Neural General Circulation Models
for Weather & Climate

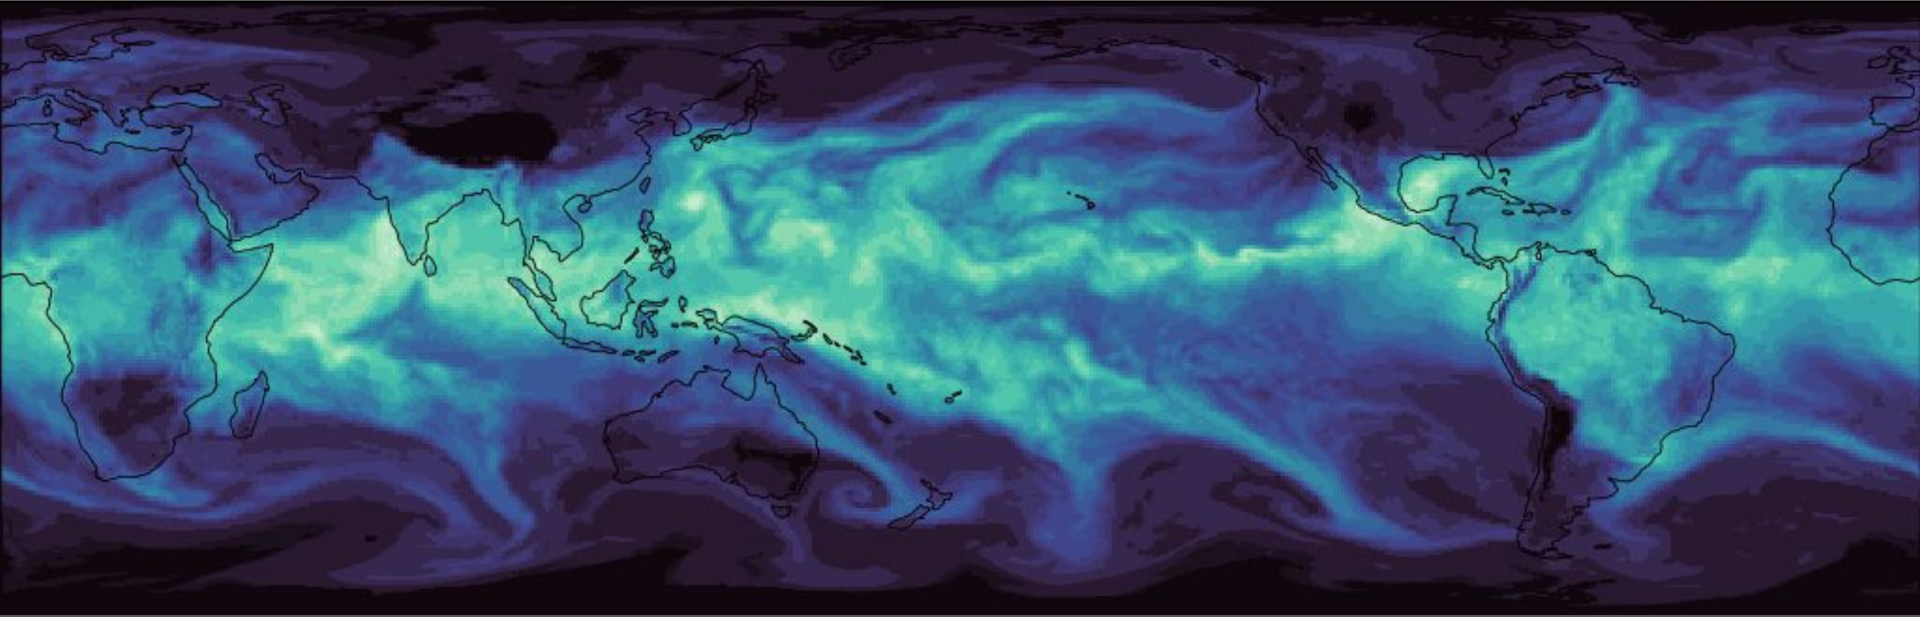

Dmitrii Kochkov
Google Research

Colorado State University
Atmospheric Science Department
4 November 2024

Neural GCM core contributors & collaborators

Dmitrii Kochkov

Janni Yuval

Ian Langmore

Peter Norgaard

Jamie Smith

Stephan Hoyer

Google Research

Google DeepMind

ECMWF

MIT

NeuralGCM collaborators:

Griffin Mooers, James Lottes, Stephan Rasp, Sam Hatfield, Peter Duben, Milan Klower, Peter Battaglia, Alvaro Sanchez-Gonzalez, Matthew Willson, Michael Brenner

Google Research

Outline

1. General Circulation Models (GCMs) for weather and climate
2. AI revolution for weather forecasting
3. Neural GCM - differentiable hybrid atmospheric model
4. Neural GCM results
   a. Weather forecasting
   b. “Climate” simulations
5. Future directions

Outline

1. General Circulation Models (GCMs) for weather and climate
2. AI revolution for weather forecasting
3. Neural GCM - differentiable hybrid atmospheric model
4. Neural GCM results
    a. Weather forecasting
    b. “Climate” simulations
5. Future directions

Simulation of weather and climate

Forecast workflow:

Data assimilation; **Forward model of the Earth system**; Post-processing

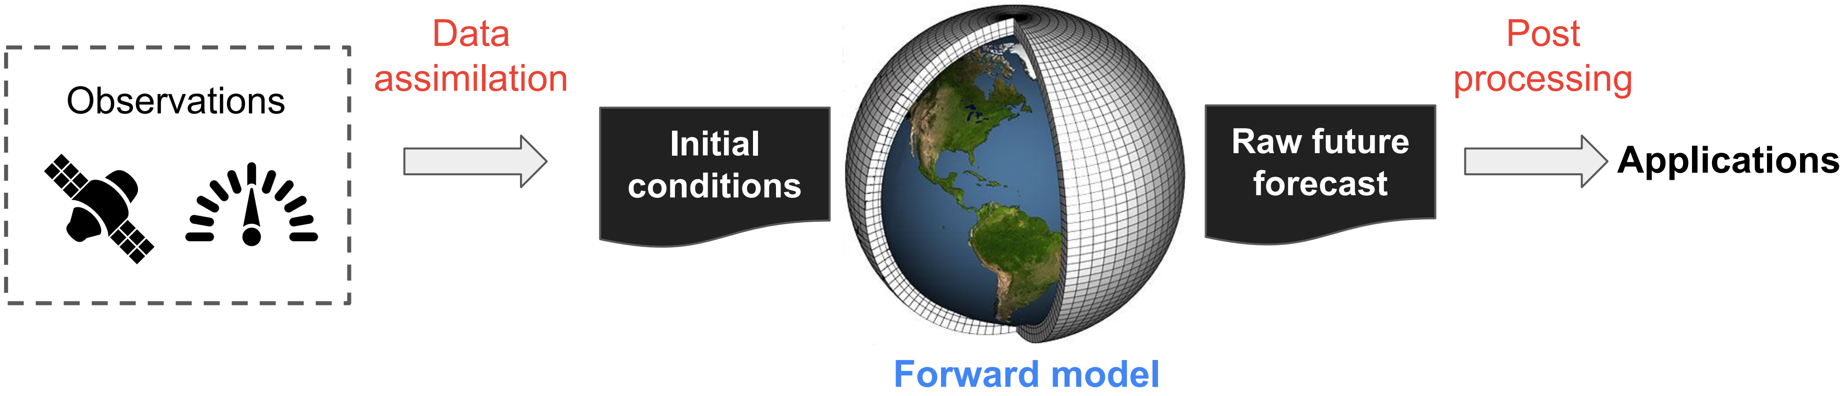

"Forecasting" queries vary in timescales:

- Will it rain in 3 hours?
- What's the weather in 3 days?
- What is return time of a class 5 hurricane?
- How warm the Earth may be in 30 years if "*"?

Initial condition

[nowcasting, medium range, …]

Boundary condition

climate variability, catastrophe risks

How do traditional General Circulation Models work?

“Dynamical core”

$$\frac{d\mathbf{u}}{dt} + f\mathbf{k}\times\mathbf{u} + \frac{1}{\rho}\nabla_z p = \mathbf{0}$$

$$\frac{\partial\rho}{\partial t} + \nabla_z\cdot(\rho\mathbf{u}) + \frac{\partial\rho w}{\partial z} = 0$$

$$\frac{dT}{dt} - \frac{\omega}{c_p\,\rho} = 0$$

$$\frac{\partial p}{\partial z} = -\rho g$$

$$p = \rho R T$$

Fluid dynamics on the surface of a rotating sphere

How do traditional General Circulation Models work?

“Dynamical core”

$$\frac{d\mathbf{u}}{dt} + f\mathbf{k}\times\mathbf{u} + \frac{1}{\rho}\nabla_z p = \mathbf{0}$$

$$\frac{\partial\rho}{\partial t} + \nabla_z\cdot(\rho\mathbf{u}) + \frac{\partial\rho w}{\partial z} = 0$$

$$\frac{dT}{dt} - \frac{\omega}{c_p\,\rho} = 0$$

$$\frac{\partial p}{\partial z} = -\rho g$$

$$p = \rho R T$$

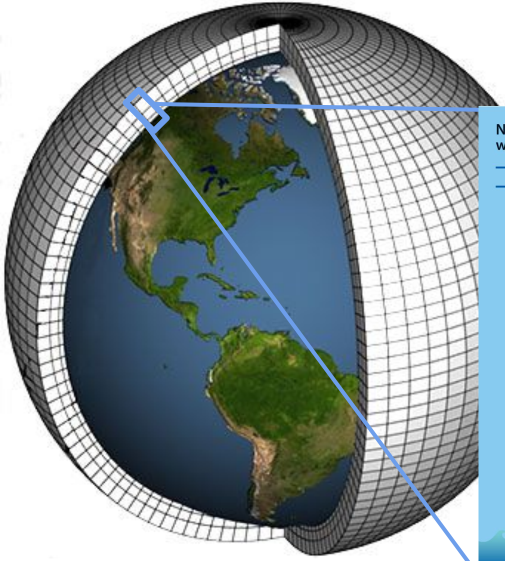

“Physics”

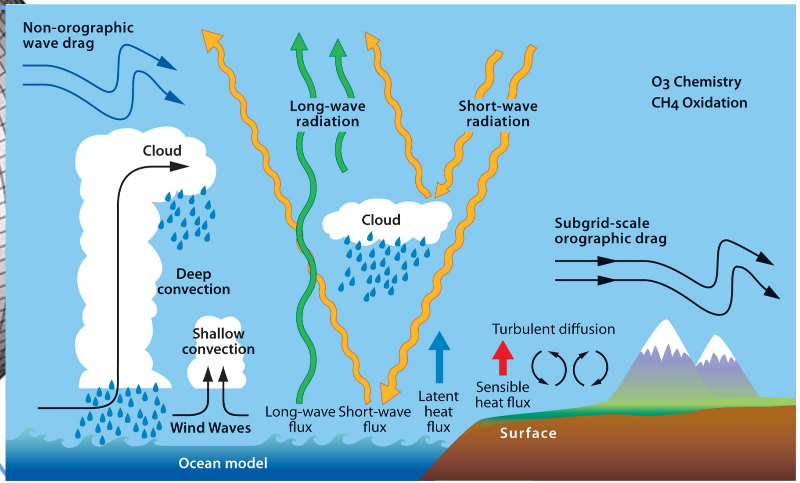

Fluid dynamics on the surface of a rotating sphere

Many PhD theses
(100k-1M lines of Fortran)

How do traditional General Circulation Models work?

“Dynamical core”

$$\frac{d\mathbf{u}}{dt} + f\mathbf{k}\times\mathbf{u} + \frac{1}{\rho}\nabla_z p = \mathbf{0}$$

$$\frac{\partial \rho}{\partial t} + \nabla_z\cdot(\rho\mathbf{u}) + \frac{\partial \rho w}{\partial z} = 0$$

$$\frac{dT}{dt} - \frac{\omega}{c_p\,\rho} = 0$$

$$\frac{\partial p}{\partial z} = -\rho g$$

$$p = \rho R T$$

Fluid dynamics on the surface of a rotating sphere

“Physics”

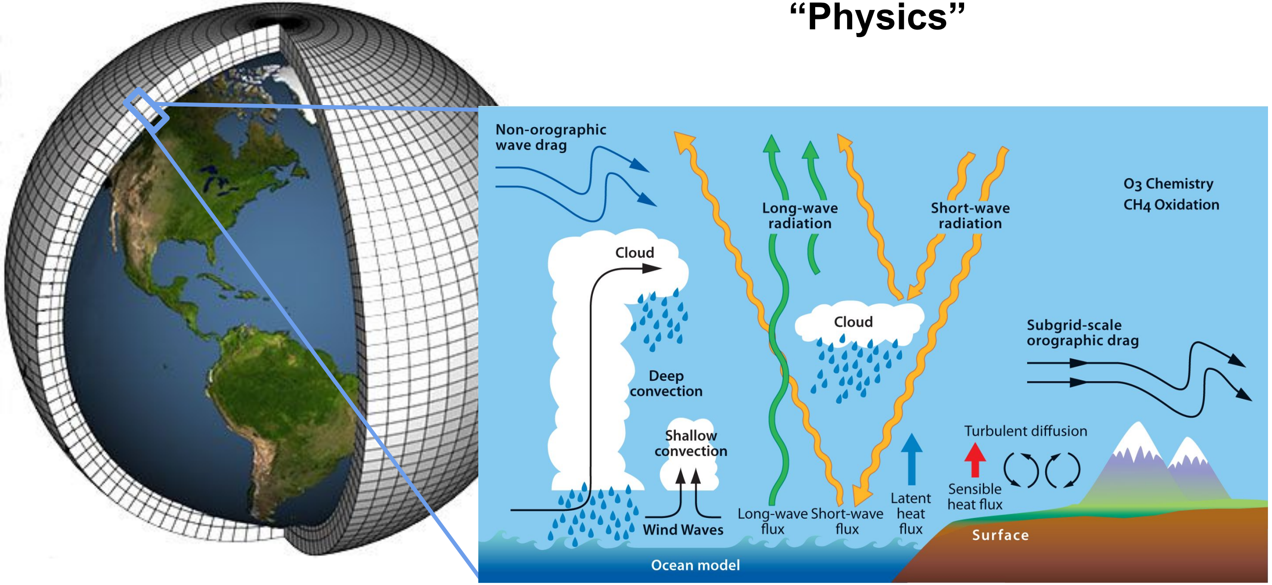

Many PhD theses
(100k-1M lines of Fortran)

Hard to tune - believed to be the largest source of inaccuracies

Success and scaling of GCMs for weather and climate

REVIEW

doi:10.1038/nature14956

The quiet revolution of numerical weather prediction

Peter Bauer[1], Alan Thorpe[1] & Gilbert Brunet[2]

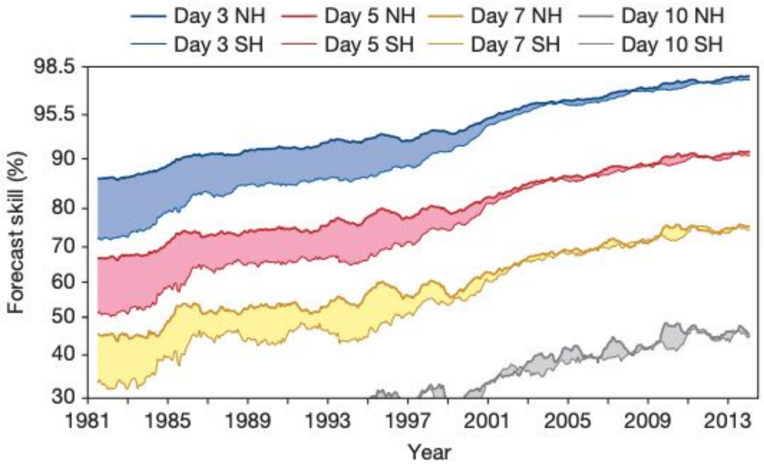

Bauer et al (2015)

COMMENTARY:

Climate goals and computing the future of clouds

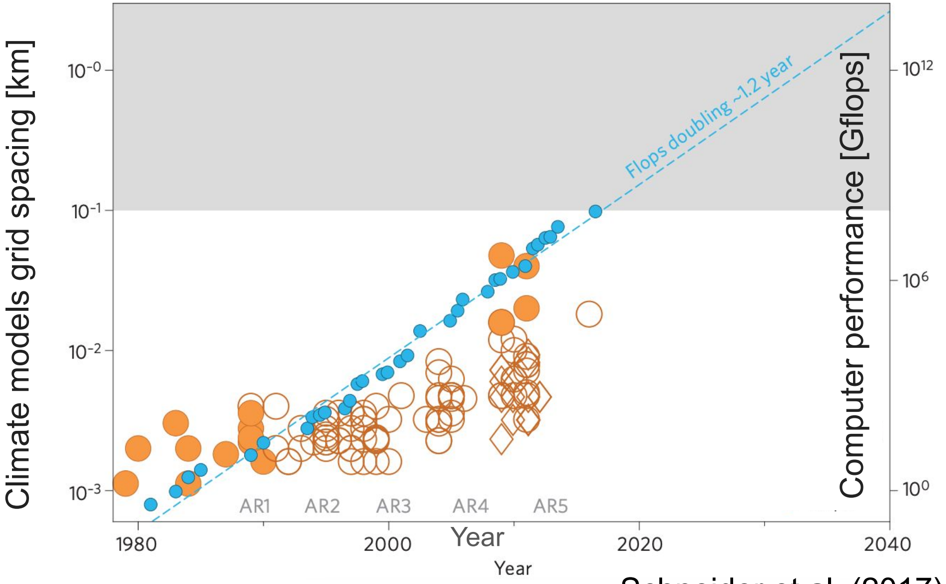

Schneider et al. (2017)

Success and scaling of GCMs for weather and climate

REVIEW

doi:10.1038/nature14956

The quiet revolution of numerical weather prediction

Peter Bauer[1], Alan Thorpe[1] & Gilbert Brunet[2]

