Aurora: A Foundation Model for the Earth System

Wessel Bruinsma

Microsoft Research AI for Science

Statistical Learning in Atmospheric Chemistry (SLAC) Community
Online, 31 Mar 2025

The Aurora Team


Paris Perdikaris
University of Pennsylvania, formerly MSR


Wessel Bruinsma
MSR


Cristian Bodnar
Silurian, formerly MSR


Ana Lučić
University of Amsterdam, formerly MSR


Megan Stanley
MSR


Richard Turner
University of Cambridge, formerly MSR


Anna Allen
University of Cambridge


Johannes Brandstetter
JKU Linz, NXAI, formerly MSR

Patrick Garvan
Formerly MSR

Maik Riechert
MSR


Max Welling
University of Amsterdam, CuspAI, formerly MSR


Elizabeth Heider
Book tour, formerly MSR

The AI Revolution in Science

AlphaFold
Protein folding

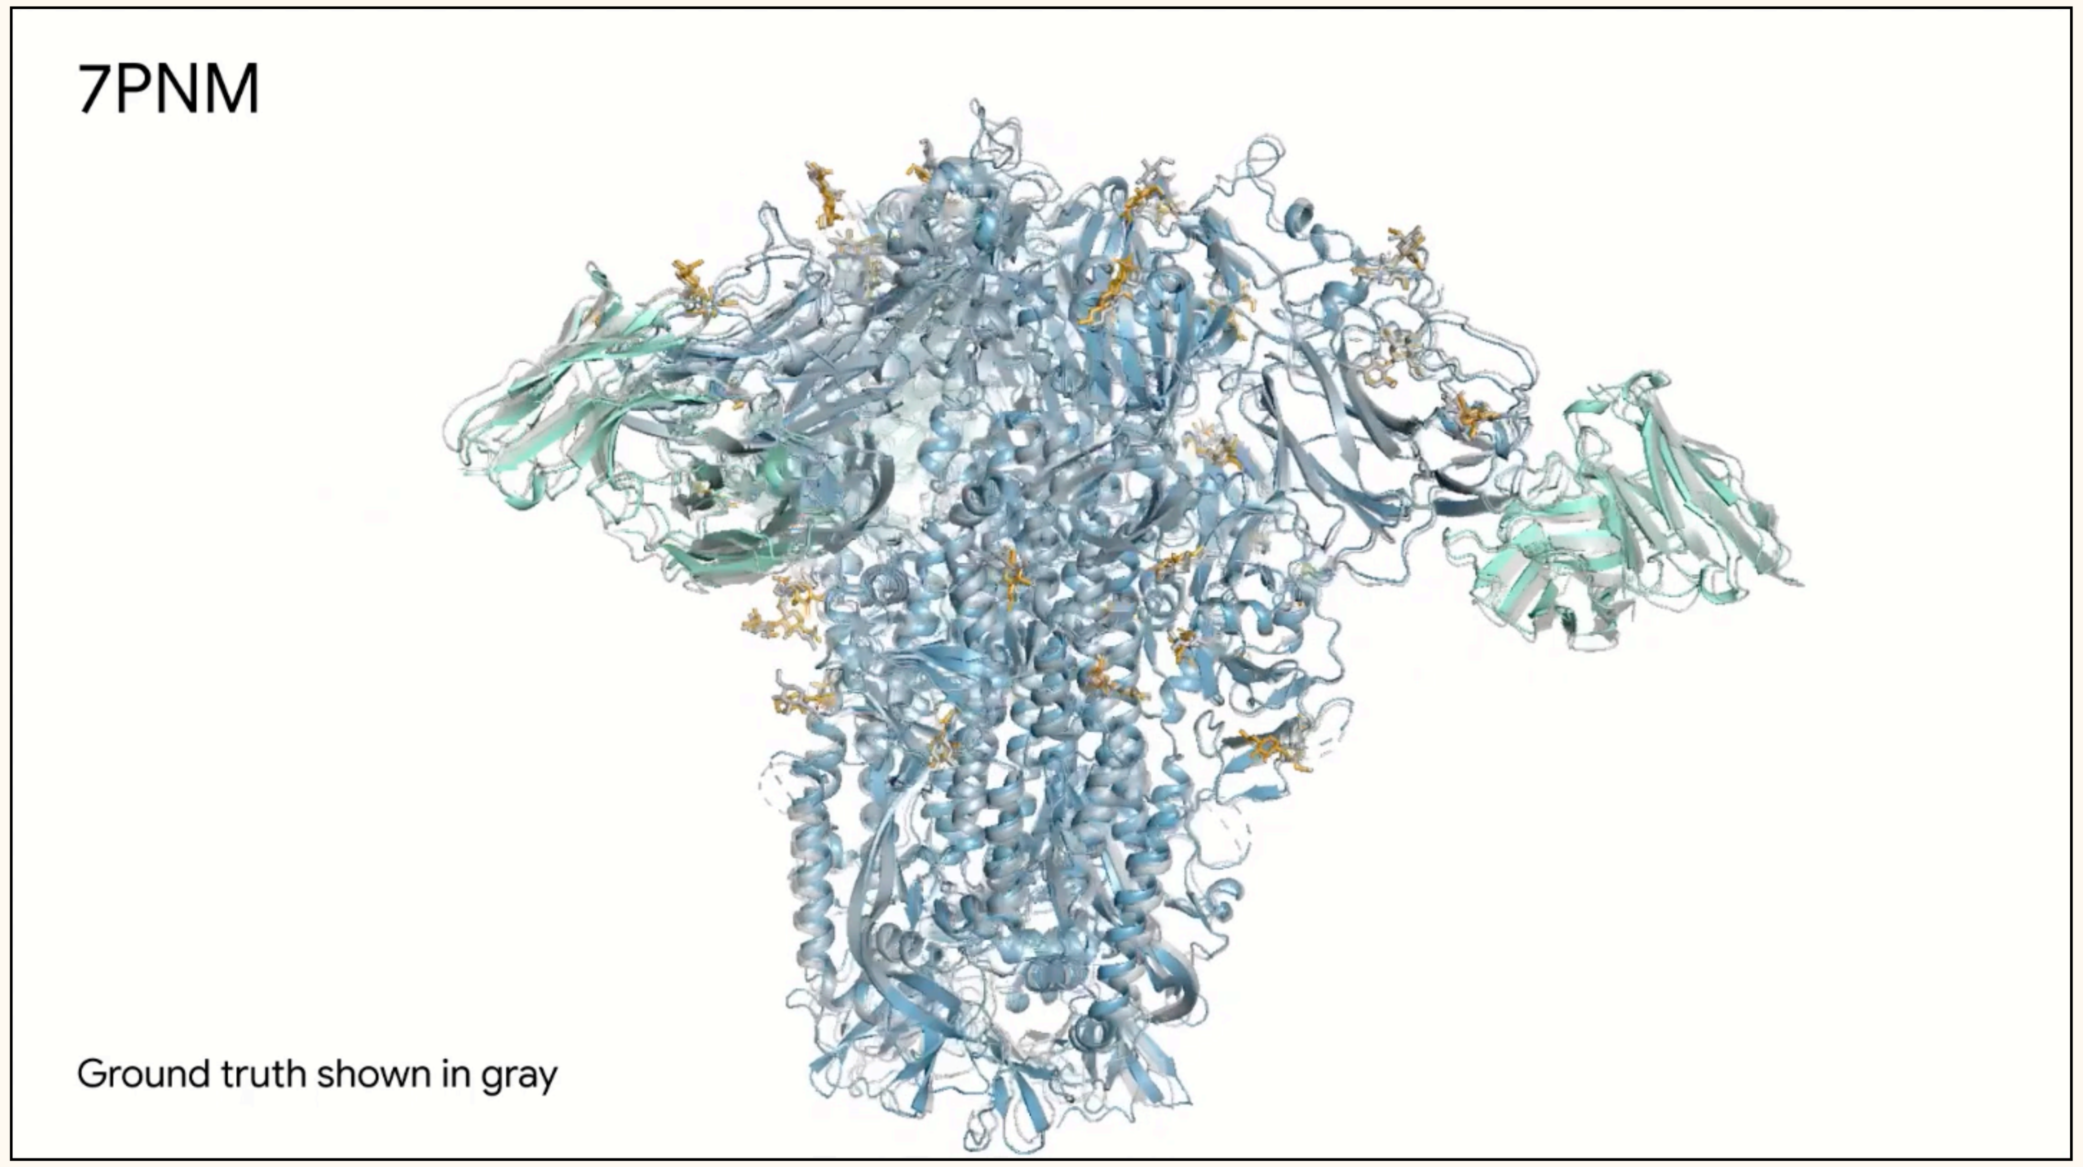

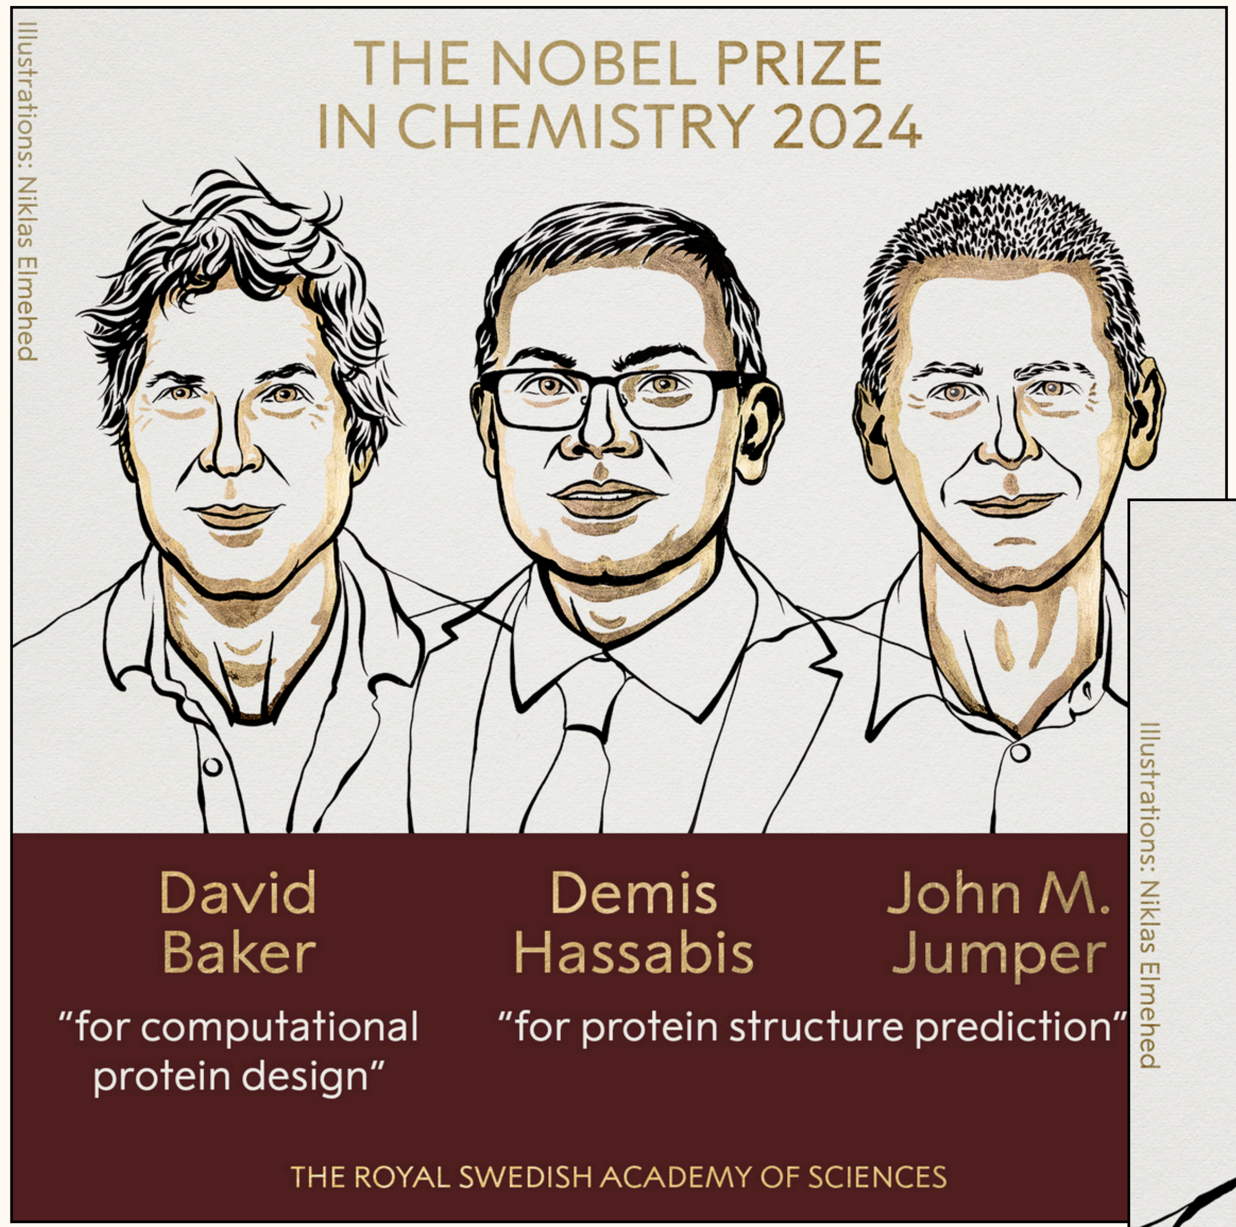

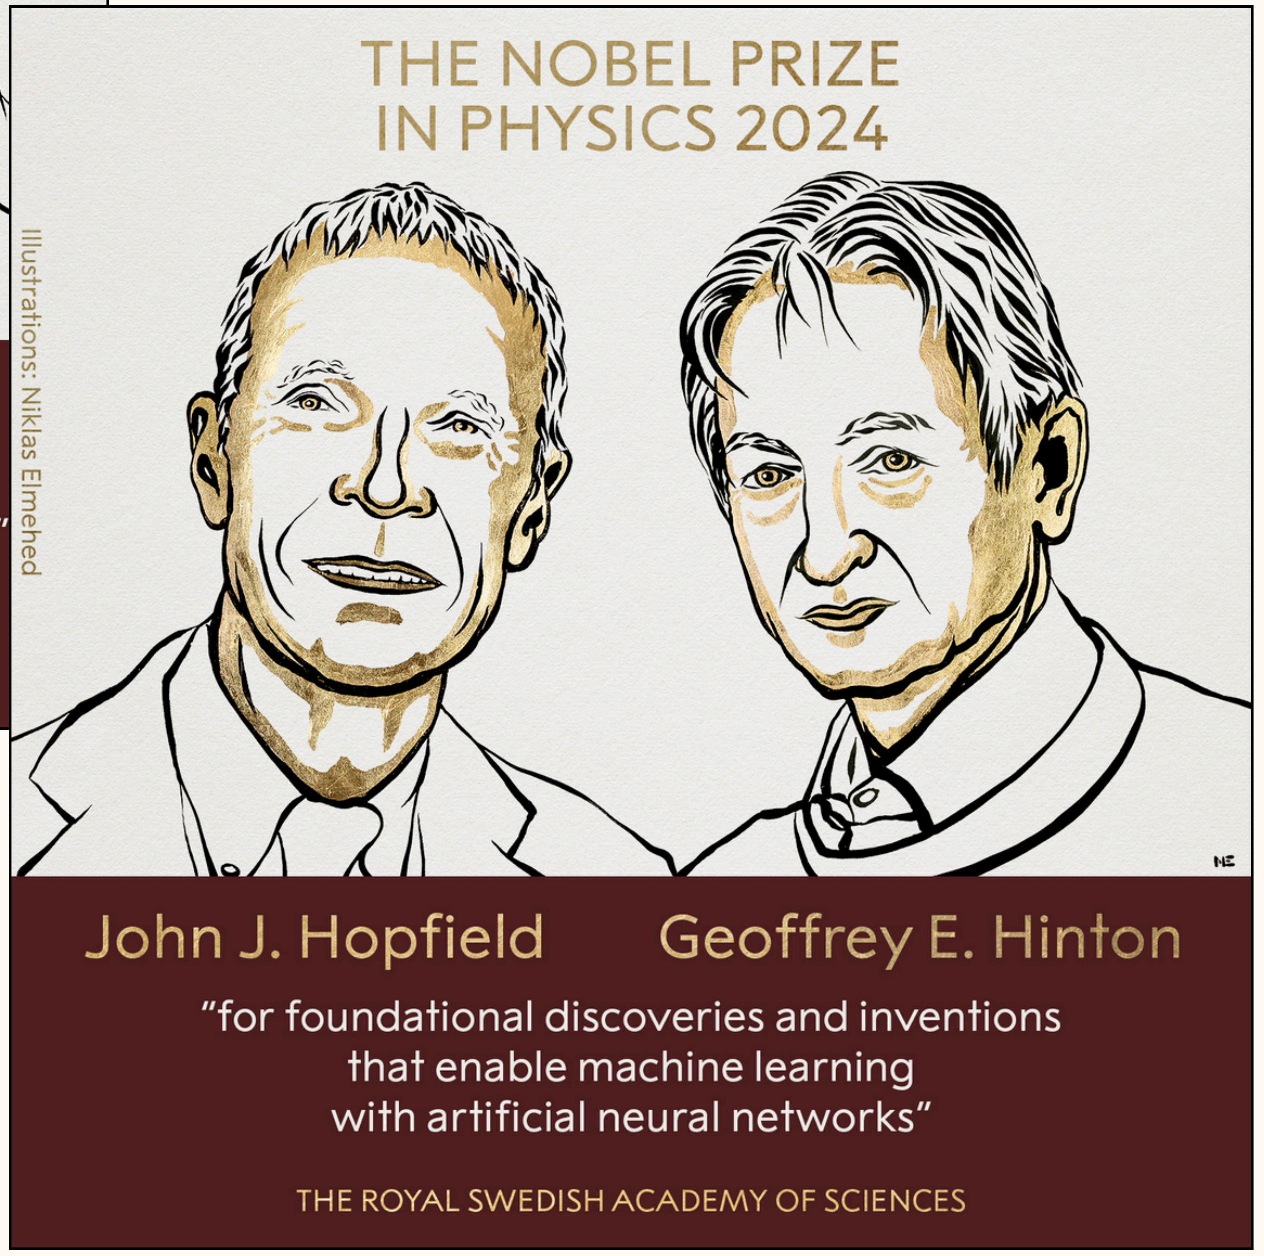

The AI Revolution in Weather Forecasting

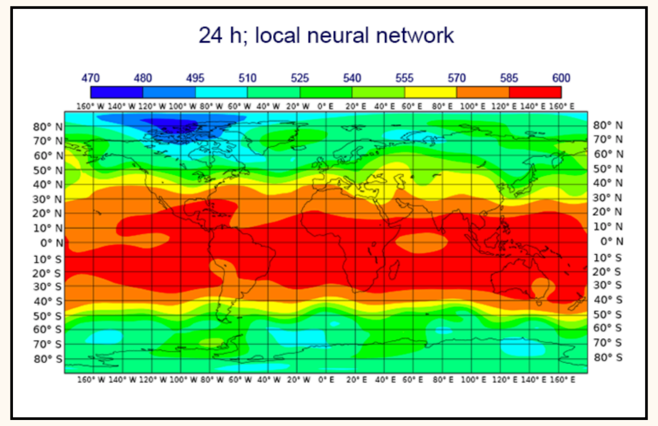

2018

First serious efforts to compare AI models to physics baselines

Dueben and Bauer (2018)

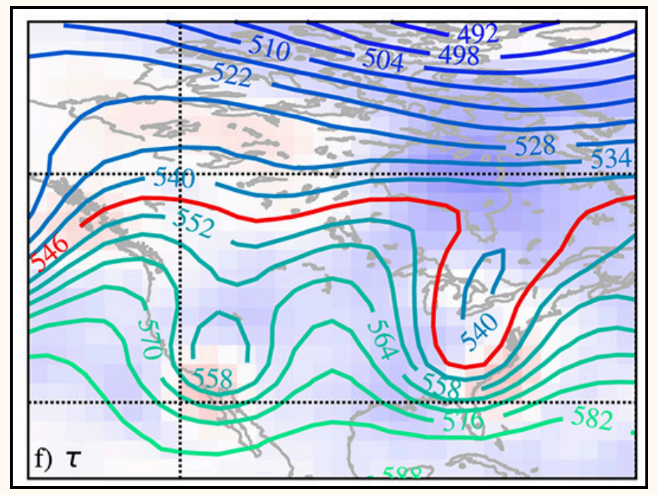

