Aurora: A Foundation Model for the Earth System

Wessel Bruinsma

The Alan Turing Institute
Work done at Microsoft Research

AI4X 2025, 11 July 2025

The AI Revolution in Weather Forecasting

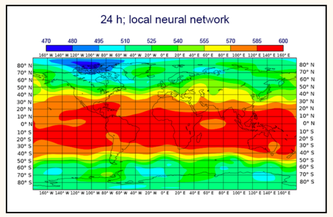

2018

First serious efforts to compare AI models to physics baselines

Dueben and Bauer (2018)

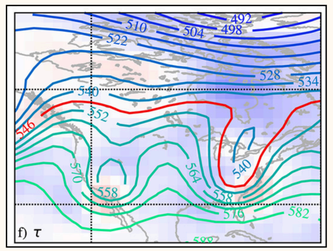

2019

AI models skillful to multiple days

Weyn et al. (2019)

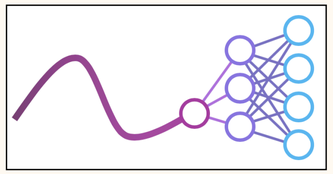

2020

WeatherBench starts to drive ML development

Rasp et al. (2020)

The AI Revolution in Weather Forecasting

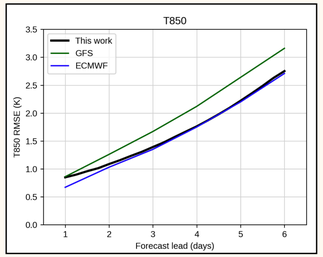

HUAWEI

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

NVIDIA

Microsoft

Google DeepMind

2022–2023

Tech companies start to work in this space

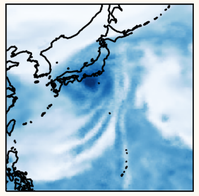

2023

GenCast outperforms IFS ensemble

Price et al. (2024)

ECMWF

2024

ECMWF launches AIFS

What About Other Forecasting Tasks?

Air pollution

Atmospheric composition

Waves

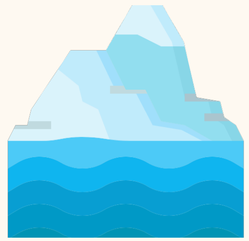

Sea ice

Ocean

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

Aurora

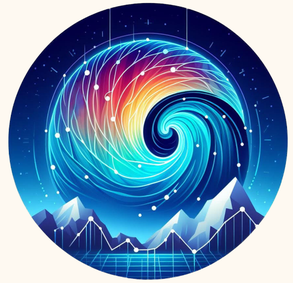

pretraining

- 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

Fine-Tuning Applications

Operational in all settings!