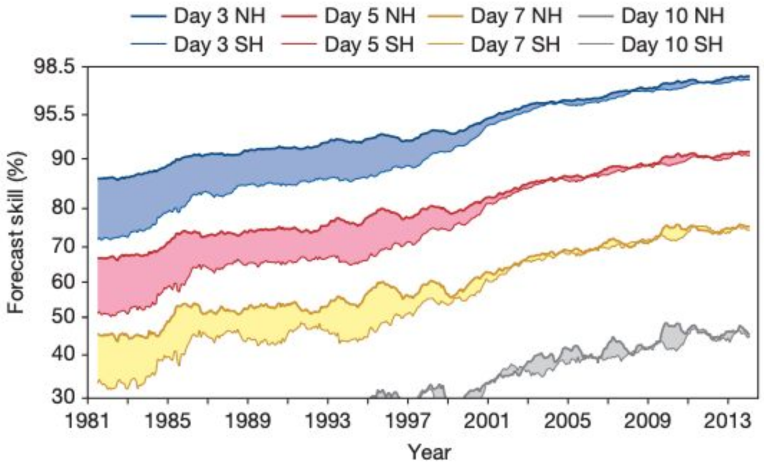

Bauer et al (2015)

COMMENTARY:

Climate goals and computing the future of clouds

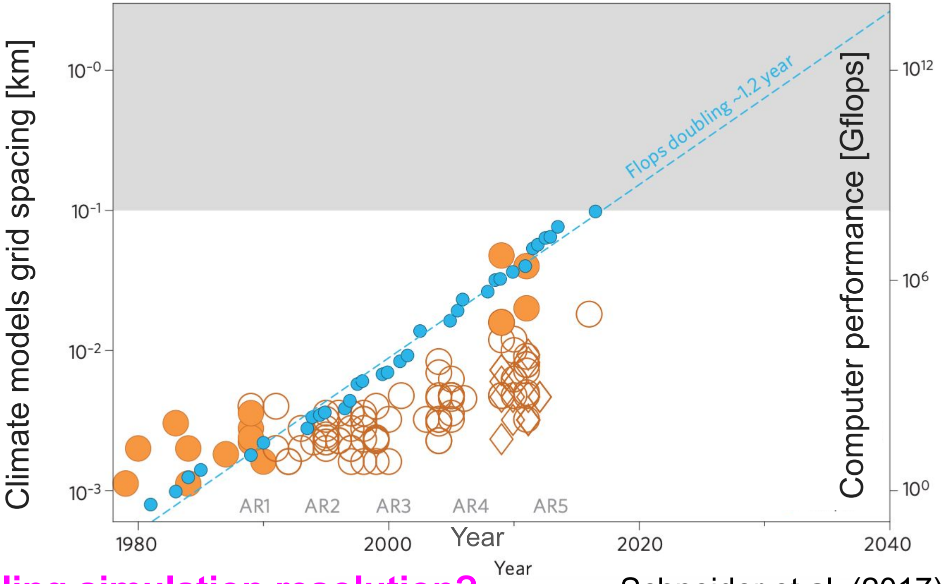

Next: continue scaling simulation resolution?

Schneider et al. (2017)

Google Research

Outline

1. General Circulation Models (GCMs) for weather and climate
2. AI revolution for weather forecasting
3. Neural GCM - differentiable hybrid atmospheric model
4. Neural GCM results
    a. Weather forecasting
    b. “Climate” simulations
5. Future directions

Pure ML models for weather forecasting

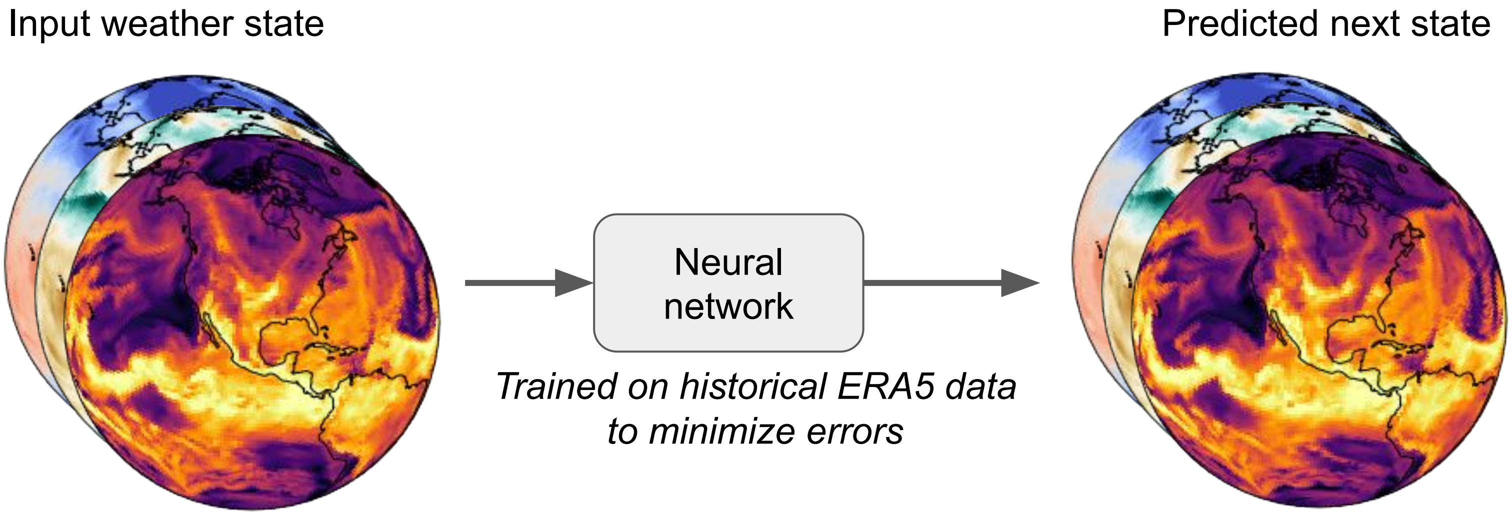

Recent disruptive results in medium-range weather forecasting

E.g., (GraphCast) [Lam & Sanchez-Gonzalez](#) et al. 2022, (Pangu weather) [Bi & Xie](#) et al. (2022)
(GenCast) [Price & Sanchez-Gonzalez](#) et al. 2024, (AIFS) [Lang & Alexe](#) et al. 2024

