



## The Cure for AI Fever

The companies that most successfully use artificial intelligence know that a simple tool is not the fix; they trust the AI process.

**By Cesar Brea and Sanjin Bicanic**

**Cesar Brea is a partner with Bain & Company in the Boston office. Sanjin Bicanic is a Bain principal in the San Francisco office; he is the practice manager of the firm's Global Advanced Analytics practice.**

---

## Executive Summary

- ▶ General artificial intelligence that can solve many broad problems most likely is still years away.
- ▶ Instead of obsessively following the latest AI fads, leading companies focus on the AI process.
- ▶ The AI process starts with a business problem and its context, then engineers the data to develop a model to address the issue, deploys it, refines it based on results and ensures its adoption.

---

Imagine this scenario: An executive's primary competitor has just hired a promising start-up, a software firm specializing in strong artificial intelligence that covers a broad range of applications, not just something specific such as translation or image recognition.

Trying to learn more about this, she looks up the start-up founder's TEDx talk but finds it impenetrable. The firm's website boasts claims for the tool bordering on science fiction and has pictures of young men and women and their dogs in a WeWork space against a backdrop of whiteboards all covered in formulas.

Natural questions arise. How solid is this technology? What is it for? How does what the people with the dogs say they do intersect with what their technology seems to be for? There is no obvious overlap between the start-up's product and any identifiable application for the executive's company.

The executive faces a decision.

Many in this circumstance cave to fear. They scramble to find somebody who knows another strong AI provider to cobble together an alternative to the version offered by the people with the dogs. In the first quarter of 2018, nearly 20% of corporate earnings calls included some mention of artificial intelligence or machine learning, and her board would have expected her to keep up with competitors.

This kind of thing is all too common, including the caving. The better choice is to treat artificial intelligence as a process, not as a tool.

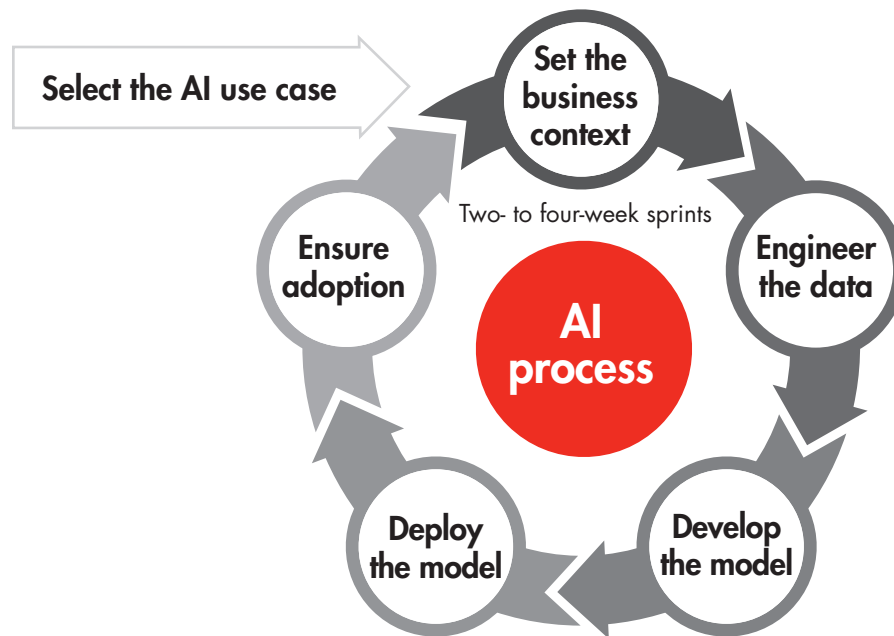
To understand the process, think about how software gets built; an AI solution advances similarly in short, tightly managed release cycles (see *Figure 1*). The goal is to get a little bit better with each cycle—to become more efficient and produce more powerful predictions.

## How the five-step AI process cycle works

**Step No. 1: Set the business context.** The AI process starts by choosing a quantifiable business objective. Picking the right problem is the single most important part of this entire process. The second most important part is setting a baseline, which then becomes a target to beat. This framing should focus on the impact on the business objective; it should not be about how well a company's AI efforts and algorithms are keeping up with the competition's AI efforts and algorithms. Translate the business objective into a machine learning objective—for example, retaining valuable but dissatisfied customers.

Figure 1

## ■ The five-step, Agile AI process



Source: Bain & Company

**Step No. 2: Engineer the data.** Evaluate the data being used today and the data that might be used tomorrow to see what is possible. Pros know that better data beats a better algorithm. Using first-party data—that is, data the company generates itself through interactions with its customers—is usually better than spending time and money getting new data from third parties.

**Step No. 3: Develop the model.** An effective AI process advances in Agile sprints. Too many projects shoot for the moon and fail, but Agile sprints, step by step, help build a ladder to the moon. At the end of each sprint, compare the performance of the new model with the baseline model that the team is trying to beat. This ties the analytics—a better targeting algorithm, for example—to the business’s ability to act. Keeping the process short and tight forces managers to be practical about what they want to build themselves and what they want to bring in from the outside and augment.

