# AI Can Do Great Things—if It Doesn't Burn the Planet

*Will Knight*

Last month, researchers at OpenAI in San Francisco revealed an algorithm capable of learning, through trial and error, how to manipulate the pieces of a Rubik's Cube using a robotic hand. It was a remarkable research feat, but it required more than 1,000 desktop computers plus a dozen machines running specialized graphics chips crunching intensive calculations for several months.

The effort may have consumed about 2.8 gigawatt-hours of electricity, estimates Evan Sparks, CEO of Determined AI, a startup that provides software to help companies manage AI projects. That's roughly equal to the output of three nuclear power plants for an hour. A spokesperson for OpenAI questioned the calculation, noting that it makes several assumptions. But OpenAI declined to disclose further details of the project or offer an estimate of the electricity it consumed.

Artificial intelligence routinely produces startling achievements, as computers learn to recognize images, converse, beat humans at sophisticated games, and drive vehicles. But all those advances require staggering amounts of computing power—and electricity—to devise and train algorithms. And as the damage caused by climate change becomes more apparent, AI experts are increasingly troubled by those energy demands.

"The concern is that machine-learning algorithms in general are consuming more and more energy, using more data, training for longer and longer," says Sasha Luccioni, a postdoctoral researcher at Mila, an AI research institute in Canada.

It's not just a worry for academics. As more companies across more industries begin to use AI, there's growing fear that the technology will only deepen the climate crisis. Sparks says that Determined.ai is working with a pharmaceutical firm that's already using huge AI models. "As an industry, it's worth thinking about how we want to combat this," he adds.

Some AI researchers are thinking about it. They're using tools to track the energy demands of their algorithms, or taking steps to offset their emissions. A growing number are touting the energy efficiency of their algorithms in research papers and at conferences. As the costs of AI rise, the AI industry is developing a new appetite for algorithms that burn fewer kilowatts.

Luccioni recently helped launch a website that lets AI researchers roughly calculate the carbon footprint of their algorithms. She is also testing a more sophisticated approach—code that can be added to an AI program to track the energy use of individual computer chips. Luccioni and others are also trying to persuade companies that offer tools for tracking the performance of code to include some measure of energy or carbon footprint. "Hopefully this will go toward full transparency," she says. "So that people will include in the footnotes 'we emitted X tons of carbon, which we offset.'"

The energy required to power cutting-edge AI has been on a steep upward curve for some time. Data published by OpenAI shows that the computing power required for key AI landmarks

over the past few years, such as DeepMind's Go-playing program AlphaZero, has doubled roughly every 3.4 months—increasing 300,000 times between 2012 and 2018. That's faster than the rate at which computing power historically increased, the phenomenon known as Moore's Law (named after Gordon Moore, cofounder of Intel.)

Recent advances in natural language processing—an AI technique that helps machines parse, interpret, and generate text—have proven especially power-hungry. A research paper from a team at UMass Amherst found that training a single large NLP model may consume as much energy as a car over its entire lifetime—including the energy needed to build it.

Training a powerful machine-learning algorithm often means running huge banks of computers for days, if not weeks. The fine-tuning required to perfect an algorithm, by for example searching through different neural network architectures to find the best one, can be especially computationally intensive. For all the hand-wringing, though, it remains difficult to measure how much energy AI actually consumes, and even harder to predict how much of a problem it could become.

The Department of Energy estimates that data centers account for about 2 percent of total US electricity usage. Worldwide, data centers consume about 200 terawatt hours of power per year—more than some countries. And the forecast is for significant growth over the next decade, with some predicting that by 2030, computing and communications technology will consume between 8 percent and 20 percent of the world's electricity, with data centers accounting for a third of that.

In recent years, companies offering cloud computing services have sought to address spiraling power consumption and offset carbon emissions with varying measures of success. Google, for example, claims "zero net carbon emissions" for its data centers, thanks to extensive renewable energy purchases. Microsoft last week announced a plan to become "carbon negative" by 2030, meaning it would offset all of the carbon produced by the company over its history. OpenAI signed a deal to use Microsoft's cloud last July.

It isn't clear how the AI boom will fit with the bigger picture of data center energy use, or how it might alter it. Cloud providers do not disclose the overall energy demands of machine-learning systems. Microsoft, Amazon, and Google all declined to comment.

Jonathan Koomey, a researcher and consultant who tracks data center energy use, cautions against drawing too many conclusions from cutting-edge AI demos. He notes that AI algorithms often run on specialized chips that are more efficient, so new chip architectures may offset some of the projected demand for compute power. He also says that the IT industry has in the past offset rising energy demands in one domain by lowering energy use in others. "People are likely to take isolated anecdotes and extrapolate to get eye popping numbers, and these numbers are almost always too high," Koomey says.

Still, as companies and other organizations increasingly use artificial intelligence, experts say it will become important to understand the technology's energy footprint, both in data centers and in other devices and gadgets. "I would agree that the analysis community needs to get a handle on it," says Eric Masanet, a professor at Northwestern University who leads its Energy and Resource Systems Analysis Laboratory.

Some AI researchers aren't waiting for the industry to wake up. Luccioni of Mila helped

organize a workshop on climate change last month at an important AI conference, NeurIPS, and she was pleased to find that the event was standing room only. “There’s a lot of interest in this,” she says.

The Allen Institute for AI, a research institute founded by the late Microsoft cofounder Paul Allen, has also called for greater awareness of AI’s environmental impact. The institute’s CEO, Oren Etzioni, says he is encouraged by the efforts of researchers, as many papers now include some account of the computational intensity of a particular algorithm or experiment.

Etzioni adds that the industry as a whole is gradually waking up to energy efficiency. Even if this is largely because of the cost involved with training large AI models, it could help prevent AI from contributing to a looming climate catastrophe. “AI is clearly moving toward lighter models and greener AI,” he says.

---