

News in focus

In February, an even bigger push came for the open-source movement when Facebook's parent company, Meta, made a model called LLaMA freely available to selected external developers. Within a week, the LLaMA code was leaked and published online for anyone to download.

The availability of LLaMA has been a game-changer for AI researchers. It is much smaller than other LLMs, meaning that it doesn't require large computing facilities to host the pretrained model or to adapt it for specialized applications. The biggest version of LLaMA consists of 65 billion parameters: the variables set during the neural network's initial, general-purpose training. This is less than half of BLOOM's 176 billion parameters, and a fraction of the 540 billion parameters of Google's latest LLM, PaLM2.

"With LLaMA, some of the most interesting innovation is on the side of efficiency," says Joelle Pineau, vice-president of AI research at Meta and a computer scientist at McGill University in Montreal, Canada.

Open-source developers have been experimenting with ways of shrinking LLaMA down even more. Some of these techniques involve keeping the number of parameters the same but reducing the parameters' precision – an approach that does not cause unacceptable drops in performance. Other ways of downsizing neural networks involve reducing the number of parameters, for example, by training a separate, smaller neural network on the responses of a large, pretrained network, rather than directly on the data.

Within weeks of the LLaMA leak, developers managed to produce versions that could fit onto laptops and even a Raspberry Pi, the bare-bones, credit-card-sized computer that is a favourite of the 'maker' community. Hugging Face is now primarily using LLaMA, and is not planning to push for a BLOOM-2.

Shrinking down AI tools could help to make them more widely accessible, says Vukosi Marivate, a computer scientist at the University of Pretoria. For example, it could help organizations such as Masakhane, a community of African researchers led by Marivate that is trying to make LLMs work for languages for which there isn't a lot of existing written text that can be used to train a model. But the push towards expanding access still has a way to go: for some researchers in low-income countries, even a top-of-the-range laptop can be out of reach. "It's been great," says Marivate, "but I would also ask you to define 'cheap'."

Looking under the hood

For many years, AI researchers routinely made their code open source and posted their results on repositories such as arXiv. "People collectively understood that the field would progress more quickly if we agreed share things with each other," says Colin Raffel, a

computer scientist at the University of North Carolina at Chapel Hill. The innovation underlying current state-of-the-art LLMs, called the Transformer architecture, was created at Google and released as open source, for example.

Making neural networks open source enables researchers to look 'under the hood' to try to understand why the systems sometimes

"People collectively understood that the field would progress more quickly if we agreed to share things."

answer questions in unpredictable ways and can carry biases over from the data they were pre-trained on, says Ellie Pavlick, a computer scientist at Brown University in Providence, Rhode Island. "One benefit is allowing many people – especially from academia – to work on mitigation strategies," she says. "If you have a thousand eyes on it, you're going to come up with better ways of doing it."

Pavlick's team has analysed open-source systems such as BLOOM and found ways to identify and fix biases that are inherited from the training data – the prototypical example

being how language models tend to associate 'nurse' with the female gender and 'doctor' with the male gender.

Pretraining bottleneck

Even if the open-source boom goes on, the push to make language AI more powerful will continue to come from the largest players. Only a handful of companies are able to create LLMs from scratch that can truly push the state of the art. Pretraining them requires massive resources – GPT-4 and Google's PaLM 2 took tens of millions of dollars' worth of computing time – and plenty of 'secret sauce'.

"We have some general recipes, but there are often small details that are not documented," says Pavlick. "It's not like someone gives you the code, you push a button and you get a model."

"Very few organizations and people can pretrain," says Louis Castricato, an AI researcher at open-source software company Stability AI in New York. "It's still a huge bottleneck."

Other researchers warn that making powerful language models accessible increases the chances that they will end up in the wrong hands. Connor Leahy, chief executive of the AI company Conjecture in London, thinks that AI will soon be intelligent enough to put humanity at risk. "We shouldn't open-source any of this," he says.

BOLD STUDY THAT GAVE PEOPLE COVID REVEALS 'SUPERSHEDDING'

A small subset of people spew huge amounts of virus into the air – even if they have mild symptoms.

By Saima Sidik

A study of people who were intentionally infected with SARS-CoV-2 has provided a wealth of insights into viral transmission – showing, for example, that a select group of people are 'supershedders' who release vastly more virus into the air than do others¹.

The publication describes data from a controversial 'challenge study', in which scientists deliberately infected volunteers with the virus that causes COVID-19 (ref. 2). Although the approach drew opposition, the work has now yielded data on questions central to public health, such as whether symptom severity correlates with how contagious people are, and whether home COVID-19 tests play a part in reducing viral spread.

The study, published on 9 June in *The Lancet*

Microbe, also suggests that human physiology, not the virus, is to blame for some of the inconsistency of COVID-19 (ref. 1).

Design with benefits

Challenge studies are "very bold", says infectious-disease physician Monica Gandhi at the University of California, San Francisco. Some people argue that it's unethical to give participants an infection that can cause severe illness, but the research design comes with benefits. Challenge studies can substantially speed up vaccine testing, and they're the only way to understand certain aspects of COVID-19, such as the stage before people test positive or develop symptoms.