The AI revolution has arrived for weather forecasting

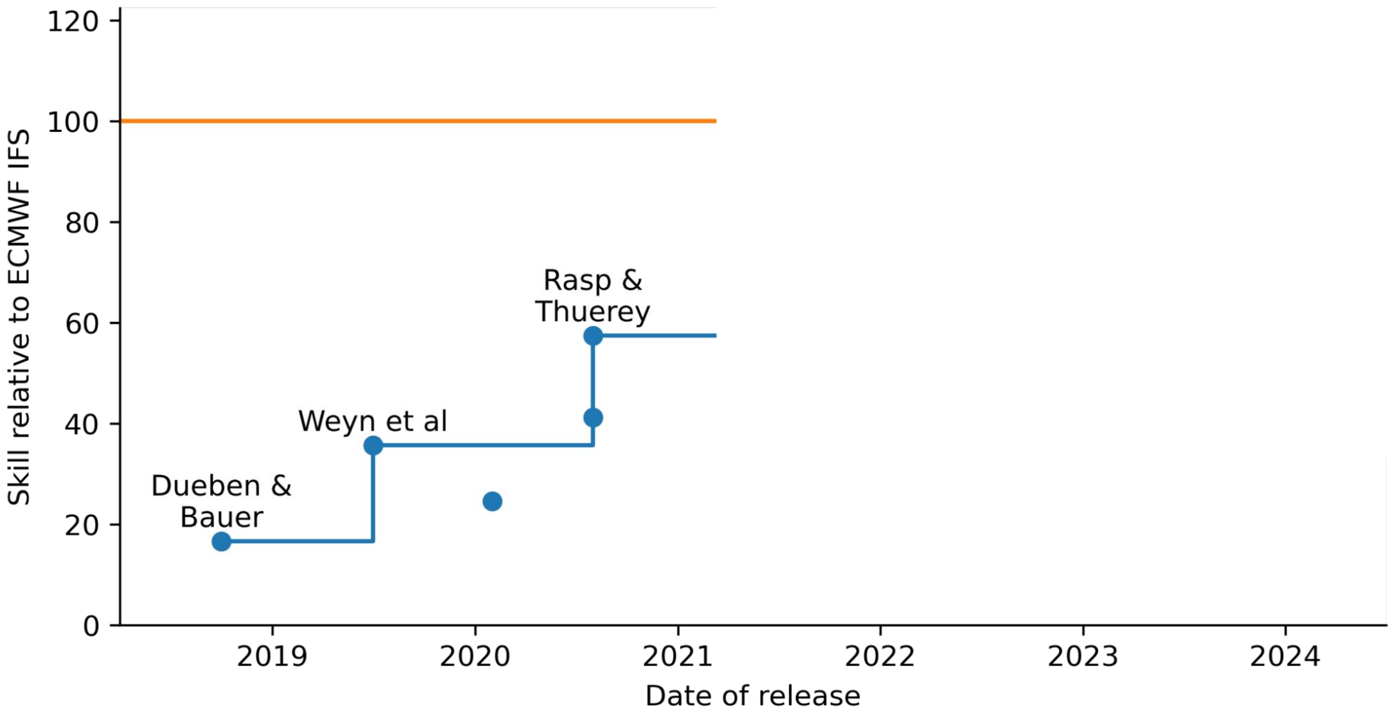

The AI revolution has arrived for weather forecasting

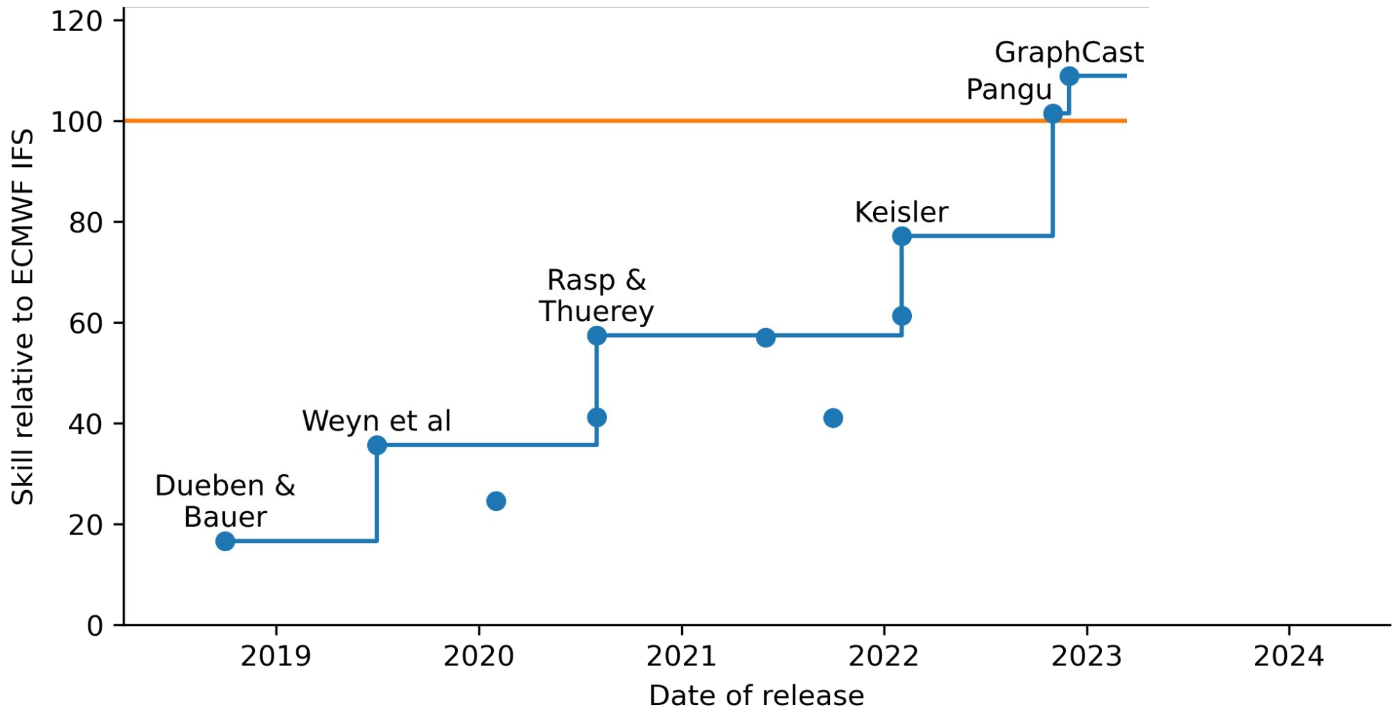

The AI revolution has arrived for weather forecasting

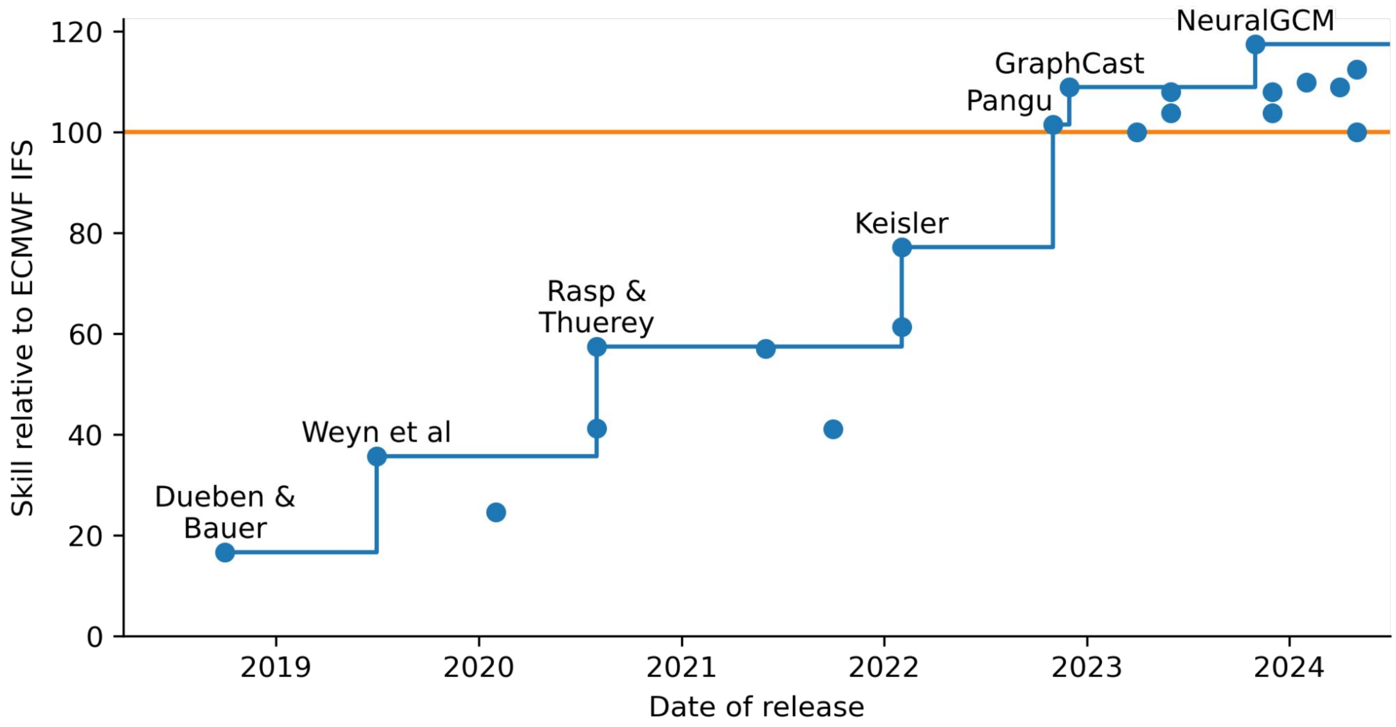

The AI revolution has arrived for weather forecasting

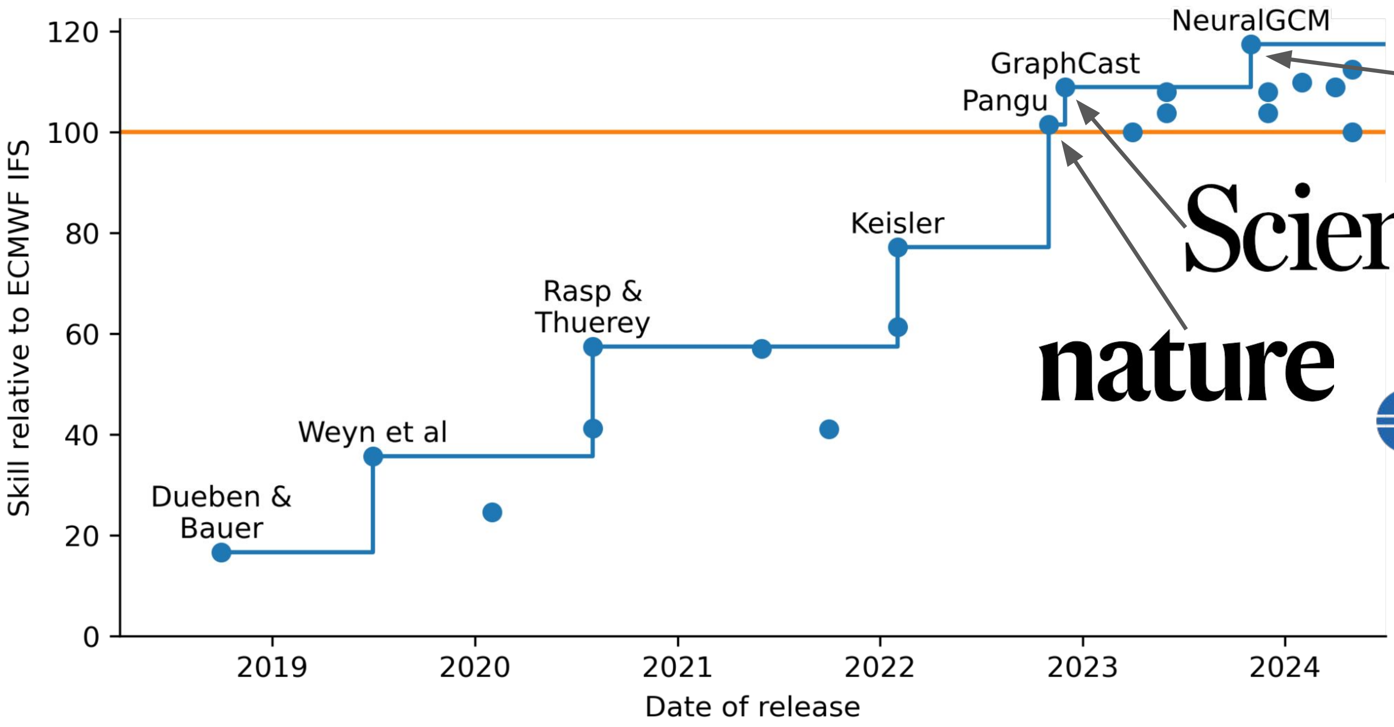

Many AI weather forecasts are skillful, but not yet fully physically realistic

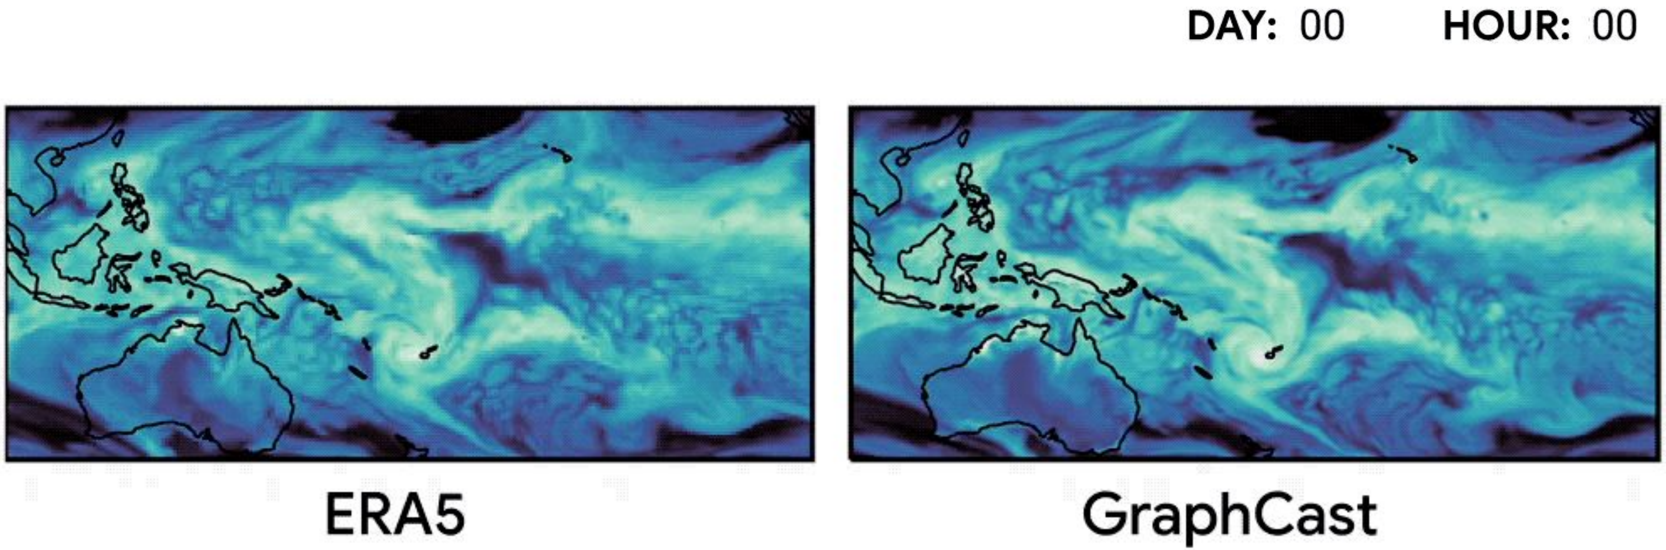

Desiderata: achieve good scores for good reasons

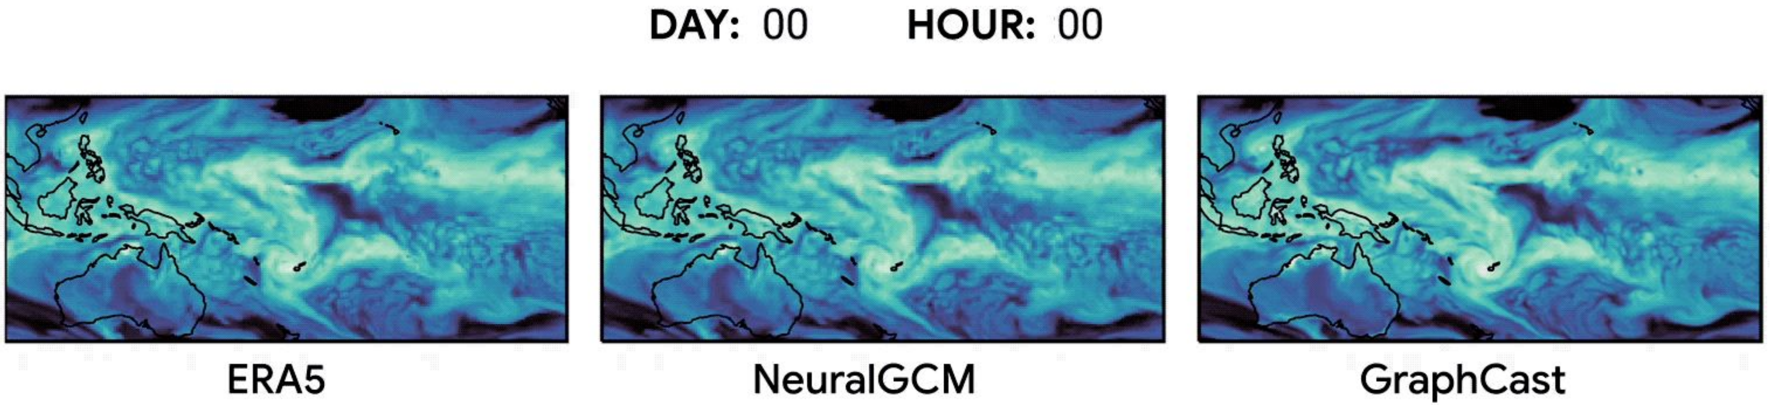

Google Research

Outline

1. General Circulation Models (GCMs) for weather and climate
2. AI revolution for weather forecasting
3. Neural GCM - differentiable hybrid atmospheric model
4. Neural GCM results
    a. Weather forecasting
    b. “Climate” simulations
5. Future directions

Hybrid modeling may offer the best of both worlds

Pure ML	*Hybrid models*	*Physics-based*
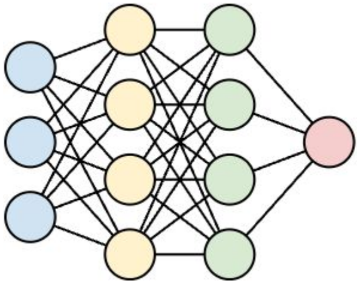	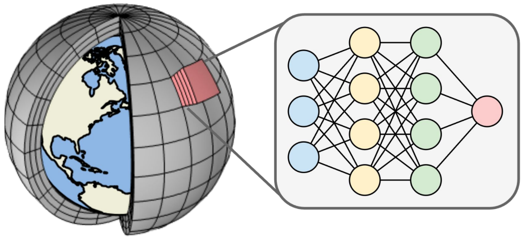	
GraphCast Pangu-Weather	NeuralGCM	Traditional NWP Climate models

Hybrid modeling may offer the best of both worlds

Pure ML

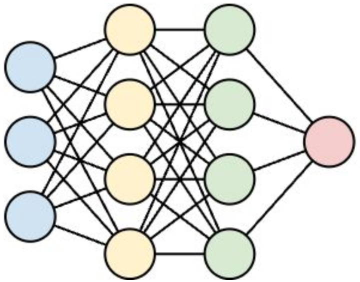

GraphCast
Pangu-Weather

Hybrid models

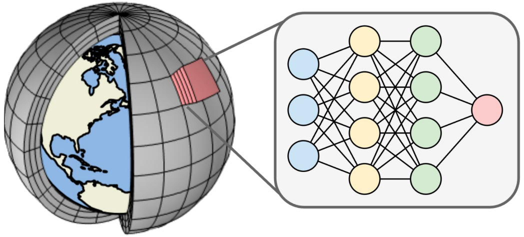

NeuralGCM

Physics-based

Traditional NWP
Climate models

Very little code
Based on data
Optimized for forecast accuracy

Hybrid modeling may offer the best of both worlds

Pure ML

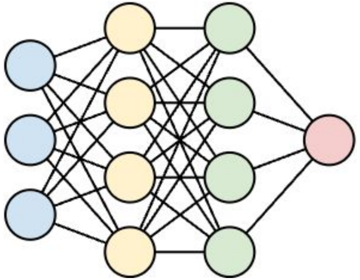

GraphCast
Pangu-Weather

Hybrid models

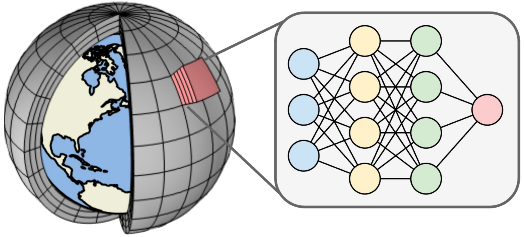

NeuralGCM

Physics-based

Traditional NWP
Climate models

Very little code
Based on data
Optimized for forecast accuracy

Complex, but interpretable
Based on physics
Designed to generalize

Traditional GCM modeling principle

“Dynamical core”

$$\frac{d\mathbf{u}}{dt} + f\mathbf{k}\times\mathbf{u} + \frac{1}{\rho}\nabla_z p = \mathbf{0}$$

$$\frac{\partial\rho}{\partial t} + \nabla_z\cdot(\rho\mathbf{u}) + \frac{\partial\rho w}{\partial z} = 0$$

$$\frac{dT}{dt} - \frac{\omega}{c_p\,\rho} = 0$$

$$\frac{\partial p}{\partial z} = -\rho g$$

$$p = \rho R T$$

Fluid dynamics on the surface of a rotating sphere

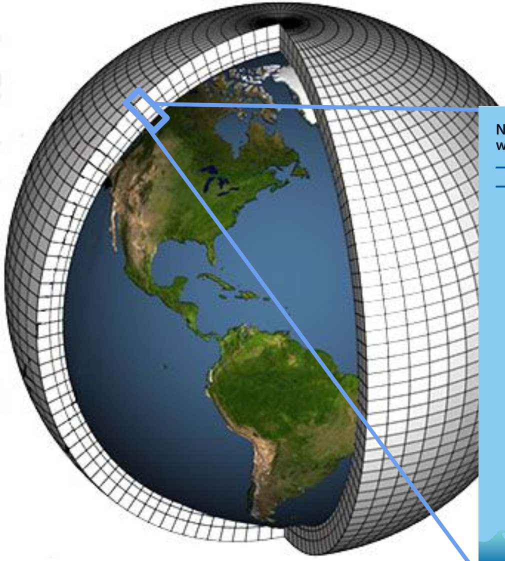

“Physics”

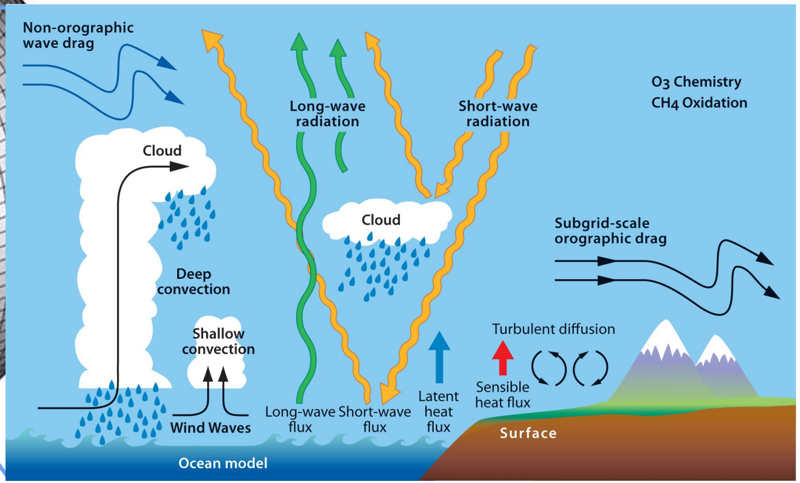

Many PhD theses
(100k-1M lines of Fortran)

Neural GCM modeling principle

"Differentiable dynamical core" (github.com/google-research/dinosaur)	**Physical processes parameterized by a NN**

$$\frac{d\mathbf{u}}{dt} + f\mathbf{k}\times\mathbf{u} + \frac{1}{\rho}\nabla_z p = \mathbf{0}$$

$$\frac{\partial \rho}{\partial t} + \nabla_z \cdot (\rho\mathbf{u}) + \frac{\partial \rho w}{\partial z} = 0$$

$$\frac{dT}{dt} - \frac{\omega}{c_p\,\rho} = 0$$

$$\frac{\partial p}{\partial z} = -\rho g$$

$$p = \rho R T$$

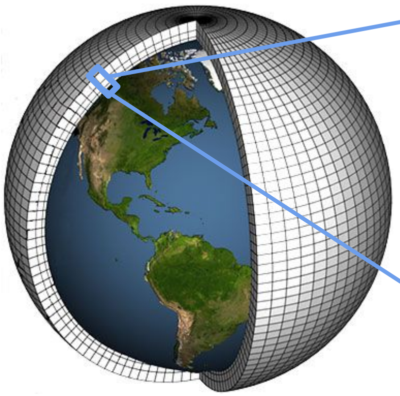

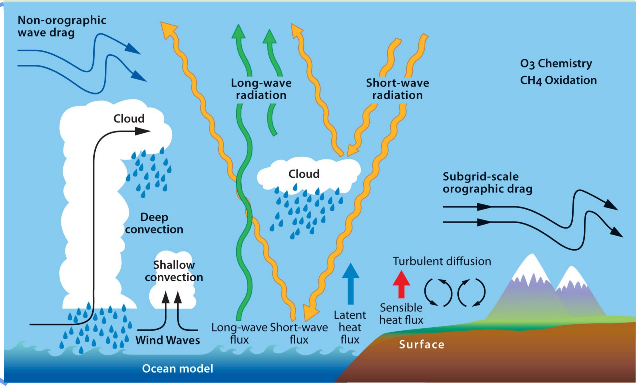

Trained end-to-end through 10-1000 of time steps (~3 simulation days)

Features & hypotheses:

- Learns mechanistic "Physics" to drive the "Dynamics"
- Accounts for feedbacks between dynamics & parameterizations