2019

AI models skillful to multiple days

Weyn et al. (2019)

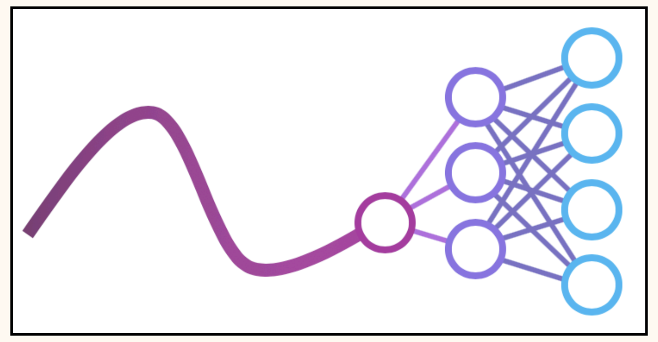

2020

WeatherBench starts to drive ML development

Rasp et al. (2018)

The AI Revolution in Weather Forecasting

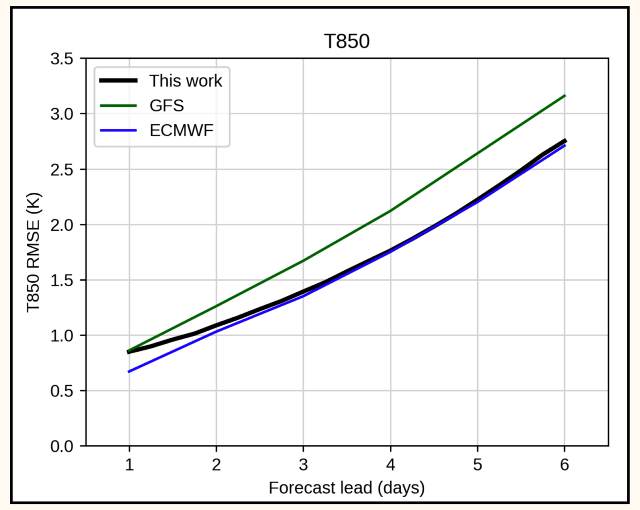

2022

GNN outperforms GFS at 1°

Keisler (2022)

HUAWEI

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

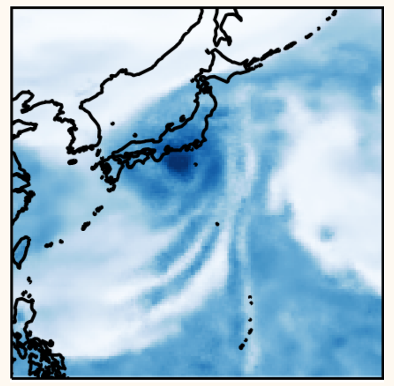

2023 GenCast outperforms IFS ensemble

Price et al. (2024)

2024 ECMWF launches AIFS

What About Other Forecasting Tasks?