Researchers inoculated 34 healthy young people by squirting a known quantity of viral particles up their noses. Eighteen developed infections and spent at least 14 days confined

to hospital rooms. Each day, researchers measured viral levels in the participants' noses and throats, in the air, and on their hands and various surfaces in the rooms.

The symptoms and severity of naturally acquired COVID-19 might vary depending on factors such as the transmission route. But in the challenge study, "we know that that was all controlled", says infectious-disease researcher Anika Singanayagam at Imperial College London, a co-author of the paper.

Of the 18 participants who developed infections, 2 shed 86% of the airborne virus detected over the course of the entire study – even though both had only mild symptoms. Previous research³ has provided evidence for the existence of superspreaders who infect large numbers of people. But whether such people are also 'supershedders' who emit copious amounts of virus, or simply have many social contacts, was up for debate, says disease ecologist Pablo Beldomenico at the Institute of Veterinary Sciences of the Coast in Esperanza, Argentina. This study "supports the existence of supershedders", he says.

Rapid tests show their value

Participants used lateral flow tests, also known as rapid antigen tests, on each day that they were in isolation. None of the participants emitted a detectable level of virus into the air before testing positive, and only a small proportion left detectable virus on their hands, on surfaces or on masks that they donned temporarily.

By the time they tested positive, most participants had already experienced mild symptoms, such as tiredness. That means that if people test as soon as they detect symptoms, rapid tests "can be a powerful tool" for controlling viral spread, says infectious-disease researcher Christopher Brooke at the University of Illinois at Urbana-Champaign.

Some researchers question the relevance of the study's results to today's world. The route of infection – drops administered through the nose – differs from that of most natural infections, says airborne-infectious-disease researcher Donald Milton at the University of Maryland in College Park. As a result, viral shedding might differ between study participants and people who become infected in the real world. The now-dominant Omicron variant spreads differently from the 2020 strain the authors used, his colleague Kristen Coleman adds.

Despite these limitations, the work "still gives us really useful information", Singanayagam says, adding that the results are in line with what she and her colleagues have observed with naturally acquired infections.

1. Zhou, J. et al. *Lancet Microbe* [https://doi.org/10.1016/S2666-5247\(23\)00101-5](https://doi.org/10.1016/S2666-5247(23)00101-5) (2023).
2. Killingley, B. et al. *Nature Med.* **28**, 1031–1041 (2022).
3. Beldomenico, P. M. *Int. J. Infect. Dis.* **96**, 461–463 (2020).

NET-ZERO CARBON PLEDGES ARE SURGING — ARE THEY SERIOUS?

Governments and firms want to cut emissions, but robust strategies to achieve their goals are lacking.

By Katharine Sanderson

The number of countries, regions, cities and companies worldwide that have set targets to reach net-zero emissions has increased significantly over the past two years, but the strategies needed to meet those pledges are lacking.

This is the conclusion of *Net Zero Stocktake 2023*, a report compiled by the climate researchers who run the Net Zero Tracker, a collaboration that aims to increase the accountability of net-zero pledges. The tracker keeps count of entities of various sizes – including nations, regions and companies – that have pledged to ensure that, by 2050, they either make no greenhouse-gas emissions or completely balance any they make through removal mechanisms. The tracker also assesses whether the commitments have legal heft. The stocktake is the third such exercise conducted by the group in the past three years.

The latest analysis shows that the proportion of countries with net-zero pledges that are either enshrined in law or in a weighty policy document has leapt from 7% in December 2020 to 75% in June 2023. "This shows that governments see actions towards net-zero as critical for the long term," says Malango Mughogho, managing director of sustainable-finance firm ZeniZeni in Johannesburg, South Africa, and former member of a United Nations net-zero group.

The stocktake focuses on the quality, or integrity, of the measures put in place to achieve the pledges. "We looked at the data

and we saw such little movement on integrity," says John Lang, who leads the Net Zero Tracker and who works at the Energy and Climate Intelligence Unit, an advisory firm in London. The group compared pledges and progress with the requirements laid out in the UN's Race to Zero campaign.

Race to Zero outlines criteria called the 'starting line', which are "the very minimum procedural requirements for a decent pledge", says Lang. The criteria include having a pledge, a plan and published evidence of action taken towards reaching the target. Most states, regions or cities that have made pledges have not met these criteria. "We saw no movement whatsoever," says Lang.

"Delivery is at the heart of net zero," says Mughogho. She adds that creating fair regulations on net-zero targets across the globe at the levels of cities, regions, businesses and financial institutions could be one of the easiest measures to implement. "It's a no-brainer," she says. The UN panel on climate change will undertake its first global stocktake of progress towards the Paris Agreement goals ahead of the COP28 climate meeting in November. Mughogho says that this will help to refine understanding of the actions that governments and companies are taking (see 'Growing goals').

The net-zero stocktake for the first time looks at entities' plans to phase out fossil fuels. Of the 114 fossil-fuel firms assessed, 67% have net-zero pledges. But none has a plan to phase out oil and gas completely. This doesn't align with scientific or policy consensus, the authors say.

GROWING GOALS

Governments and companies are increasingly setting net-zero targets, but plans to implement the goals are lacking in almost all cases.

