NVIDIA

Earth-2: Digital Twins for Weather and Climate

Karthik Kashinath, Principal Engineer and Scientist, AI-HPC

Engineering Lead, NVIDIA Earth-2 Initiative

Agenda

- NVIDIA's Earth-2 initiative: The Big Picture
- FourCastNet & latest accomplishments in Weather.
- Beyond Weather, towards Climate.
- Digital Twinning platforms

The Future Under Climate Change will be Harsh

We urgently need better tools to prepare for it

WILDFIRE PREVENTION

WATER MANAGEMENT

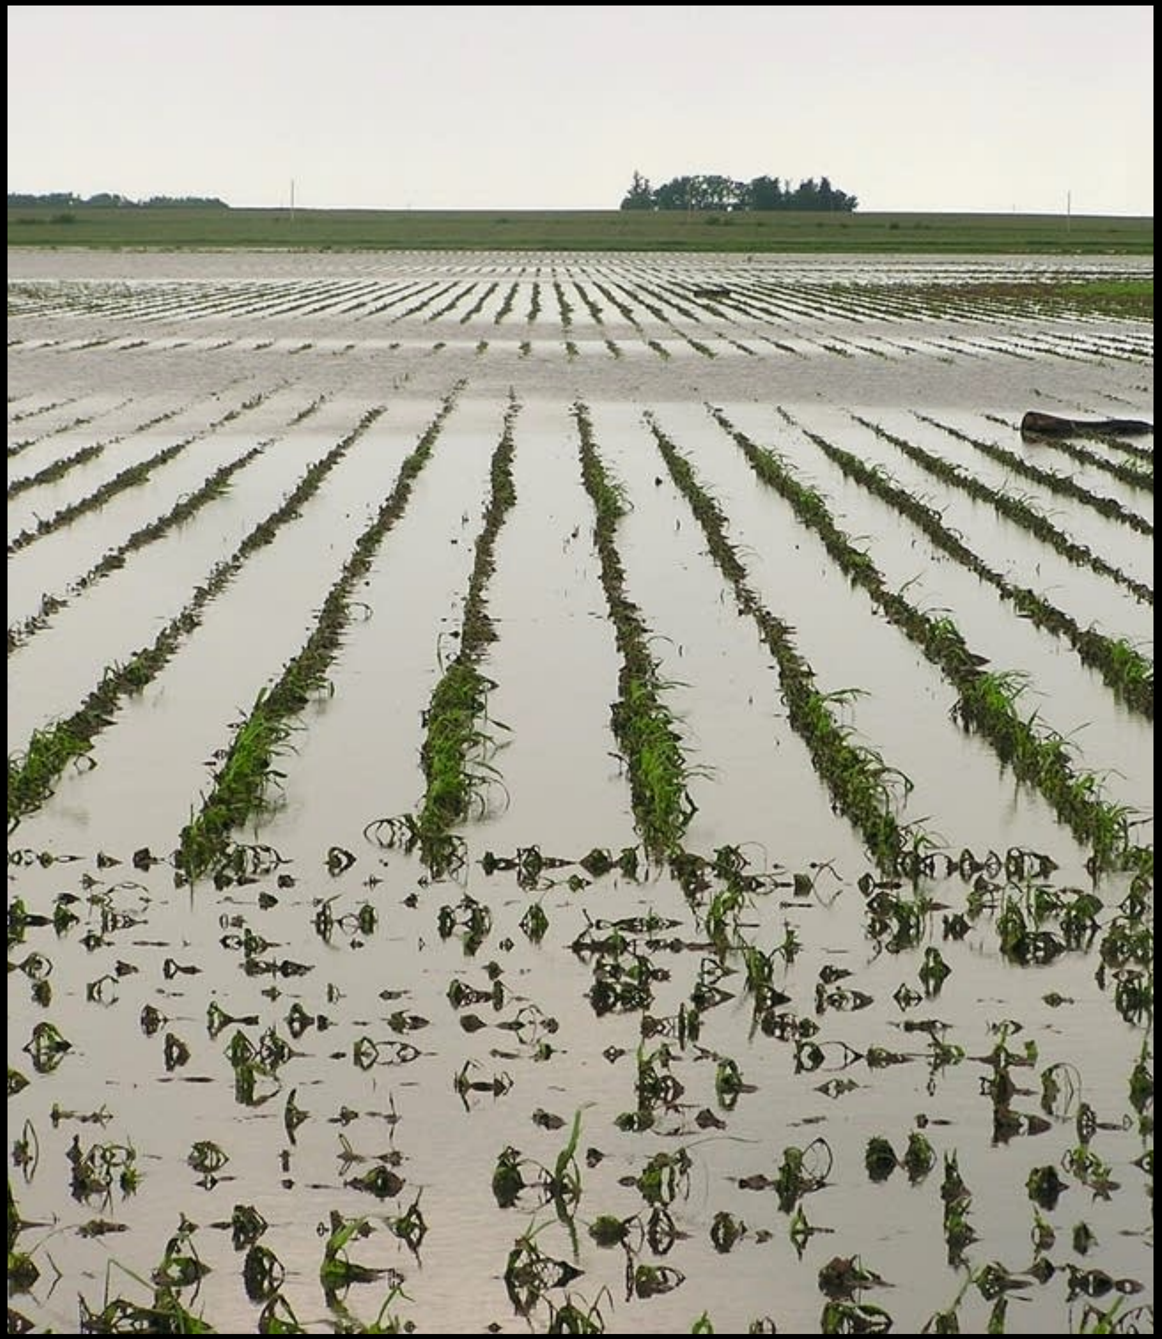

CROP FORECASTING

High-Resolution Climate Prediction is a Computational Challenge

Today's climate models are too low resolution. Brute force numerical solvers are decades away from what is needed.

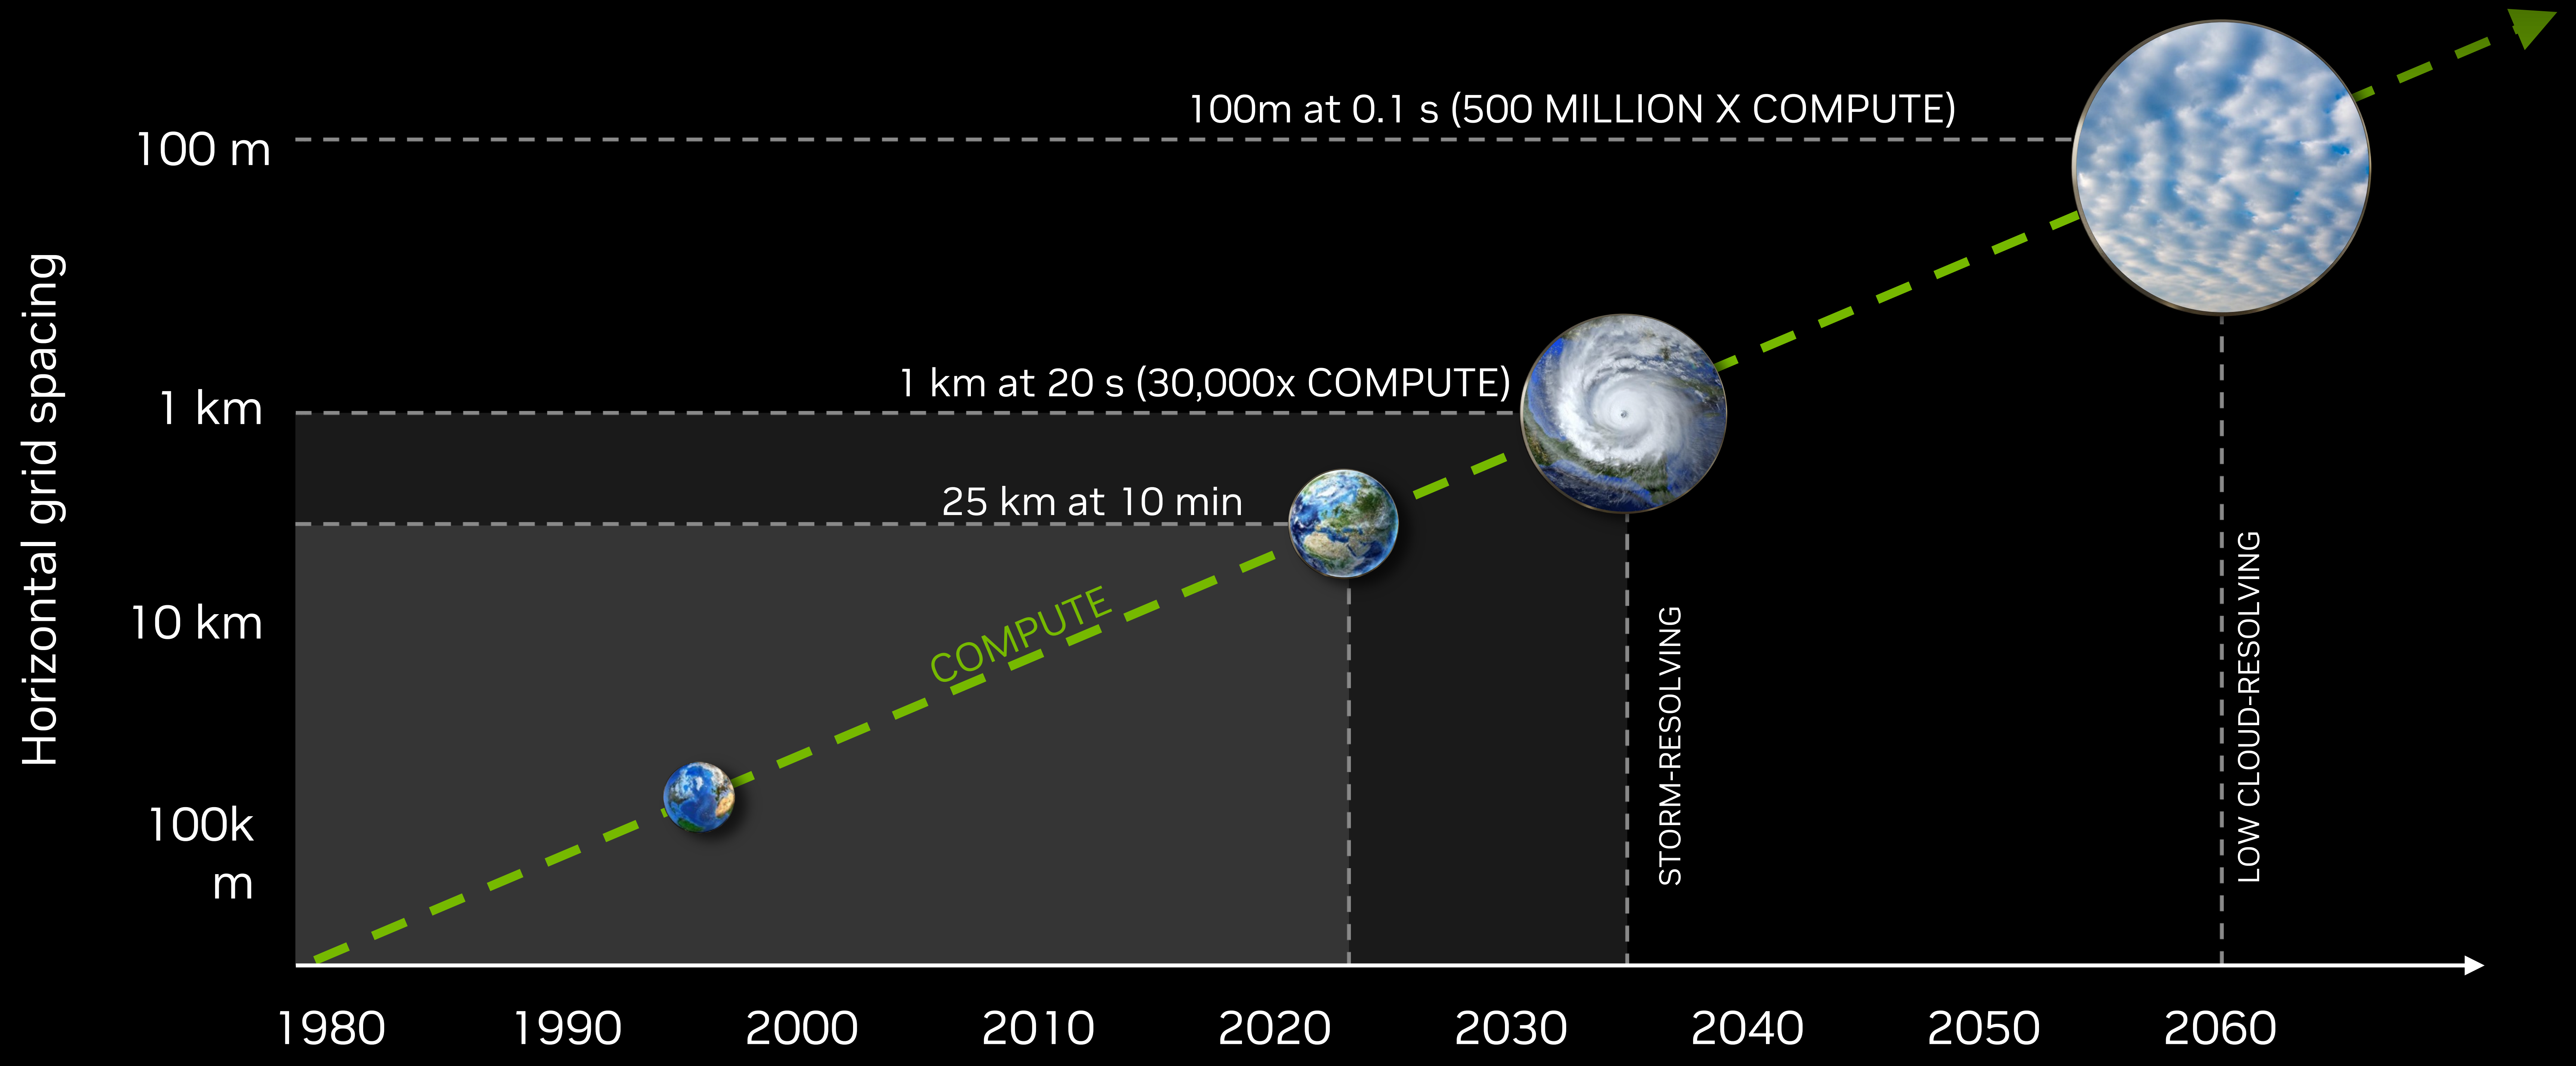

Figure adapted from: Schneider, T., Teixeira, J., Bretherton, C. *et al.* "*Climate goals and computing the future of clouds*". *Nature Climate Change* **7,** 3–5 (2017)

It is Hard to Interact with High-Resolution Climate Predictions

"We can compute km-scale predictions, but can't effectively extract information content, let alone interact with it"

-- Prof. Dr. Bjorn Stevens.

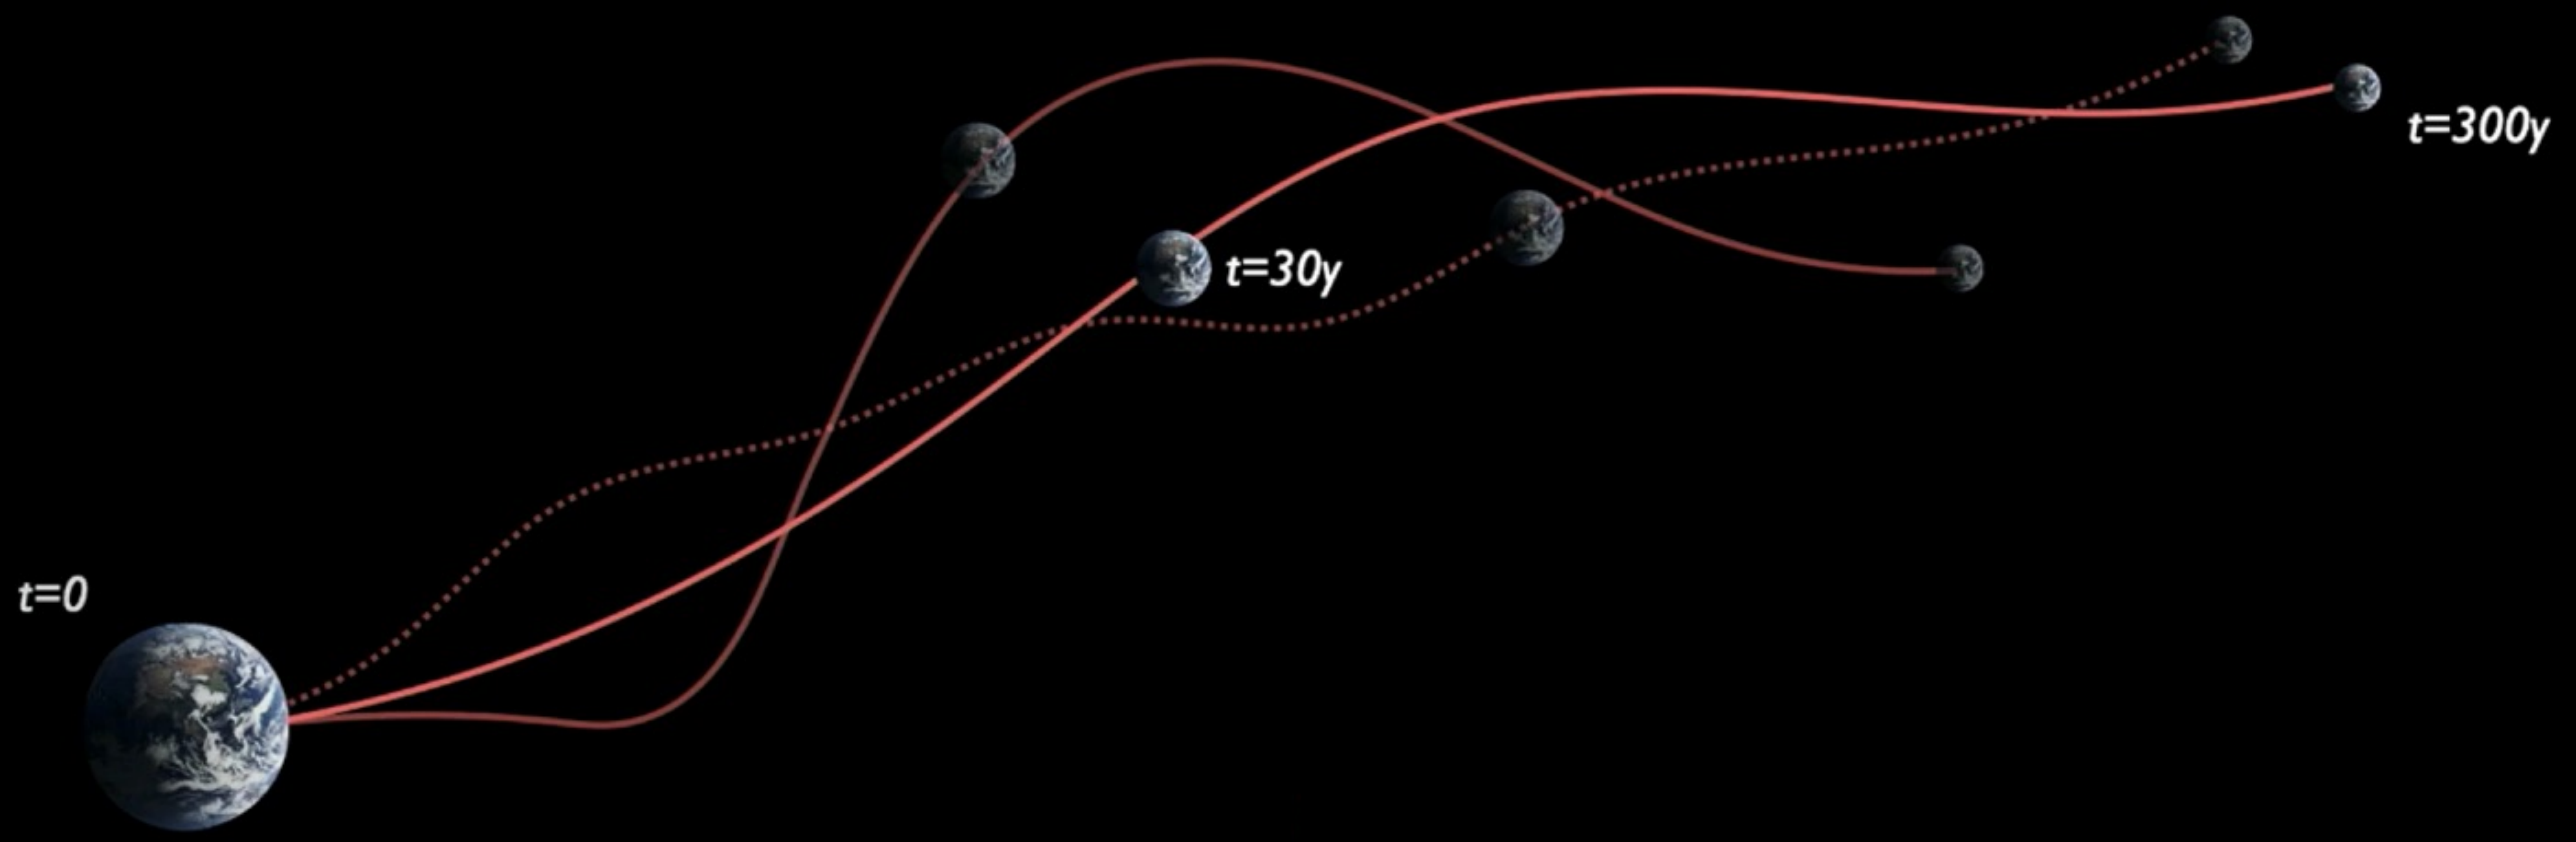

NVIDIA

EARTH-2: Began as a vision

Of a highly interactive climate information system for serving society with next-gen climate predictions.

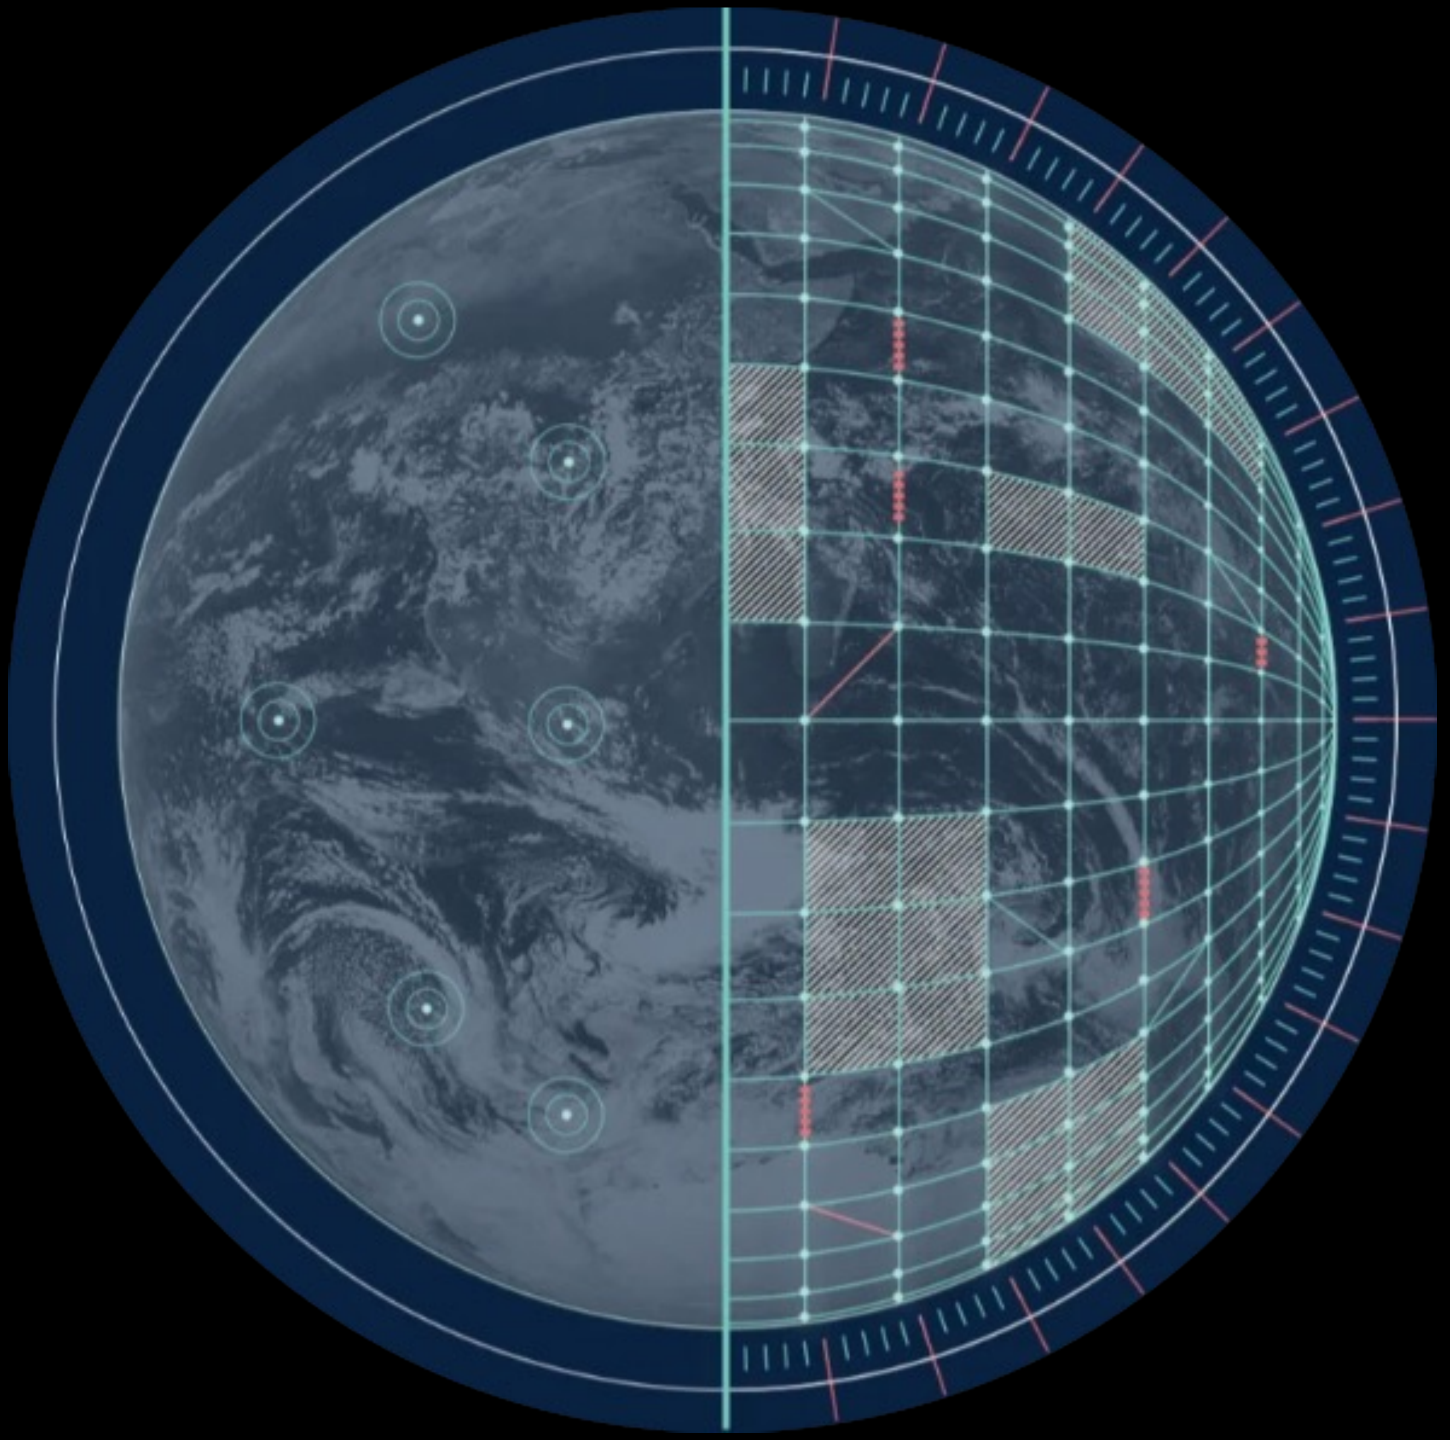

Drawing on Destination Earth: Bauer, Stevens, Hazeleger. "A *Digital Twin of Earth for the Green Transition*". *Nature Climate Change* **11,** 80–83 (2021)

Imagine you could Select a Region of the Planet...

