NVIDIA

AI-Accelerated Physical Modelling for Weather, Climate, and Engineering at NVIDIA

Ira Shokar, PhD | Applied Scientist | NVIDIA | ishokar@nvidia.com

CUDA-X Accelerates Every Industry

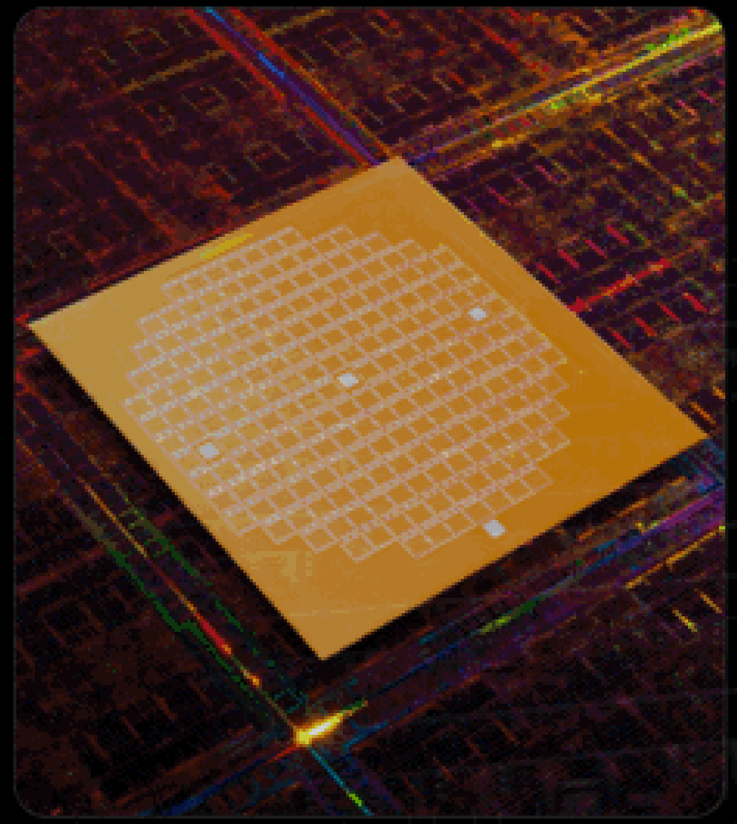

cuLitho
Computational Lithography

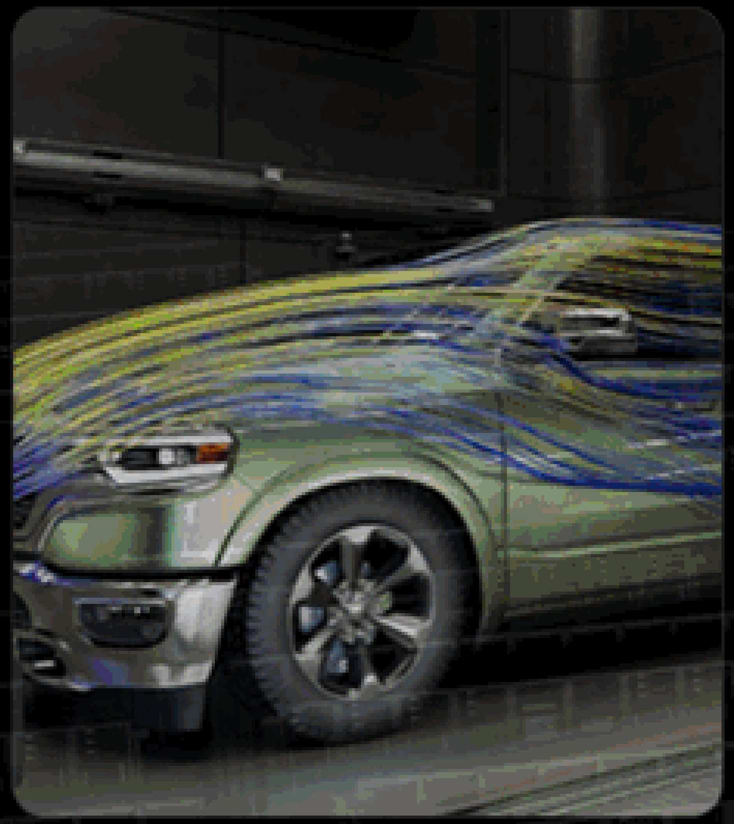

cuDSS
CAE

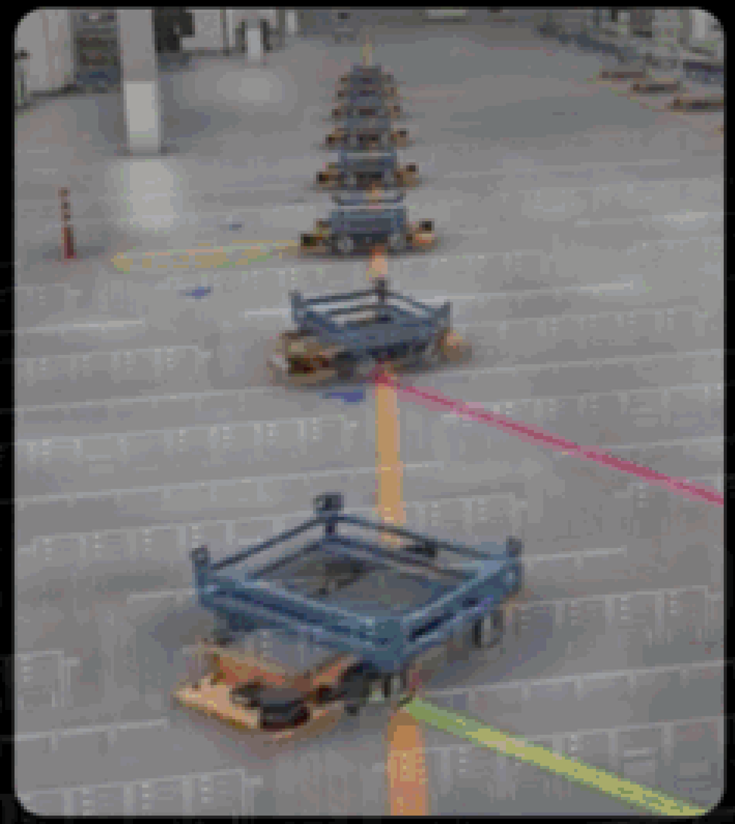

cuOpt
Decision Optimization

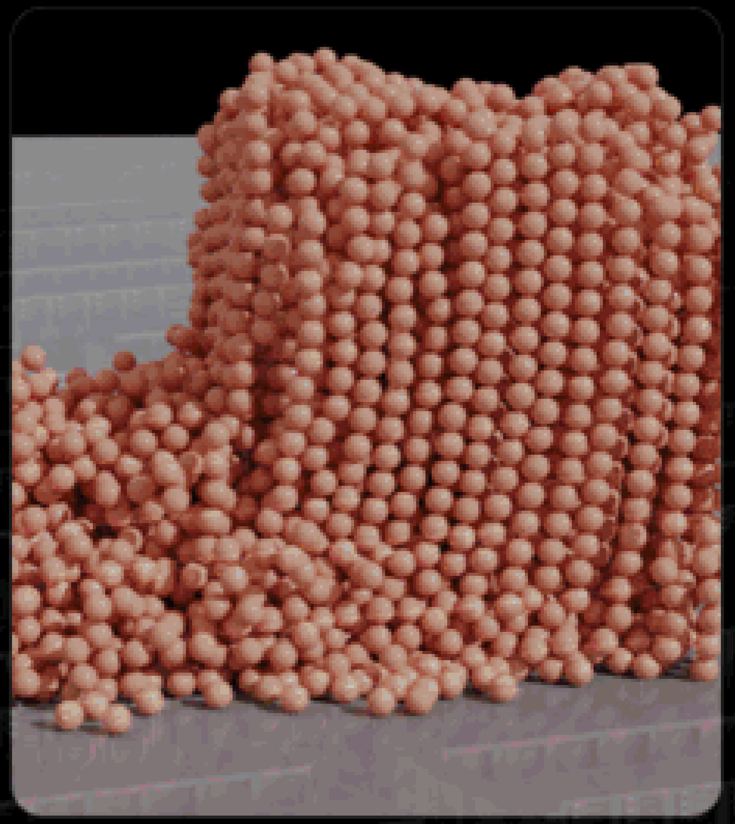

Warp
Physical Simulation

