NVIDIA HPC Technologies Enabling the Earth-2 Digital Twin Initiative

Peter Messmer, David Hall, Karthik Kashinath, Mike Pritchard, Stan Posey

NVIDIA Earth-2 Announcement - GTC Nov 2021

NVIDIA to Build Earth-2 Supercomputer to See Our Future

November 12, 2021 by JENSEN HUANG https://blogs.nvidia.com/blog/2021/11/12/earth-2-supercomputer/

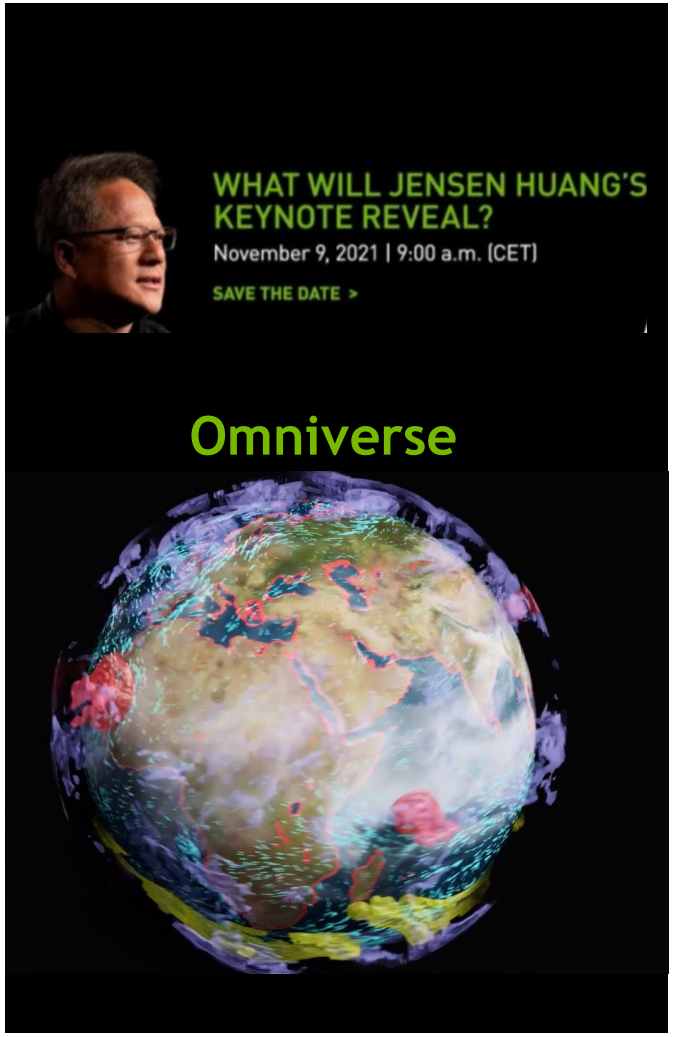

- "NVIDIA plans to build the world's most powerful AI supercomputer dedicated to predicting climate change. Named Earth-2, or E-2, the system would create a digital twin of Earth in Omniverse."
- "We will dedicate ourselves and our significant resources to direct NVIDIA's scale and expertise in computational sciences, to join with the world's climate science community."
- "All the technologies we've invented up to this moment are needed to make Earth-2 possible. I can't imagine a greater or more important use."

-- Jensen Huang, CEO NVIDIA

NVIDIA Earth-2 Announcement - GTC Nov 2021

NVIDIA to Build Earth-2 Supercomputer to See Our Future

November 12, 2021 by JENSEN HUANG https://blogs.nvidia.com/blog/2021/11/12/earth-2-supercomputer/

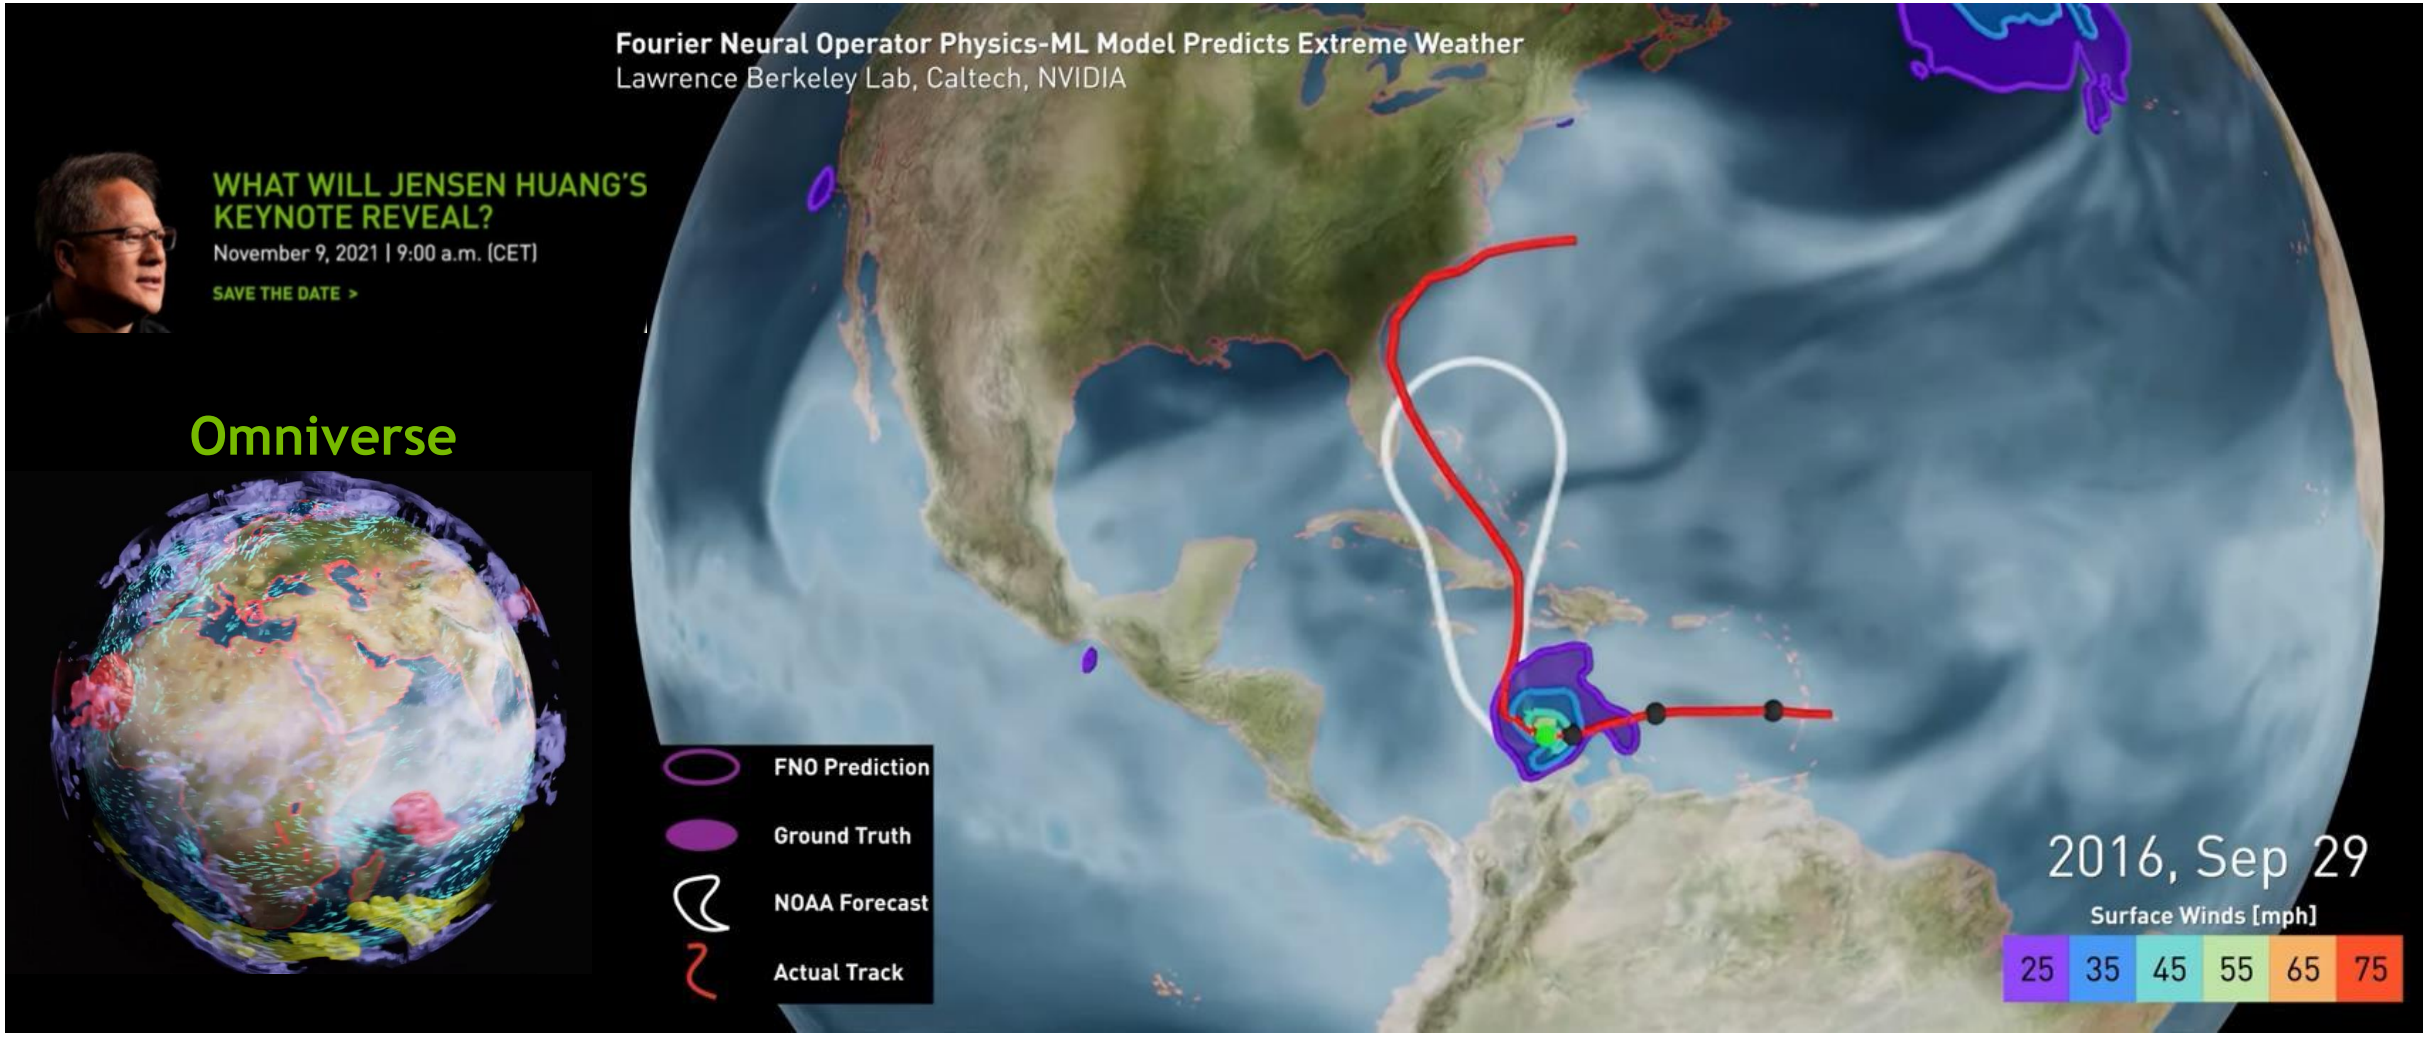

https://www.youtube.com/watch?v=JnGPxZ9glVk

NVIDIA Earth-2 Early Inspiration and Vision

- **SC20 Keynote Lecture:**
  - From Deep Thoughts to (Destination) Earth: Climate Science in the Age of Exascale
    - *--Prof. Dr. Bjorn Stevens of the Max-Planck-Institute for Meteorology*

- **EC's Destination Earth:**
  - Destination Earth and digital twins - a European opportunity for HPC
    - *--Dr. Peter Bauer, ECMWF, Director of Destination Earth*

ECMWF

- **CliMA Lecture:**
  - AI-Accelerated Climate Modeling **(16 Jul 21)**
    - *--Dr. Tapio Schneider, CalTech and NASA JPL*

- **AI2 Collaboration:**
  - AI2 Climate Modeling
    - *--Dr. Chris Bretherton, et al.*

AI2 Allen Institute for AI

NVIDIA Earth-2 is Collaborative and Open Science

- **NVIDIA Key Contributors:**
  - *Mike Pritchard, Director of Climate Science Research (also UC Irvine)*
  - *Anima Anandkumar, Sr. Director of ML Research (also CalTech)*
  - *Peter Messmer, Director of Developer Technology*
  - *Karthik Kashinath, Principal Engineer and Scientist, AI-HPC*
  - *David Hall, Solutions Architect, AI Developments*
  - *Several others . . .*
- **Earth-2 Domain Science Advisors:**
  - ***Dr. Peter Bauer, ECMWF, Director of Destination Earth***
  - ***Prof. Dr. Bjorn Stevens, MPI for Meteorology, Managing Director***
  - ***Dr. Peter Dueben, ECMWF, Chief Scientist and AI Coordinator***
  - ***Dr. Nils Wedi, ECMWF, Destination Earth Technical Lead***
  - ***Dr. F. (Paco) Doblas-Reyes, BSC, Director of Earth System Science***
- **Climate Community HPC Guidance:**
  - ***Dr. Thomas Schulthess, CSCS, SC Center Director***
  - ***Dr. Tsengdar Lee, NASA HQ, Program Director***
  - ***Dr. Thomas Hauser, NCAR, NCAR Lab Director***
  - *Several other reviews and discussions ongoing . . .*