**Step No. 4: Deploy the model.** Each cycle creates some value, so put that to work. Simultaneously, test the results with real users, and use the results to inform how the business process delivers value. During this step, the requirements for a variety of capability enhancements are determined—including changes to governance (specifically, how challenges are selected); changes to data platforms; changes to infrastructure and processes for testing as well as operational execution at scale; changes to the talent needed and the organization of that talent. It will not all be figured out during the first cycle, but over time, it will be. It is valuable to figure out what is needed before managers start writing seven-figure checks to vendors. Unfortunately, this is precisely the opposite of how most people approach artificial intelligence and a big reason why their AI efforts fail.

**Step No. 5: Ensure adoption.** A great AI model that ends up sitting on a virtual shelf and unused by the front line is just Shelfware 2.0. Test models with the actual end user to understand what works and what does not. Identify both model end users and those who will inherit model maintenance. Paint the picture of the metaphorical beach that AI will enable your organization to reach, and propagate that image throughout the sponsorship spine—both the leaders who sponsor change and those who must implement it in the target organizations—to ensure adoption.

While the competition is distracted by AI buzzwords, leaders use this five-step AI process cycle to deliver real business impact. And while others talk about their point of arrival, this group is lapping them—both in what they have already accomplished and the business performance they have to show for it.

A large mobile carrier recently faced an existential threat when a major competitor started a price war. Customer churn spiked. At first, the carrier estimated it had 10 weeks to respond. Then 10 turned to 5 when another big player followed suit. Using a week-by-week release cycle, the team focused on slightly better targeting models and slightly smarter offers during each cycle and rolled those out in waves across their customer base, with controls to ensure that they drew the right conclusions from the results.

Their initial churn model could distinguish a true cherner from a false positive with a rate of accuracy only slightly better than a coin flip, but bit by bit, it improved. They started with open-source tools and classification models fed with existing data and moved on to more sophisticated techniques. Even more important, they brought in network performance variables gathered through a method called deep packet inspection, a technique used by telcos to gather user behavior data at a very granular level to complement customer account records. By combining better customer targeting and better creative and pricing tactics, the company reduced churn by 60%, returning it to historic levels compared with control groups.

This analytical work did not really require AI. The team considered but held off on deploying neural nets—in part, because it was important for the marketers using the model to have confidence in it. They established that it will be easier to bring users along by starting simply and transparently. Over time, as momentum and confidence continue to build, it will be easier to move toward such opaque deep learning solutions that may improve the performance even more.

---

The goal is to get a little bit better with each cycle—to become more efficient and produce more powerful predictions.

---

---

## The Cure for AI Fever

With analytics, people and operating models need to change along with data and algorithms. A major consumer products firm grappled with this when trying to address excessive returns. A relatively simple model produced its demand forecast, and while it made good predictions across the full country, its recommendations for individual items in a specific store were often poor. To compensate, multiple parties would pad the recommended order, including the tens of thousands of field reps who had final say over the orders for their routes. That ended up creating a lot of overordering and high rates of waste.

To tackle this, the company built a database containing several years' worth of order and return histories, and it used open-source libraries to build early forecast models.


In tandem, the company developed a new interface through which field reps could consume the forecast and begin to trust it, and it engineered a model pipeline that could efficiently predict demand for millions of combinations of items and stores across the country. In closely managed pilot environments, the new tools were tested and refined, confirming the expected benefit.

This new forecasting approach has reduced returns and is expected to result in a nine-figure annual payback. The company is building an internal organization to continuously improve it, and artificial intelligence may be deployed next year if model performance in this part of the business is still the bottleneck. Alternatively, the team may focus elsewhere, perhaps deeper in the supply chain.

The experiences of both these companies, the mobile carrier and the consumer products manufacturer, are consistent with what pros operating at the bleeding edge of machine learning know very well: There is often plenty of room for improvement before applying the AI arsenal.

Google's Martin Zinkevich, in his widely read paper "Rules of Machine Learning: Best Practices for ML Engineering," called this rule No. 1.

**Rule No. 1: Don't be afraid to launch a product without machine learning. Machine learning is cool, but it requires data. ... If you think that machine learning will give you a 100% boost, then a heuristic will get you 50% of the way there.**

So the next time a competitor collaborates with a strong or general AI software company, hold on to your wallet. Most experts agree that general artificial intelligence that can solve many broad problems is still years away. Instead of falling prey to AI tool fever, do what the leaders do: Trust the AI process. 

## *Shared Ambition, True Results*

**Bain & Company is the management consulting firm that the world's business leaders come to when they want results.**

Bain advises clients on strategy, operations, technology, organization, private equity and mergers and acquisitions. We develop practical, customized insights that clients act on and transfer skills that make change stick. Founded in 1973, Bain has 56 offices in 36 countries, and our deep expertise and client roster cross every industry and economic sector. Our clients have outperformed the stock market 4 to 1.

### **What sets us apart**

We believe a consulting firm should be more than an adviser. So we put ourselves in our clients' shoes, selling outcomes, not projects. We align our incentives with our clients' by linking our fees to their results and collaborate to unlock the full potential of their business. Our Results Delivery® process builds our clients' capabilities, and our True North values mean we do the right thing for our clients, people and communities—always.



## Key contacts in Bain's Advanced Analytics practice

### Americas

**Chris Brahm** in San Francisco (*chris.brahm@bain.com*)

**Cesar Brea** in Boston (*cesar.brea@bain.com*)

### Asia-Pacific

**James Anderson** in Sydney (*james.anderson@bain.com*)

### Europe, Middle East and Africa

**Florian Mueller** in Munich (*florian.mueller@bain.com*)

For more information, visit [www.bain.com](http://www.bain.com)