cuDF
Data Processing

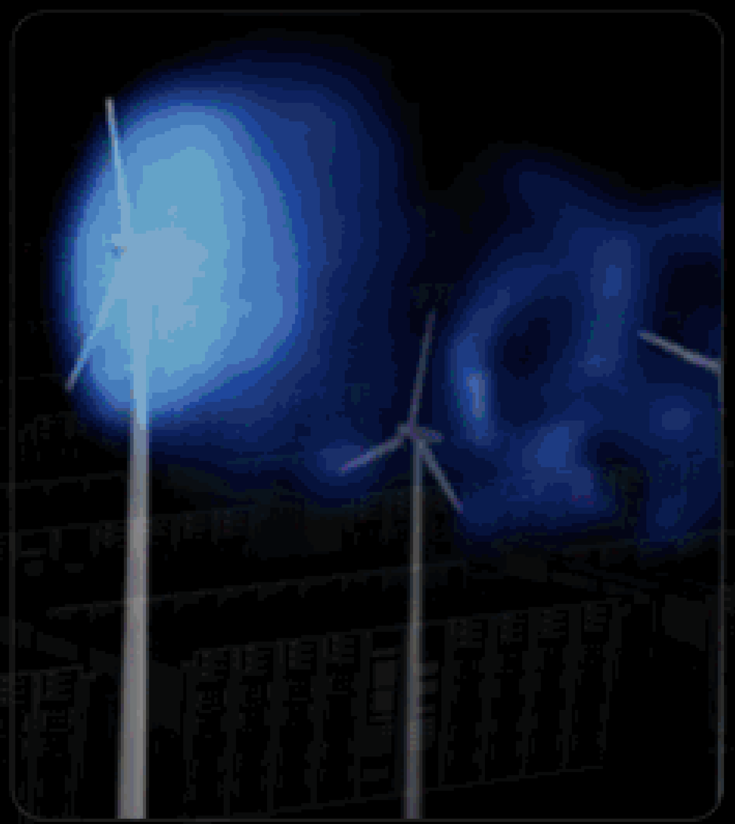

PhysicsNeMo
AI Physics

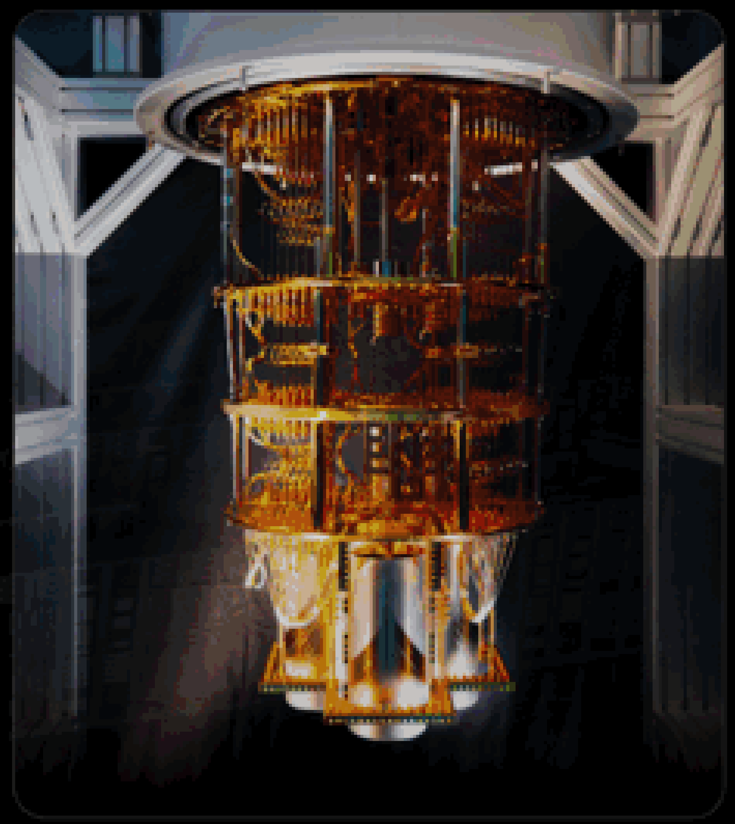

CUDA-Q
Quantum Computing

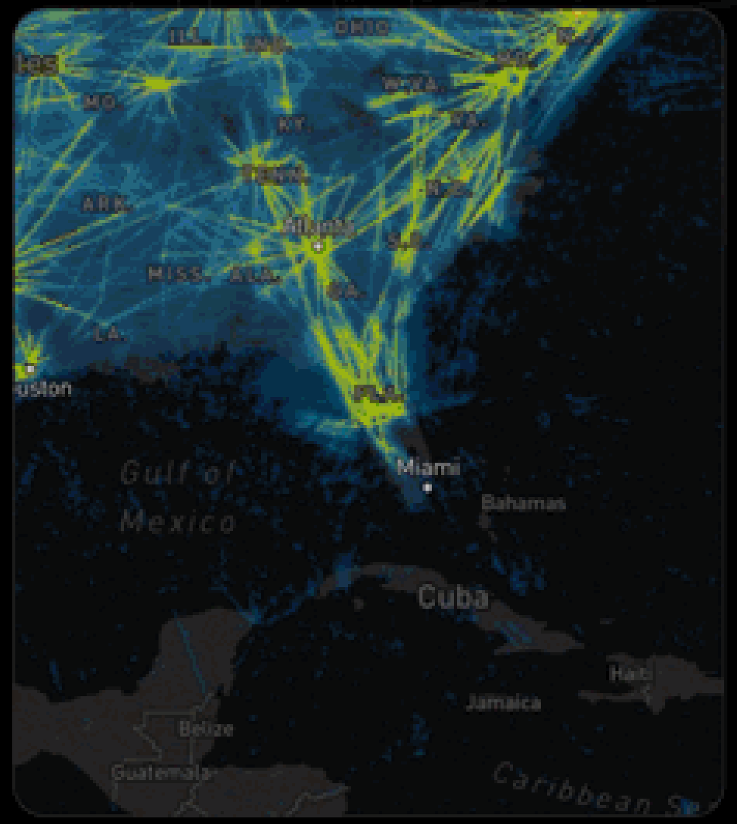

cuEquivariance
Drug & Materials Discovery

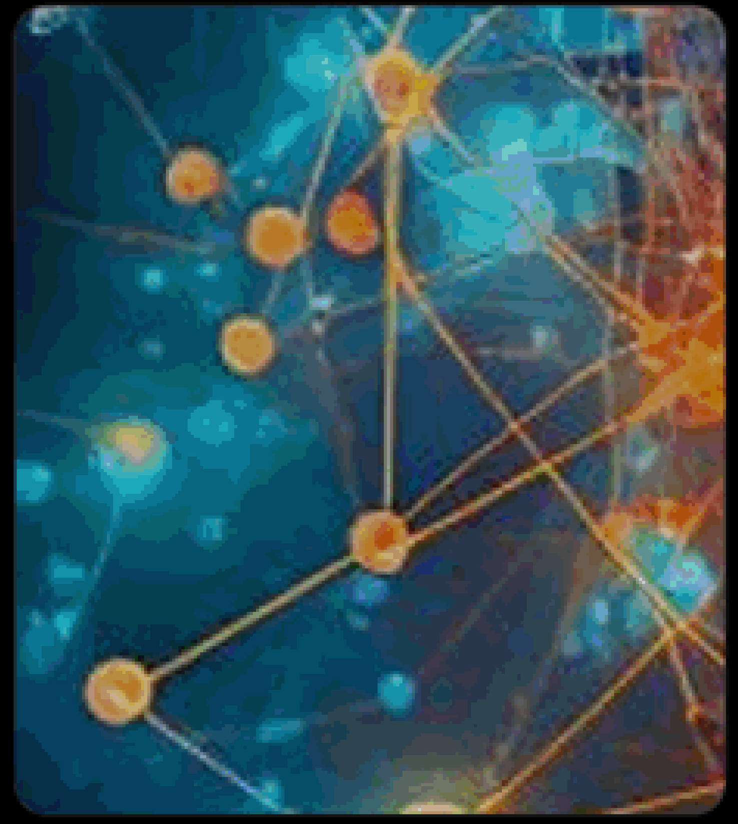

ALCHEMI
AI Materials Science

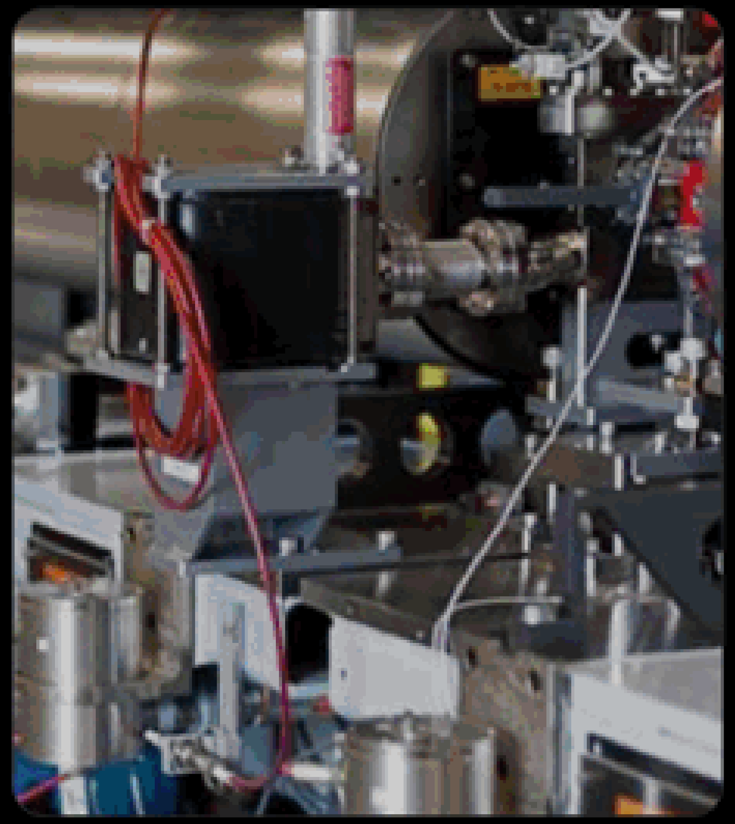

Holoscan
Edge HPC

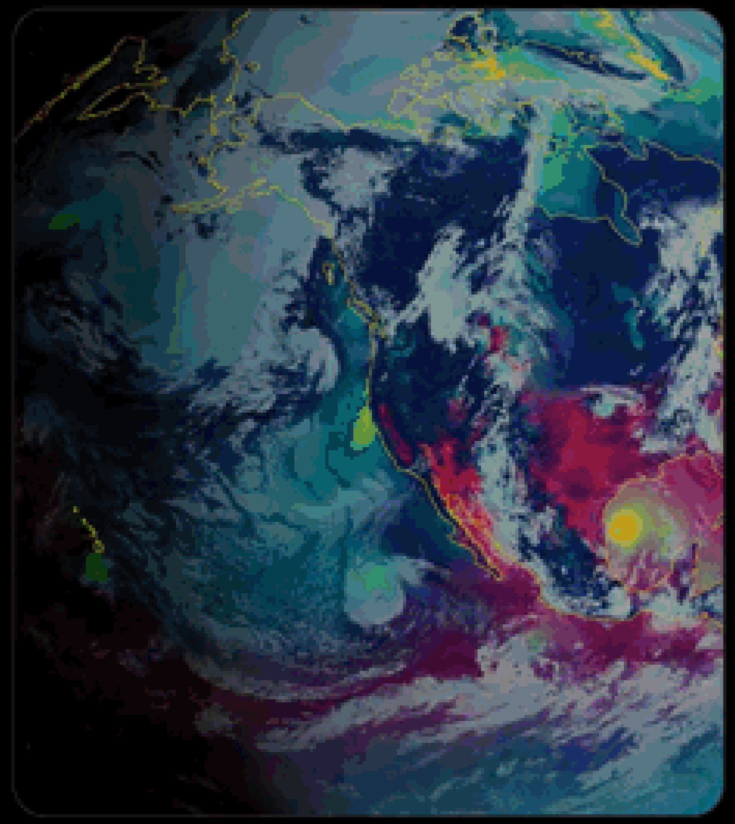

Earth-2
Weather Analytics

Parabricks
