mouse brain for areas that both receive neuronal projections from the rACC and showed neural activity during the expectation of analgesia. To do this, they used a genetic technique called targeted recombination in active populations (TRAP) to identify expression of the gene *Fos*, which occurs after neurons have been active. Three brain areas were identified: the striatum, the thalamic and subthalamic nuclei, and, to the researchers' surprise, a pair of structures in the brainstem called the pontine nuclei (Pn).

A causal role of the rACC-to-Pn pathway in placebo was established, using calcium imaging (to examine neuronal activity in awake behaving mice), electrophysiology (to record neuronal firing in brain slices) and artificial 'optogenetic' activation and inhibition of the rACC-to-Pn circuit, using light-controlled proteins. The researchers also used single-cell RNA sequencing approaches to look at gene expression, and showed that the relevant Pn neurons were excitatory and expressed the gene encoding δ-opioid receptors, supporting the idea that placebo analgesia is indeed opioid-mediated.

The Pn act as a liaison between the cerebral cortex and the cerebellum. Although the Pn have been observed to be activated by pain[8], they are not recognized as part of the network of brain areas that typically respond to and process pain — referred to as the pain matrix. Because the cerebellum is the main target of Pn neurons, and because an insistent reviewer spurred the work on, the authors next examined the activity of the main neurons of the cerebellar cortex — Purkinje cells — in mice experiencing expectation of analgesia. Chen *et al.* identified a particular cluster of Purkinje cells that encode the expectation of pain relief, and found that the activity of this cluster was driven by the rACC-to-Pn circuit (Fig. 1b).

This is hardly the first animal model of placebo analgesia; nor is it the first use of conditioning to study pain[9]. But Chen *et al.* have certainly delivered the most in-depth neural explanation of placebo analgesia so far, using essentially all of the modern, high-resolution techniques currently available for defining neural circuits in mice. Both the placebo effect and its evil twin 'nocebo' — in which an individual's negative expectation of treatment causes their symptoms to worsen — are hugely important for such explication owing to their ubiquity and robustness in the mediation of disease and treatment. It is worth noting that the placebo effect is probably more complicated in humans than in mice because, in humans, it involves expectation that is based on verbal instruction and accompanying social influences, in addition to conditioning.

The really interesting finding here is that, of all possible brain areas, it is the Pn and the cerebellum that are responsible for producing expectation, a concept that might be assumed to be mediated by much 'higher' levels of the central nervous system that coordinate complex cognitive functions. This demonstration, along with other recent circuit-tracing efforts[10], could lead to a re-evaluation of pain neuroanatomy.

Jeffrey S. Mogil is at the Departments of Psychology and Anesthesia, Faculties of Science, Medicine, and Dental Medicine, and Alan Edwards Centre for Research on Pain, McGill University, Montreal, QC H3A 2B1, Canada.
e-mail: jeffrey.mogil@mcgill.ca

1. Chen, C. *et al. Nature* **632**, 1092–1100 (2024).
2. Hróbjartsson, A. & Gøtzsche, P. C. *N. Engl. J. Med.* **344**, 1594–1602 (2001).
3. Colloca, L., Noel, J., Franklin, P. D. & Senevratne, C. (eds) *Placebo Effects: Through the Lens of Translational Research* (Oxford Univ. Press, 2023).
4. Levine, J. D., Gordon, N. C. & Fields, H. L. *Lancet* **312**, 654–657 (1978).
5. Zunhammer, M., Spisák, T., Wager, T. D., Bingel, U. & The Placebo Imaging Consortium. *Nature Commun.* **12**, 1391 (2021).
6. Li, C. N., Keay, K. A., Henderson, L. A. & Mychasiuk, R. *J. Neurosci.* **44**, e1538232024 (2024).
7. Petrovic, P., Kalso, E., Petersson, K. M. & Ingvar, M. *Science* **295**, 1737–1740 (2002).
8. Sclocco, R. *et al. Phil. Trans. R. Soc. A* **374**, 20150189 (2016).
9. Sufka, K. J. *Pain* **58**, 355–366 (1994).
10. Gan, Z. *et al. Science* **378**, 1336–1343 (2022).

The author declares no competing interests.
This article was published online on 21 August 2024.