Air pollution

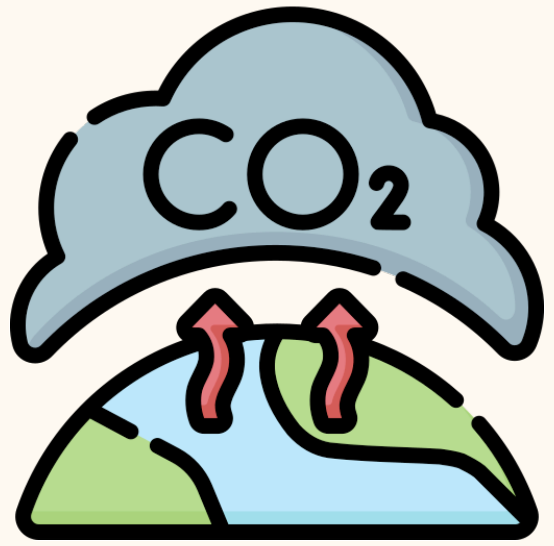

Atmospheric composition

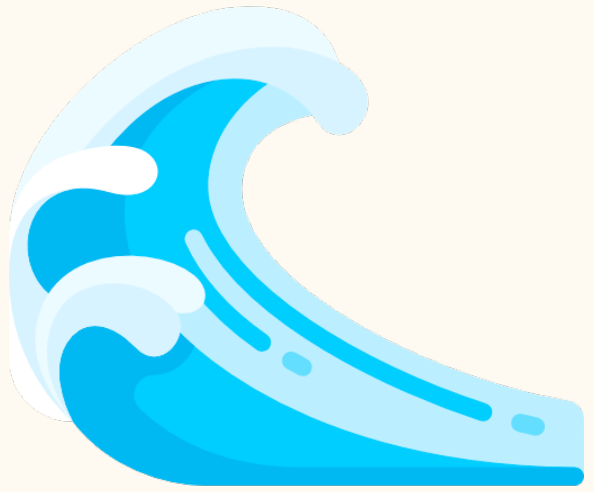

Waves

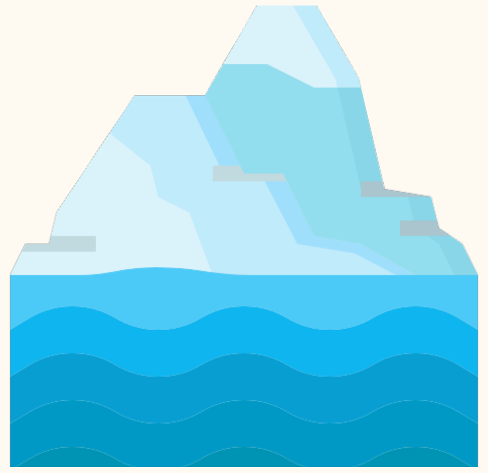

Sea ice

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

Aurora

Pretraining

Variable	Units	Description
Surface-level meteorological variables		
2T	K	Temperature at 2 m above surface of land or sea
U10	$\mathrm{m\,s^{-1}}$	Eastward component of wind at 10 m
V10	$\mathrm{m\,s^{-1}}$	Southward component of wind at 10 m
WS	$\mathrm{m\,s^{-1}}$	Wind speed at 10 m; equal to $(\mathrm{U10}^2 + \mathrm{V10}^2)^{1/2}$
MSL	Pa	Air pressure at mean sea level
Atmospheric meteorological variables		
U	$\mathrm{m\,s^{-1}}$	Eastward component of wind
V	$\mathrm{m\,s^{-1}}$	Southward component of wind
T	K	Temperature
Q	$\mathrm{kg\,kg^{-1}}$	Specific humidity
Z	$\mathrm{m^2\,s^{-2}}$	Geopotential

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

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

Cost:

- 150 000 steps
- 32 A100s
- 3 weeks

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

Aurora

Fine-Tuning Applications

Operational in all settings!