Gene Sequencing

cuPyNumeric
Numerical Computing

NVIDIA

NVIDIA AI Accelerated Computing Platform

Hardware and Software Acceleration Across Every Workload and Vertical

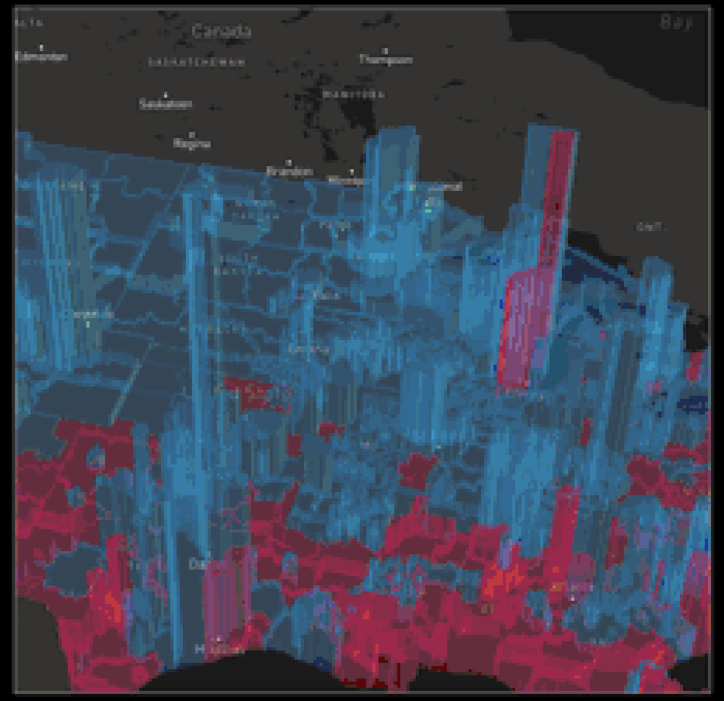

Data Processing

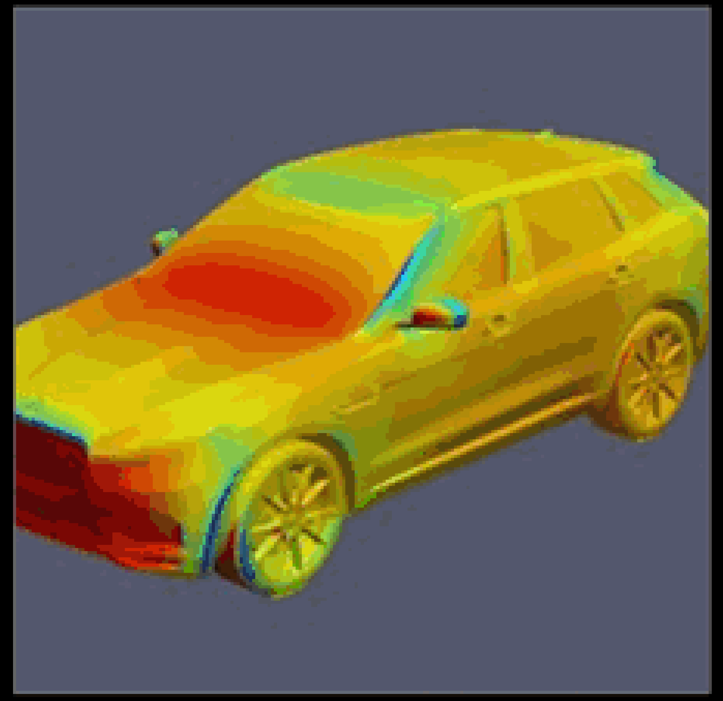

CAD, CAE, SDA

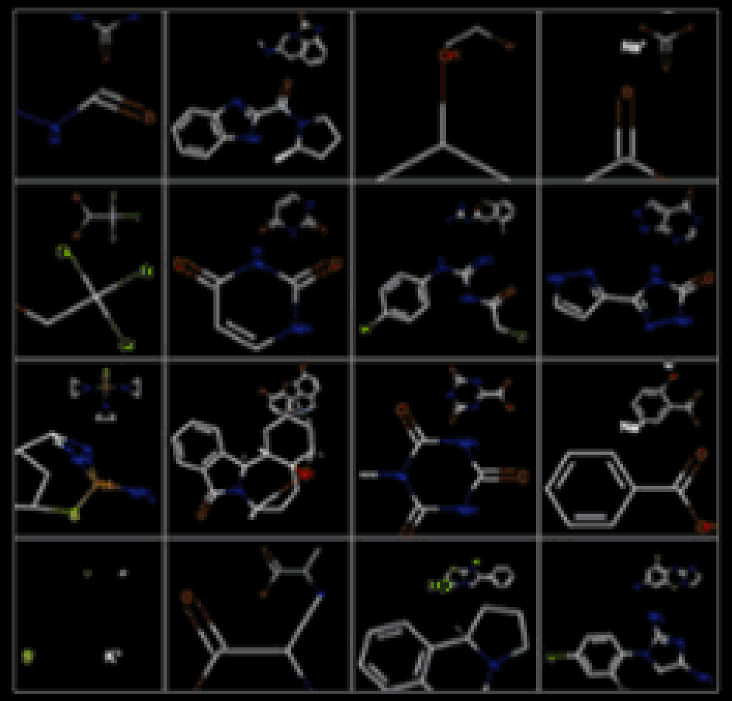

Computer-aided Drug Design

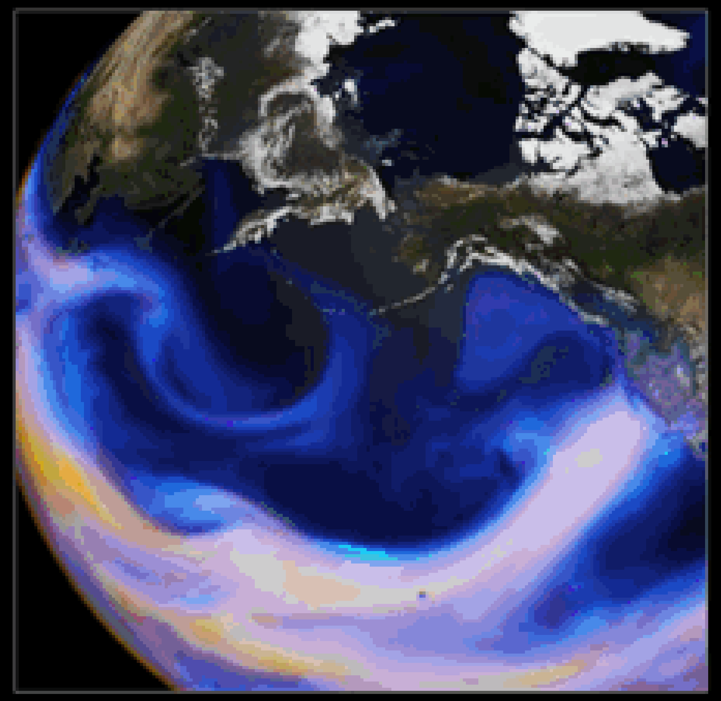

Climate Simulation

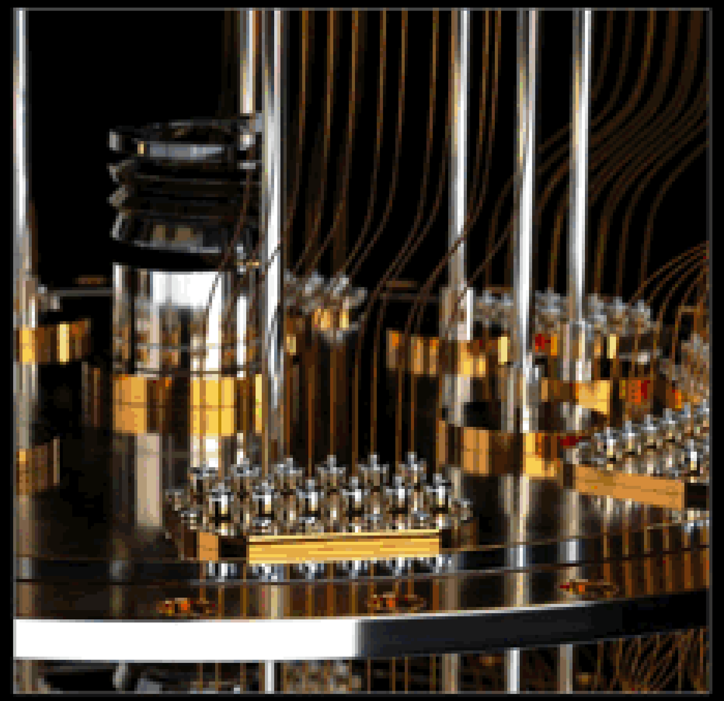

