
Where AI Can Help Your Business (and Where It Can't)



By **Phanish Puranam** , INSEAD

Your firm produces data, so surely it can benefit from applying AI, right? Wrong. Here are five questions to ask yourself about whether a business problem is “AI-solvable”.

Machine learning, the latest incarnation of artificial intelligence (AI), works by detecting complex patterns in past data and using them to predict future data. Since almost all business decisions ultimately rely on predictions (about profits, employee performance, costs, regulation etc.) it would seem obvious that machine learning (ML) could be useful whenever “big” data are available to support business decisions. But that isn’t quite right.

The reality in most organisations is that data may be captured but it is stored haphazardly. Its quality is uneven, and integrating it is problematic as it sits in disparate locations and jurisdictions. But even when data is cleaned up and stored properly, it’s not always appropriate for the questions or decisions that management has in mind. So, how do you know whether applying predictive analytics through AI techniques to a particular business problem is worthwhile? While every organisation and context is different,

here are five general principles that should be useful in answering the question.

1. Data must be representative

If you want to make predictions in the context of a business process, make sure you have data that represents the process. Let's look at an example. You have data about which of your employees are stars and want to use this to better predict whom to hire. This may sound like a reasonable objective, but it's not. Hiring involves screening, say, 1000 applications and recruiting 100. You then look at the data on the 100 to identify the "stars". But these are stars conditional on having made it into the 100. The data lacks information about who would have been stars among the 1000 applications you received. What you ideally need is data on the 900 you didn't hire (maybe tracking their careers on LinkedIn?), who nonetheless may have gone on to become stars elsewhere. [\[1\]](#)

Remedy: For any decision to which you want to apply AI techniques, aim to have data that truly represent the decision-making process, including the alternatives that were considered as well as those selected (also see [previous blog post](#) on this theme). This applies as much to customer targeting, supplier selection or site selection as it does to hiring. In each case, using data only on the outcomes of those selected is a misleading basis for deciding what or whom to select.

2. Representative data does not necessarily mean there is an authentic pattern

If you have lots of representative data, this could still just be a lot of noise. You can't assume that throwing tons of data at fancy algorithms will necessarily produce actionable insights. I suspect most companies who use ML are happier talking, if at all, about the (few) projects that succeed, rather than the (perhaps) many that fail. It's called data mining for a reason – not every seam in the cave will contain gold. For example, one might assume that in-company suggestion box schemes are fertile grounds for ML-based predictive analytics, containing valuable correlations between people's attributes and their success at innovation. But I have examined such data for companies where, despite our best efforts, we sometimes found little of interest. Success or failure in those contexts seemed to have little to do with what we had data on and may have just been dependent on random unobservable factors such as the contributor's (or evaluator's) mood at the

relevant times.

Remedy: Prioritise projects in domains where the difference in decision quality between experts and novices is significant. For instance, if you have experienced senior managers who seem to have a knack for pricing strategy, supplier selection or hiring, this shows a pattern to be detected in these activities in your organisation, because human decision makers also work ultimately by pattern detection. It's just that their algorithms are sitting between their ears and their datasets are called "life experience".

Another solution is to outsource pattern detection to those who don't have to meet a profit objective, i.e. academics. Researchers are paid to search for interesting patterns in data and explain them, and they can afford to conduct many more unsuccessful searches than a corporation with a profit objective typically can. As we have found at INSEAD, it's often good training for budding researchers (PhD students) to play with some data anyway, so this is a win-win for both sides.

3. Patterns must be stable

Every ML-based prediction algorithm in existence works by assuming that the world tomorrow will be similar to the world yesterday. To the extent that ML works, it can often surprise us by showing that there is a stable pattern even where we as humans don't see it. But, again, this is hardly a given. If the ways things work in a certain domain seem to be changing a lot, then even if you have a lot of data, there is no guarantee you can detect a useful pattern. Suppose we were looking at data on foreign direct investment strategies by multinational companies from 2000 to 2015. This might be an invaluable resource for historians and academics, but is unlikely to provide accurate predictions on what MNCs do in 2019 when they invest around the world. Too much has changed since then and is still changing.

Remedy: Prioritise projects in domains that seem fairly stable (or change in ways that themselves have a pattern). If you are not sure if these properties apply, rely on others to do the exploration for you (see above). Alternately, slice the data into smaller periods within which there might be more stability: re-train your models frequently.

4. The pattern should not perpetuate a socially unacceptable process

Amazon famously **stopped using machine learning for hiring** because the algorithm was making gender-biased hiring recommendations. But the bias possibly lay in the process the algorithm was trained to emulate, not the data or the algorithm itself. Perhaps the algorithm accurately reflected the fact that hiring and evaluation practices in the past had been prone to gender discrimination. The social bias in the process (hiring and promotion) made it dangerous to use ML solutions on this data to predict whom to hire, because these solutions can and will replicate the bias within the data-set they are trained on, in this case male-dominated hiring practices. This problem would not go away even if we had representative data with a stable pattern – if the stable pattern is one of industry-wide gender discrimination.

Remedy: It's doubtful whether Amazon's execs would have even found out about their "gender problem" without the algorithm doggedly demonstrating it to them. No HR manager will own up to discriminating against women, but when quizzed properly, algorithms are more likely to "confess". Therefore, the lesson to draw from stories like Amazon's is not to stop using algorithms. Instead, when applying ML to HR data, a necessary step is the consideration of possible social biases in the processes that produced the data. Sophisticated ML users have developed a suite of measures to check for fairness in predictive analytics, and this should become an integral part of any AI applications that use people data.

5. All predictions based on patterns are not equally valuable

The success of an ML algorithm is measured in prediction accuracy. One obvious application has been to build systems that are at least as accurate as humans, but which are substantially cheaper. This is straightforward automation and the principles listed above all apply.

A more ambitious use of ML would be to create competitive advantage by generating *more accurate* predictions than humans. But even when that is possible, it's not always worth doing. Not all differences *make* a difference, because the value of accuracy varies by context. For example, consumers may be unlikely to pay more for a weather prediction engine that was 2 percent more accurate than competitors in forecasting the chances of rain. But if you could build a search engine whose results were 2 percent better than Google's, you would swiftly corner the market.

Remedy: Think in terms of the value-accuracy curve and its steepness. Ideally you want to apply ML when a marginal increase in accuracy by using

it relative to human decision making produces disproportionately large benefits – a case of “increasing returns” in accuracy. The key strategic challenge is to identify the key parts of your business where the returns in increased accuracy are steep.

In general, as with all business improvement projects, you should look for machine learning projects with high impact and high feasibility to prioritise. But the principles above add texture to this intuition.

[1] This is the well-known problem of “sample selection” in statistics. For instance see Heckman, J. (1979). Sample Selection Bias as a Specification Error. *Econometrica*. **47**(1): 153–61.

Find article at

<https://knowledge.insead.edu/operations/where-ai-can-help-your-business-and-where-it-cant>

About the author(s)

Phanish Puranam is the Roland Berger Chaired Professor of Strategy and Organisation Design at INSEAD.