ECMWF

Max-Planck-Institut für Meteorologie

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

CSCS

NASA

NCAR

Recommended Earth-2 Talks from NVIDIA Scientists

Motivation (I): Limitations of Climate Predictions and HPC

Current climate models are low resolution and decades away from what's needed

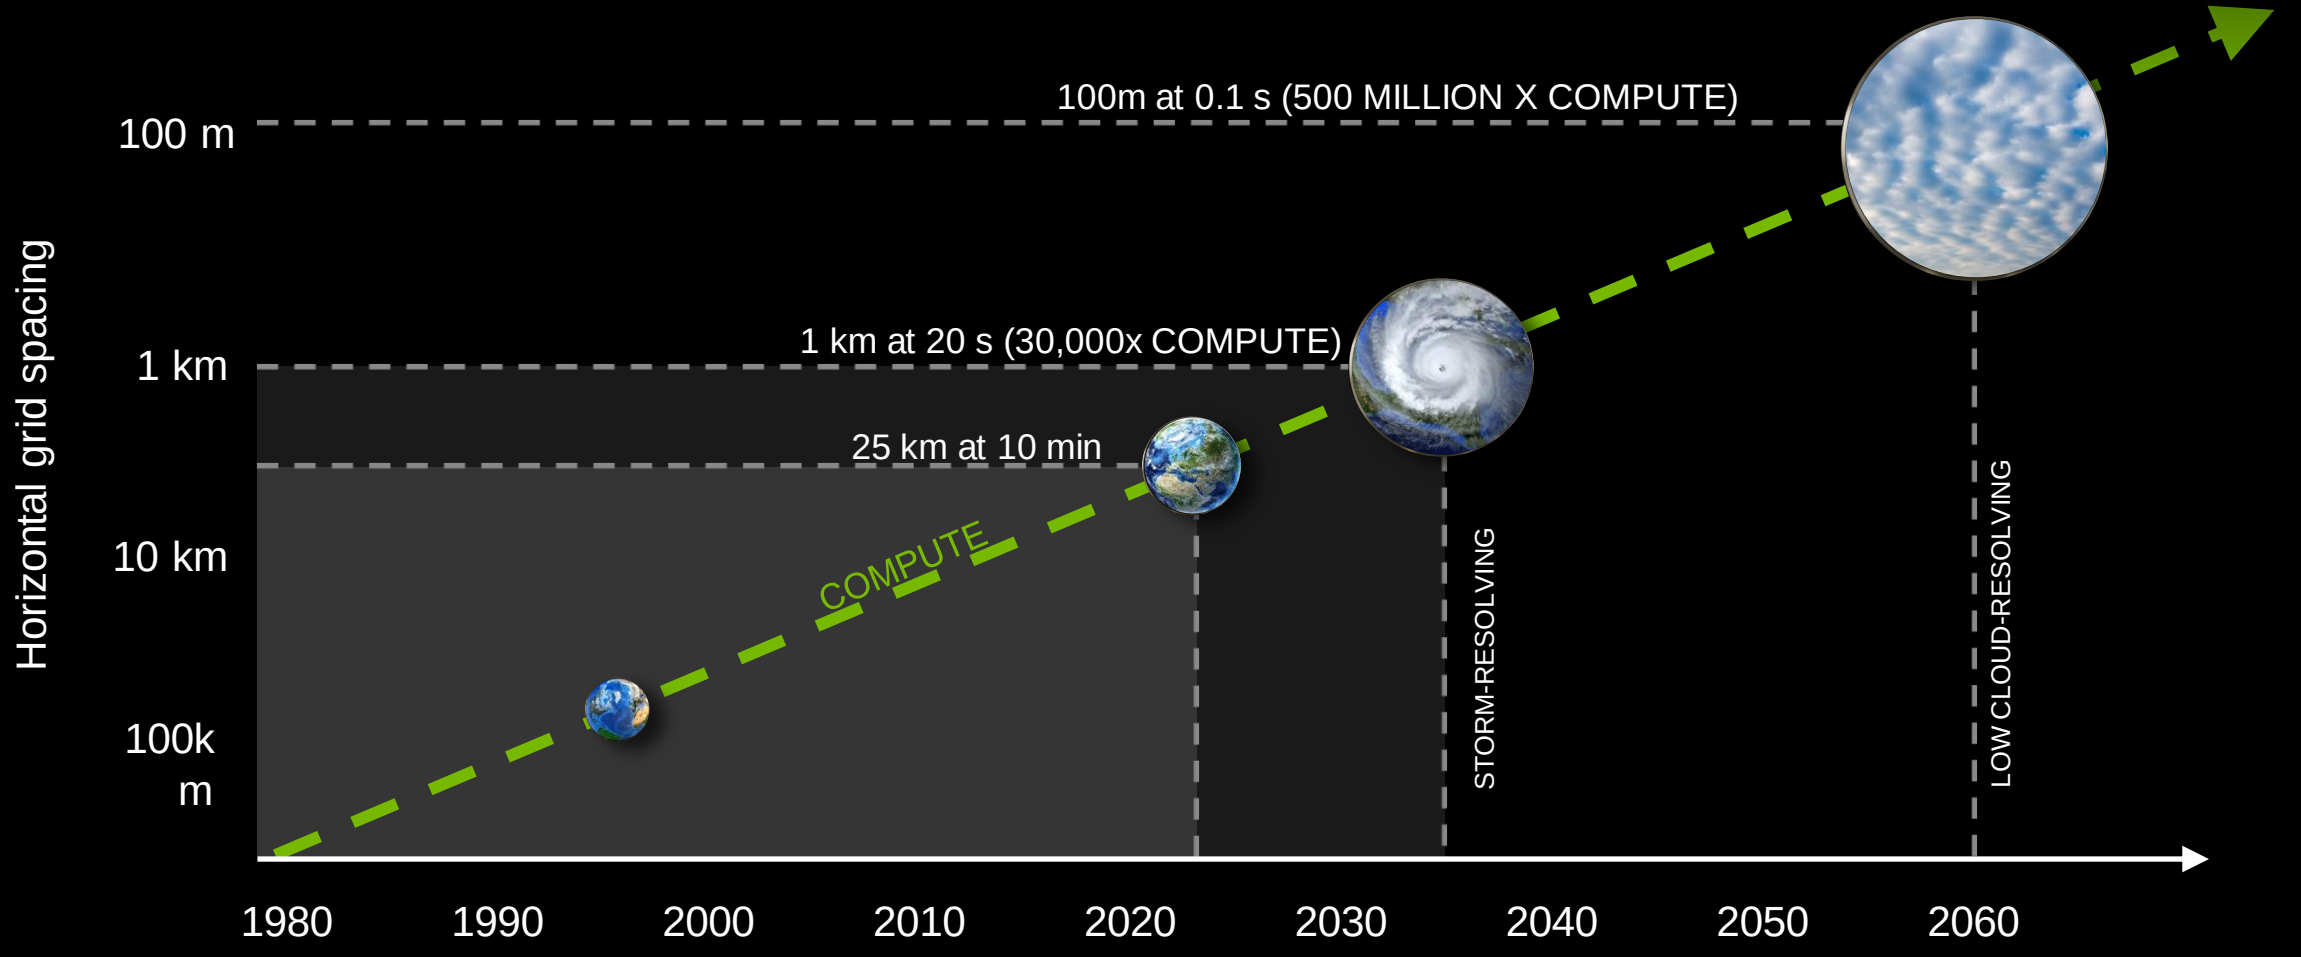

Figure adapted from: Schneider, T., Teixeira, J., Bretherton, C. *et al.* "*Climate goals and computing the future of clouds*". *Nature Climate Change* **7,** 3–5 (2017)

Motivation (II): Ever Increasing Computational Demand

Climate model resolution is only one axis driving computational demand

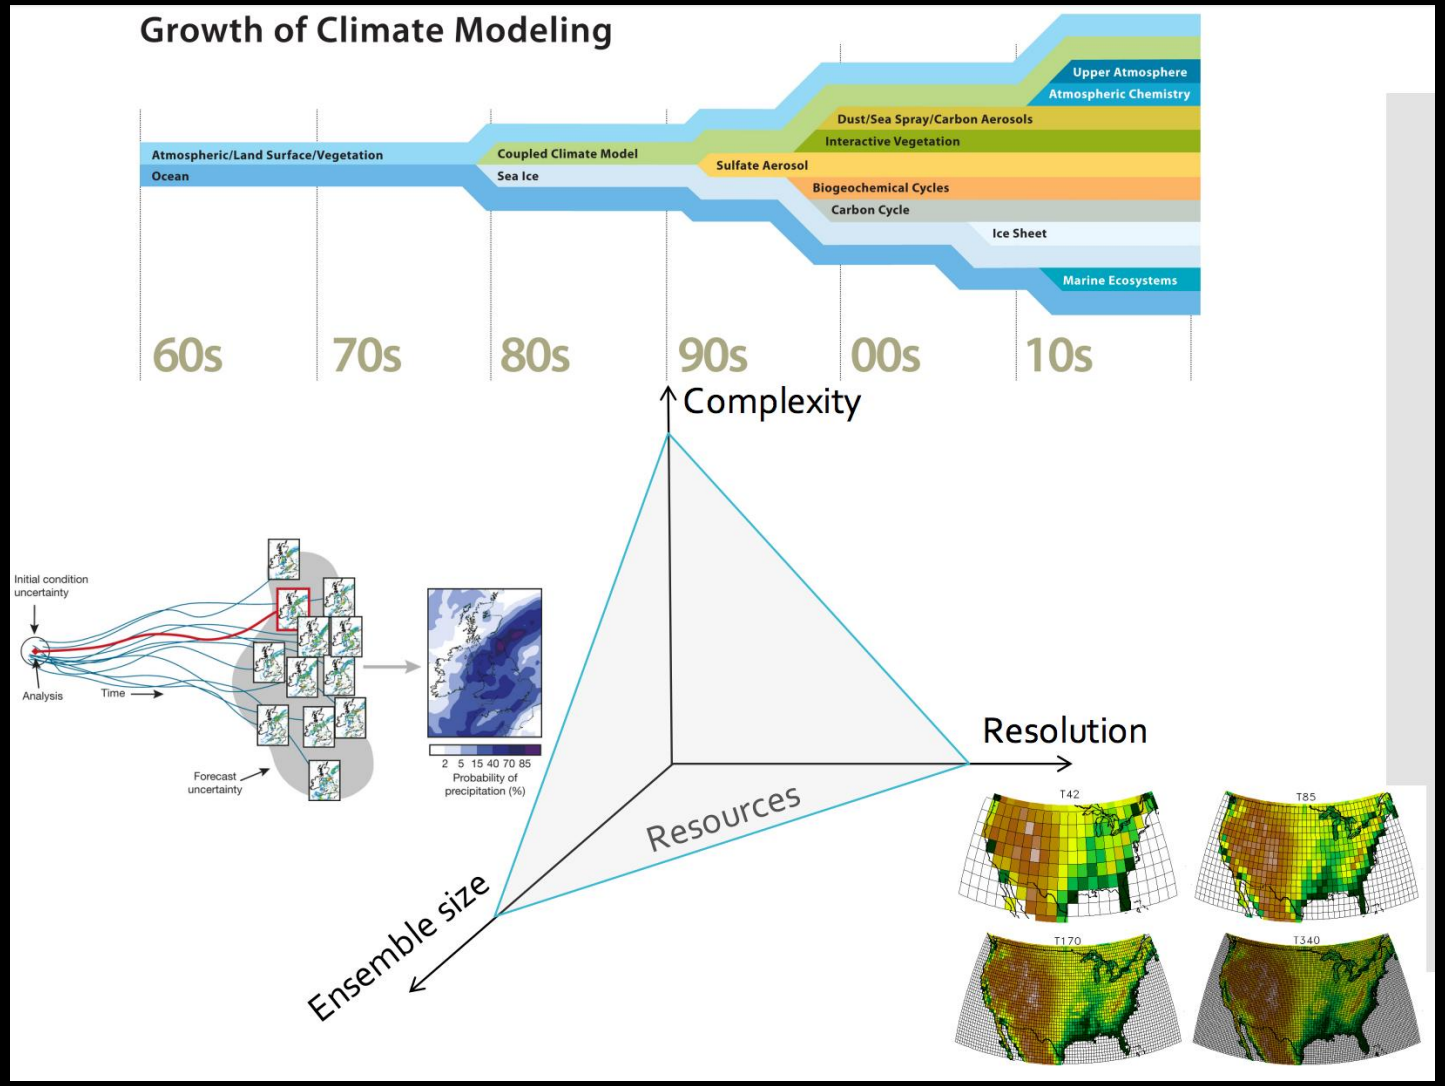

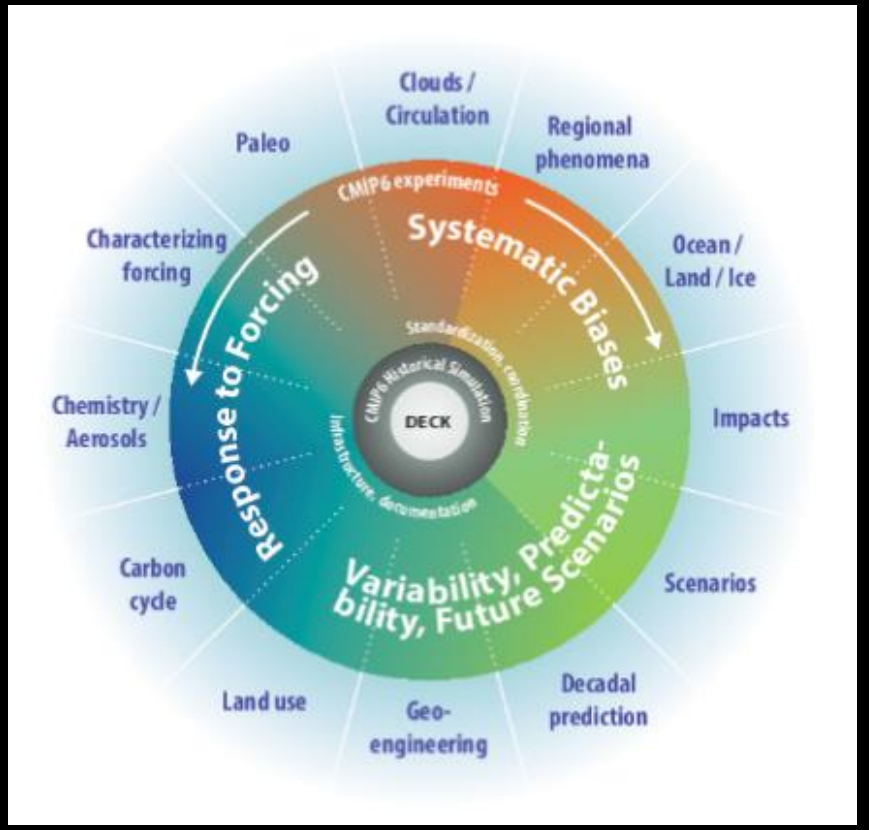

Motivation (III): Need for Interactivity with Climate Prediction Data

Coupling fast and accurate models can enable end-to-end regional analyses and climate action

Climate Twin

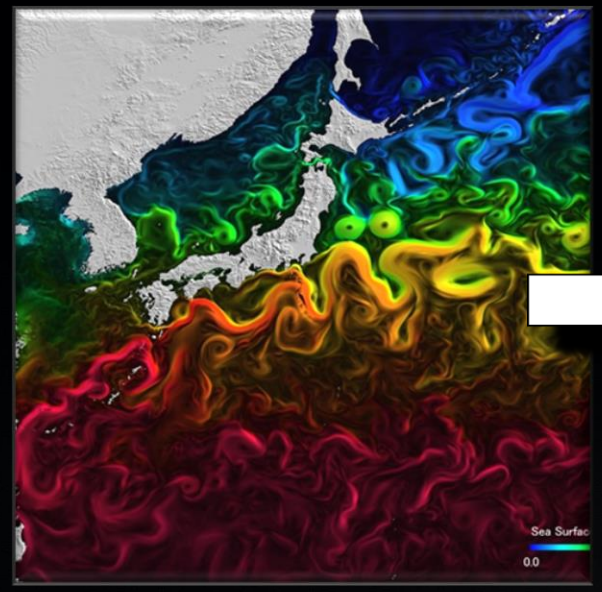

http://www.jamstec.go.jp/gallery/e/simulation/weather/002.html

Weather Twin

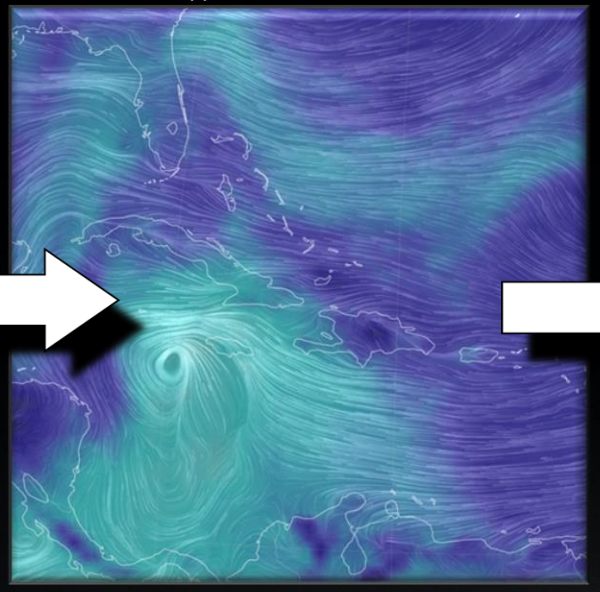

https://earth.nullschool.net/#current/wind/surface/level/overlay=total_precipitable_water/orthographic=-64.71,19.12,2322

Powerplant Twin

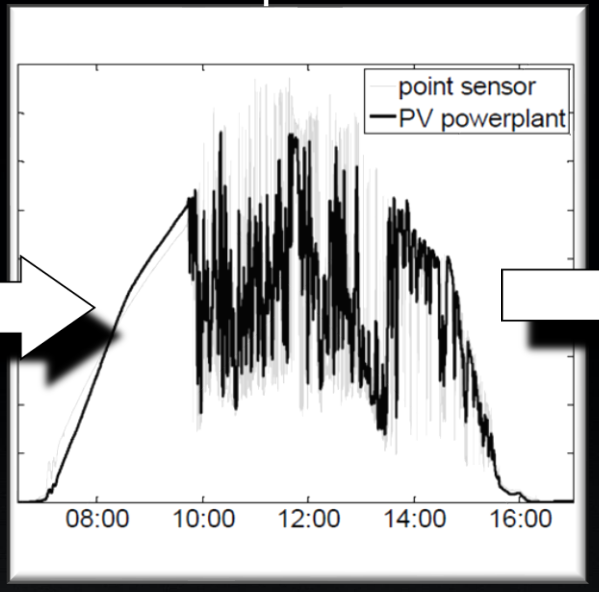

https://www.esig.energy/wiki-main-page/simulating-solar-power-plant-variability-a-review-of-current-methods/

Electric Grid Model

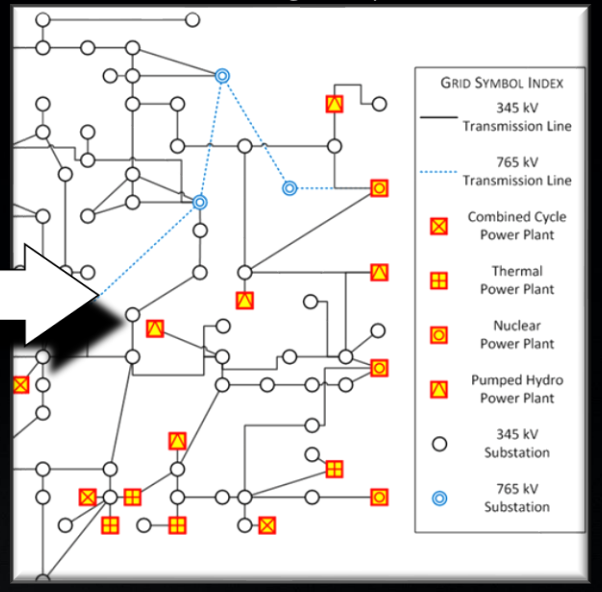

https://www.researchgate.net/figure/Overall-configuration-of-electric-power-grid-modelling_fig2_311217737

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

-- Bjorn Stevens, MPI-M

The Climate Community is Responding to these Challenges

- Project DestinE envisions what Earth system modeling could be
- NVIDIA contributing to DestinE:
  - ESM development collaboration on GPUs (IFS, ICON, etc.)
  - HPC systems Leonardo (CINECA) and MareNostrum5 (BSC)