Quantum Simulation

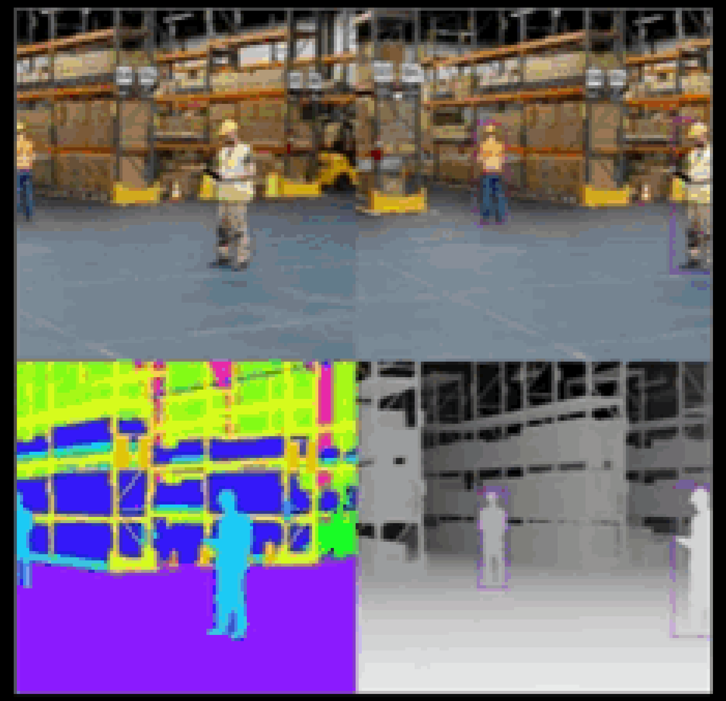

Robotics & Industrial Digital Twins

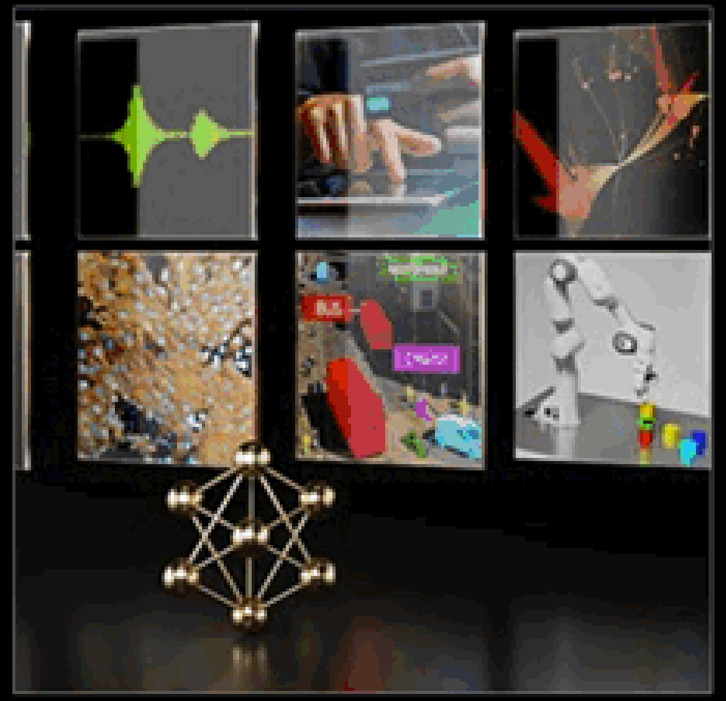

Enterprise AI

NVIDIA AI Enterprise

OMNIVERSE

CUDA-X Libraries

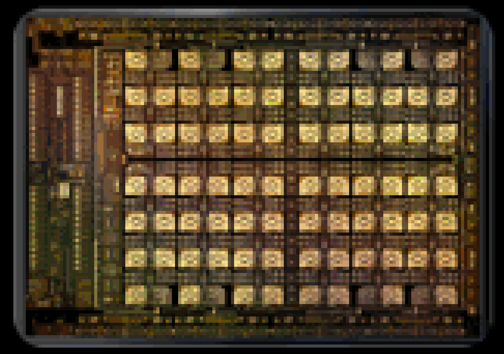

CPU

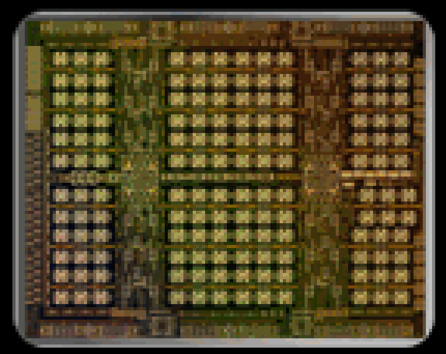

GPU

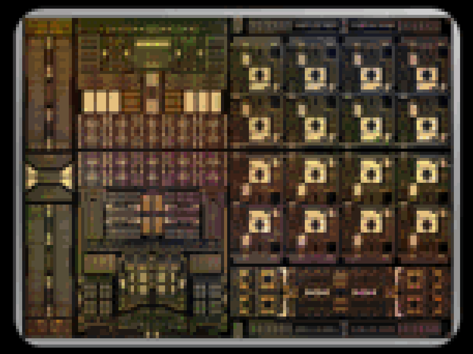

DPU

Accelerated Computing

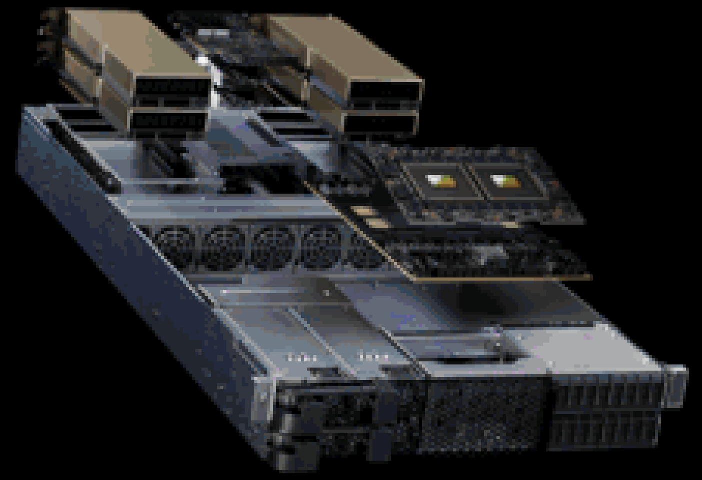

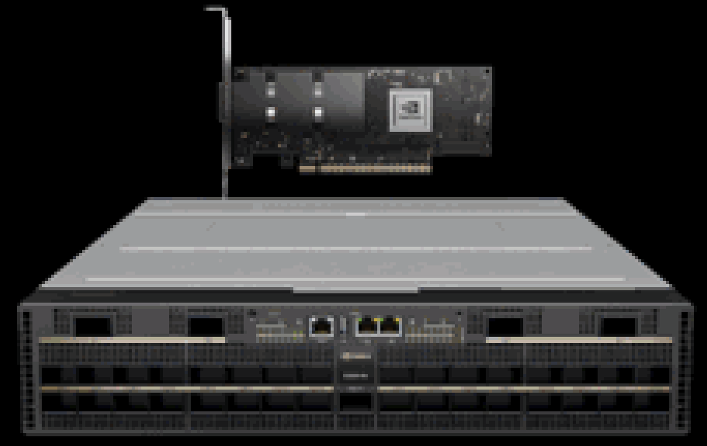

