

Artificial Intelligence: Intelligent or just artificial? Submitted by Michael Durrie

If your company is like most, you're exploring how to integrate AI into your business processes – or you're already implementing it. And for good reason: AI is revolutionizing eCommerce, supply chains and logistics, diagnostics, customer service and risk management, to name just a few examples. AI is advancing at a dizzying pace, as is its adoption. A 2019 [survey](#) by the research institute O'Reilly found that 85% of respondents were planning AI implementation or already using, with over 50% identifying as mature adopters.

So should all companies embrace AI on all levels possible or risk being left behind? Not so fast.

A closer look reveals that AI implementation has been a money-losing proposition for many players. It can also increase risk of error and misleading results to practically unmanageable levels. Success in AI adoption belongs to those who fully understand what the technologies can do in the context of their organizations, and which areas of implementation to steer clear of (at least for the foreseeable future).

Enter Chat GPT

Recent advances in AI-based coding and copywriting have triggered lively discussions. But how much of a game-changer is open AI beyond the make and mirrors? What should we make of the move by industry leaders including Elon Musk to slow down AI development?

First of all, let's all slow down and put AI into perspective. It's a tool that – like any tool – can be used for good or evil. The technology is not likely to save the world or take over the world any time soon. That said, AI holds great promise in terms of enhancing human performance. Instead of struggling to create a first draft of a paper or article, which we then proceed to revise beyond recognition, we can now use AI to kick-start our efforts and leapfrog over the initial draft phase. Rather than rendering human authors obsolete, this raises the bar. We now have to be better than the machine and connect the dots in a way AI can't.

Back to basics

to explore AI with the aim of identifying its potential benefits, we need to begin by unpacking exactly what AI is. According to the [Encyclopedia Britannica](#), AI is “the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings.” Think back to some of the first amazing things we heard about AI in past decades: in the 1950s, computer scientists used AI software to train a computer to play chess with remarkable skill. This feat was followed in the 60s by a computerized “psychotherapist” named Eliza, which used the reflective questioning psychology method. Users were convinced that the software had feelings.

And of course, this was just the beginning: Fast-forward to 2016 and you'll find examples of [AI-generated music](#) that mimics the composition style of Johann Sebastian Bach so well that even most Bach experts can't tell the difference. But as great an accomplishment as this is, it's telling to note that the software developers chose to test their algorithms on Bach, a composer who used a highly sophisticated *yet consistent* counterpoint methodology in his compositions. Suffice it to say, we're not going to hear true-to-life imitations of [Giacomo Puccini's](#) style anytime soon.

So what does this mean in terms of implementing AI in your business?

Today's most successful and reliable implementations of AI and ML include translation, voice and handwriting recognition and predictions in data-rich contexts like weather forecasting or some areas of radiology. Note that these are disciplines in which millions of labelled data examples are available and where causality is practically unambiguous.

In other words, AI and ML can deliver actionable results to your business provided you can clearly define what you are looking for and how to interpret those results once you have them. That means you need a good understanding of the overall system and the long-term causal relations between datasets. If these criteria are met, an ML model can adapt and improve to remain useful in the future.

The technology's performance can go dangerously sideways if these criteria are not met, e.g. when patterns are inconsistent or correlations signal false causalities (spurious correlation).

In the case of [melanoma diagnosis](#) or detection of diabetes-related disease of the retina, ML has proven highly effective. The patterns are consistent. In contrast, ML is unreliable in detection of the retinal disease neurofibromatosis type 2 because the pattern is inconsistent. Similarly, a recent [study](#) showed that radiologists outperform the technology in breast cancer screening. (It should be noted that experienced medical experts have achieved significantly improved diagnostic accuracy by drawing on deep learning AI/ML technology to augment their expertise.)

In one example of spurious correlation, the AI/ML software failed to distinguish a wolf from husky in photos. Researchers wondered what patterns were confusing the algorithm – snout? Markings? Paws? Only after digging deeply into the math did researchers realize that it was the snow in the background that was confusing the ML. The technology simply couldn't tell the difference between the two animals, so it guessed based on a correlating feature of both photos.

This last example demonstrates how dangerously stupid AI can be. In practical use, the technology is a black box. Unlike researchers testing its capabilities in identifying different canine species or its effectiveness in medical diagnostics, business decision-makers are more likely to take the AI findings at face value – after all, efficiency is the whole point. In terms of risk

assessment, HR-related decisions or market analyses, such correlation errors can have extremely costly fallout – not to mention potential reputational damage.

AI hazard avoidance – Mind the gap

To bridge the gap between what AI can and can't do as well as to identify what applications are simply too unreliable and risky, an AI hazard avoidance assessment should be carried out before committing to any AI implementation project. The process, designed to evaluate a proposed or existing AI project against a set of criteria, helps you avoid wasting time and money or – worse yet – making detrimental decisions based on flawed AI-generated findings.

To be clear, an AI application for business can deliver value if – and only if – the description of the system is encapsulated (free of incalculable external variables) and its purpose well defined and understood. Systems like stock prices, fraud or user preferences in music and films may display certain patterns and trends, but external factors – not least, human behavior – can drastically alter predicted trajectories.

Ideally, a cross-disciplinary team of stakeholders who own and know the dataset(s) should be assemble. This makes it possible to gather a good understanding of the system and its behavior. Together, you need to ask:

- How well have we defined the real problem we want to solve?
- Are we asking the right questions?
- What findings do we need the solution to deliver so that we can confidently answer those questions?

This should be followed by a product-gap analysis to build on the findings and flesh out the benefits that can be realistically expected from the project based on the proven track record of data science. With these pieces in place, the AI/ML solution itself can be developed at an accelerated pace and with a high level of confidence.

Stay tuned: Moore to come

Considering the dynamics of Moore's Law, we will see AI vastly improve and overcome many of its shortcomings within the not-too-distant future. But for those wishing to profit from its capabilities in whatever way possible and feasible now – which is by any measure a good idea – it's important to see through the hype and proceed systematically. Don't let your first step into the brave new world of AI be your last.