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

Earth-2 Will Contribute to Digital Twins in Different Ways

Heavy focus on AI research and latest methods, and large-scale AI training

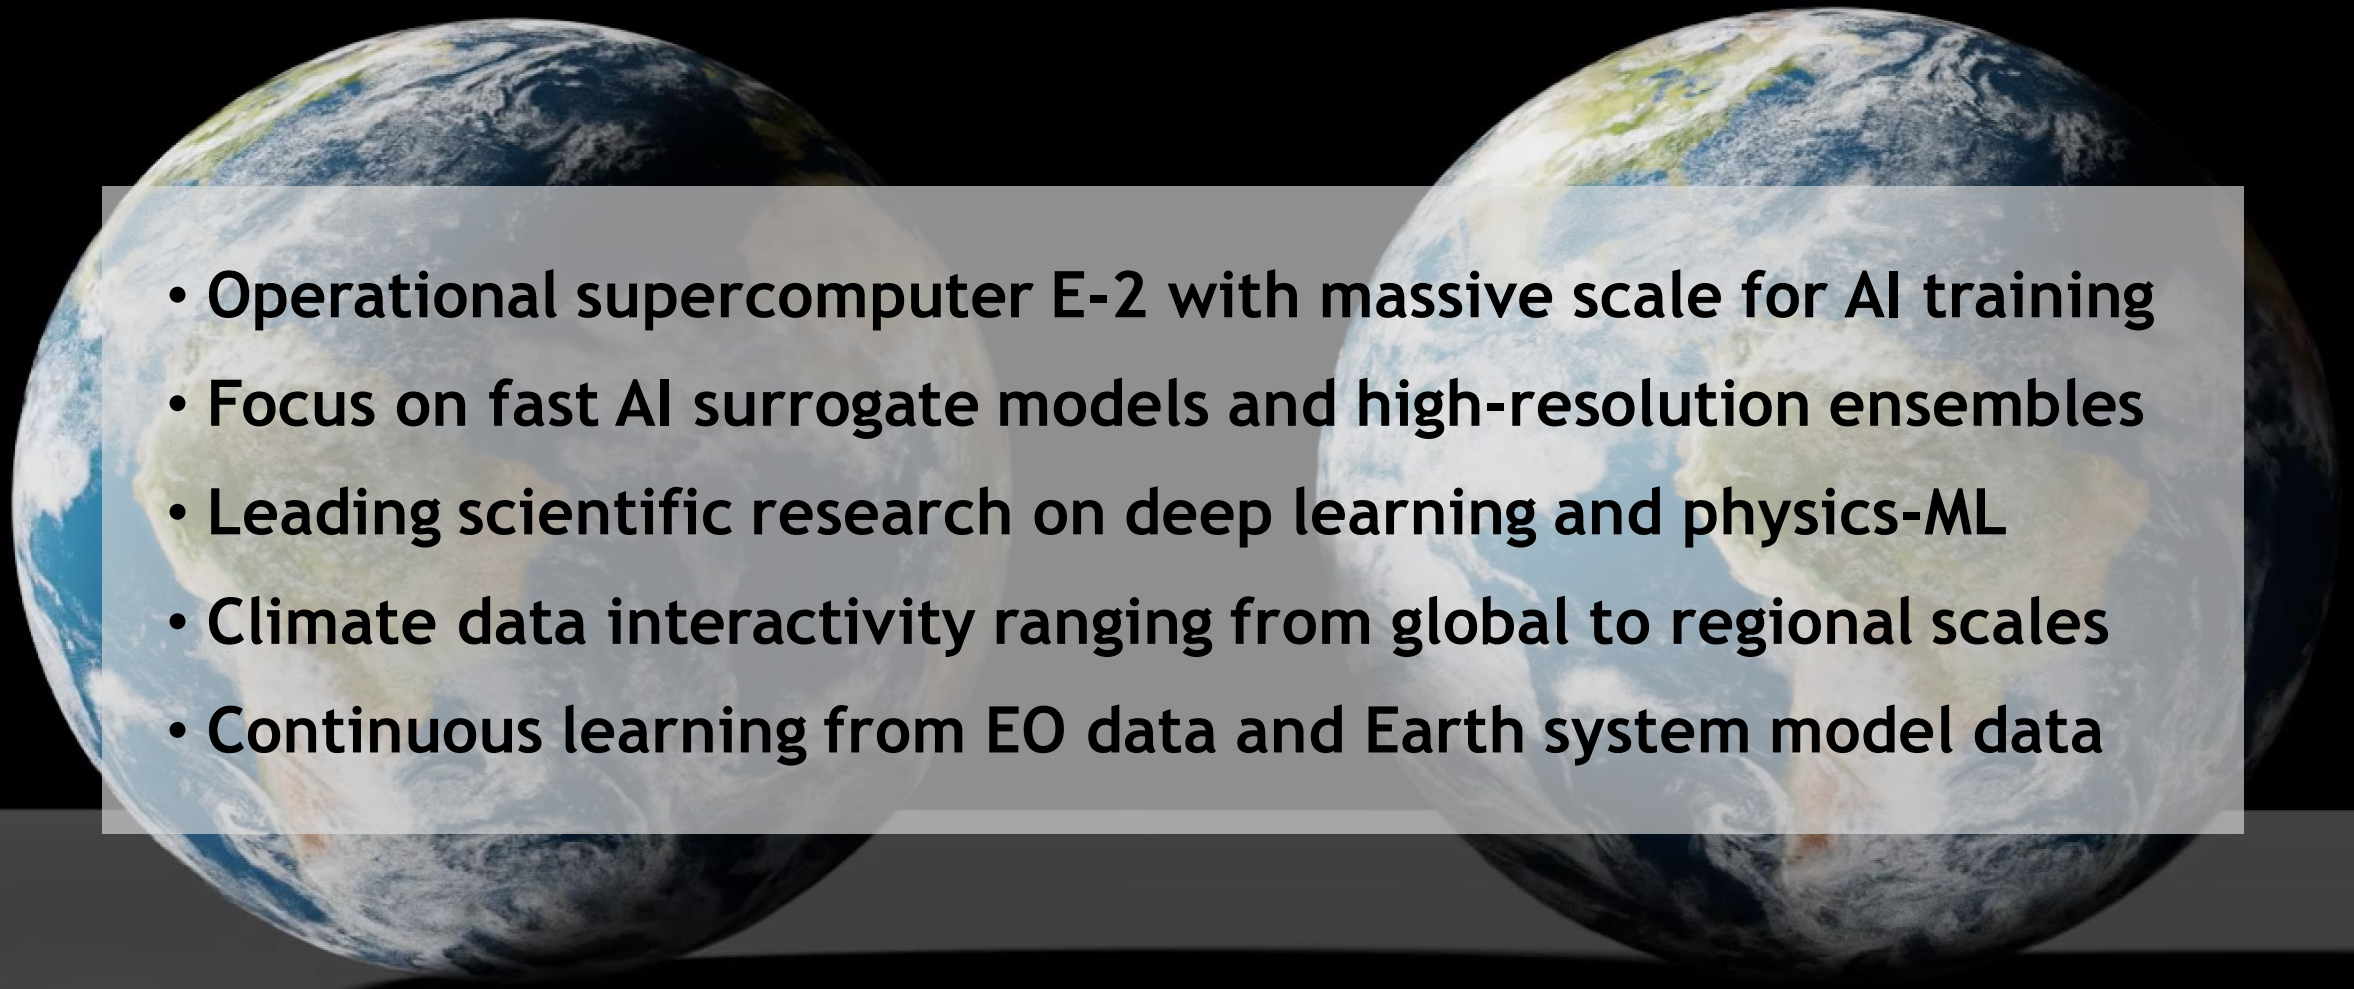

- Operational supercomputer E-2 with massive scale for AI training
- Focus on fast AI surrogate models and high-resolution ensembles
- Leading scientific research on deep learning and physics-ML
- Climate data interactivity ranging from global to regional scales
- Continuous learning from EO data and Earth system model data

Ultimately Earth-2 Will Tether to High Resolution Predictions

Cloud feedbacks and storm dynamics from km-scale simulators matter to predicting regional climate risk

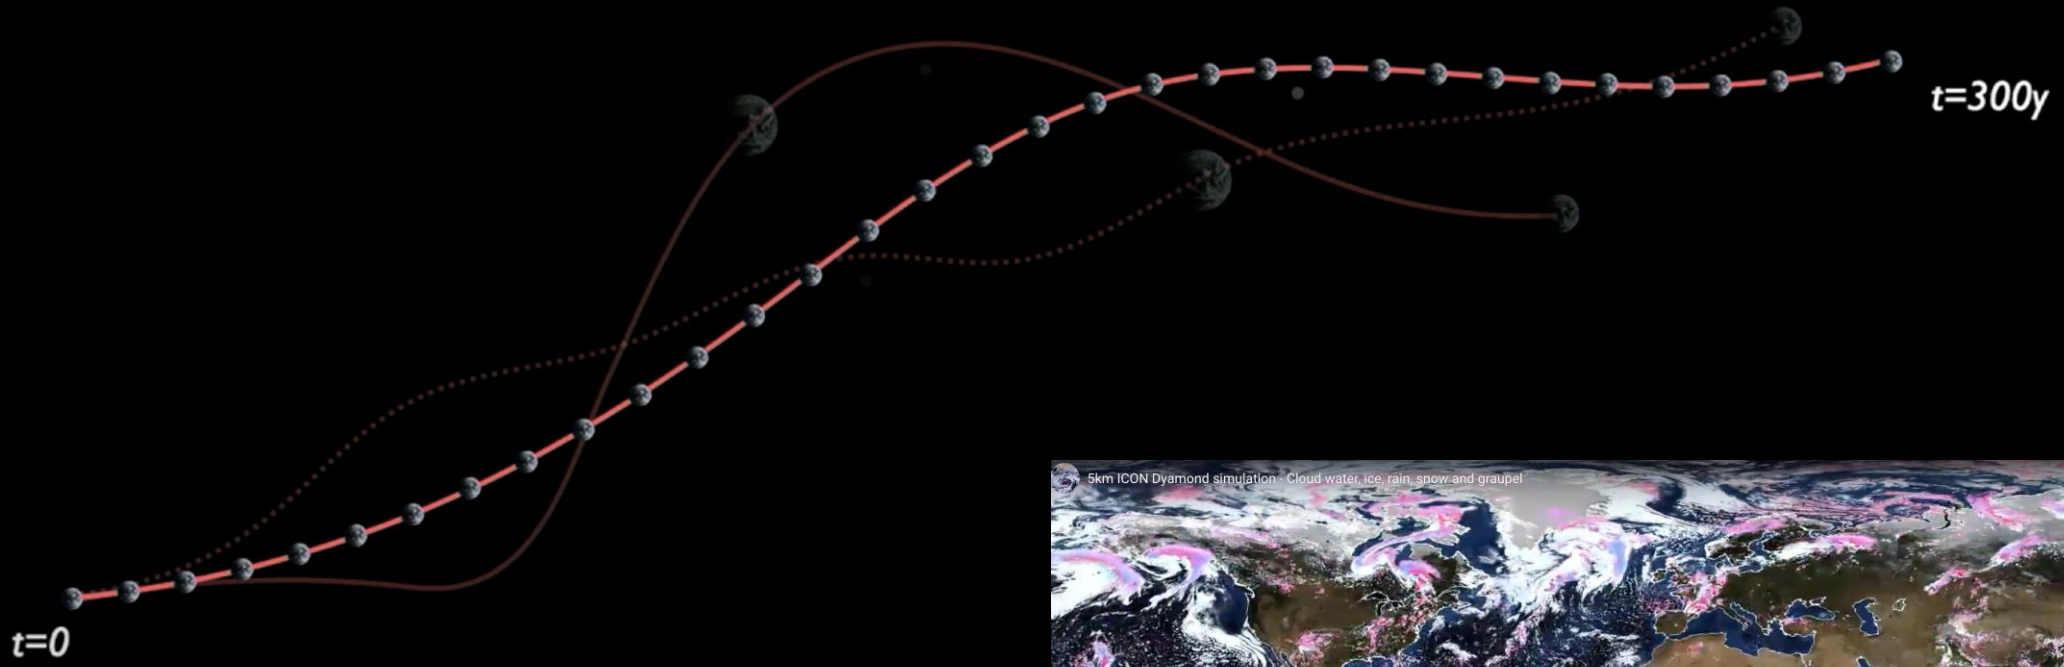

AI nimbly generates details between "checkpoints" saved only infrequently from physics-based climate simulations

-- Bjorn Stevens, GTC 2021

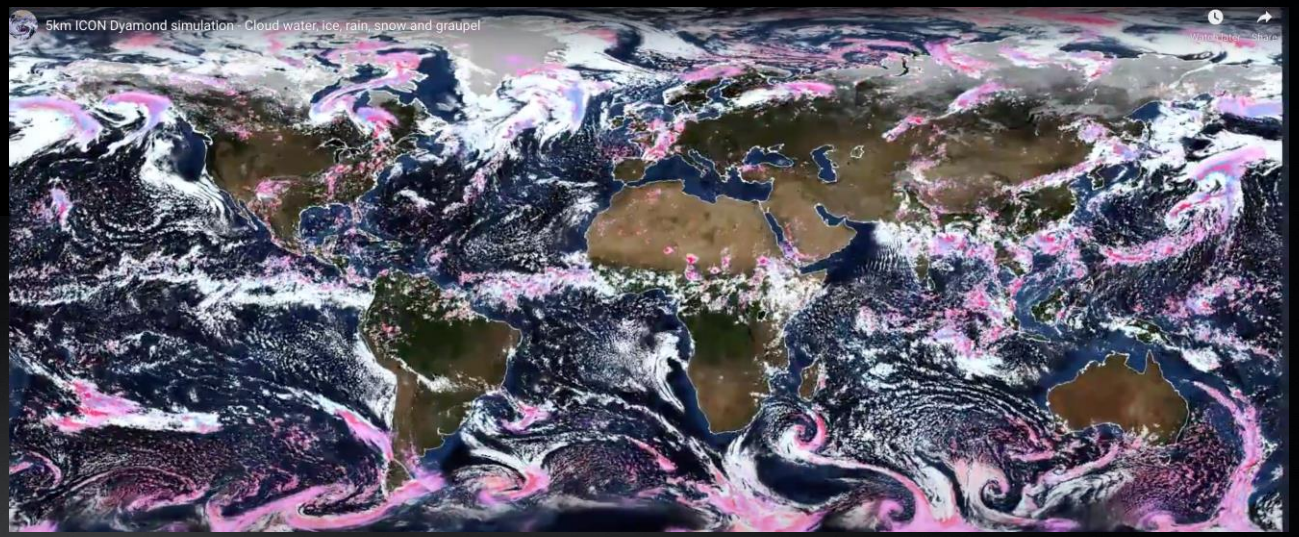

NVIDIA Collaborations With Atmospheric Models

	Model	Organizations	Funding Source
Global:			
	E3SM, MMF, HOMEXX, SCREAM	US DOE: ORNL, SNL	E3SM, ECP
	MPAS-A	NCAR, UWyo, IBM	WACA II
	FV3 Dycore	NOAA, AI2	SENA, AI2
	NUMA/NEPTUNE	US Naval Res Lab, NPS	ONR
	IFS	ECMWF	ESCAPE, US DOE
	GungHo/LFRic	MetOffice, STFC	PSyclone
	ICON	DWD, MPI-M, CSCS, MCH	PASC ENIAC
	GEOS-5	NASA GMAO	NASA
	CLIMA/NUMA	CLiMA (NASA JPL, MIT, NPS)	Private, US NSF
Regional:			
	AROME	Meteo France	MF/CNRS
	COSMO	MCH, CSCS, DWD	PASC GridTools
	AceCAST-WRF	TempoQuest	Venture backed

IFS Scaling to 11,500+ GPUs on ORNL Summit

ECMWF

Batched matrix multiplication: speedup in ESCAPE1 dwarf

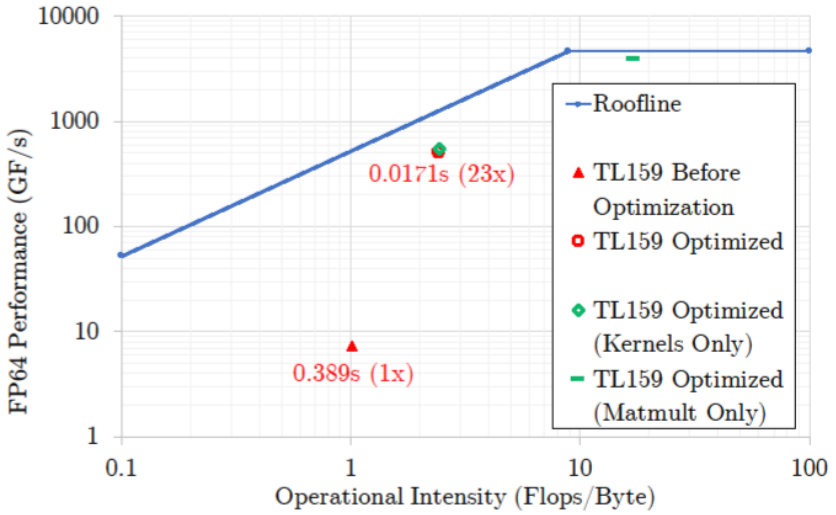

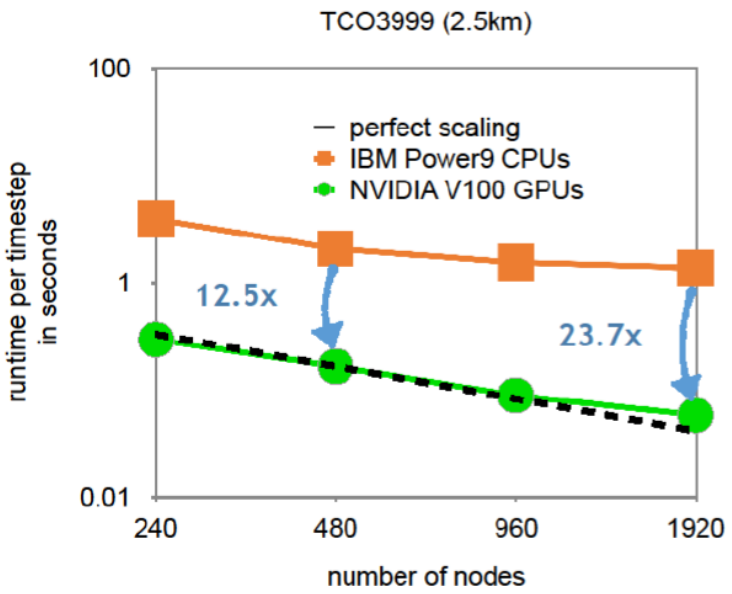

- added zero operations increase operational intensity
- overall huge speedup
- should also have strong positive effect on strong scaling

Spectral Harmonics Dwarf = 23.7x

Batched Legendre Transform (GEMM)

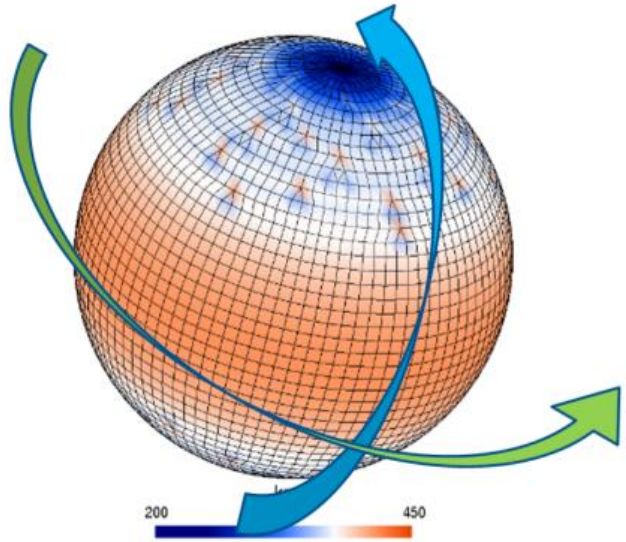

https://www.ecmwf.int/en/about/media-centre/science-blog/2019/oak-ridge-leadership-computing-facility-user-workshop-2019