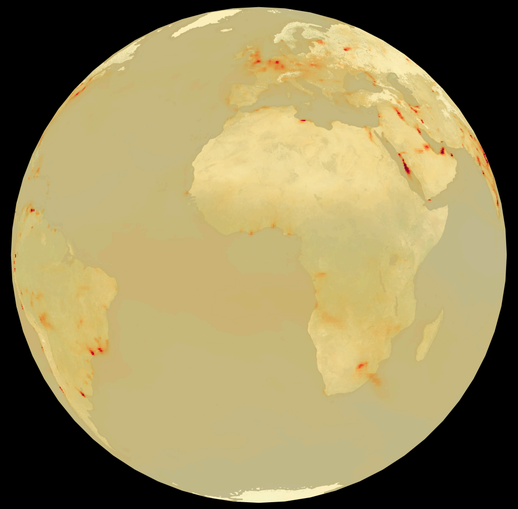

Atmospheric comp. and air pollution

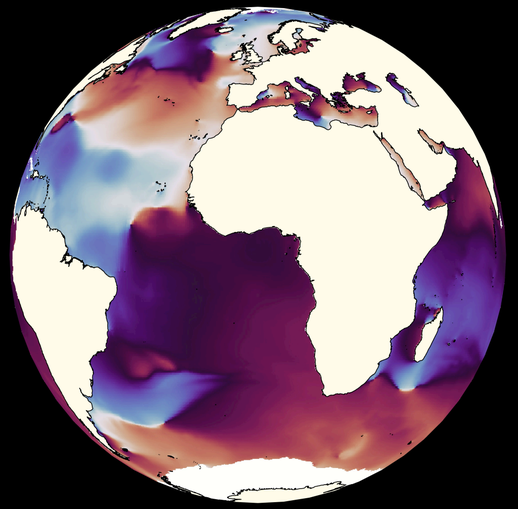

Ocean waves

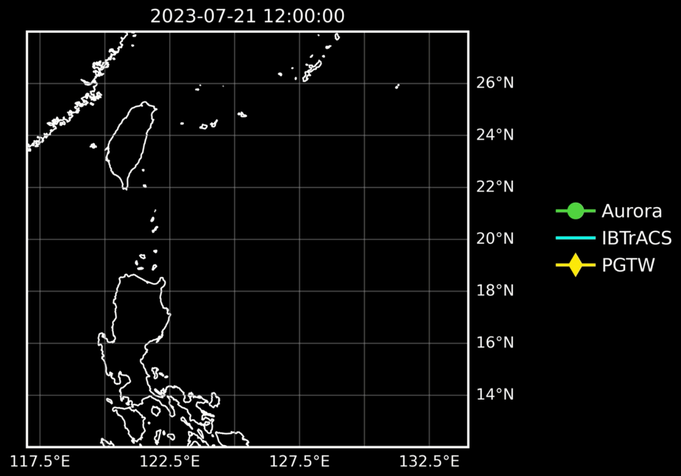

Tropical cyclone tracks

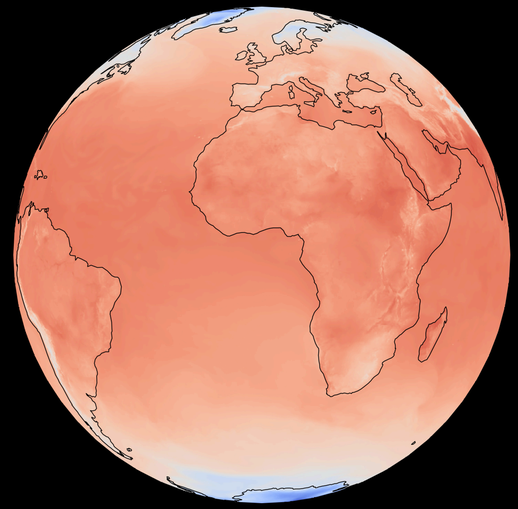

High-resolution weather

This talk

Megan Stanley
AI for Climate & Weather (Lecture Theater 53)

Animations created by Megan Stanley

Air Pollution Forecasting

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

- **Setup:** model 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

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

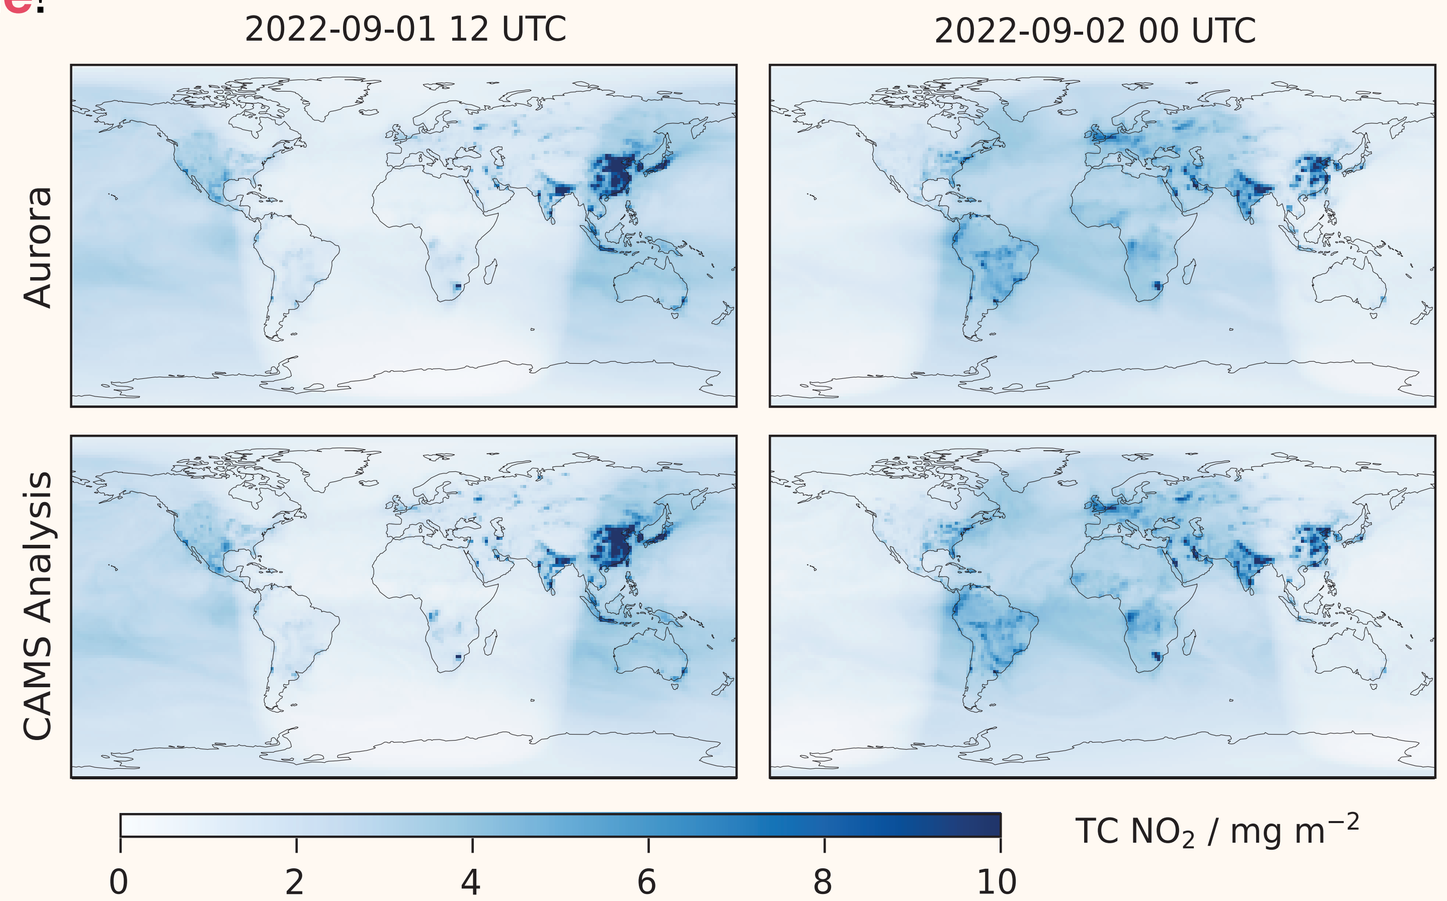

Overall:

Competitive on 95%
(≤ 20% RMSE)

Better on 75%

Three days:

Competitive on 100%
(≤ 20% RMSE)