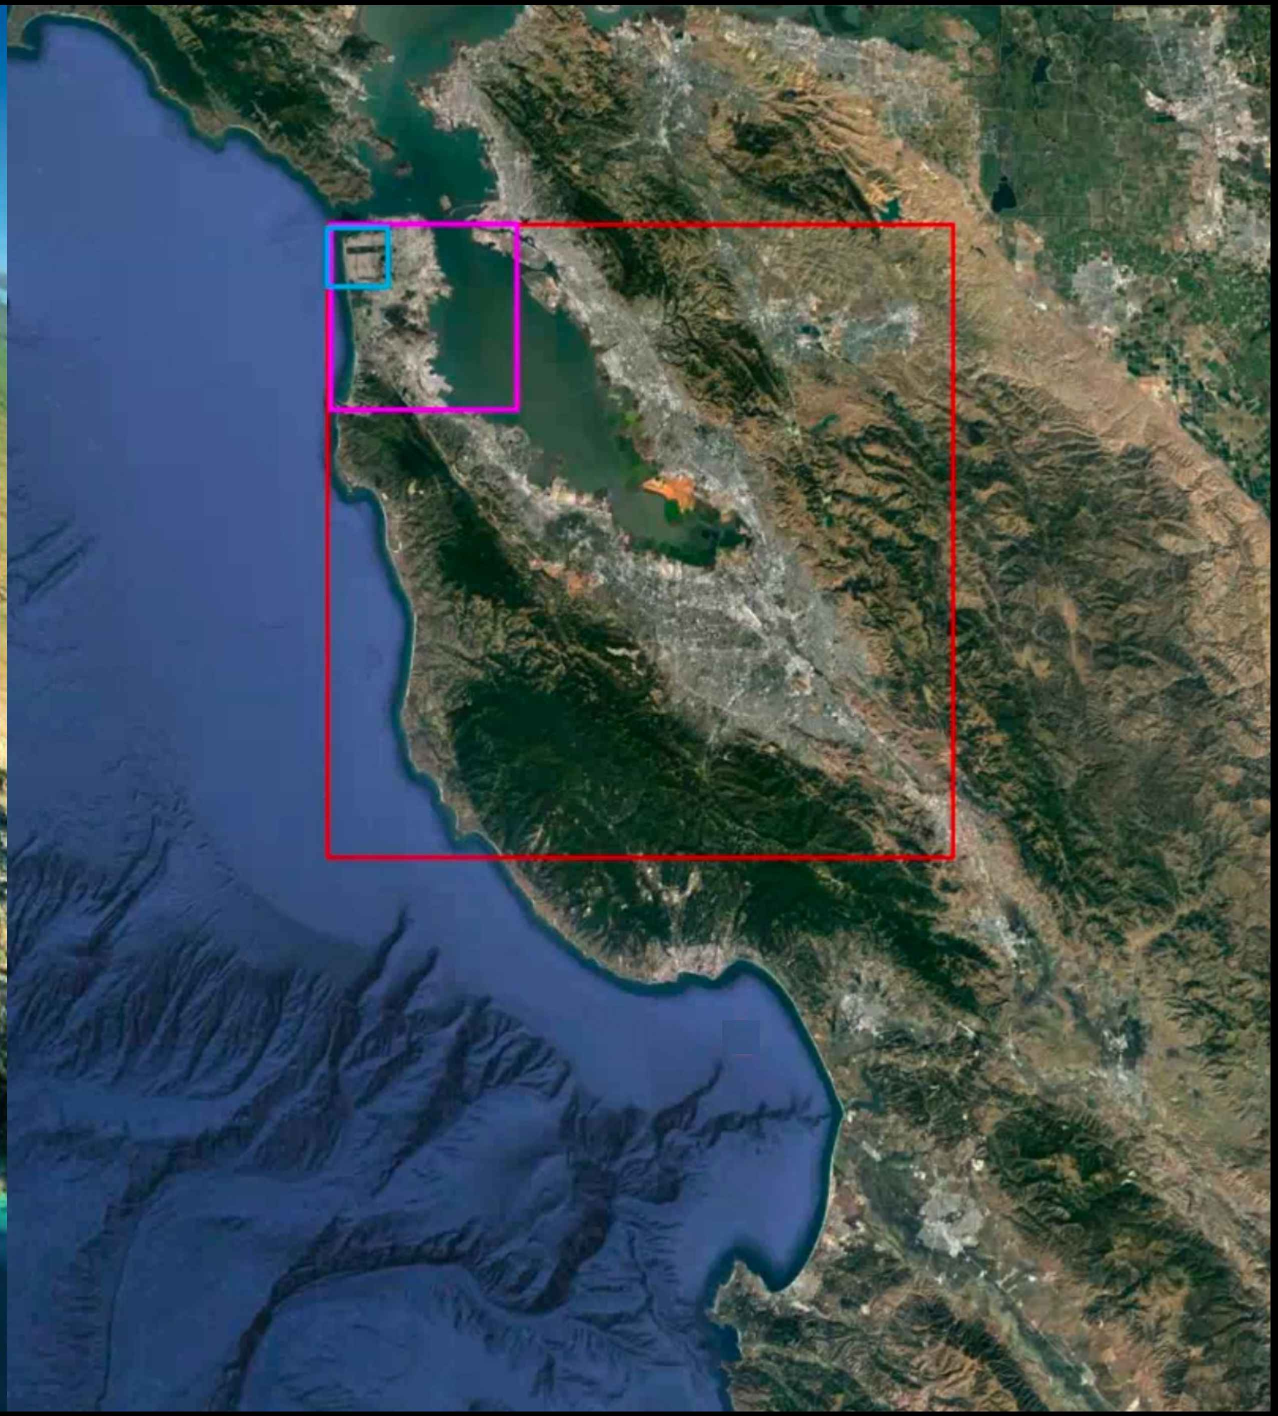

… Ask Questions about Climate Change's Impacts…

On Food, Health, Infrastructure, Energy systems, and more…

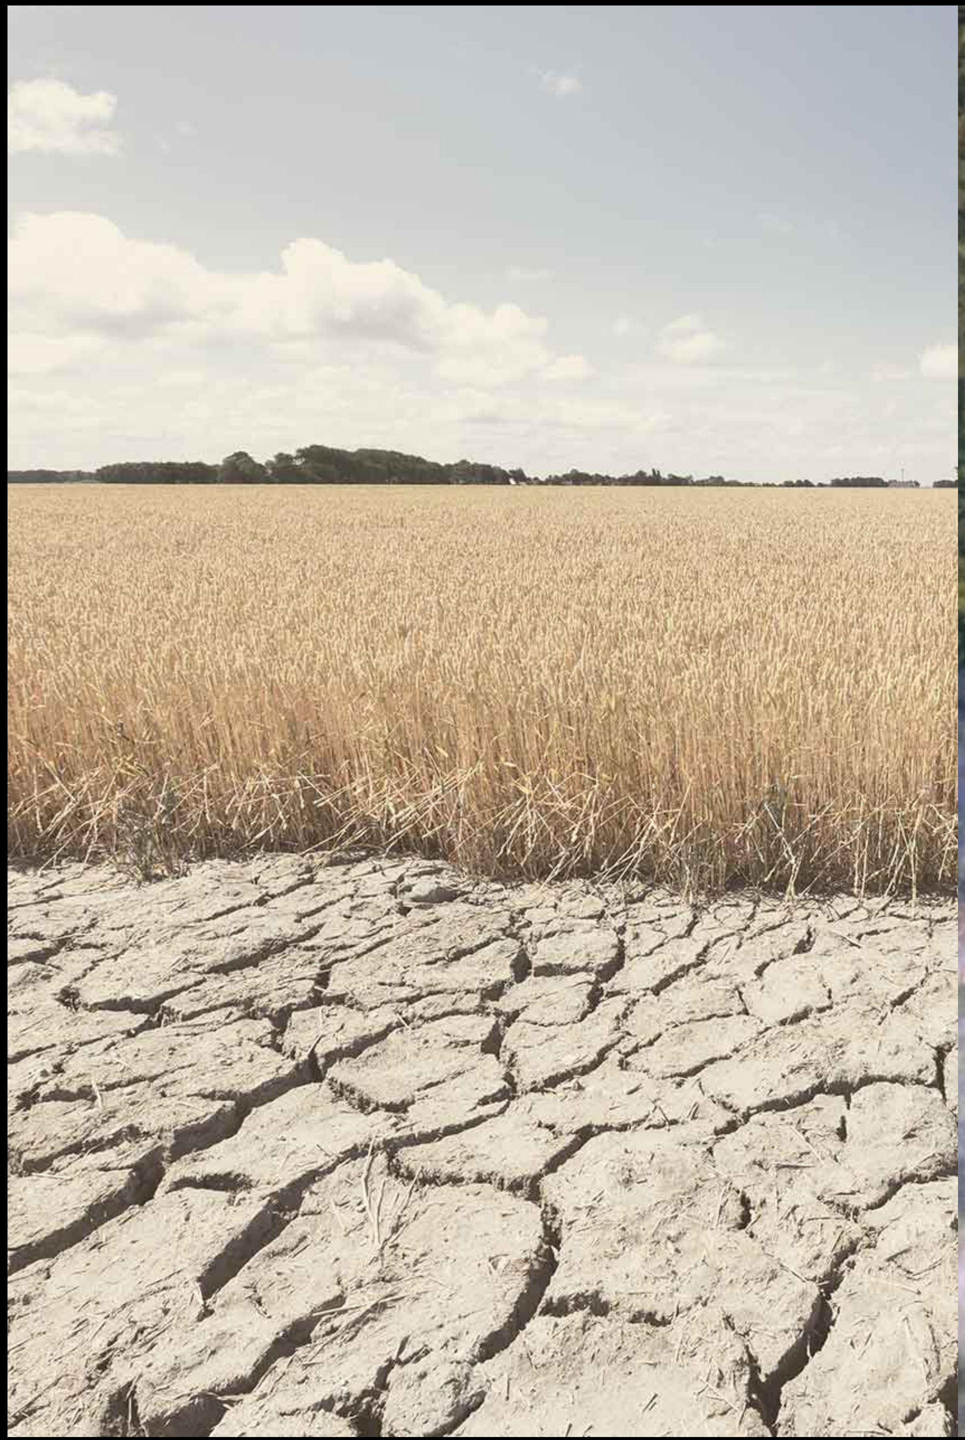

… and Receive Useful, Visual & Statistical Guidance?

From a Highly Interactive Future Climate Information System, at High Resolution, that Serves Society…

Earth

Twin Earth

Imagine the System Evolving in Scientific Fidelity and Computational Ambition

Eventually fed by a new library of climate predictions so high-resolution they seem impossible today.

Earth-2 Mission #2

Achieving Next-Gen Climate Predictions using Hybrid Physics, Machine Learning & HPC.

The World is Already Working Hard on these Problems

Example: Project Destination Earth envisions what Digital Twins could be.

Interactive Collaborative Platform

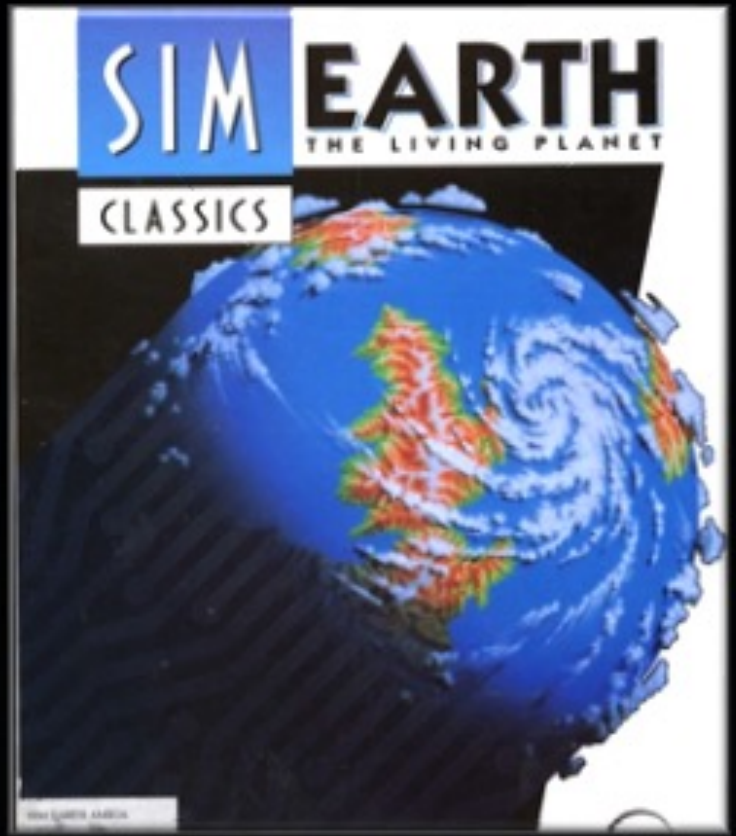

Data-driven Models

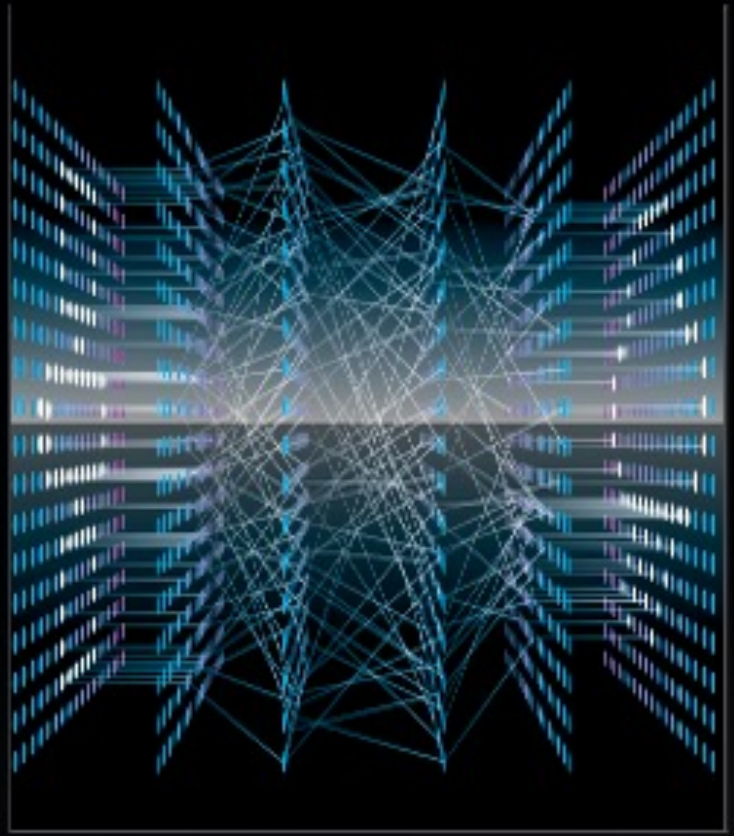

Storm-resolving Models

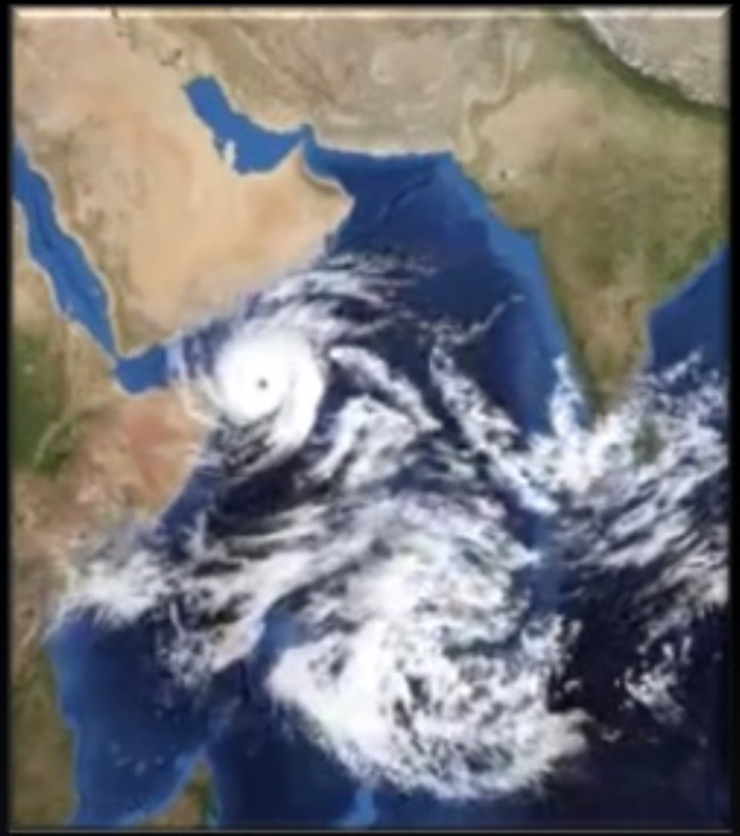

Unified Observations

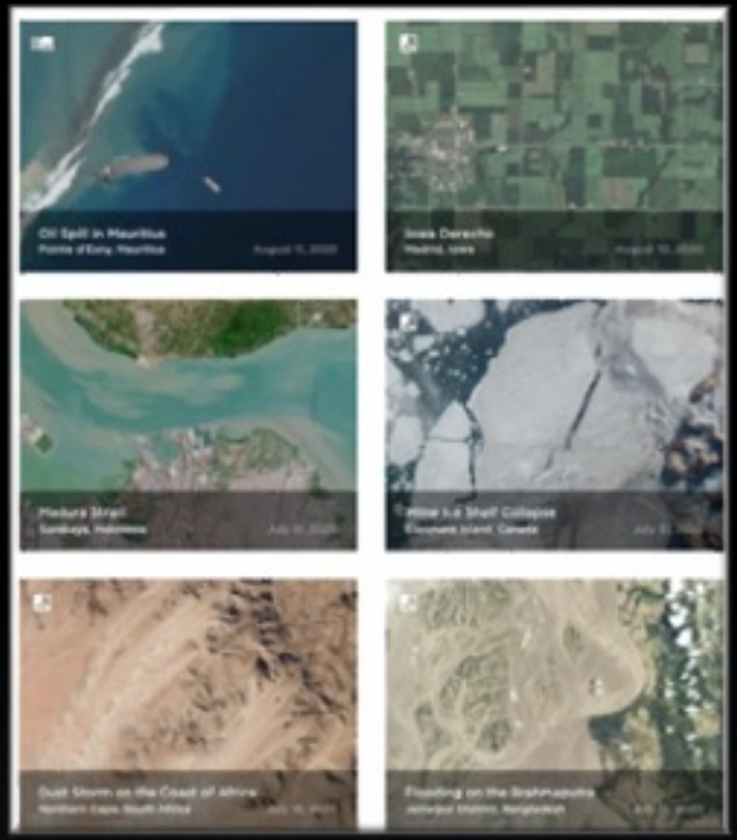

Exascale Compute

https://digital-strategy.ec.europa.eu/en/library/destination-earth

NVIDIA's technical know-how can make a big difference

Earth-2 harnesses NVIDIA's full-stack technologies to make Earth digital twins a reality

Interactive Collaborative Platform

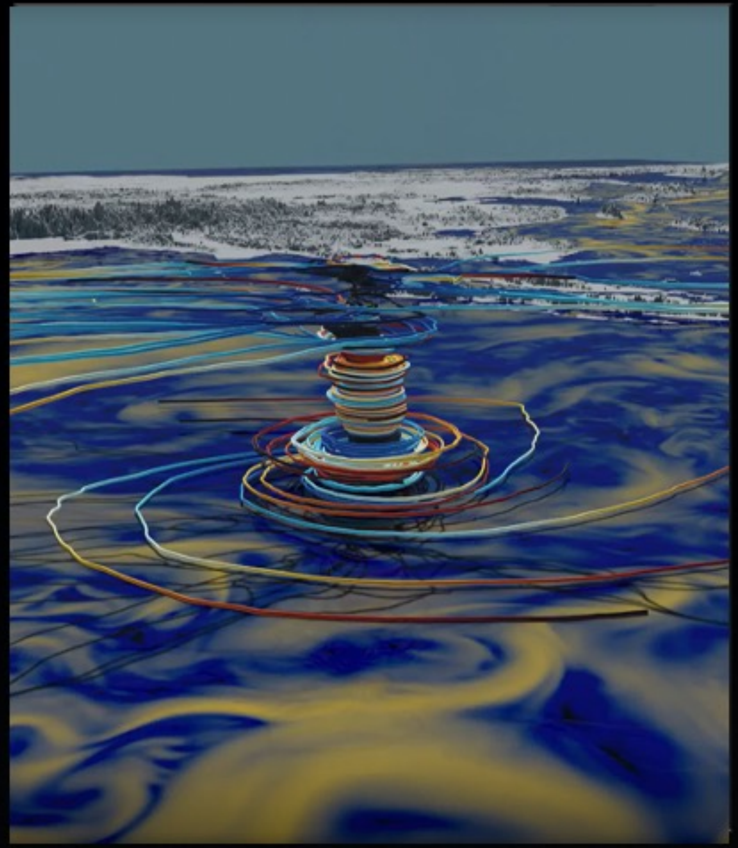

OMNIVERSE

Data-driven Models

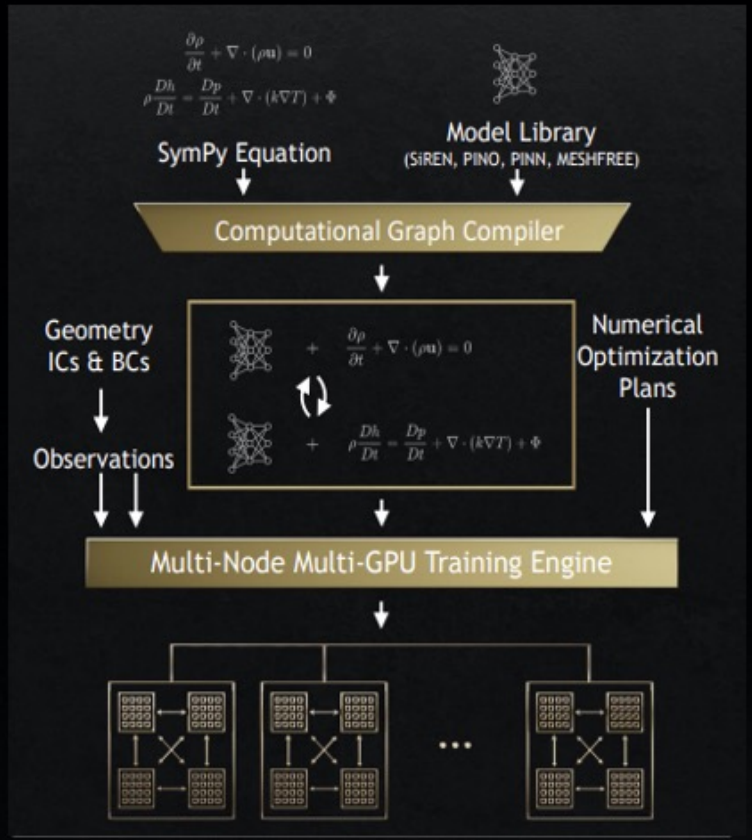

PHYSICS-ML / MODULUS

Storm-resolving Models

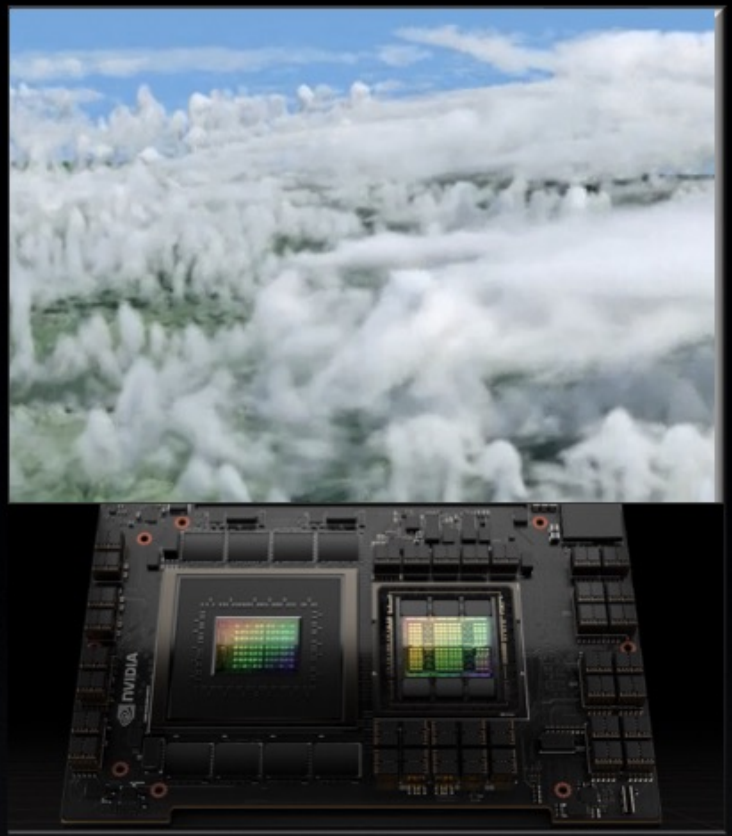

GPU-ACCELERATION

Unified Observations

OMNIVERSE NUCLEUS

Exascale Compute

OVX SUPERPOD

Earth-2 is in Collaboration with International Climate Science

NVIDIA's AI, engineering & full-stack expertise complement research capacity in academia & government.

MAX-PLANCK-INSTITUT FÜR METEOROLOGIE

ECMWF
EUROPEAN CENTRE FOR MEDIUM RANGE WEATHER FORECASTS

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

NASA

NOAA
NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION
U.S. DEPARTMENT OF COMMERCE

DEPARTMENT OF ENERGY
UNITED STATES OF AMERICA

USMILE

MeteoSwiss

Ai2

LEAP

Agenda

- NVIDIA's Earth-2 initiative: The Big Picture
- FourCastNet & latest accomplishments in Weather.
- Beyond Weather, towards Climate.
- Digital Twinning platforms

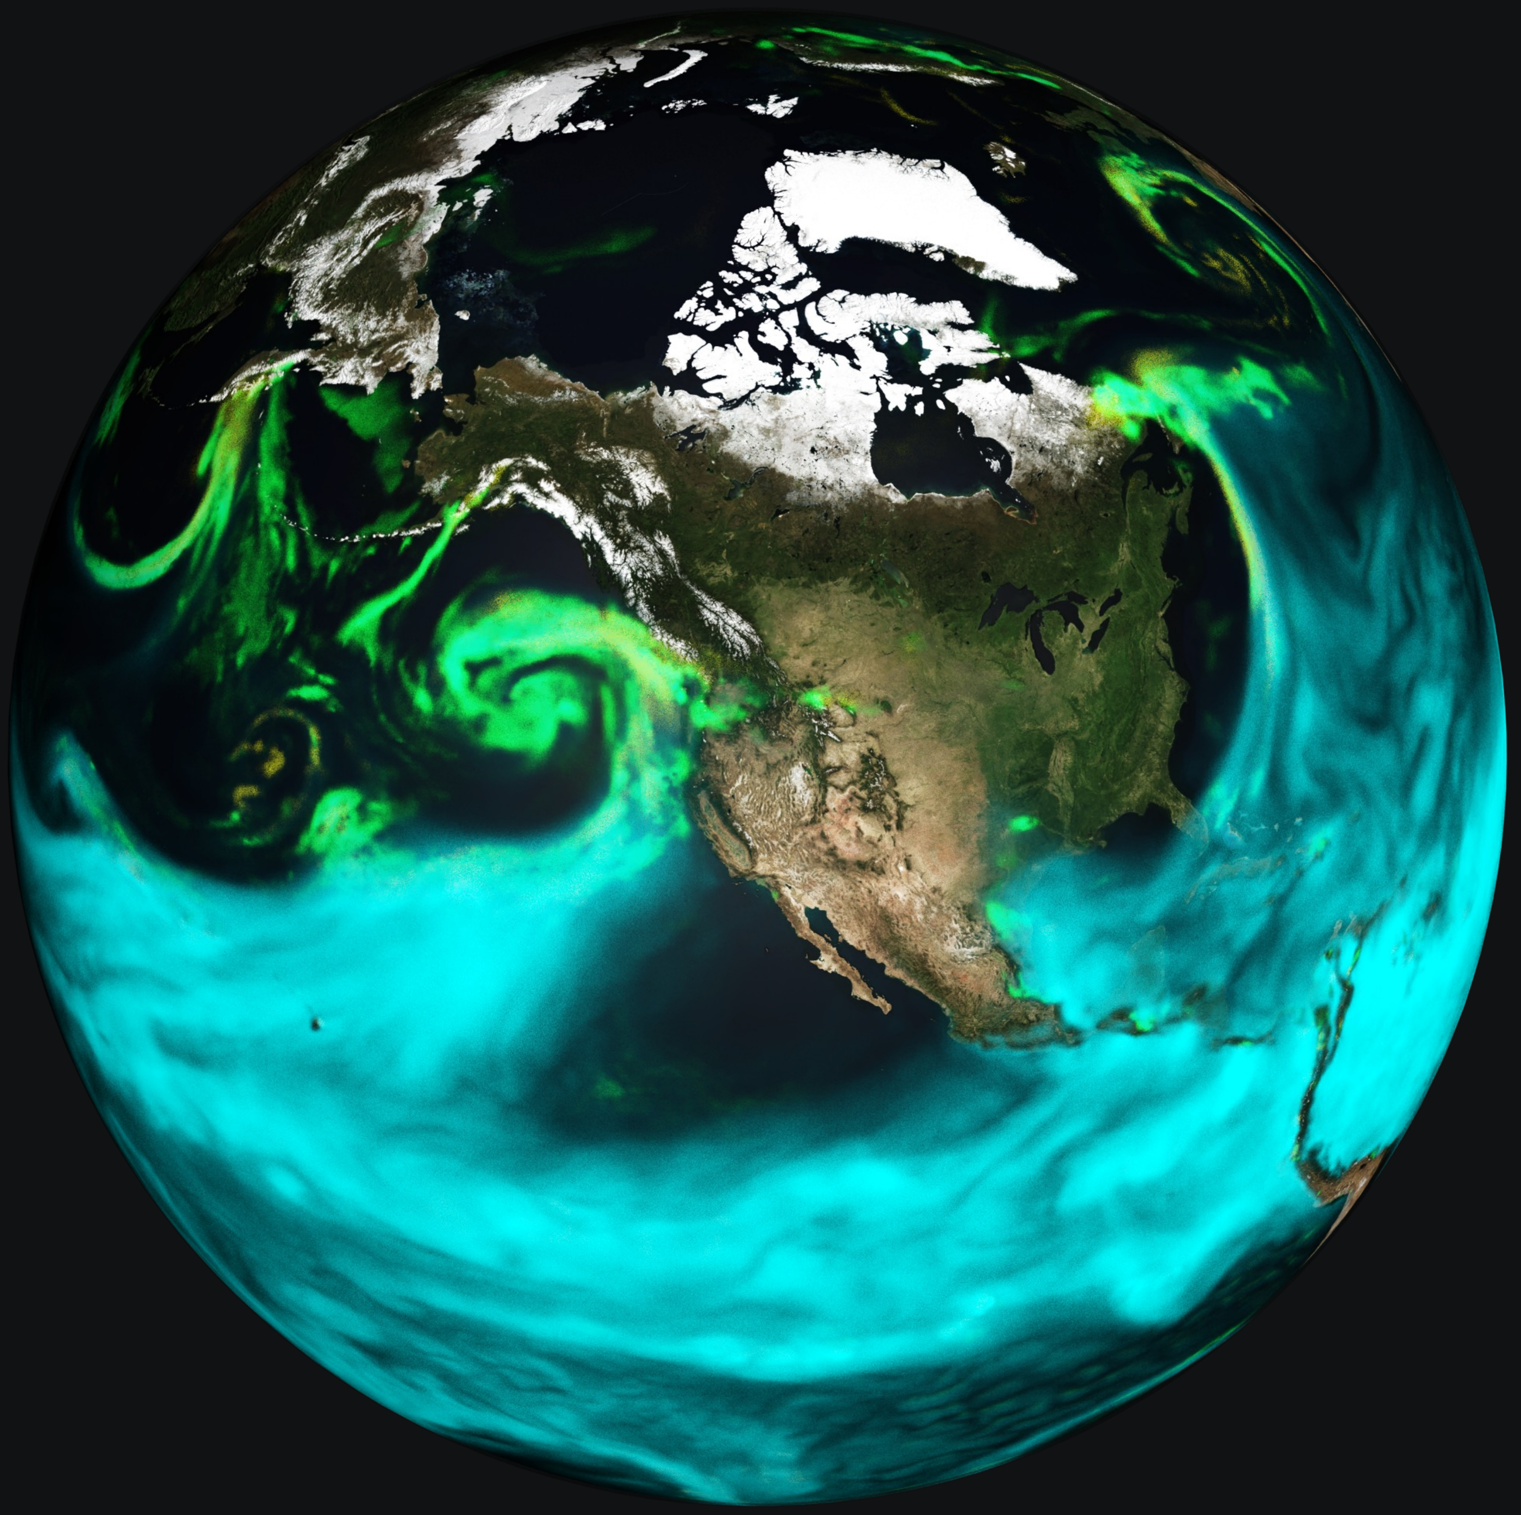

FourCastNet pushing the frontier of AI-Driven Digital Twins

Scope	Global
Model Type	Full-Atmosphere AI Surrogate
Architecture	Fourier Neural Operator + Transformer
Resolution	25km
Training Data	ERA5 Reanalysis
Initial Condition	ERA5 / GFS / UFS
Inference Time	0.5 sec (2-week forecast)
Calibration	IC + Bayesian model uncertainty
Speedup vs NWP	O(10,000 – 100,000)
Power Savings	O(10,000)
Max Stable Rollout	250+ days
Project Type	Open-source

