


Why AI can't solve everything

May 25 2018, by Vyacheslav Polonski



Credit: artificial-intelligence-503593_1920/flickr

The hysteria about the future of artificial intelligence (AI) is everywhere. There seems to be no shortage of sensationalist news about how AI could [cure diseases](#), [accelerate human innovation](#) and [improve human creativity](#). Just looking at the media headlines, you might think

that we are already living in a future where AI has infiltrated every aspect of society.

While it is undeniable that AI has opened up a [wealth of promising opportunities](#), it has also led to the emergence of a mindset that can be best described as "AI solutionism". This is the philosophy that, given enough data, machine learning algorithms can [solve all of humanity's problems](#).

But there's a big problem with this idea. Instead of supporting AI progress, it actually jeopardises the value of machine intelligence by disregarding important AI safety principles and setting unrealistic expectations about what AI can really do for humanity.

AI solutionism

In only a few years, AI solutionism has made its way from the technology evangelists' mouths in Silicon Valley to the minds of government officials and policymakers around the world. The pendulum has swung from the dystopian notion that [AI will destroy humanity](#) to the utopian belief that our [algorithmic saviour](#) is here.

We are now seeing governments pledge support to national AI initiatives and compete in a technological and rhetorical [arms race](#) to dominate the burgeoning machine learning sector. For example, the UK government has [vowed to invest £300m in AI research](#) to position itself as a leader in the field. Enamoured with the transformative potential of AI, the French president Emmanuel Macron [committed to turn France](#) into a global AI hub. Meanwhile, the Chinese government is increasing its AI prowess with a national plan to create a [Chinese AI industry worth US\\$150 billion by 2030](#). AI solutionism is on the rise and it is here to stay.

Neural networks – easier said than done

While many political manifestos tout the transformative effects of the looming "[AI revolution](#)", they tend to understate the complexity around deploying advanced machine learning systems in the real world.

One of the most promising varieties of AI technologies are [neural networks](#). This form of machine learning is loosely modelled after the neuronal structure of the human brain but on a much smaller scale. Many AI-based products use neural networks to infer patterns and rules from large volumes of data. But what many politicians do not understand is that simply adding a neural network to a problem will not automatically mean that you'll find a solution. Similarly, adding a neural network to a democracy does not mean it will be instantaneously more inclusive, fair or personalised.

Challenging the data bureaucracy

AI systems need a lot of data to function, but the public sector typically does not have the appropriate [data infrastructure](#) to support advanced machine learning. Most of the data remains stored in offline archives. The few digitised sources of data that exist tend to be buried in bureaucracy. More often than not, data is spread across different government departments that each require special permissions to be accessed. Above all, the public sector typically lacks the human talent with the right technological capabilities to fully reap the [benefits of machine intelligence](#).

For these reasons, the sensationalism over AI has attracted many critics. Stuart Russell, a professor of computer science at Berkeley, has long advocated a more realistic approach that focuses on [simple everyday applications of AI](#) instead of the hypothetical takeover by super-

intelligent robots. Similarly, MIT's professor of robotics, [Rodney Brooks](#), writes that "almost all innovations in robotics and AI take far, far, longer to be really widely deployed than people in the field and outside the field imagine".

One of the many difficulties in deploying machine learning systems is that AI is extremely susceptible to [adversarial attacks](#). This means that a malicious AI can target another AI to force it to make wrong predictions or to behave in a certain way. Many [researchers](#) have warned against the rolling out of AI without appropriate [security standards and defence mechanisms](#). Still, AI security remains an often overlooked topic.

Machine learning is not magic

If we are to reap the benefits and minimise the potential harms of AI, we must start thinking about how machine learning can be meaningfully applied to specific areas of government, business and society. This means we need to have a discussion about [AI ethics](#) and the distrust that many people have [towards machine learning](#).

Most importantly, we need to be aware of the limitations of AI and where humans still need to take the lead. Instead of painting an unrealistic picture of the power of AI, it is important to take a step back and separate the [actual technological capabilities of AI from magic](#).

For a long time, [Facebook believed](#) that problems like the spread of misinformation and hate speech could be algorithmically identified and stopped. But under recent pressure from legislators, the company quickly pledged to replace its algorithms with an army of over [10,000 human reviewers](#).

The medical profession has also recognised that AI cannot be considered a solution for all problems. The [IBM Watson for Oncology](#) programme

was a piece of AI that was meant to help doctors treat cancer. Even though it was developed to deliver the best recommendations, human experts found it difficult to trust the machine. As a result, the AI programme [was abandoned](#) in most hospitals where it was trialled.

Similar problems arose in the legal domain when [algorithms were used in courts](#) in the US to sentence criminals. An algorithm calculated risk assessment scores and advised judges on the [sentencing](#). The system was found to amplify structural racial discrimination and was later abandoned.

These examples demonstrate that there is no AI solution for everything. Using AI simply for the sake of AI may not always be productive or useful. Not every problem is best addressed by applying machine intelligence to it. This is the crucial lesson for everyone aiming to boost investments in national AI programmes: all solutions come with a cost and not everything that can be automated should be.

This article was originally published on [The Conversation](#). Read the [original article](#).

Provided by The Conversation

Citation: Why AI can't solve everything (2018, May 25) retrieved 23 April 2024 from <https://phys.org/news/2018-05-ai.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--