Better on 86%

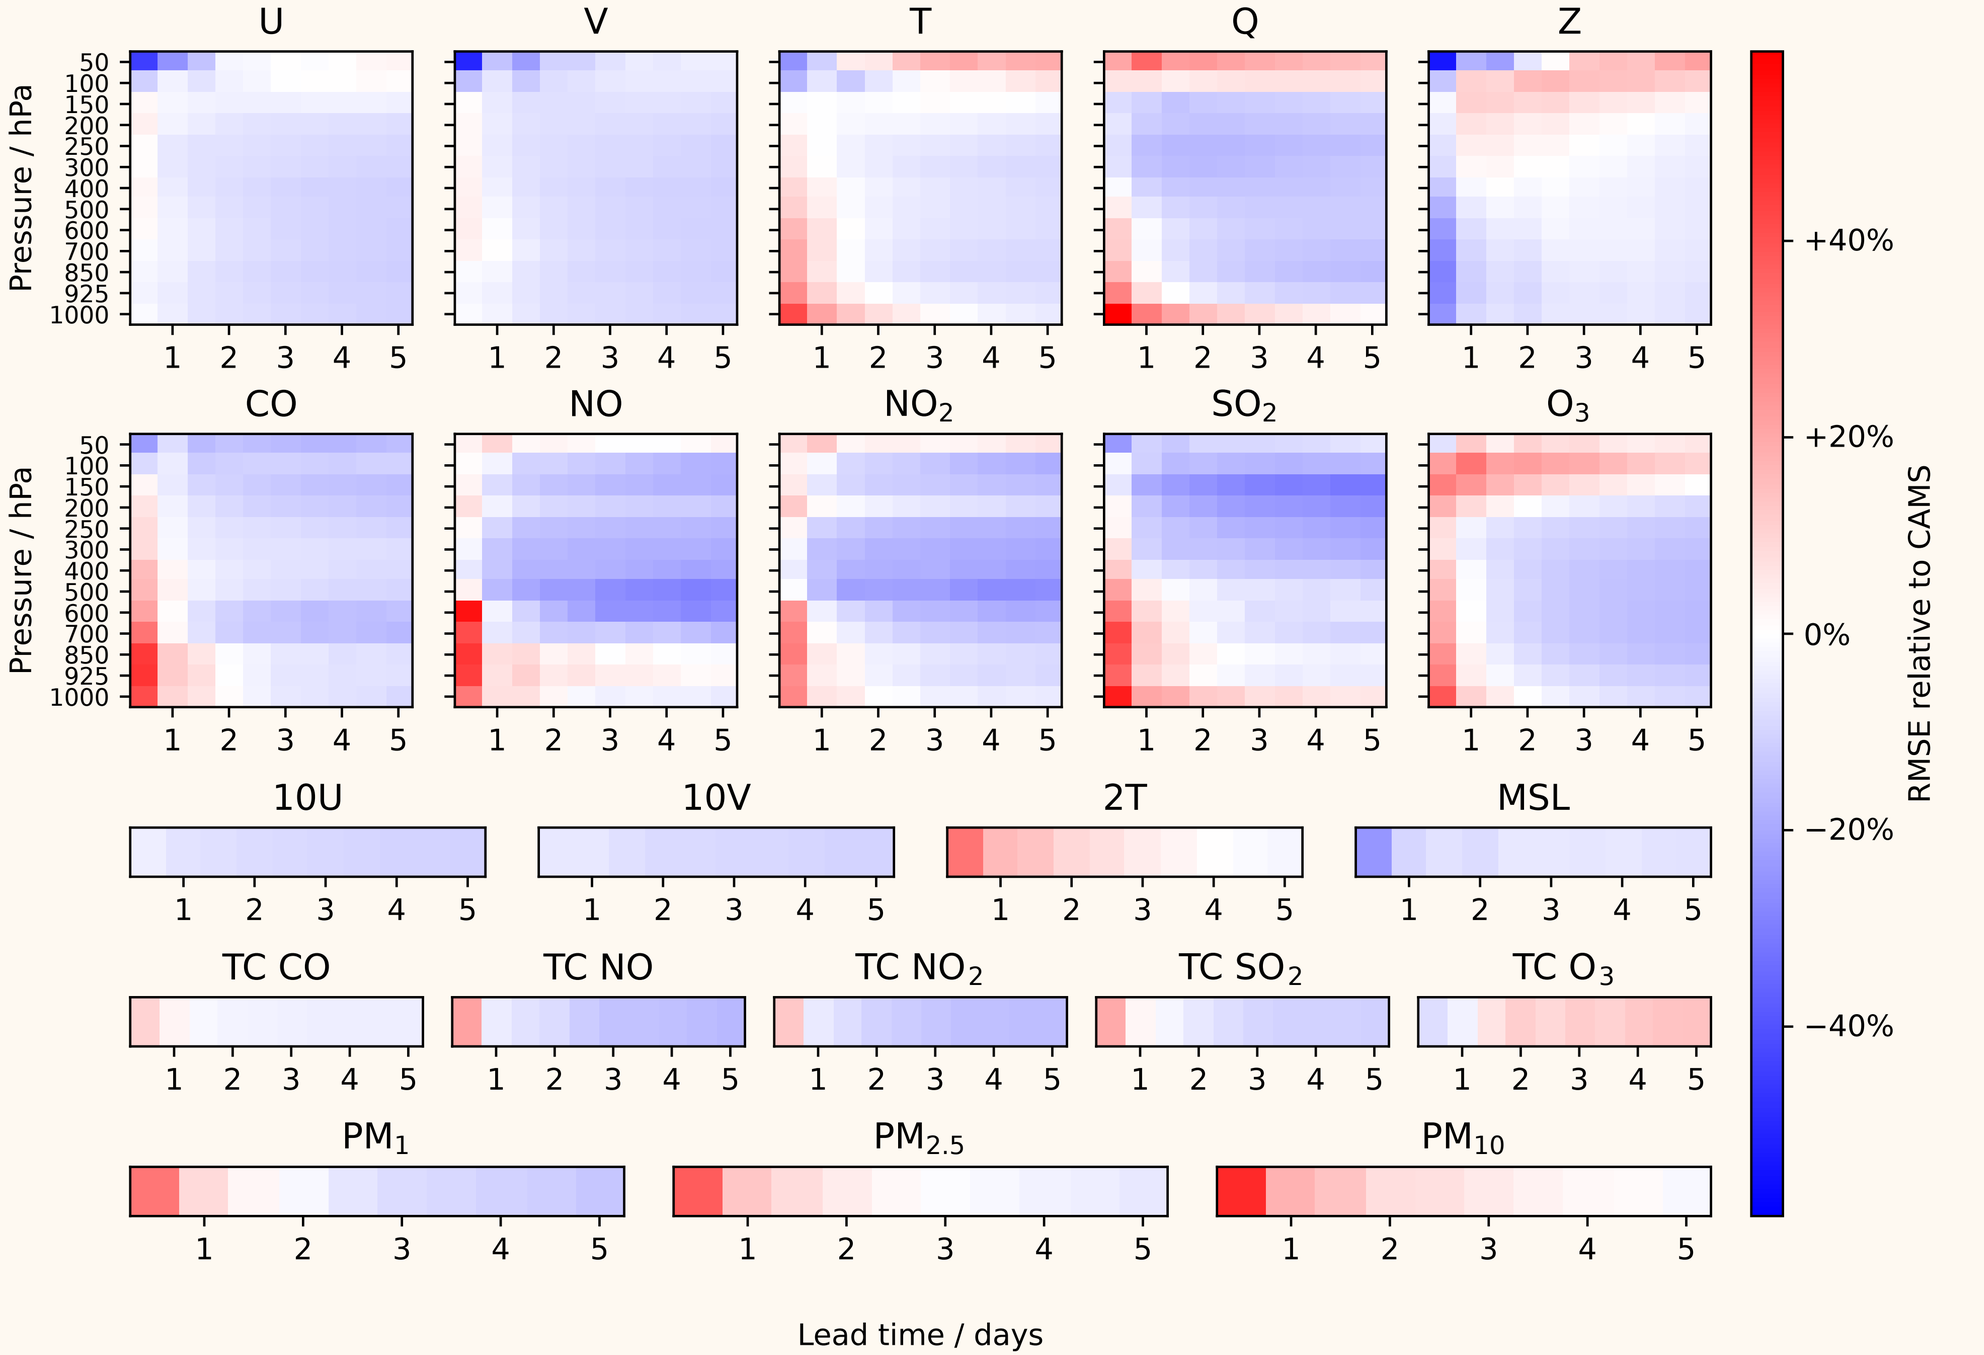

Ocean Wave Forecasting

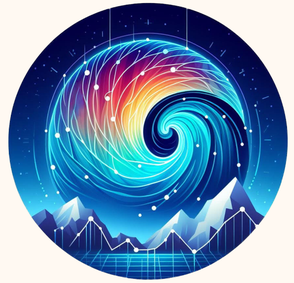

- **Setup:** model height, direction, and period of wave components
- **Data:** HRES-WAM analysis
- **Baseline:** HRES-WAM forecasts