IFS Spectral Transform at Global 1km on 512 GPUs

ECMWF

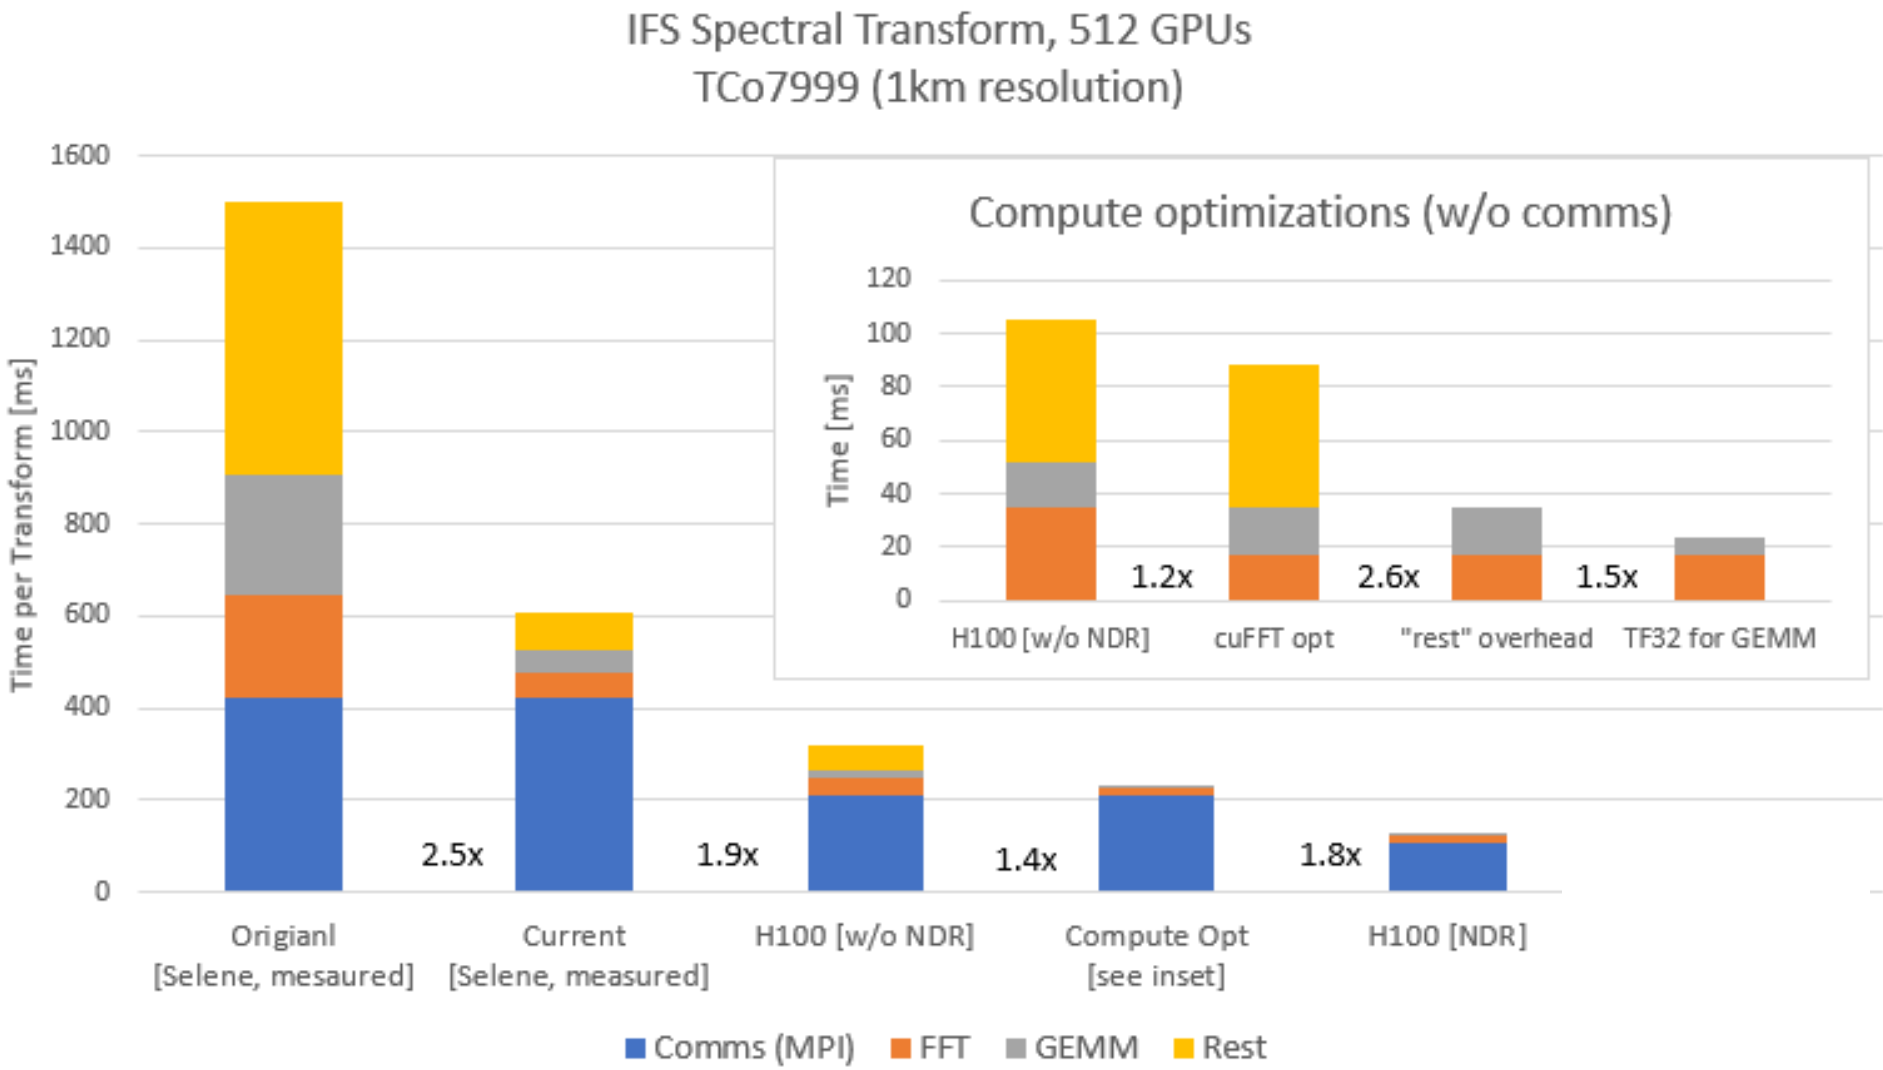

IFS 1km 512 GPUs

- Global 1km
- Selene = A100
- H100 results preliminary

NVIDIA

HiRes Climate Drives CSCS Next-Gen System Alps

CSCS MIGRATION FROM PIZ DAINT TO ALPS
GRACE (ARM) + HOPPER (H100 GPU) SYSTEM

20 Exaflops of AI

Accelerated w/ NVIDIA Grace CPU and NVIDIA Hopper (H100) GPU

HPC and AI For Scientific Computing

Advanced Weather and Climate: Applications based on IFS and ICON

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Hewlett Packard Enterprise

NVIDIA

NVIDIA Next-Gen GPU H100 and Bandwidth

- Connect to other NVIDIA GPUs
- Program with NCCL
- *Mostly* used in AI applications

Custom TSMC 4N Process
4.9 TB/s Total External B/W

- Connect to x86 processors
- PCI Gen 5
- Not cache coherent

- Connect to Grace Arm Processor
- Can sustain full NVLINK BW into large host memory
- *Cache coherent*

ICON GPU Improvement ~6x Over 2 Year Span

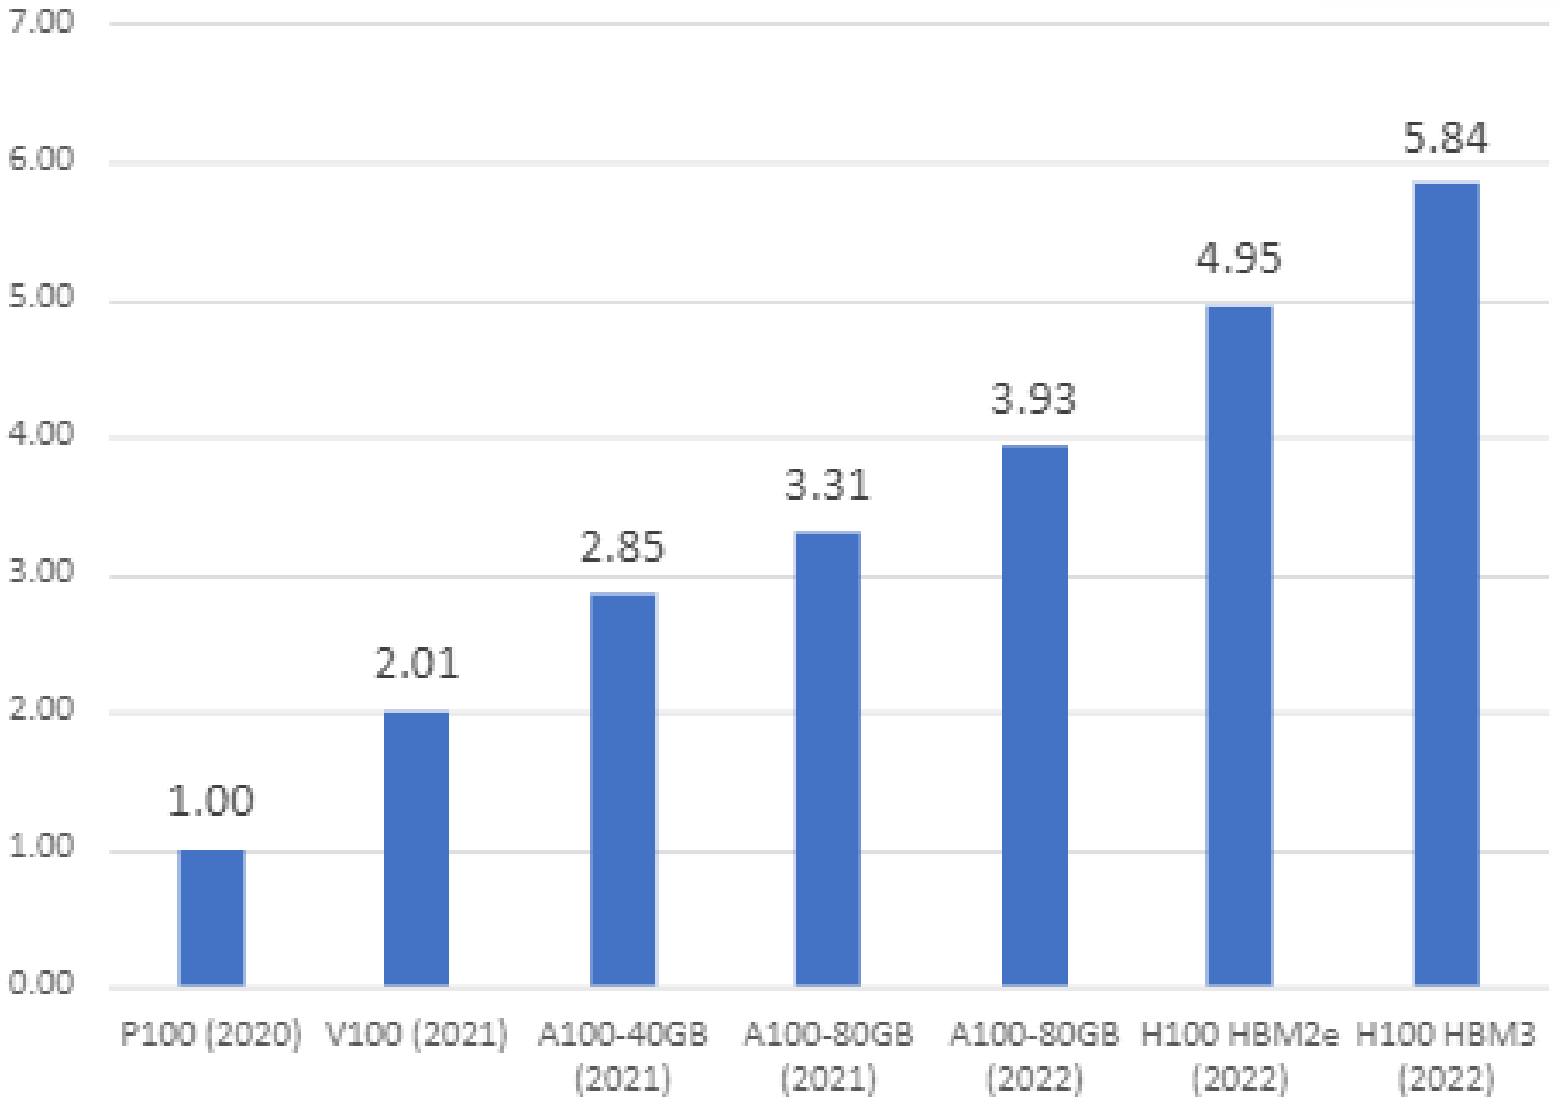

ICON ~6x from P100 to H100

- Spans 4 generations
- Single GPU results
- Coarse 160km, 80v
- H100 preliminary

ICON Strong Scaling on Large GPU Systems

QUBICC – R02B07 – 20km

5 dynamics substeps, horizontal diffusion and transport, graupel microphysics, vertical diffusion and JSBACH land; and RTE-RRTMGP radiation

Using internal timer report for values, excludes IO

GPU results are with NVHPC 21.2 except Piz Daint (PGI 20.1 is the latest there)

CPU results use Intel compiler and best values of ranks per node, nproma and radiation chunk

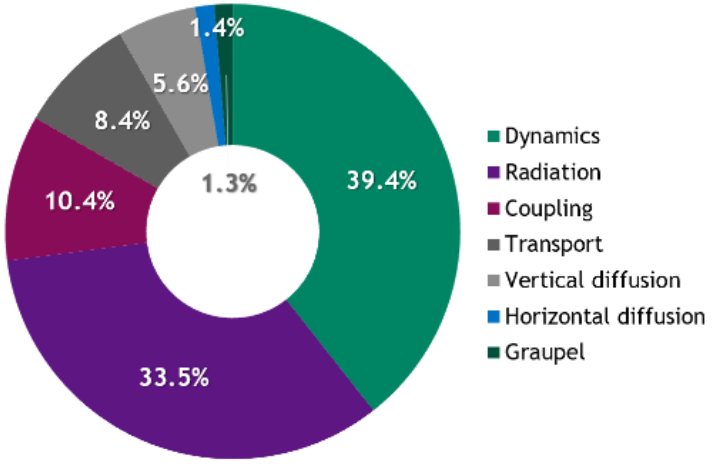

Time distribution per step on 2xEPYC 7742

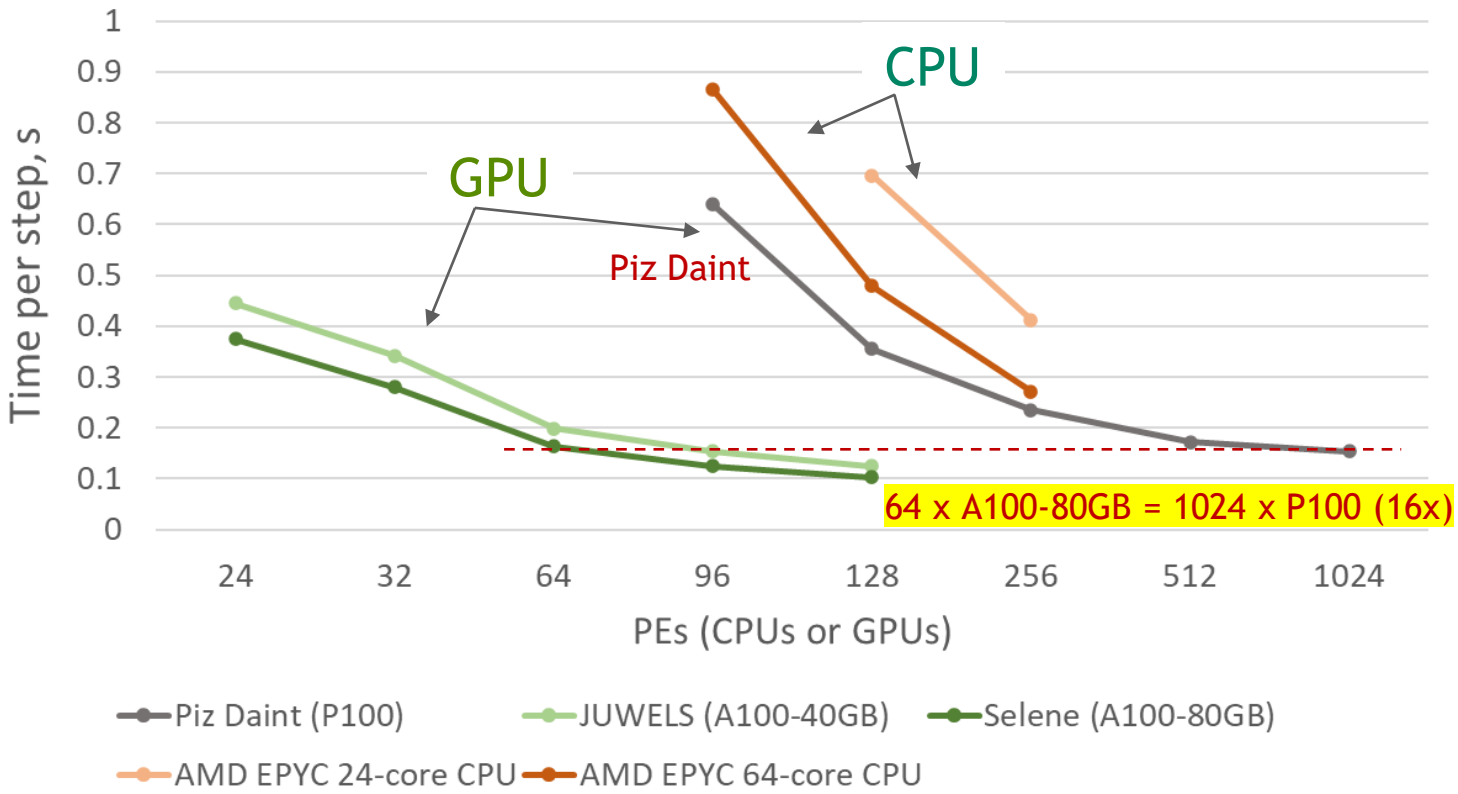

Source: NVIDIA, March 2021 Dr. Dmitry Alexeev

To Begin Earth-2 Will Train on Existing Climate Predictions

Reanalysis data and IPCCs current library of climate predictions

- Training demonstrated on ERA5 reanalysis for medium range weather scale (later slides)
- Current focus IPCC CMIP-6 library of 100km and HiResMIP 25km data for climate scale
- Plans for progression to higher resolutions and diversification of models:
  - Reanalysis data from NCAR, NASA, etc.
  - NCAR CESM, GPU-MPAS convection-permitting at 3.75km global resolution
  - US DOE E3SM MMF and SCREAM SR model
  - DYAMOND I (IFS, ICON, MPAS, GEOS, etc.) and II (coupled ocean)
  - Ultimately SR-ESMs towards 1km resolution (nextGEMS: ICON and IFS)
  - Etc.
- Earth observation data for calibration

