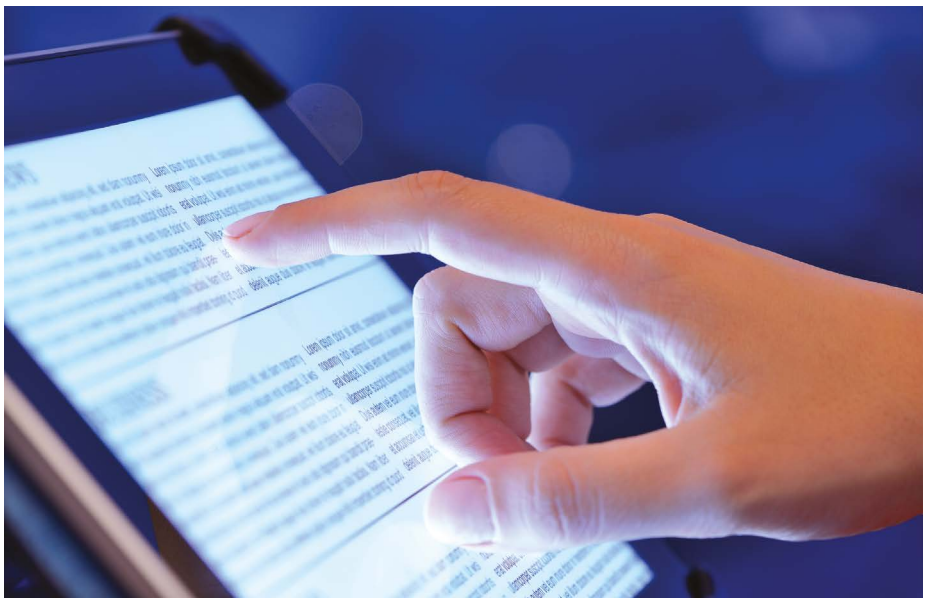

Researchers have warned against possible harm from AI that processes and generates text.

GETTY

OPEN-SOURCE LANGUAGE AI CHALLENGES BIG TECH'S MODELS

BLOOM aims to address the biases that machine-learning systems inherit from the texts they train on.

By Elizabeth Gibney

An international team of around 1,000 volunteers, most of them academics, has tried to break big tech's stranglehold on natural-language processing and reduce the harm it can cause. Trained with US$7-million-worth of publicly funded computing time, the BLOOM language model will rival in scale those made by firms Google and OpenAI, but will be open-source. BLOOM will also be the first model of its scale to be multilingual.

The collaboration, called BigScience, launched an early version of the model on 17 June, and hopes that it will ultimately help to reduce harmful outputs of artificial intelligence (AI) language systems. Big tech firms increasingly use models that recognize and generate language in applications ranging from chat bots to translators, and the models can sound so eerily human that a Google engineer this month claimed that the firm's AI model was sentient (Google strongly denies that the AI possesses sentience). But such models also have serious practical and ethical flaws, such as parroting human biases. These are difficult to tackle because the inner workings of most such models are closed to researchers.

As well being a tool for exploring AI, BLOOM will be open for a range of research uses, such as extracting information from historical texts and making classifications in biology. "We think that access to the model is an essential step to do responsible machine learning," says Thomas Wolf, co-founder of Hugging Face, a company that hosts an open-source platform for AI models and data sets, and has helped to spearhead the initiative.

Learning machines

Large language models are algorithms that learn statistical associations between billions of words and phrases to perform tasks such as generating summaries, translating, answering questions and classifying text. Built using brain-inspired architectures known as neural networks, the models learn through adjusting values, called parameters, by blanking out words and comparing their predictions with reality. BLOOM has 176 billion parameters, on a par with GPT-3, one of the best-known such models, which was created by the non-profit firm OpenAI and licensed by Microsoft.

Although such models are sometimes impressive – generating poetry or correctly answering trivia questions – they have no sense of the meaning of language, which causes them to create gibberish, too. More worryingly, they can also promote abuse or self-harm, and echo existing racist or sexist associations that are sown throughout the human-written text they learn on, such as linking 'Islam' with terrorism. The models generally cost millions of dollars to train and have an enormous carbon footprint (BigScience plans to reveal its carbon emissions eventually).

Whereas most natural-language models are built by small in-house teams, BLOOM was the work of many hundreds of researchers – mostly academics – including ethicists, legal scholars and philosophers, but also some employees of Facebook and Google, working in a personal capacity. To train BLOOM, BigScience was granted free access to France's national Jean Zay supercomputer facility outside Paris. The model is currently in the final few weeks of its three-month training period.

Hand-picked text

Models are only as good as the data sets they are based on, so a major task was selecting what texts the model should learn from, says Yacine Jernite, a machine-learning researcher at Hugging Face. Most major models rip language directly from the web, including sites such as Reddit. Instead, the BigScience researchers hand-picked nearly two-thirds of their 341-billion-word data set from 500 sources. Among them was Semantic Scholar, an AI-backed search engine for academic publications that also includes content such as *Nature* news articles. The sources were suggested during a series of workshops, including some with community groups, such as the African natural-language-processing community Masakhane, LatinX in AI and Machine Learning Tokyo. "We wanted to make sure people with proximity to the data, their country, the language they speak, had a hand in choosing what language came into the model's training," says Jernite.

