

Aurora: A Foundation Model for the Earth System


Wessel Bruinsma

The Alan Turing Institute

Work was done at Microsoft Research

UCL AI Centre, London, 24 July 2025




The
Alan Turing
Institute




The AI Revolution in Science



AlphaFold
Protein folding




The AI Revolution in Weather Forecasting



2018

First serious efforts to compare AI models to physics baselines


Dueben and Bauer (2018)



2019

AI models skillful to multiple days

Weyn et al. (2019)




2020

WeatherBench starts to drive ML development

Rasp et al. (2020)

The AI Revolution in Weather Forecasting



2022

GNN outperforms GFS at 1°
Keisler (2022)



2022


Pangu-Weather outperforms HRES at 0.25°
Bi et al. (2023)

The AI Revolution in Weather Forecasting



2022–2023

Tech companies start to work in this space



2023

GenCast outperforms IFS ensemble


Price et al. (2024)




2024

ECMWF launches AIFS

What About Other Forecasting Tasks?




Air
pollution




Atmospheric
composition



Waves




Sea ice



Ocean

- Current models are impressive, but **limited to one setting**.
- Unified approach?

Aurora



pretraining

- Train a single neural network a *large* body of Earth system data
- Learn general-purpose representation of dynamics that govern atmospheric and oceanic flow
- Slow and data hungry

fine-tuning


- Leverage learned representation to **efficiently adapt to new domains!**
- Fast and data efficient

Aurora: a **foundation model** for the Earth system

The Model



The Model



- Predict global state of **any variables** at **any resolution** 6 h ahead:


$$\hat{X}^{t+6\text{ h}} = \Phi(X^t, X^{t-6\text{ h}}),$$

$$\hat{X}^{t+12\text{ h}} = \Phi(\hat{X}^{t+6\text{ h}}, X^t),$$



$$\hat{X}^{t+18\text{ h}} = \Phi(\hat{X}^{t+12\text{ h}}, \hat{X}^{t+6\text{ h}}),$$

⋮



- Transformer-based encoder–decoder architecture:





Encoder



Backbone




Decoder



The background of the slide features a dramatic landscape under a dark, cloudy sky. In the foreground, there's a green field with some small white structures or trees visible. The overall mood is somber and atmospheric.

Pretraining



Pretraining




Name	Resolution	Timeframe	Surf. variables	Atmos. variables	Levels	Steps	Size
ERA5	0.25° × 0.25°	1979–2020	2T, 10U, 10V, MSL	U, V, T, Q, Z	13	368.18 k	105.50 TB
HRES-0.25 forecasts	0.25° × 0.25°	2016–2020	2T, 10U, 10V, MSL	U, V, T, Q, Z	13	149.81 k	42.93 TB
IFS-ENS-0.25	0.25° × 0.25°	2018–2020	2T, 10U, 10V, MSL	U, V, T, Q, Z	3	6.69 M	527.54 TB
IFS-ENS-0.25 mean	0.25° × 0.25°	2018–2020	2T, 10U, 10V, MSL	U, V, T, Q, Z	3	133.71 k	10.55 TB
GFS forecasts	0.25° × 0.25°	Feb 2015–2020	2T, 10U, 10V, MSL	U, V, T, Q, Z	13	354.40 k	101.56 TB
GFS T0	0.25° × 0.25°	Feb 2015–2020	2T, 10U, 10V, MSL	U, V, T, Q, Z	13	8.64 k	2.48 TB
GEFS reforecasts	0.25° × 0.25°	2000–2019	2T, MSL	U, V, T, Q, Z	7	2.96 M	454.61 TB
CMCC-CM2-VHR4	0.25° × 0.25°	1950–2014	2T, 10U, 10V, MSL	U, V, T, Q	7	94.96 k	12.62 TB
ECMWF-IFS-HR	0.45° × 0.45°	1950–2014	2T, 10U, 10V, MSL	U, V, T, Q, Z	7	94.96 k	4.75 TB
MERRA-2	0.625° × 0.50°	1980–2020	2T, 10U, 10V, MSL	U, V, T, Q	13	119.81 k	5.58 TB
Total						10.97 M	1,268.12 TB

- 150 000 steps on 32 GPUs (A100)
- The magic: **data scaling** and **model scaling!**


Data Scaling



Model Scaling



Economical to train big models!





