


ANALYTICS

How to Make Your Company Machine Learning Ready

by James Hodson

NOVEMBER 07, 2016



In recent years, there has been a staggering surge in interest in intelligent systems as applied to everything from customer support to curing cancer. Simply sprinkling the term “AI” into startup pitch decks seems to increase the likelihood of getting access to funding. The media

continuously reports that AI is going to steal our jobs, and the U.S. government seems as worried about the prospect of super-intelligent killer robots as it is about addressing the highest wealth disparity in the country's history. Comparatively, there has been very little discussion of what artificial intelligence *is*, and where we should expect it to actually affect business.

When people talk about AI, machine learning, automation, big data, cognitive computing, or deep learning, they're talking about the ability of machines to learn to fulfill objectives based on data and reasoning. This is tremendously important, and is already changing business in practically every industry. In spite of all the bold claims, there remain several core problems at the heart of Artificial Intelligence where little progress has been made (including learning by analogy, and natural language understanding). Machine learning isn't magic, and the truth is we have neither the data nor the understanding necessary to build machines that make routine decisions as well as human beings.

That may come as a disappointment to some, and potentially disrupt some very expensive marketing campaigns. But the likelihood of self-directed, super-intelligent computational agents emerging in the foreseeable future is extremely low – so keep it out of the yearly business plan for now. Having said that, an enormous amount can already be achieved with the machinery we have today. And that's where forward-thinking managers should be focusing.

Over the next five to 10 years, the biggest business gains will likely stem from getting the right information to the right people at the right time. Building upon the business intelligence revolution of the past years, machine learning will turbocharge finding patterns and automate value extraction in many areas. Data will increasingly drive a real-time

economy, where resources are marshaled more efficiently, and the production of goods and services becomes on-demand, with lower failure rates and much better predictability. This will mean different things for different industries.

In services, we will not only get better at forecasting demand, but will learn to provide the right product on a hyper-individualized basis (the Netflix approach).

In retail we will see more sophisticated supply chains, a deeper understanding of consumer preferences, and the ability to customize products and purchase experiences both on- and off-line. Retailers will focus on trend creation and preference formation/brand building.

In manufacturing there will be an evolution towards real-time complete system monitoring, an area known as “anomaly detection.” The components will become increasingly connected, allowing for streams of real-time data that machine learning algorithms can use to reveal problems before they happen, optimize the lifetime of components, and reduce the need for human interventions.

In agriculture, data will be used to decide which crops to grow, in what quantities, in what locations, and will render the growing process more efficient year after year. This will create more efficient supply chains, better food, and more sustainable growth with fewer resources.

In short, AI may be a ways off, but machine learning already offers huge potential. So how can managers incorporate it into daily decision-making and longer-term planning? How can a company become *ML-ready*?

First, catalogue your business processes. Look for procedures and decisions that are made frequently and consistently, like approving or denying a loan application. Make sure you're collecting as much data as is feasible about how the decision was made, along with any data used to make it. And make sure to collect the decision itself. In the hypothetical loan example, you want to record whether the loan was approved; the data used to make that decision; and any other information about the circumstances behind the decision. (Who made it? At what time of day? How confident did they feel in the decision?) This is the kind of data that can be used to fuel machine learning in the future.

Second, focus on simple problems. Automation and machine learning will work well where the problem is well defined and well understood, and where the available data exemplifies the information necessary to make a decision. A good problem for machine learning is identifying a fraudulent transaction. The question "What makes customers feel happy?" is vaguer, more challenging, and not the place to start.

Third, don't use machine learning where standard business logic will suffice. Machine learning is useful when the set of rules is unclear, or follows complex, non-linear patterns. If you want transparency and reliability, go for the simplest possible approach that meets your performance criteria.

Fourth, if a process is complicated, use machine learning to create decision support systems. If the objective is too unclear to define with respect to the data, try to create intermediate results to help your teams be more effective. You can think of machine learning as part of the hierarchical decision-making path, and it will engender a better understanding of the problem in future.

The point is that there is a lot that can be done without needing to dig very deep. The majority of your workforce will continue to have a job, and you can help them to be more productive, and work on more interesting and demanding (read: *more valuable*) tasks by digitizing more of the mechanical parts of your business. For now, artificial intelligence cannot turn a business's performance from bad to good, but it *can* make some aspects of a *good* business *great*.

If you run out of low hanging fruit – though I'll wager you won't – it may be time to consider building a team to attack more complex problems with machine learning. Be patient, as this investment will not pay off immediately. If you do decide to create such a team, be open, engage with the research community, and you will be contributing to building tomorrow's economy.

James Hodson is CEO of the AI for Good Foundation, Chairman of the Financial Data Science Association, and an early-stage researcher in Artificial Intelligence and Finance at the Jozef Stefan Institute AI Lab.

This article is about ANALYTICS

 FOLLOW THIS TOPIC

Related Topics: TECHNOLOGY

Comments

Leave a Comment


POST

4 COMMENTS

david mata a year ago

Hi, this is a great post. I am currently working in a medium-big hotel company (we have 5000 people working in our biggest complex and we have more people working in other complexes) How can I use machine learning to help improve my company? We have registered everything because all guests pay with their room number. For example, I can think of how often a guest comes and in what dates so I can predict when they are going to travel and call them to sell a vacation in the hotel.

REPLY



▼ [JOIN THE CONVERSATION](#)

POSTING GUIDELINES

We hope the conversations that take place on HBR.org will be energetic, constructive, and thought-provoking. To comment, readers must sign in or register. And to ensure the quality of the discussion, our moderating team will review all comments and may edit them for clarity, length, and relevance. Comments that are overly promotional, mean-spirited, or off-topic may be deleted per the moderators' judgment. All postings become the property of Harvard Business Publishing.