FourCastNet: A data-driven weather predictor of unusually high resolution

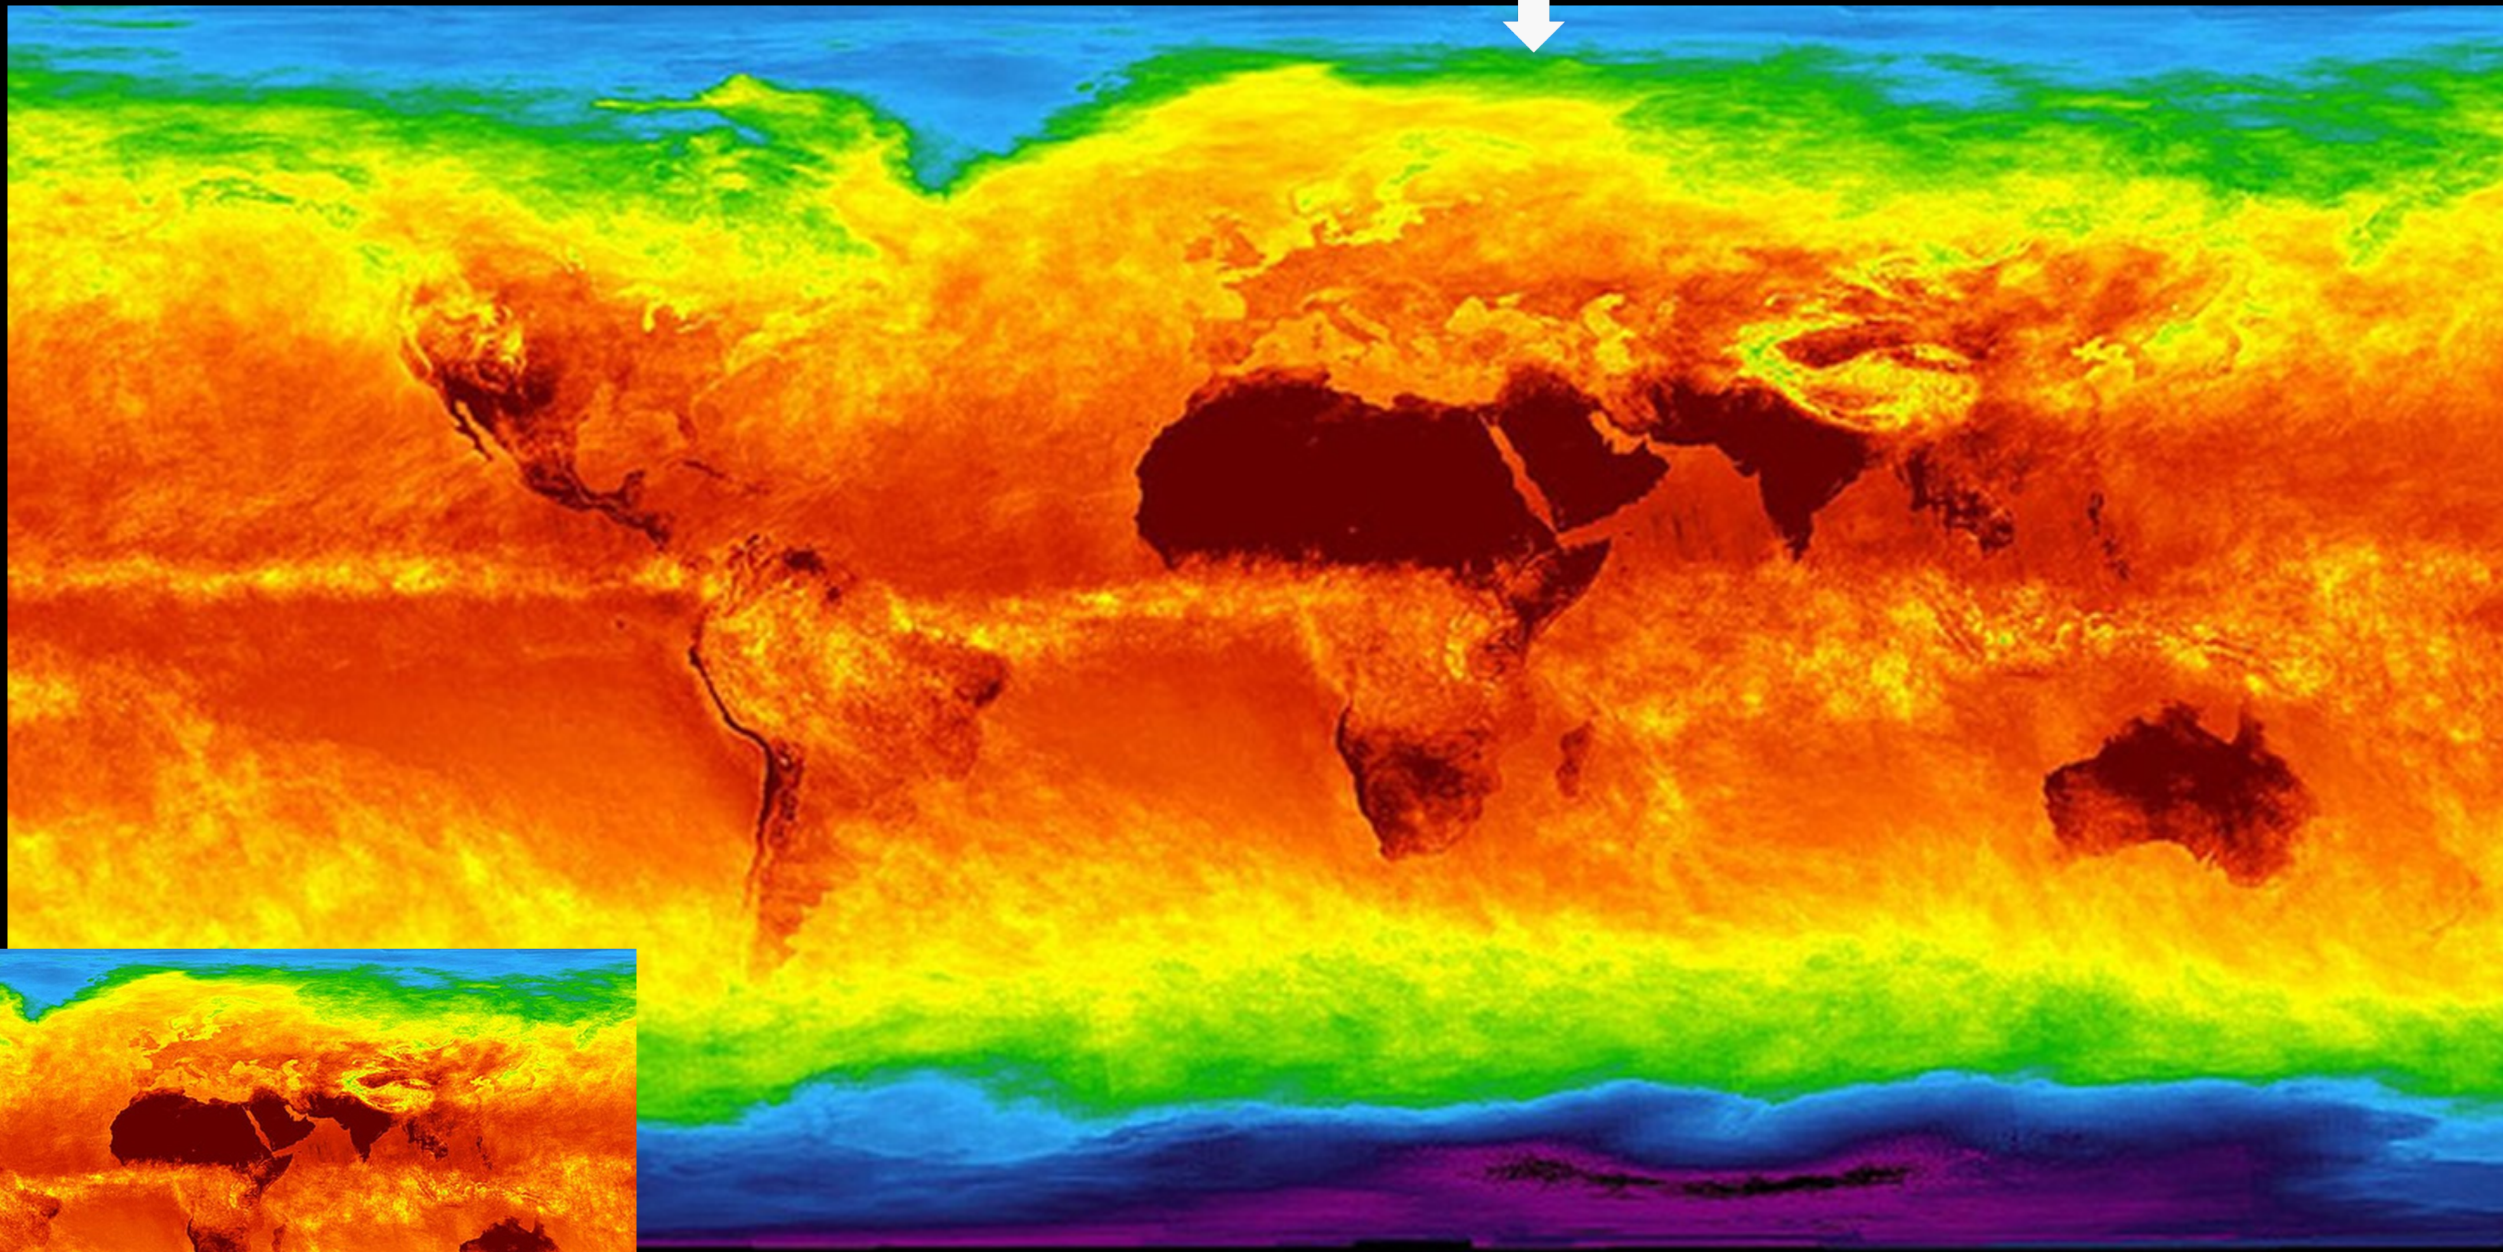

FourCastNet, Pathak et al. (2022), 0.25°, ~1,000,000 Pixels, ViT+AFNO

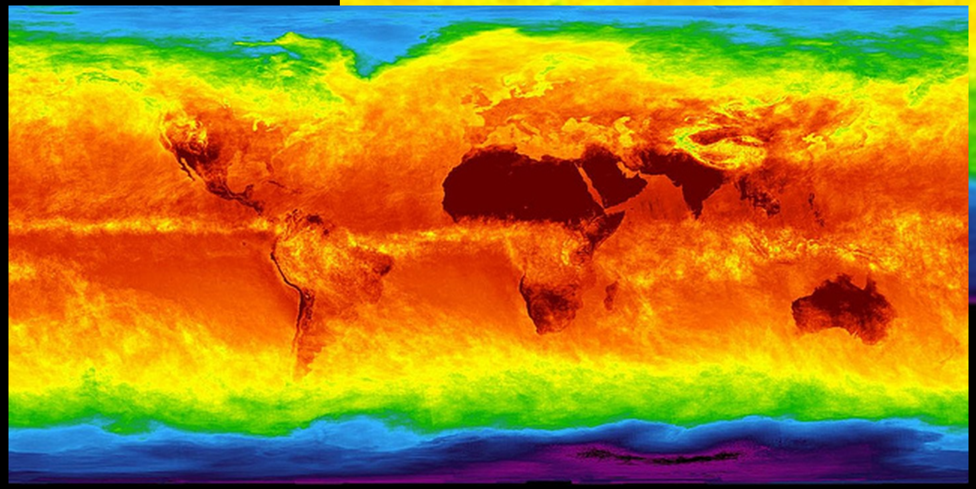

GNN, Keisler et al. (2022), 1°, 64,000 Pixels, Graph Neural Networks

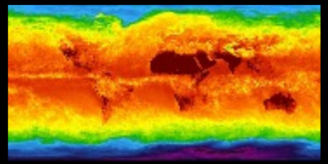

DLWP, Weyn et al. (2020). 2°, 16K pixels, Deep CNN on Cubesphere/(2021) ResNet

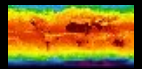

Weyn et al. (2019), 2.5° N.H only, 72x36, 2.6k pixels, ConvLSTM

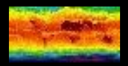

WeatherBench, Rasp et al. (2020). 5.625°, 64x32, 2K pixels, CNN

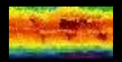

Deuben & Bauer (2018), 6° , 60x30, 1.8K pixels, MLP

FourCastNet uses a novel transformer architecture

With Fourier Neural Operator Blocks - in search of grid-free, high-resolution, machine-learnt simulations.

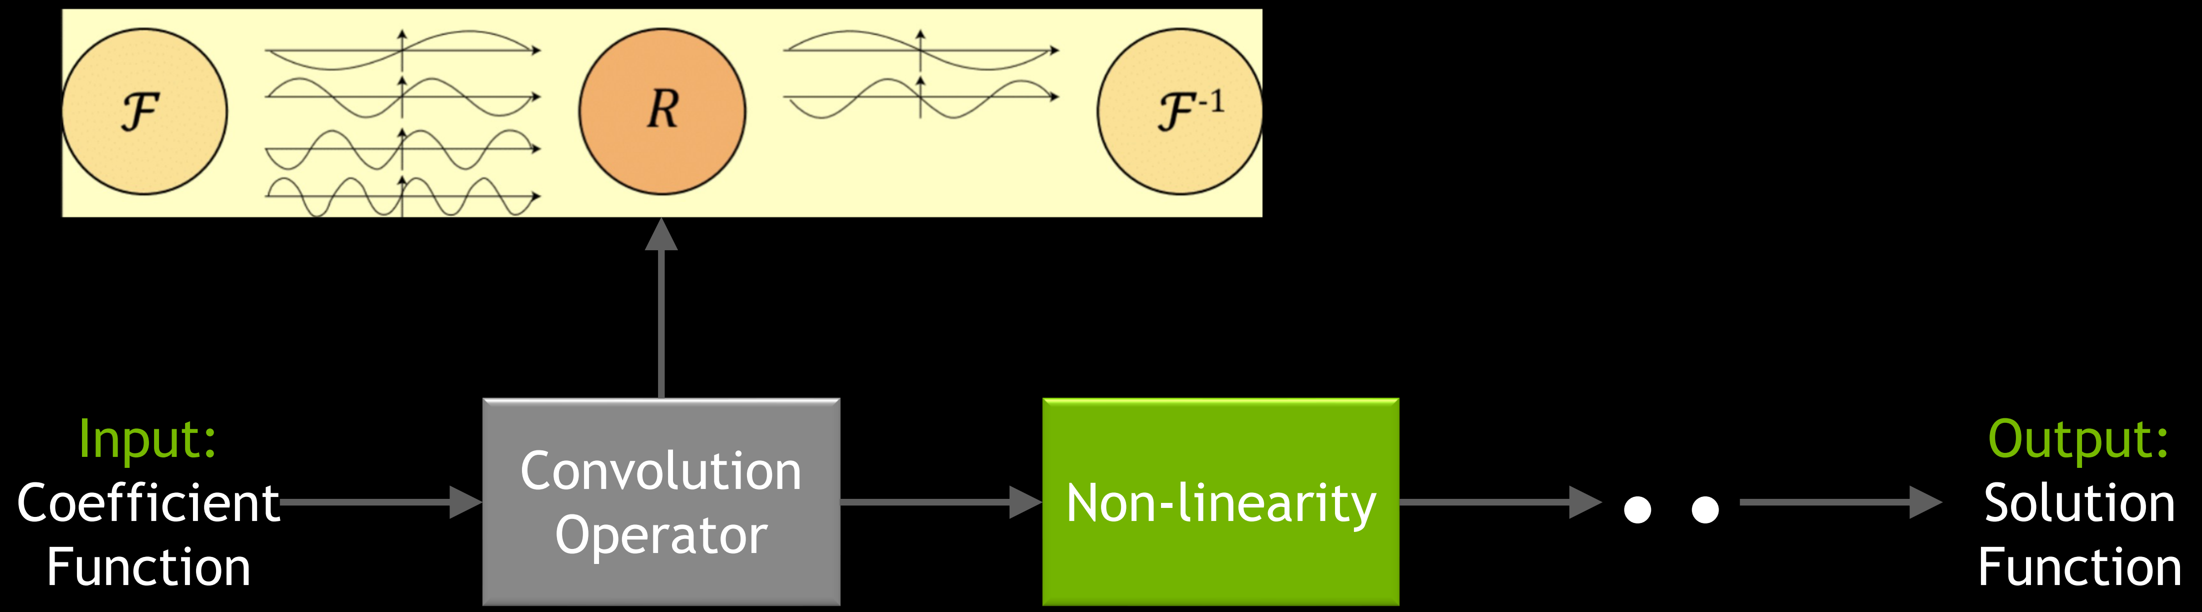

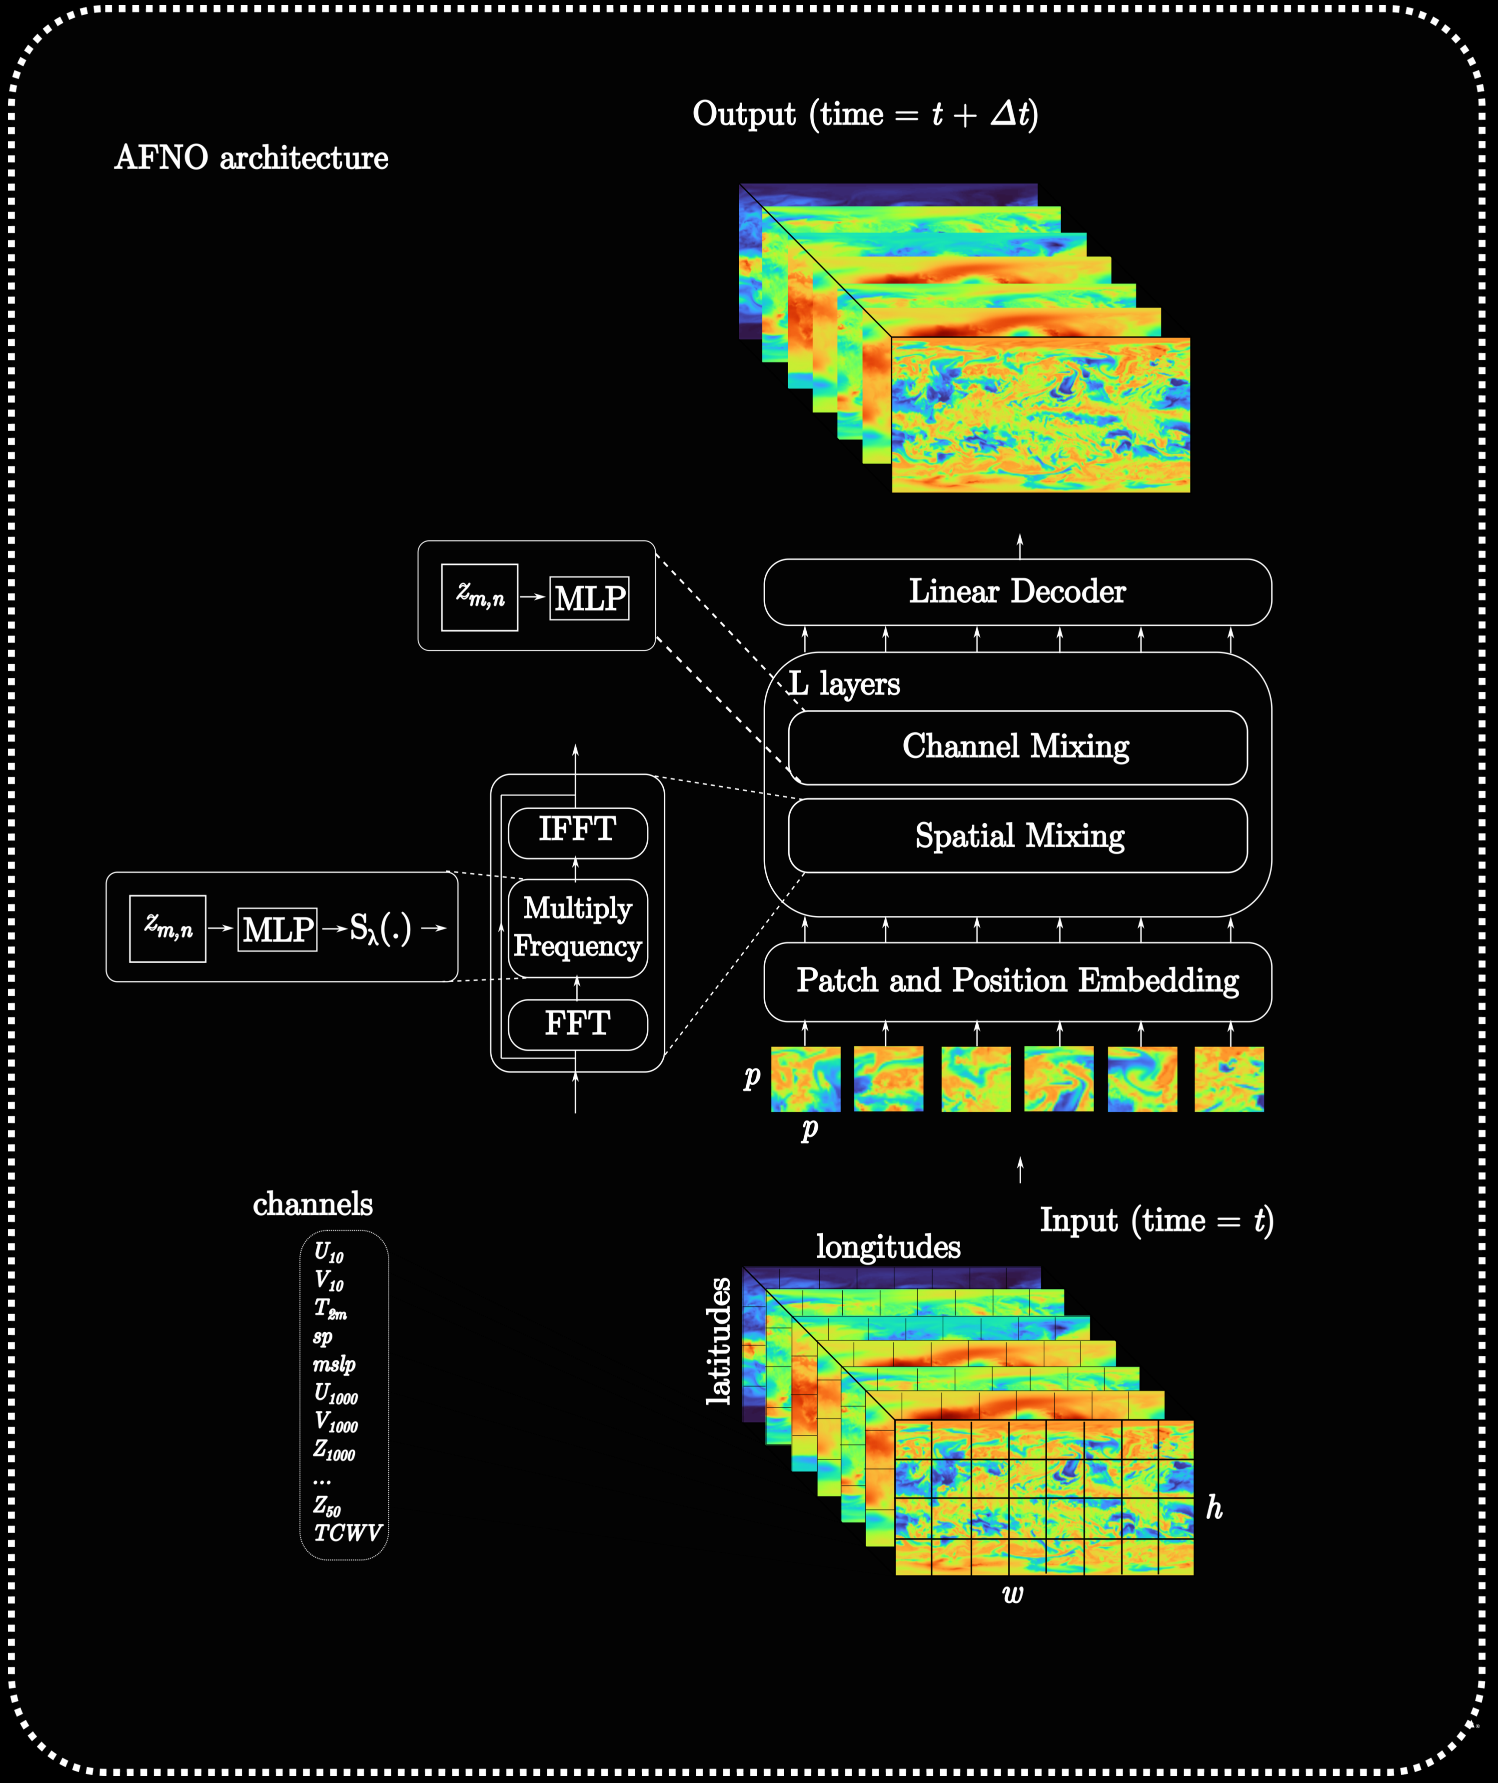

Guibas et al. (2022), Adaptive Fourier Neural Operators: Efficient Token Mixers for Transformers, arXiv:2111.13587

FourCastNet (FCN) is trained on 0.25-degree ERA5 data

With 26 channels (2D fields) of surface and atmospheric variables.

- Surface: U, V, T, MSLP
- 5 vertical levels: U, V, T, Z, RH
- Integrated column water vapor

Extending to include radiation, surface and TOA fluxes, vapor transport, clouds

Training set: 1979 to 2015

Validation set: 2016, 2017

Held out: 2018 onwards

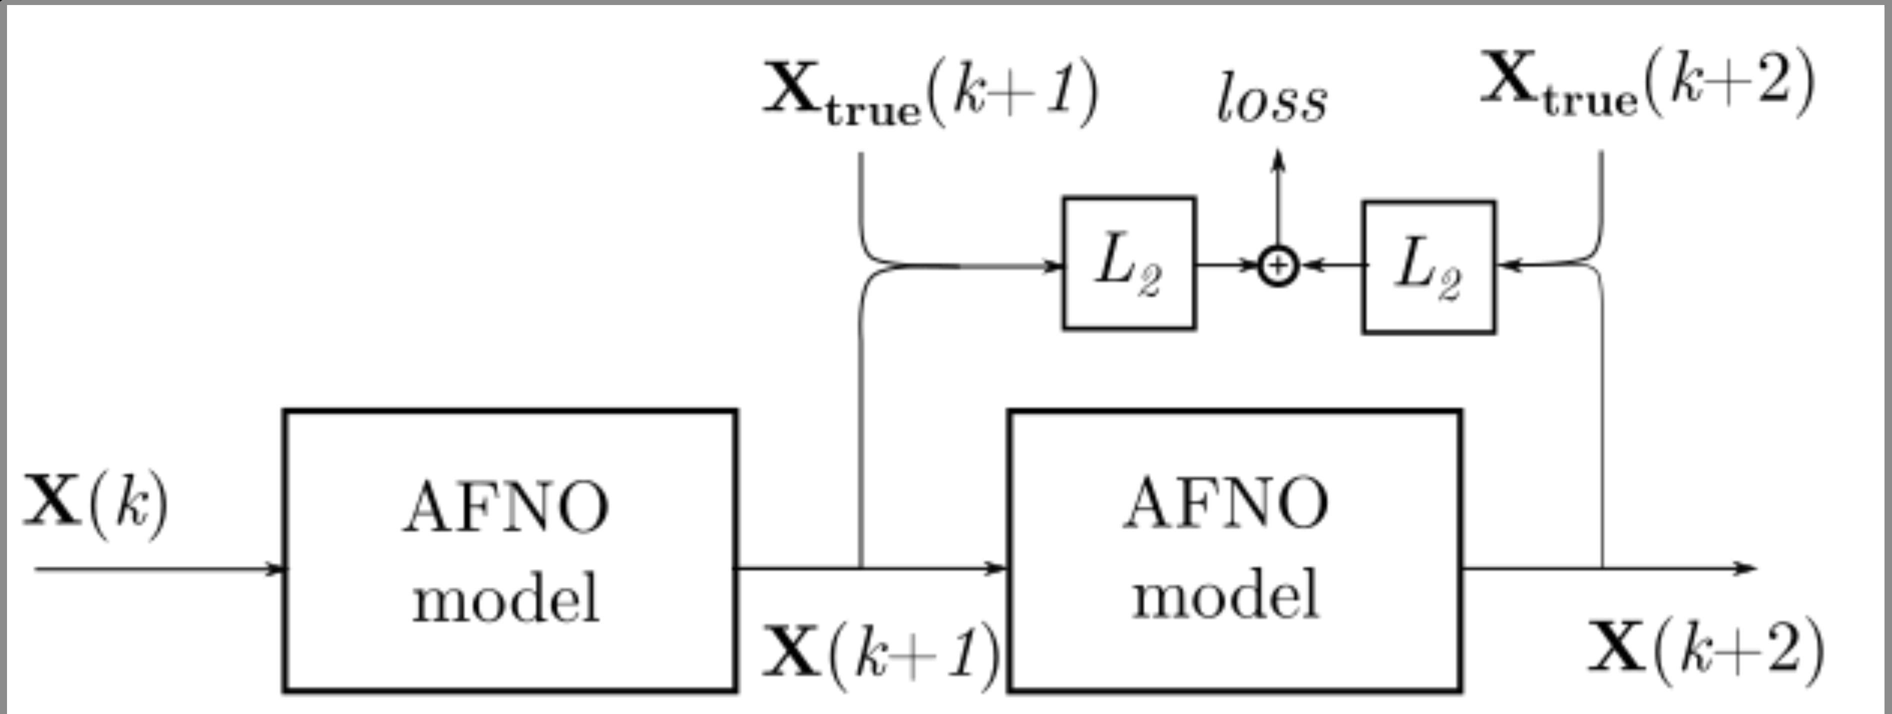

Two step fine-tuning

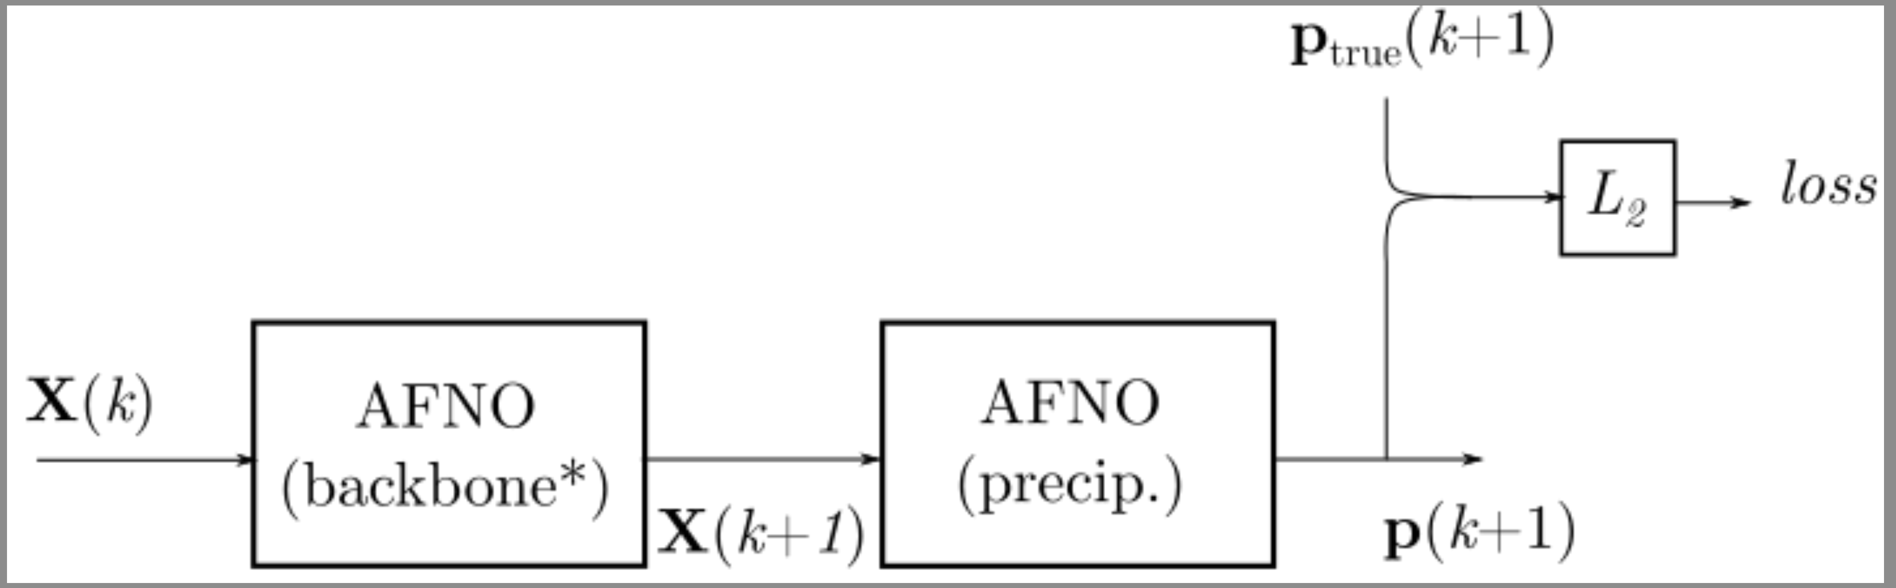

Diagnosed precip

Pathak et al. (2022), FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators, arXiv:2202.11214

FCN medium-range weather forecast skill improving with training ambition.

Could it one day outperform deterministic models? We don't yet know the limit.

