

# Five Rules for Fixing AI in Business

**FEBRUARY 25, 2022**

By [David Galley](#), [Nicolas de Bellefonds](#), and [Sylvain Duranton](#)

**READING TIME: 12 MIN**

*Companies struggle to use AI to improve the way their businesses operate, a challenge that can be mitigated by following five rules.*

There are two faces to the machine-learning component of [artificial intelligence](#) algorithms. There are the gee-whiz competitions in which data grand masters, given fixed rules, design algorithms that set in motion a race to solve aspects of complex problems like HIV or traffic forecasting. Open-source libraries have been created to house the winning solutions, allowing parts of them to be used again in future contests. Because speed is crucial, automated data manipulation and pattern identification are frequently employed.

The other face of machine learning involves the more real-world challenge of how to fix an obsolescent business process or address a business challenge using algorithms and programs. The business environment is inconsistent and full of uncertainty. It changes frequently and lacks the relatively prescribed borders of a machine-learning competition. Moreover, business issues are diverse, ranging from predicting rare events, such as equipment breakdowns, to making ultragranular decisions about personalized products for individual customers.

Not surprisingly then, using machine learning in a business setting has by and large been discouraging. A 2020 study conducted by BCG GAMMA, the BCG Henderson Institute, and MIT Sloan Management Review found that only one out of ten companies has enjoyed significant benefits from AI. Try as they may, a company's data scientists cannot often repeat the successes of machine-learning competitions in day-to-day operations in which decision making must be forward-looking and reliable.

To address the divide between the promise of machine learning and the disappointing results in real-time decision making and other business applications, BCG has distilled five rules. They range from determining which machine-learning applications to focus on to ensuring that an application is sufficiently robust.

## Targeting Business Value

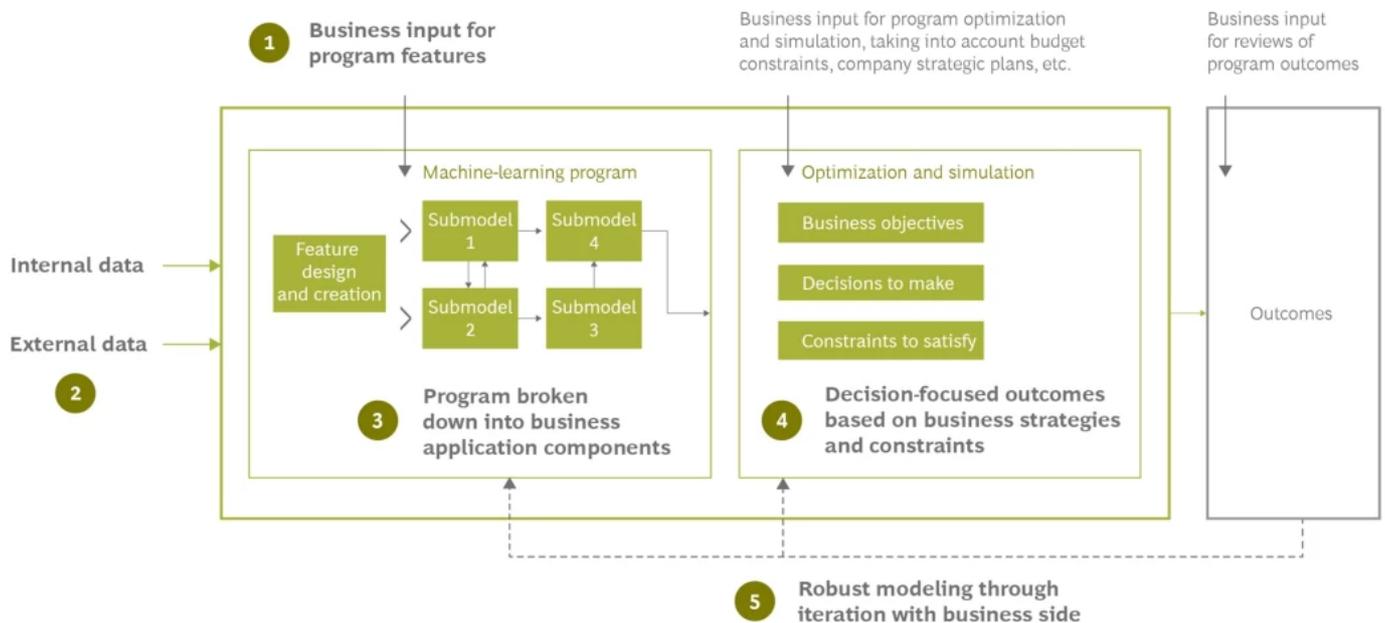
At many companies, machine learning and AI—the umbrella term for algorithms running on a data and technology infrastructure of which machine learning is the self-learning segment—are viewed as technical problems. But for AI to have real value in an organization, both its potential and its limitations must be understood by the business side of the company as well, so that expectations and outcomes can be realistically established.

