are much larger than neutron stars, so their moments of inertia can be more than 100,000 times larger. They also rotate much more slowly than do neutron stars, which is consistent with the observed 21-minute period. However, although thousands of white dwarfs have been observed in our Galaxy, many of which are much closer to Earth than GPM J1839−10, only one has shown even remotely comparable radio emission. That object, known as Ar Sco, has a pulsation period of two minutes, and is around 1,000 times less luminous than GPM J1839−10 (ref. 8).

Another possibility, which is perhaps less speculative, is that the source is a magnetar, an extreme form of neutron star thought to bear the Universe's strongest known magnetic fields[9,10]. Magnetars have rotation rates that are slow compared with those of radio pulsars[11] — although nowhere near as slow as that of GPM J1839−10. They can also have radio emission at least as luminous as that of GPM J1839−10 (ref. 12). But magnetars commonly undergo sudden episodes in which they emit lots of X-ray bursts for a few weeks, and then go quiet. Hurley-Walker *et al.* found no evidence for bursts in their X-ray observations of GPM J1839−10 while it was emitting radio pulses.

And although magnetars constantly emit X-rays, they typically produce radio emission that appears suddenly, at the same time as an X-ray outburst (see, for example, ref. 13), and then fades on a timescale of months. This was true of the 18-minute source, GLEAM XJ162759.5−523504.3, which faded after just three months. By contrast — and amazingly — Hurley-Walker *et al.* show that GPM J1839−10 has been emitting radiation at radio frequencies for the past three decades, much longer than any bona fide magnetar found so far.

The puzzling long-term activity of this newly recognized source constrains any models invoked to explain it. And it might have gone unnoticed, were it not for the foresight of radio astronomers who meticulously archived and made public their voluminous data, in the hope that doing so would serve scientists in the future. Radio observations are not special in this respect. Astronomical data from across the electromagnetic spectrum have long been carefully catalogued and made freely available, resulting in a multitude of discoveries akin to that reported here. In this way, astronomy had set a high standard for open science well before other fields made it a priority.

The bounty yet hidden in astronomical archives will continue to be tapped into, and will no doubt help to answer many more questions. One key issue raised by Hurley-Walker *et al.* is whether sources such as GPM J1839−10 and GLEAM XJ162759.5−523504.3 are unusual, or whether there exists a substantial population of extremely slow pulsars awaiting discovery in the Milky Way. The astronomical archive will surely be of great assistance in answering this question. Only time will tell what else lurks in these data, and what observations across many astronomical timescales will reveal.

Victoria M. Kaspi is in the Department of Physics, McGill University, Montreal, Quebec H3A 2T8, Canada, and the Trottier Space Institute, McGill University, Montreal.
e-mail: victoria.kaspi@mcgill.ca

1. Lorimer, D. R. & Kramer, M. *Handbook of Pulsar Astronomy* (Cambridge Univ. Press, 2012).
2. Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J. & Crawford, F. *Science* **318**, 777–780 (2007).
3. The CHIME/FRB Collaboration *et al. Astrophys. J. Suppl. Ser.* **257**, 59 (2021).
4. Hurley-Walker, N. *et al. Nature* **619**, 487–490 (2023).
5. Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F. & Collins, R. A. *Nature* **217**, 709–713 (1968).
6. Tan, C. M. *et al. Astrophys. J.* **866**, 54 (2018).
7. Hurley-Walker, N. *et al. Nature* **601**, 526–530 (2022).
8. Marsh, T. R. *et al. Nature* **537**, 374–377 (2016).
9. Olausen, S. A. & Kaspi, V. M. *Astrophys. J. Suppl. Ser.* **212**, 6 (2014).
10. Kaspi, V. M. & Beloborodov, A. M. *Annu. Rev. Astron. Astrophys.* **55**, 261–301 (2017).
11. The CHIME/FRB Collaboration. *Nature* **587**, 54–58 (2020).
12. Bochenek, C. D. *et al. Nature* **587**, 59–62 (2020).
13. Camilo, F. *et al. Nature* **442**, 892–895 (2006).

The author declares no competing interests.

Meteorology

The outlook for AI weather prediction

Imme Ebert-Uphoff & Kyle Hilburn

Two models demonstrate the enormous potential that artificial intelligence holds for weather prediction. But the risks involved demand that meteorologists learn to design, evaluate and interpret such systems. **See p.526 & p.533**

Daily news headlines document the use and misuse of generative artificial intelligence (AI) — a type of AI model that can produce realistic content, such as text and videos. Public enthusiasm for these models has been tempered with trepidation, and the broader discussion about this technology is mirrored by a similar one among atmospheric scientists, some of whom have started to incorporate generative AI into their weather-forecasting models. Two groups now report such models: Bi *et al.*[1] (page 533) present a model for forecasting weather up to seven days in the future, and Zhang *et al.*[2] (page 526) describe one for predicting precipitation up to three hours ahead of time. Both studies are impressive, and together they provide a timely opportunity to examine the benefits and risks of these new developments.