Our dynamical core solves the moist primitive equations with spectral methods

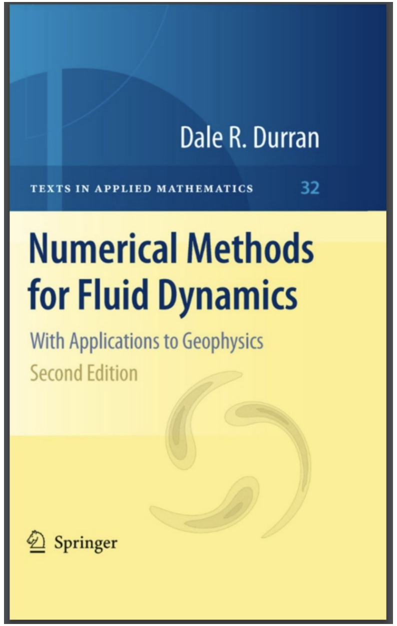

Written in JAX and runs fast on GPUs and Google TPUs

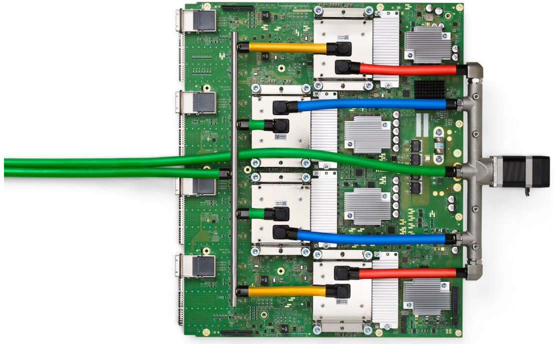

Up to 16x model parallelism

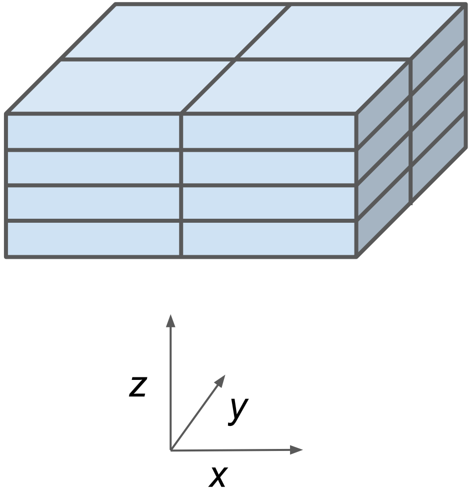

Neural GCM model overview

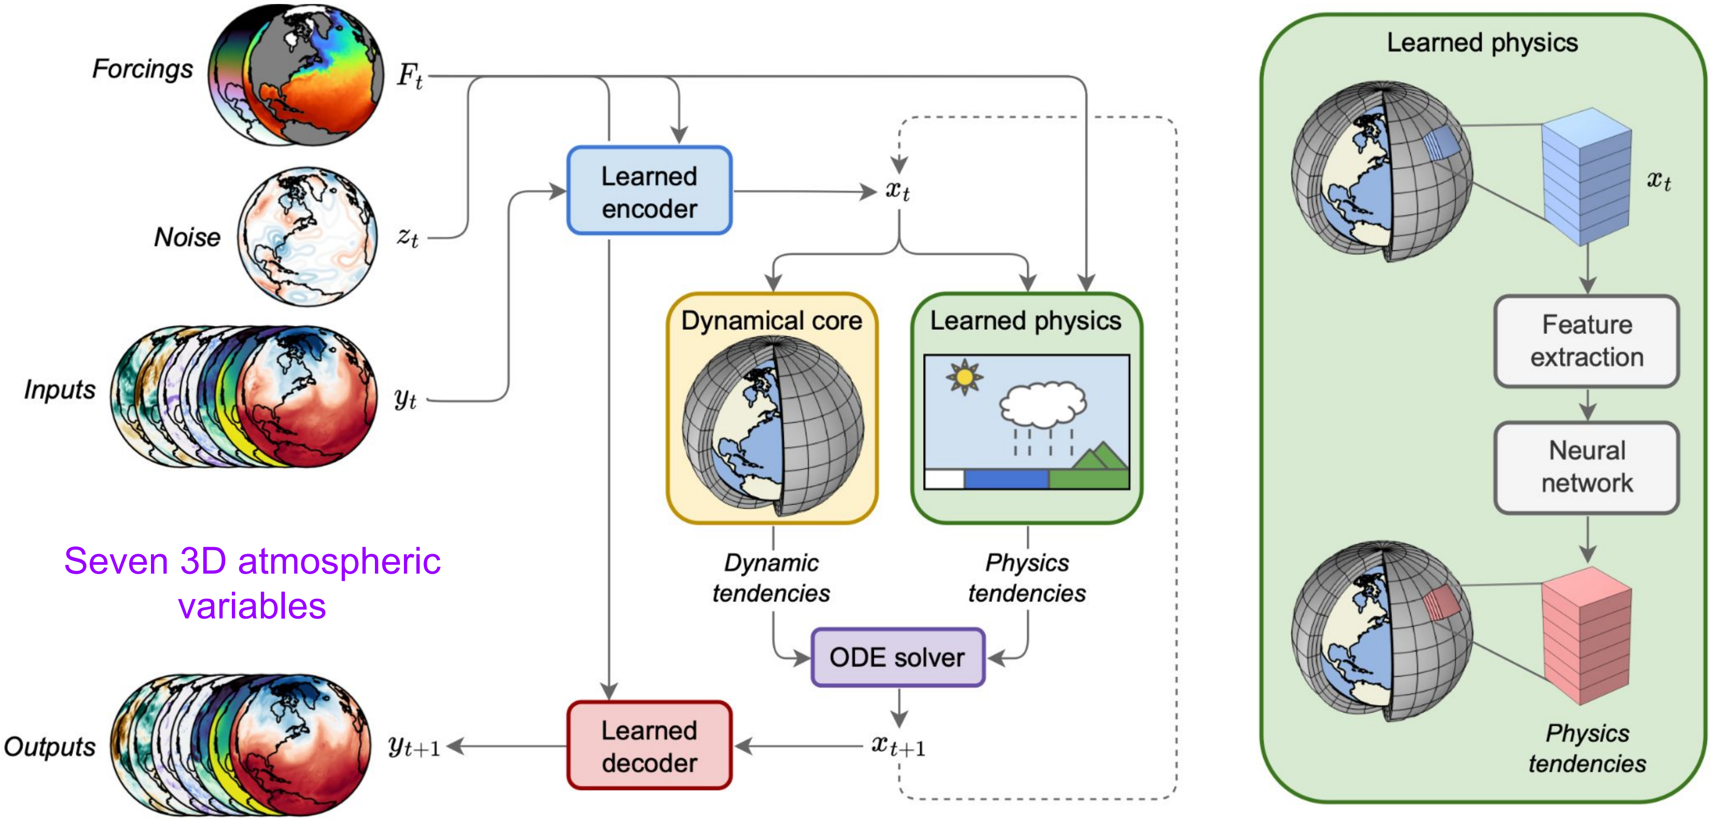

Seven 3D atmospheric variables

Model trained to minimize discrepancies between *Outputs* and ERA5 data

Outline

1. General Circulation Models (GCMs) for weather and climate
2. AI revolution for weather forecasting
3. Neural GCM - differentiable hybrid atmospheric model
4. Neural GCM results
    a. Weather forecasting
    b. “Climate” simulations
5. Future directions

NeuralGCM achieves state-of-the-art results both for weather forecasting and climate simulation

(1) Competitive 1-15 day ***ensemble weather*** forecasts with ECMWF

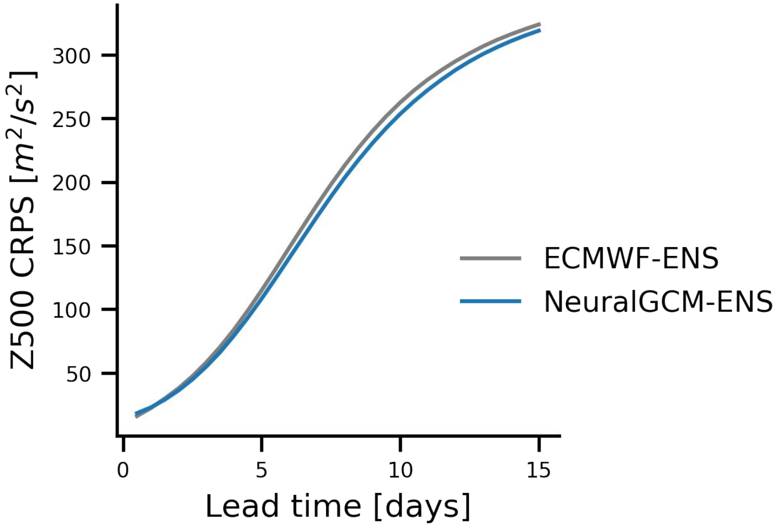

(2) Realistic **year-to-decades runs,** competitive with atmosphere only (AMIP) climate models

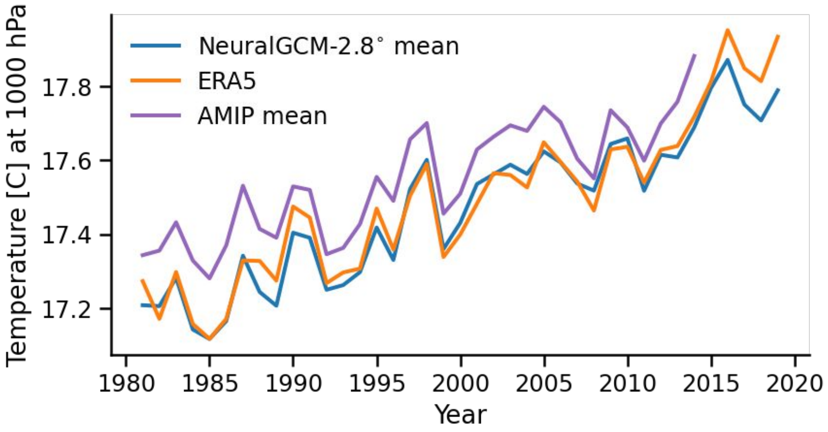

Weather “Turing test”: which one is ERA5?

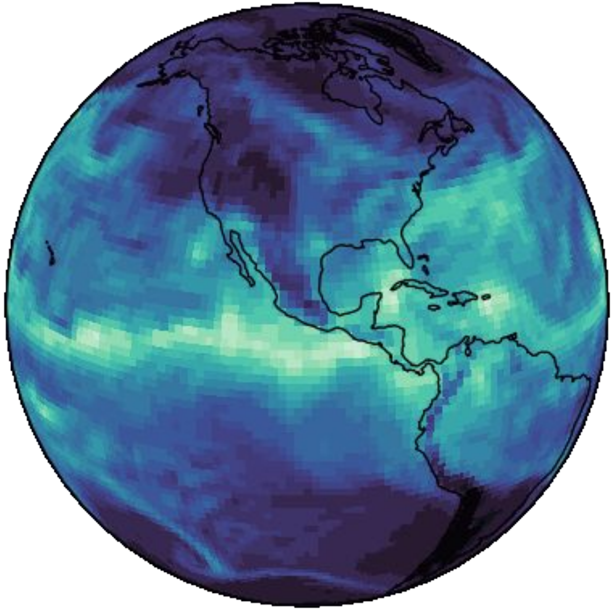

Option A

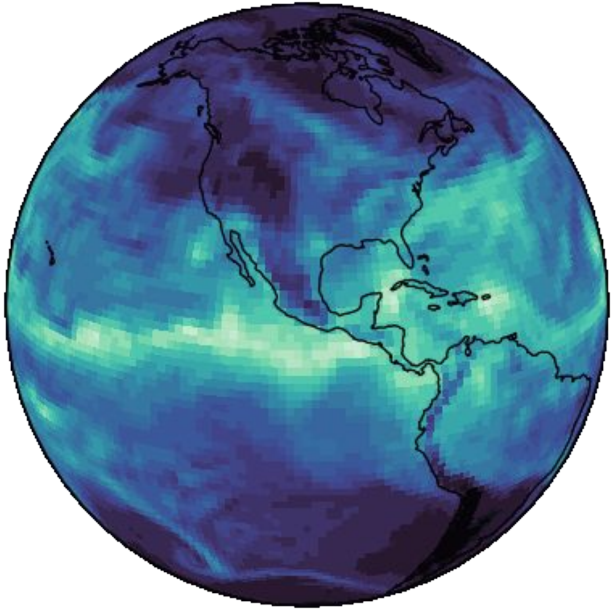

Option B

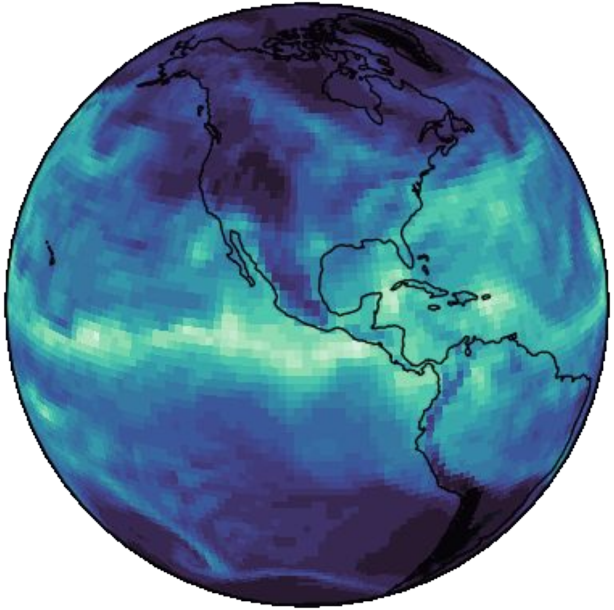

Option C

Weather “Turing test”: which one is ERA5?

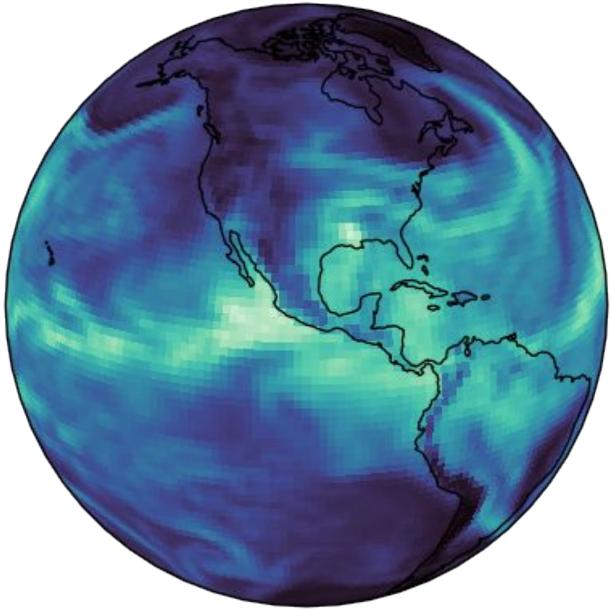

Option A

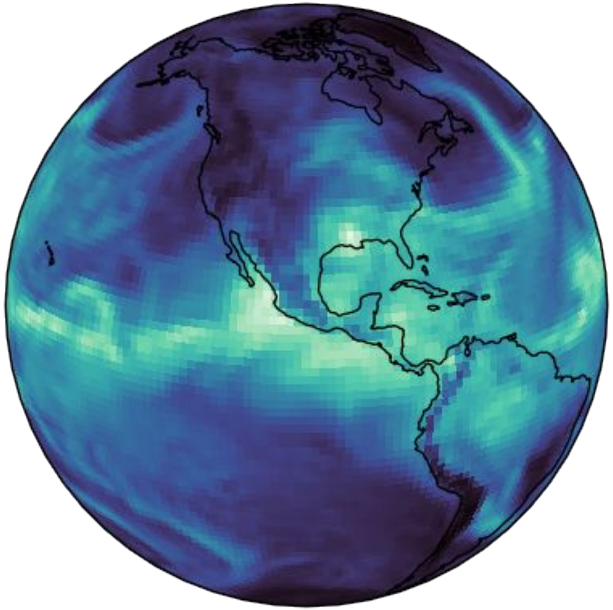

Option B

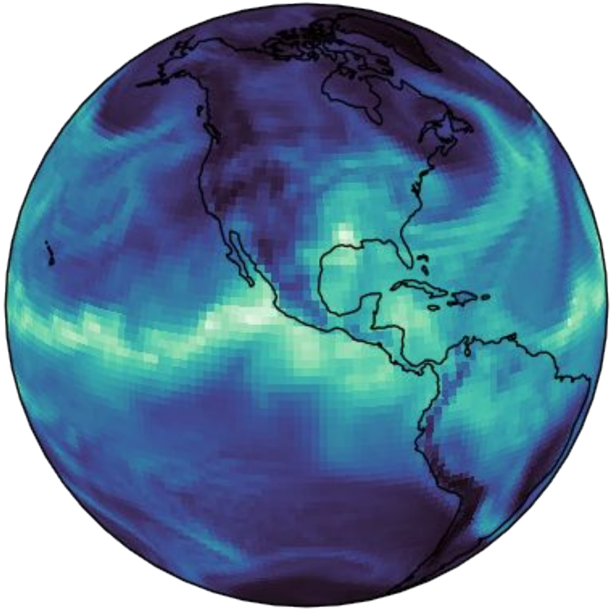

Option C

Weather “Turing test”: which one is ERA5?

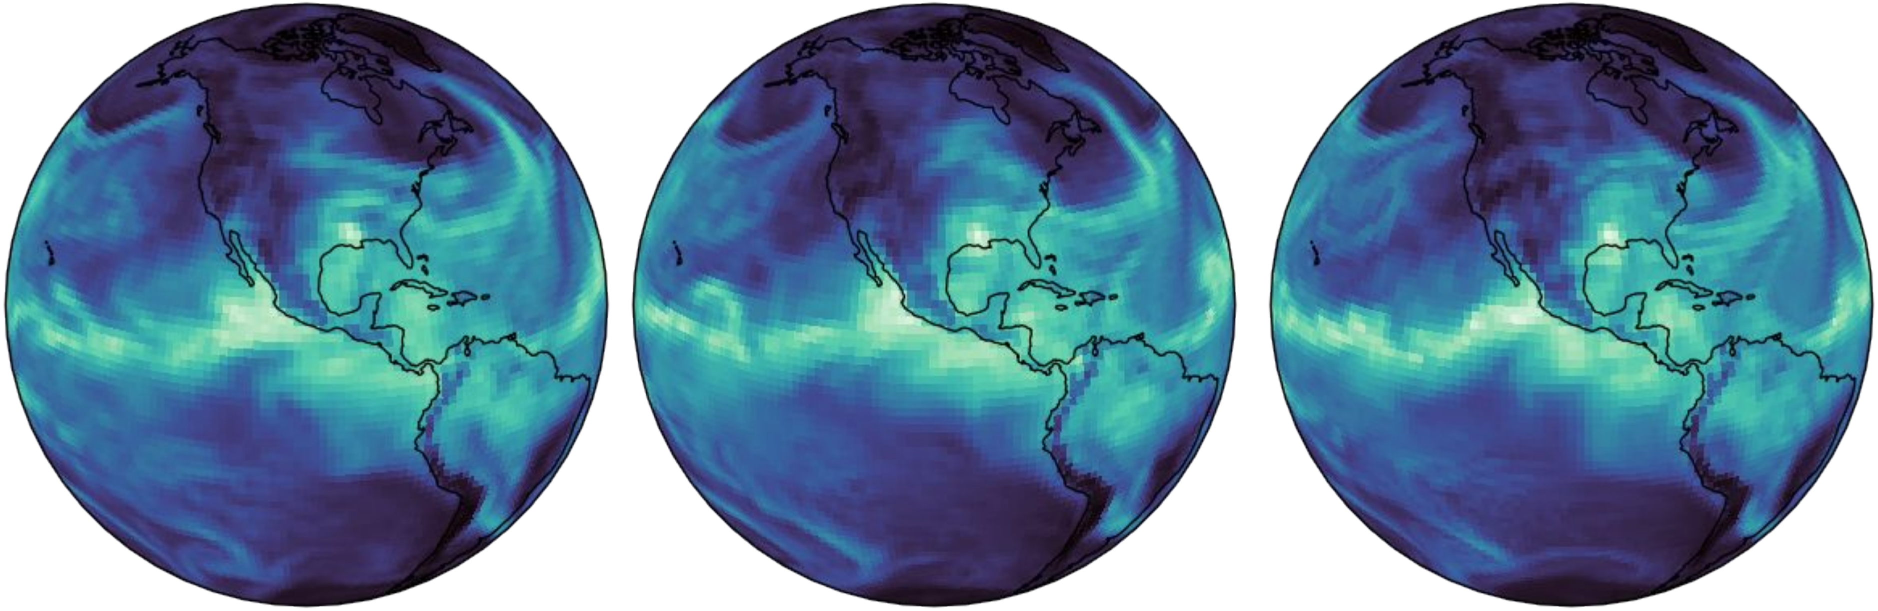

NeuralGCM NeuralGCM ERA5

Precipitation climate “Turing test”: which one is ERA5?