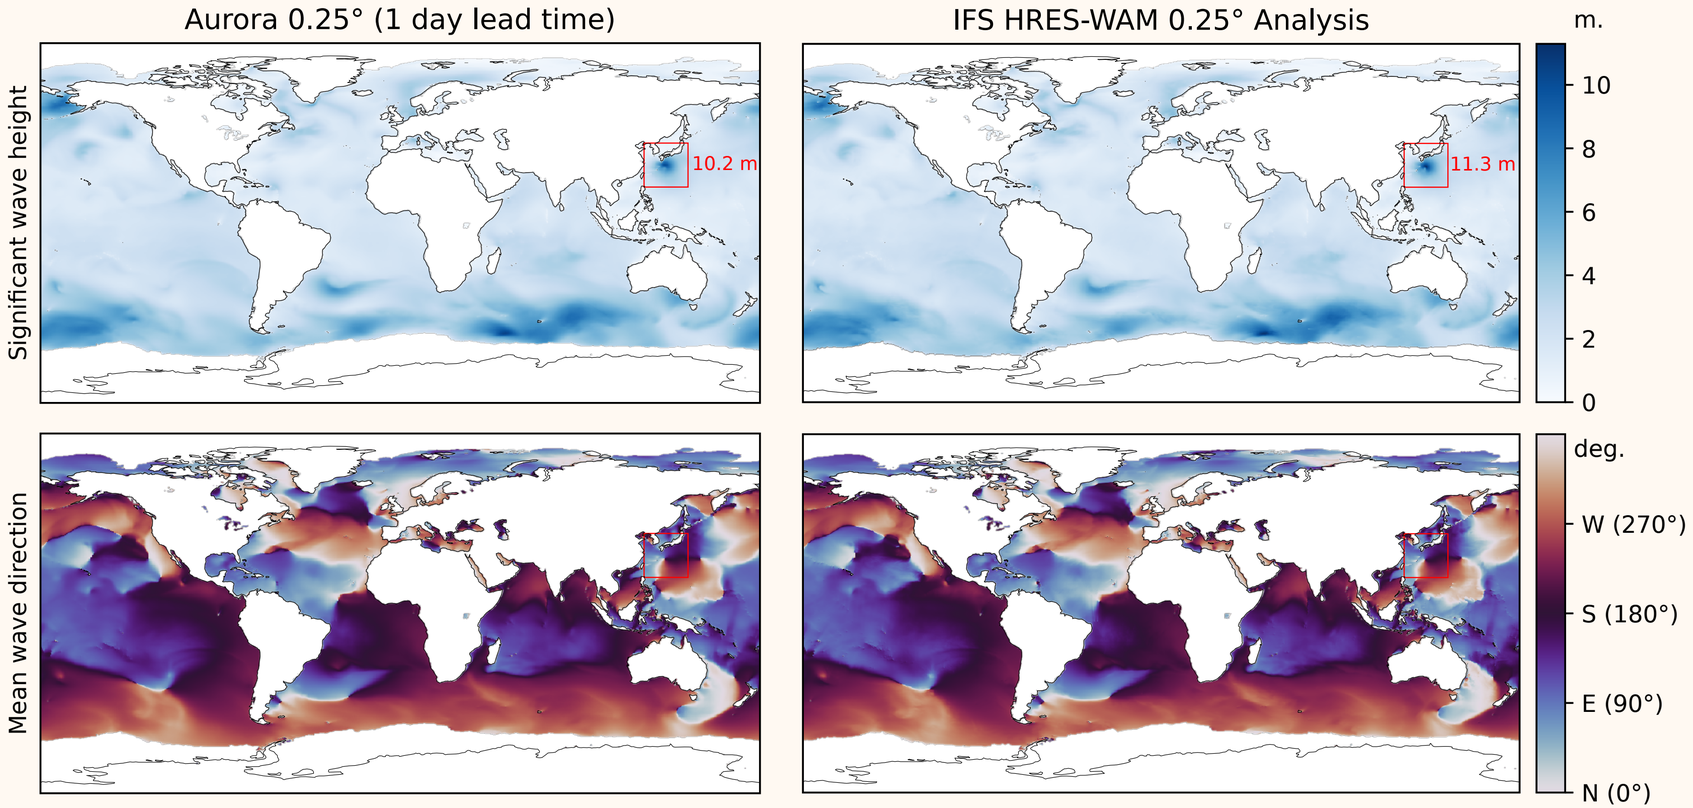

Overall:

Competitive on 96%
(≤ 20% RMSE)

Better on 86%

Three days:

Competitive on all but PP1D
(≤ 20% RMSE)

