempt to optimize.

6. This sounds awfully like automation replacing humans in the long run. Are we any nearer to knowing whether machines will replace managers?

It's true that change is coming (and data are generated) so quickly that human-in-the-loop involvement in all decision making is rapidly becoming impractical. Looking three to five years out, we expect to see far higher levels of artificial intelligence, as well as the development of distributed autonomous corporations. These self-motivating, self-contained agents, formed as corporations, will be able to carry out set objectives autonomously, without any direct human supervision. Some DACs will certainly become self-programming.

One current of opinion sees distributed autonomous corporations as threatening and inimical to our culture. But by the time they fully evolve, machine learning will have become culturally invisible in the same way technological inventions of the 20th century disappeared into the background. The role of humans will be to direct and guide the algorithms as they attempt to achieve the objectives that they are given. That is one lesson of the automatic-trading algorithms which wreaked such damage during the financial crisis of 2008.

No matter what fresh insights computers unearth, only human managers can decide the essential questions, such as which critical business problems a company is really trying to solve. Just as human colleagues need regular reviews and assessments, so these "brilliant machines" and their works will also need to be regularly evaluated, refined—and, who knows, perhaps even fired or told to pursue entirely different paths—by executives with experience, judgment, and domain expertise.

The winners will be neither machines alone, nor humans alone, but the two working together effectively.

7. So in the long term there's no need to worry?

It's hard to be sure, but distributed autonomous corporations and machine learning should be high on the C-suite agenda. We anticipate a time when the philosophical discussion of what intelligence, artificial or otherwise, might be will end because there will be no such thing as intelligence—just processes. If distributed autonomous corporations act intelligently, perform intelligently, and respond intelligently, we will cease to debate whether high-level intelligence other than the human variety exists. In the meantime, we must all think about what we want these entities to do, the way we want them to behave, and how we are going to work with them.

About the Authors

Dorian Pyle is a data expert in McKinsey's Miami office, and **Cristina San Jose** is a principal in the Madrid office.