AI tools produce dazzling results—but do they really have 'intelligence'?

February 14 2024, by Paul Compton

Credit: AI-generated image

Sam Altman, chief executive of ChatGPT-maker OpenAI, is reportedly trying to find up to US$7 trillion of investment to manufacture the enormous volumes of computer chips he believes the world needs to run artificial intelligence (AI) systems. Altman also recently said the world will need more energy in the AI-saturated future he envisions—so much

more that some kind of technological breakthrough like nuclear fusion may be required.

Altman clearly has big plans for his company's technology, but is the future of AI really this rosy? As a long-time "artificial intelligence" researcher, I have my doubts.

Today's AI systems—particularly generative AI tools such as ChatGPT—are not truly intelligent. What's more, there is no evidence they can become so without fundamental changes to the way they work.

What is AI?

One definition of AI is a computer system that can "[perform tasks commonly associated with intelligent beings](#)."

This definition, like many others, is a little blurry: should we call spreadsheets AI, as they can carry out calculations that once would have been a high-level human task? How about factory robots, which have not only replaced humans but in many instances surpassed us in their ability to perform complex and delicate tasks?

While spreadsheets and robots can indeed do things that were once the domain of humans, they do so by following an algorithm—a process or set of rules for approaching a task and working through it.

One thing we can say is that there is no such thing as "an AI" in the sense of a system that can perform a range of intelligent actions in the way a human would. Rather, there are many different AI technologies that can do quite different things.

Making decisions vs. generating outputs

Perhaps the most important distinction is between "discriminative AI" and "generative AI".

Discriminative AI helps with making decisions, such as whether a bank should give a loan to a [small business](#), or whether a doctor diagnoses a patient with disease X or disease Y. AI technologies of this kind have existed for decades, and bigger and better ones are [emerging all the time](#).

Generative AI systems, on the other hand—ChatGPT, Midjourney and their relatives—generate outputs in response to inputs: in other words, they make things up. In essence, they have been exposed to billions of [data points](#) (such as sentences) and use this to guess a likely response to a prompt. The response may often be "true", depending on the source data, but there are no guarantees.

For generative AI, there is no difference between a "hallucination"—a false response invented by the system—and a response a human would judge as true. This appears to be an inherent defect of the technology, which uses a kind of neural network called a transformer.

AI, but not intelligent

Another example shows how the goalposts of "AI" are constantly moving. In the 1980s, I worked on a computer system designed to provide expert medical advice on laboratory results. It was written up in the US research literature as [one of the first four](#) medical "expert systems" in [clinical use](#), and in 1986 an Australian government report described it as the most successful expert system developed in Australia.

I was pretty proud of this. It was an AI landmark, and it performed a task that normally required highly trained medical specialists. However, the system wasn't intelligent at all. It was really just a kind of look-up table which matched lab test results to high-level diagnostic and patient

management advice.

There is now technology which makes it very easy to build such systems, so there are thousands of them in use around the world. (This technology, based on research by myself and colleagues, is provided by an Australian company called Beamtree.)

In doing a task done by highly trained specialists, they are certainly "AI", but they are still not at all intelligent (although the more complex ones may have thousands and thousands of rules for looking up answers).

The transformer networks used in generative AI systems still run on sets of rules, though there may be millions or billions of them, and they cannot easily be explained in human terms.

What is real intelligence?

If algorithms can produce dazzling results of the kind we see from ChatGPT without being intelligent, what is real intelligence?

We might say intelligence is insight: the judgment that something is or is not a good idea. Think of Archimedes, leaping from his bath and shouting "Eureka" because he had had an insight into the principle of buoyancy.

Generative AI doesn't have insight. ChatGPT can't tell you if its answer to a question is better than Gemini's. (Gemini, until recently known as Bard, is Google's competitor to OpenAI's GPT family of AI tools.)

Or to put it another way: generative AI might produce amazing pictures in the style of Monet, but if it were trained only on Renaissance art it would never invent Impressionism.

Generative AI is extraordinary, and people will no doubt find widespread and very valuable uses for it. Already, it provides extremely useful tools for transforming and presenting (but not discovering) information, and tools for turning specifications into code are already in routine use.

These will get better and better: Google's just-released Gemini, for example, appears to try to [minimize the hallucination problem](#), by using search and then re-expressing the search results.

Nevertheless, as we become more familiar with generative AI systems, we will see more clearly that it is not truly intelligent; there is no insight. It is not magic, but a very clever magician's trick: an algorithm that is the product of extraordinary human ingenuity.

This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Provided by The Conversation