Conventional weather-prediction models are based on physical equations that are implemented using numerical models — an approach known as numerical weather prediction. Generative AI weather models work differently: instead of making predictions on the basis of an understanding of physics, they forecast weather patterns that are statistically plausible given historical measurements. This approach has proved so promising that it has raised the possibility of a paradigm shift, in which AI-based models could replace numerical weather prediction completely.

At the heart of a numerical weather-prediction model is the dynamical core or 'dycore', in which numerical equations encode the underlying physical constraints: conservation of momentum, mass and energy. However, these equations take a long time to solve, even with the fastest computers, and they result in predictions with a resolution of only about 28 kilometres between grid points (see go.nature.com/3cyh4ck), which is too coarse to model small-scale physical processes, such as clouds, radiation and turbulence.

This problem can be circumvented by expressing the state of the physical system as a parameter, or set of parameters, but this replacement introduces a source of forecast error. An alternative approach, proposed almost two decades ago[3], is to keep the dycore, but to replace parameterizations with much faster AI models. Bi *et al.* and Zhang *et al.* have both taken an even more radical approach, by replacing the entire numerical weather-prediction system with an AI model. Bi and colleagues' AI model is trained entirely on observations, whereas Zhang and co-workers' AI model is trained on both physical equations and observations.

Bi and colleagues' model is called Pangu-Weather, and it forecasts temperature, wind speed and pressure, as well as other variables. The model produces predictions about 10,000 times faster than numerical weather-prediction models at the same

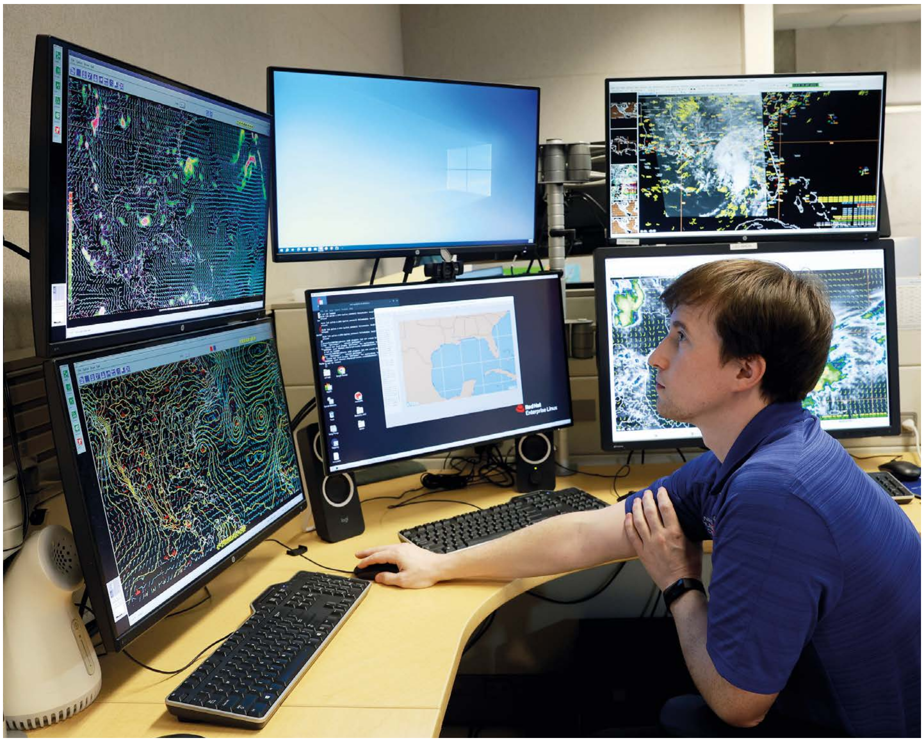

JOE RAEDLE/GETTY

Figure 1 | Weather forecasters must vet artificial-intelligence weather-prediction models. There is a pressing need for scientists to learn how to use and evaluate weather-prediction models that use artificial intelligence to ensure that they meet the needs of forecasters and thus contribute to — and not endanger — public safety.

spatial resolution, and with comparable accuracy. Pangu-Weather provides forecasts for a larger number of discrete height levels above Earth's surface than do its AI predecessors, such as FourCastNet[4]. It also uses a 3D model to ensure that predictions are consistent between these levels, and to reliably capture atmospheric states at different pressures, thereby improving accuracy.