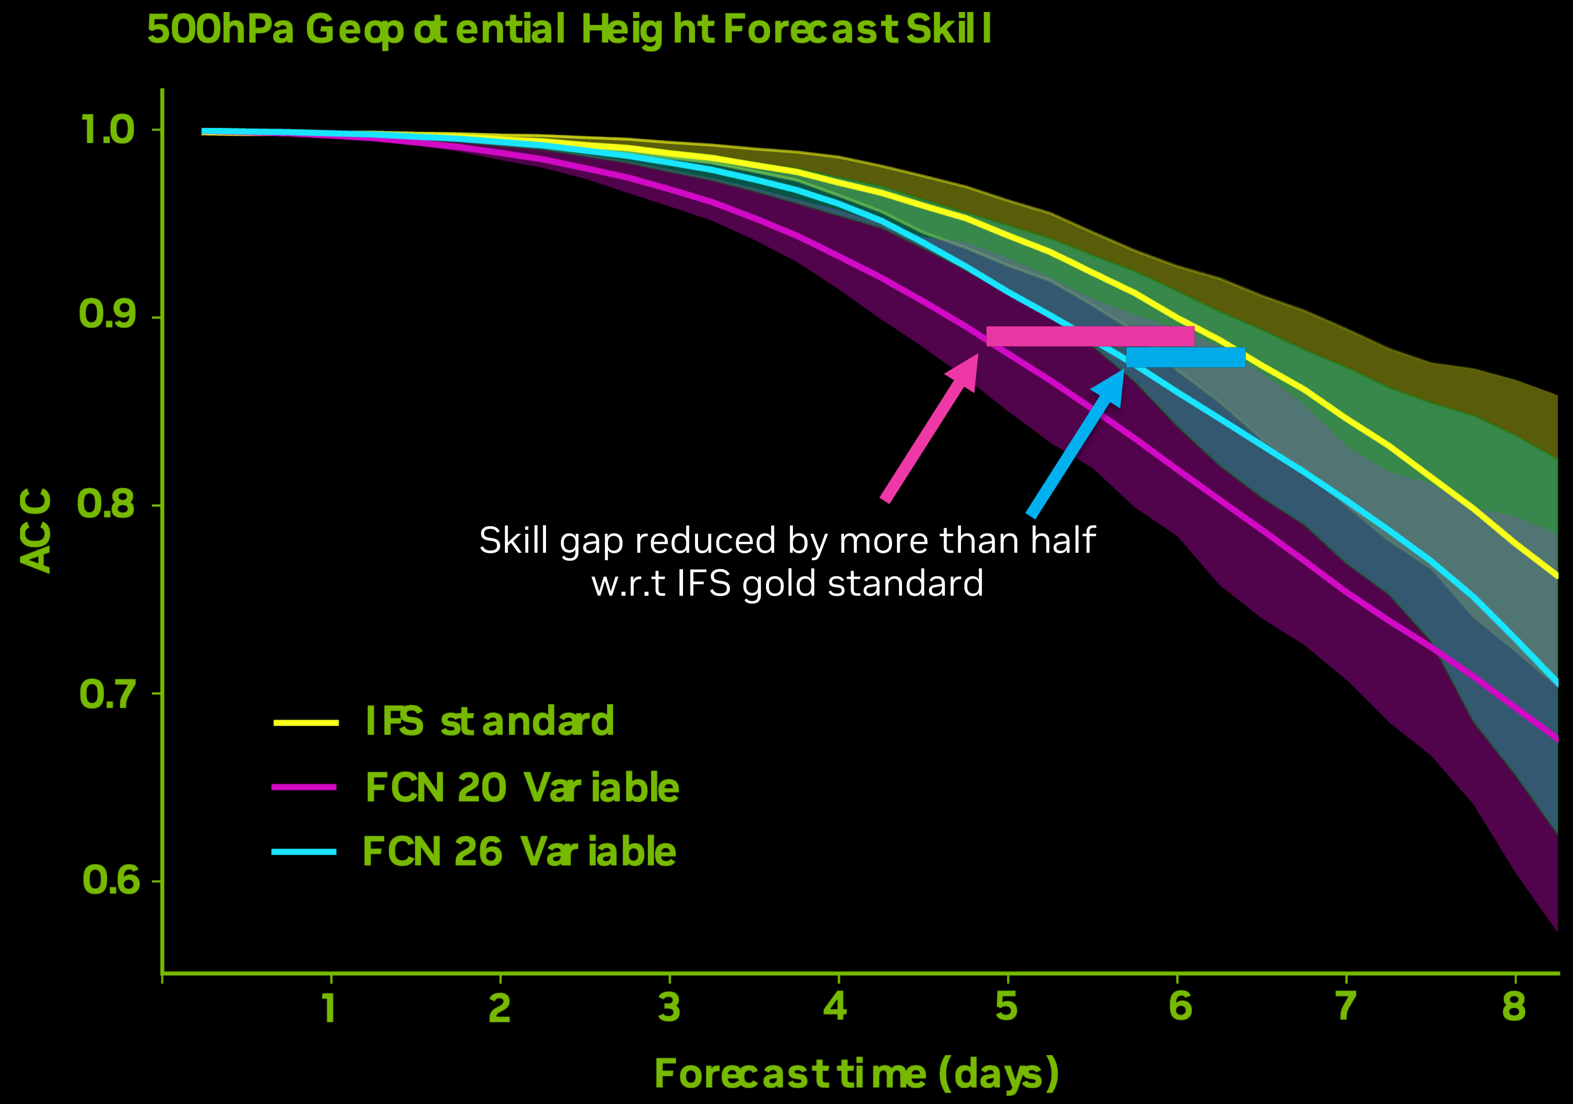

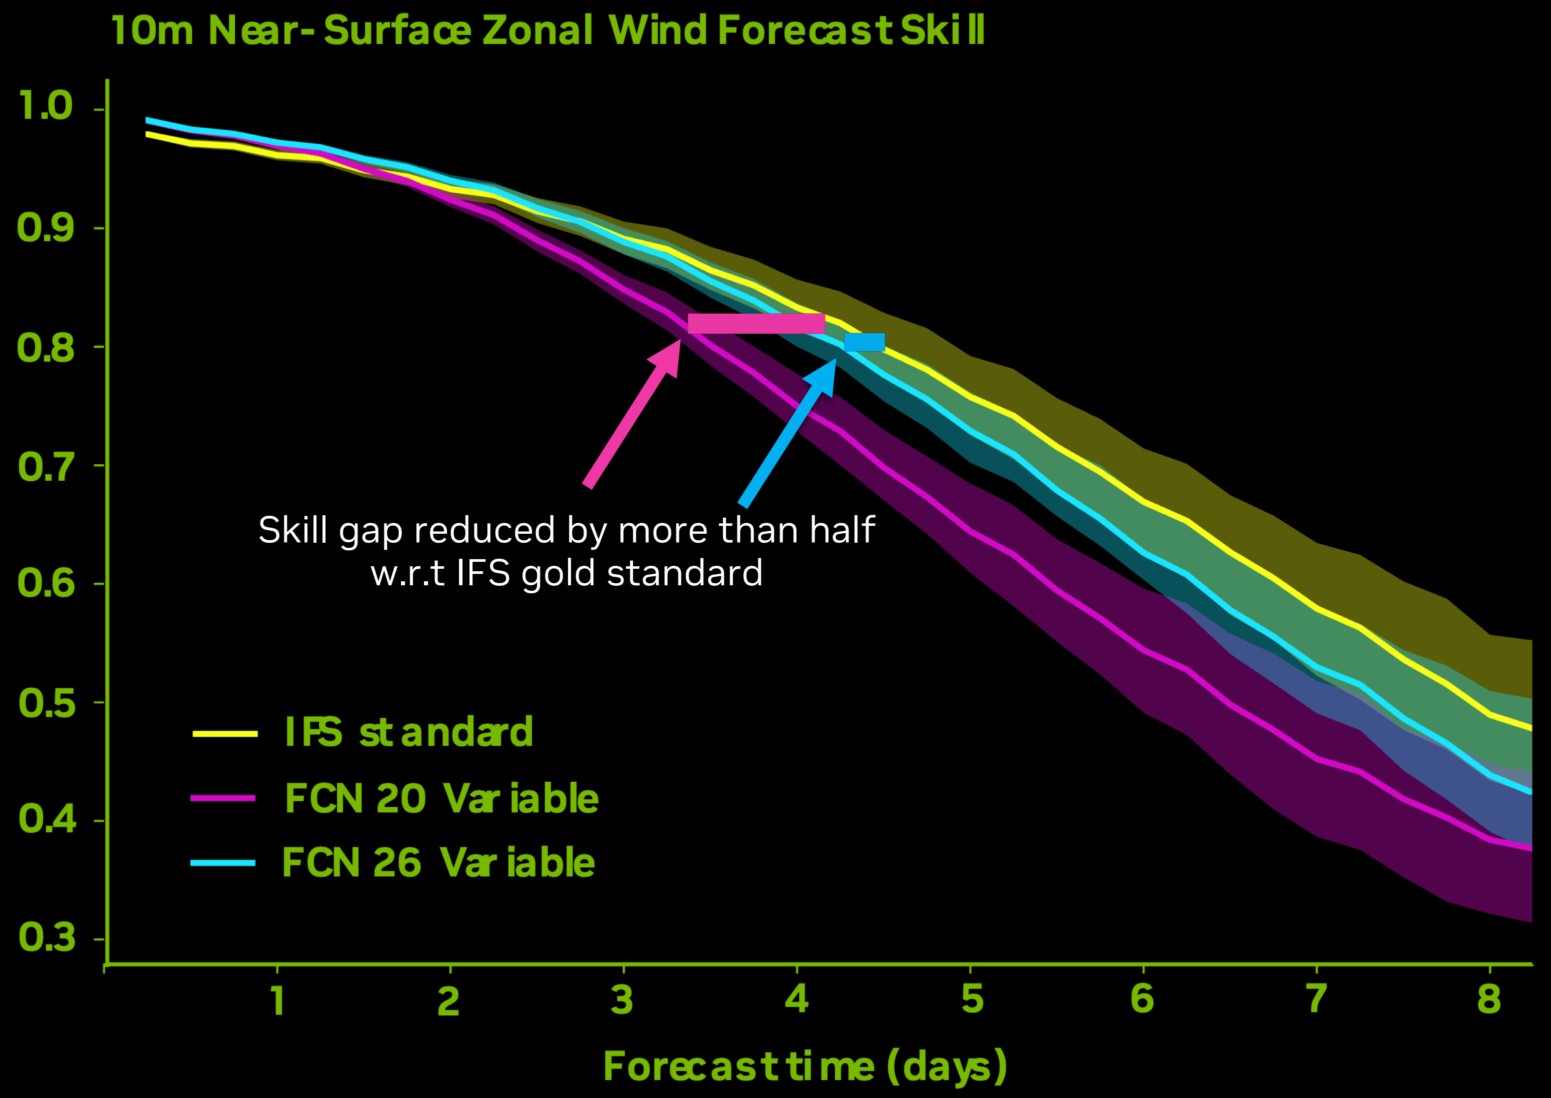

Acronym Alert:
ACC: Anomaly Correlation Coefficient (metric of weather skill)
IFS: The Integrated Forecast System, a gold standard weather model
FCN: FourCastNet, our digital twin of weather.

NVIDIA

Can FCN be initialized with real-time conditions?

Yes. Zero-shot skill transfer using initial conditions from a separate US dataset that FCN was not directly trained on.

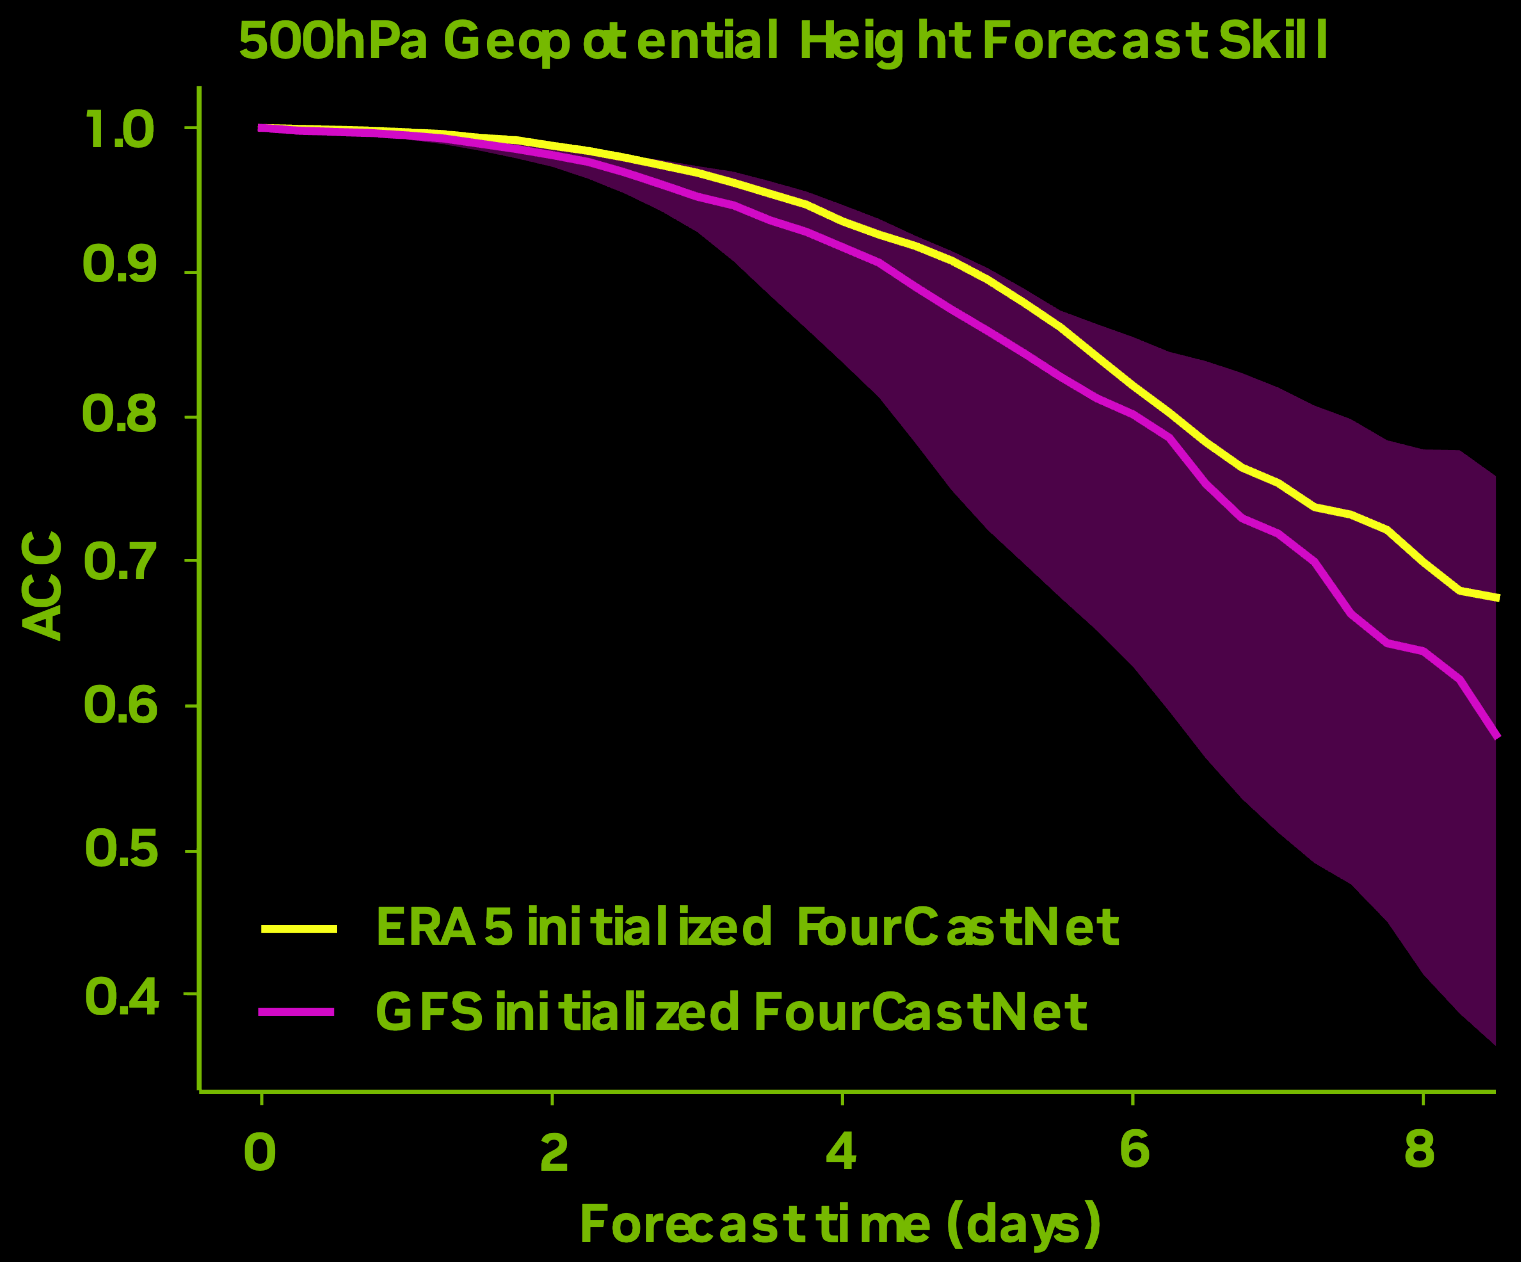

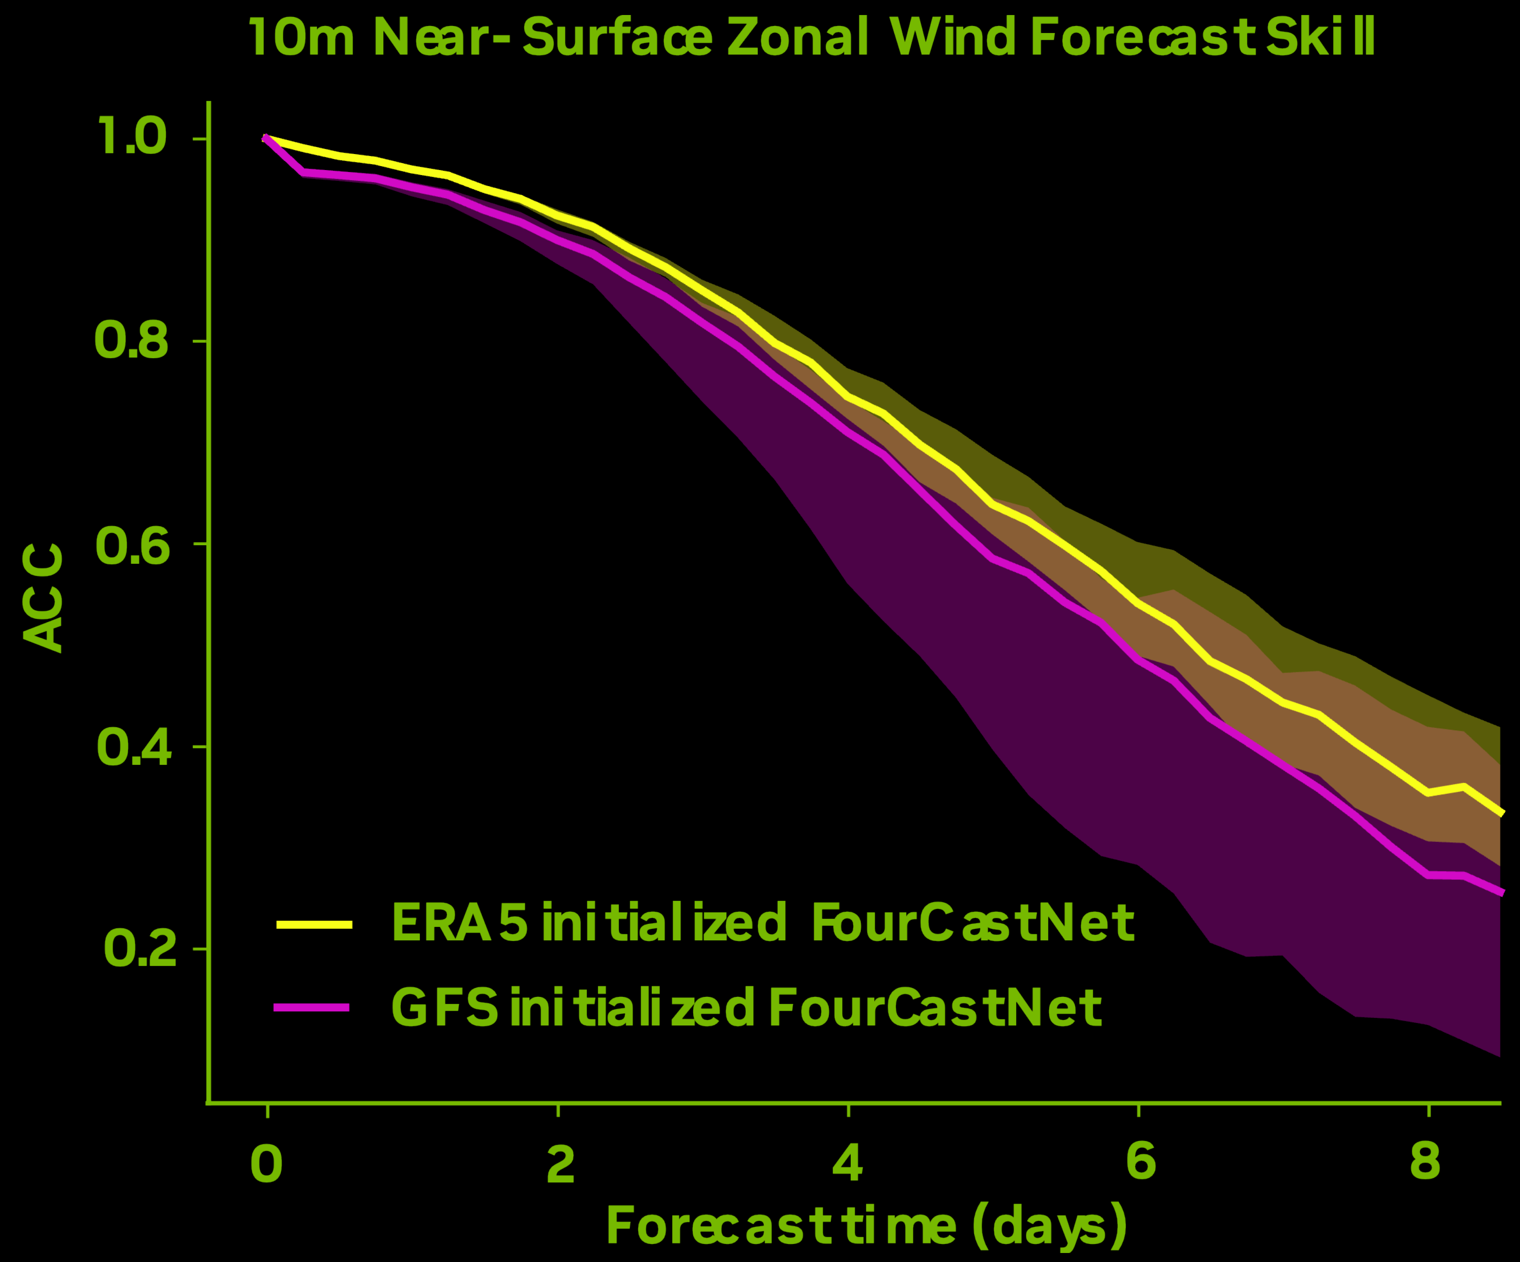

Probabilities: Spread matters as much as Skill

FCN's ensembles calibrated using initial condition uncertainty and model uncertainty (Bayesian SWA-G).

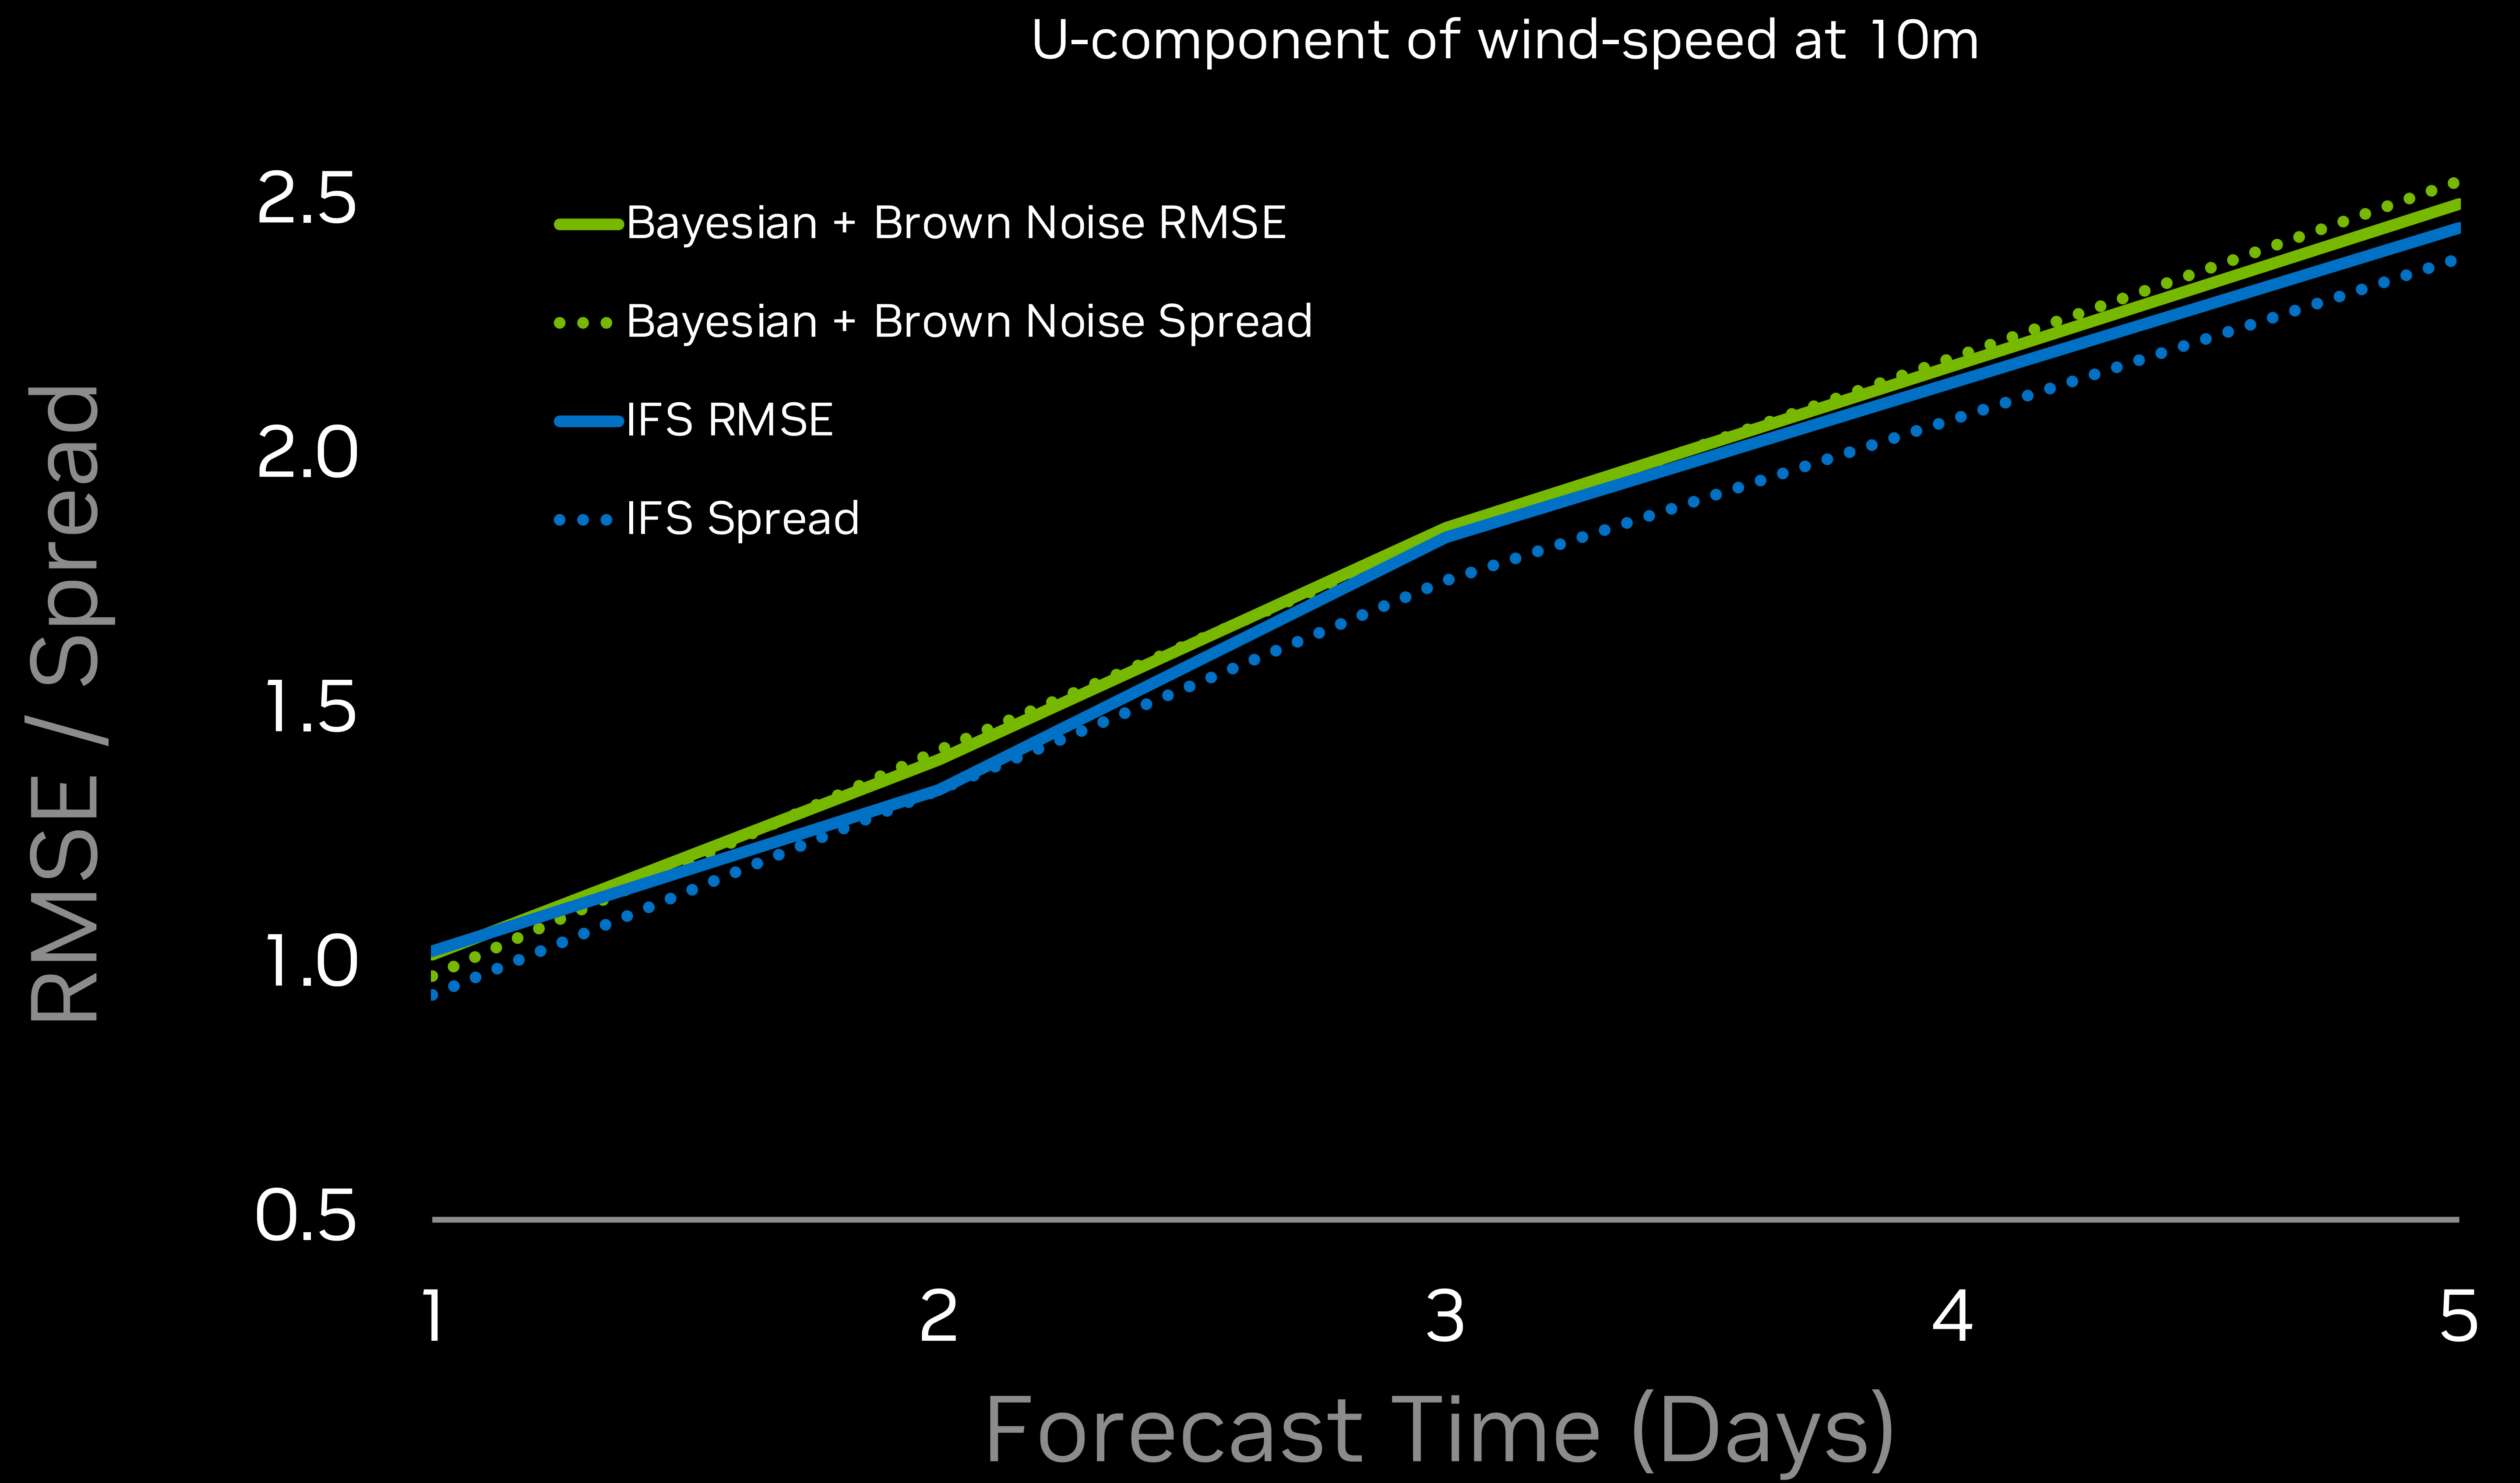

Continuous Ranked Probability Score competitive with IFS standard

FCN's ensembles calibrated using initial condition uncertainty and model uncertainty (Bayesian SWA-G).