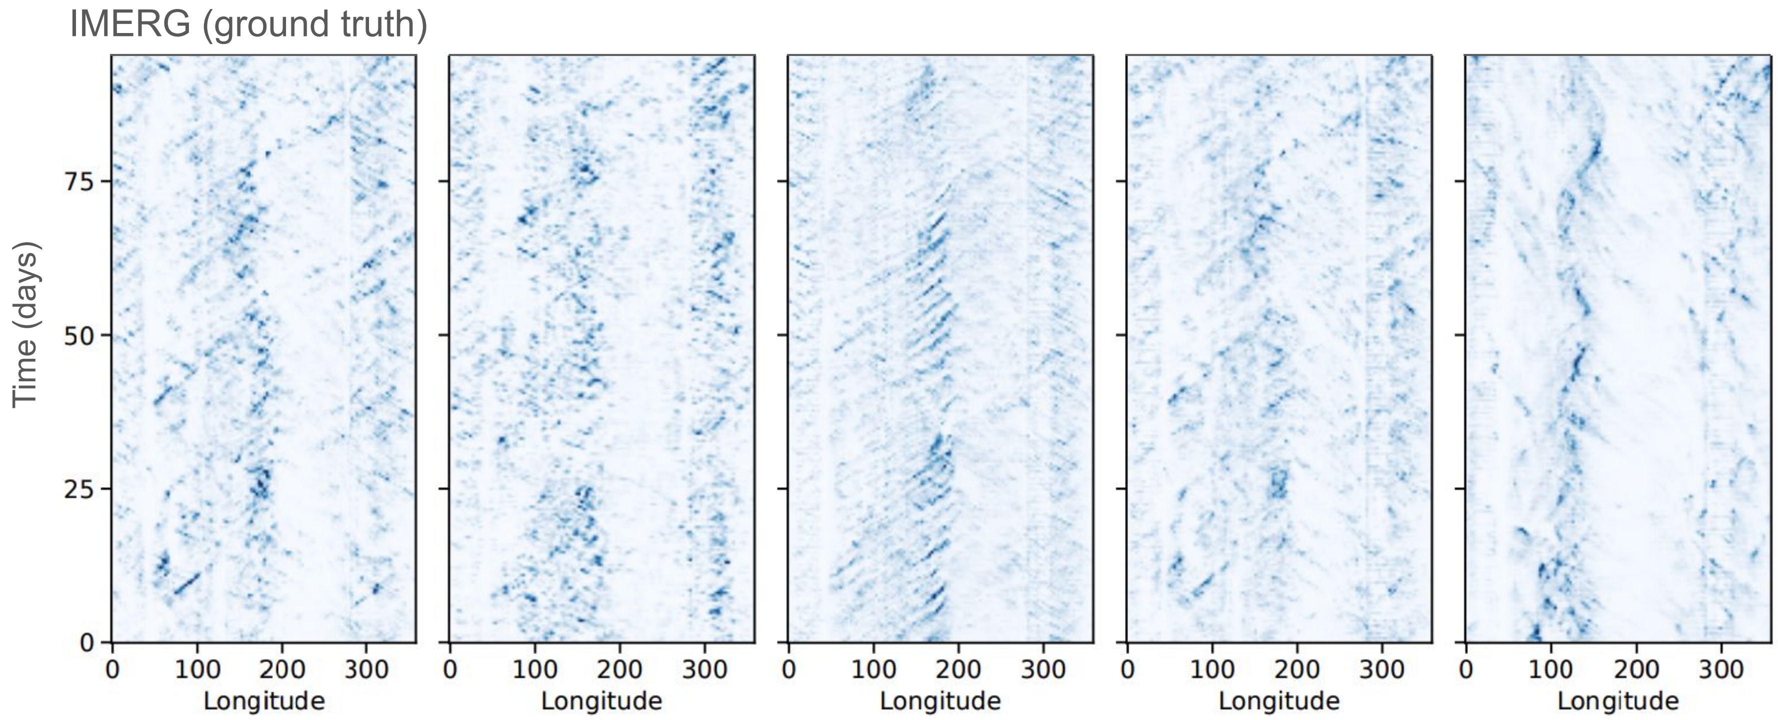

From NeuralGCM trained on combined ERA5 & IMERG data

Precipitation climate “Turing test”: which one is ERA5?

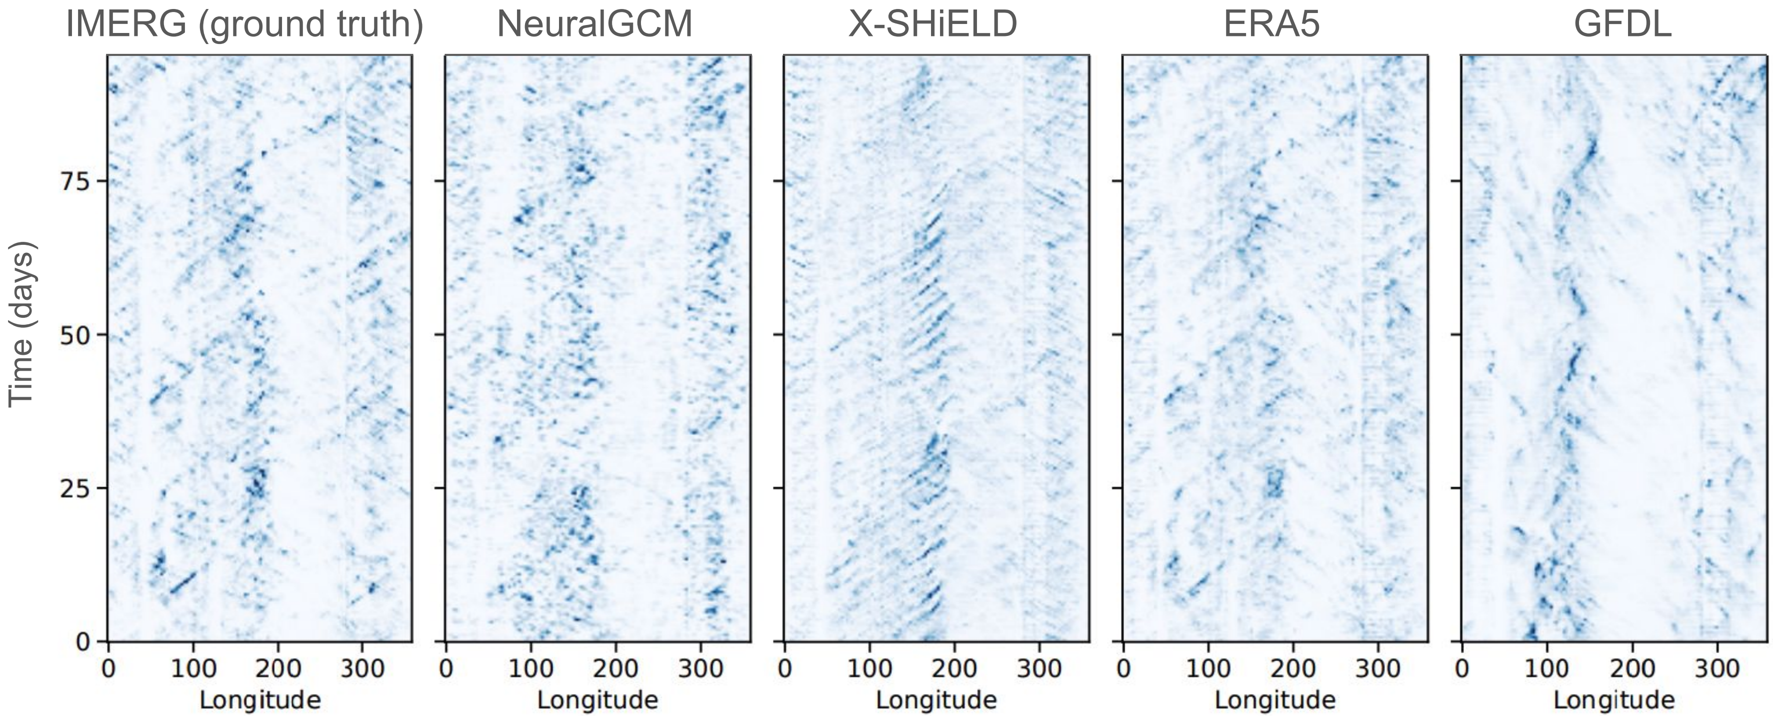

From NeuralGCM trained on combined ERA5 & IMERG data

Google Research

Outline

1. General Circulation Models (GCMs) for weather and climate
2. AI revolution for weather forecasting
3. Neural GCM - differentiable hybrid atmospheric model
4. Neural GCM results
    a. Weather forecasting
    b. “Climate” simulations
5. Future directions

Evaluation of weather forecasts

Setup:

1. Models trained using historical data until 2020
2. Evaluate 7-15 day forecasts issued for 2020 initialized every 12h

Competing approaches:

- Operational physics models
  - ECMWF HRES
  - ECMWF ENS
- ML models
  - GraphCast
  - Pangu

Criteria, driving questions and (metrics):

1. Forecast accuracy — Does the forecast track weather patterns accurately? (RMSE, CRPS)
2. Physical consistency — Does it look like weather? (spectral density, biases)

WeatherBench 2

Stephan Rasp et al

github.com/google-research/weatherbench2

RMSE and CRPS scores

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(Y_i - \hat{Y}_i)^2$$

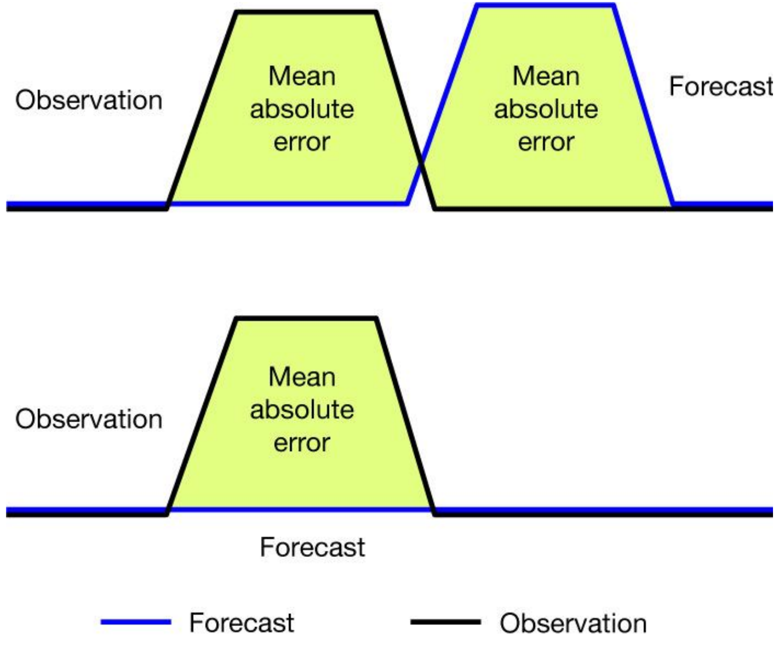

Figure credit: ECMWF

$$CRPS = \frac{1}{n}\sum_{i=1}^{n}(|Y_i - \hat{Y}_i| - \frac{1}{2}|Y_i - Y_i^{'}|)$$

Encourages accuracy w.r.t. truth

Encourages spread within ensemble

When minimized, one expects similar deviations between the ensemble members (Y, Y') and the ground truth.

Accuracy: RMSE scores

- At short times **NeuralGCM-0.7°** and **GraphCast** achieve lowest errors
- At 5-7 days ensemble mean of **NeuralGCM-ENS** and **ECMWF-ENS** perform best

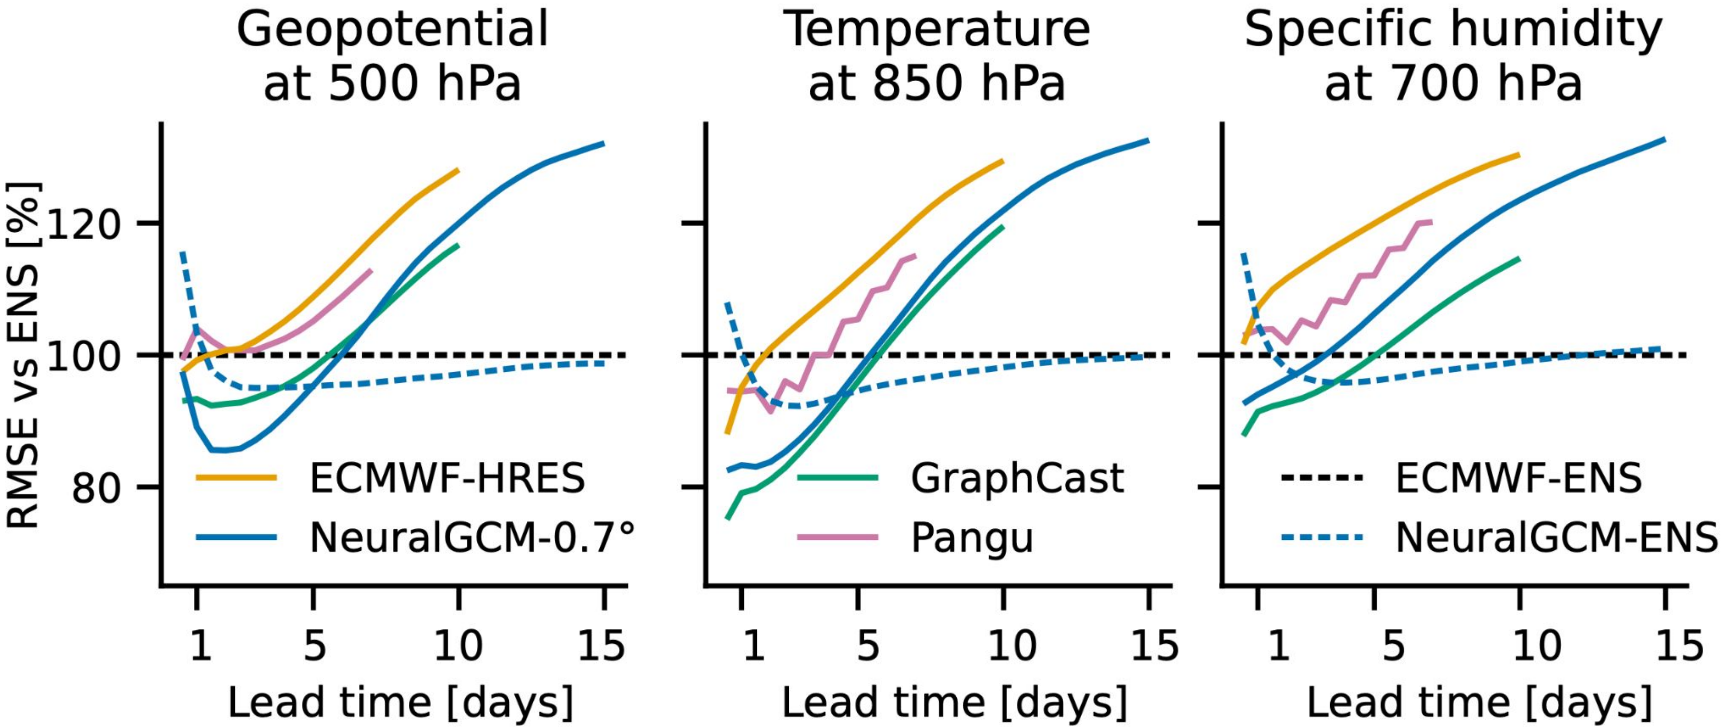

Accuracy: CRPS scores

- **NeuralGCM-ENS** slightly outperforms **ECMWF-ENS** in Continuous Rank Probability Score