Better on 91%

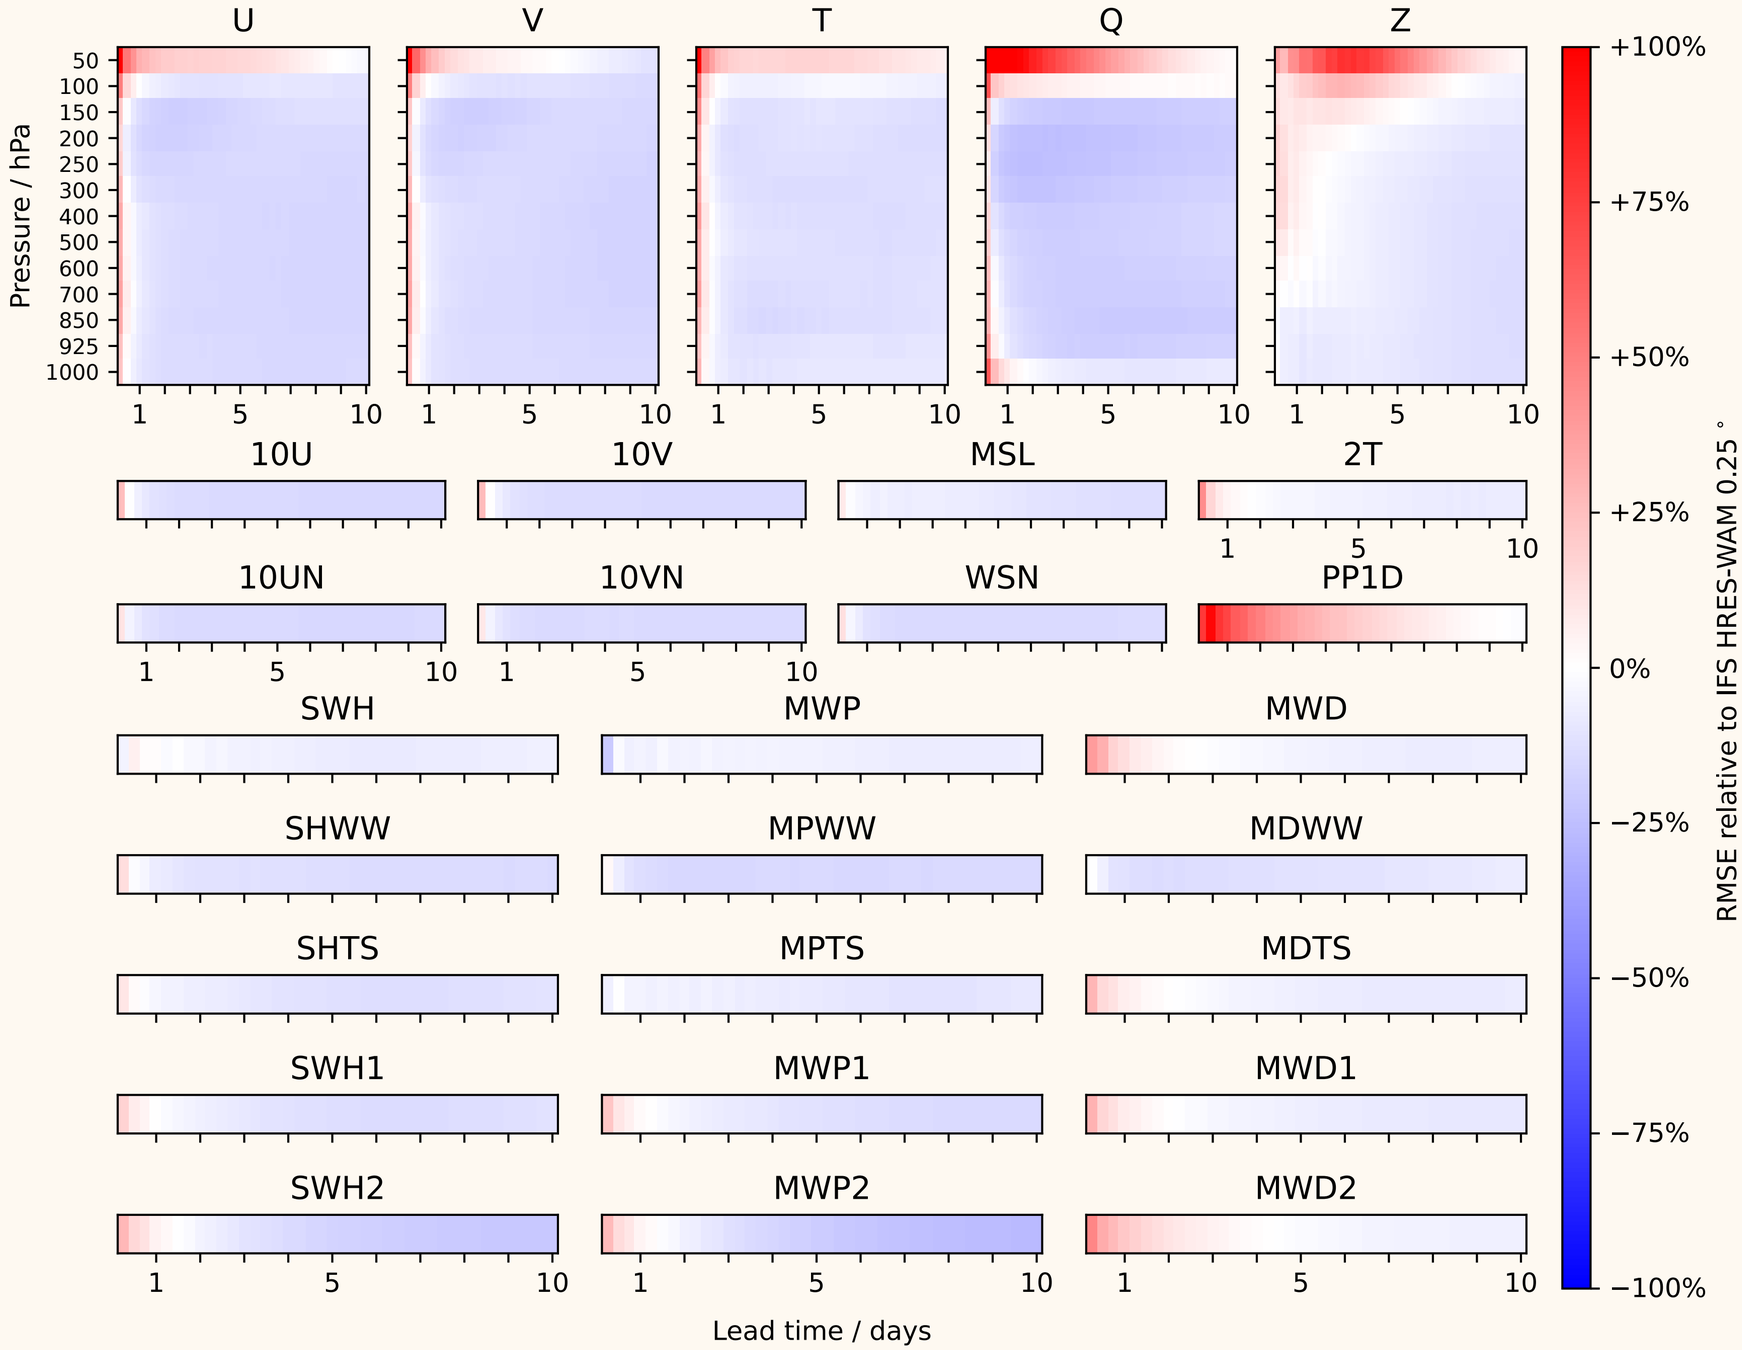

Open Source

- All models open source under MIT licence!
- Details docs with examples

```python
import torch

from aurora import Aurora, Batch, rollout

model = Aurora()
model.load_checkpoint()

model.eval()
model.to("cuda")

batch = Batch(...)

with torch.inference_mode():
    for prediction in rollout(model, batch, steps=10):
        ...  # Do something with `prediction`.

