AI was born at a US summer camp 68 years ago—here's why that event still matters today

August 28 2024, by Sandra Peter

Dartmouth College. Credit: Pixabay/CC0 Public Domain

Imagine a group of young men gathered at a picturesque college campus

in New England, in the United States, during the northern summer of 1956.

It's a small casual gathering. But the men are not here for campfires and nature hikes in the surrounding mountains and woods. Instead, these pioneers are about to embark on an experimental journey that will spark countless debates for decades to come and change not just the course of technology—but the course of humanity.

Welcome to the Dartmouth Conference—the birthplace of artificial intelligence (AI) as we know it today.

What transpired here would ultimately lead to ChatGPT and the many other kinds of AI which now help us diagnose disease, detect fraud, put together playlists and write articles (well, not this one). But it would also create some of the many problems the field is still trying to overcome. Perhaps by looking back, we can find a better way forward.

The summer that changed everything

In the mid 1950s, rock'n'roll was taking the world by storm. Elvis's Heartbreak Hotel was topping the charts, and teenagers started embracing James Dean's rebellious legacy.

But in 1956, in a quiet corner of New Hampshire, a different kind of revolution was happening.

The [Dartmouth Summer Research Project on Artificial Intelligence](#), often remembered as the Dartmouth Conference, kicked off on June 18 and lasted for about eight weeks. It was the brainchild of four American computer scientists—John McCarthy, Marvin Minsky, Nathaniel Rochester and Claude Shannon—and brought together some of the brightest minds in [computer science](#), mathematics and cognitive

psychology at the time.

These scientists, along with some of the 47 people they invited, set out to tackle an ambitious goal: to make intelligent machines.

As McCarthy put it in the conference proposal, they aimed to find out "how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans."

The birth of a field—and a problematic name

The Dartmouth Conference didn't just coin the term "artificial intelligence"; it coalesced an entire field of study. It's like a mythical Big Bang of AI—everything we know about machine learning, neural networks and deep learning now traces its origins back to that summer in New Hampshire.

But the legacy of that summer is complicated.

Artificial intelligence won out as a name over others proposed or in use at the time. Shannon preferred the term "automata studies," while two other conference participants (and the soon-to-be creators of the first AI program), Allen Newell and Herbert Simon, continued to use "complex information processing" for a few years still.

But here's the thing: having settled on AI, no matter how much we try, today we can't seem to get away from comparing AI to human intelligence.

This comparison is both a blessing and a curse.

On the one hand, it drives us to create AI systems that can match or exceed human performance in specific tasks. We celebrate when AI

outperforms humans in games such as chess or Go, or when it can detect cancer in medical images with greater accuracy than human doctors.

On the other hand, this constant comparison leads to misconceptions.

When a [computer beats a human at Go](link), it is easy to jump to the conclusion that machines are now smarter than us in all aspects—or that we are at least well on our way to creating such intelligence. But AlphaGo is no closer to writing poetry than a calculator.

And when a large language model sounds human, [we start wondering if it is sentient](link).

But ChatGPT is no more alive than a talking ventriloquist's dummy.

The overconfidence trap

The scientists at the Dartmouth Conference were incredibly optimistic about the future of AI. They were convinced they could solve the problem of machine intelligence in a single summer.

This overconfidence has been a recurring theme in AI development, and it has led to several cycles of hype and disappointment.

[Simon stated in 1965](link) that "machines will be capable, within 20 years, of doing any work a man can do." [Minsky predicted in 1967](link) that "within a generation […] the problem of creating '[artificial intelligence](link)' will substantially be solved."

Popular futurist [Ray Kurzweil now predicts](link) it's only five years away: "we're not quite there, but we will be there, and by 2029 it will match any person."

Reframing our thinking: New lessons from Dartmouth

So, how can AI researchers, AI users, governments, employers and the broader public move forward in a more balanced way?

A key step is embracing the difference and utility of machine systems. Instead of focusing on the race to "artificial general intelligence," we can focus on the unique strengths of the systems we have built—for example, the enormous creative capacity of image models.

Shifting the conversation from automation to augmentation is also important. Rather than pitting humans against machines, let's focus on how AI can assist and augment human capabilities.

Let's also emphasize ethical considerations. The Dartmouth participants didn't spend much time discussing the ethical implications of AI. Today, we know better, and must do better.

We must also refocus research directions. Let's emphasize research into AI interpretability and robustness, interdisciplinary AI research and explore new paradigms of intelligence that aren't modeled on human cognition.

Finally, we must manage our expectations about AI. Sure, we can be excited about its potential. But we must also have realistic expectations, so that we can avoid the disappointment cycles of the past.

As we look back at that summer camp 68 years ago, we can celebrate the vision and ambition of the Dartmouth Conference participants. Their work laid the foundation for the AI revolution we're experiencing today.

By reframing our approach to AI—emphasizing utility, augmentation, ethics and realistic expectations—we can honor the legacy of Dartmouth while charting a more balanced and beneficial course for the future of AI.

After all, the real intelligence lies not just in creating smart machines, but in how wisely we choose to use and develop them.

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

Provided by The Conversation