

Not everything we call AI is actually ‘artificial intelligence’. Here’s what you need to know

George Siemens, *University of South Australia*

In August 1955, a group of scientists made a funding request for US\$13,500 to host a summer workshop at Dartmouth College, New Hampshire. The field they proposed to explore was artificial intelligence (AI).

While the funding request was humble, the conjecture of the researchers was not: “every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it”.

Since these humble beginnings, movies and media have romanticised AI or cast it as a villain. Yet for most people, AI has remained as a point of discussion and not part of a conscious lived experience.

AI has arrived in our lives

Late last month, AI, in the form of ChatGPT, broke free from the sci-fi speculations and research labs and onto the desktops and phones of the general public. It’s what’s known as a “generative AI” – suddenly, a cleverly worded prompt can produce an essay or put together a recipe and shopping list, or create a poem in the style of Elvis Presley.

While ChatGPT has been the most dramatic entrant in a year of generative AI success, similar systems have shown even wider potential to create new content, with text-to-image prompts used to create vibrant images that have even won art competitions.

AI may not yet have a living consciousness or a theory of mind popular in sci-fi movies and novels, but it is getting closer to at least disrupting what we think artificial intelligence systems can do.

Researchers working closely with these systems have swooned under the prospect of sentience, as in the case with Google’s large language model (LLM) LaMDA. An LLM is a model that has been trained to process and generate natural language.

Generative AI has also produced worries about plagiarism, exploitation of original content used to create models, ethics of information manipulation and abuse of trust, and even “the end of programming”.

At the centre of all this is the question that has been growing in urgency since the Dartmouth summer workshop: does AI differ from human intelligence?

What does ‘AI’ actually mean?

To qualify as AI, a system must exhibit some level of learning and adapting. For this reason, decision-making systems, automation, and statistics are not AI.

AI is broadly defined in two categories: artificial narrow intelligence (ANI) and artificial general intelligence (AGI). To date, AGI does not exist.

The key challenge for creating a general AI is to adequately model the world with all the entirety of knowledge, in a consistent and useful manner. That’s a massive undertaking, to say the least.

Most of what we know as AI today has narrow intelligence – where a particular system addresses a particular problem. Unlike human intelligence, such narrow AI intelligence is effective *only* in the area in

which it has been trained: fraud detection, facial recognition or social recommendations, for example.

AGI, however, would function as humans do. For now, the most notable example of trying to achieve this is the use of neural networks and “deep learning” trained on vast amounts of data.

Neural networks are inspired by the way human brains work. Unlike most machine learning models that run calculations on the training data, neural networks work by feeding each data point one by one through an interconnected network, each time adjusting the parameters.

As more and more data are fed through the network, the parameters stabilise; the final outcome is the “trained” neural network, which can then produce the desired output on new data – for example, recognising whether an image contains a cat or a dog.

The significant leap forward in AI today is driven by technological improvements in the way we can train large neural networks, readjusting vast numbers of parameters in each run thanks to the capabilities of large cloud-computing infrastructures. For example, GPT-3 (the AI system that powers ChatGPT) is a large neural network with 175 billion parameters.

What does AI need to work?

AI needs three things to be successful.

First, it needs **high-quality, unbiased data**, and lots of it. Researchers building neural networks use the large data sets that have come about as society has digitised.

Co-Pilot, for augmenting human programmers, draws its data from billions of lines of code shared on GitHub. ChatGPT and other large language models use the billions of websites and text documents stored online.

Text-to-image tools, such as Stable Diffusion, DALLÉ-2, and Midjourney, use image-text pairs from data sets such as LAION-5B. AI models will continue to evolve in sophistication and impact as we digitise more of our lives, and provide them with alternative data sources, such as simulated data or data from game settings like Minecraft.

AI also needs **computational infrastructure** for effective training. As computers become more powerful, models that now require intensive efforts and large-scale computing may in the near future be handled locally. Stable Diffusion, for example, can already be run on local computers rather than cloud environments.

The third need for AI is **improved models and algorithms**. Data-driven systems continue to make rapid progress in domain after domain once thought to be the territory of human cognition.

However, as the world around us constantly changes, AI systems need to be constantly retrained using new data. Without this crucial step, AI systems will produce answers that are factually incorrect, or do not take into account new information that’s emerged since they were trained.

Neural networks aren’t the only approach to AI. Another prominent camp in artificial intelligence research is symbolic AI – instead of digesting huge data sets, it relies on rules and knowledge similar to the human process of forming internal symbolic representations of particular phenomena.

But the balance of power has heavily tilted toward data-driven approaches over the last decade, with the “founding fathers” of modern deep learning recently being awarded the Turing Prize, the equivalent of the Nobel Prize in computer science.

Data, computation and algorithms form the foundation of the future of AI. All indicators are that rapid

progress will be made in all three categories in the foreseeable future. .

George Siemens, Co-Director, Professor, Centre for Change and Complexity in Learning, *University of South Australia*

This article is republished from The Conversation under a Creative Commons license. Read the original article.