Using AI in Science and Engineering

PhysicsNeMo Case-Studies: Physics AI across different applications

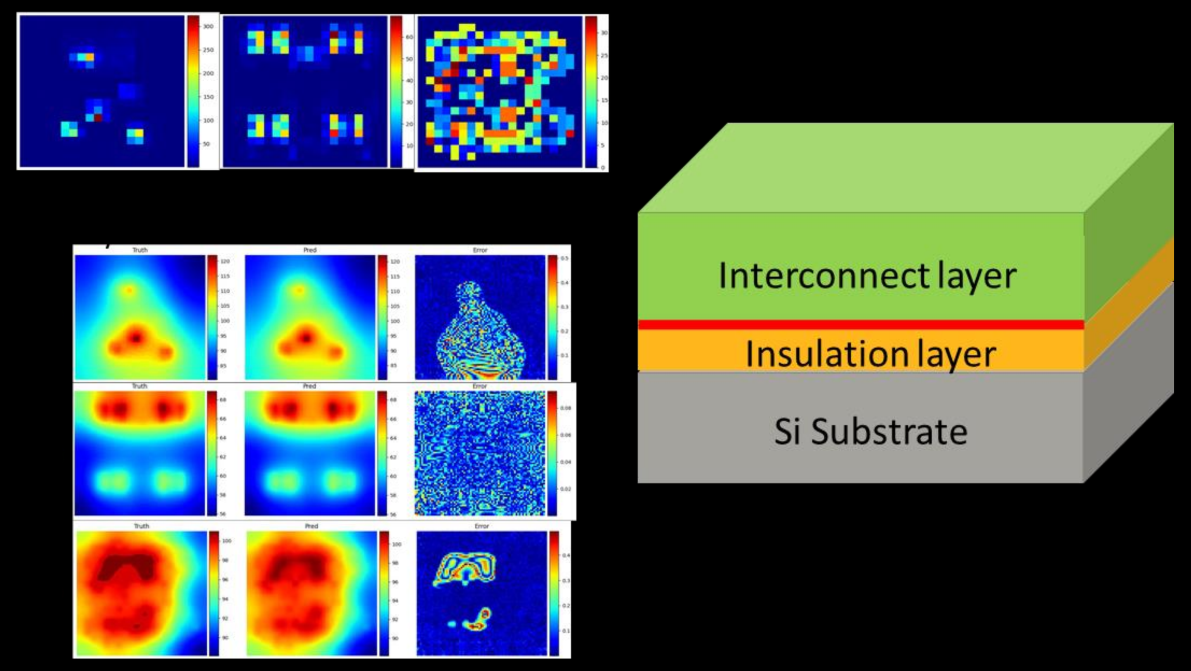

Electro-thermal cooling

Blog: Link

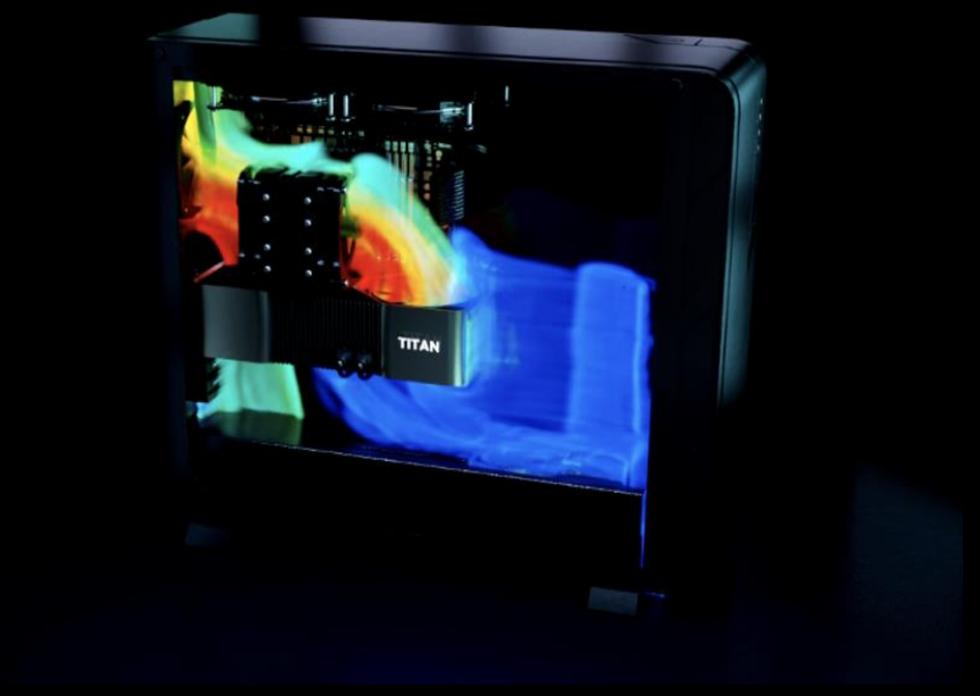

RTX 4090 heat sink design

Demo: Link

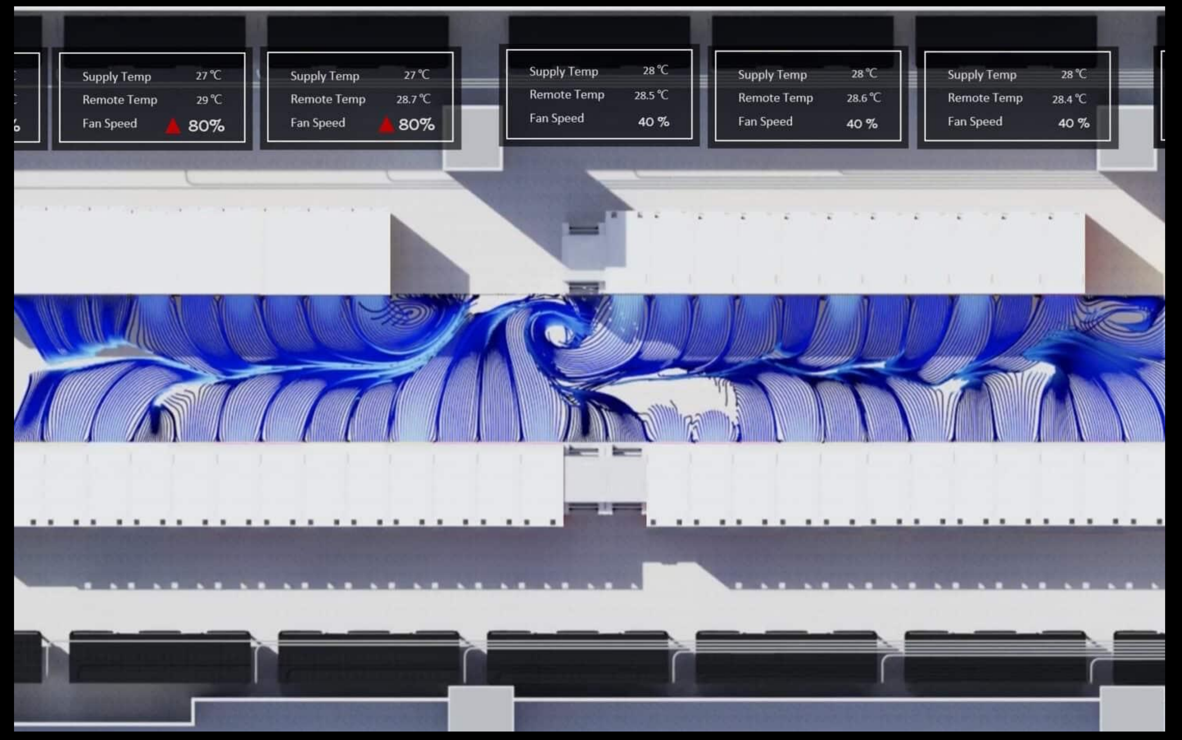

Data Center Digital Twin

Blog: Link, GTC Session: Link

Additive Manufacturing: 3D Printing

Blog: Link

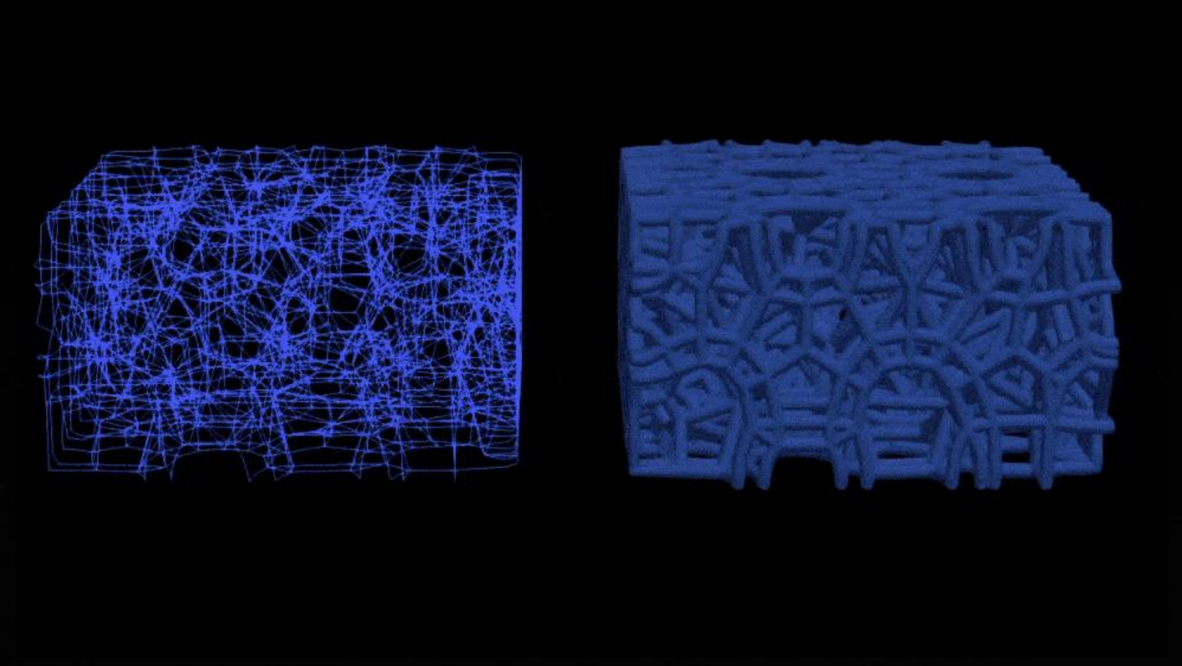