Earth-2 Results at Weather Scale Predictions: FourCastNet

Fourier Forecasting Network

- Collaboration between NVIDIA and several US academics and National labs
- Data-driven ML surrogate model for global medium-range weather prediction
- Trained on 10TB of ~35 years ECMWF ERA5 (30km, 137 levels) reanalysis data
- Adaptive FNO architecture trained on 20 input variable per grid point
- Highest resolution data driven model ever trained
- Inference time of .25 sec for 2 week forecast
  - Speedup vs NWP: $O(10^4\text{-}10^5)$
  - Power savings vs. NWP: $O(10^4)$
- Exceptional skill for extreme weather prediction

Input Variables (20)

Vertical Level	Variables
Surface	$U_{10}, V_{10}, T_{2m}, sp, mslp$
$1000hPa$	U, V, Z
$850hPa$	T, U, V, Z, RH
$500hPa$	T, U, V, Z, RH
$50hPa$	Z
Integrated	$TCWV$

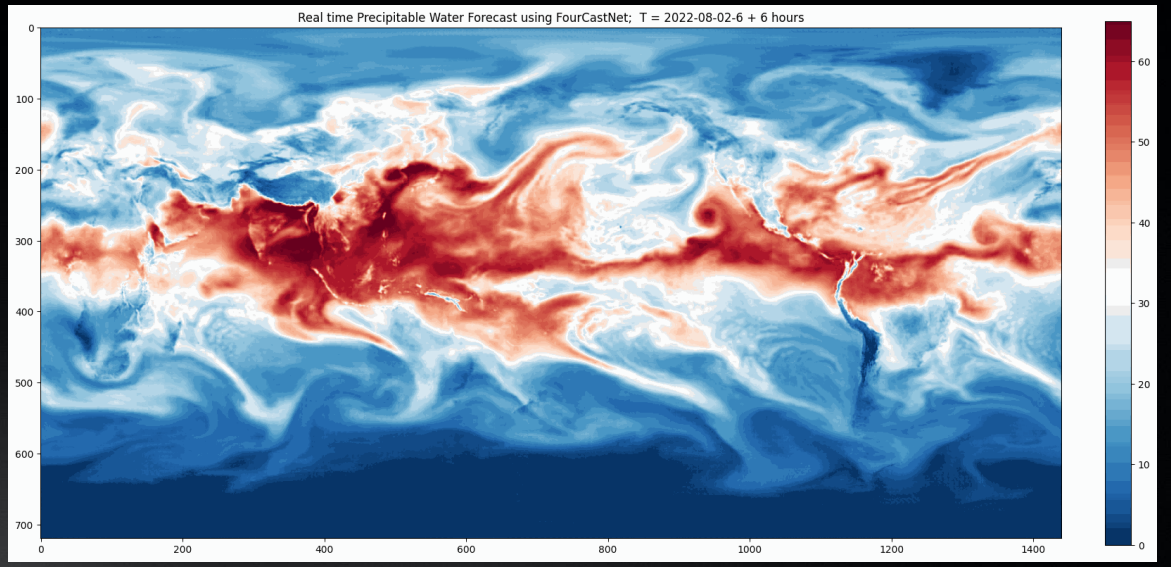

Pathak et al. (2022), FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators, arXiv:2202.11214

NVIDIA

Weather Forecast Skill Demonstrated in a Fraction of a Second

Expensive HPC model demonstrated as an interactive surrogate with similar skill

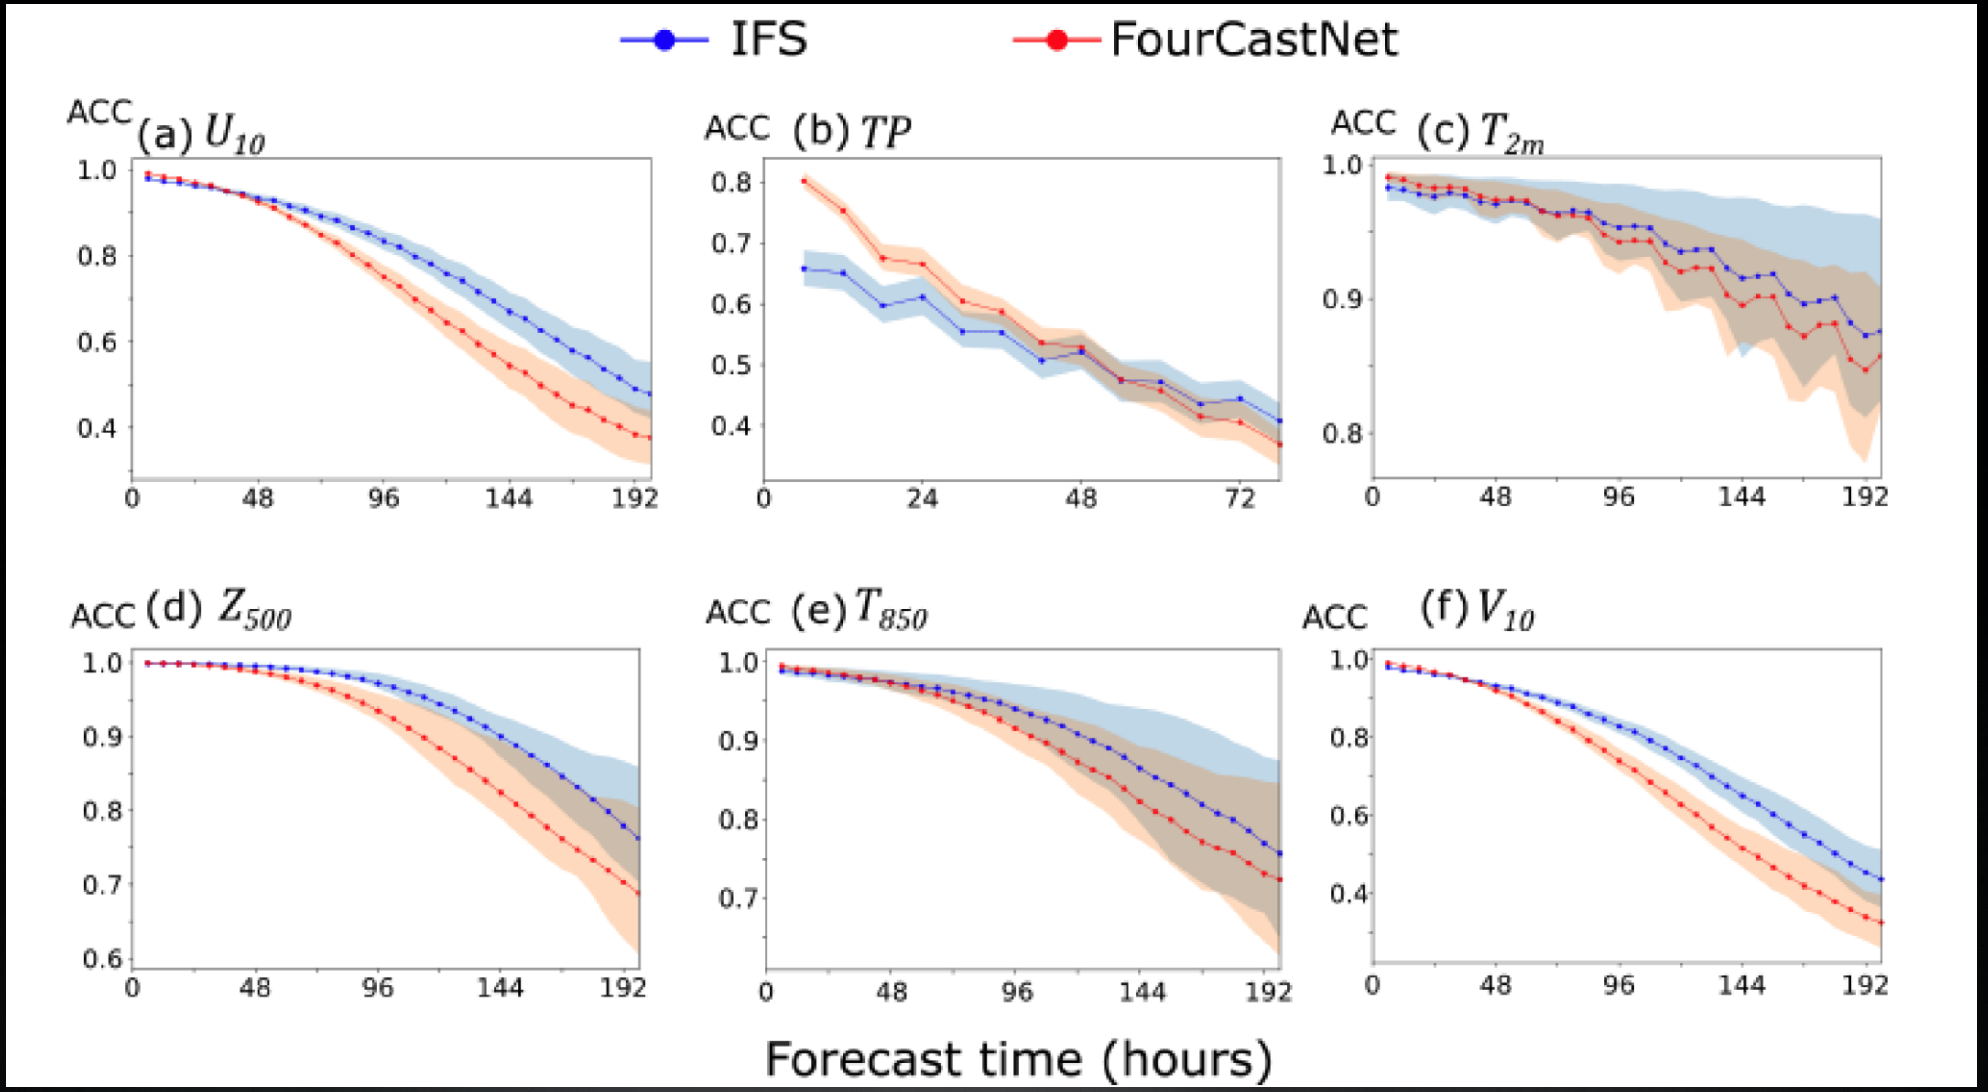

FourCastNet Highest Resolution of Data-Driven Weather Models

Comparison of resolutions for data-driven weather models since 2018 (Dueben & Bauer)

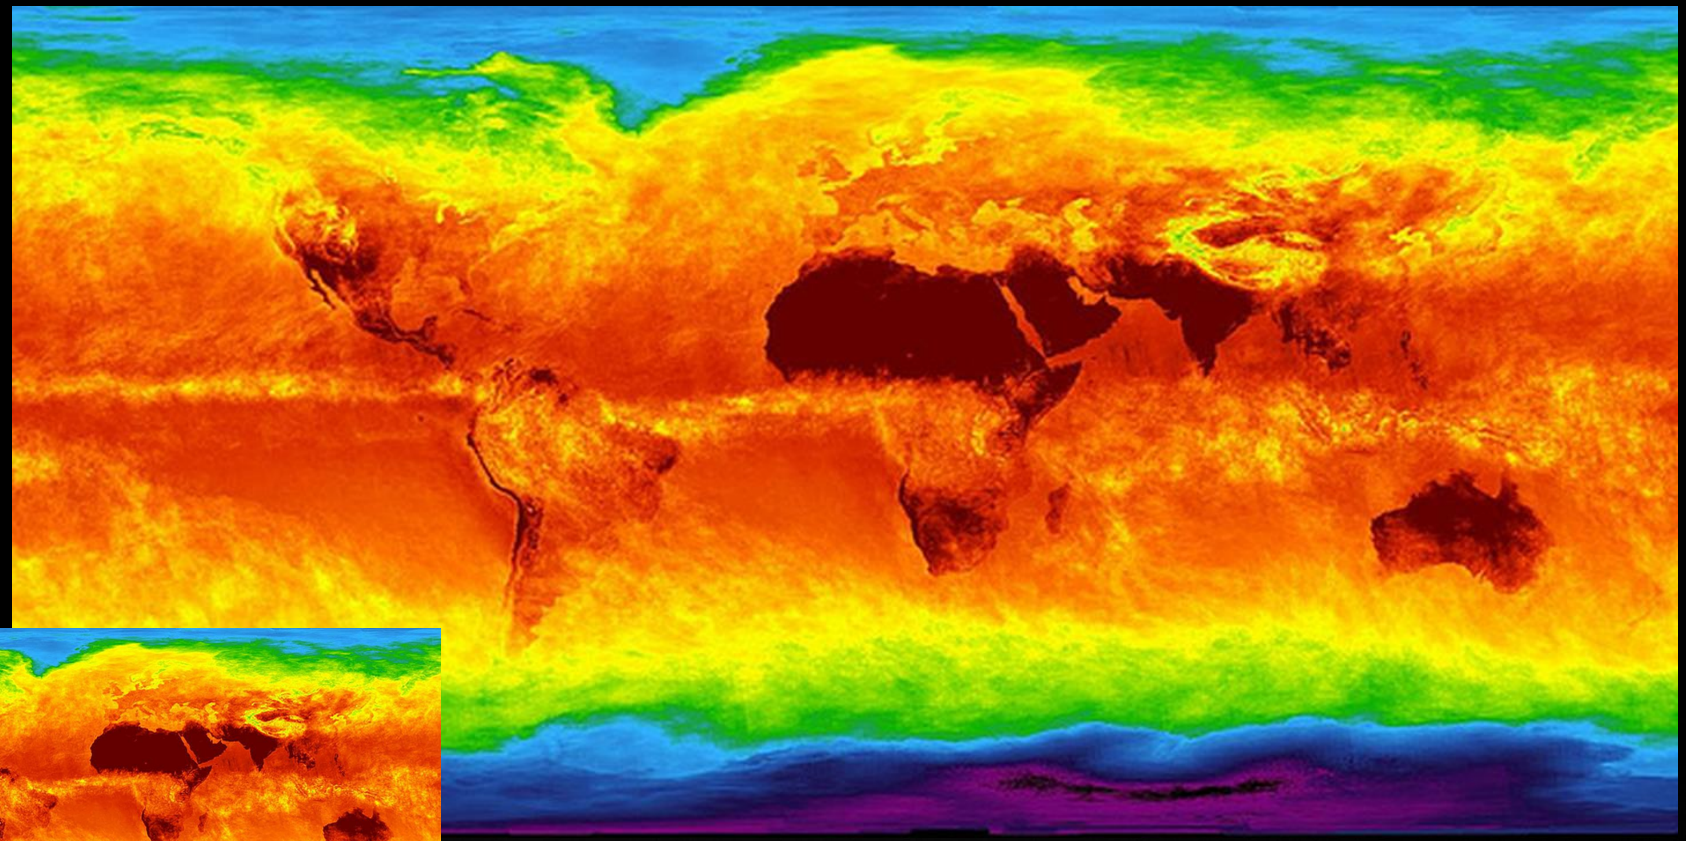

FourCastNet, Pathak et al. (2022), 0.25°, ~1,000,000 Pixels, ViT+AFNO

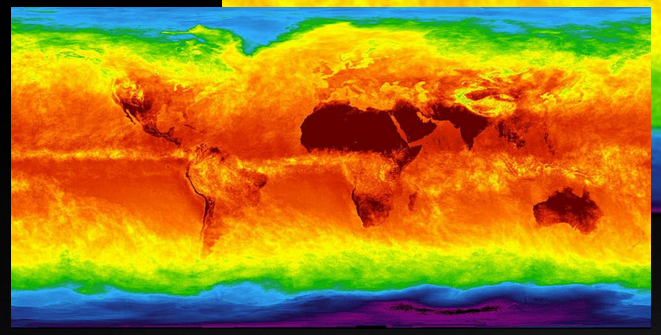

GNN, Keisler et al. (2022), 1°, 64,000 Pixels, Graph Neural Networks

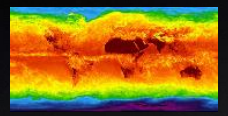

DLWP, Weyn et al. (2020). 2°, 16K pixels, Deep CNN on Cubesphere/(2021) ResNet

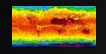

Weyn et al. (2019), 2.5° N.H only, 72x36, 2.6k pixels, ConvLSTM

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

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

FourCastNet Strong Scaling for Training on Large Scale Systems

FourCastNet training scaled efficiently up to ~ 4000 GPUs on Top500 supercomputing systems