The background of the image is a dramatic landscape featuring a vast, flat field in the foreground. In the middle ground, a line of trees stands against a dark, overcast sky filled with heavy, billowing clouds. The overall atmosphere is somber and moody.


Fine-Tuning

Applications

Operational in all settings!




Atmospheric comp.
and air pollution



Ocean
waves

Tropical cyclone
tracks



High-resolution
weather

Air Pollution Forecasting




- **Setup:** model PM₁, PM_{2.5}, PM₁₀, CO, NO, NO₂, SO₂, O₃
- **Data:** Copernicus Atmospheric Monitoring Service (CAMS) analysis
- **Baseline:** CAMS forecasts


Coupled to IFS, ~10x more expensive:
~16 node-hours per hour lead time!

Aurora: **~0.5 s per hour lead time**

Air Pollution Forecasting



- Heterogeneous and spiky
- Anthropogenic factors
- Scarce
- Non-stationary



Overall:


Competitive on
95%
(≤ 20% RMSE)

Better on 75%

Three days:

Competitive on
100%
(≤ 20% RMSE)


Better on 86%




Ocean Wave Forecasting




- **Setup:** model height, direction, and period of wave components
- **Data:** HRES-WAM analysis
- **Baseline:** HRES-WAM forecasts



Ocean Wave Forecasting



- Angle-valued variables
- Absence of wave components



Overall:


Competitive on
96%
(≤ 20% RMSE)

Better on 86%

Three days:

Competitive on all
but PP1D
(≤ 20% RMSE)


Better on 91%



Open Source

- All models open source under **MIT licence!**


```
1 import torch
2
3 from aurora import Aurora, Batch, rollout
4
5 model = Aurora()
6 model.load_checkpoint()
7
8 model.eval()
9 model.to("cuda")
10
11 batch = Batch(...)
12
13 with torch.inference_mode():
14     for prediction in rollout(model, batch, steps=10):
15         ... # Do something with `prediction`.
```



<https://github.com/microsoft/aurora>
pip install microsoft-aurora

The Weather Forecasting Pipeline

Aardvark-Weather



The Aurora Team



Paris Perdikaris

University of Pennsylvania,
formerly MSR



Richard Turner

U. of Cambridge, The Alan
Turing Institute, formerly MSR



Max Welling

University of Amsterdam,
CuspAI, formerly MSR



Wessel Bruinsma

The Alan Turing Institute,
formerly MSR



Anna Allen

University of Cambridge, The
Alan Turing Institute



Elizabeth Heider

Book tour, formerly MSR




Cristian Bodnar

Silurian, formerly MSR




Johannes Brandstetter

JKU Linz, Emmi AI, formerly MSR



Ana Lučić

University of Amsterdam,
formerly MSR




Patrick Garvan

IONQ, formerly MSR




Megan Stanley

MSR



Maik Riechert

MSR



Conclusion

- Medium-term weather forecasting has seen incredible progress
- **Pretraining–fine-tuning paradigm** to extend these advancements to other domains
- Aurora only scratches the surface!

 wessel.ai/pdf/aurora

 wessel.ai/pdf/aardvark

 hi@wessel.ai

Bodnar, C., Bruinsma, W.P., Lučić, A., Stanley M., Allen, A. et al. A foundation model for the Earth system. *Nature* **641**, 1180–1187 (2025).

Allen, A., Markou, S. et al. End-to-end data-driven weather prediction. *Nature* **641**, 1172–1179 (2025).