NewScientist

INSIGHT 10 October 2018

AI's dirty secret: Energy-guzzling machines may fuel global warming

Advances in artificial intelligence could lead to massive growth in energy use as smart machines push into every corner of our lives

AI is used in cars (above), smart speakers (below) and for playing games (bottom)
Christie Hemm Klok / New York Times / Redux / eyevine

By **Michael Le Page**

ARTIFICIAL intelligence breakthroughs have become a regular occurrence in recent years. One of the most impressive achievements so far was in 2016, when Google DeepMind's AlphaGo AI beat champion Lee Sedol at one of the world's most complex games, Go.

The feat made headlines around the world as an example of machines besting humans, but in some sense it wasn't a fair fight. Sedol's brain would have been consuming around 20 watts of power, with only a fraction of that being used for the game itself. By contrast, AlphaGo was using some 5000 watts.

It isn't widely appreciated how incredibly energy hungry AI is. If you ran AlphaGo non-stop for a year, the electricity alone would cost about £6000. That doesn't matter for one-off events, like an epic Go showdown. But it does matter if billions of people want their smartphones to be truly smart, or have their cars drive themselves.

“Deep learning involves performing ever more computations on ever more data”

Many potential uses of AI simply won’t become widespread if they require too much energy. On the flip side, if the uses are so desirable or profitable that people don’t care about the costs, it could lead to a surge in electricity consumption and make it even harder to limit further warming of the planet.

AI consumes so much energy because the technique behind these recent breakthroughs, deep learning, involves performing ever more computations on ever more data. “The models are getting deeper and getting wider, and they use more and more energy to compute,” says Max Welling of the University of Amsterdam, the Netherlands.

Take image recognition, for instance. This is done by training neural nets that mimic the brain. Images are split up and fed into the first layer of the net, which might look for edges. The results get passed to the next layer, which might look for simple shapes such as a dog’s tail or a letter, and so on.

The more you want the net to recognise, the more layers it must have and the more complex it has to be. Identifying different breeds of dogs, say, is a much harder task than simply recognising dogs.

20 per cent
The additional energy consumption AI adds to a self-driving car

The most sophisticated neural nets have grown “outrageously large”, as a recent paper put it. Their complexity can be measured by the number of parameters in the net – and some now have more than 100 billion parameters.

This is important for two reasons, says Welling. First, energy is money. “There is an economic ceiling to what is a useful AI technology,” he says. For big online services like Facebook and Google, the cost per use of AI has to be tiny – in the region of a thousandth of a penny.

Second, mobile devices can only use so much energy. It is not just about battery life, but the limit that comes from ensuring devices don’t get too hot, Welling says.

This is why voice assistants like Siri or Alexa need an internet connection for full functionality – your phone or smart speaker doesn’t have the processing power needed to run the AI locally, or the space to store all the associated data.

All this means that our use of AI can’t keep expanding indefinitely unless we reduce the energy requirements. “What matters is how much intelligence we can squeeze out per joule,” says Welling.

Battery busters

Take self-driving cars. These require all sorts of extra systems, from cameras to radar, that use power and also increase weight and drag, further increasing energy use. But the single largest consumer of energy besides the engine is the processor running the AI.

According to a study out earlier this year, self-driving cars could use up to 20 per cent more energy than conventional cars. That is a big issue for a battery-powered car, limiting its range and increasing running costs. What’s more, the study assumes the AI processor consumes about 200 watts, even

Michael Leckie for The times T2

though current prototypes consume in excess of 2000 watts.

For taxi companies using AI to directly replace human drivers, the savings in wages would probably far outweigh the higher energy costs. But for ordinary car owners this would be a major issue.

5 per cent
The estimated fraction of global electricity used by computing

Thankfully, people are already working to improve the efficiency of deep-learning algorithms, by getting rid of unnecessary precision in the calculations that greatly increases the processing required, for example.

“Even if you use fewer bits, you can train a neural network to get the same results,” says Avishek Biswas of the Massachusetts Institute of Technology.

Many big tech companies are also developing specialised hardware for running AI. Nvidia, for instance, has produced a chip just for self-driving cars – although the latest version still uses 500 watts.

Google, meanwhile, has created what it calls tensor processing units, or TPUs, designed to run its TensorFlow machine-learning framework, and is renting them out via the cloud.

There are more revolutionary designs in the works, too. Shunting data back and forth between the memory and processor wastes a great of deal of energy, says Biswas. So he has developed a chip

intended for smartphones that slashes energy use by around 95 per cent by carrying out key operations in the memory itself.

Lee Jin-man/AP/REX/Shutterstock

The pace of progress is astonishing. The latest version of AlphaGo, called AlphaGo Zero, runs on just four TPUs, and consumes around 400 watts.

But will these new chips actually bring down the overall energy use of AI? As AI becomes more energy efficient, we will build better systems and use it to do more, meaning there may be no net gain. "Certainly energy use will keep on going up with more sophisticated AI," says Welling.

In fact, AI is set to play a major part in the world economy. Global spending on AI systems is forecast to reach $60 billion in 2021, according to market research company IDC. Meanwhile, another research company, Gartner, predicts that the value of AI-derived business will reach up to $4 trillion by 2022.

$4 trillion
Forecast value of AI-derived business in 2022

But we have no idea what this means for global energy use. It is estimated that computing and the internet use about 5 per cent of electricity globally, and AI is clearly a fraction of this at the moment.

"Not much is known about it at this stage," says Anders Andrae of Huawei Technologies, who studies the energy consumption of information technology. "But it is very dangerous to say that this is nothing to care about."

That's because what matters is not just how much energy AI itself uses, but also the knock-on effects. To go back to the example of self-driving cars, in theory they could operate in a way that reduces overall energy use despite their higher energy requirements. For instance, they could line up behind each other to reduce drag and communicate with each other so they don't need to stop at intersections.

But most of these savings require a high proportion of the cars on the roads to be autonomous – just as a social network like Facebook can only really take off once it reaches a critical mass of people. "The network-dependent effects won't happen in the early phases," says Zia Wadud at the University of Leeds, UK, who studies transport and energy.

"Energy use will keep going up as we develop more sophisticated artificial intelligence"

People with self-driving cars are likely to commute longer distances, because they won't need to focus on the road and can use the time for browsing the internet or watching videos. Wadud thinks this could double energy use – and his study doesn't count the AI part, or the massive increase in data traffic these vehicles are expected to generate.

In recent years, electricity demand has remained flat despite soaring internet traffic, thanks to big improvements in efficiency. But Andrae thinks we have reached the limits of the easy gains in efficiency.

If the amount of data we generate rises without big efficiency gains, information and communication technology energy use could soar to 20 per cent of all electricity by 2025 and account for a 20th of all carbon emissions, Andrae warns. "This could be a very huge problem," he says.

High costs make this worst-case scenario unlikely. But Andrae's work didn't factor in AI and could even be underestimating data growth, for instance from lots of AI cars talking to each other.

"I think AI will generate even more data than we have seen," he says. "If more data means money, they will use artificial intelligence to make more data, and more data means more electricity." And until our electricity systems are 100 per cent renewable, more electricity means more global warming.

This article appeared in print under the headline "Knowledge means power"

Magazine issue 3199, published 13 October 2018

GET THE APP

FOLLOW US

Back to top