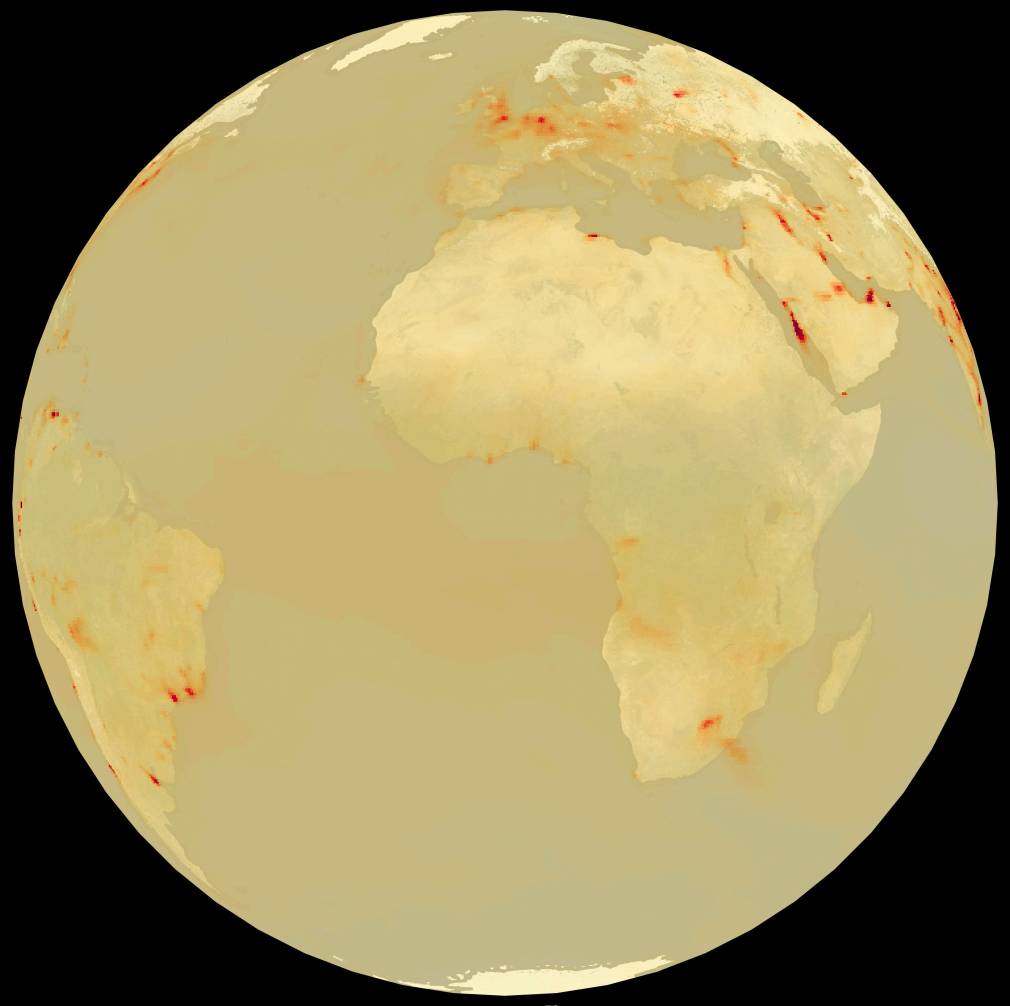

Atmospheric comp.
and air pollution

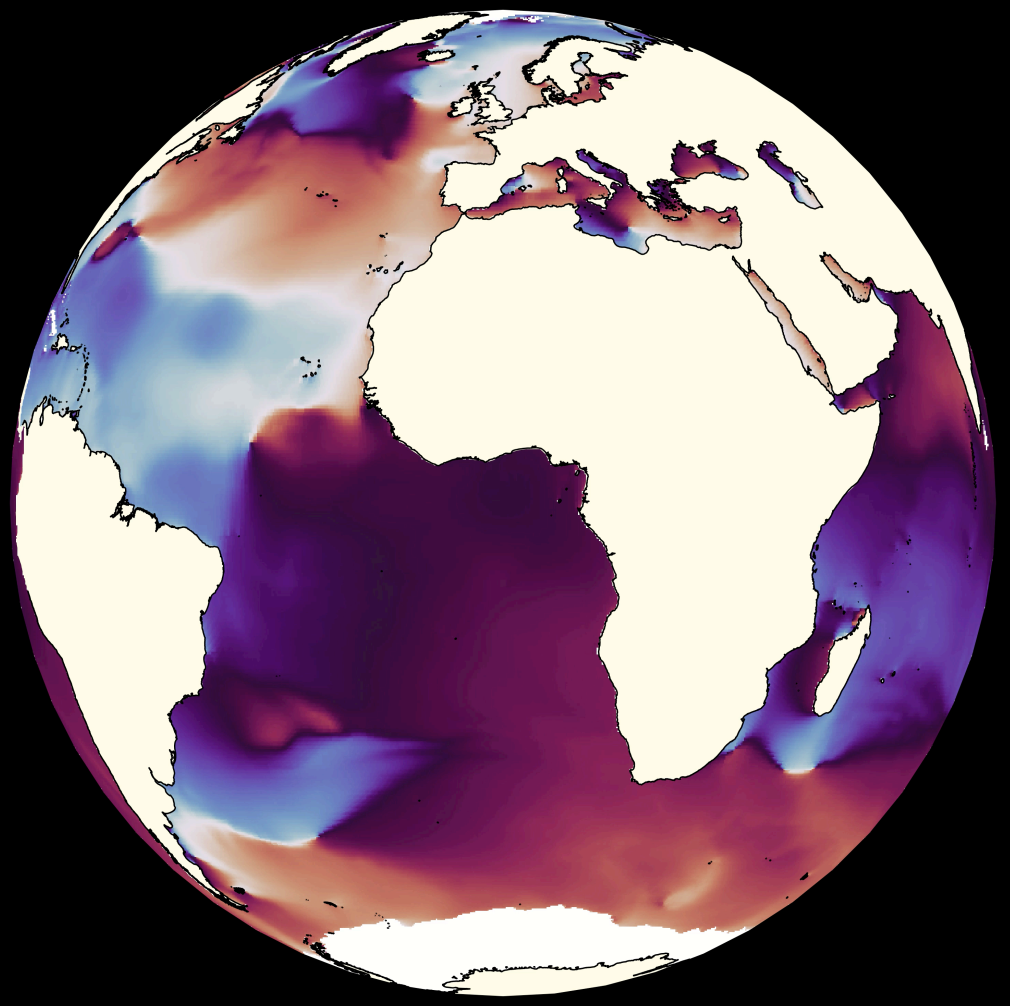

Ocean
waves

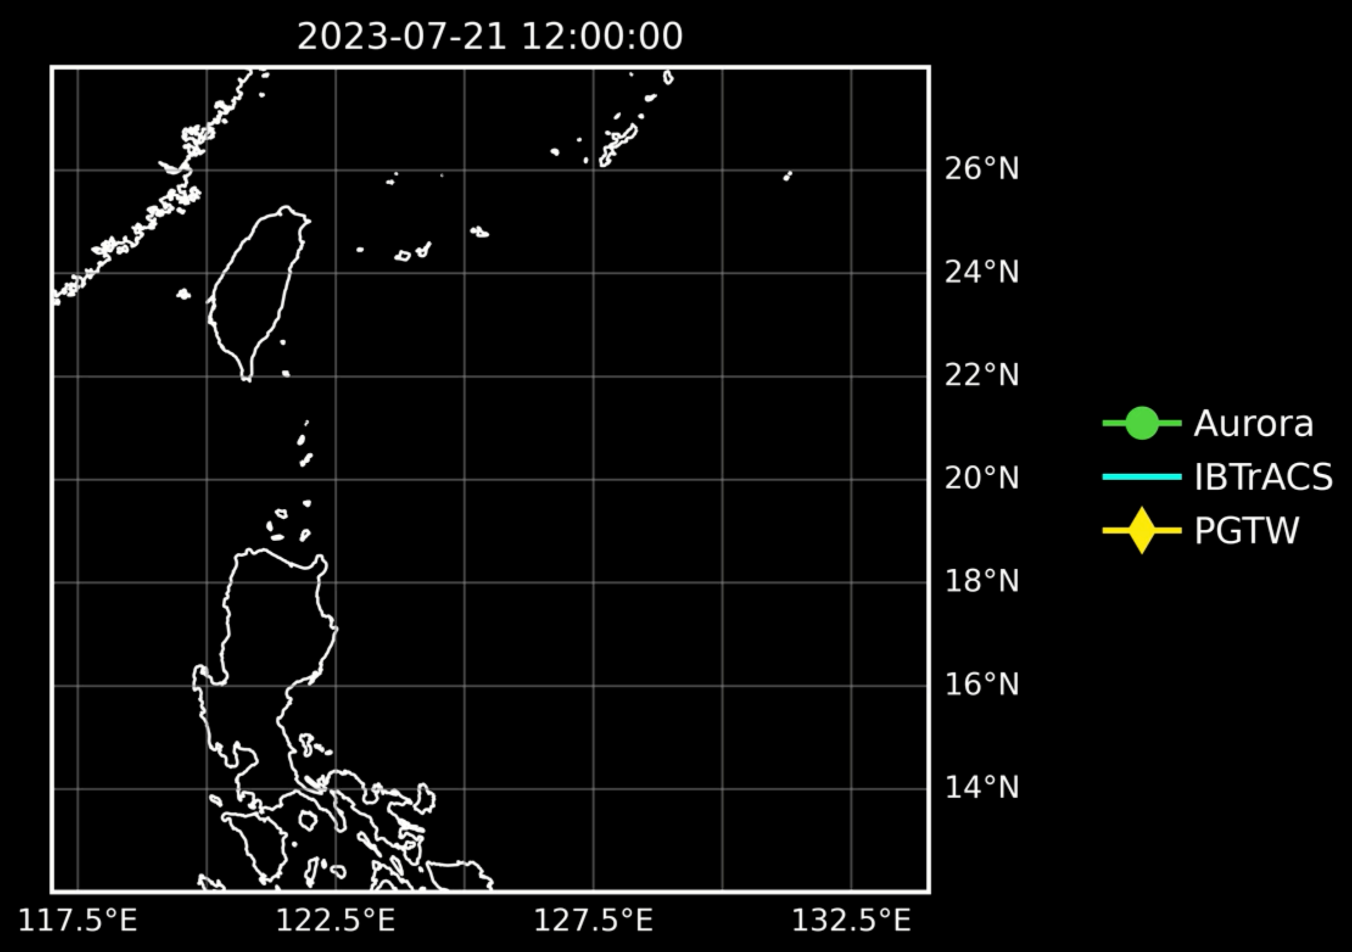

Tropical cyclone
track

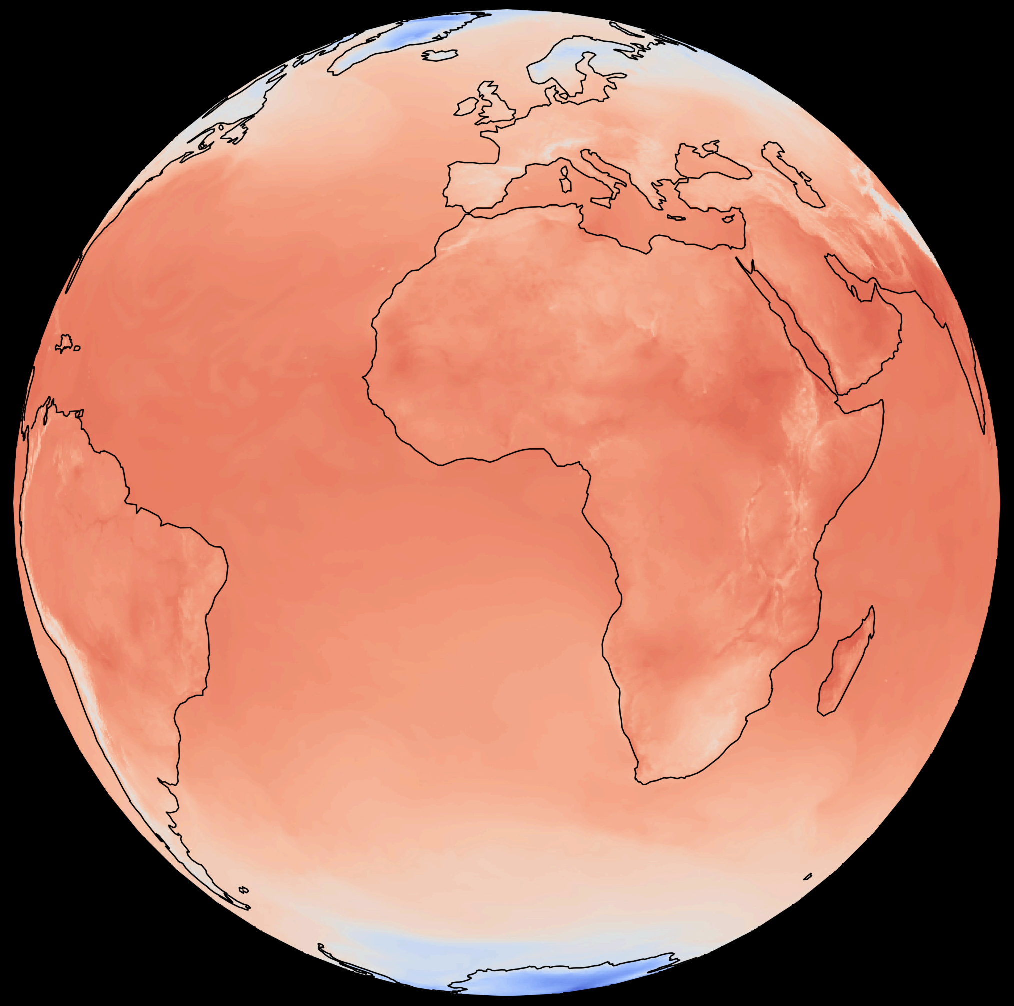

High-resolution
weather

Animations created by Megan Stanley

Aurora

Air Pollution Forecasting

- **Setup:** model levels of PM_1, $PM_{2.5}$, PM_{10}, CO, NO, NO_2, SO_2, O_3
- **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**