Additive Manufacturing: Lattice Simulation

Blog: Link

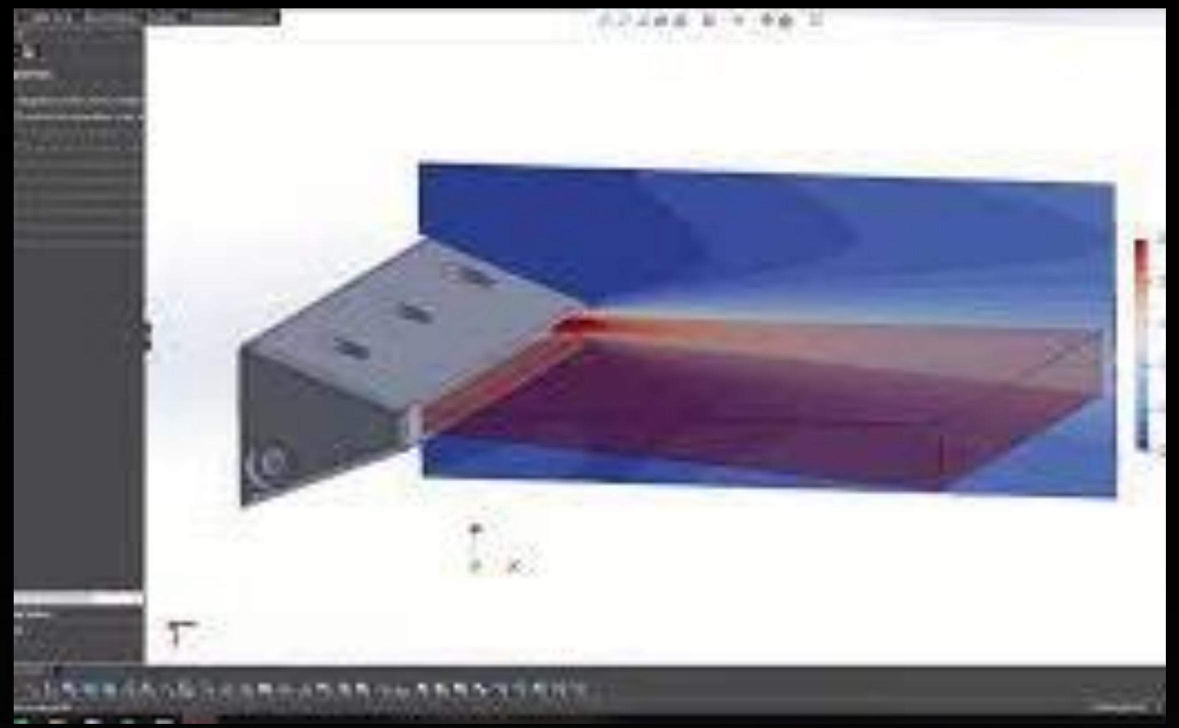

Design optimization

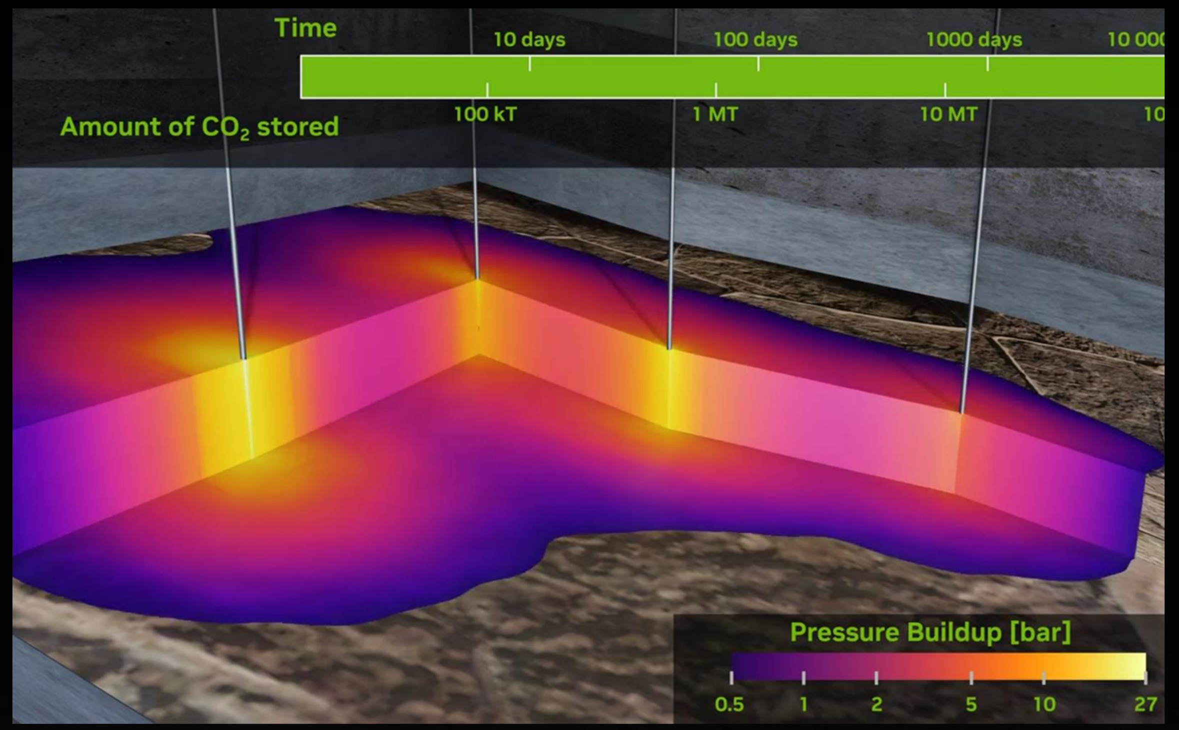

Carbon capture and storage

Demo: Link, Blog: Link

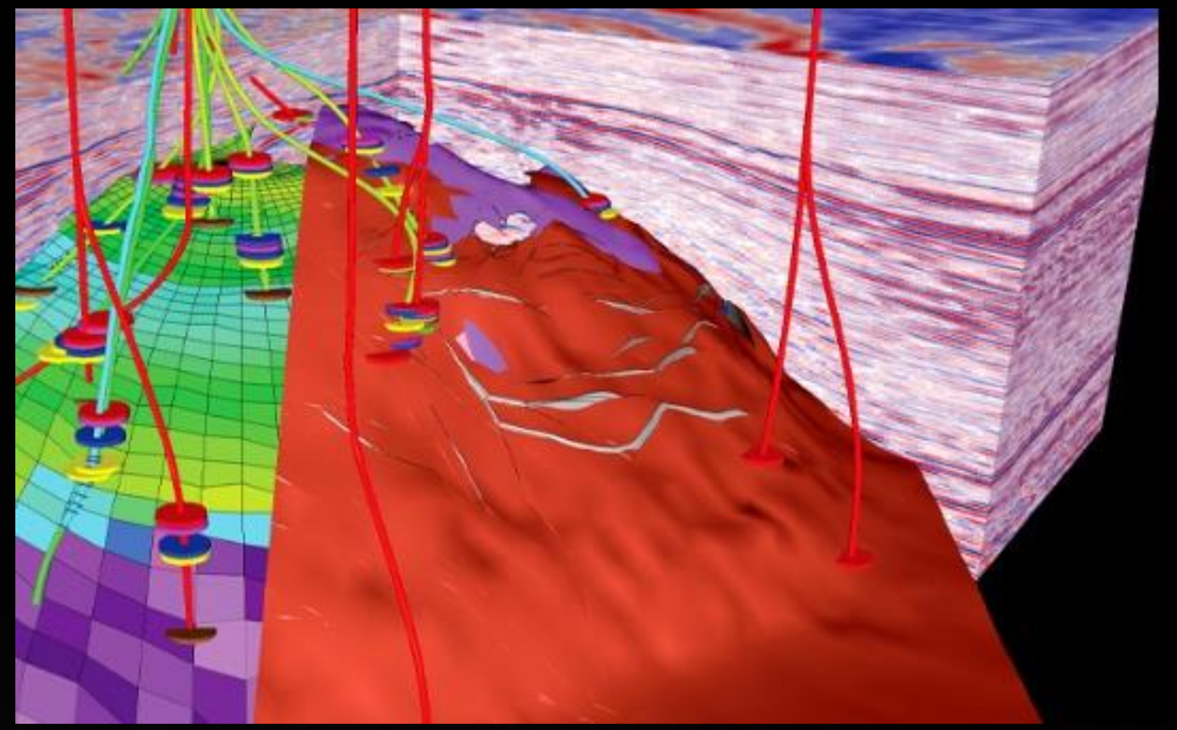

Sub surface simulations

Resource: Link

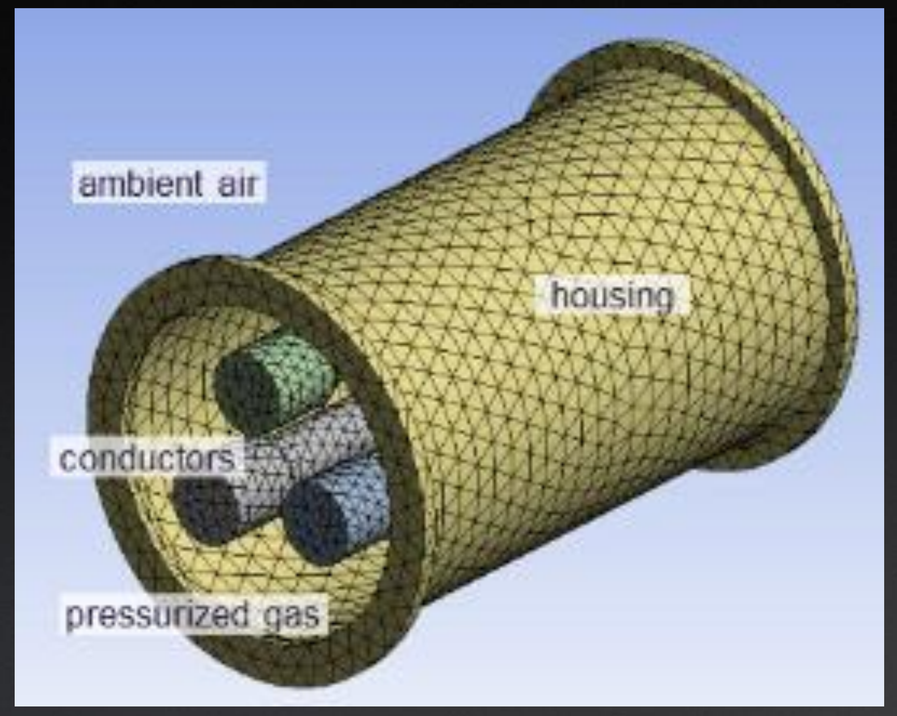

HRSG Digital Twin

Blog: Link, GTC Session: Link

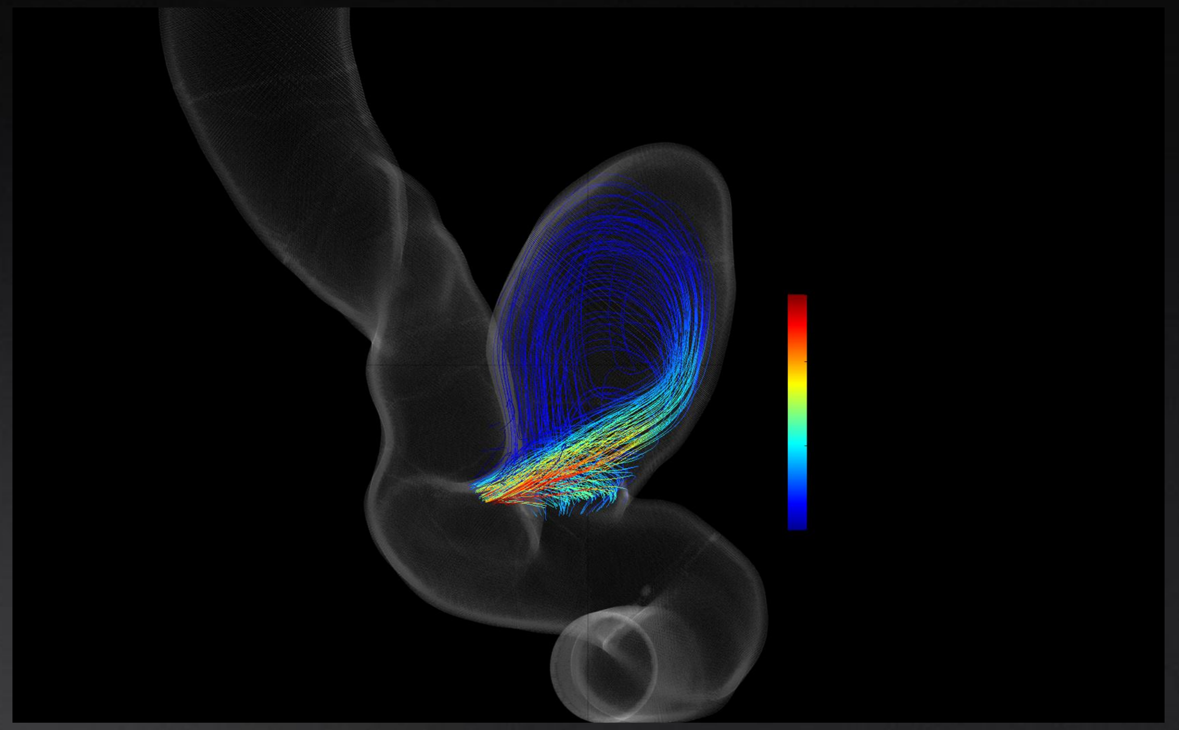

Brain Aneurysm Simulation