Enlarging on that point, global AI expert Kai-Fu Lee, chairman of Sinovation Ventures, argued recently that creative vertical implementations of AI—that is, industry-specific, business-focused applications customized for, say, auto manufacturing or banking—will drive companies' competitive gains and digital transformation in the coming years. By contrast, horizontal, cross-industry AI solutions, such as adding off-the-shelf image recognition to medical equipment or cameras, will increasingly be less valuable.

For this reason, companies should assess the tangible business value of AI through the lens of a 10-20-70 formula. That is, 10% of the effort will lie in building an adequate machine-learning model—an algorithmic set of rules or instructions to help the system learn on its own; 20% will involve high-quality data and technology implementation and innovation; and 70% will focus on developing new business processes or transforming the way business functions operate.

To adopt this formula and get the biggest performance boosts from AI, companies must take a series of tactical and strategic steps keeping five rules in mind. (See the exhibit.)

# Five Rules for Designing a Successful Machine-Learning Program



Source: BCG analysis.

## Rule 1. The AI Application Must Match the Desired Business-Related Outcomes

Even within a single category, there are numerous options for the types of raw data to feed a machine-learning algorithm. For instance, if the system is intended to create a sales forecast, data about product prices can be input using absolute currency value; however, discount versus reference prices can often be more relevant, particularly when demand elasticity is a priority. On the other hand, if cannibalization is a concern, relative price differences among similar products sold by other companies may matter more. Many of these data transformations can now be obtained via open-source libraries, such as AutoML. But often companies choose that option instead of devoting the time needed upfront to fully understand the input data required by the specific business problem.

Collaboration with the people most responsible for dealing with the business issue at hand is essential to building the right features into the system. For example, a retailer that sells luxury fashion products (primarily women's handbags and clothing) recently wanted to generate a list of high-potential customers. But since relatively few purchases of luxury products are made, an AI system using traditional classification information—demographic profiles or purchase history by category, price, or shopping location—would not likely produce helpful results.

After lengthy discussions with sales and marketing experts, the retailer developed an AI model using such data as people who visited a store just after the release of a new product or people who browsed the retailer's website and afterwards visited a store to make a purchase. Adding these wrinkles to the data input into the AI system resulted in a list of potential customers that was 20 percentage points more effective at predicting actual shopping behavior than a similar list of customers identified using standard approaches. Additionally, the people singled out by the machine-learning system spent 80% more money on the retailer's products than those in the other group.

## **Rule 2. Use External Data to Amplify the Business Impact**

Many companies mistakenly assume that data generated internally is sufficient to drive machine-learning models. This is akin to thinking that a company can operate in isolation from the world around it—which is not an option in today's dynamic landscape. Of course, collecting exhaustive, accurate, and relevant external data for machine learning requires a lot of effort, but it also provides a clear competitive advantage because it allows for decision making based on dynamic changes in the environment.

For instance, in an engagement with a major airline, BCG was asked to build a forecasting tool to predict consumer demand across different routes as travel picks up post-pandemic. Rather than relying solely on the carrier's reams of historical data, we looked at external information as well, including travel search patterns across demographics, transaction data across airlines, actual and likely travel restrictions in specific regions, and consumer economic activity by region. With this data, machine-learning forecasts of the number of passengers per route were 20% more accurate than those provided by older predictive systems, resulting in improved flight and crew planning and some \$60 million in cost savings in 2020. The airline's AI system drove a decision support tool that helped optimize routes according to real-time demand by airport, route, destination, day, and time and helped airline personnel react to scheduling changes and match flight and crew planning to shifting demand.



**For AI to have real value, both its potential and its limitations must be understood by the business side, so that expectations and outcomes can be realistically established.**

## **Rule 3: Conquer Complexity by Breaking Down the AI Model into Its Smallest Parts**

In the world of machine-learning competitions, complexity reigns. Data scientists assemble vast and impenetrable algorithmic systems to show off how smart a machine can be. The problem is that these systems are relatively useless for dealing with the issues that a business might face. Rather than try to replicate the successes of these competitions, companies can put an AI model's agility to work by designing simpler and more interpretable submodels targeted at the business logic of relatively circumscribed problems.

To illustrate, a B2B distributor in the construction materials and equipment sector wanted to build an AI-based product recommender system for its sales teams serving small to midsize subcontractors. Typical recommender systems, such as those used by Netflix or Amazon, will analyze people's purchasing and browsing habits and produce a broad list of suggested items that might match an individual's interests. In these setups, the consumer hardly notices if a few of the recommendations are totally off base, and this will rarely affect whether he or she ends up making a purchase or continuing to engage with the website. But in a B2B setting, where a salesperson must make accurate and relevant face-to-face recommendations or possibly lose the customer, there is little room for multiple misguided suggestions.