Aurora

Air Pollution Forecasting: Challenges

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

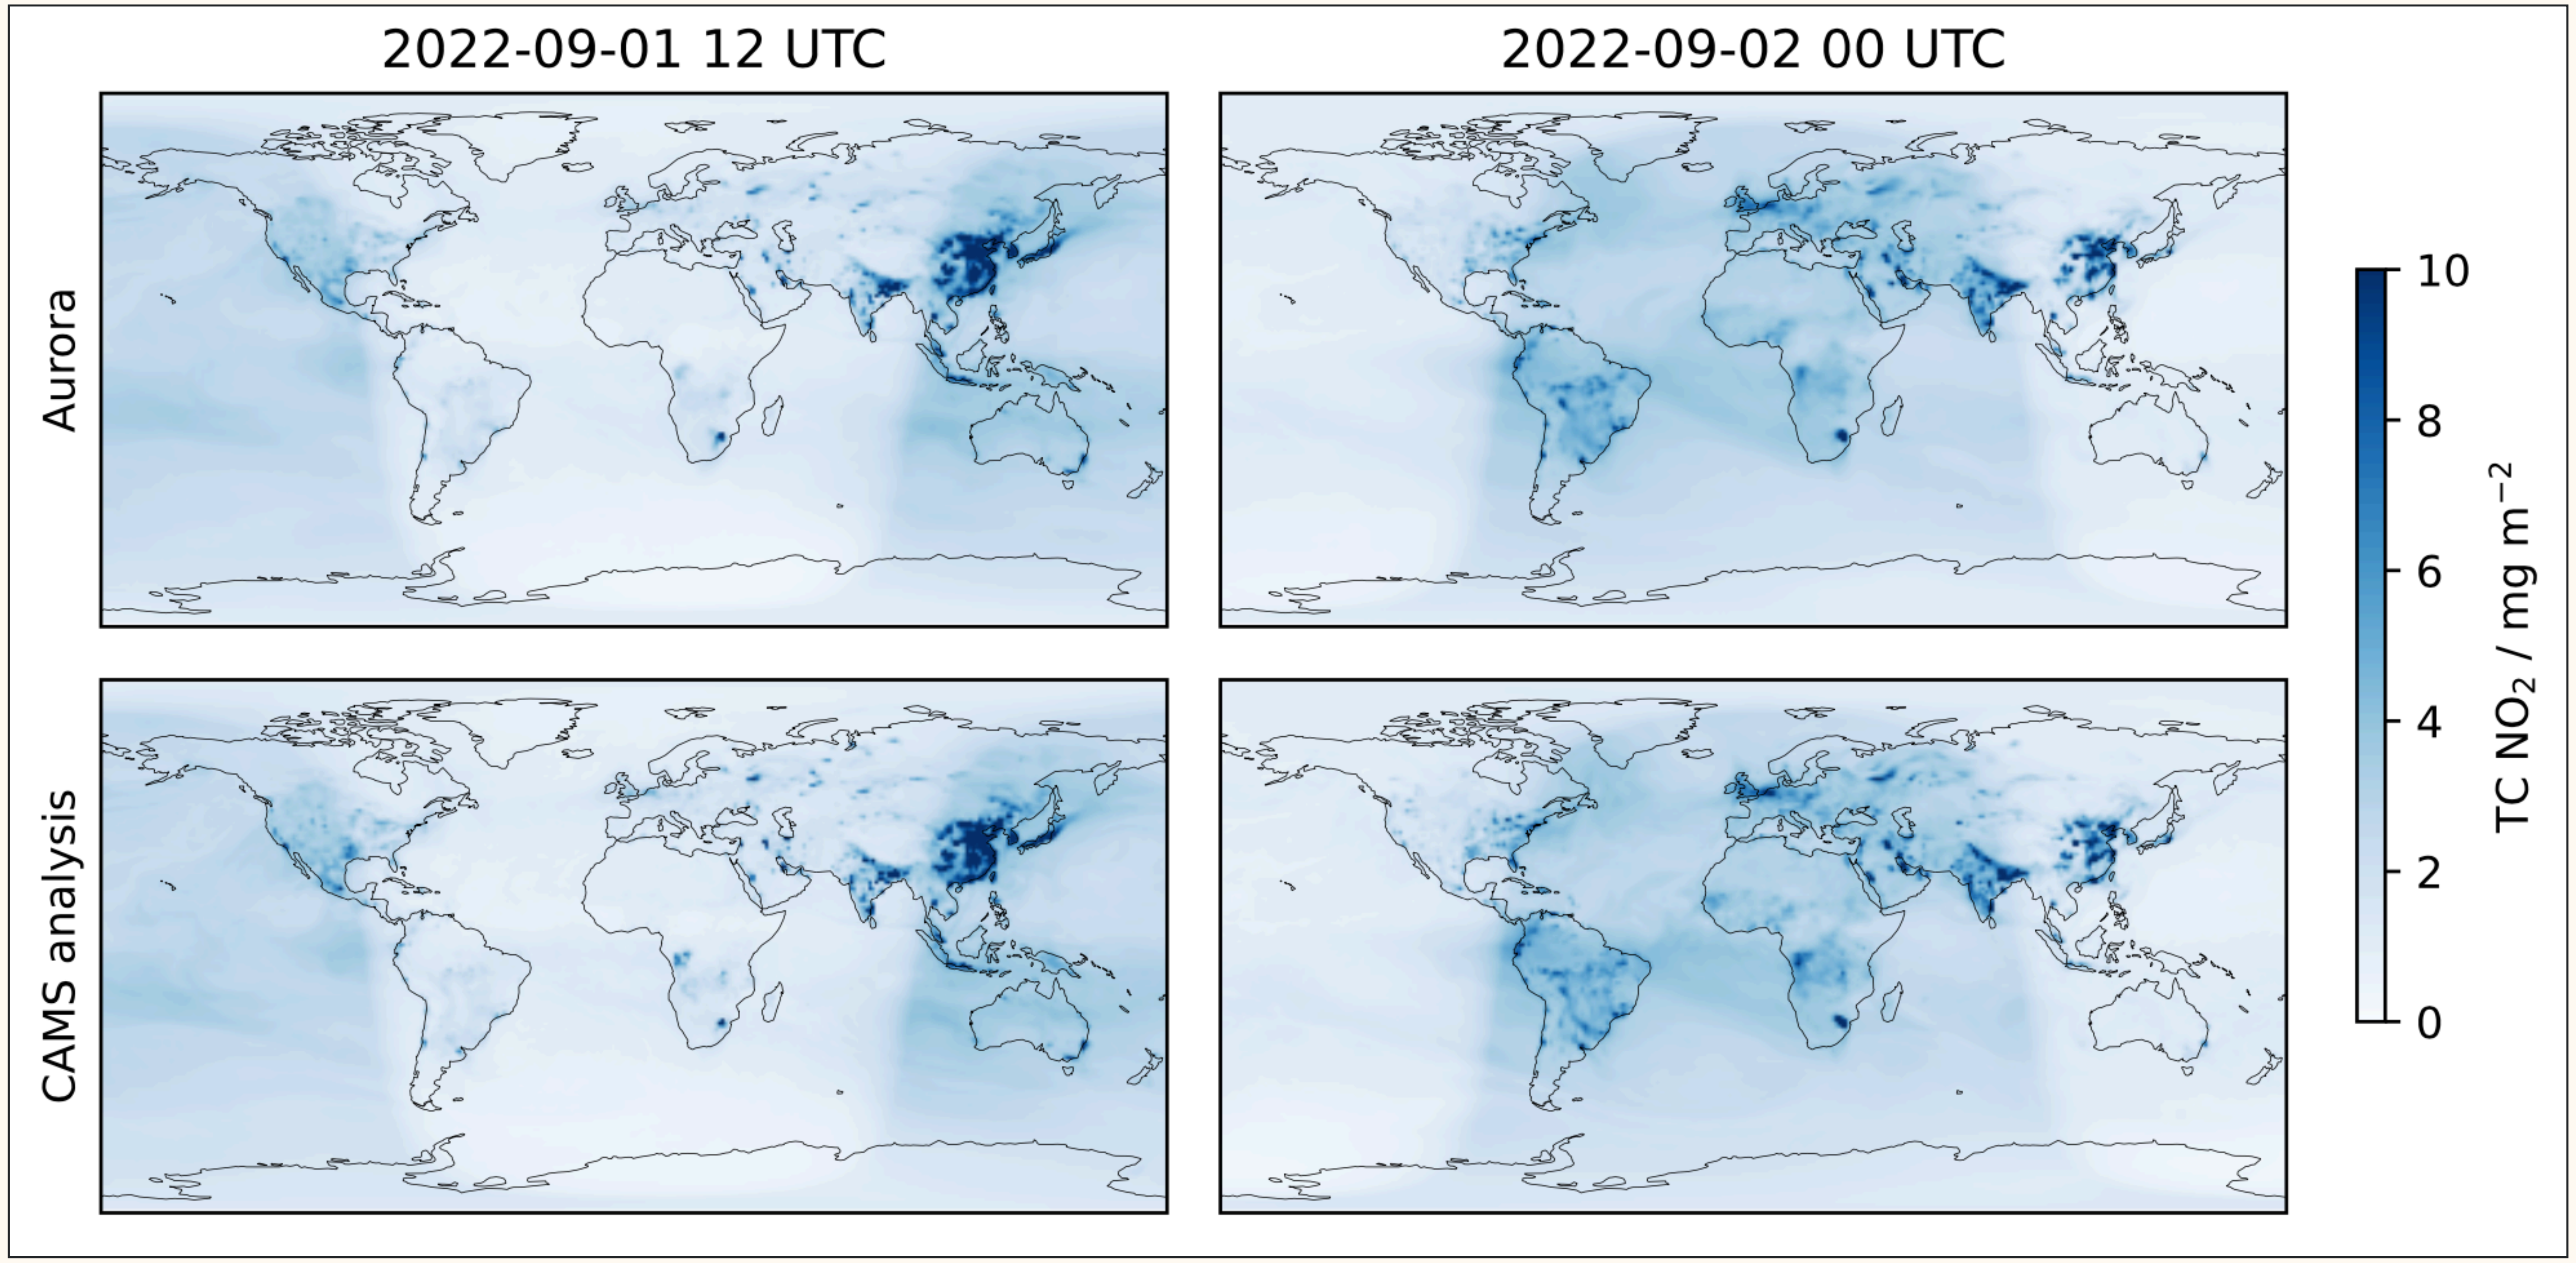

Overall:

Competitive on 95%
(≤ 20% RMSE)

Better on 75%

Three days:

Competitive on 100%
(≤ 20% RMSE)