Demo: Link

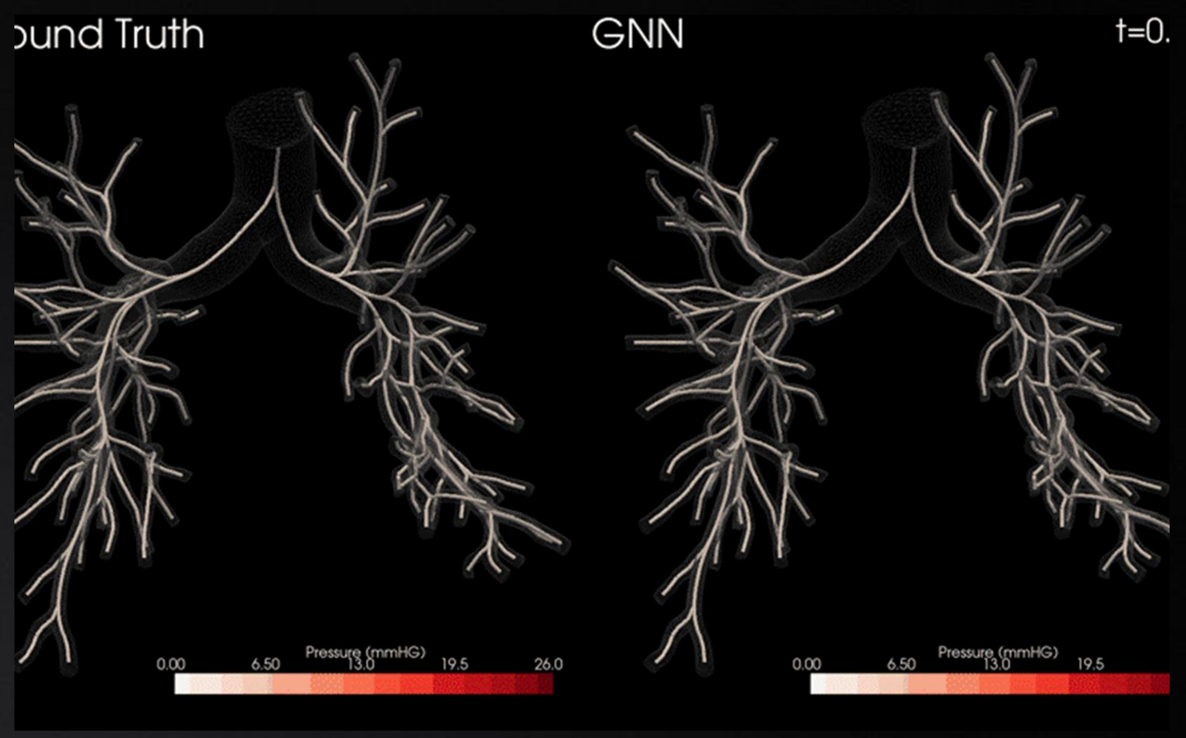

Cardiovascular Simulation

Blog: Link

BioNeMo Framework Supports Optimized Biomolecular Models

Proteins | Small Molecules | Genomics

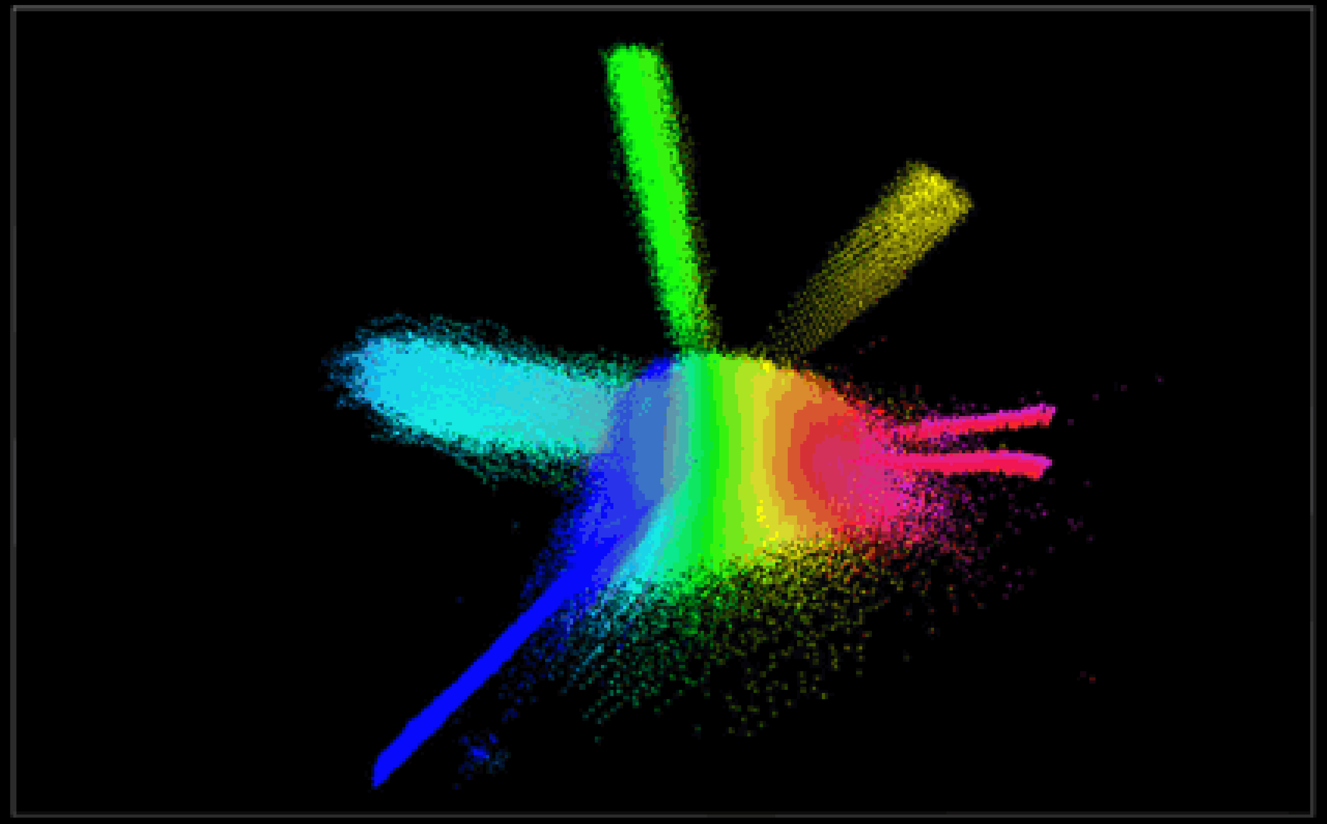

ESM-1 | ESM-2
Protein LLMs

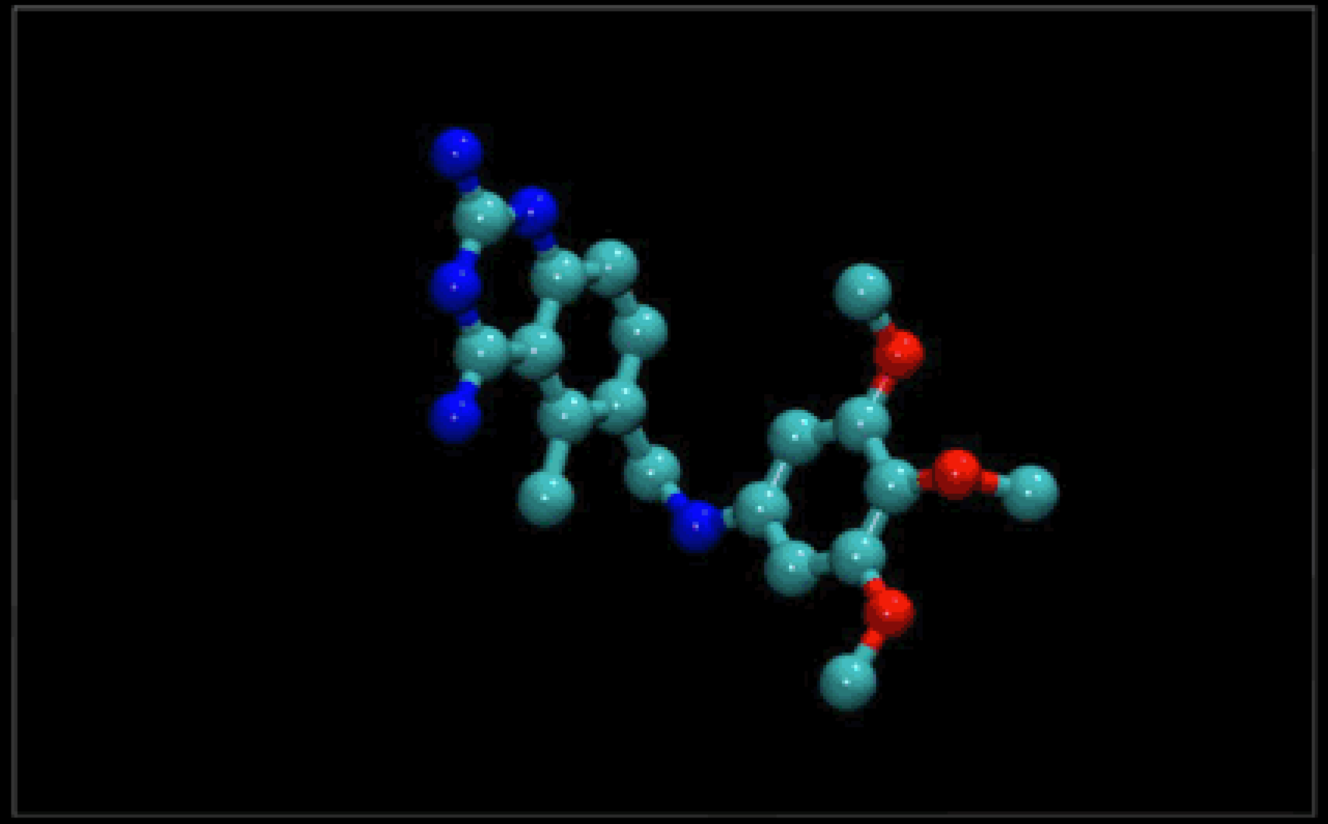

MegaMolBART
Generative Chemistry Model

ProtT5
Protein Sequence Generation

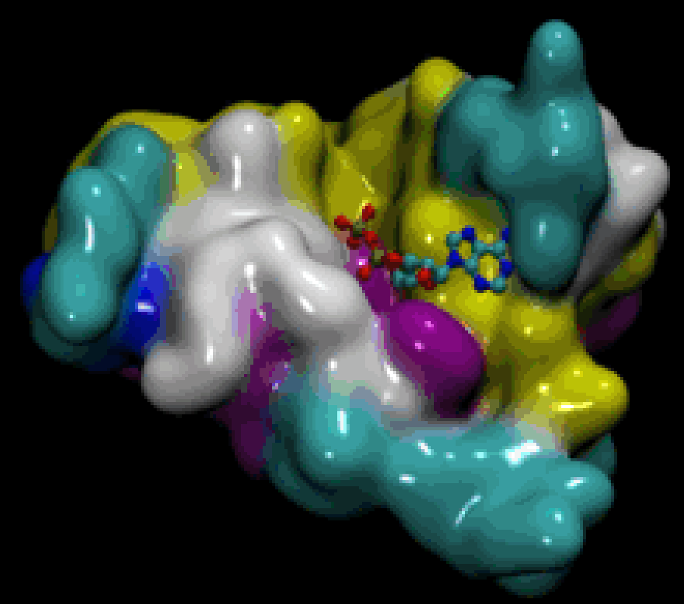

NEW: DiffDock | EquiDock