To make full use of the computing power available, the team topped up the data trove using a multilingual web crawl, filtered for quality and with some redaction for privacy. The collaboration also attempted to reduce the usual over-representation of porn sites (which can lead to sexist associations in the model) but without excluding keywords that would remove content associated with frank discussion of sexuality in communities that are often under-represented.

Jernite acknowledges that BLOOM will not be free of biases. But by providing it with multicultural and high-quality sources, the team hopes to improve on existing models. Crucially, because the code and data set behind the model are open, researchers can try to understand the roots of harmful behaviours, which could improve future iterations, says Wolf.

Evaluation of the model will also differ from the usual benchmarks, says Ellie Pavlick, a

natural-language-learning researcher at Brown University in Providence, Rhode Island. As well as comparing BLOOM with other models in its abilities to, for example, answer questions, researchers want to look at more diverse metrics, such as how strongly it makes certain stereotyped associations or how biased its abilities are towards a specific language. Pavlick hopes that because the model has been trained to be multilingual, it might have a deeper understanding of language, which could help in its ability to generalize to a diversity of tasks.

Free to use

The fully trained BLOOM model will be available to download for researchers who want to experiment with it or train it on new data for specific applications. But downloading and running it require significant hardware capacity. Because that's available to so few research teams, BigScience will also publish smaller, less-hardware-intensive versions and create a distributed system that allows laboratories to share the model across their servers. In addition, Hugging Face has released a web application that will enable anyone to query BLOOM without downloading it.

BLOOM could find uses in research outside AI. Francesco de Toni, a linguist at the University of Western Australia in Perth, jointly leads a BigScience working group that is looking at using models to extract information from collections of historical texts that are too large to be gone through by hand. Models can, for example, extract all the names or goods mentioned in a collection of letters by Renaissance merchants.

CRUCIAL BIODIVERSITY SUMMIT WILL GO AHEAD IN CANADA, NOT CHINA

After pandemic delays, researchers say countries must agree on ambitious targets.

By Natasha Gilbert

Scientists are relieved that a pivotal summit to finalize a new global agreement to protect the environment will go ahead this year, after two years of delays because of the COVID-19 pandemic. But they say the hard work of negotiating an ambitious deal lies ahead.

The United Nations Convention on Biological Diversity (CBD) announced on 21 June that the meeting will move from Kunming in China to Montreal, Canada. The second part Fifteenth meeting of the Conference of the Parties (COP15) will bring together representatives from almost 200 CBD member states from 5 to 17 December. China will continue as president of COP15, and Huang Runqiu, China's minister of ecology and environment, will remain the chair.

Conservation and biodiversity scientists were becoming increasingly concerned that China's strict 'zero COVID' strategy, which uses measures such as lockdowns to quash all SARS-CoV-2 infections, would force the host nation to delay the meeting again. Researchers warned that another setback to the agreement, which aims to halt the alarming rate of species extinctions and to protect vulnerable ecosystems, would be disastrous for countries' abilities to meet the proposed, ambitious targets to protect biodiversity over the next decade.

"We are relieved and thankful that we have a firm date for these critically important biodiversity negotiations within this calendar year," says Andrew Deutz, a specialist in biodiversity law and finance at the Nature Conservancy, a conservation group in Arlington, Virginia. "The global community is already behind in agreeing, let alone implementing, a plan to halt and reverse biodiversity loss by 2030," he says.

IVAN VALENCIA/BLOOMBERG/GETTY

Deforestation in places such as the Amazon contributes to biodiversity loss.

With the date now set, Anne Larigauderie, executive secretary of the Intergovernmental Platform on Biodiversity and Ecosystem Services, says key to the success of the global biodiversity agreement will be to focus on the direct and indirect drivers of nature loss, and the behaviours that underpin them. "Policy should be led by science, action [should be] adequately resourced and change should be transformative," she adds.

The decision to move the meeting came after discussions by the representatives of the decision-making body of the COP. China and Canada then thrashed out the details of how the move would work. The CBD has provisions that if a host country is unable to hold a COP, the meeting changes to the home of the convention's secretariat, Montreal.

Announcing the decision, Elizabeth Mrema, executive secretary of the CBD, said in a statement, "I want to thank the government of China for their flexibility and continued commitment to advancing our path towards an ambitious post-2020 global biodiversity framework."

In a statement, Runqiu said, "China would like to emphasize its continued strong commitment, as COP president, to ensure the success of the second part of COP15, including the adoption of an effective post-2020 global biodiversity framework, and to promote its delivery throughout its presidency."

China also agreed to pay for ministers from the lowest-income countries and some small island states to travel to Montreal to participate in the meeting.

Work ahead

Paul Matiku, an environmental scientist and head of Nature Kenya, a conservation organization in Nairobi, says that the move "is a welcome decision" after "the world lost patience after a series of postponements".

But he says that wealthy nations need to reach deeper into their pockets to help low- and middle-income countries — which are home to much of the world's biodiversity — to implement the deal. Disputes over funding already threaten to stall the agreement.

"The extent to which the CBD is implemented will depend on the availability of predictable, adequate financial flows from developed nations to developing-country parties," says Matiku.

Further talks were held in Nairobi from 21 to 26 June. Deutz hopes that countries can find common ground on key issues such as financing before heading to Montreal. Having a firm date set for COP15 will help to push negotiations forward, he says. "Negotiators only start to compromise when they are up against a deadline," he adds. "Now they have one."