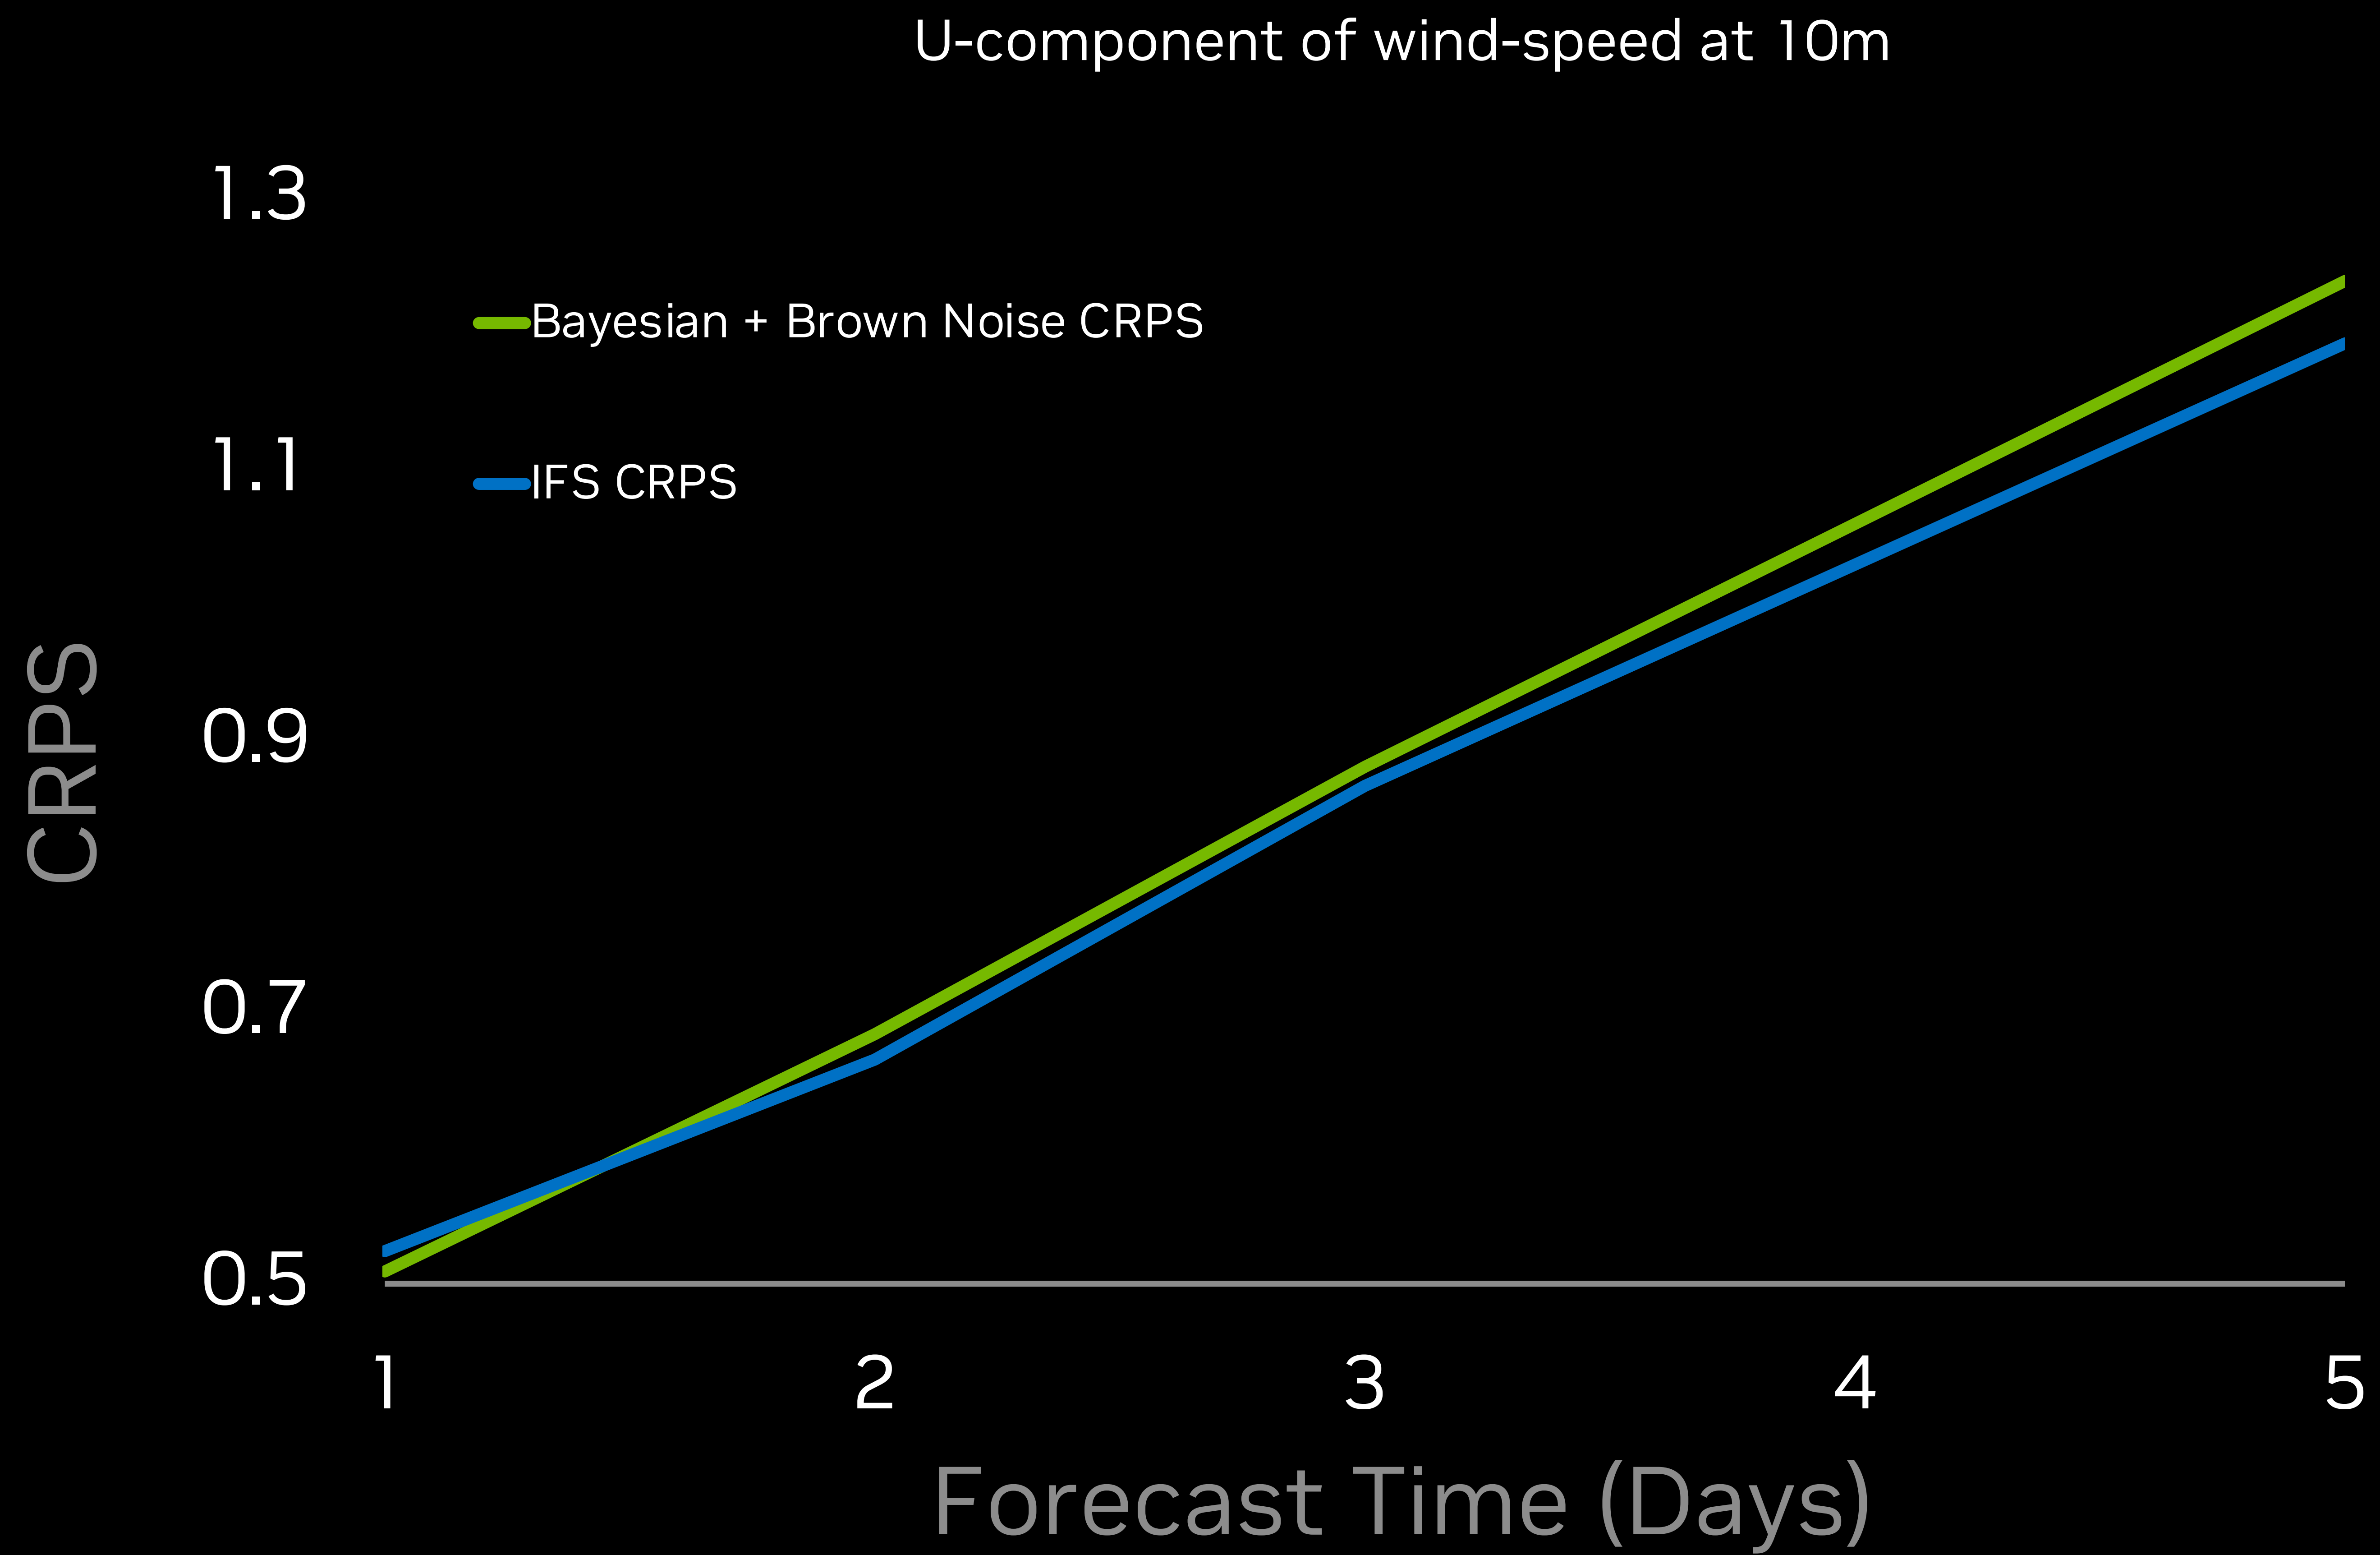

FCN has impressive skill on forecasting extremes.

Including tropical cyclones, extra-tropical cyclones, and atmospheric rivers.

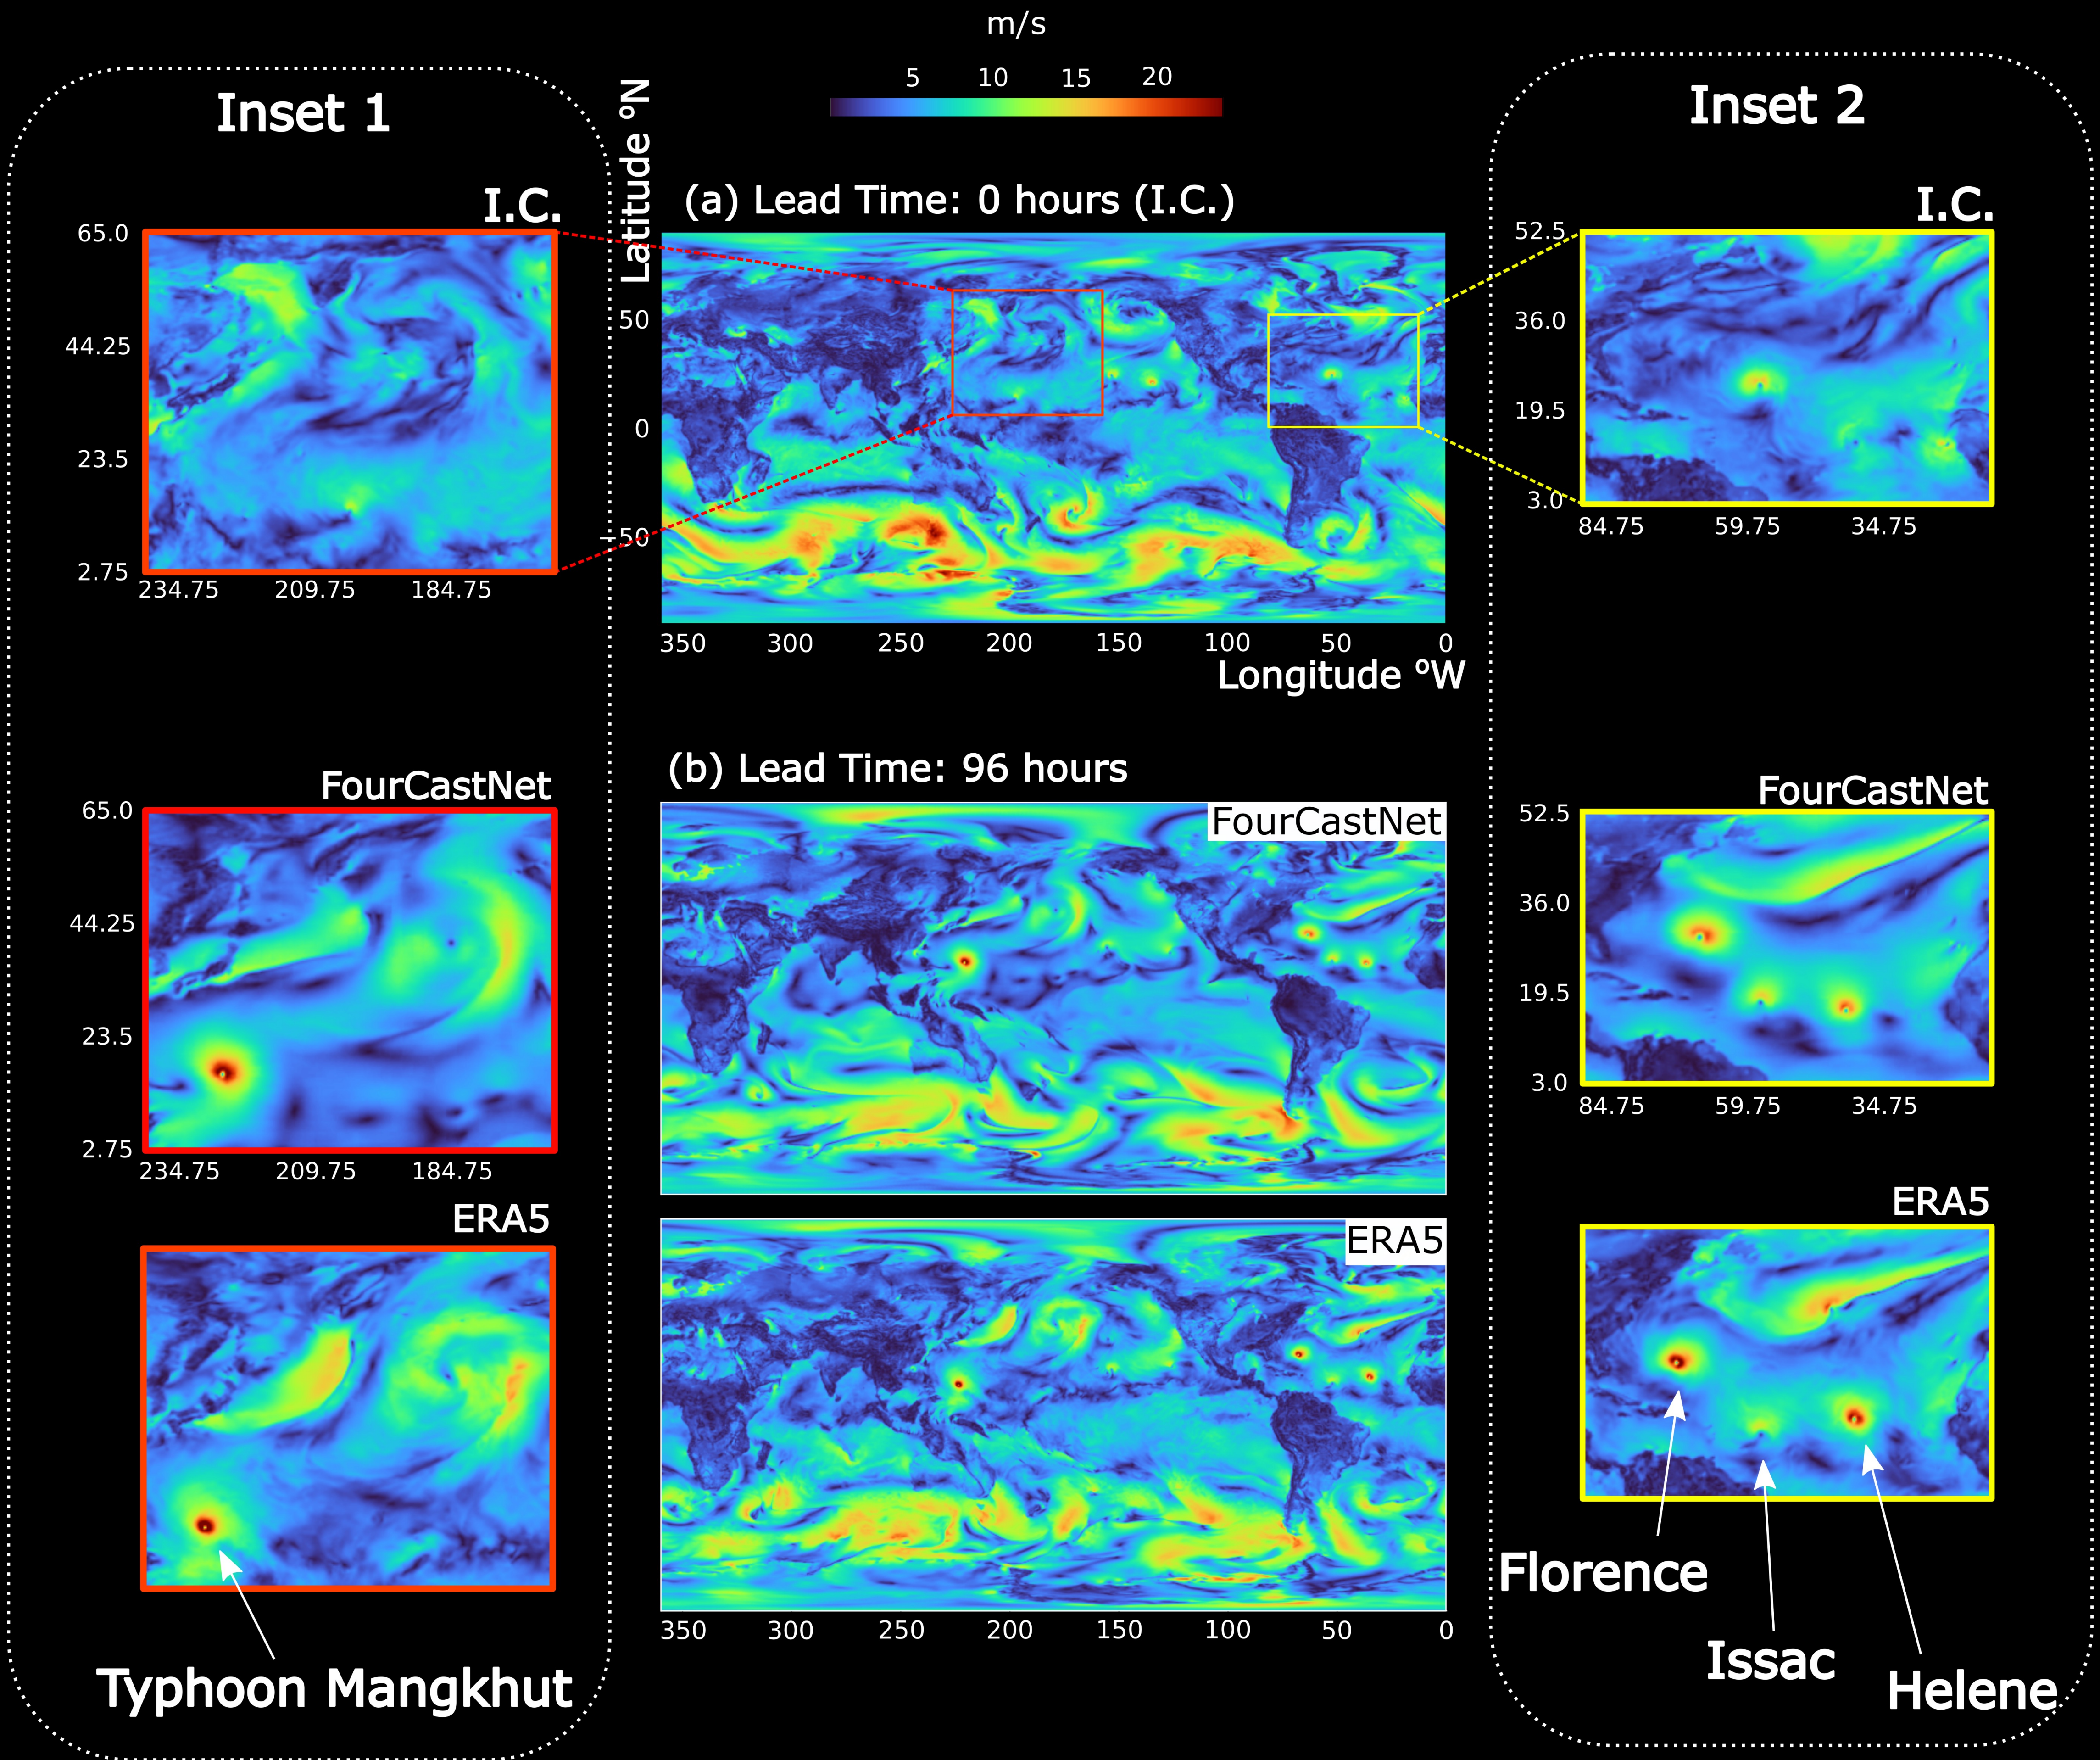

Can generative approaches improve regional extremes?

Yes. Adversarial loss improves fine scale detail skillfully.

Forecast lead time of 18 hours

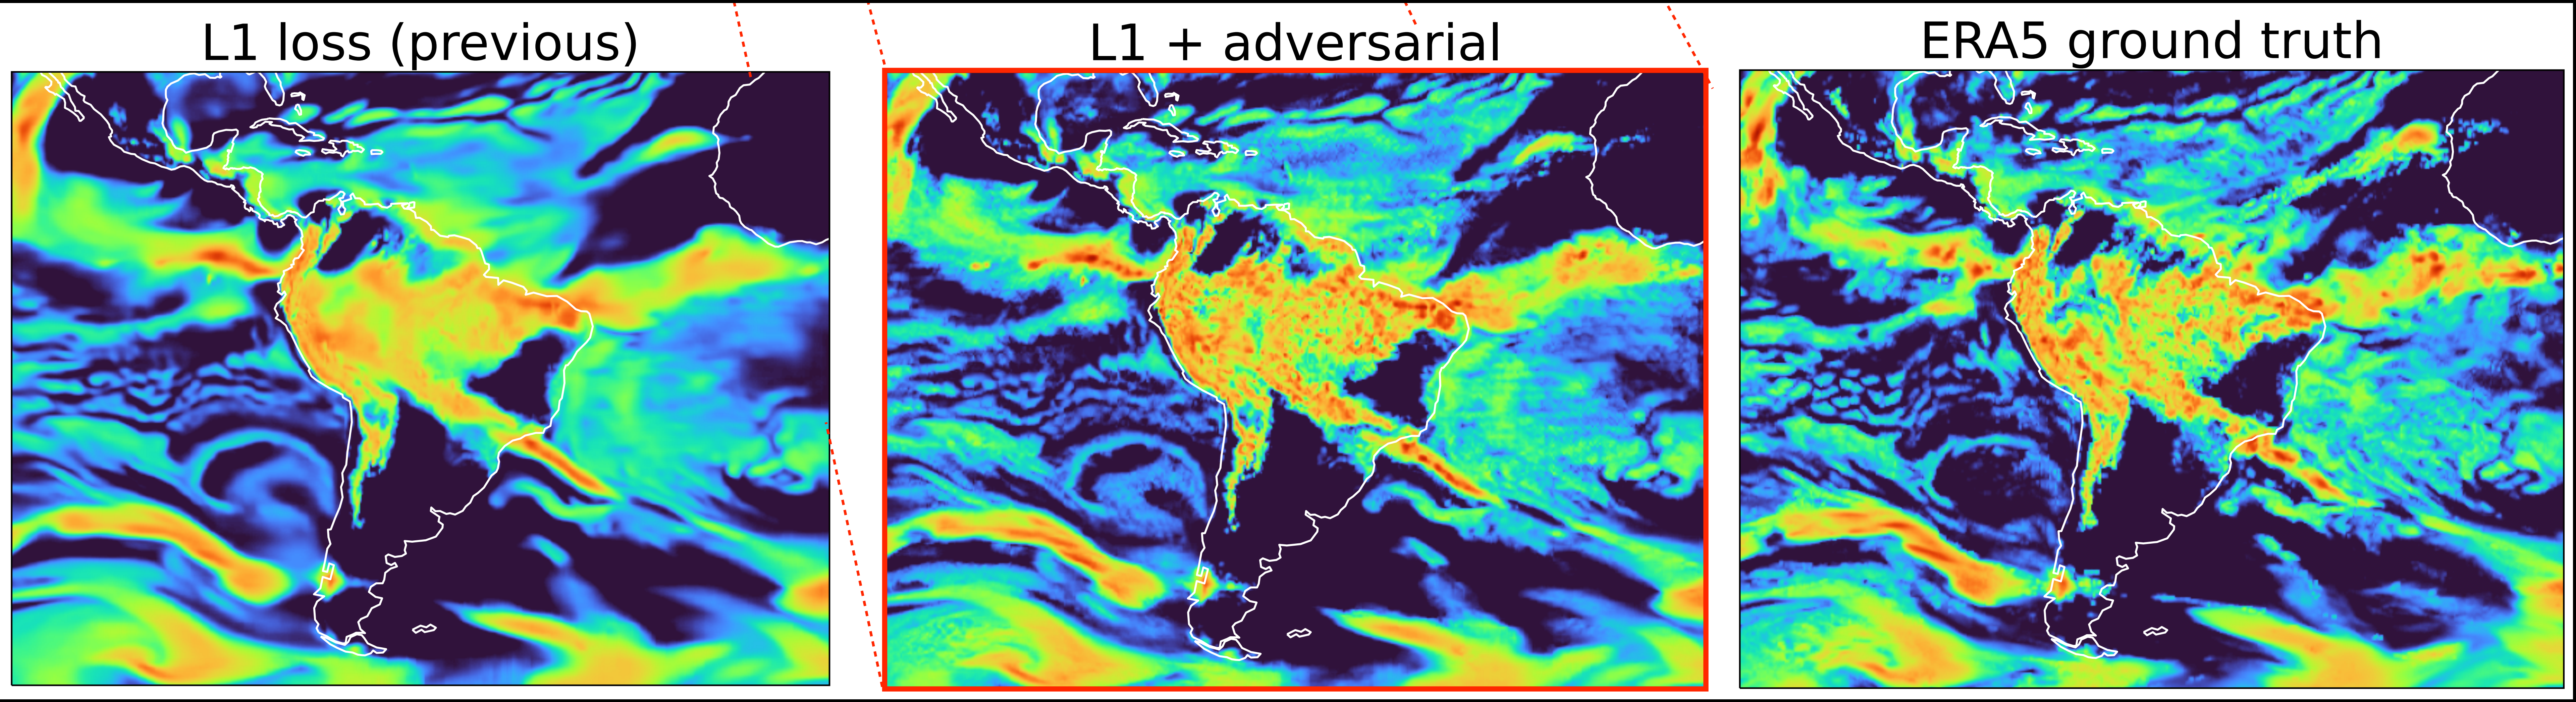

BERKELEY LAB

NERSC National Energy Research Scientific Computing Center

Duncan et al. (2022), Generative Modeling of High-resolution Precipitation Forecasts, arXiv:2210.12504

Progress in capturing extreme precipitation statistics

Adding generative adversarial loss improves predictions of rarest, most intense rainfall events.