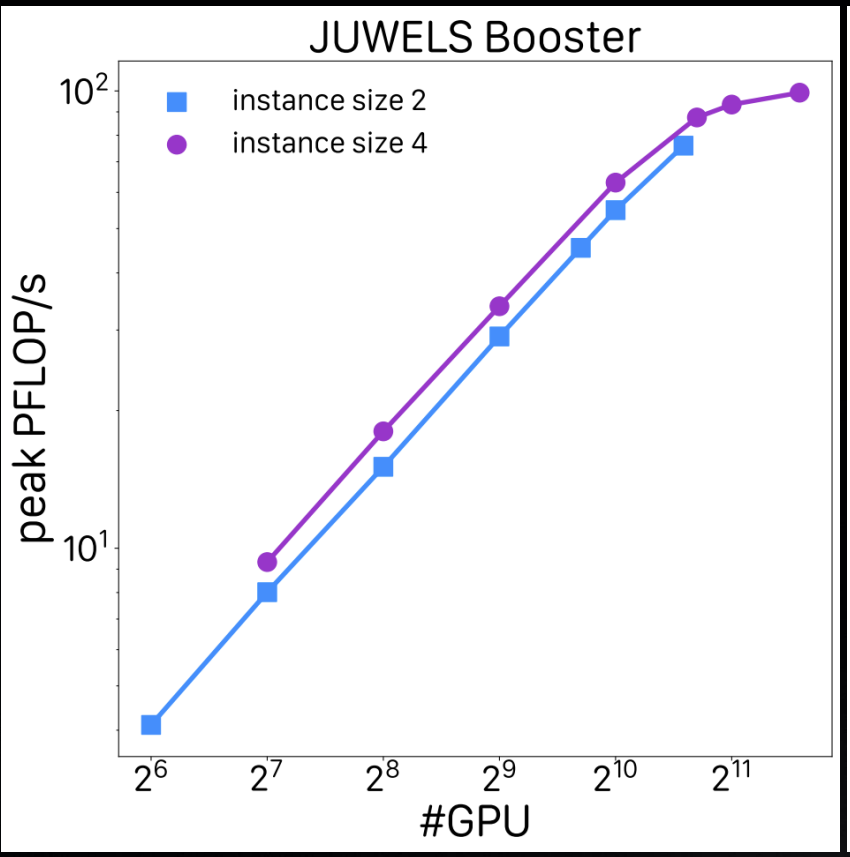

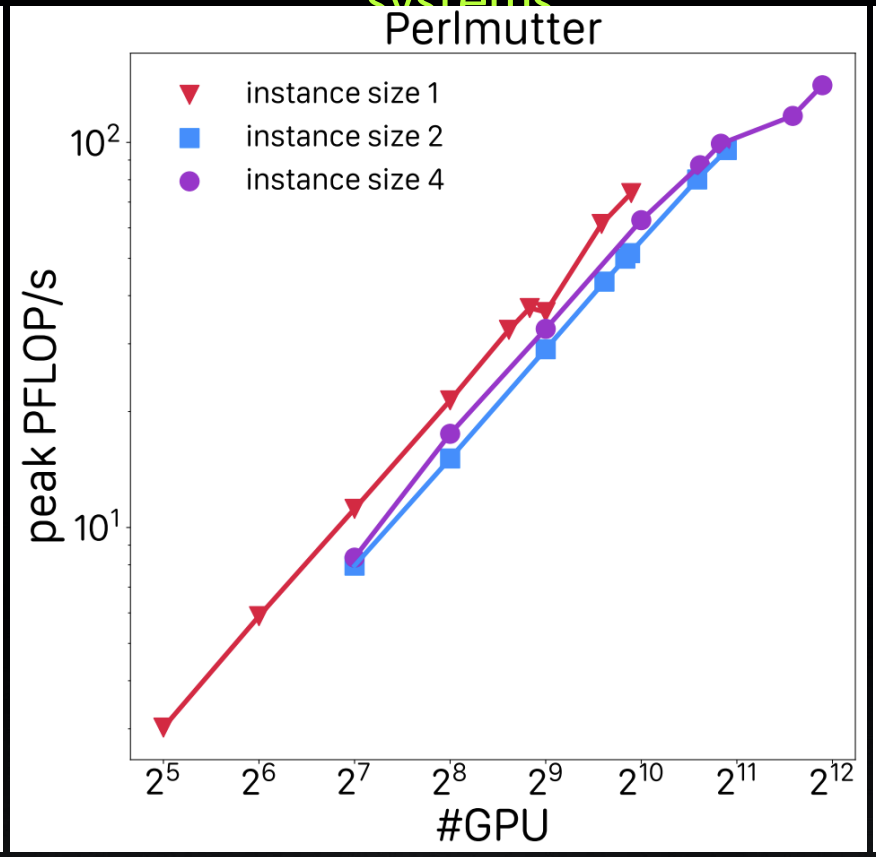

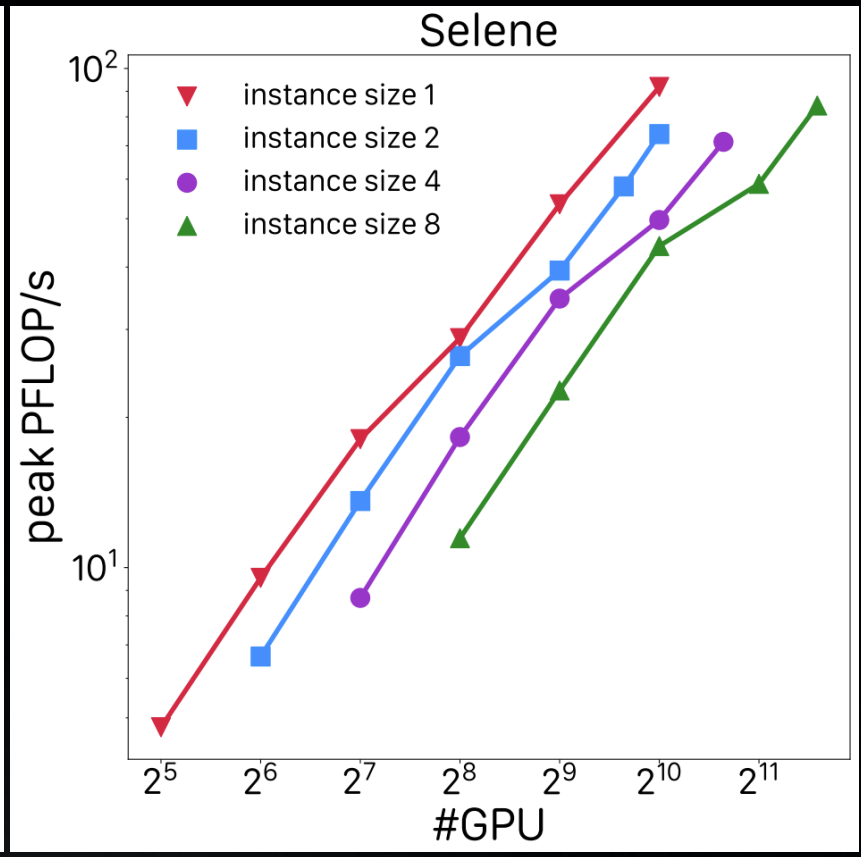

Peak performance is 141 petaFLOPS in mixed precision (averaged over a full epoch)