CRPS is the training objective

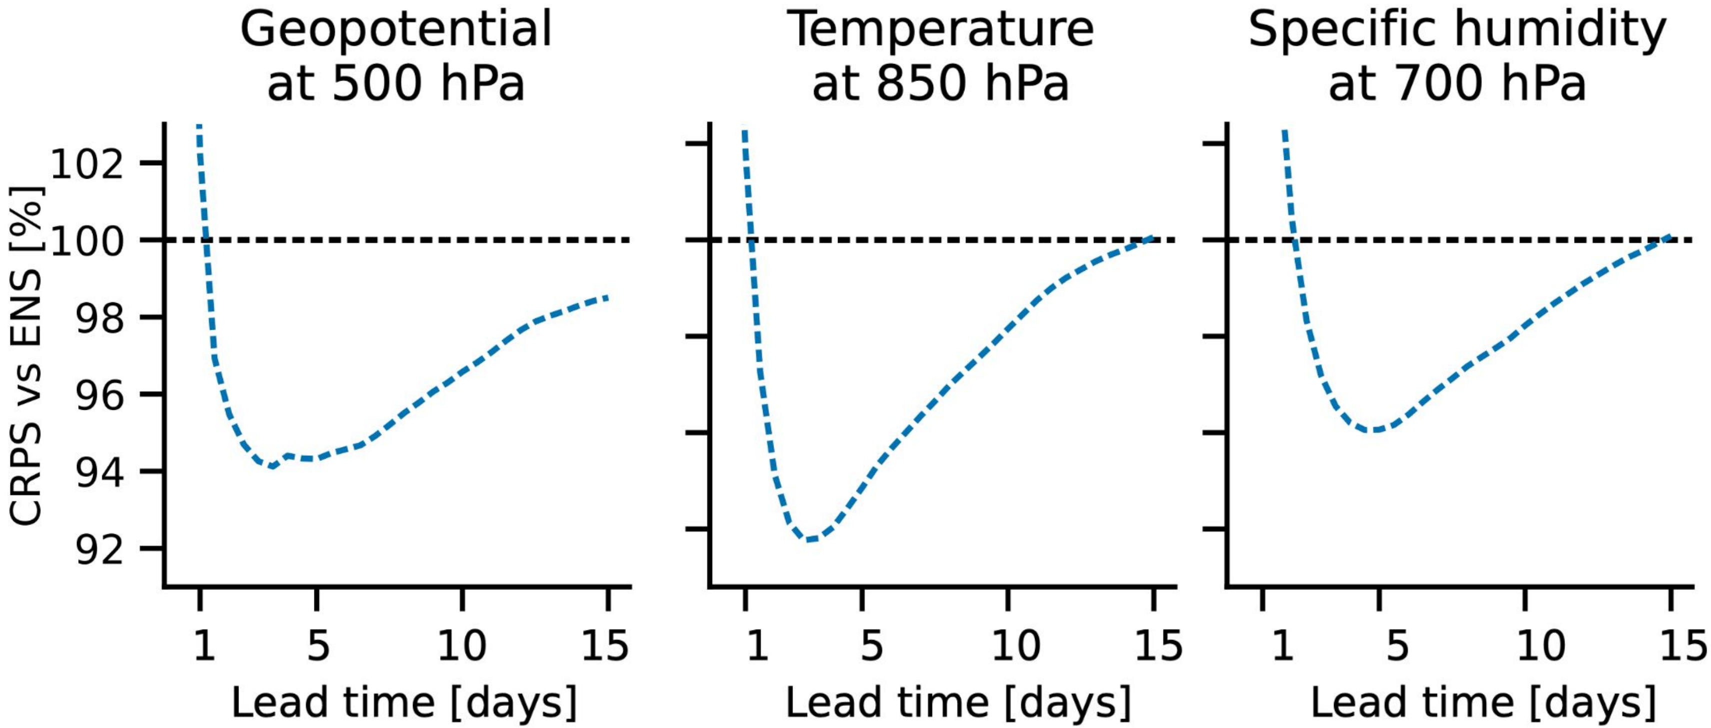

Consistency: Spectral density and RMSB

Spectral density
Evaluates sharpness and stationarity

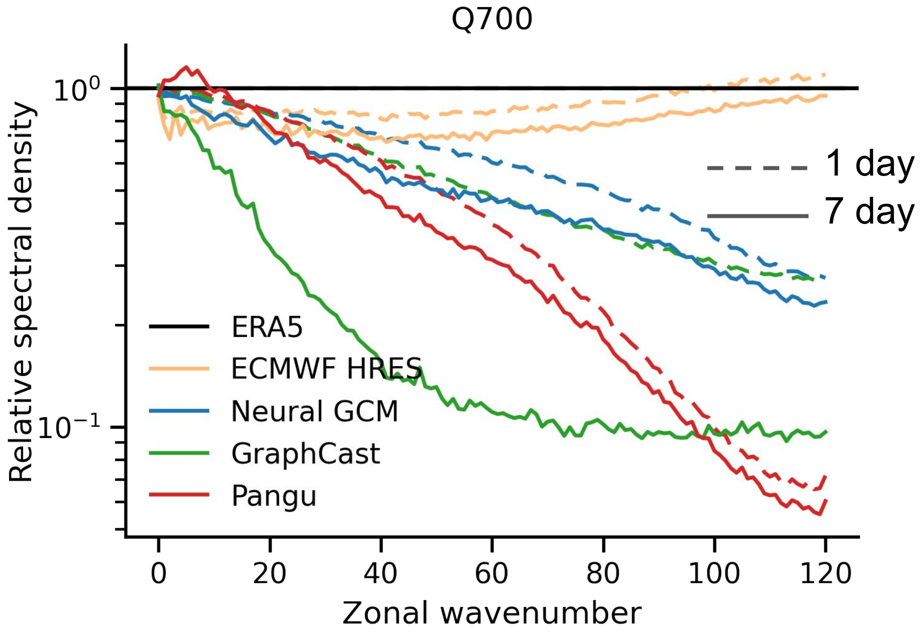

Bias RMSE
Evaluates persistent error accumulation

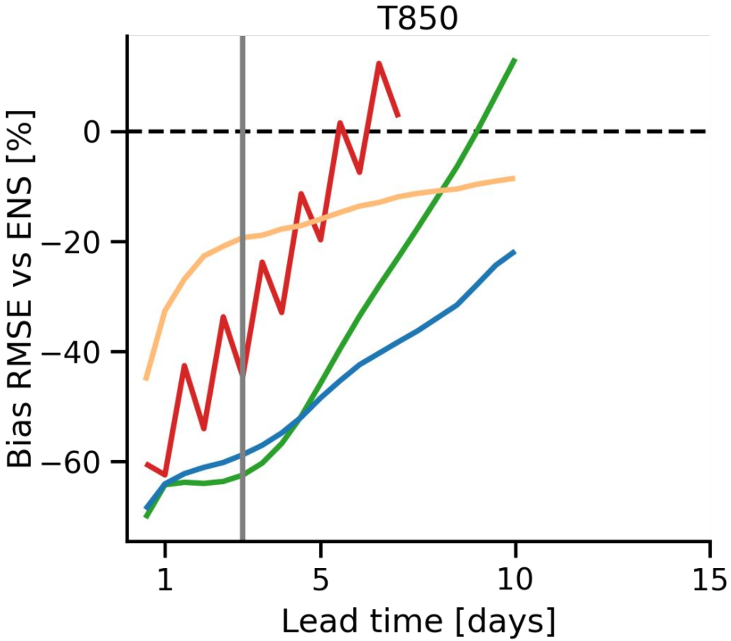

Consistency: bias and spread-skill spatial distributions

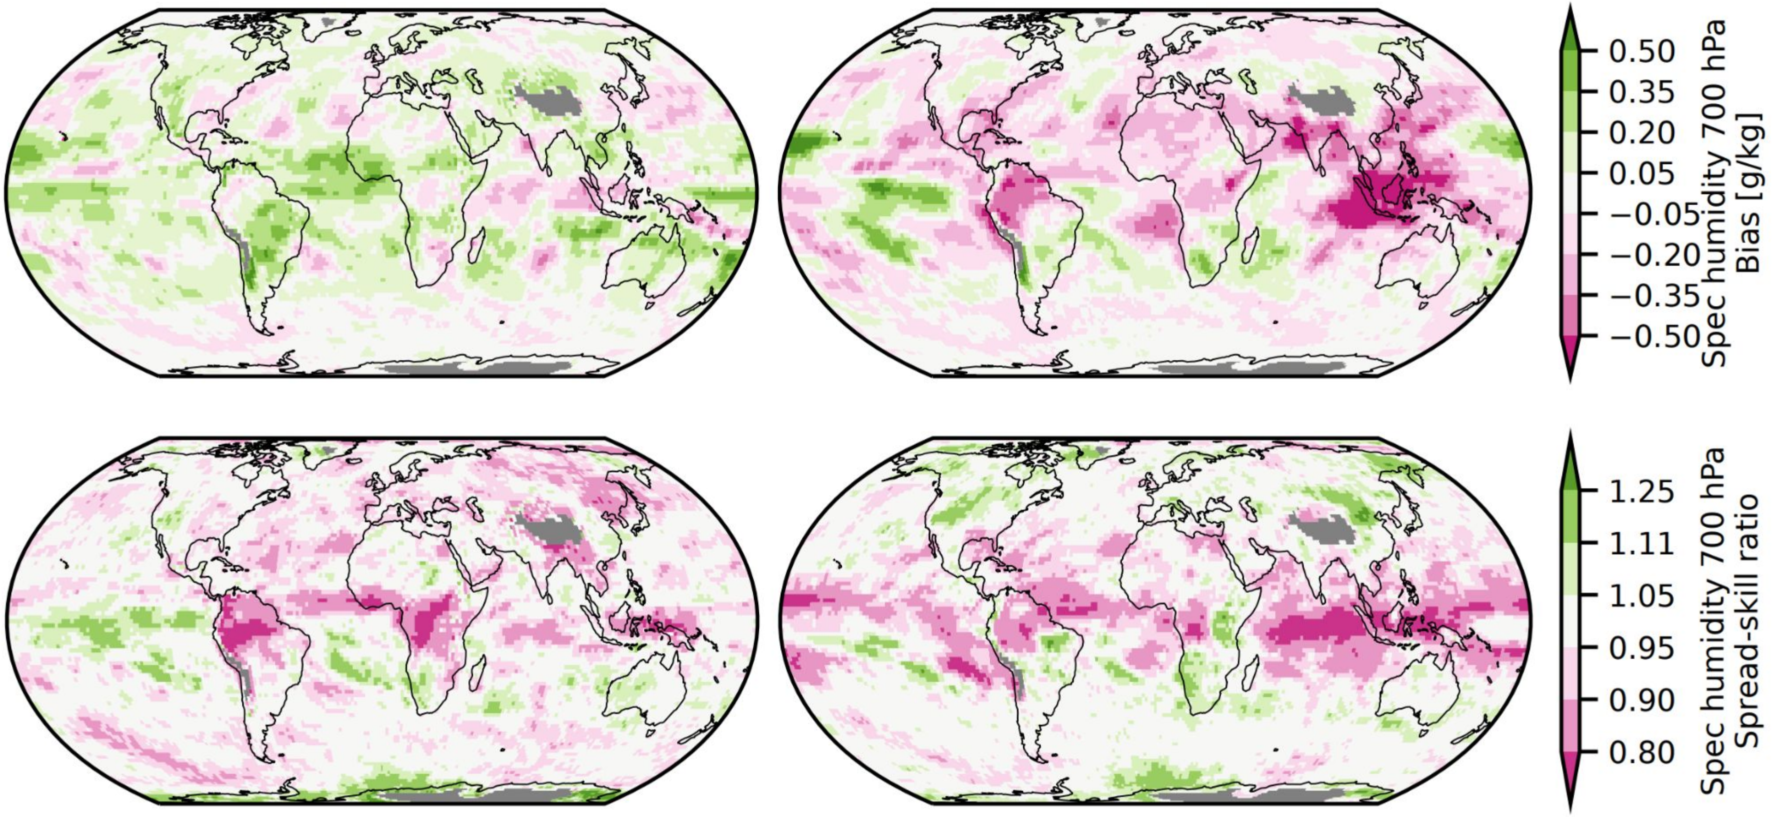

Consistency: water budget from model architecture

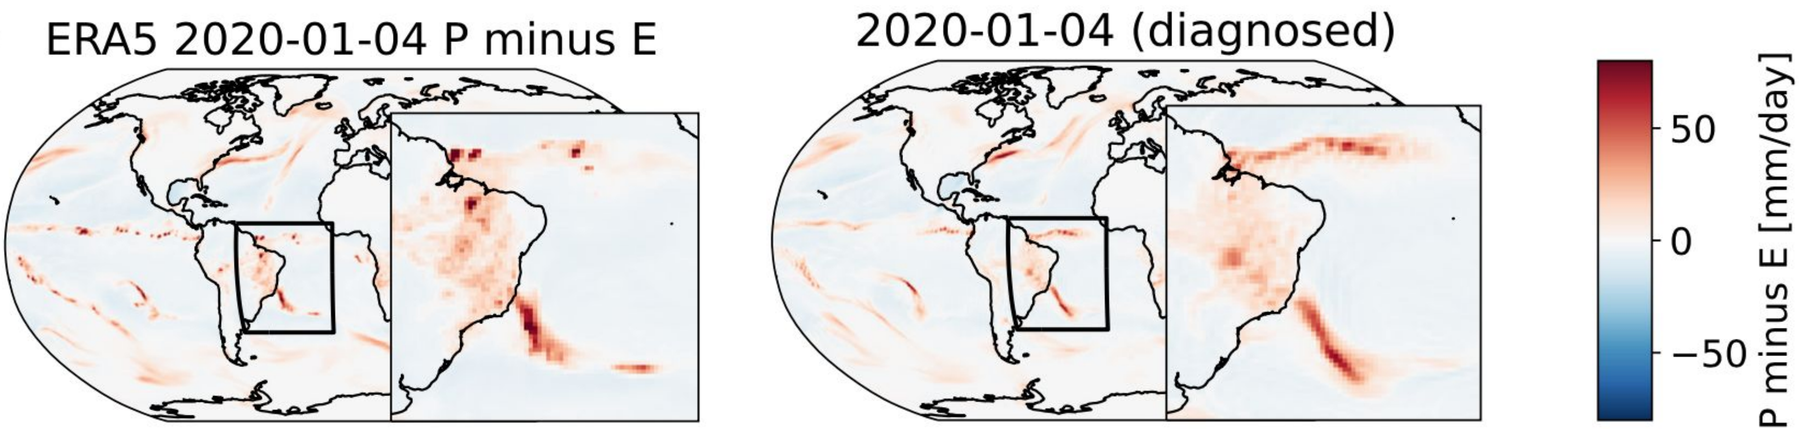

NeuralGCM At 3 day lead time

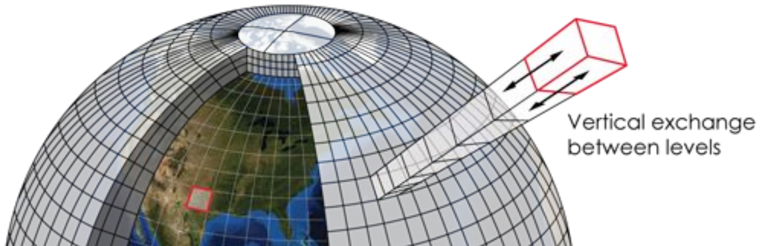

"Dynamics" + "Physics" separation enables us to directly diagnose changes in moisture (precipitation - evaporation)

Case study: Ensemble of tropical cyclone tracks

ECMWF ensemble

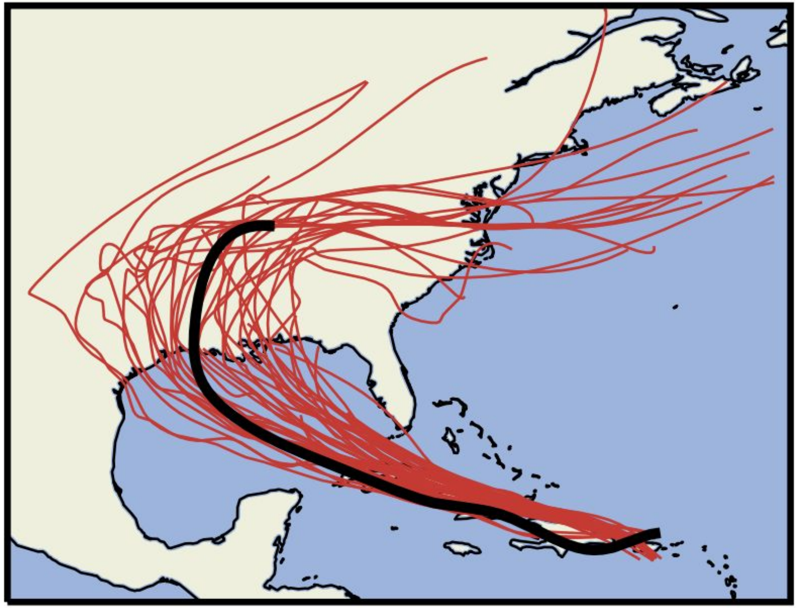

NeuralGCM ensemble

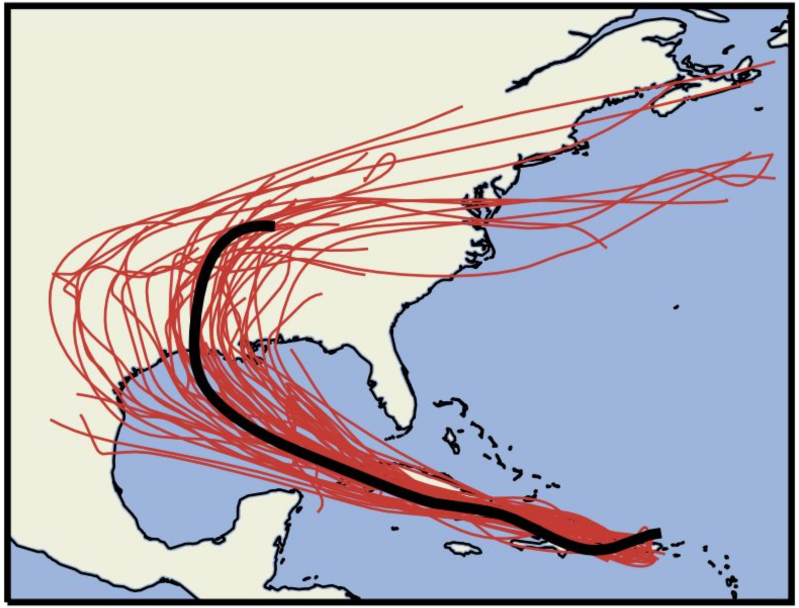

+5 day forecasts of 2020’s Hurricane Laura

Outline

1. General Circulation Models (GCMs) for weather and climate
2. AI revolution for weather forecasting
3. Neural GCM - differentiable hybrid atmospheric model
4. Neural GCM results
    a. Weather forecasting
    b. “Climate” simulations
5. Future directions

Emergent long-term behaviors in Neural GCM

Multi-year long simulations with **prescribed ocean**

Neural-GCM trained on 3-day rollouts

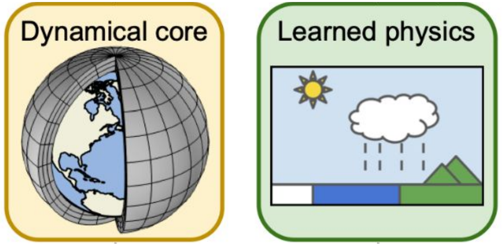

Monsoon wind reversal

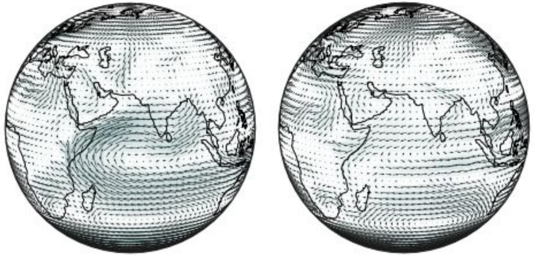

Seasonal temperature cycles

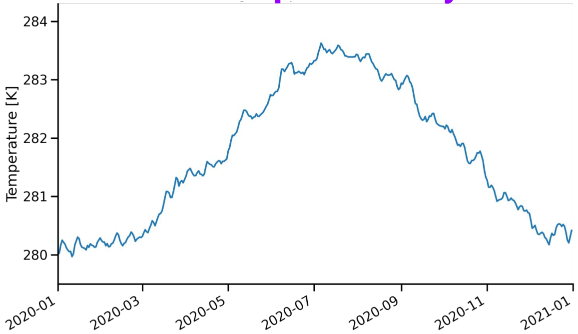

Global warming trends

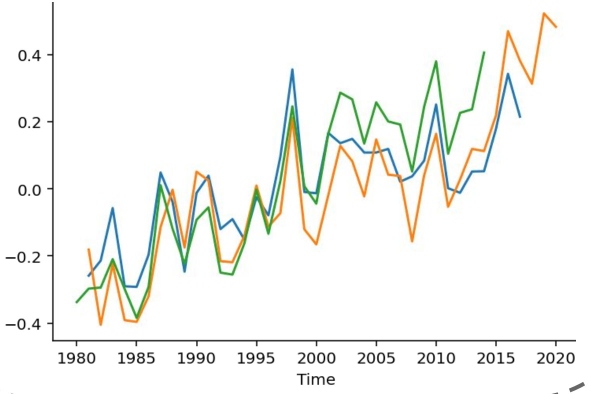

Low biases in temperature & moisture

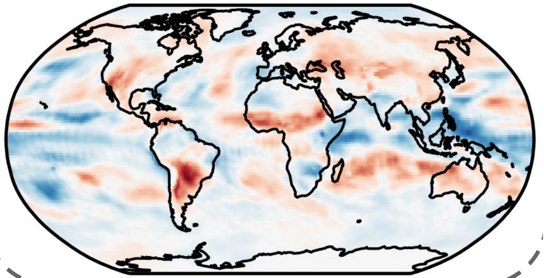

Generation of realistic tropical cyclones

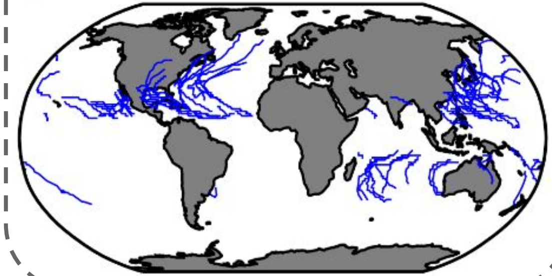

Evaluation of climate simulations

Setup:

1. Models trained on historical data until 2017
2. 2-year simulations initialized throughout 2019 use 1.4° (140 km) NeuralGCM
3. 40-year run initialized in 1980 use 2.8° (280 km) NeuralGCM

Inference is **AMIP** setup – prescribed historical sea surface temperature

Reference models:

- X-SHiELD - state of the art cloud resolving model (3 km resolution)
- CESM - state of the art climate model
- Climatology - predict average climate

Desiderata:

1. Climate-like variability
2. No significant climate drift or bias

NeuralGCM runs fast on modern hardware

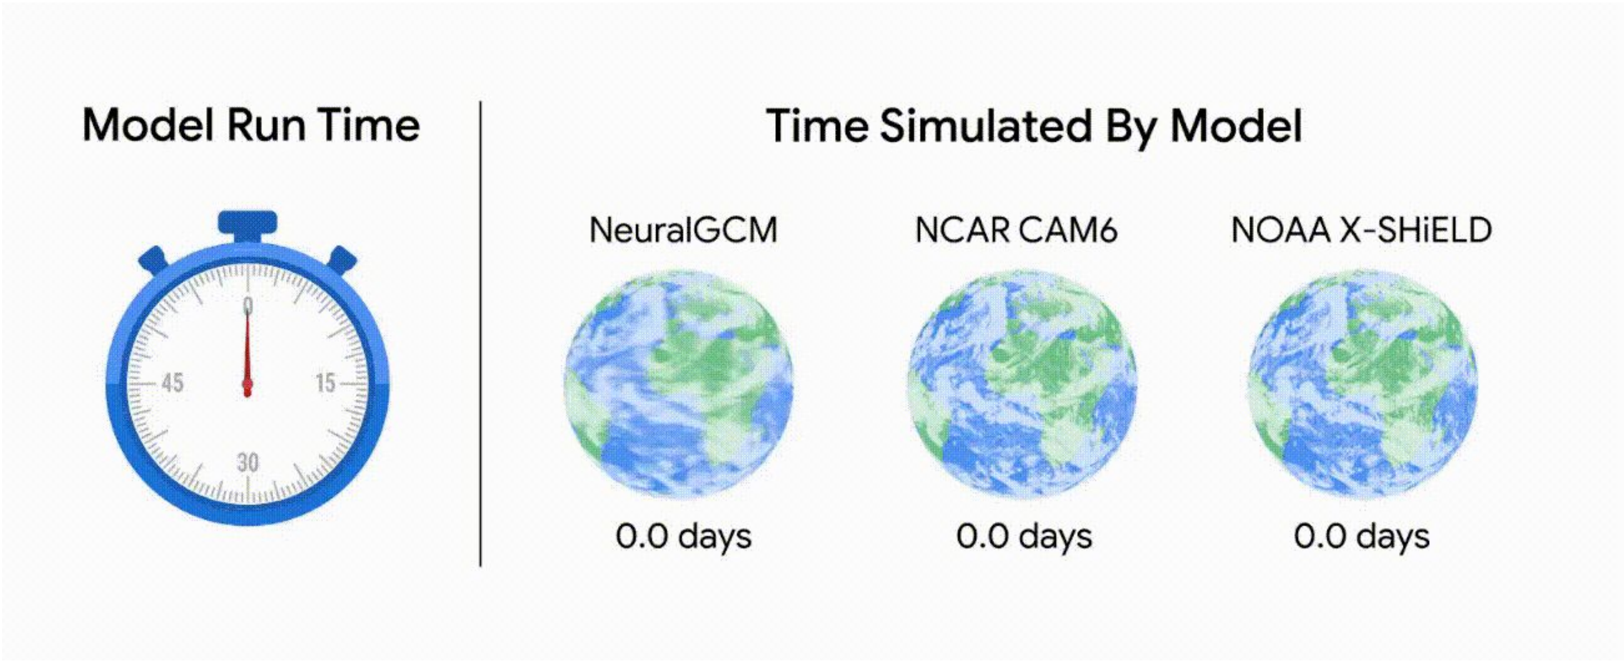

This is largely the consequence of reduced resolution

Neural GCM captures near-term climate in 1+ year forecasts

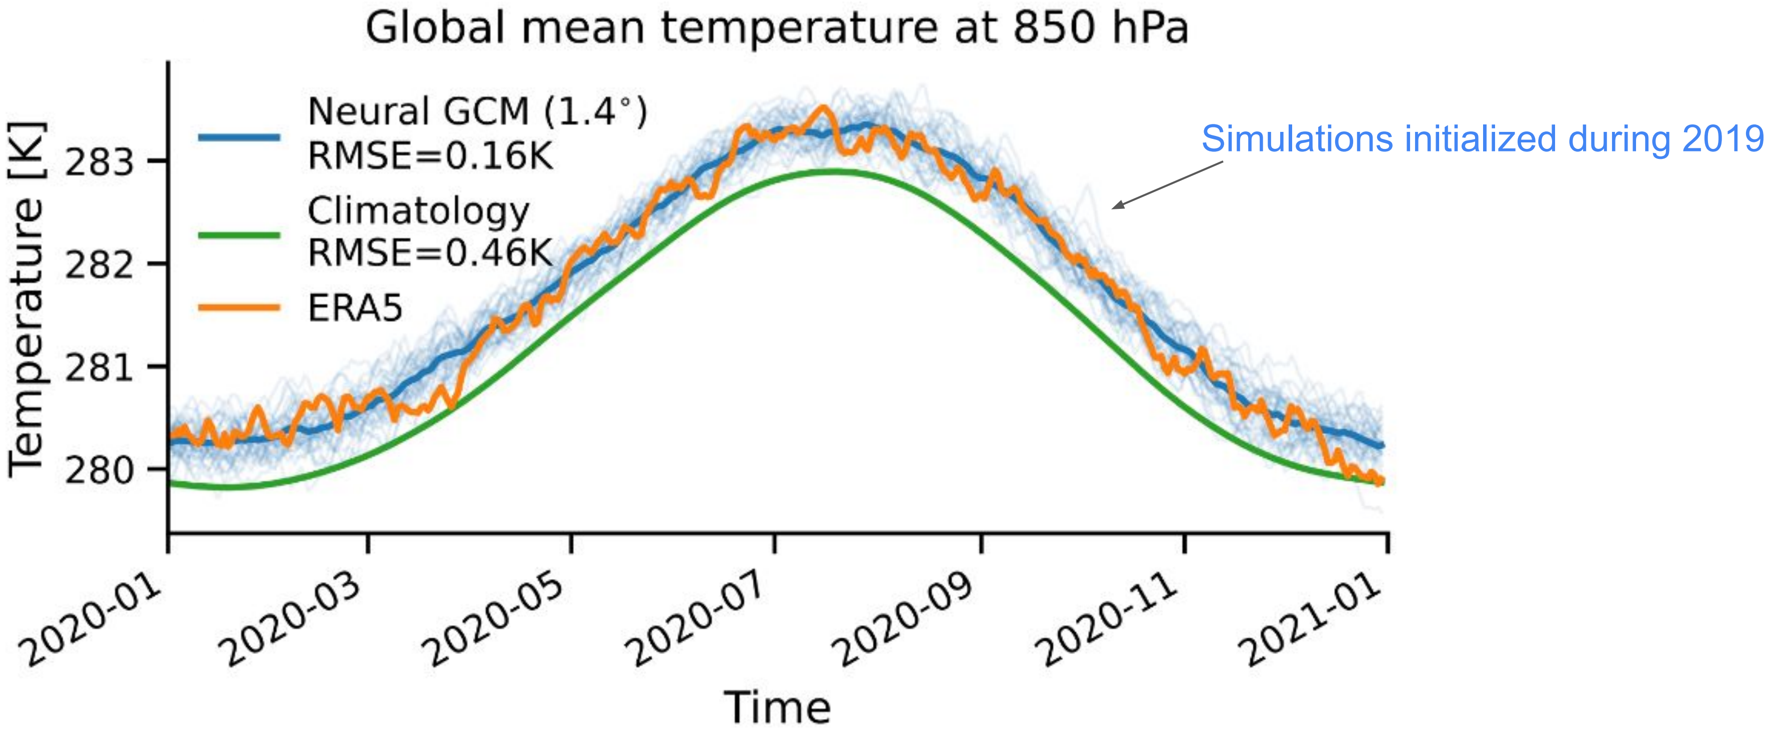

NeuralGCM reproduces near-term climate more accurately than global storm resolving models