Forecast lead time of 18 hours

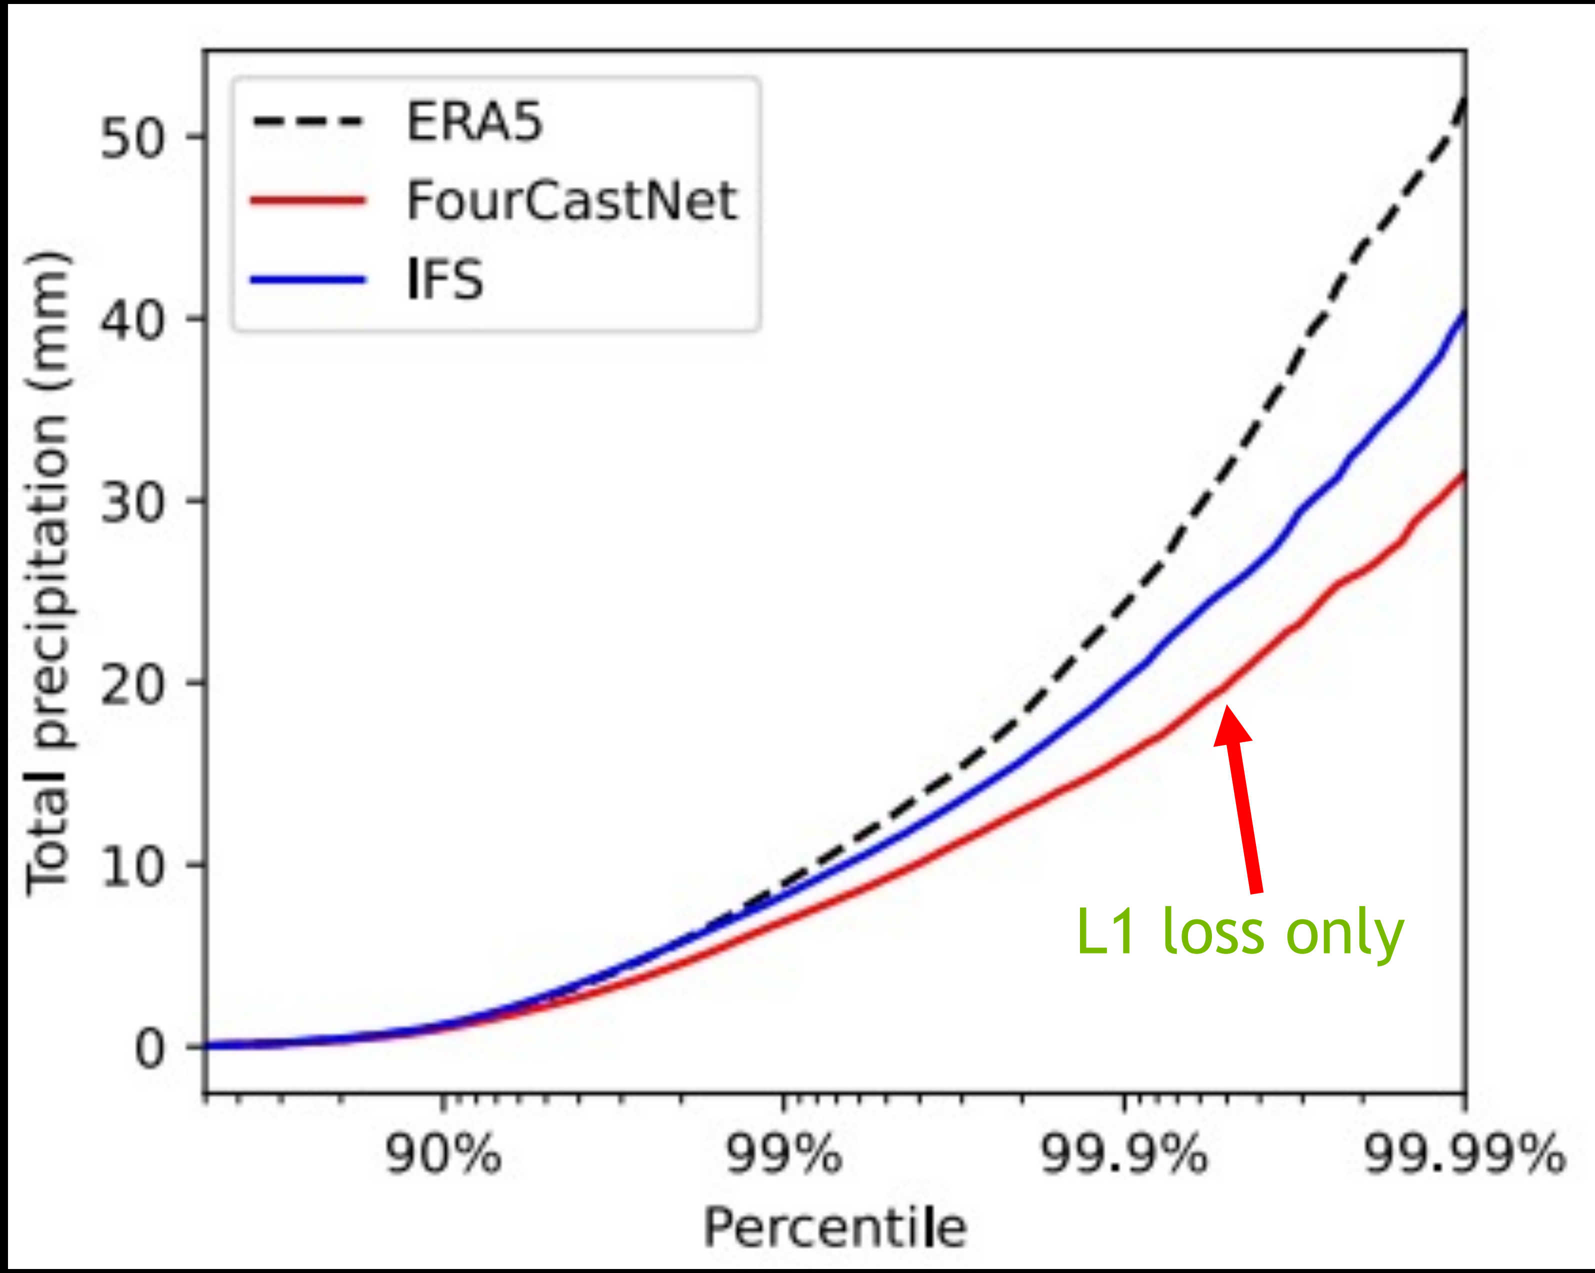

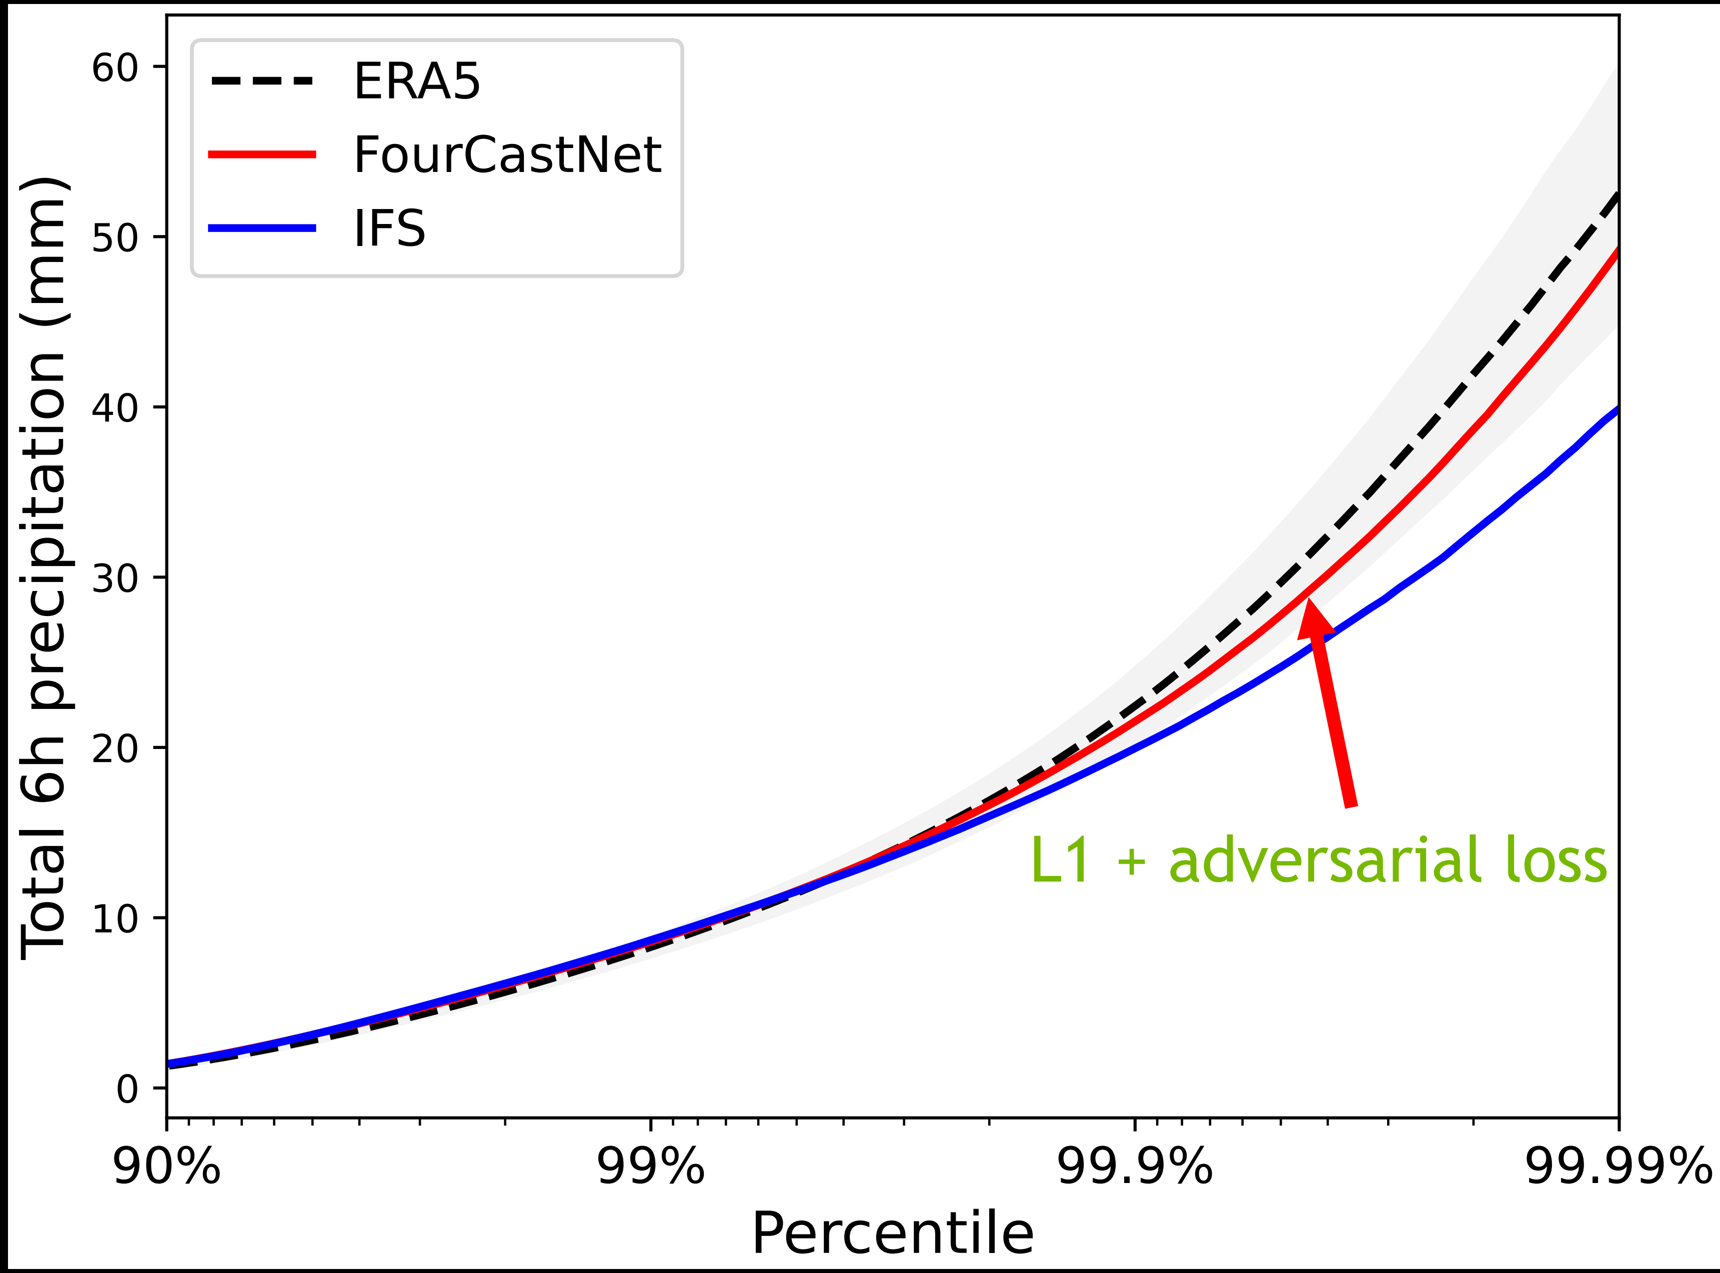

Duncan et al. (2022), Generative Modeling of High-resolution Precipitation Forecasts, arXiv:2210.12504

FCN trained on ambitious amounts of data scales efficiently up to ~ 4000 GPUs on three supercomputing systems

Thanks to full-stack AI + HPC expertise we train on a growing amount of the world's petabytes of past weather data.

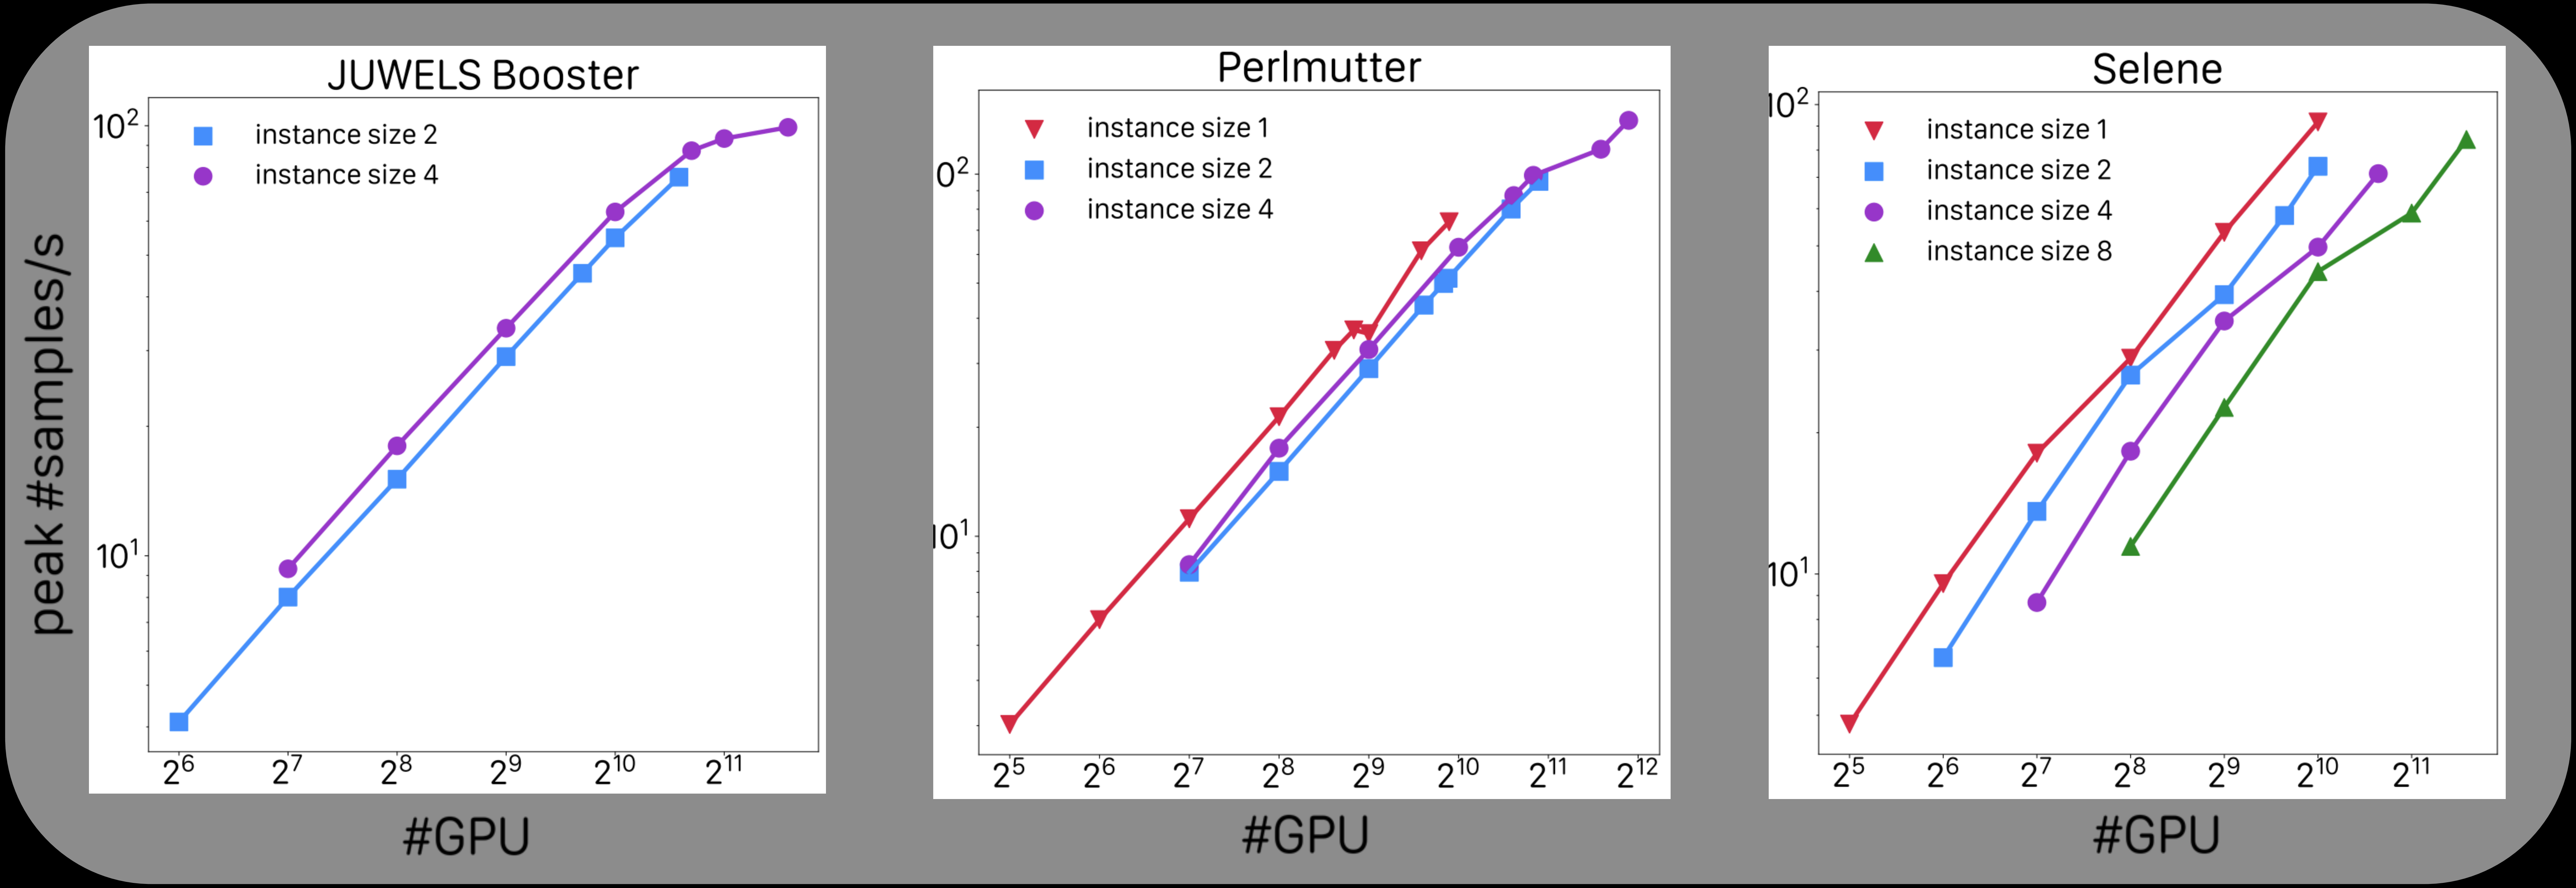

Peak performance is 140.8 petaFLOPS in mixed precision (averaged over a full epoch).
Time to solution decreased from 24+ hours to 67 minutes with model and data parallelism

FourCastNet: Accelerating Global High-Resolution Weather Forecasting using Adaptive Fourier Neural Operators, Kurth et al. (2022), https://arxiv.org/abs/2208.05419

FCN's 50,000x speedup w.r.t NWP enables massive ensembles in seconds

With over 10,000x smaller energy footprint

Computational and energy costs of 100-member ensemble forecast

	IFS (18km)	FCN (25km)	IFS / FCN
• Nodes Required	3,060	1	1530
• Latency (node-seconds)	984,000	7	44727
• Energy consumed (kJ)	271,000	7	12318

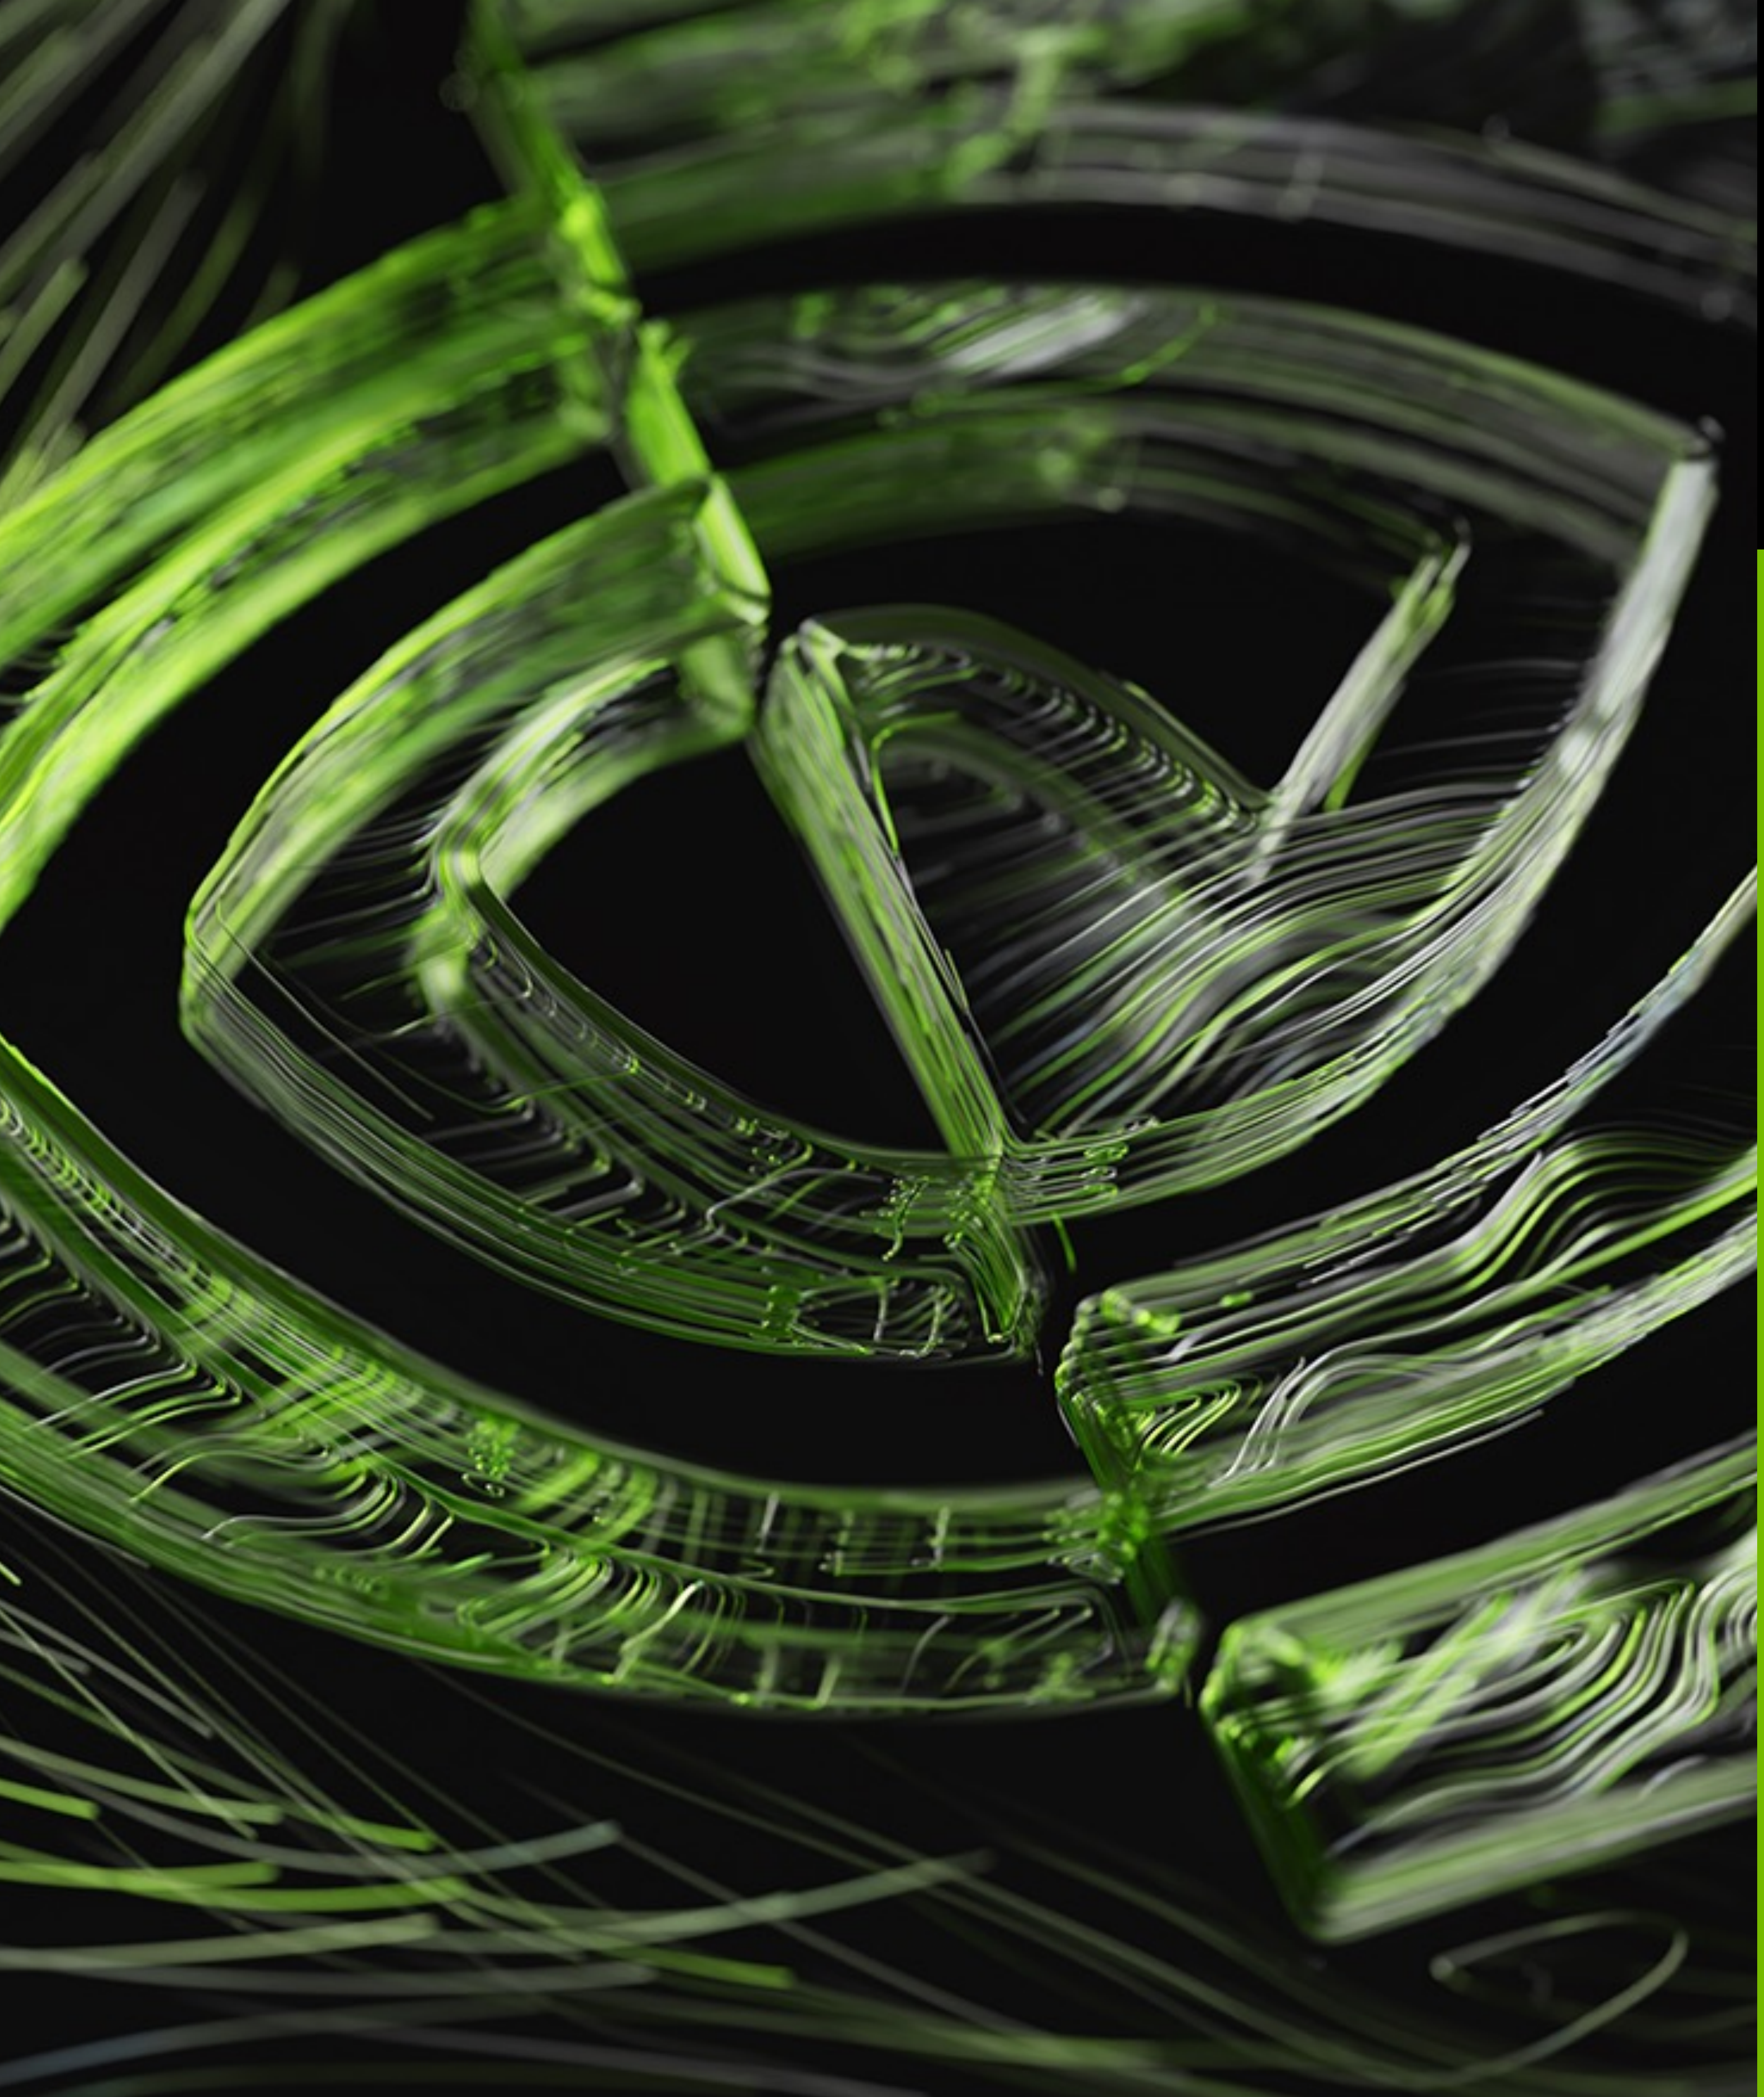

Agenda

- NVIDIA's Earth-2 initiative: The Big Picture
- FourCastNet & latest accomplishments in Weather.
- Beyond Weather, towards Climate.
- Digital Twinning platforms

Given Future Data, FourCastNet's Speed Allows Fast Tethering

AI nimbly generates details between "checkpoints" saved only infrequently from physics-based climate simulations

-- Bjorn Stevens, GTC 2021

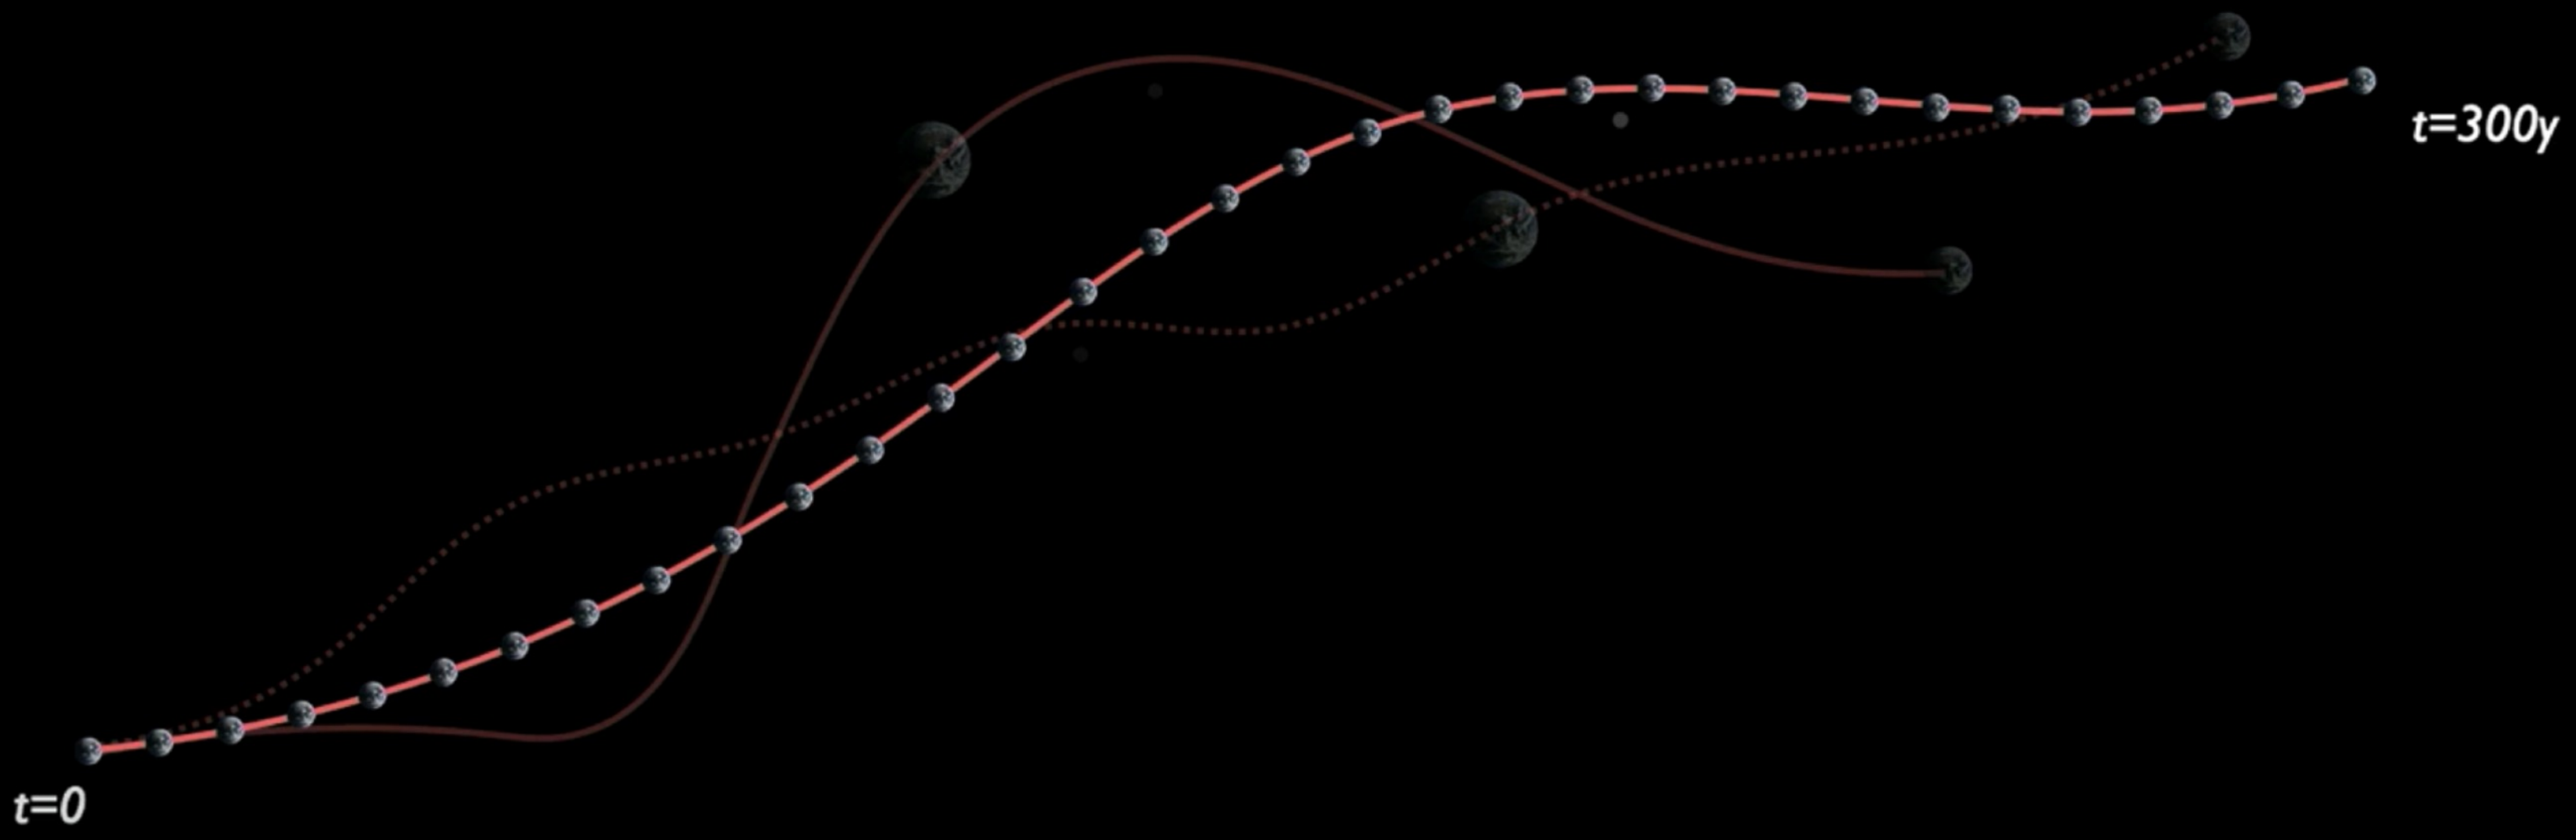

Open Question

For how long can full-AI models like FourCastNet be trusted to "tether" between climate checkpoints?