Model training time: ~ 24 hrs reduced to 67 min (~21x) on 3,072 GPUs

FOURCASTNET PREDICTION OF ATMOSPHERIC RIVERS

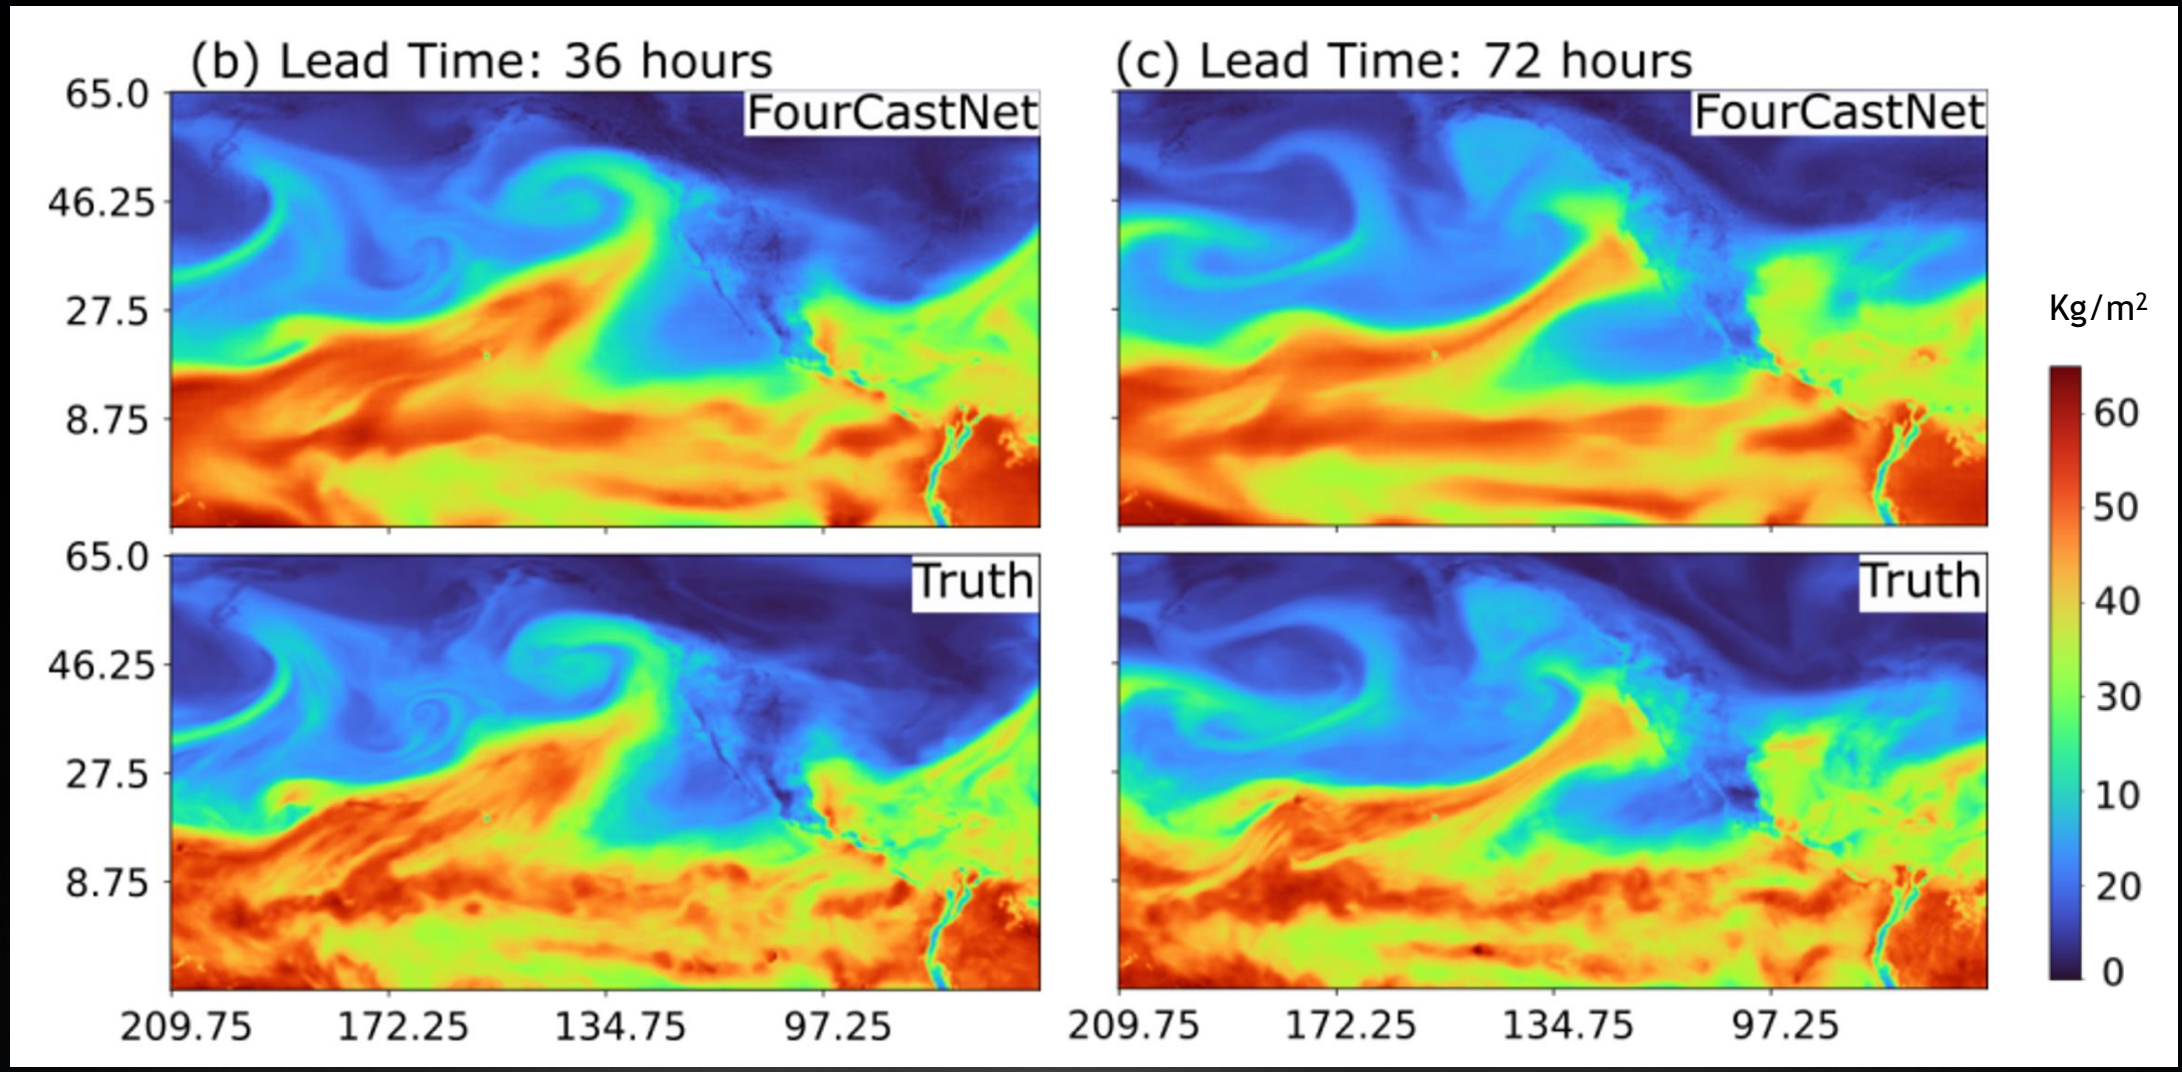

FOURCASTNET PRECIPITATION AND SURFACE WINDS

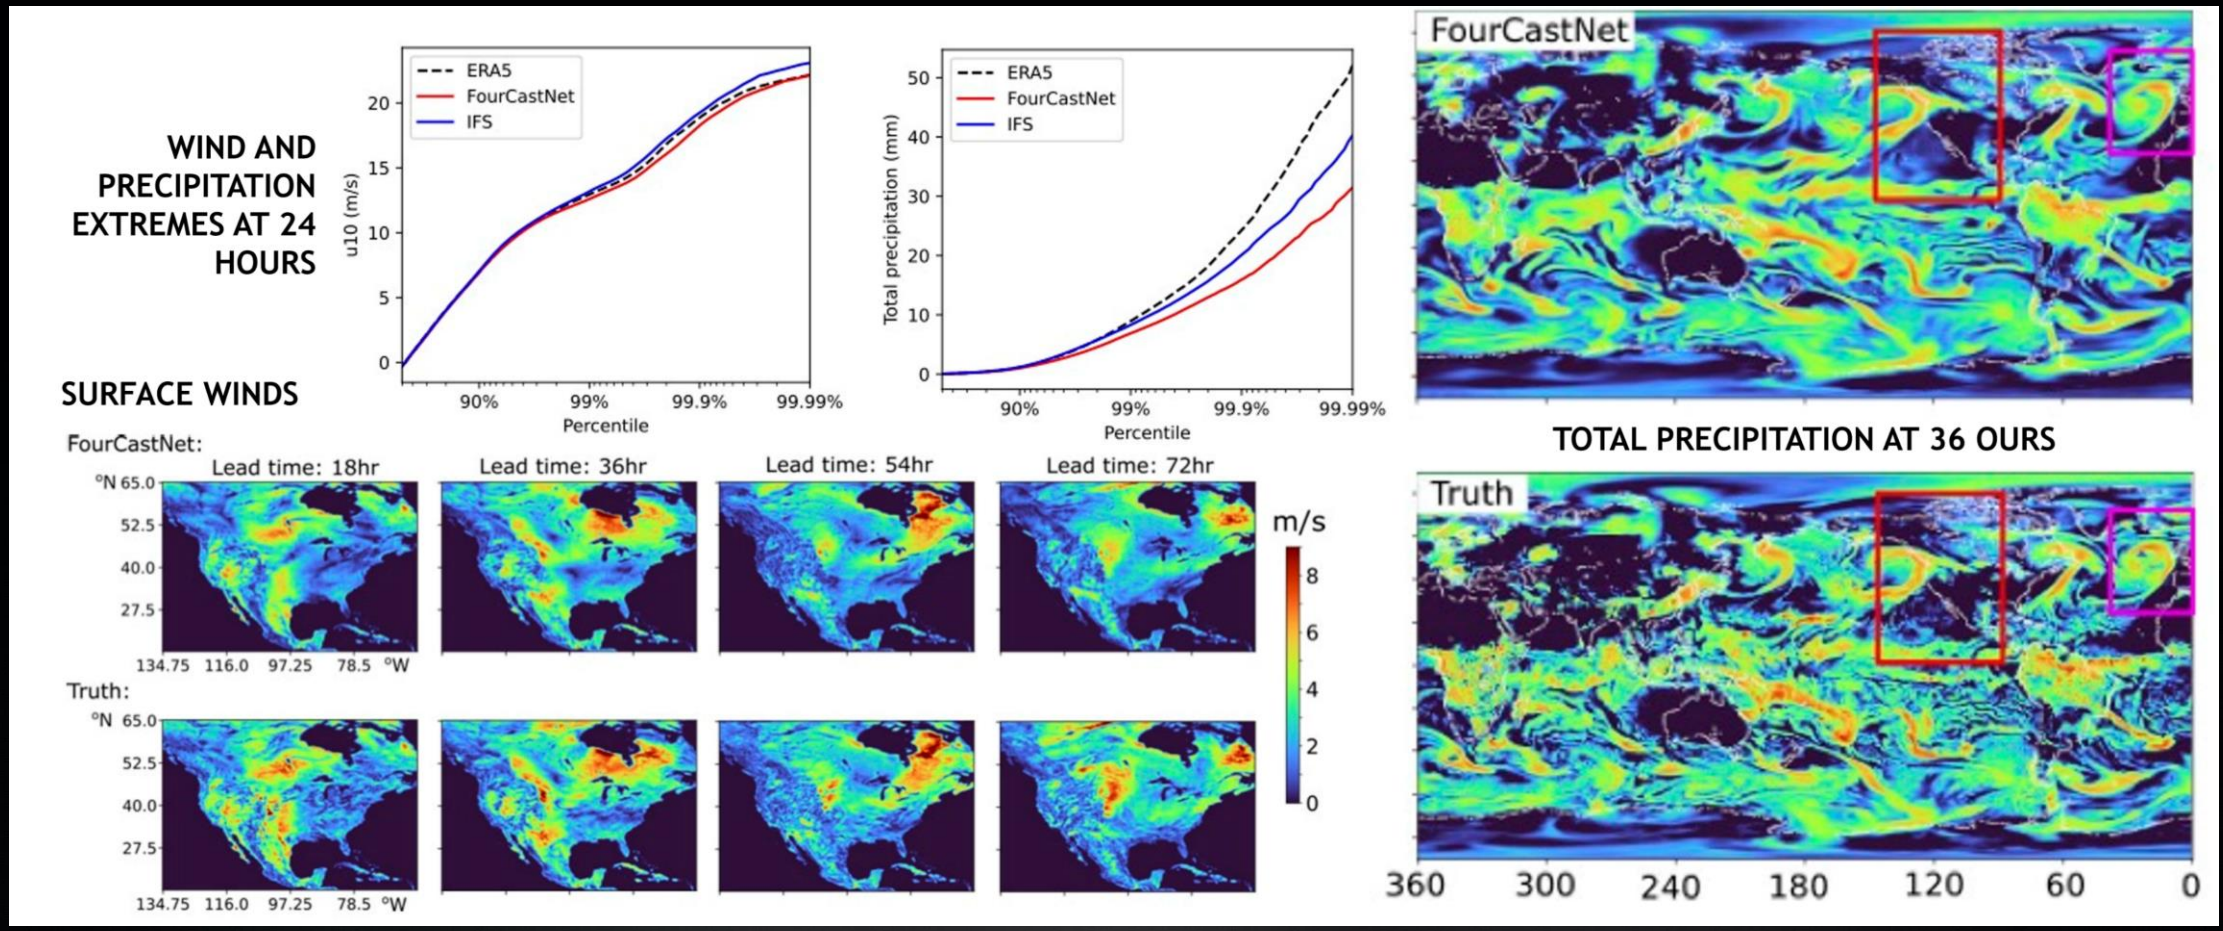

FOURCASTNET PREDICTS HURRICANE PATHS AND INTENSITIES

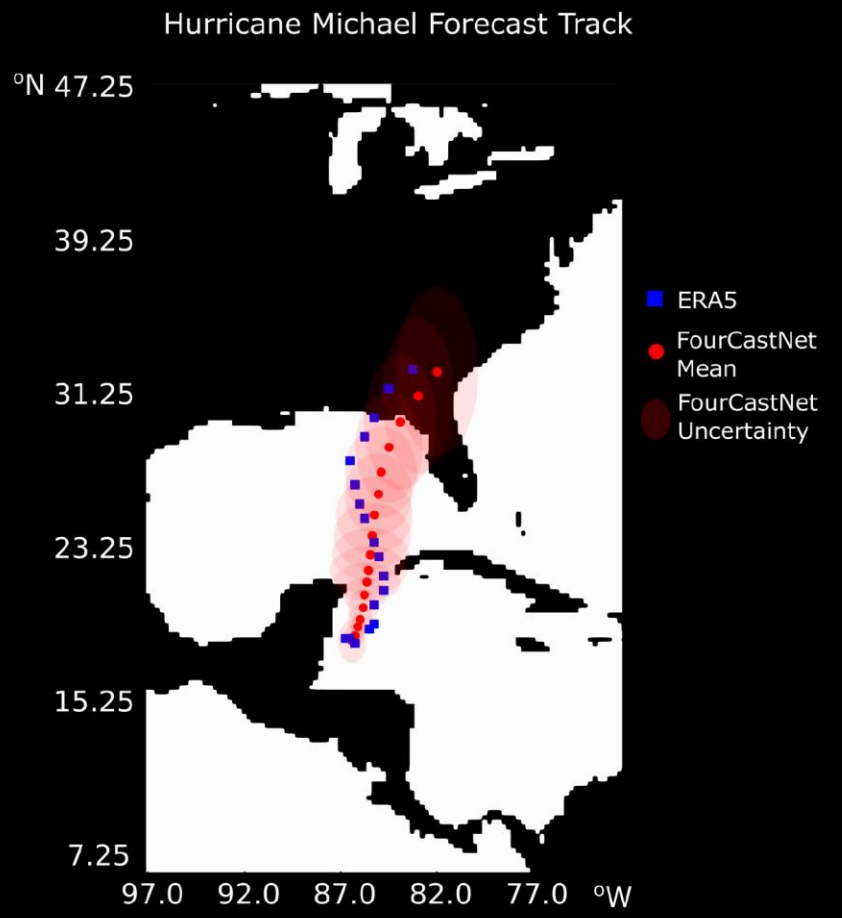

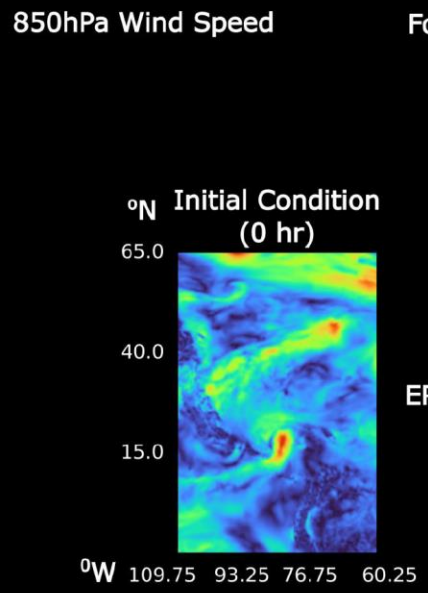

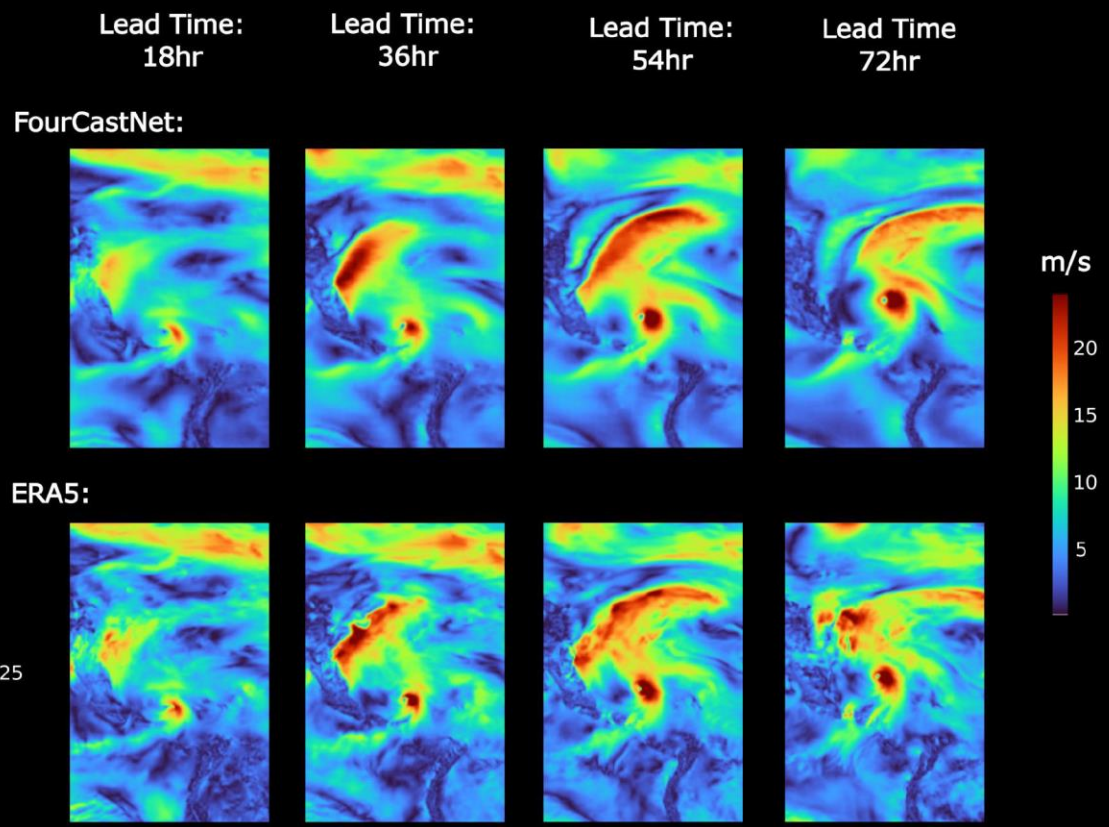

FourCastNet Skill Improves with More Variables

FourCastNet skill Improved significantly from 20 to 26 variables on medium range forecast skill

We've not yet reached the skill limit

First 20 Variables

Vertical Level	Variables
Surface	$U_{10}, V_{10}, T_{2m}, sp, mslp$
$1000hPa$	U, V, Z
$850hPa$	T, U, V, Z, RH
$500hPa$	T, U, V, Z, RH
$50hPa$	Z
Integrated	$TCWV$

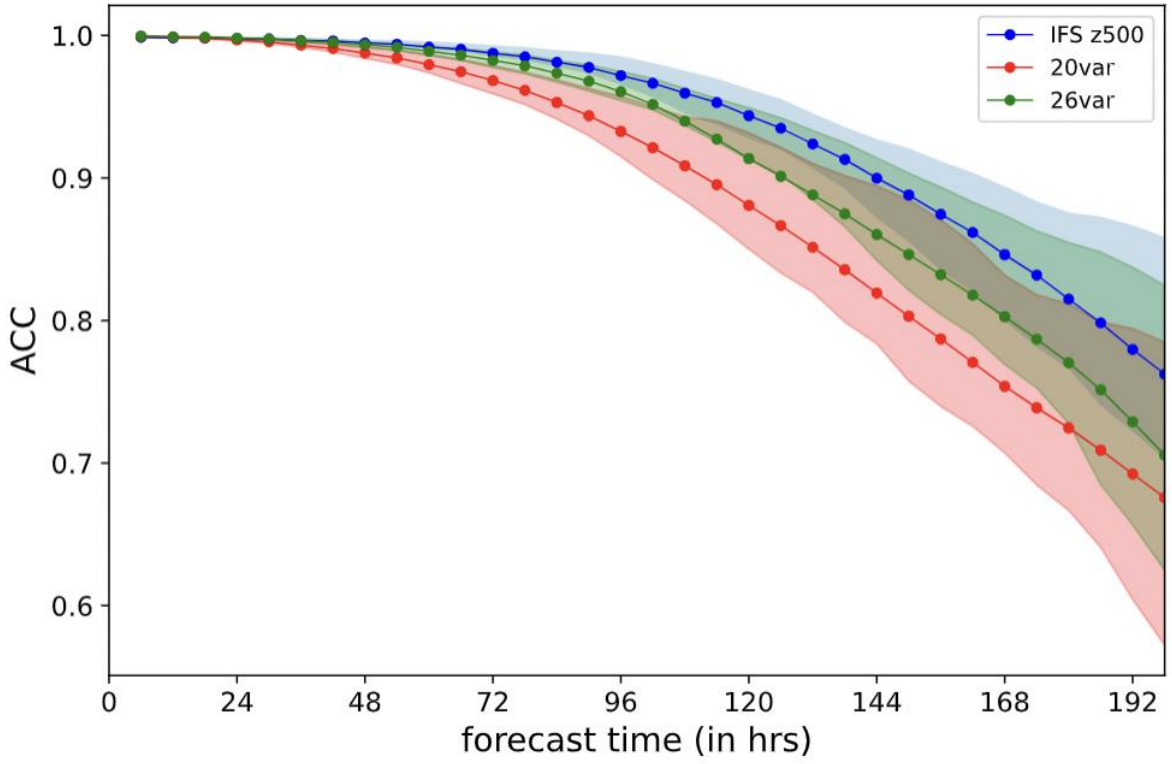

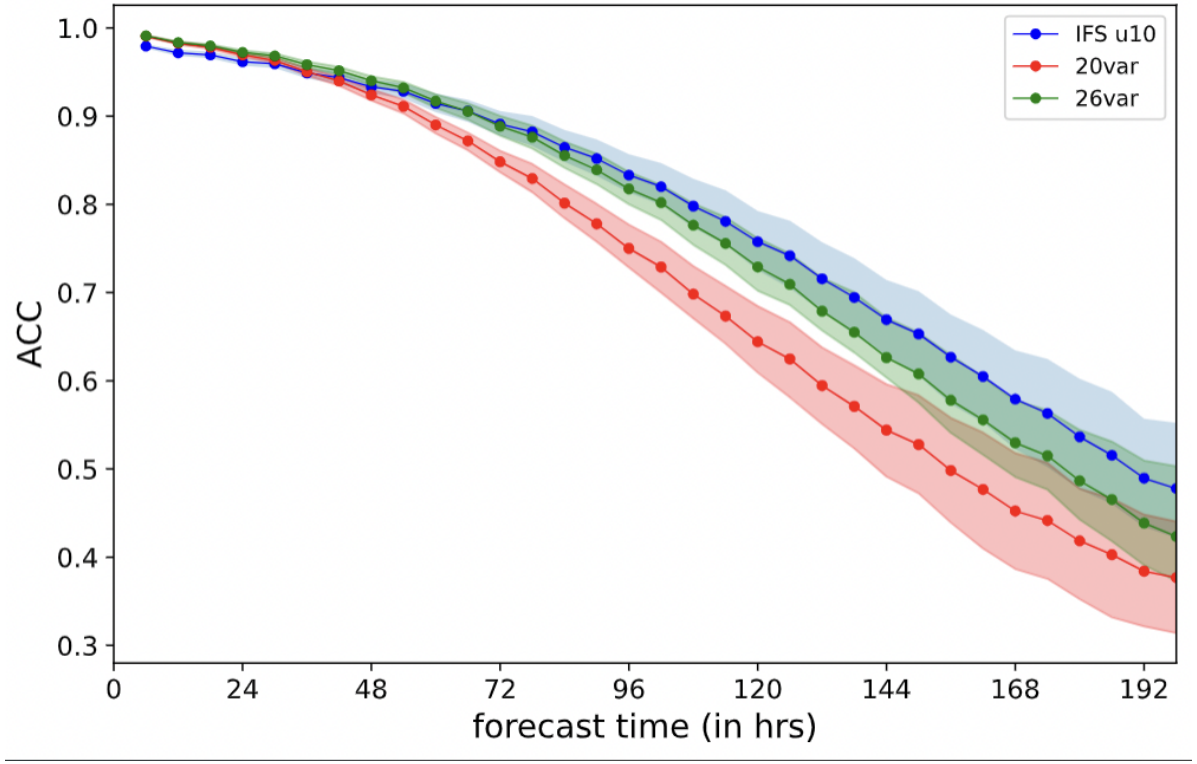

FourCastNet Initialized with GFS Forecast

FCN

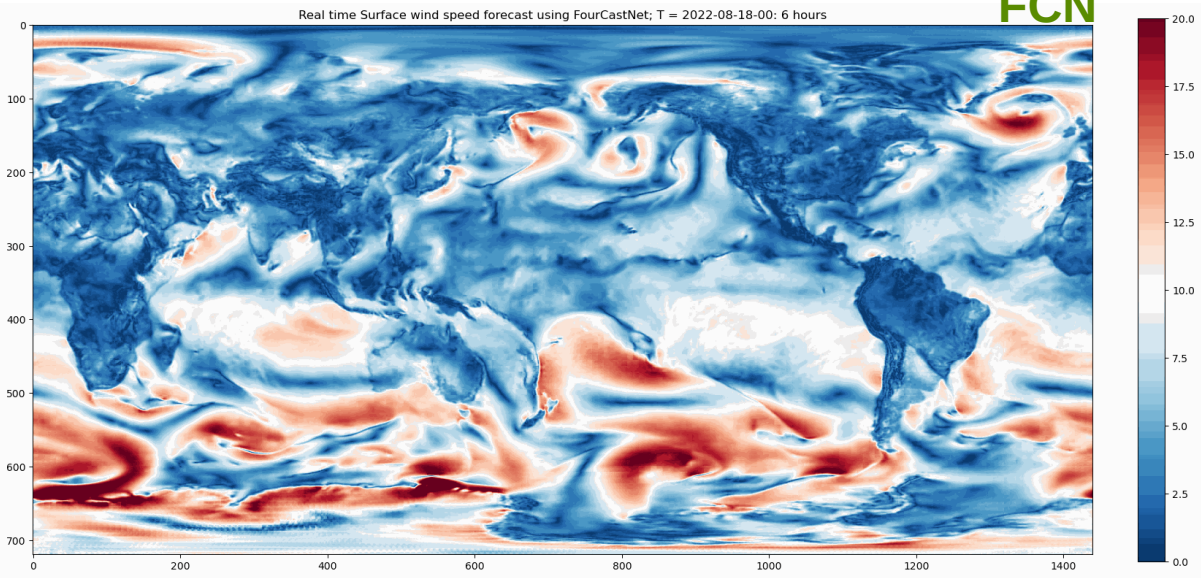

GFS

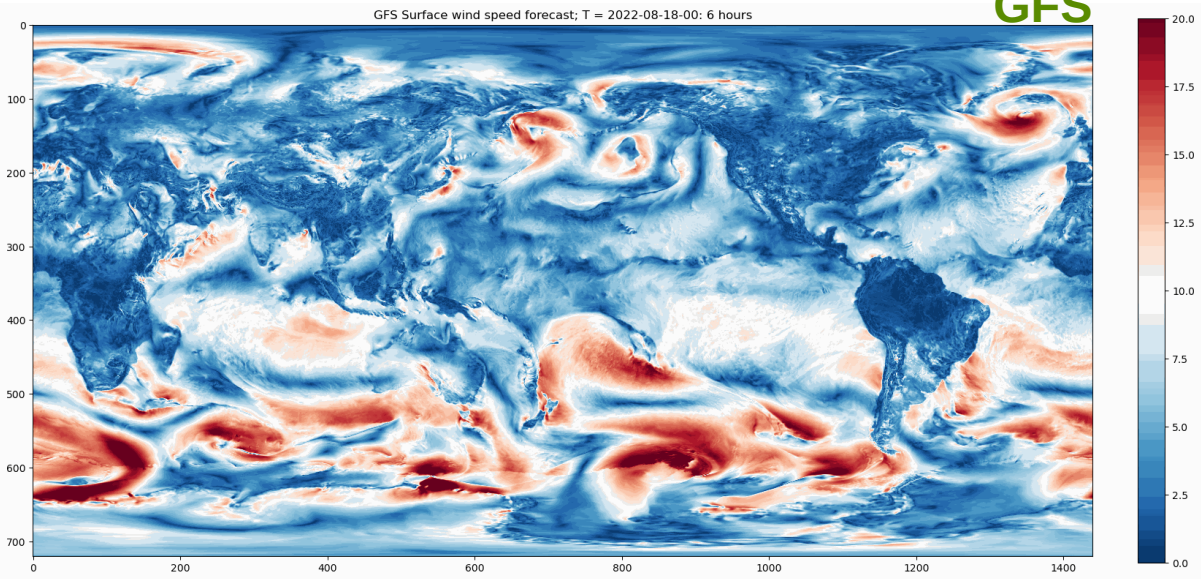

- **Surface wind speed forecast at 10m week of 18 Aug 22**
- **FourCastNet prediction upper; GFS (ground truth) lower**
- FCN trained on ERA5 and initialized with GFS demonstrating zero-shot skill transfer of a GFS medium-range forecast
- **Ensembles**: FCN can generate 1000s of calibrated ensemble forecasts in minutes to aid in forecast uncertainty quantification, or used as part of a DA pipeline
- **Local Area Model**: Regional version of FCN in development
- **Larger input vector**: Plans to forecast a much larger vector of atmospheric variables – from the current 26 to 250

FourCastNet Available as Open Source

Join NVIDIA efforts in pushing the frontiers of data-driven weather prediction

https://github.com/NVlabs/FourCastNet

README.md

FourCastNet

This repository contains the code used for "FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators" [paper]

The code was developed by the authors of the preprint: Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, Pedram Hassanzadeh, Karthik Kashinath, Animashree Anandkumar

FourCastNet, short for Fourier Forecasting Neural Network, is a global data-driven weather forecasting model that provides accurate short to medium-range global predictions at 0.25∘ resolution. FourCastNet accurately forecasts high-resolution, fast-timescale variables such as the surface wind speed, precipitation, and atmospheric water vapor. It has important implications for planning wind energy resources, predicting extreme weather events such as tropical cyclones, extra-tropical cyclones, and atmospheric rivers. FourCastNet matches the forecasting accuracy of the ECMWF Integrated Forecasting System (IFS), a state-of-the-art Numerical Weather Prediction (NWP) model, at short lead times for large-scale variables, while outperforming IFS for variables with complex fine-scale structure, including precipitation. FourCastNet generates a week-long forecast in less than 2 seconds, orders of magnitude faster than IFS. The speed of FourCastNet enables the creation of rapid and inexpensive large-ensemble forecasts with thousands of ensemble-members for improving probabilistic forecasting. We discuss how data-driven deep learning models such as FourCastNet are a valuable addition to the meteorology toolkit to aid and augment NWP models.