Better on 86%

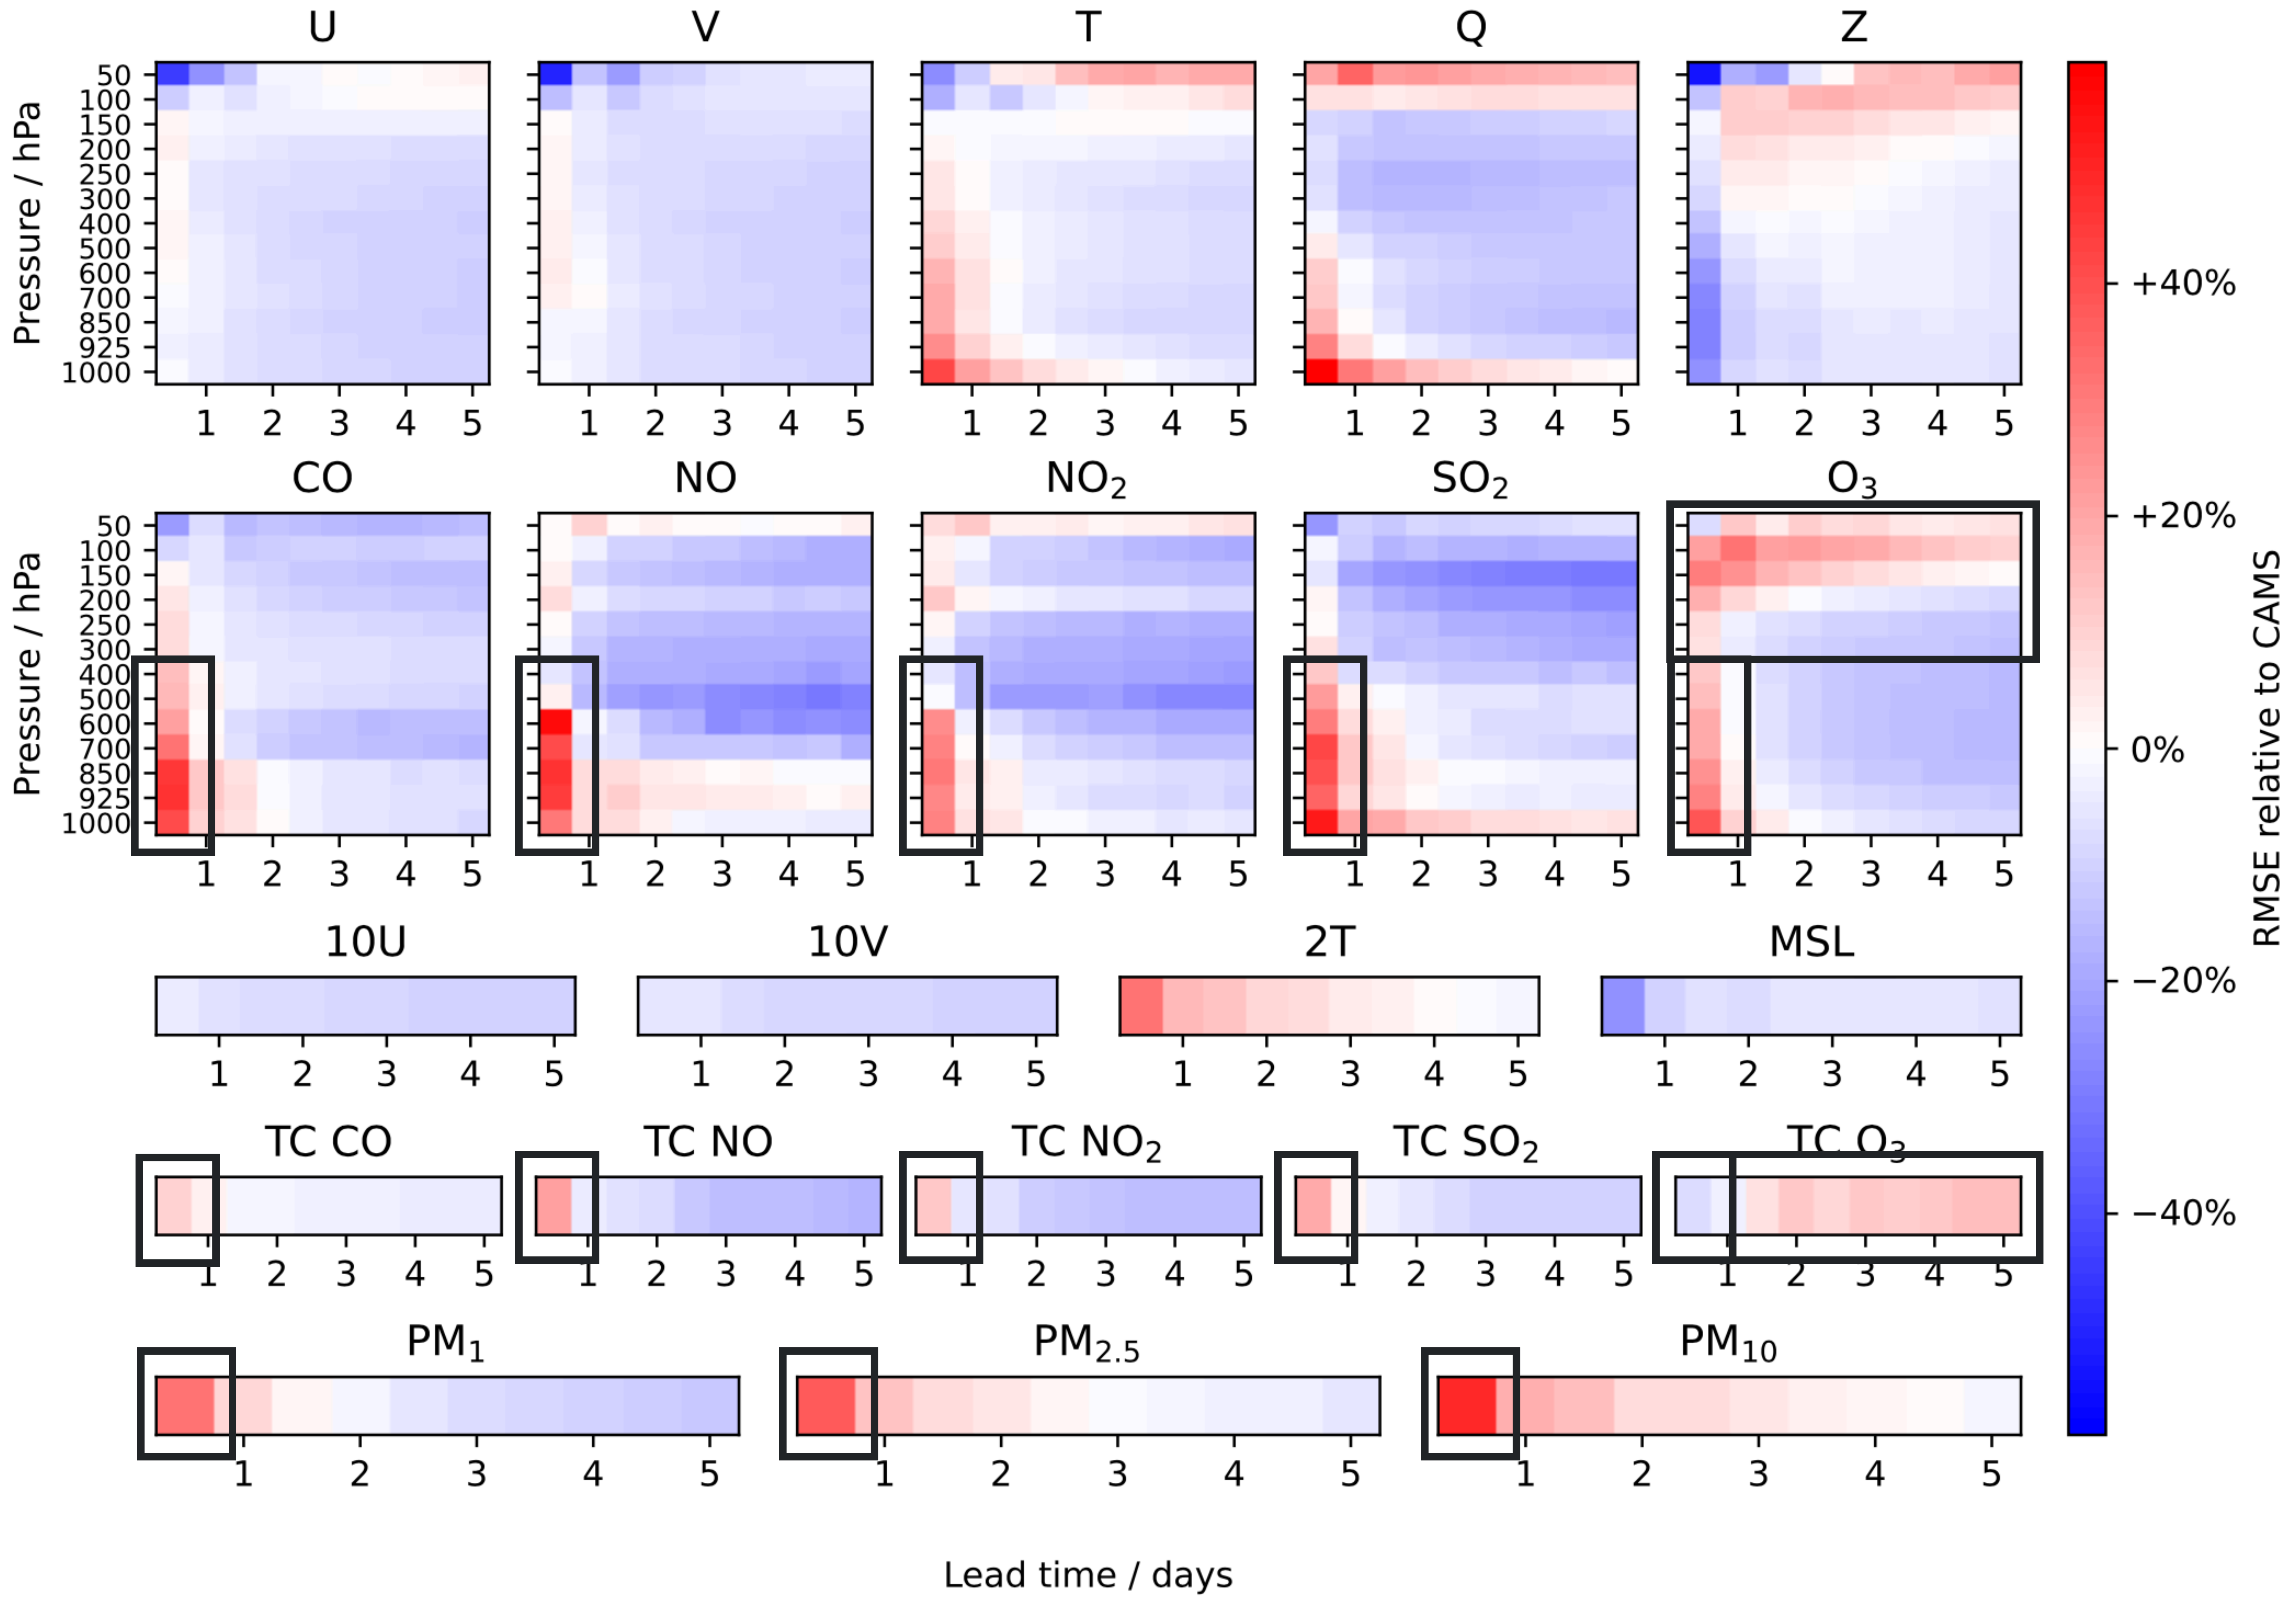

Aurora

Air Pollution Forecasting: Some Details

- Use pretrained 12-hour model

$$\hat{X}^{t+12\,\mathrm{h}} \;=\; \Phi_{\mathrm{pred}}(X^t, X^{t-12\,\mathrm{h}})$$

- "Mean-max" normalisation: $\mathrm{scale}_v = \mathrm{mean}_{\mathrm{time}}(\max_{\mathrm{space}}(X_v^t))$

- Transformation of inputs:

$$X_{\mathrm{transformed}} = c_1 \min(X, 2.5) + c_2 \frac{\log(\max(X, 10^{-4})) - \log(10^{-4})}{-\log(10^{-4})}$$

- Various architectural modifications

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
hi@wessel.ai