Atmospheric science

Weather and climate predicted by hybrid model

Oliver Watt-Meyer

A cutting-edge model of the atmosphere combines machine learning with a numerical model based on the laws of physics. This 'hybrid' system accurately predicts the weather — and even shows promise for climate simulations. **See p.1060**

Numerical models of Earth's climate system, based on the laws of physics, are essential tools for predicting changes in weather and climate in our warming world. However, persistent uncertainties in the representation of small-scale processes, such as cloud formation, have resulted in only slow improvements in the fidelity of these models. Interest in using machine learning to provide a data-driven representation of these uncertain processes has therefore skyrocketed. Yet challenges have persisted in coupling these representations to conventional atmospheric models. On page 1060, Kochkov *et al.*[1] report a hybrid global atmospheric model that overcomes these challenges, combining physics-based approaches with machine learning. This model not only predicts the weather up to 15 days ahead with state-of-the-art accuracy, but also simulates historical temperature trends over a 40-year period, demonstrating its ability to assess climate.

Conventional models of the atmosphere separate the representation of large-scale motions from that of processes occurring on smaller scales, such as turbulence and the formation of clouds. Large-scale motions exceed the horizontal resolution of the 3D grid that represents Earth in the model, which is typically in the range of 10–100 kilometres. These large-scale motions have been well characterized, and can be described by equations that are approximated using standard numerical techniques in a part of the model code known as the dynamical core, resulting in representations that have low uncertainties.

However, representations of smaller, subgrid-scale processes are much more uncertain, and so empirical estimates (known as parameterizations) of the effects of these processes are used instead. In the past five years, concerted efforts have been made to use machine learning to systematically predict such effects[2,3]. Progress has been made, but breakthroughs have been slower to arrive than in the rapidly advancing field of machine learning for weather prediction[4,5].

Kochkov *et al.* have taken a considerable step forward in the use of machine learning for parameterization. As in previous works[2,3], an artificial neural network is used to represent the effects of subgrid-scale processes. The key technical advance is that Kochkov and colleagues have written a dynamical core in JAX, a programming language that enables automatic differentiation[6] — that is, it allows the gradients of outputs with respect to model parameters to be calculated accurately and efficiently. This allows the optimization of the neural network to account for the all-important feedbacks between the large-scale, resolved atmospheric dynamics and the

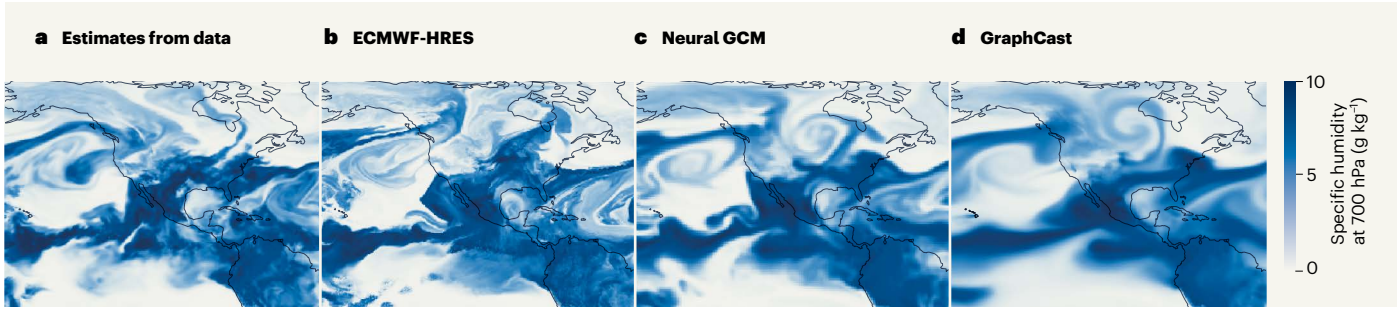


Figure 1 | Comparison of a hybrid atmospheric model with conventional and machine-learning models. Kochkov *et al.*[1] report NeuralGCM, a general circulation model of the atmosphere that combines machine learning with a numerical model of atmospheric dynamics. This hybrid model accurately forecasts weather and shows promise for predicting climate. In this example, estimates of the true specific humidity, based on observational data over North America and the northeast Pacific Ocean on 1 September 2020 (**a**), are compared with 10-day forecasts made by ECMWF-HRES (**b**, a purely numerical weather model), NeuralGCM (**c**) and GraphCast (**d**, a weather model that uses machine learning alone). Neural GCM provides a less blurry forecast than GraphCast, and has accuracy comparable to that of ECMWF-HRES, yet requires much less computational power than does ECMWF-HRES. hPa, hectopascal.

unresolved subgrid processes.