FourCastNet is based on the vision transformer architecture with Adaptive Fourier Neural Operator (AFNO) attention proposed in Guibas-Mardani et al. [paper], [code].

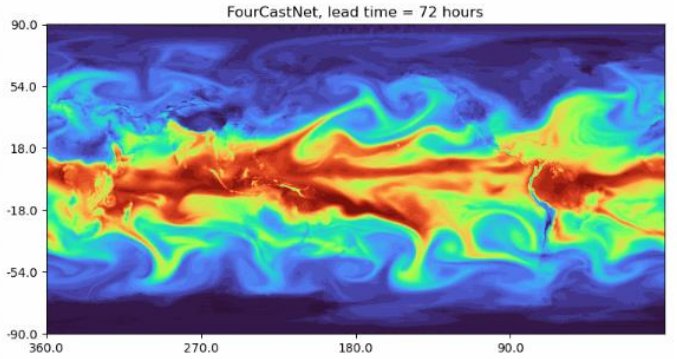

Future Work Towards Climate Digital Twins

NVIDIA is pursuing various strategies to improve the Earth-2 climate digital twin

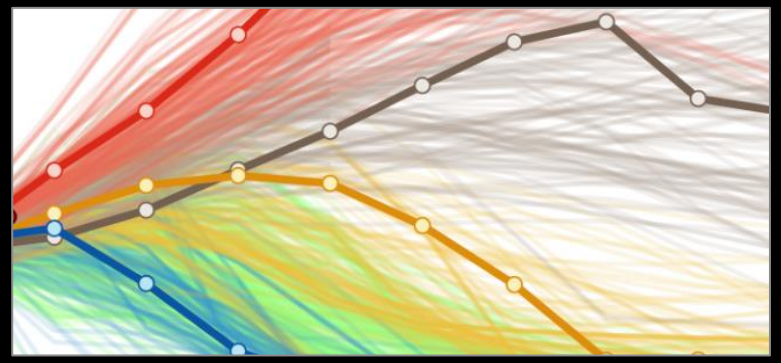

CMIP-6 Initialization or Surrogate

Source: Fuss et al., 2014

ICON 1km tethering

source: dyamond proj

OVX hardware / software co-design

ovx superpods

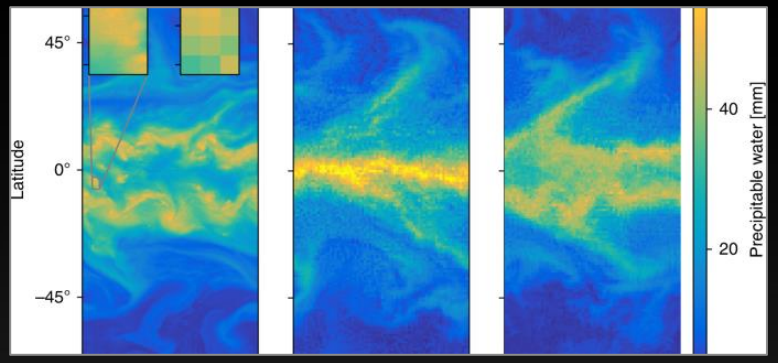

1km → 5 km AI sub-grid emulation

source: nature.com

Regional Fine-Tuning and Downscaling

source

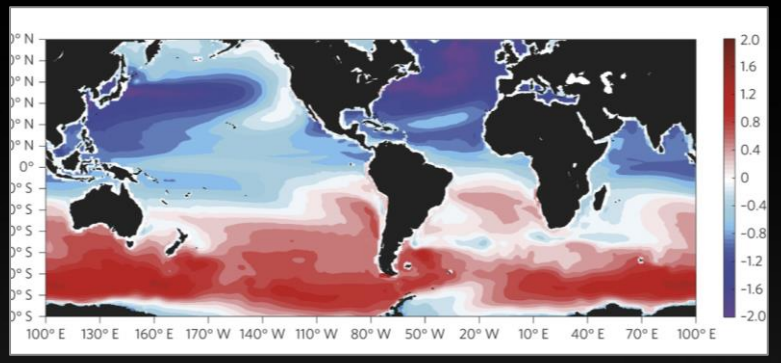

Model Auto-Calibration via RL

source: nature

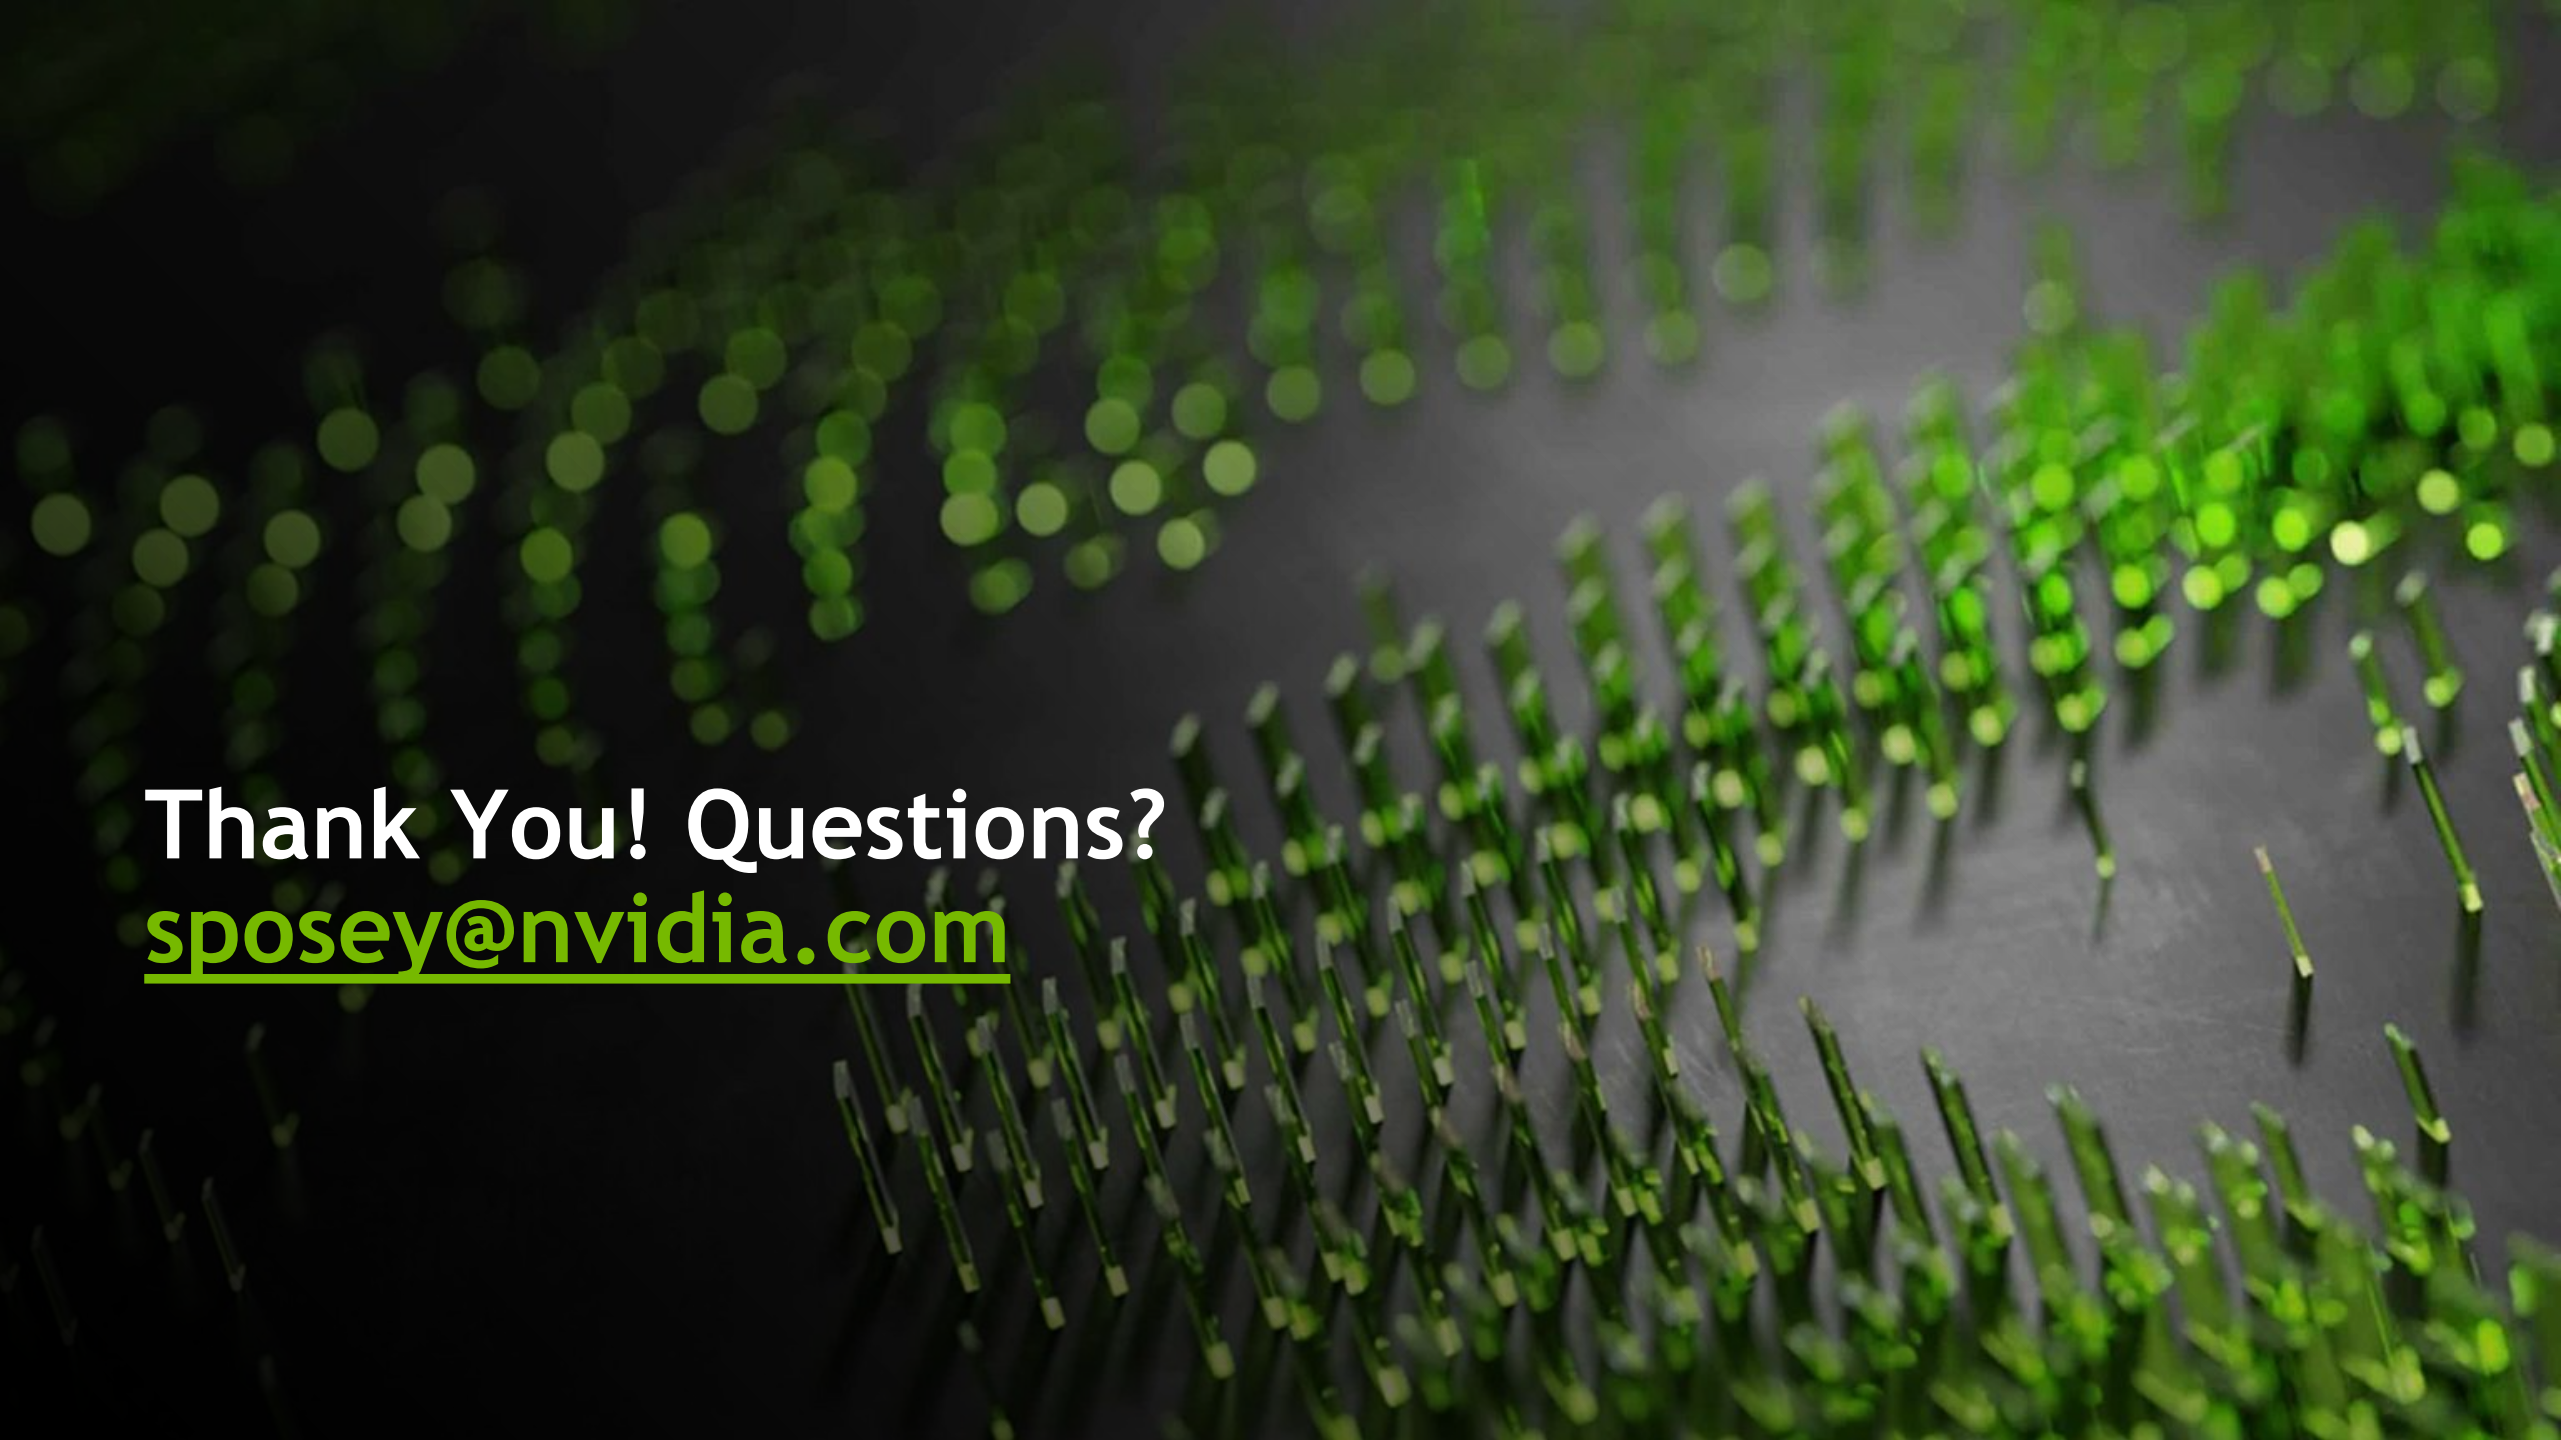
Thank You! Questions?
sposey@nvidia.com