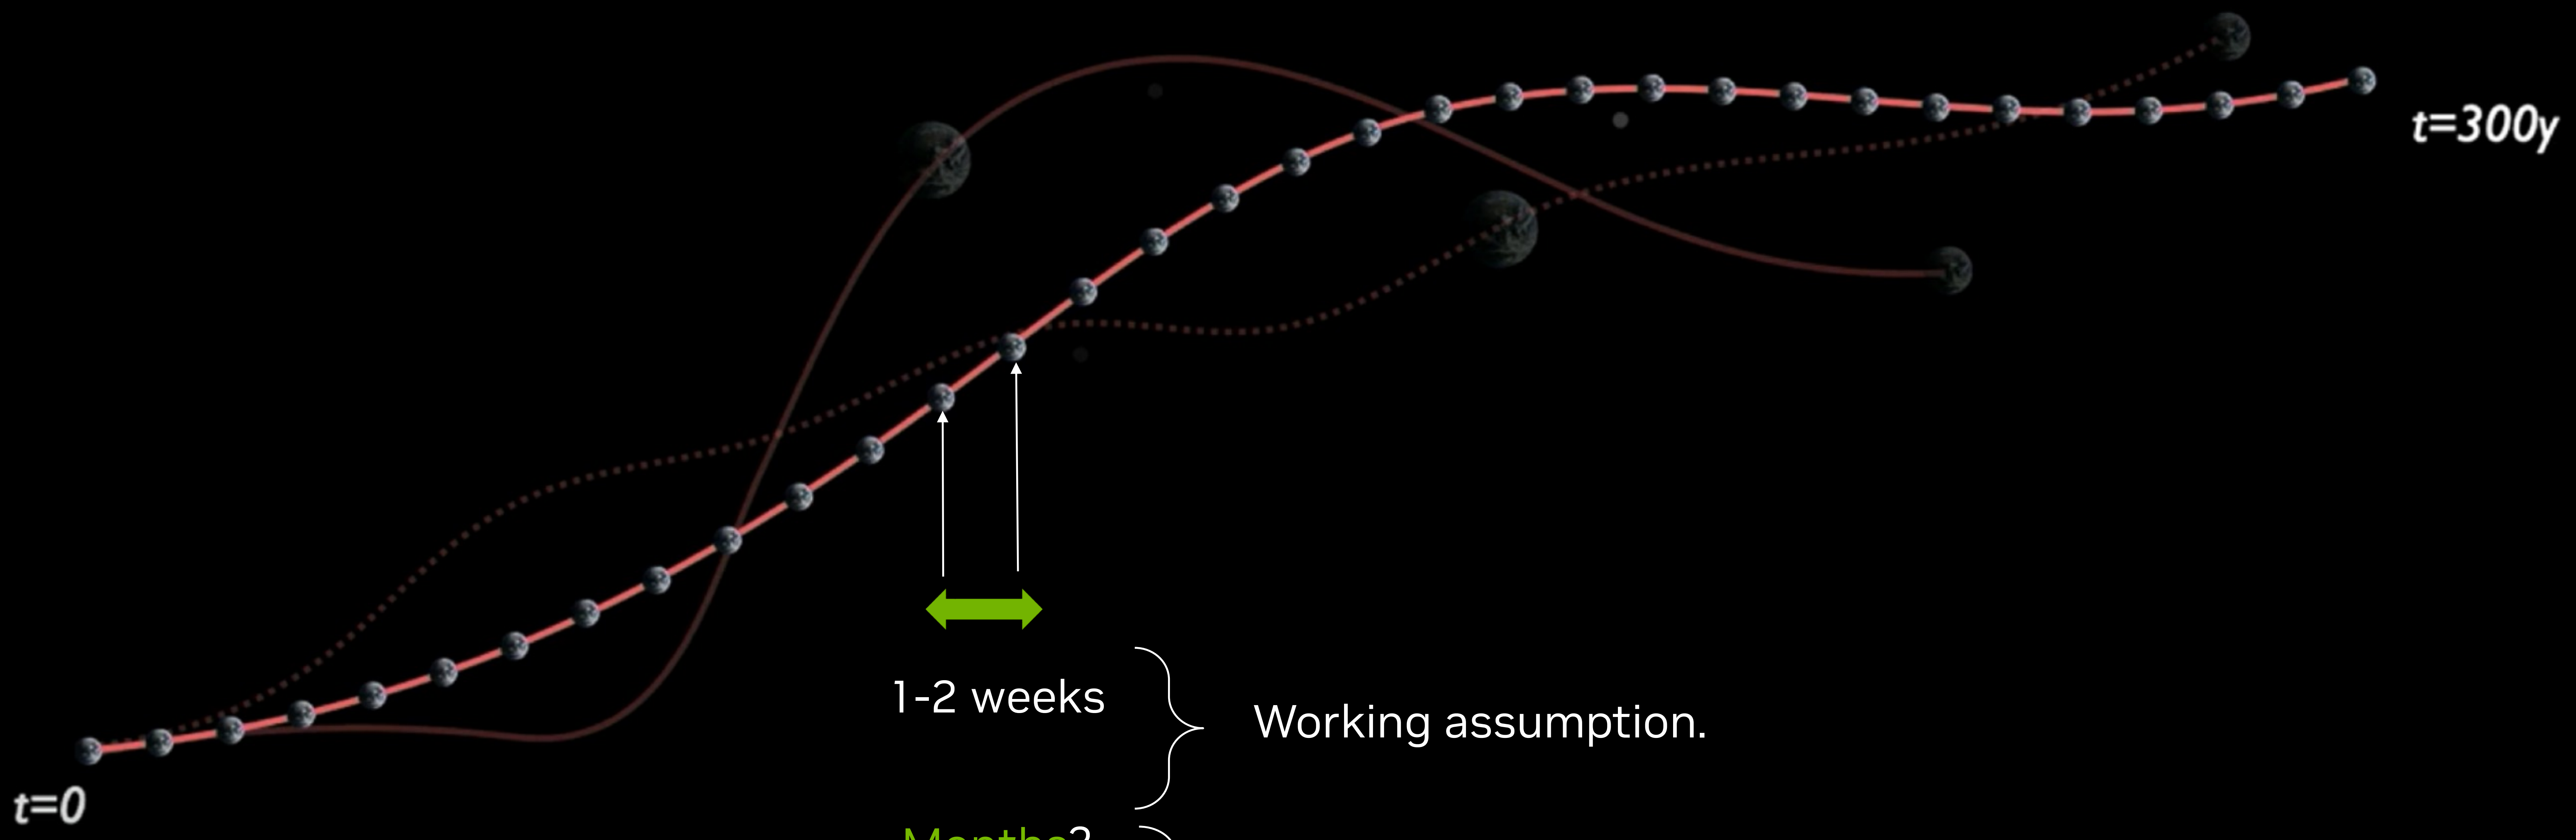

1-2 weeks — Working assumption.

Months?
Years?? — Very interesting if possible... stay tuned.

FourCastNet Roadmap:
Beyond weather, towards climate.

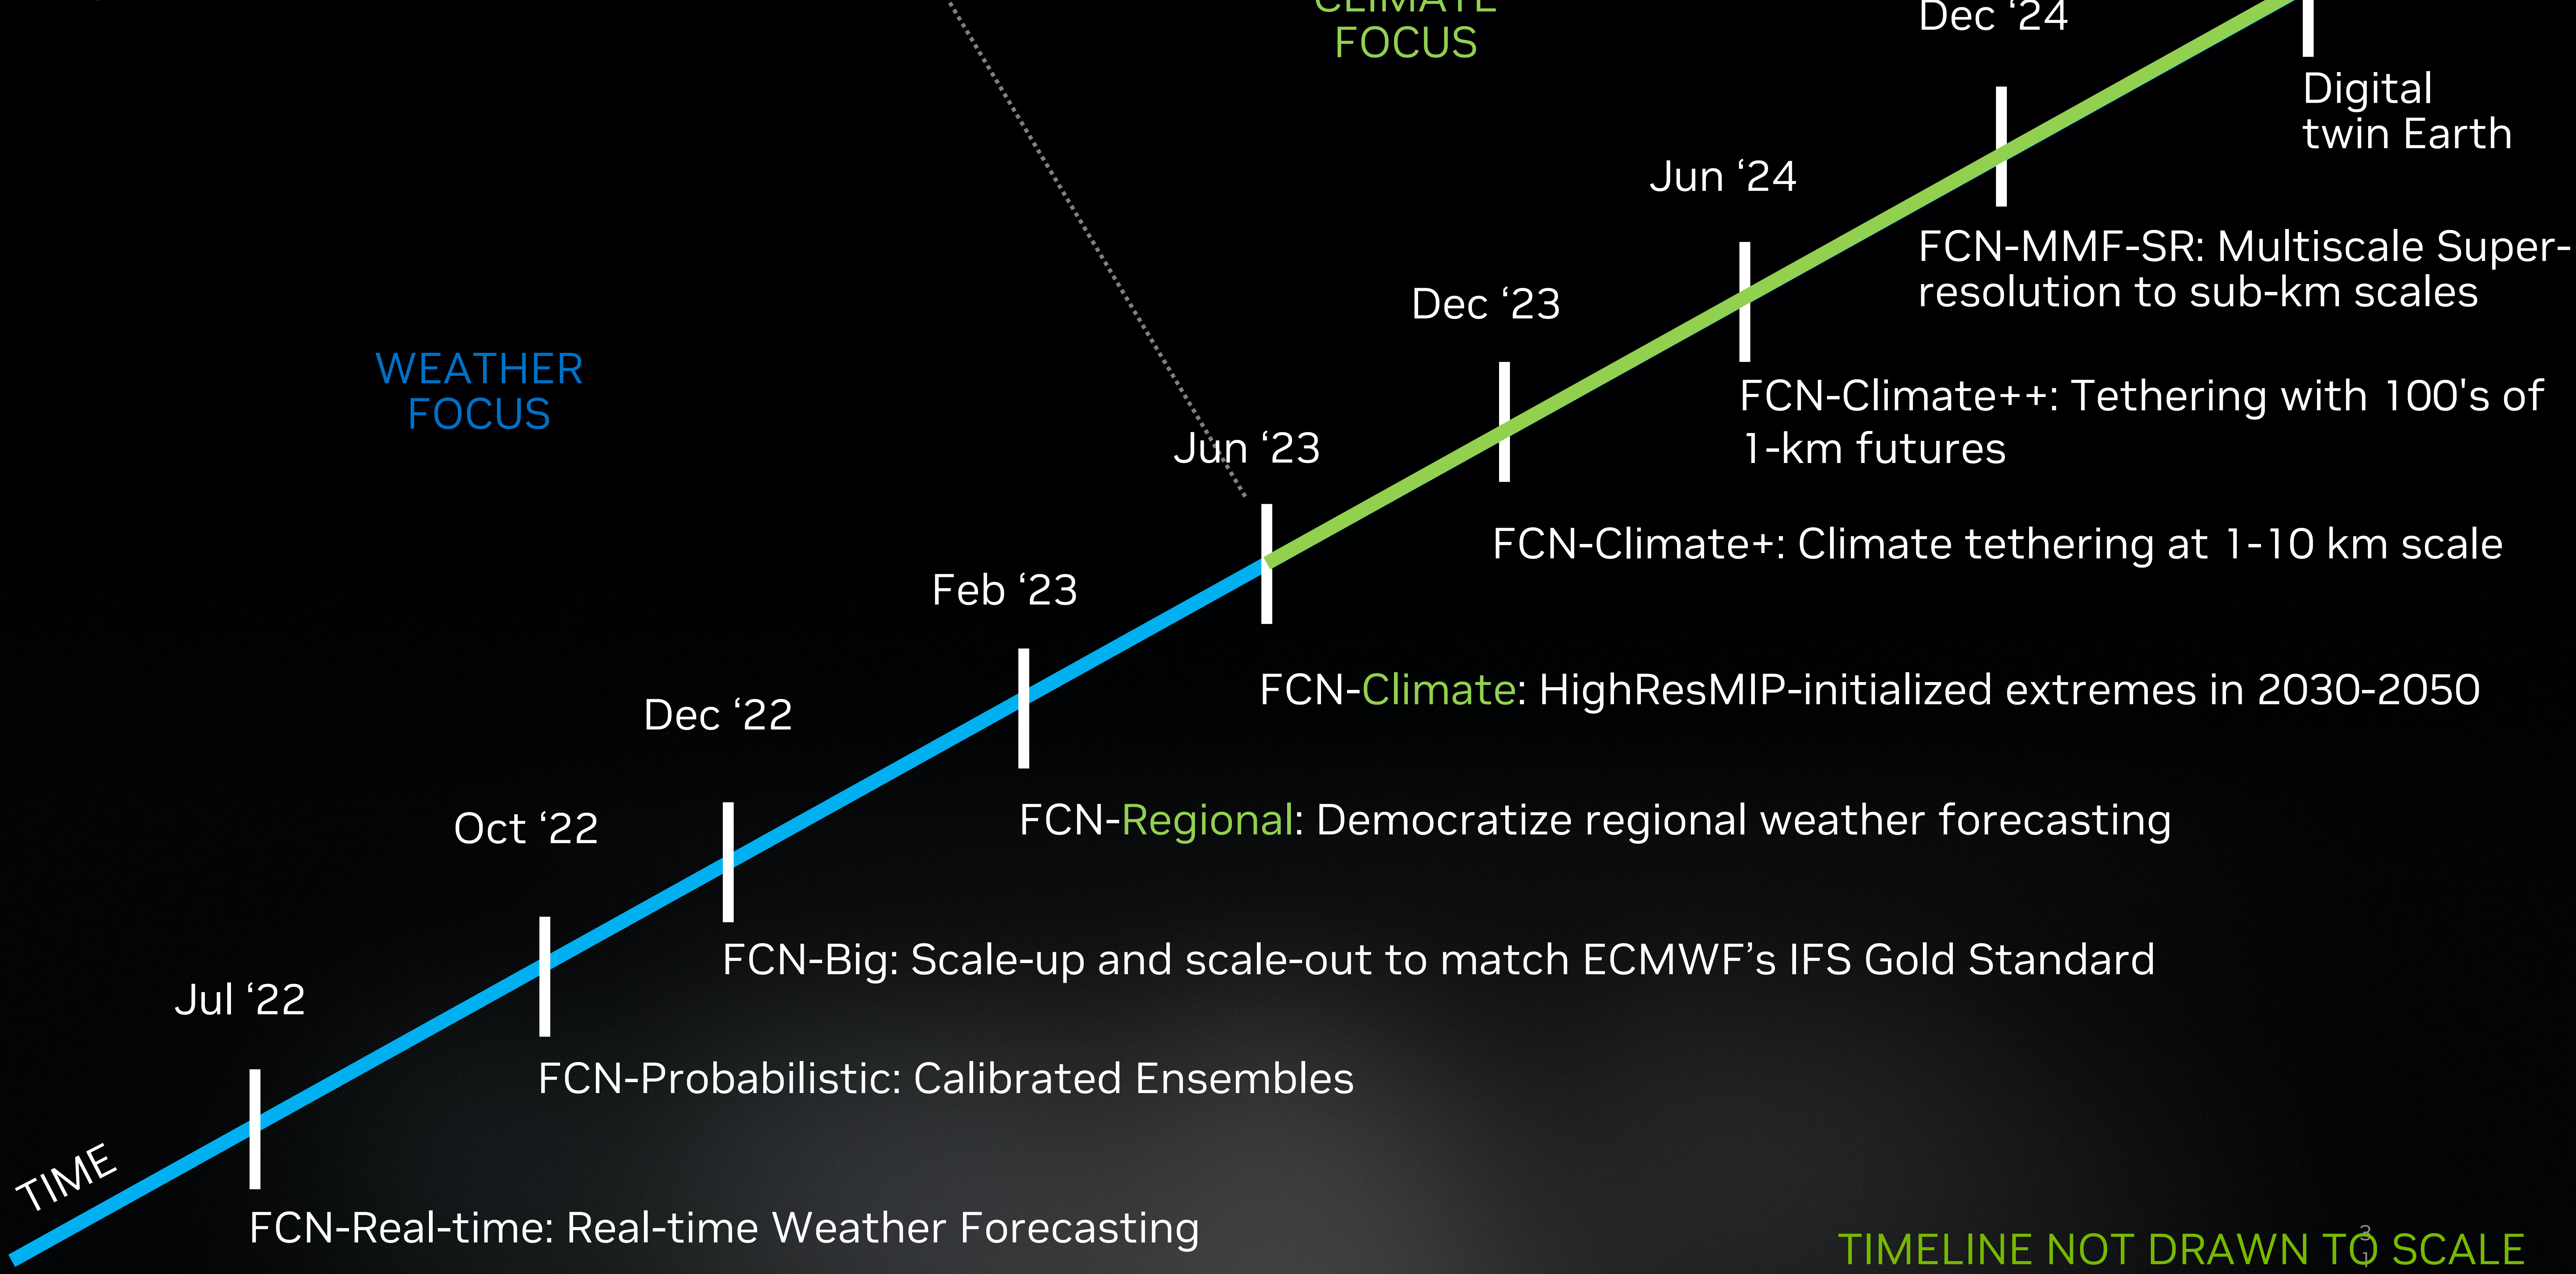

To Begin, we can Tether to Existing Climate Predictions.

Using the world's current data library of 25-km resolution HighResMIP climate predictions.

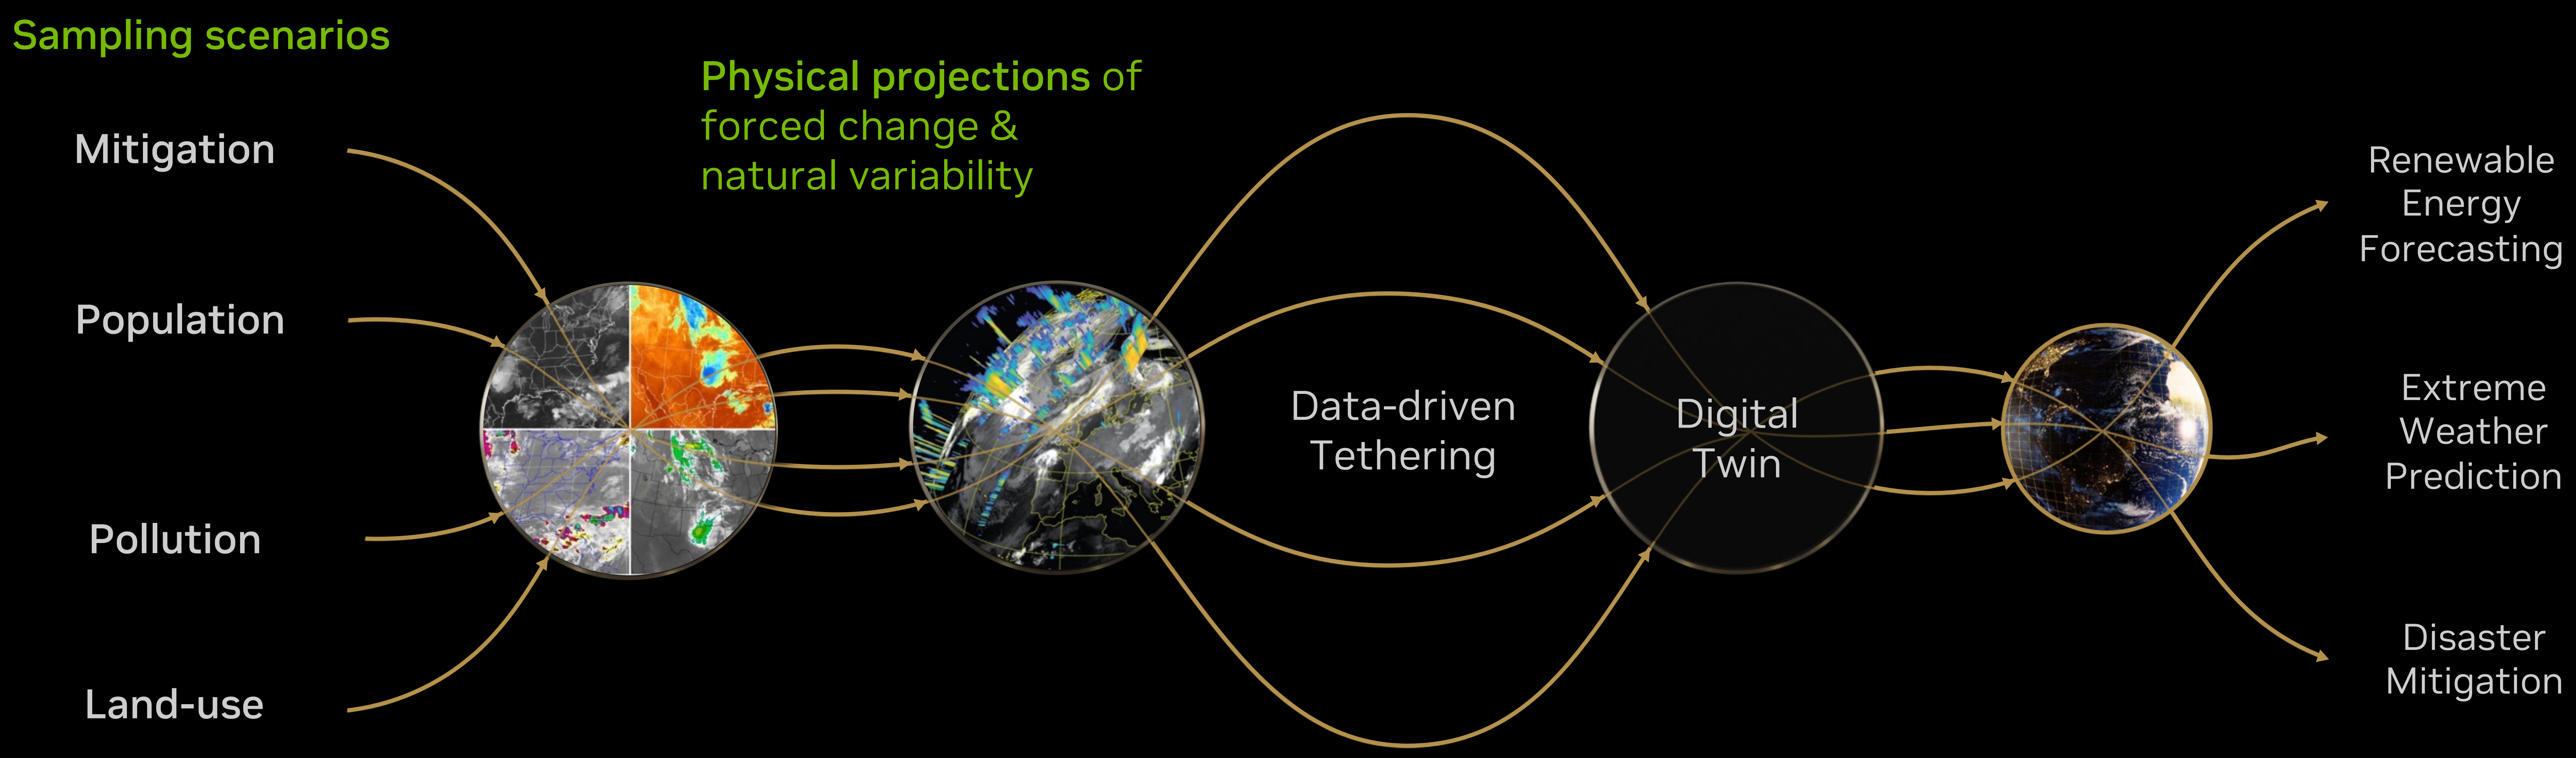

But Eventually we want to Tether to km-scale Predictions.

Because credible cloud feedbacks and storm dynamics from km-scale simulators matter to predicting regional risk.

