Google DeepMind predicts weather more accurately than leading system

This article is more than 3 months old

AI program GenCast performed better than ENS forecast at predicting day-to-day weather and paths of hurricanes and cyclones

For those who keep an eye on the elements, the outlook is bright: researchers have built an artificial intelligence-based weather forecast that makes faster and more accurate predictions than the best system available today.

GenCast, an AI weather program from Google DeepMind, performed up to 20% better than the ENS forecast from the European Centre for Medium-Range Weather Forecasts (ECMWF), widely regarded as the world leader.

In the near term, GenCast is expected to support traditional forecasts rather than replace them, but even in an assistive capacity it could provide clarity around future cold blasts, heatwaves and high winds, and help energy companies predict how much power they will generate from windfarms.

In a head-to-head comparison, the program churned out more accurate forecasts than ENS on day-to-day weather and extreme events up to 15 days in advance, and was better at predicting the paths of destructive hurricanes and other tropical cyclones, including where they would make landfall.

“Outperforming ENS marks something of an inflection point in the advance of AI for weather prediction,” said Ilan Price, a research scientist at Google DeepMind. “At least in the short term, these models are going to accompany and be alongside existing, traditional approaches.”

Traditional physics-based weather forecasts solve vast numbers of equations to produce their predictions, but GenCast learned how global weather evolves by training on 40 years of historic data generated between 1979 and 2018. This included wind speed, temperature, pressure, humidity and dozens more variables at different altitudes.

Given the latest weather data, GenCast predicts how conditions will change around the planet in squares of up to 28km by 28km for the next 15 days in 12-hour steps.

While a traditional forecast takes hours to run on a supercomputer with tens of thousands of processors, GenCast takes only eight minutes on a single Google Cloud TPU, a chip designed for machine learning. Details are published in Nature.

Google has released a string of AI-powered weather forecasts in recent years, the fruits of researchers dabbling with different approaches. In July, the firm announced NeuralGCM, which combines AI and traditional physics for long range forecasts and climate modelling.

In 2023, Google DeepMind unveiled GraphCast, which produces one single best-guess forecast at a time. GenCast builds on GraphCast by generating an ensemble of 50 or more forecasts, assigning probabilities for different weather events ahead.

Weather forecasters welcomed the advance. Steven Ramsdale, a Met Office chief forecaster with responsibility for AI, said the work was “exciting”, while a spokesperson for the ECMWF called it “a significant advance”, adding that components of GenCast were being used in one of its AI forecasts.

“Weather forecasting is on the brink of a fundamental shift in methodology,” said Sarah Dance, professor of data assimilation at the University of Reading.

“This opens up the possibility for national weather services to produce much larger ensembles of forecasts, providing more reliable estimates of forecast confidence, particularly for extreme events.”

But questions remain. “The authors have not answered whether their system has the physical realism to capture the ‘butterfly effect’, the cascade of fast-growing uncertainties, which is critical for effective ensemble forecasting,” Prof Dance said.

“There is still a long way to go before machine learning approaches can completely replace physics-based forecasting,” she added.

The data GenCast trained on combines past observations with physics-based “hindcasts” that need sophisticated maths to fill gaps in historic data, she said.

“It remains to be seen whether generative machine learning can replace this step and go straight from the most recent unprocessed observations to a 15-day forecast,” Dance said.

The performance is promising, but is a “Michael Fish moment” lurking on the horizon? “Will AI forecasting be immune?” said Price. “All prediction models would have the chance of making an error and GenCast is no different.”