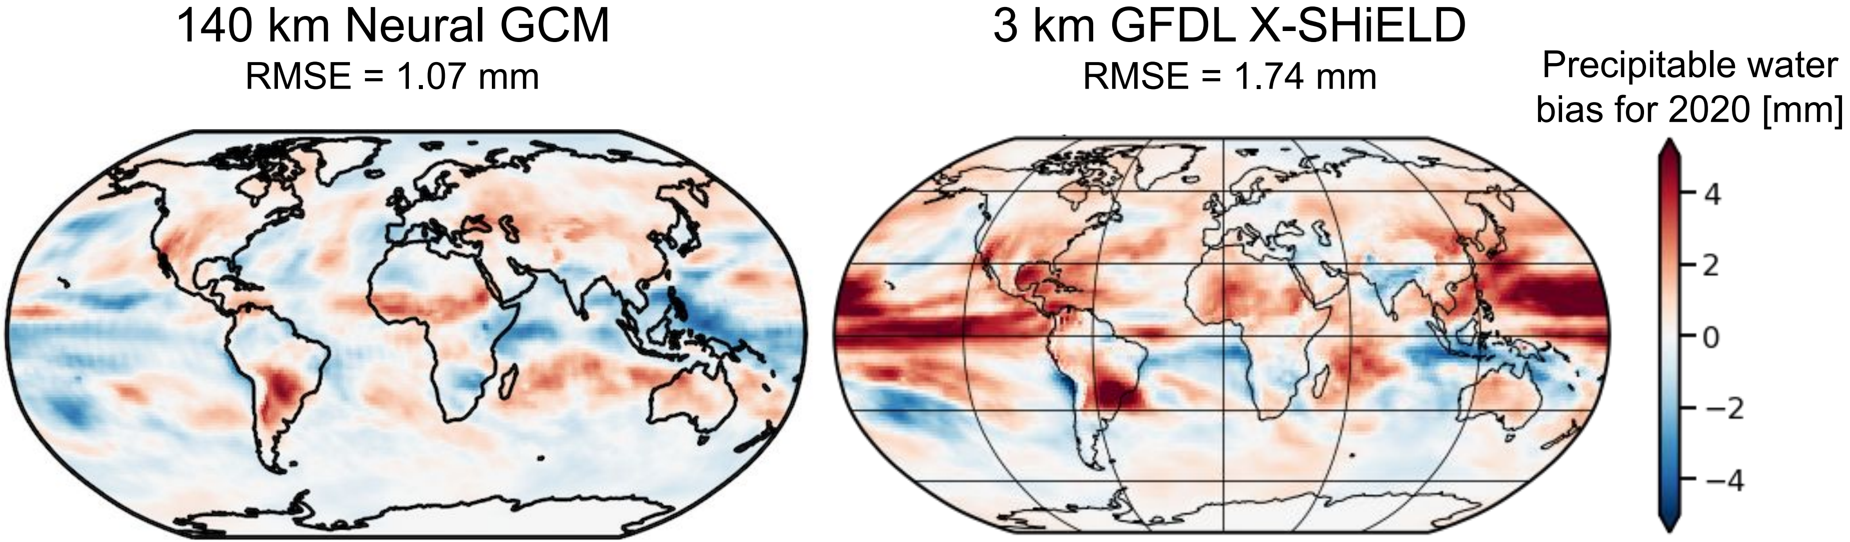

Neural GCM generates realistic tropical cyclones

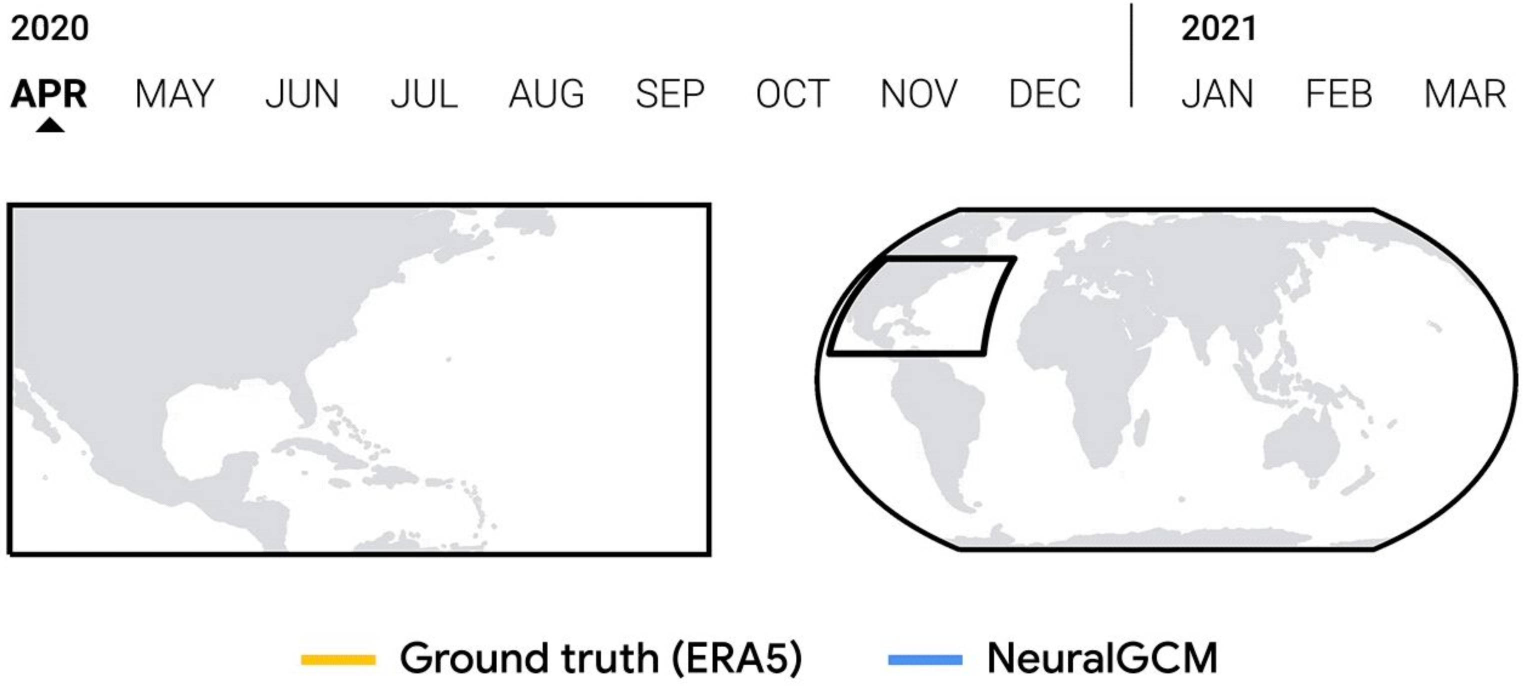

Instability and climate drift can occur in decadal predictions

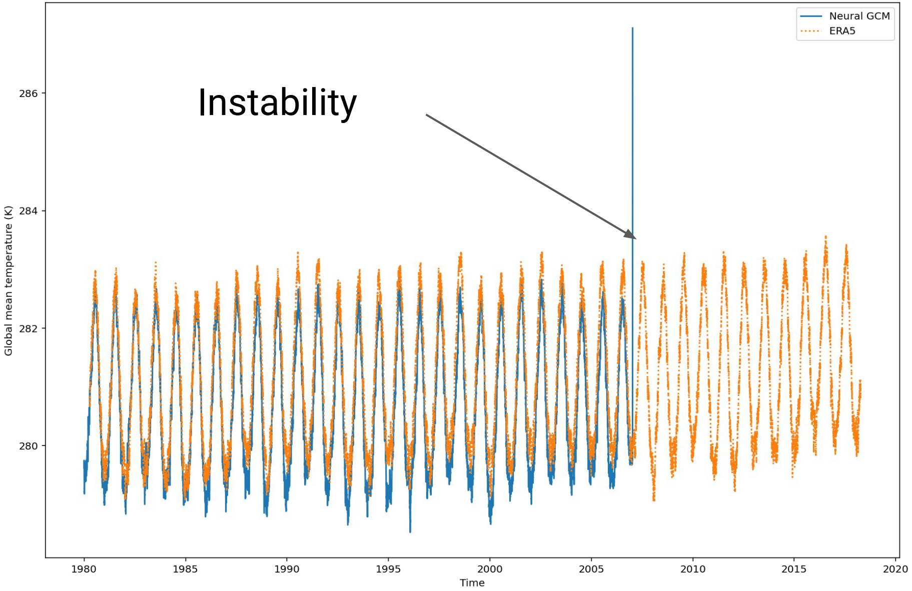

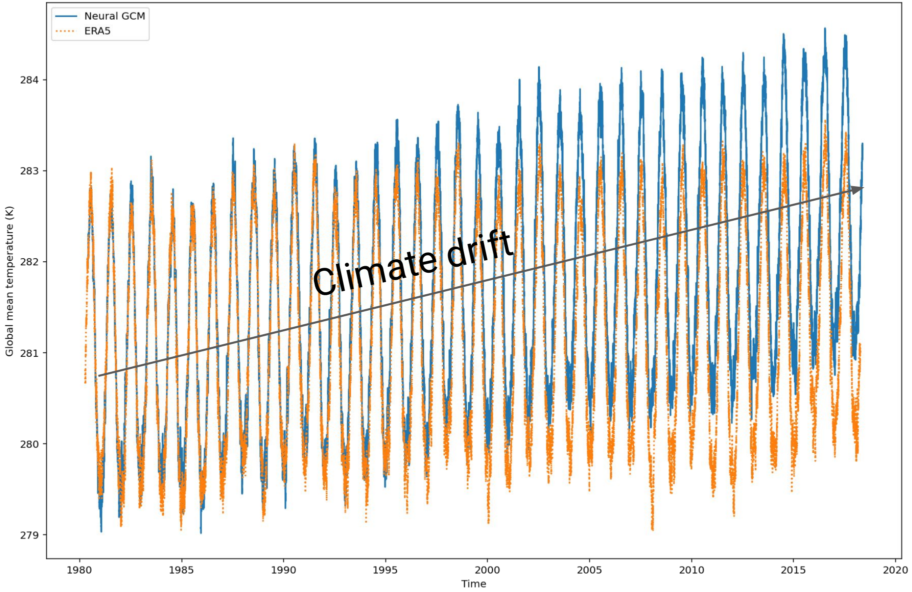

Some NeuralGCM models are stable for decadal runs

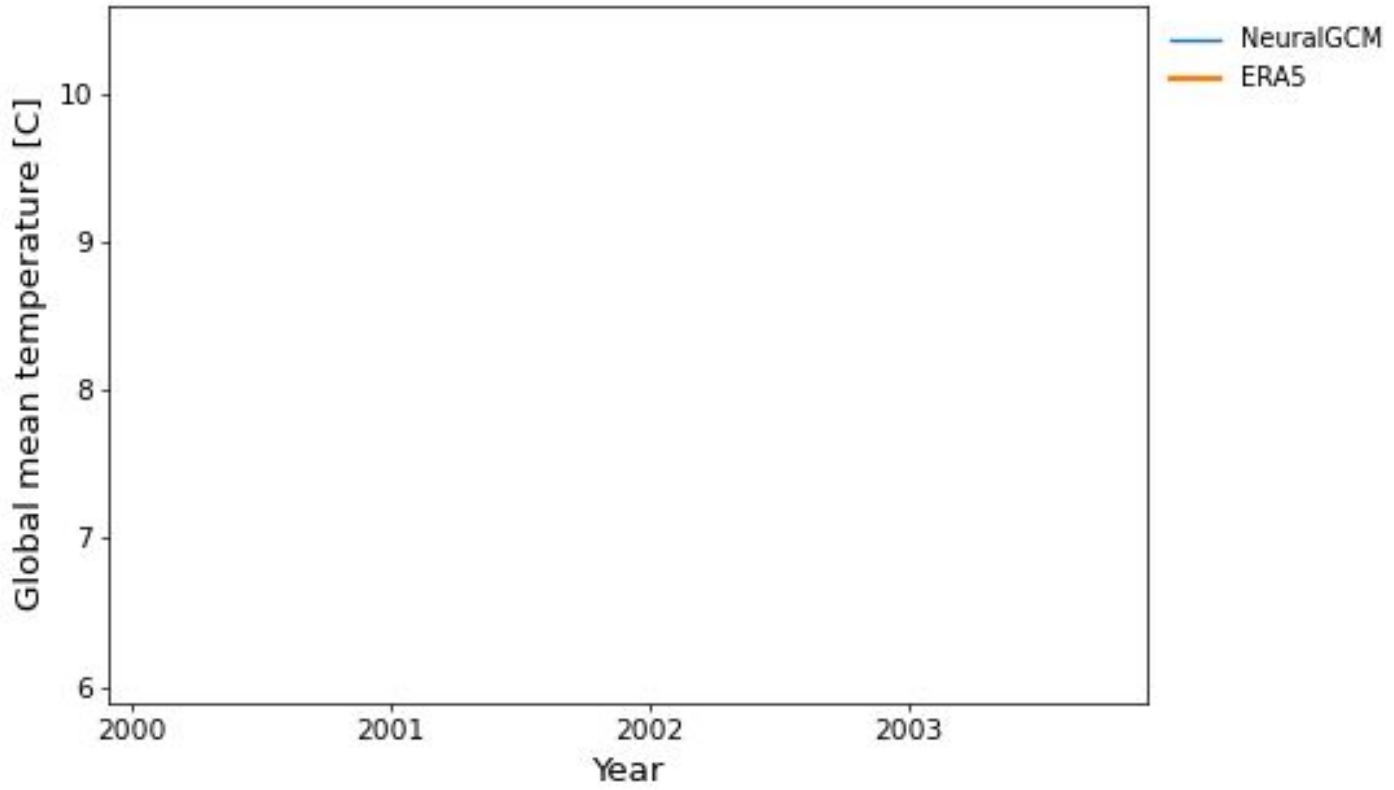

Neural GCM can sometimes capture warming trends with comparable accuracy to CESM

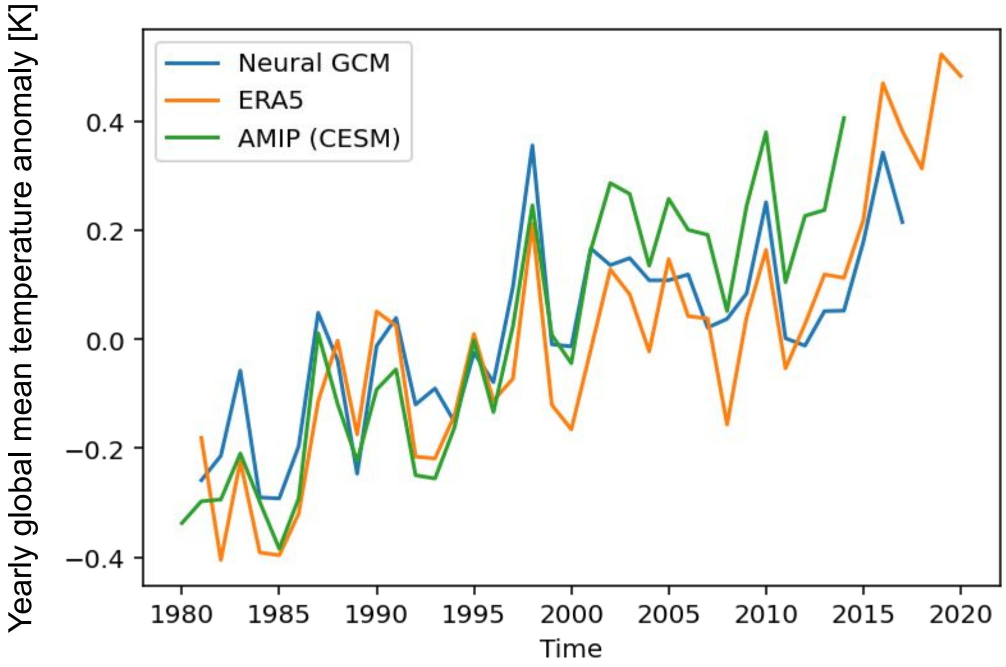

Here NeuralGCM infers global warming signal from the provided ocean temperature

For real world applications one would need to incorporate CO^2, other gas species and run coupled land-ocean-atmosphere simulations

Google Research

Outline

1. General Circulation Models (GCMs) for weather and climate
2. AI revolution for weather forecasting
3. Neural GCM - differentiable hybrid atmospheric model
4. Neural GCM results
   a. Weather forecasting
   b. “Climate” simulations
5. Future directions

Full Earth System modeling system vision

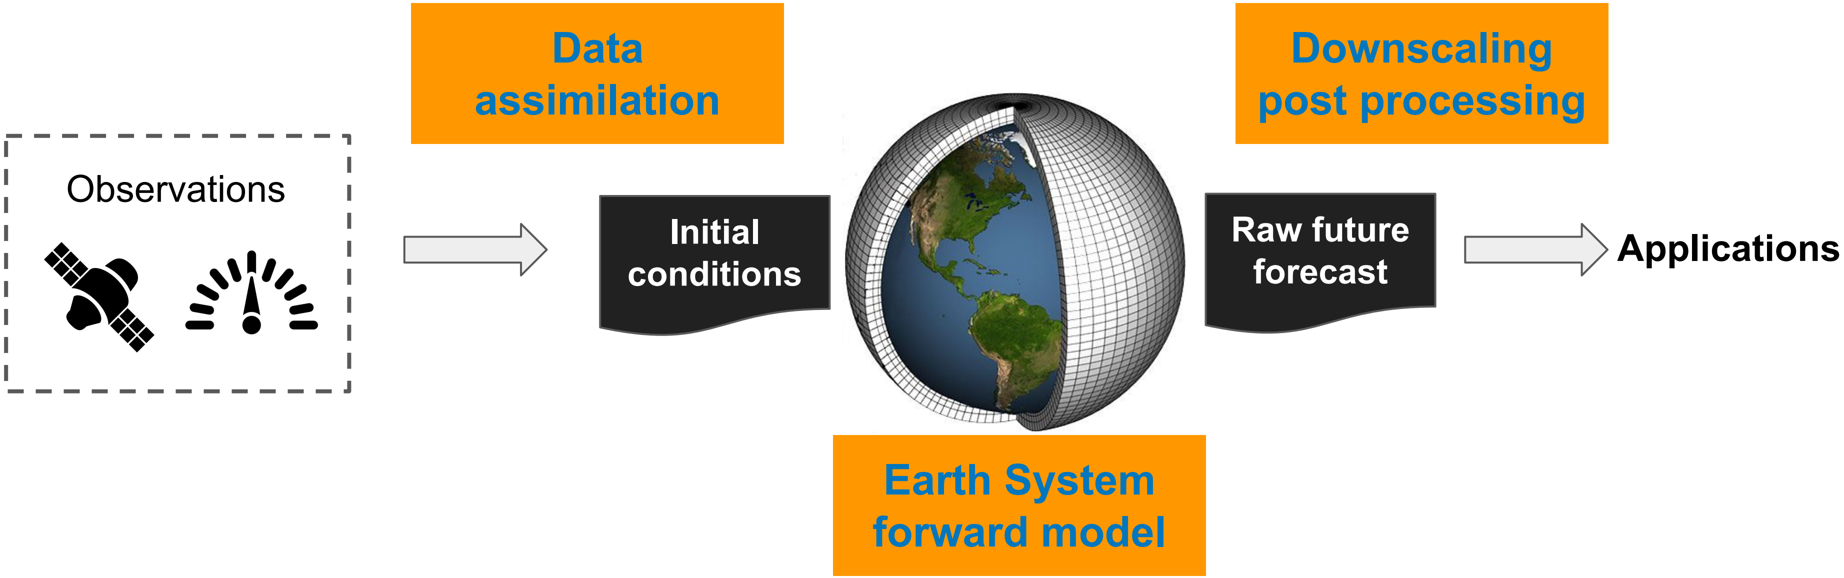

- **Provide high quality nowcast everywhere**
- **Predict severe weather days ahead of time**
- **Help plan & prepare for seasonal changes**
- **Understand & adapt to warming climate**

Initial condition
[nowcasting, medium range, …]

Boundary condition
climate variability, catastrophe risks

Current focus: Coupled components & Observation data

Coupled modeling

Promising results on data-driven ocean emulation & work towards coupled models

Learning from observations

Improved precipitation modeling by training on IMERG data

Open source code

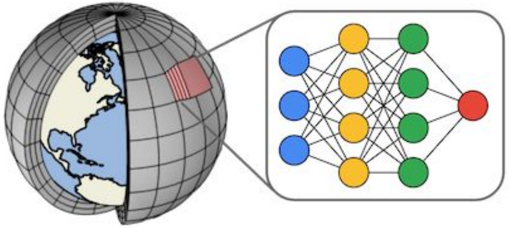

NeuralGCM

Updating code to support new research on weather & climate (DA, DS, coupled models)

We hope to provide meaningful insights via S2S forecasting and quantitative climate variability estimates

Near term Earth System Modeling goals

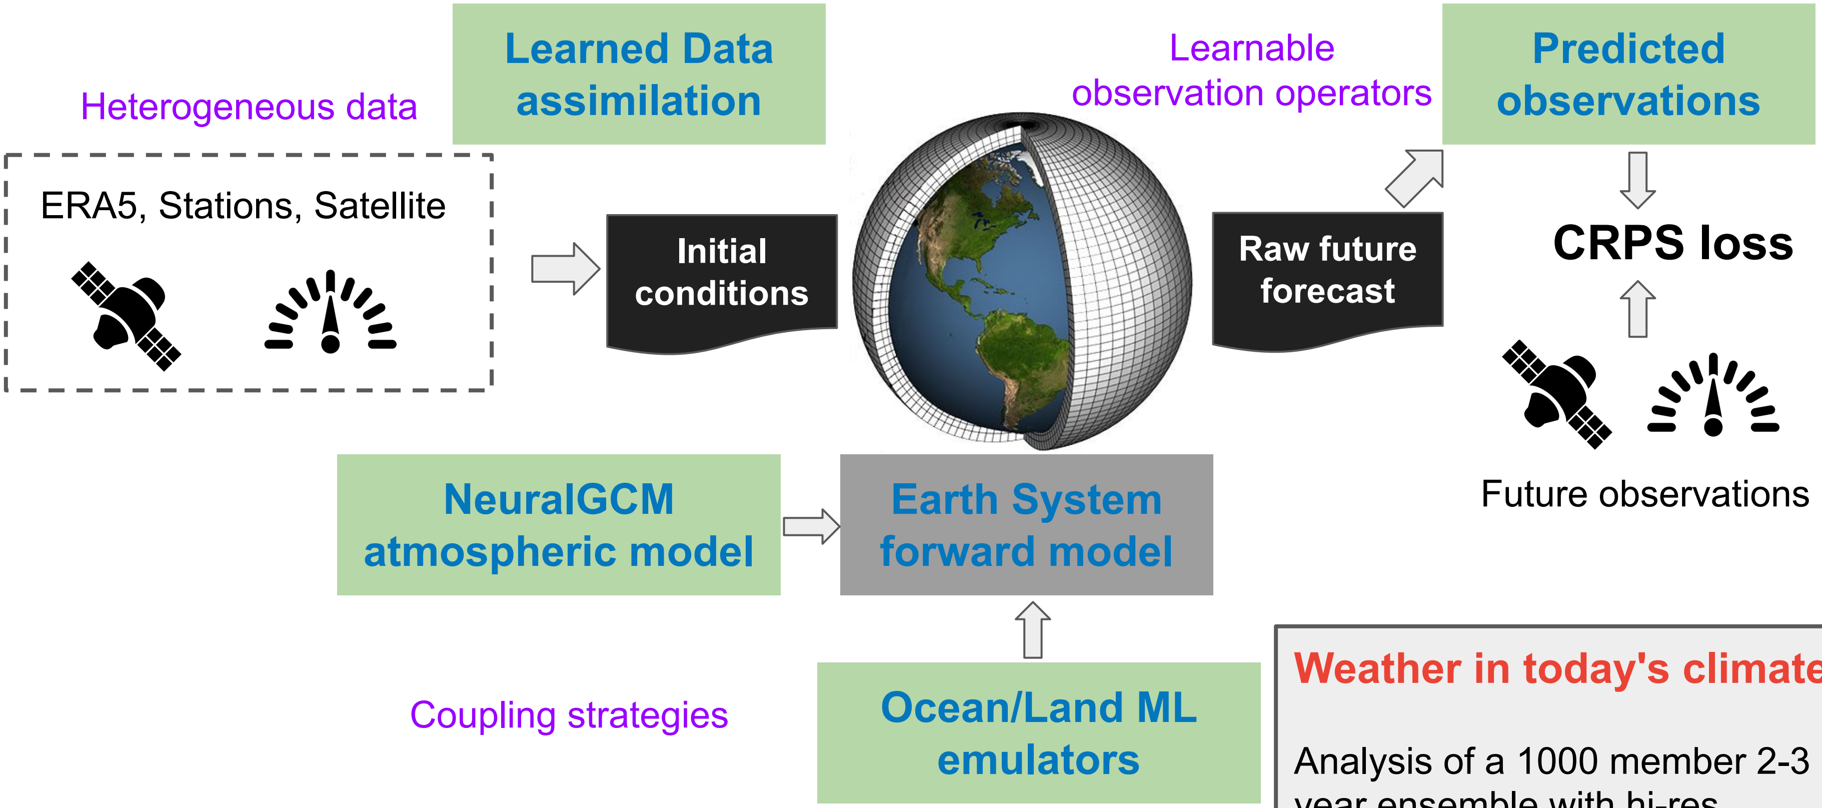

Weather in today's climate?

Analysis of a 1000 member 2-3 year ensemble with hi-res downscaling in areas of interest

Summary

NeuralGCM is an open, fast efficient model that generates realistic ensembles of weather forecasts and features relevant emergent phenomena at longer time integrations

We are hoping to increase the breadth of AI-for-climate research and enable larger community to improve upon our models

Thanks to the Neural GCMs core team and collaborators:

Dmitrii Kochkov

Janni Yuval

Ian Langmore

Stephan Hoyer

NeuralGCM collaborators:
Peter Norgaard, Jamie Smith, Griffin Mooers, James Lottes, Stephan Rasp, Sam Hatfield, Peter Duben, Milan Klower, Peter Battaglia, Alvaro Sanchez-Gonzalez, Matthew Willson, Michael Brenner