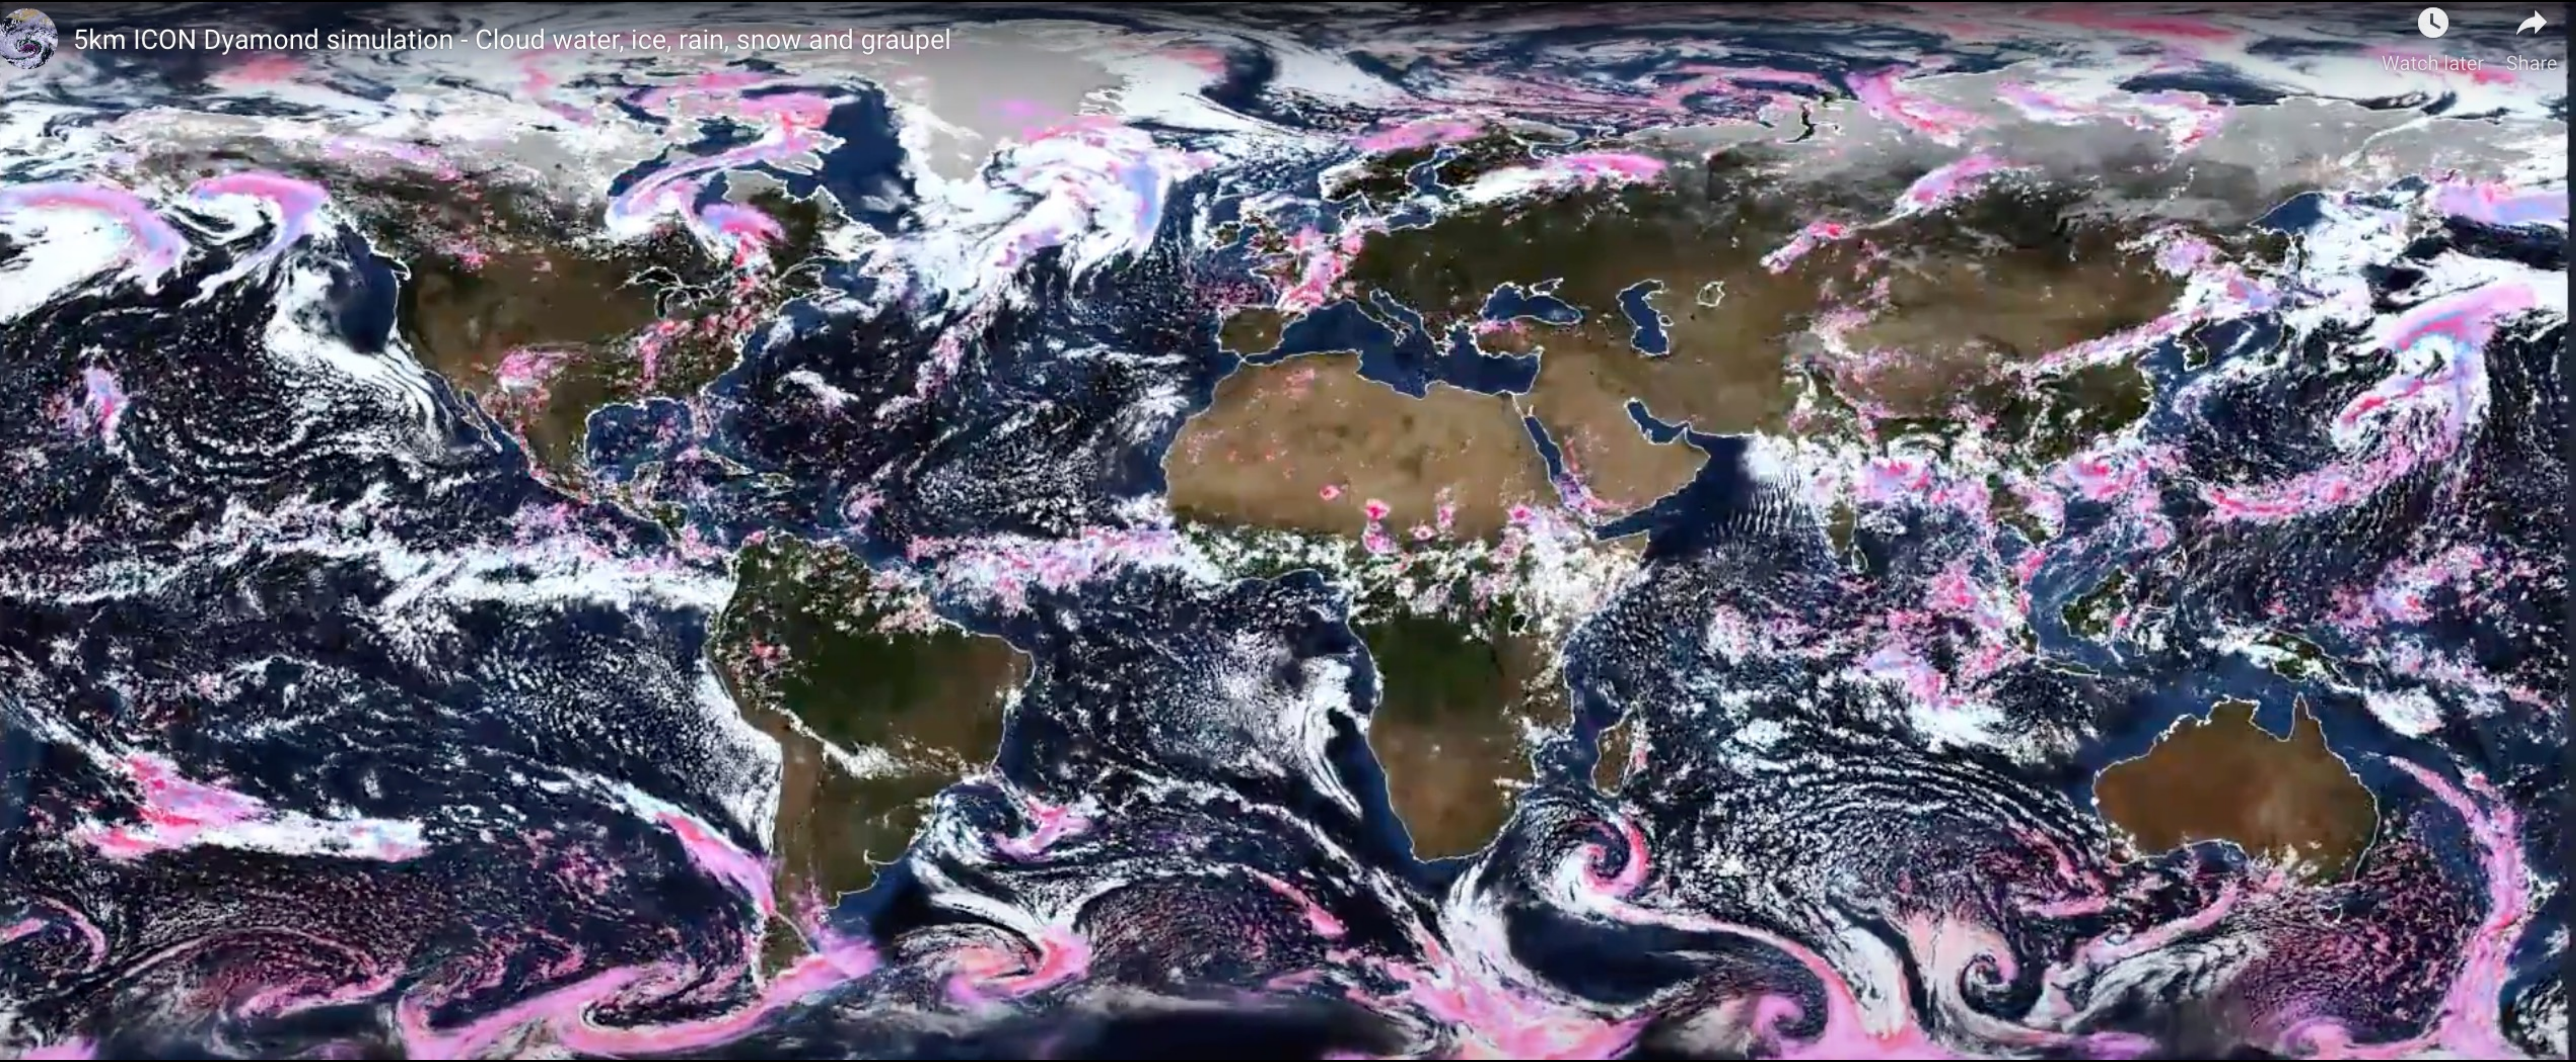

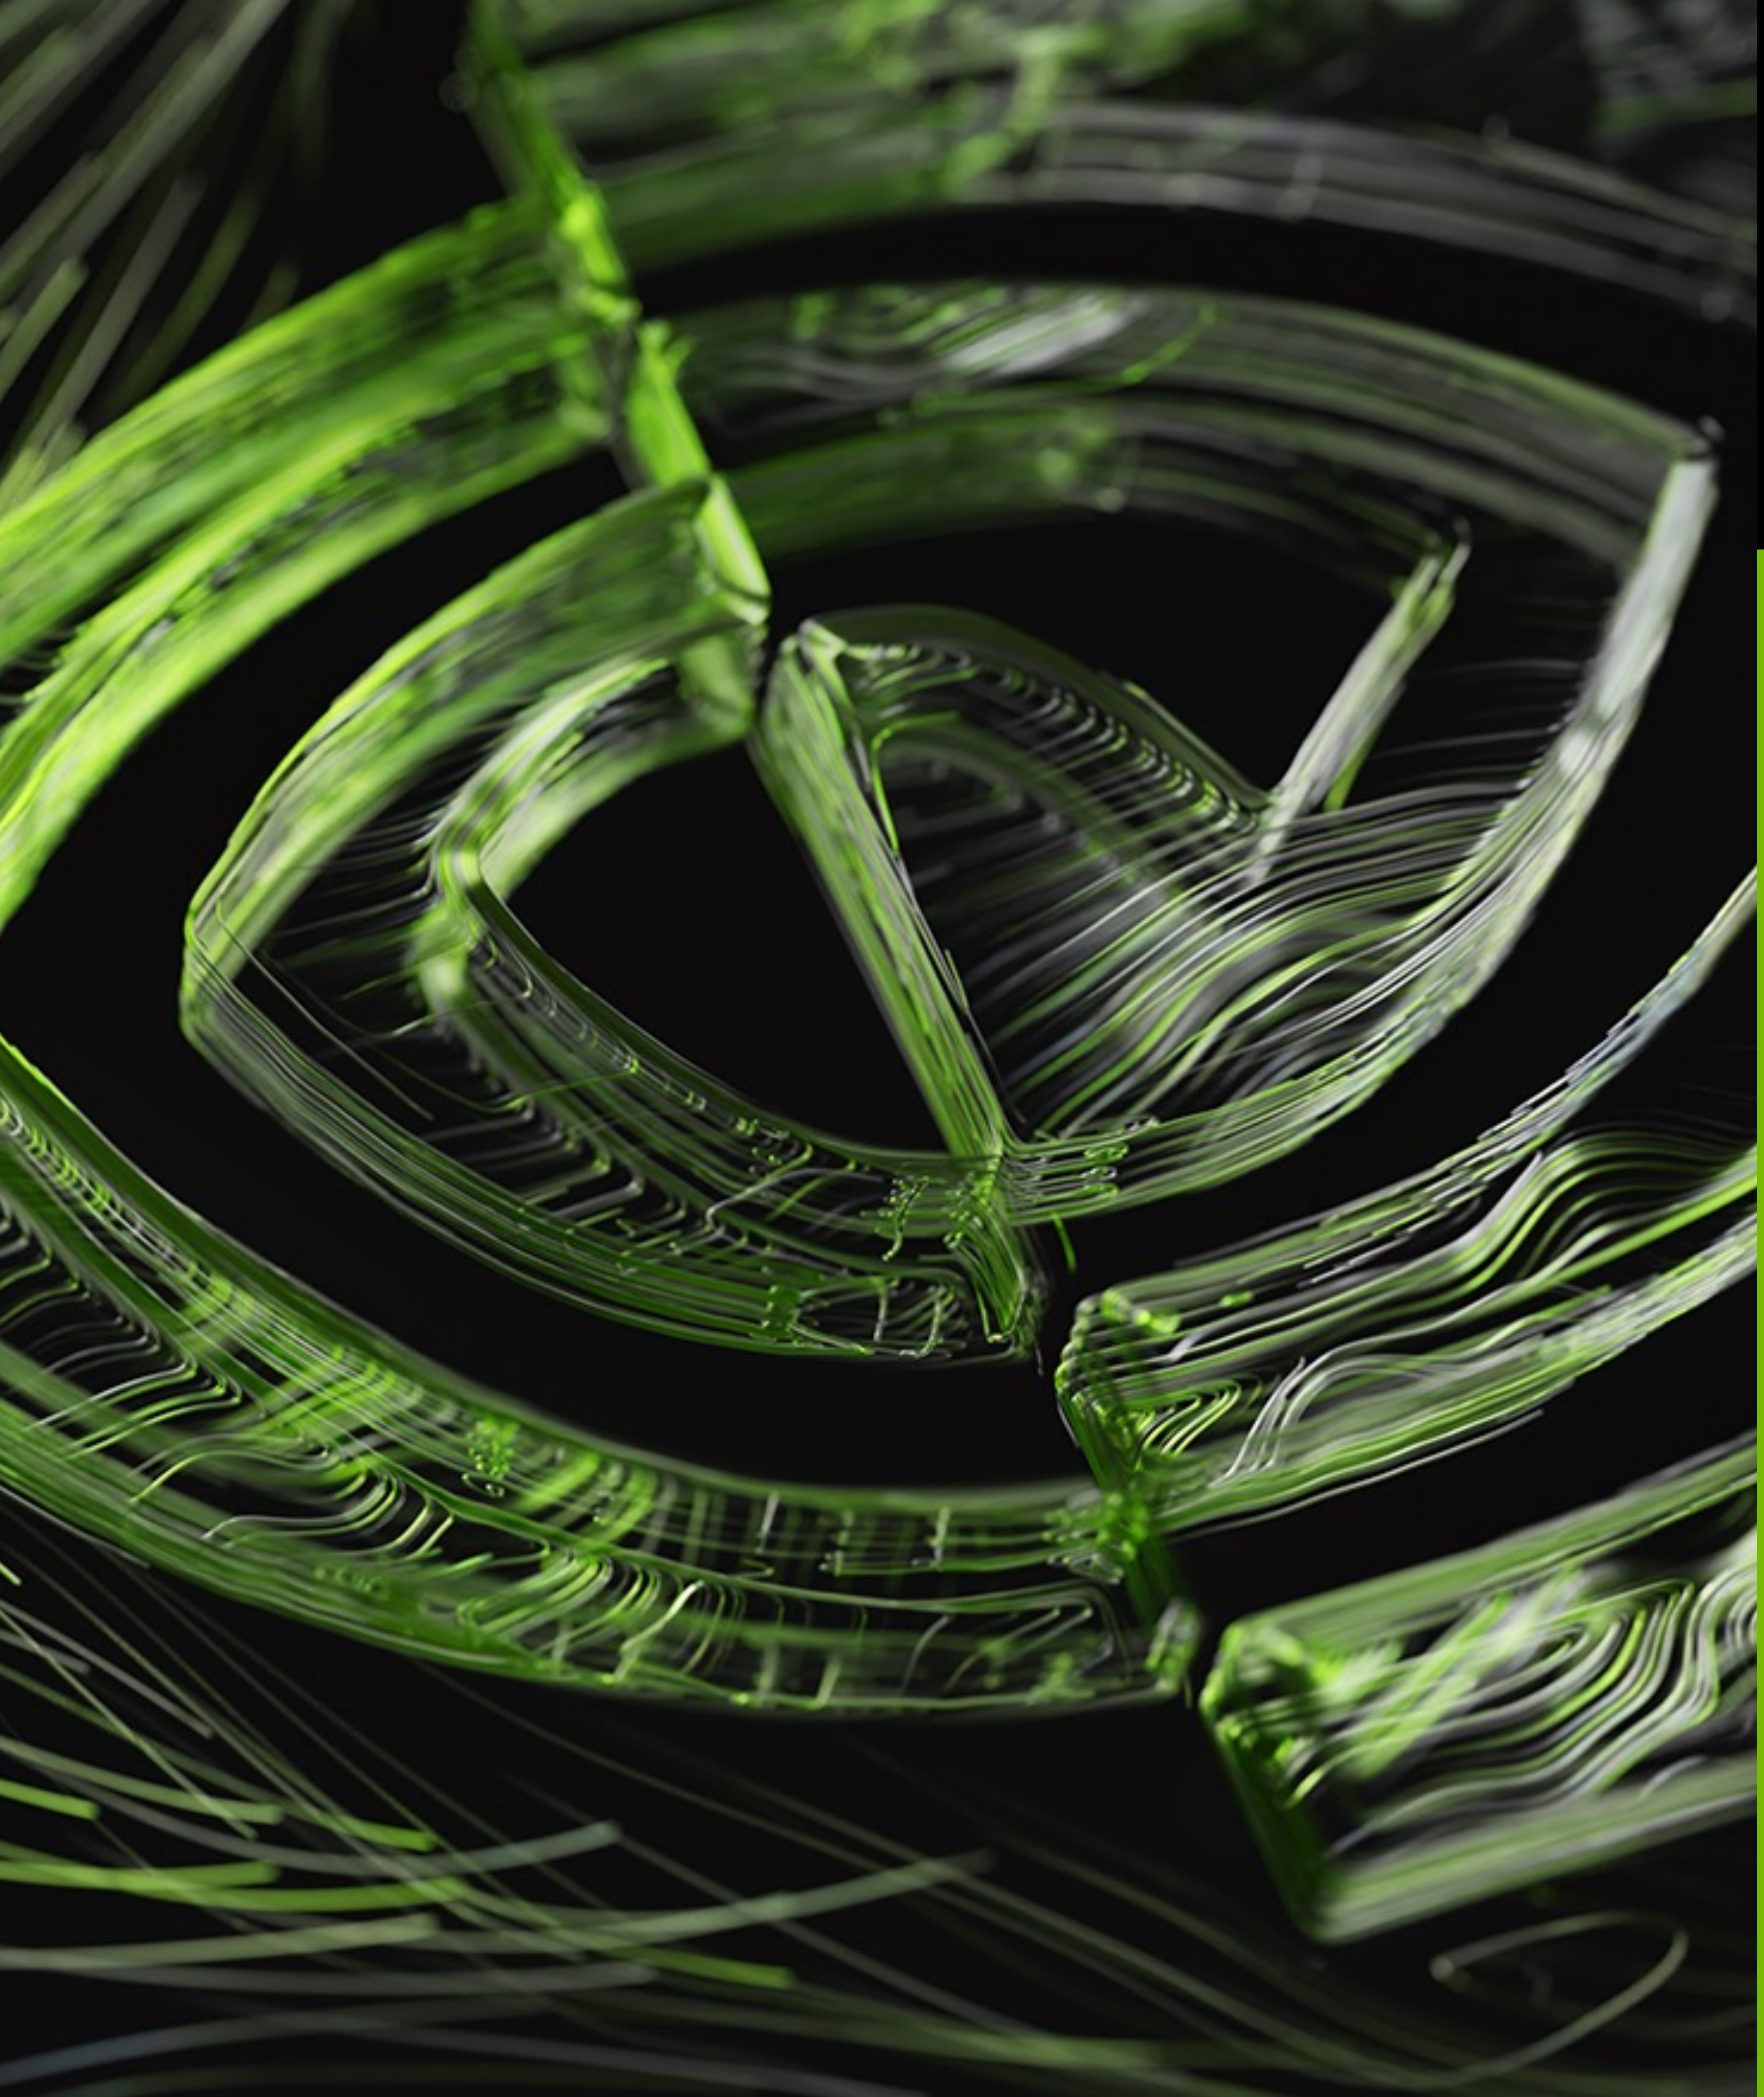

Agenda

- NVIDIA's Earth-2 initiative: The Big Picture
- FourCastNet & latest accomplishments in Weather.
- Beyond Weather, towards Climate.
- Digital Twinning platforms

Nimbus will enable end-to-end MLOps

Cloud-native end-to-end MLOps pipeline to push frontiers of AI-driven weather and climate research

- End-to-end MLOPs pipeline for data ingest, processing, training, inference, deployment
- NVIDIA Super-Cloud hosted
- API for easy access to popular data sources: models, reanalyses, and observations
- Leveraging Modulus Physics-ML framework and NVIDIA AI and performance optimization tools
- Fast inference and deployment
- Pre-trained models and transfer learning
- Recipes for model development and fine-tuning for regional prediction, specific phenomena (cyclones, heat waves, etc.)

Omniverse will enable scientists to create digital twins together

Nucleus: A shared space where models, data, tools, services, and applications synchronize

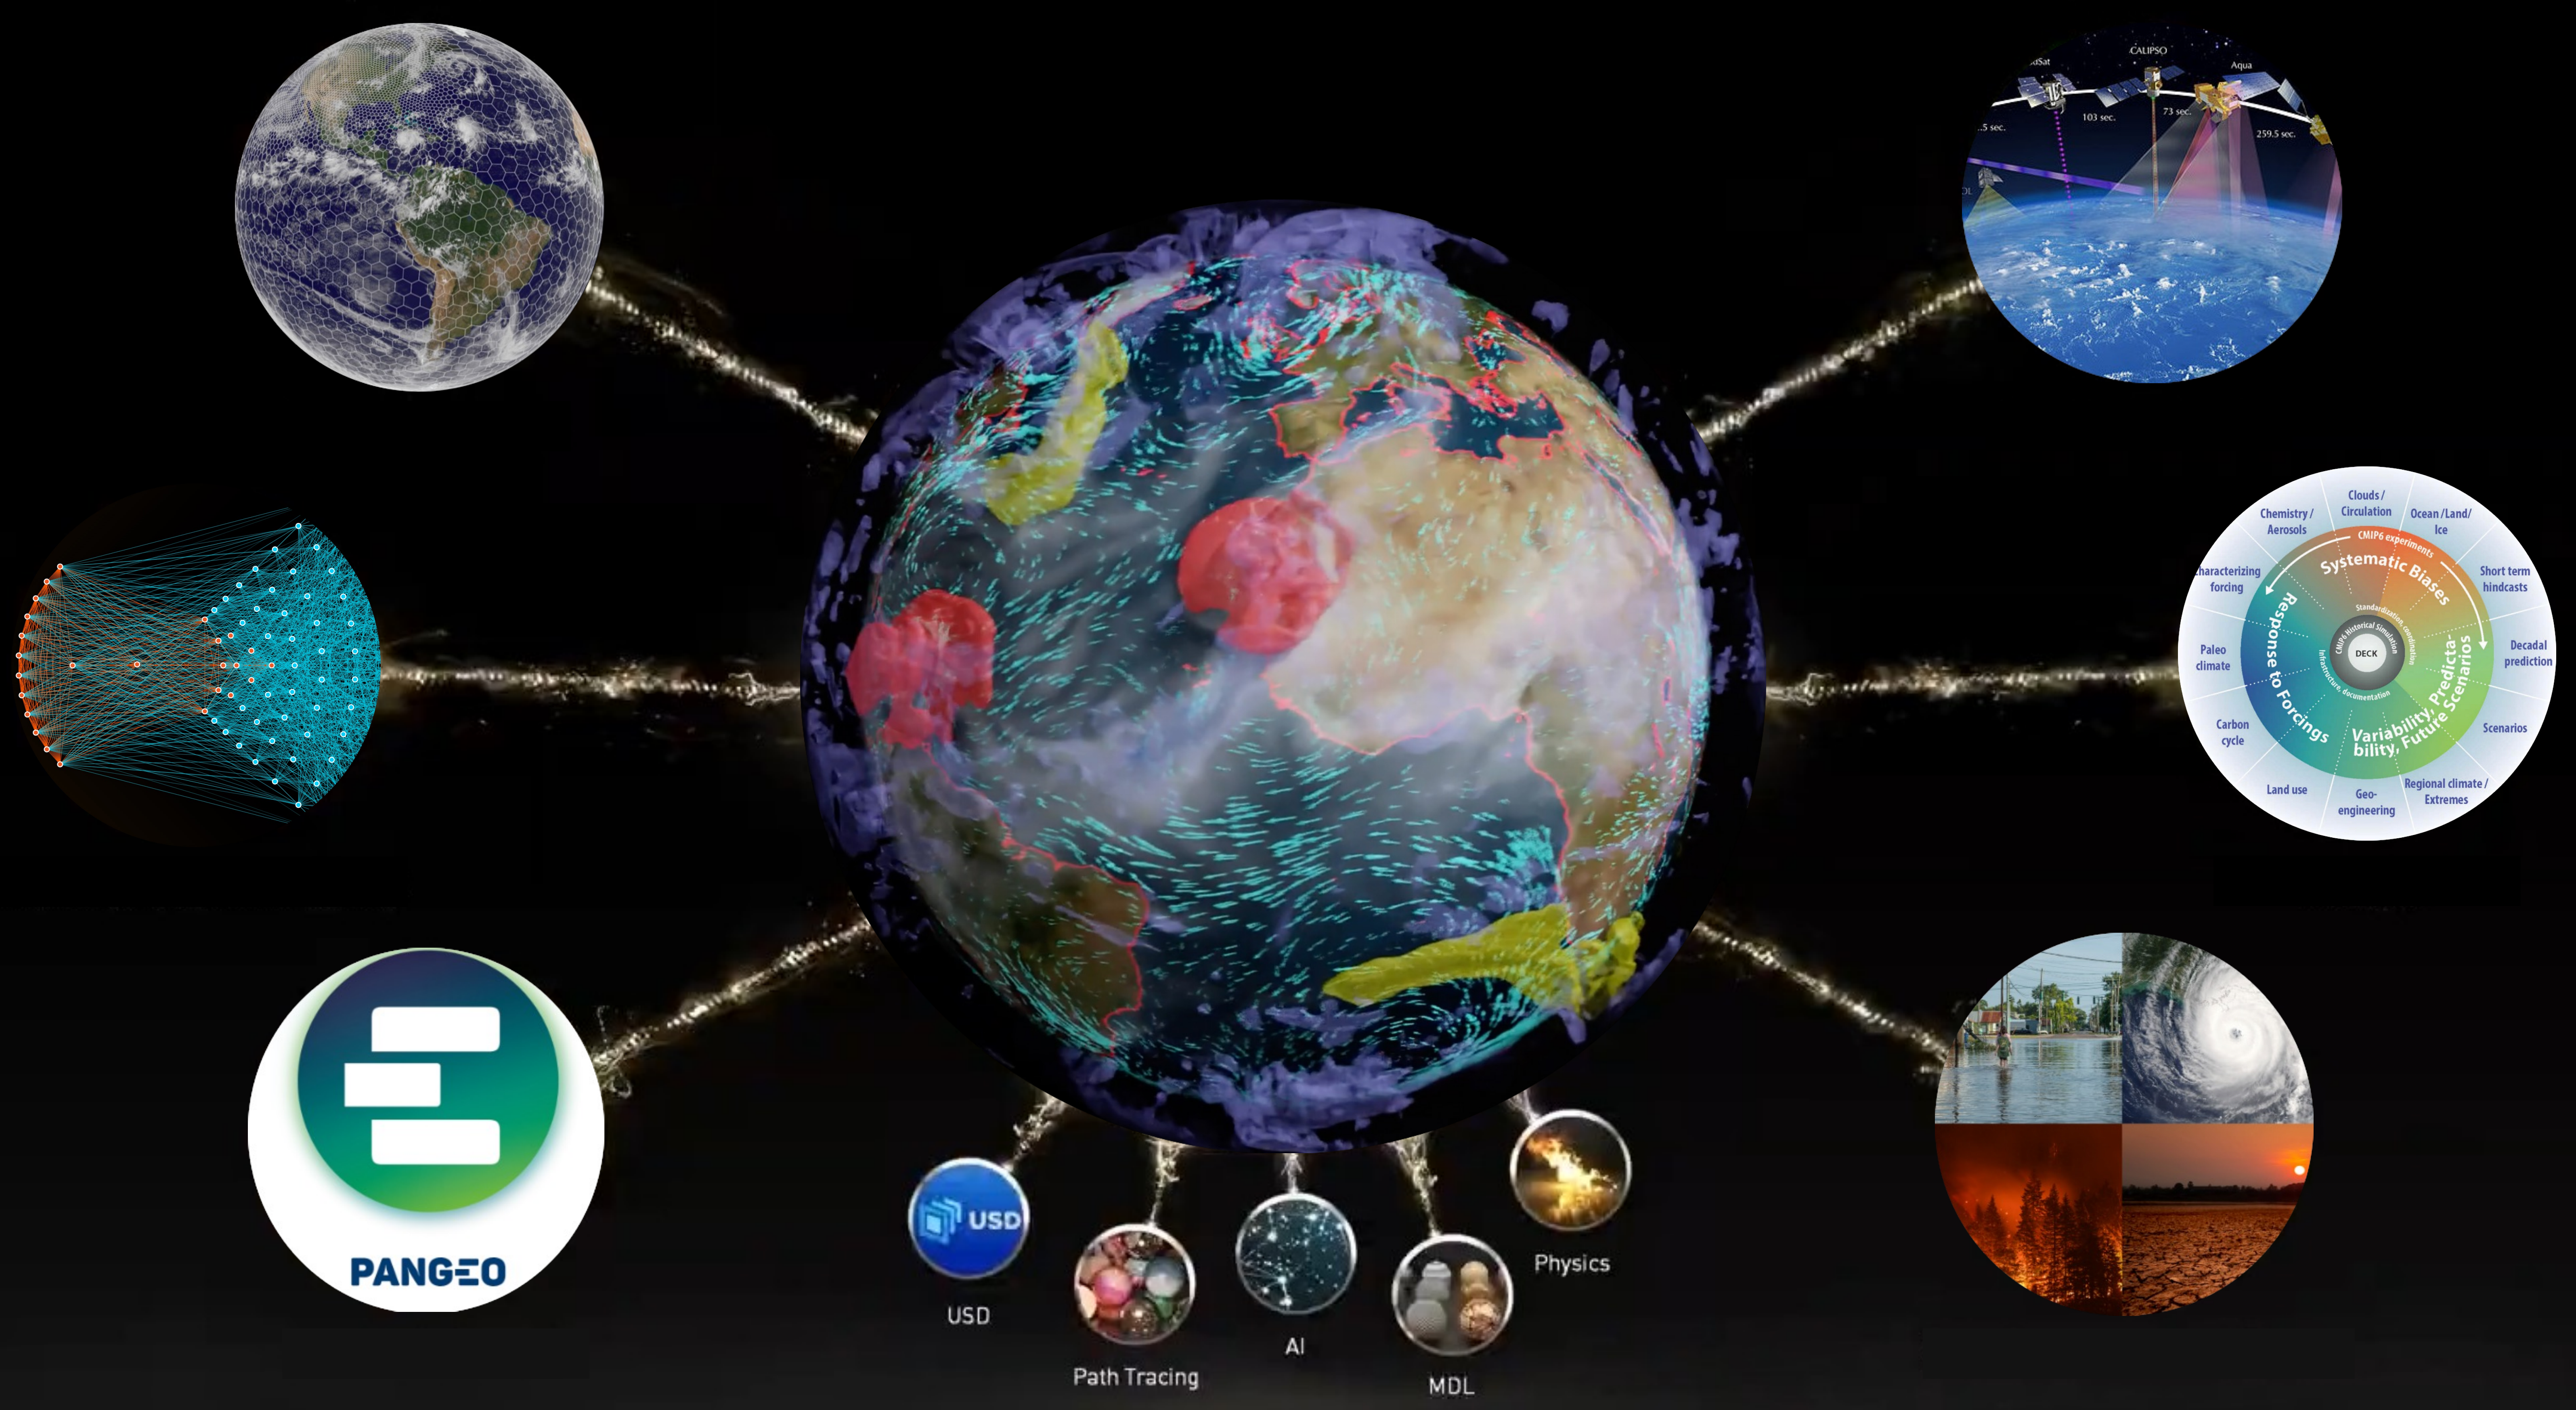

The Vision of Earth-2

Is Beginning to Take Shape

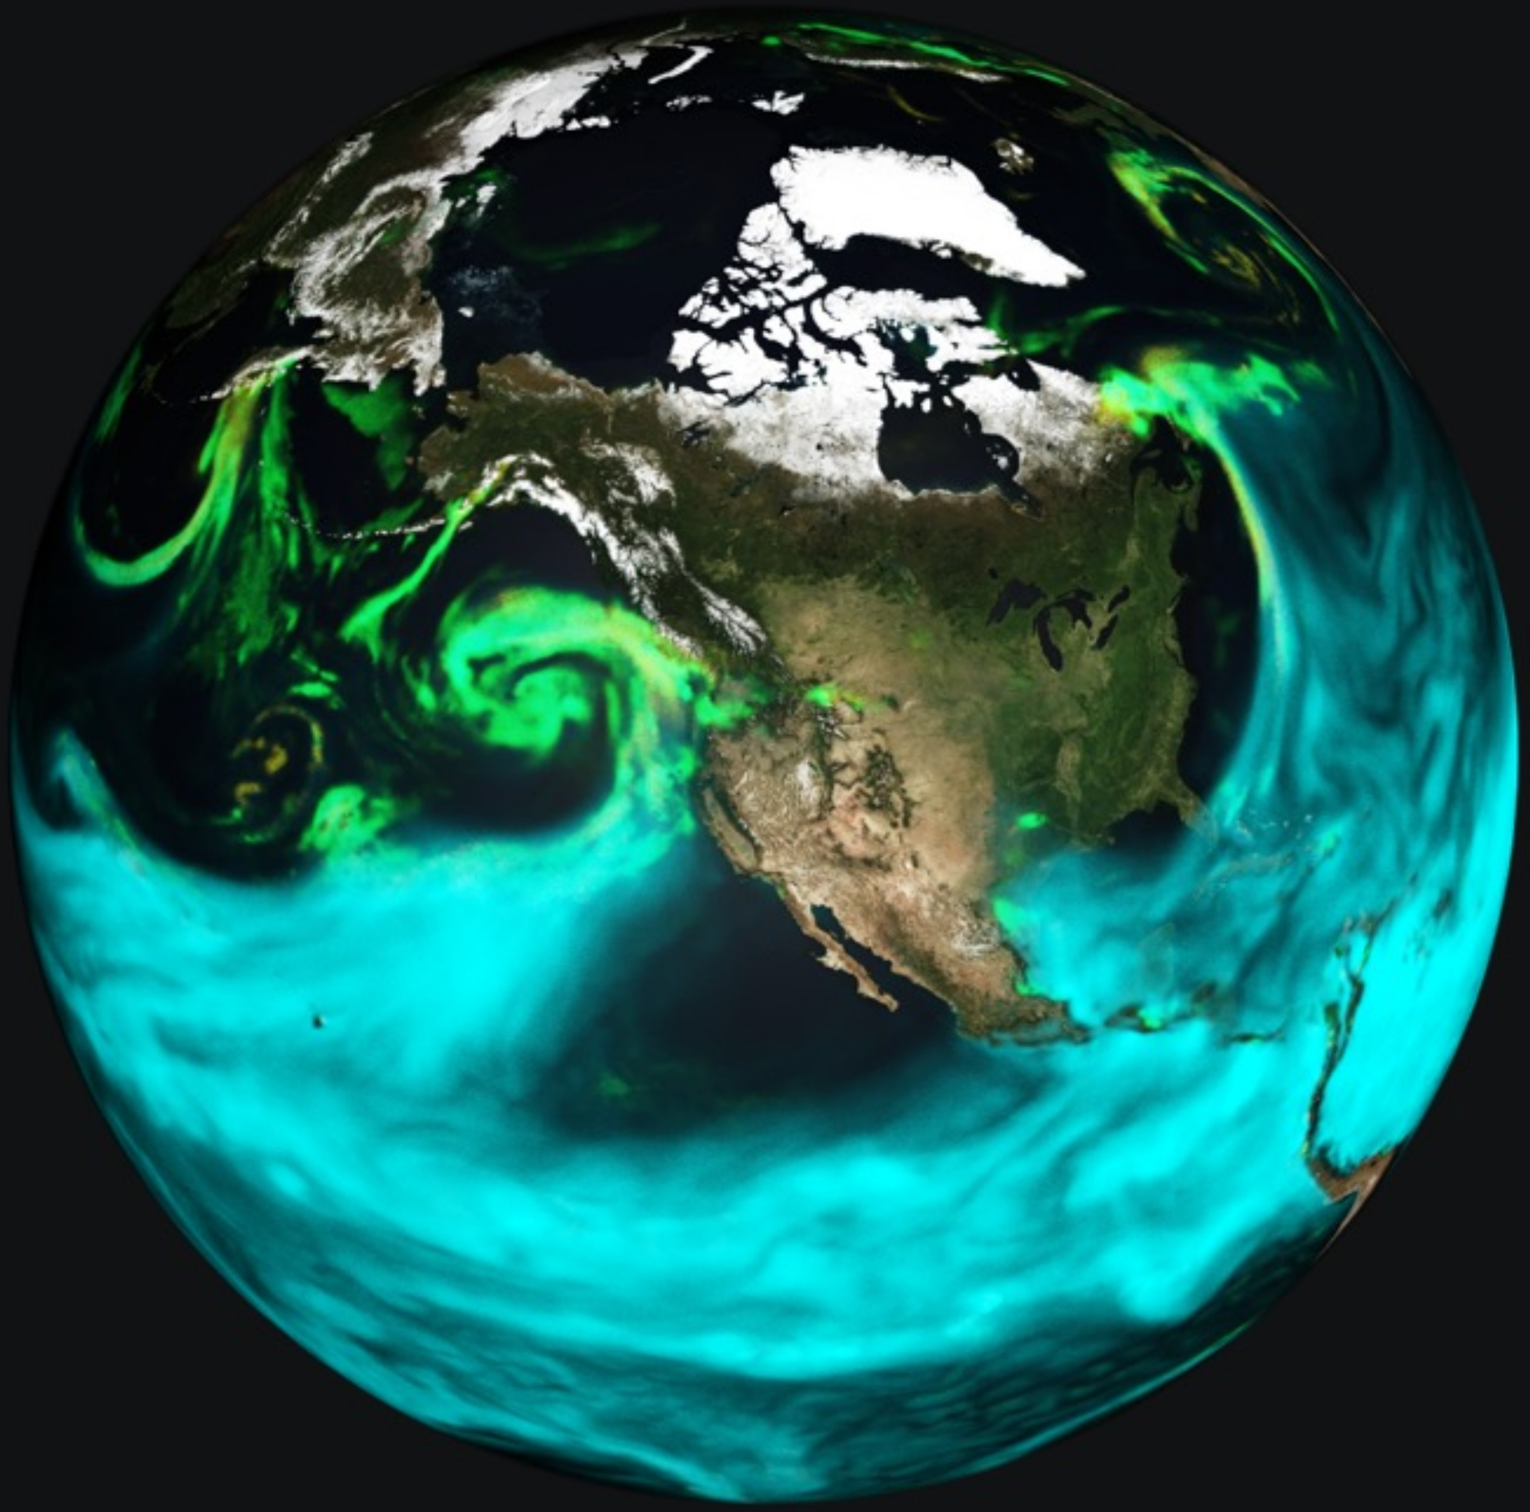

Acknowledging: Mike Pritchard, Anima Anandkumar, David Hall, Jaideep Pathak, Noah Brenowitz, Yair Cohen, Thorsten Kurth, Boris Bonev, Christian Hundt, Andre Graubner, Peter Messmer, Stan Posey, Akshay Subramaniam, Sanjay Choudhry, Farah Hariri, Niklas Roebler, Ram Cherukuri, Nicholas Geneva, Mathias Hummel, Christopher Lamb, Mike Houston, Kamyar Azizzadenesheli, Jean Kossaifi, Steffen Roemer, Marius Koch & David Appelhans, many more NV staff & our **generous external climate science advisors Bjorn Stevens, Peter Deuben, Peter Bauer, Nils Wedi, Thomas Schulthess, and Francisco Doblas-Reyes**.

NVIDIA