Pangu-Weather is adept at generating medium-range forecasts, but the model does not attempt to predict precipitation, which is the most difficult weather variable to forecast[5]. This challenge is taken up by Zhang and colleagues, with their short-term forecasting model to predict rain on a timescale of hours. The model, known as NowcastNet, focuses exclusively on this task, and succeeds in producing sharper and more realistic meteorological features than is possible with its AI-based predecessors, such as PredRNN[6].

In principle, increases in computational speed, such as those reported by Bi and colleagues, could yield immense benefits. Agencies responsible for numerical weather predictions have limited budgets for computing resources. Being able to do more with less will allow these agencies to address forecast priorities that are currently out of reach, such as fire spread, atmospheric chemistry and smoke patterns, and vegetation changes. Speed increases could also lead to higher-resolution models, and to an expansion of the use of global models in place of regional ones to reduce the impact of numerical errors arising at the boundaries between regions. They could allow forecasters to generate large ensembles of predictions that represent a range of future weather possibilities, and to integrate physical processes (such as the spread of fire) that have pronounced effects on air quality and human health, but take a long time to run on standard computers.

However, AI also presents potential risks for both nowcasting and global weather predictions. Three of these risks pertain to extreme events, which are more likely to occur in a changing climate. First, depending on the duration of the data records used to train the AI model, extreme events, such as 'monster storms' that currently occur only a few times a century, might be undersampled. Second, AI models for weather forecasting are typically optimized by taking locally accurate error measurements and averaging them across large regions. This could lead to problems in predicting meteorological features, such as severe storms, fronts or tropical cyclones. And third, the behaviour of an AI system is often unpredictable when the program operates under conditions that it has never encountered before[7]. An extreme weather event might therefore trigger highly erratic predictions.

Other issues are more technical in nature. In building models that predict several variables, such as Pangu-Weather, researchers must take extra precautions to consider dependencies between those variables. Numerical weather-prediction models have these dependencies built in, but AI models do not. Furthermore, many AI models are still only proof-of-concept and do not include all the variables that a forecaster would want to see, such as precipitation type — for example, rain, hail or snow — or the physical factors involved in precipitation[8]. Finally, complex models, such as Pangu-Weather, require substantial computational resources, and only large companies can currently afford to develop them.

Given both the potential benefits and risks associated with AI models for weather prediction, we would like to issue a call to action. Now is the time for weather forecasters (Fig. 1) to get involved in ensuring that AI-based weather-prediction models are well suited to their tasks, and to learn how to interpret their predictions. This point is crucial, because AI models behave differently from models based on physics, so understanding their predictions requires specialized training. And although such complex models are not trivial to develop, they can be run easily on standard computers.

For example, we are using code that has been made available by the developers of FourCastNet (https://github.com/NVlabs/FourCastNet)[4] to generate real-time forecasts that can be compared with numerical weather-prediction forecasts and satellite observations to provide feedback to the developers. A key requirement of such initiatives is that publications are accompanied by easy-to-run code. On this basis, we think that journal editors should mandate the availability of such code. We hope that researchers will take advantage of the access to the Pangu-Weather and NowcastNet code to further evaluate the models, provide feedback and help meteorologists to decide on the appropriate use of these models with public safety in mind.

Imme Ebert-Uphoff and **Kyle Hilburn** are at the Cooperative Institute for Research in the Atmosphere, Colorado State University, Fort Collins, Colorado 80521, USA.
e-mails: iebert@colostate.edu; kyle.hilburn@colostate.edu

1. Bi, K. *et al. Nature* **619**, 533–538 (2023).
2. Zhang, Y. *et al. Nature* **619**, 526–532 (2023).
3. Krasnopolsky, V. M. & Fox-Rabinovitz, M. S. *Neural Networks* **19**, 122–134 (2006).
4. Pathak, J. *et al.* Preprint at https://arxiv.org/abs/2202.11214 (2022).
5. Sun, J. *et al. Bull. Am. Meteor. Soc.* **95**, 409–426 (2014).
6. Wang, Y. *et al. IEEE Trans. Pattern Anal. Mach. Intell.* **45**, 2208–2225 (2022).
7. Buiten, M. C. *Eur. J. Risk Regul.* **10**, 41–59 (2019).
8. Doswell, C. A. III, Brooks, H. E. & Maddox, R. A. *Weather Forecast.* **11**, 560–581 (1996).

The authors declare no competing interests.
This article was published online on 5 July 2023.