```

https://github.com/microsoft/aurora

pip install microsoft-aurora

The Weather Forecasting Pipeline

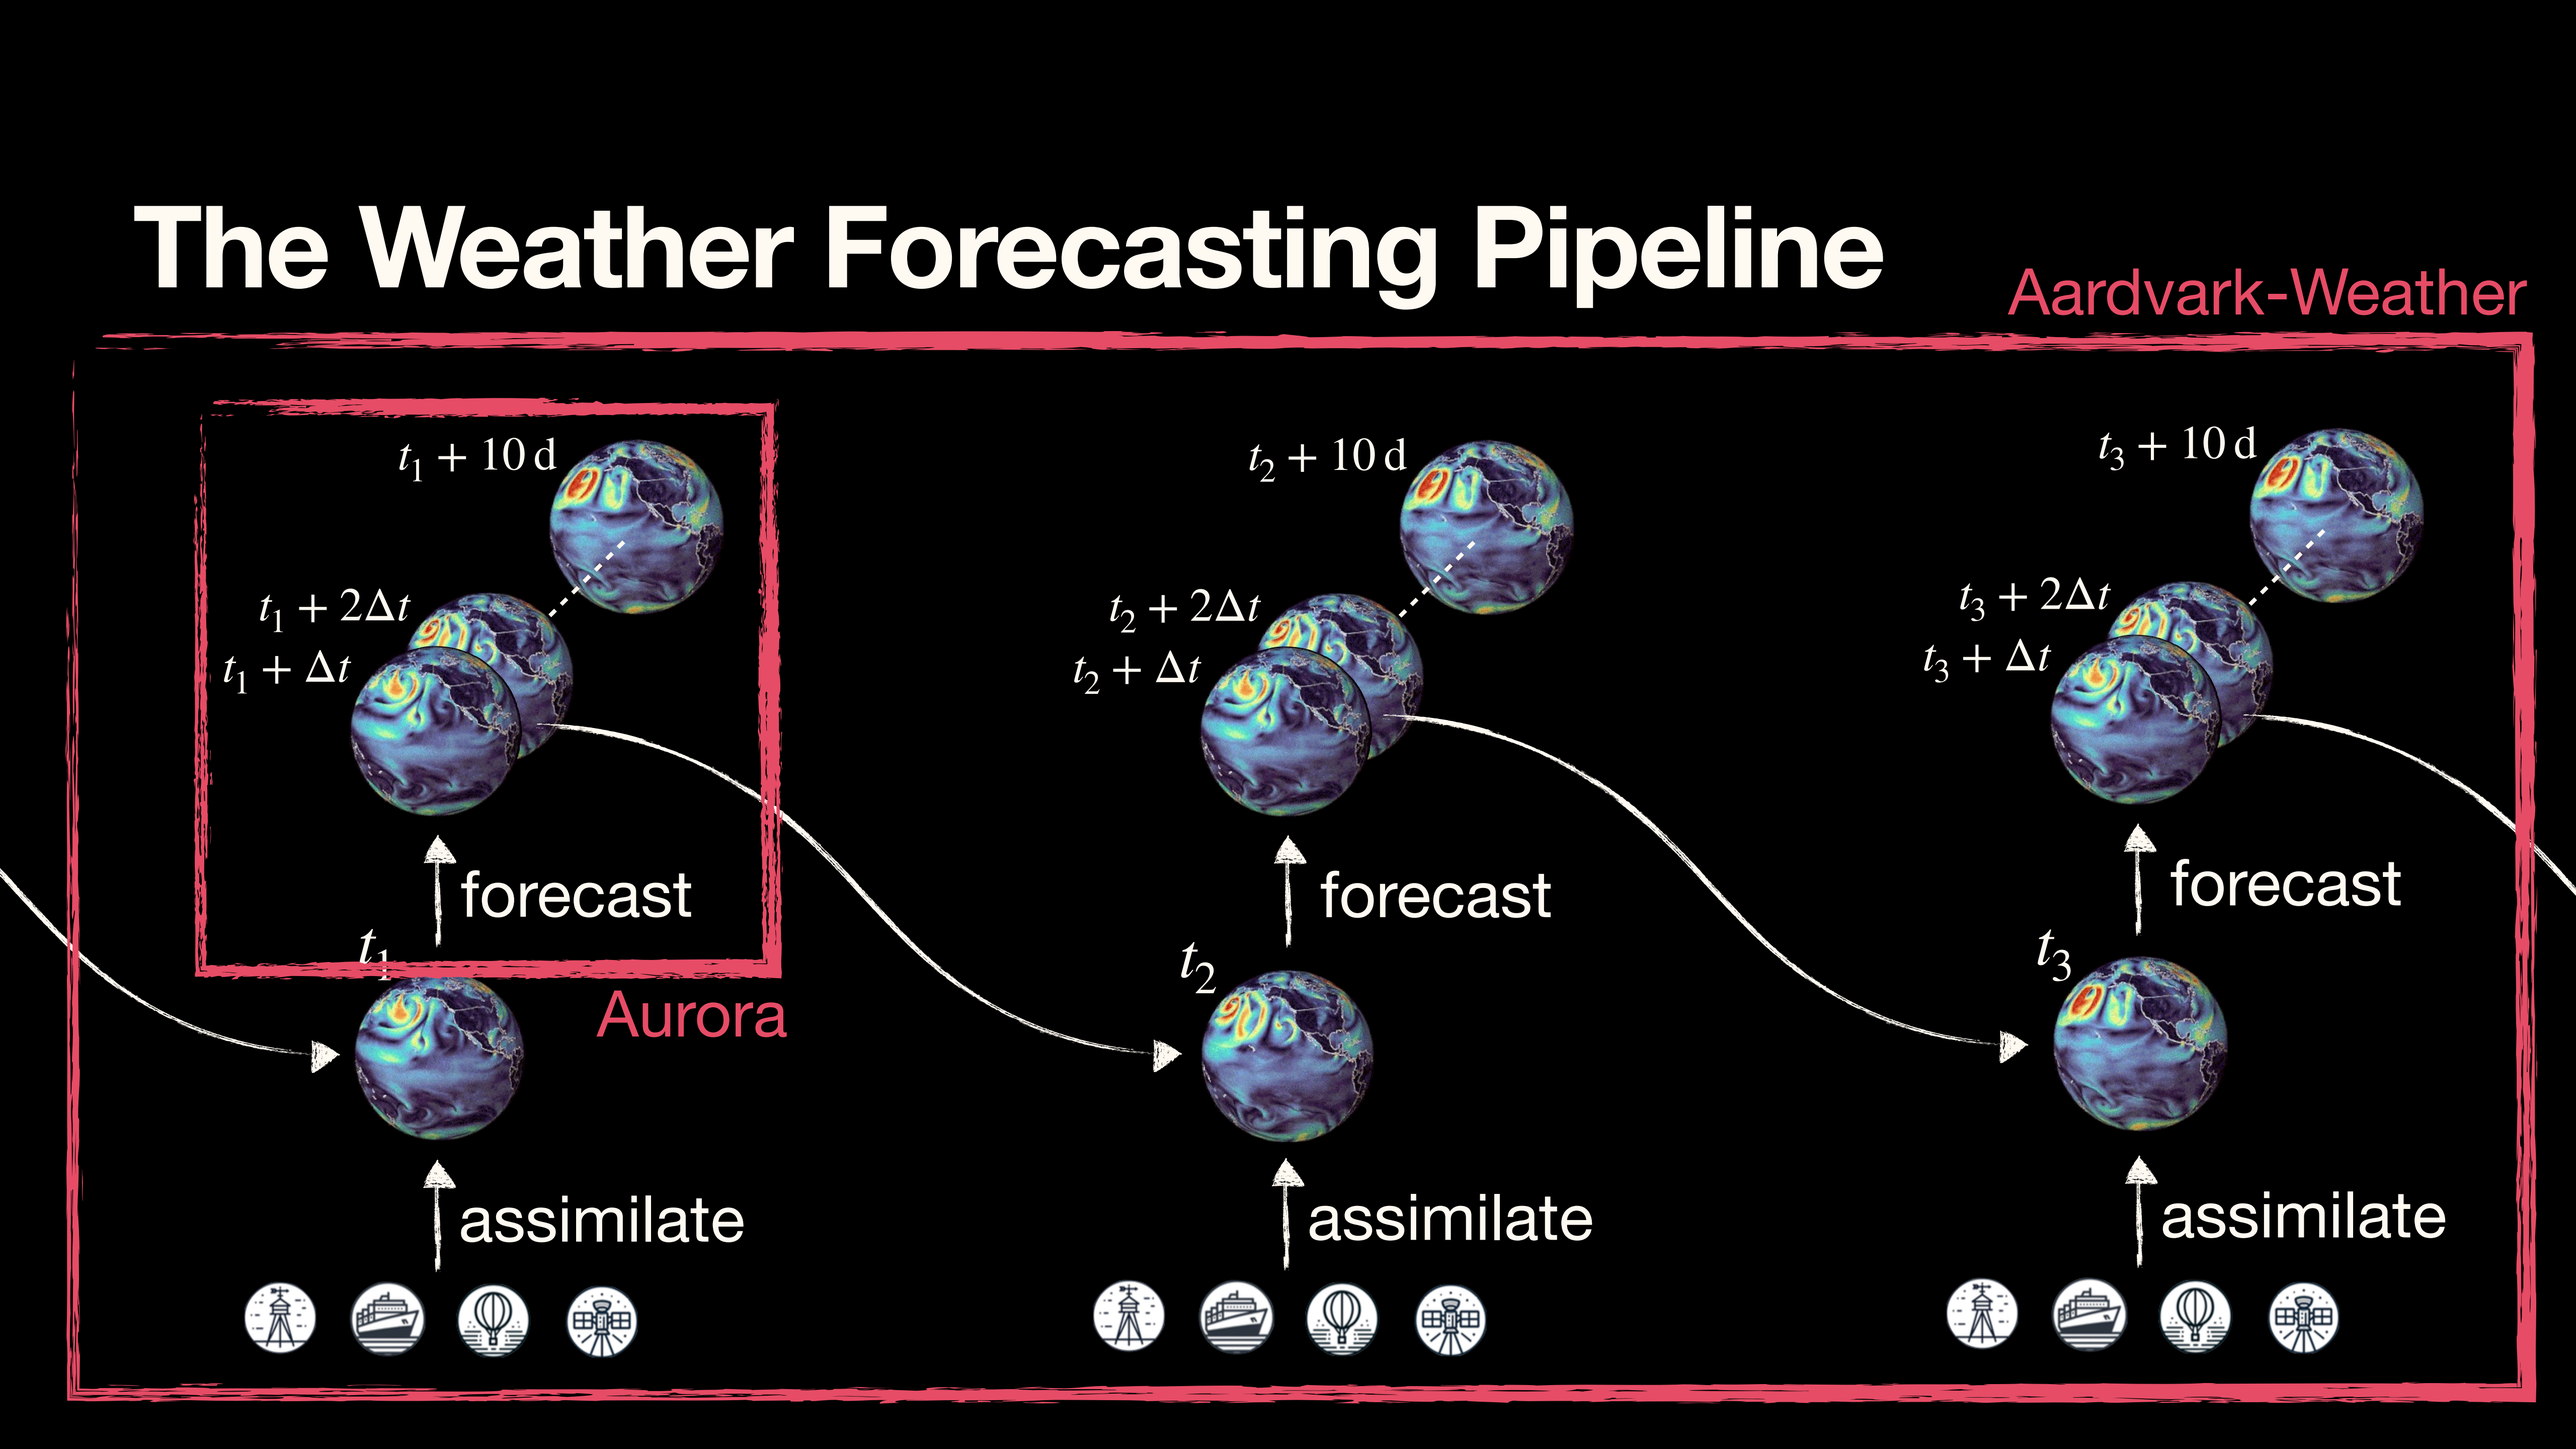

The Aurora Team

Paris Perdikaris

University of Pennsylvania, formerly MSR

Wessel Bruinsma

The Alan Turing Institute, formerly MSR

Megan Stanley

MSR

Cristian Bodnar

Silurian, formerly MSR

Ana Lučić

University of Amsterdam, formerly MSR

Richard Turner

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

Anna Allen

University of Cambridge, The Alan Turing Institute

Johannes Brandstetter

JKU Linz, Emmi AI, formerly MSR

Patrick Garvan

IONQ, formerly MSR

Maik Riechert

MSR

Max Welling

University of Amsterdam, CuspAI, formerly MSR

Elizabeth Heider

Book tour, formerly MSR

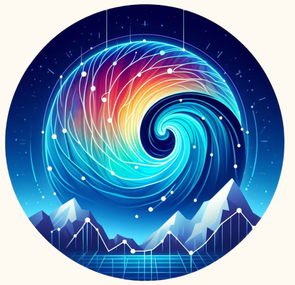

Conclusion

- 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).