When the authors trained their hybrid system using past observations of atmospheric behaviour, they found that it generated highly accurate and well-calibrated ensemble weather forecasts. These are collections of forecasts produced by making slight variations to the starting conditions and to the model itself, to provide a range of possible future weather conditions. The ensemble forecasts were for up to 15 days into the future, and throughout most of this period they were more accurate than the leading physics-based ensemble forecasting system. Individual forecasts (rather than ensemble predictions) also compare favourably with those produced by purely numerical or purely machine-learning models (Fig. 1).

The hybrid model could also run longer simulations (40-year projections) to estimate climate statistics, although some of these simulations (15 out of 37) were not stable throughout the 40-year period — a problematic feature that is not typical of conventional models. However, the stable simulations exhibited time-averaged errors that were small compared with those produced by conventional models, when evaluated against the historical record. Furthermore, the hybrid system makes its predictions efficiently: running on a single tensor processing unit (a specialized chip used for machine learning), it can simulate about 200 years per day at a horizontal resolution of about 150 kilometres. This is much faster than conventional models of a similar resolution can manage.

One of the main applications for climate models is the consideration of hypothetical scenarios, such as what would have happened to the ozone hole and climate in a world in which emissions of ozone-depleting substances were not stopped[7], and what will happen in the future as greenhouse-gas concentrations rise or fall[8]. However, a fundamental limitation of machine learning is that it cannot make accurate predictions when scenarios involve ranges for variables that far exceed the ranges used for training. Indeed, when perturbed with a strong, uniform sea surface warming of 4 °C, the model developed by Kochkov and colleagues predicts unrealistic changes in atmospheric temperature — disagreeing with the expectations of fundamental physics and with conventional atmospheric models.

How can this issue be resolved? One approach is to rescale variables in such a way that their distributions alter less with climate change — for example, by using relative humidity instead of specific humidity[9]. But not all variables are amenable to this type of rescaling. It is more likely that, to ensure robust predictions, machine-learning models will need to be trained on the output from existing climate models, which can generate training data across a wide range of climate states[10].

"Running on a single tensor processing unit, the system is much faster than conventional models of a similar resolution."

The promise of machine learning for weather prediction is clearly being realized. But what about its promise for climate modelling and climate prediction? As I noted earlier, generalizing machine-learning models to all climate scenarios is a key challenge. More broadly, the complexity of the Earth system, with its interacting realms (such as the land, the ocean and the atmosphere) is daunting. Conventional climate models have tackled this problem by incorporating more physical processes, such as those that couple the atmosphere to the ocean, and carbon-cycle processes that involve interactions between realms[11]. Machine-learning climate models will probably follow a similar, but accelerated, path — although the limited amount of data available for realms other than the atmosphere might cause difficulty. Nevertheless, a long-duration, coupled atmosphere–ocean simulation based on machine-learned components will probably soon become available. This will enable the exploration of phenomena such as the El Niño–Southern Oscillation, specifically of its predictability and variability.

As machine-learning climate models rapidly advance, scientists should give them a spin. Many are already publicly available (including Neural GCM; see go.nature.com/3m2m9x6) and are easier to install and run than conventional climate models, which often require the use of particular supercomputers. Only when climate scientists engage fully with these models and their developers will machine-learning methods reach their maximum potential.

Oliver Watt-Meyer is at the Allen Institute for Artificial Intelligence, Seattle, Washington 98103, USA.
e-mail: oliverwm@allenai.org

1. Kochkov, D. *et al. Nature* **632**, 1060–1066 (2024).
2. Rasp, S., Pritchard, M. S. & Gentine, P. *Proc. Natl Acad. Sci. USA* **115**, 9684–9689 (2018).
3. Brenowitz, N. D. & Bretherton, C. S. *Geophys. Res. Lett.* **45**, 6289–6298 (2018).
4. Bi, K. *et al. Nature* **619**, 533–538 (2023).
5. Lam, R. *et al. Science* **382**, 1416–1421 (2023).
6. Frostig, R., Johnson, M. J. & Leary, C. In *Proc. Syst. Mach. Learn.* abstr. 146 (2018).
7. Banerjee, A., Fyfe, J. C., Polvani, L. M., Waugh, D. & Chang, K.-L. *Nature* **579**, 544–548 (2020).
8. IPCC. *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (eds Masson-Delmotte, V. *et al.*) (Cambridge Univ. Press, 2023).
9. Beucler, T. *et al. Sci. Adv.* **10**, eadj7250 (2024).
10. Clark, S. K. *et al. J. Adv. Model. Earth Syst.* **14**, e2022MS003219 (2022).
11. Eyring, V. *et al. Geosci. Model Dev.* **9**, 1937–1958 (2016).

The author declares no competing interests.
This article was published online on 13 August 2024.