Why artificial intelligence has disappointed in the pandemic

Nathan Benaich

In late February, an artificial intelligence algorithm created by MIT researchers made front-page news for discovering a powerful new antibiotic that could save millions of lives by treating drug-resistant diseases. Yet for many what was notable about the moment was that it wasn't surprising: we have become accustomed to a steady drumbeat of AI breakthroughs.

Just three weeks later, the World Health Organization declared Covid-19 a pandemic. Geeks around the world sprang into action. A company called C3.ai compiled data sets for AI researchers to crunch. The AI-first drug discovery company Recursion released a morphological image data set of Covid-19 infected cells as they react to thousands of drugs. DeepMind used its AlphaFold AI system to predict and publish structures associated with coronavirus. AI seemed on track to crush the new pathogen, just as it mastered the Chinese board game Go. And then — nothing.

At a time when we expect AI to solve most problems — since that has been the industry promise for years — why has it been so ineffective in the fight against the biggest global crisis in a generation?

There are inherent limitations to AI today. Current systems learn by finding patterns in data. In general, the more data you feed in, the smarter the AI. GPT-3, an algorithm by San Francisco-based OpenAI, can write coherent paragraphs on any topic with only a few word prompts. It learnt to do this by analysing almost half a trillion words. Yet such data also restricts AI.

To achieve superhuman performance, systems must typically be trained with high-quality inputs that model desired behaviours. That is easy to pull off in games such as Go — staged situations with clear parameters — but far harder in less predictable real-life scenarios. AI often flounders in the transition to real-world applications. In the February MIT study, AI only uncovered the new antibiotic with the help of human researchers, who ensured the computer avoided killing people by helping it distinguish between bad bacteria and species that are beneficial to healthy living cells.

Covid-19 is so new and complex that the data needed to train AI to combat it does not exist. Scientists do not fully understand the disease and while efforts to collect data have begun — the UK-based health science company ZOE has launched the Covid Symptom Study to track the health status of 4m people — high-quality clinical and biological data sets are few and far between.

AI has also failed to support government efforts — in France, the US and UK among others — to build effective contact-tracing systems because, again, the necessary raw ingredients are missing: a clear and well-defined task, and a large, high-quality and accurate data set. In the UK, the lack of systematic data gathering to track and trace Covid-19 cases has all but eliminated the possibility of AI-powered contact-tracing interventions in the near future.

But before we give up on AI, maybe we should give it more time. The history of AI is characterised by years of incremental improvements punctuated by bursts of innovation. Pioneers in the 1950s believed they could go a long way to build machines that mimic human intelligence in a single summer. But it wasn't until "big data" arrived in the 2000s that AI learnt how to build spam filters, drive autonomously and much more. In the mid-2010s, most researchers believed AI winning at Go was at least a decade away — and were stunned when it did so in 2017.

AI may, so far, be losing the race to solve Covid-19 to classical biology and chemistry. But that does not mean it isn't ready for primetime. As the MIT antibiotic research shows, AI can unlock medical mysteries under the right conditions, and it may yet play a role in the coronavirus crisis. If this is not the first pandemic to be solved by AI, then it will surely be the last to be solved without it.

The writer is a biologist, general partner of Air Street Capital and co-author of State of AI Report

AI seemed on track to crush the new pathogen, just as it mastered Go. And then — nothing