With these constraints in mind, the B2B company adopted a three-step AI algorithm. The first one identifies the scope of the project—for instance, whether the initial order indicates that the client requires maintenance or a whole new installation. A second algorithm identifies which product categories are needed to complete the project. And a third algorithm identifies which SKUs best meet the project's specific requirements. Importantly, this hierarchy of recommendations naturally supports the sales pitch: first the specific project, then other building activities related to the project, and finally the individual products needed to execute those additional parts of the job. This machine-learning approach increased the distributor's sales by about 2%.

## **Rule 4: Machine Learning Should Help in Making Concrete Business Decisions**

Machine-learning systems are often judged by how accurate their view of the future is, but perfect predictions do not necessarily produce optimal decisions. Instead, performance should be measured against the incremental value that the company gains from decisions driven by the system. Generating a demand forecast without an associated decision support tool that adjusts supply to demand will not add maximum value.

This is where optimization algorithms enter the picture. An optimization engine uses the AI knowledge base to suggest decisions that best suit an objective, such as improving margins, limiting costs, or meeting specific lead times. Optimization algorithms explore all possible decisions and identify the ones that will have the most impact while still accommodating business constraints. They allow users to build simulation and planning tools that open a dialogue between the AI system and people on the company's business side; this, in turn, often leads to the adoption of advanced analytics that further expand the company's knowledge about its operations and the best ways to achieve optimal performance.

A consumer products company recently installed an AI system to assess its promotional programs in stores. At its most basic level, the machine-learning component provides forecasts of the volume sales improvement per product generated by different discounting and advertising campaigns, taking into account time of year, extent of the promotional coverage, competitor activity, and product attractiveness. But these forecasts have little value without associated decisions about the types and scheduling of promotional plans per retailer and product that could potentially produce a maximum profit above the investment in the AI system.

Building a promo plan like this is cumbersome, first because of the huge number of possibilities and second because of the operational constraints it must satisfy regarding promotion intensity, interaction with other marketing activities, brand strategic plans, retailer contracts, and other legal issues. The optimization algorithms in the company's AI program can easily handle these computations, while business managers can focus on evaluating the cost and efficacy of various decisions about the promotional plans. With this system in gear, the consumer products company enjoyed a 10- to 20-point increase in promotional return on investment, depending on the location and size of the campaign.



**Collaboration with the people most responsible for dealing with the business issue at hand is essential to building the right features into the system.**

## **Rule 5: Avoid Machine-Learning Outcomes That Seem Accurate but May Not Prove Useful**

When leading a business, you act today to create tomorrow's growth. Likewise, when building a machine-learning system, it's a mistake to focus on the accuracy of its past outcomes, which relate to

prior conditions at the company and are not necessarily relevant to contexts and problems that may arise in the future.

A better approach is to assess the future credibility and usefulness of the system's conclusions by running sensitivity analyses against all types of possible conditions that could affect the company's short- and long-term performance. Even better, perhaps, would be to input a robust set of information in the development stage. Often AI results are faulty because the data is too thin and myopic, culled, for instance, solely from manufacturing processes or marketing campaigns that have remained essentially unchanged for at least a few years. Algorithms assembled from this data will likely omit knowledge about possible nuances and variations that may govern the business process—and only reflect a relatively inert past. To mitigate this problem, algorithms should be combined from the outset with a well-rounded representation of the company's business processes, details that can only be obtained from people on the business side who manage day-to-day operations.

The value of this approach is evident from an AI suite that we installed for a client operating large-scale industrial facilities. The engine integrates all the empirical rules from the factory floor so intuitively that the system is now used to test hypotheses for potential changes in those rules. When adopted, the changes further enrich the system's own training database, subsequently allowing for even more refined and insightful recommendations. With this robust, self-learning AI suite, the company has reached record performance levels, including a 15% gain in output in an industry that celebrates when productivity is up just 3%.

---

The five rules we offer require significant collaboration between data scientists and so-called domain experts on the business side of the company. For that reason, to truly harness the power of machine learning, companies should build multidisciplinary teams that meld technology and operations so that the organization can enjoy the best of two worlds: data scientists with strong business-side knowledge and business experts who understand what they can achieve with machine learning. These teams will be able to make optimal use of AI for business gain and, in the process, allow the business to take advantage of the less geeky but more workaday face of machine learning—the one that can ultimately have the biggest impact on business performance.

# Authors



**David Galley**

**MANAGING DIRECTOR AND PARTNER, BCG X**

Paris

---



**Nicolas de Bellefonds**

**MANAGING DIRECTOR & SENIOR PARTNER**

Paris

---



**Sylvain Durantou**

**MANAGING DIRECTOR & SENIOR PARTNER; GLOBAL LEADER, BCG X**

Paris

---

## ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

© Boston Consulting Group 2023. All rights reserved.

For information or permission to reprint, please contact BCG at [permissions@bcg.com](mailto:permissions@bcg.com). To find the latest BCG content and register to receive e-alerts on this topic or others, please visit [bcg.com](https://bcg.com). Follow Boston Consulting Group on [Facebook](#) and [X \(formerly Twitter\)](#).