Docking Prediction

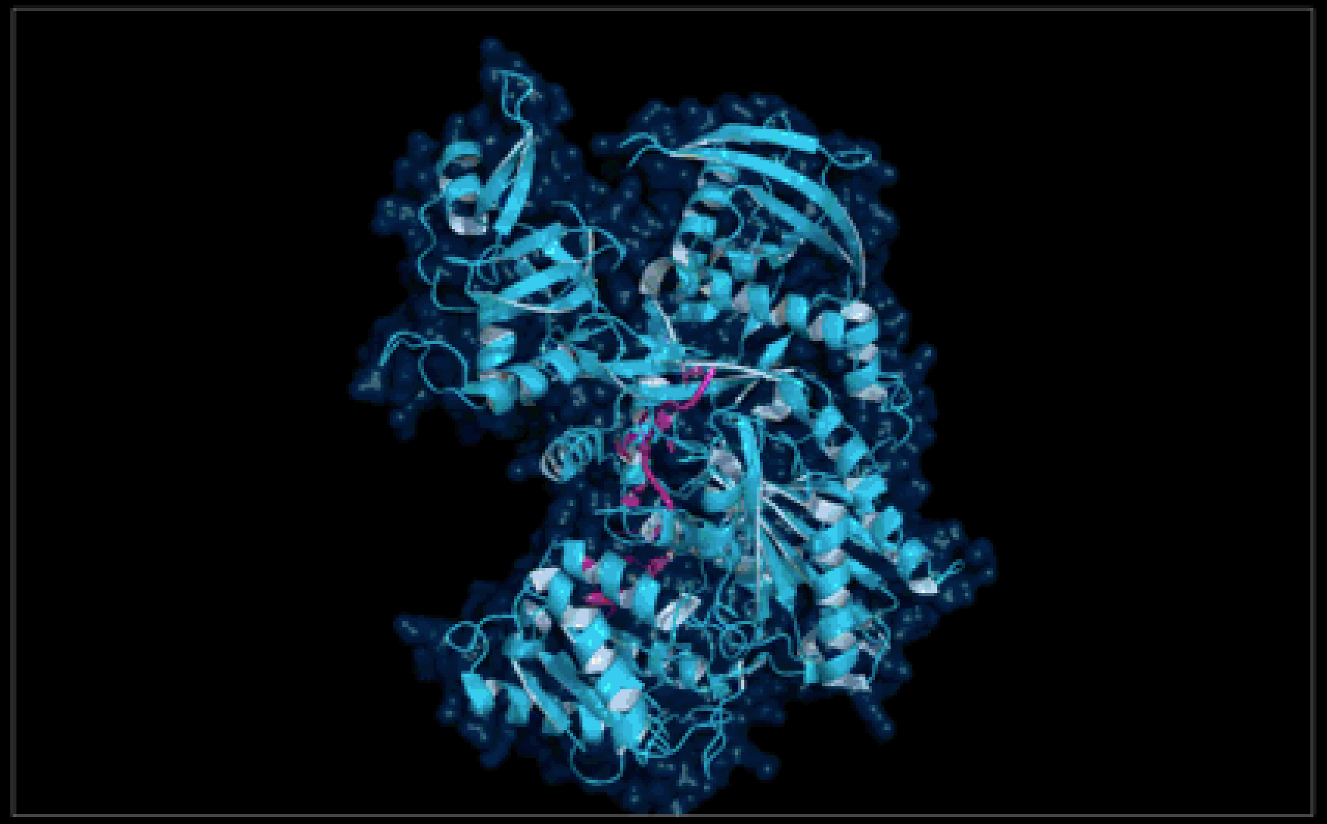

NEW: OpenFold
3D Protein Structure Prediction

NEW: DNABERT
DNA Sequence Model

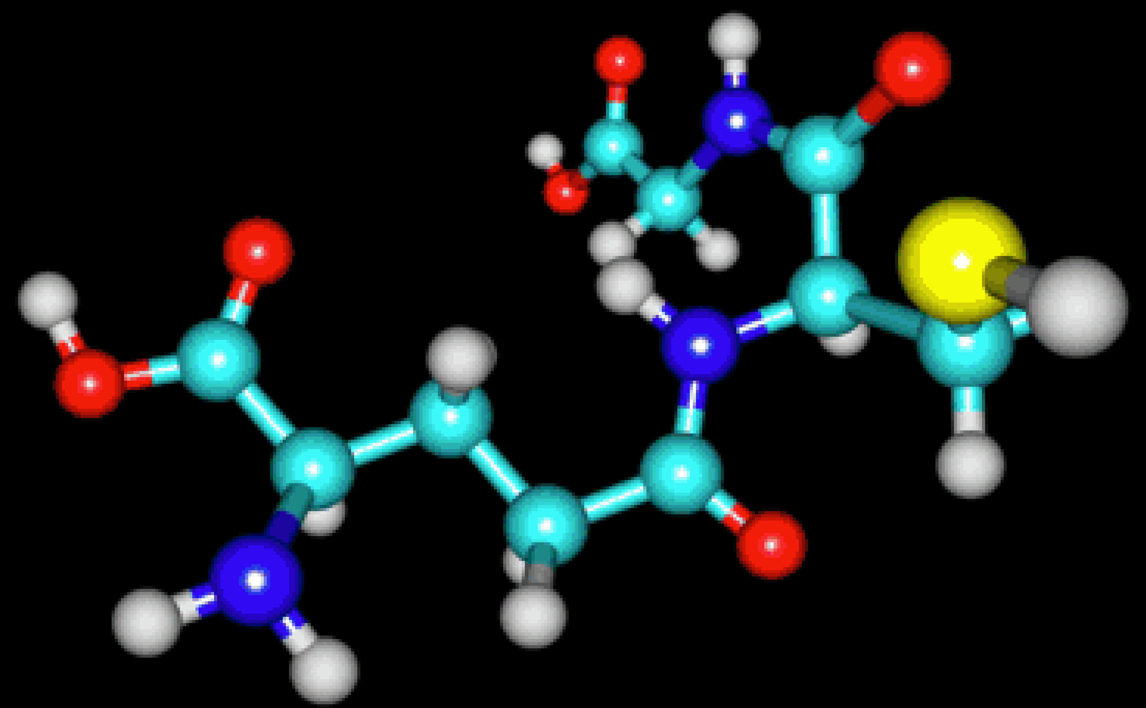

COMING SOON: MolMIM
Molecular Generation

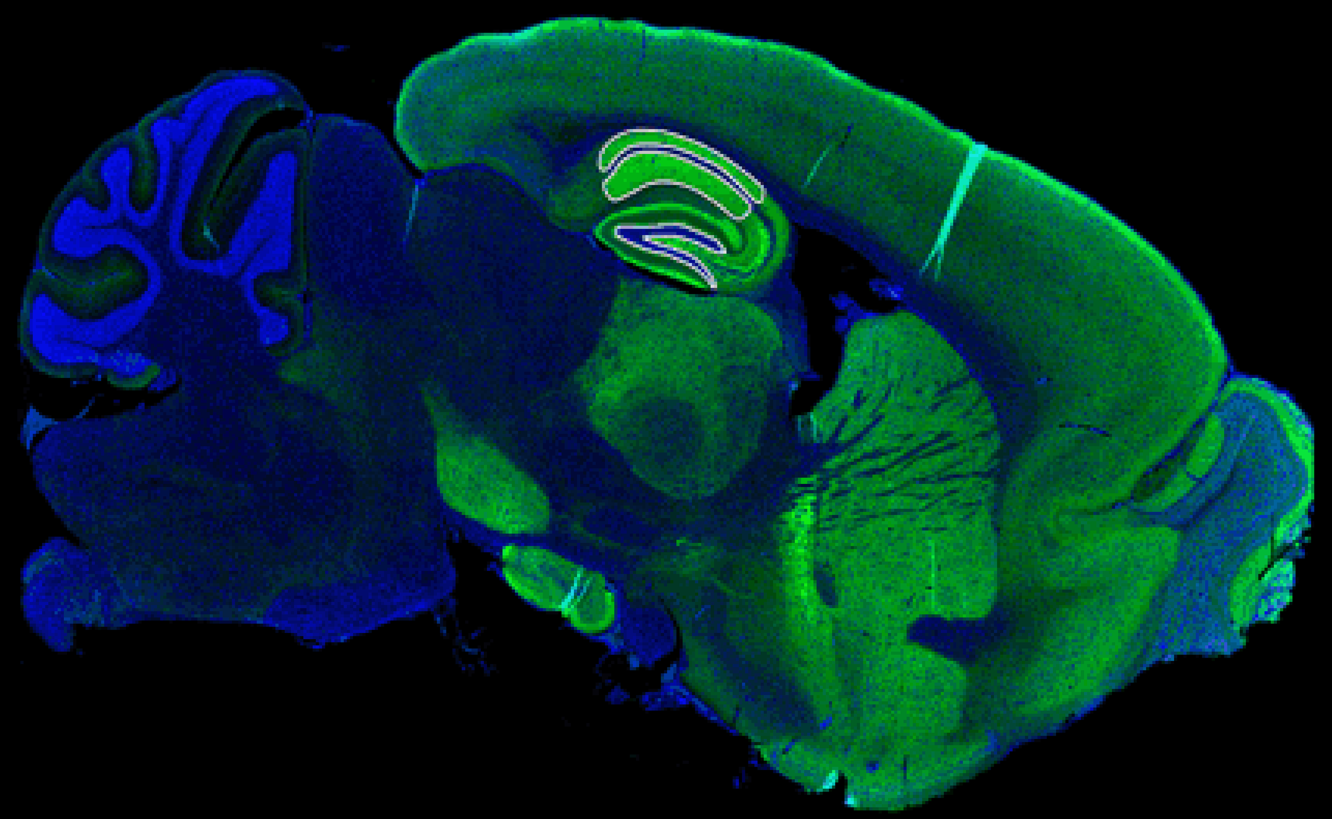

COMING SOON: Single Cell BERT
Single Cell Expression Model

NVIDIA SOFTWARE FOR PHYSICS ACCELERATION

Libraries, APIs, and microservices to facilitate the acceleration of physics workflows

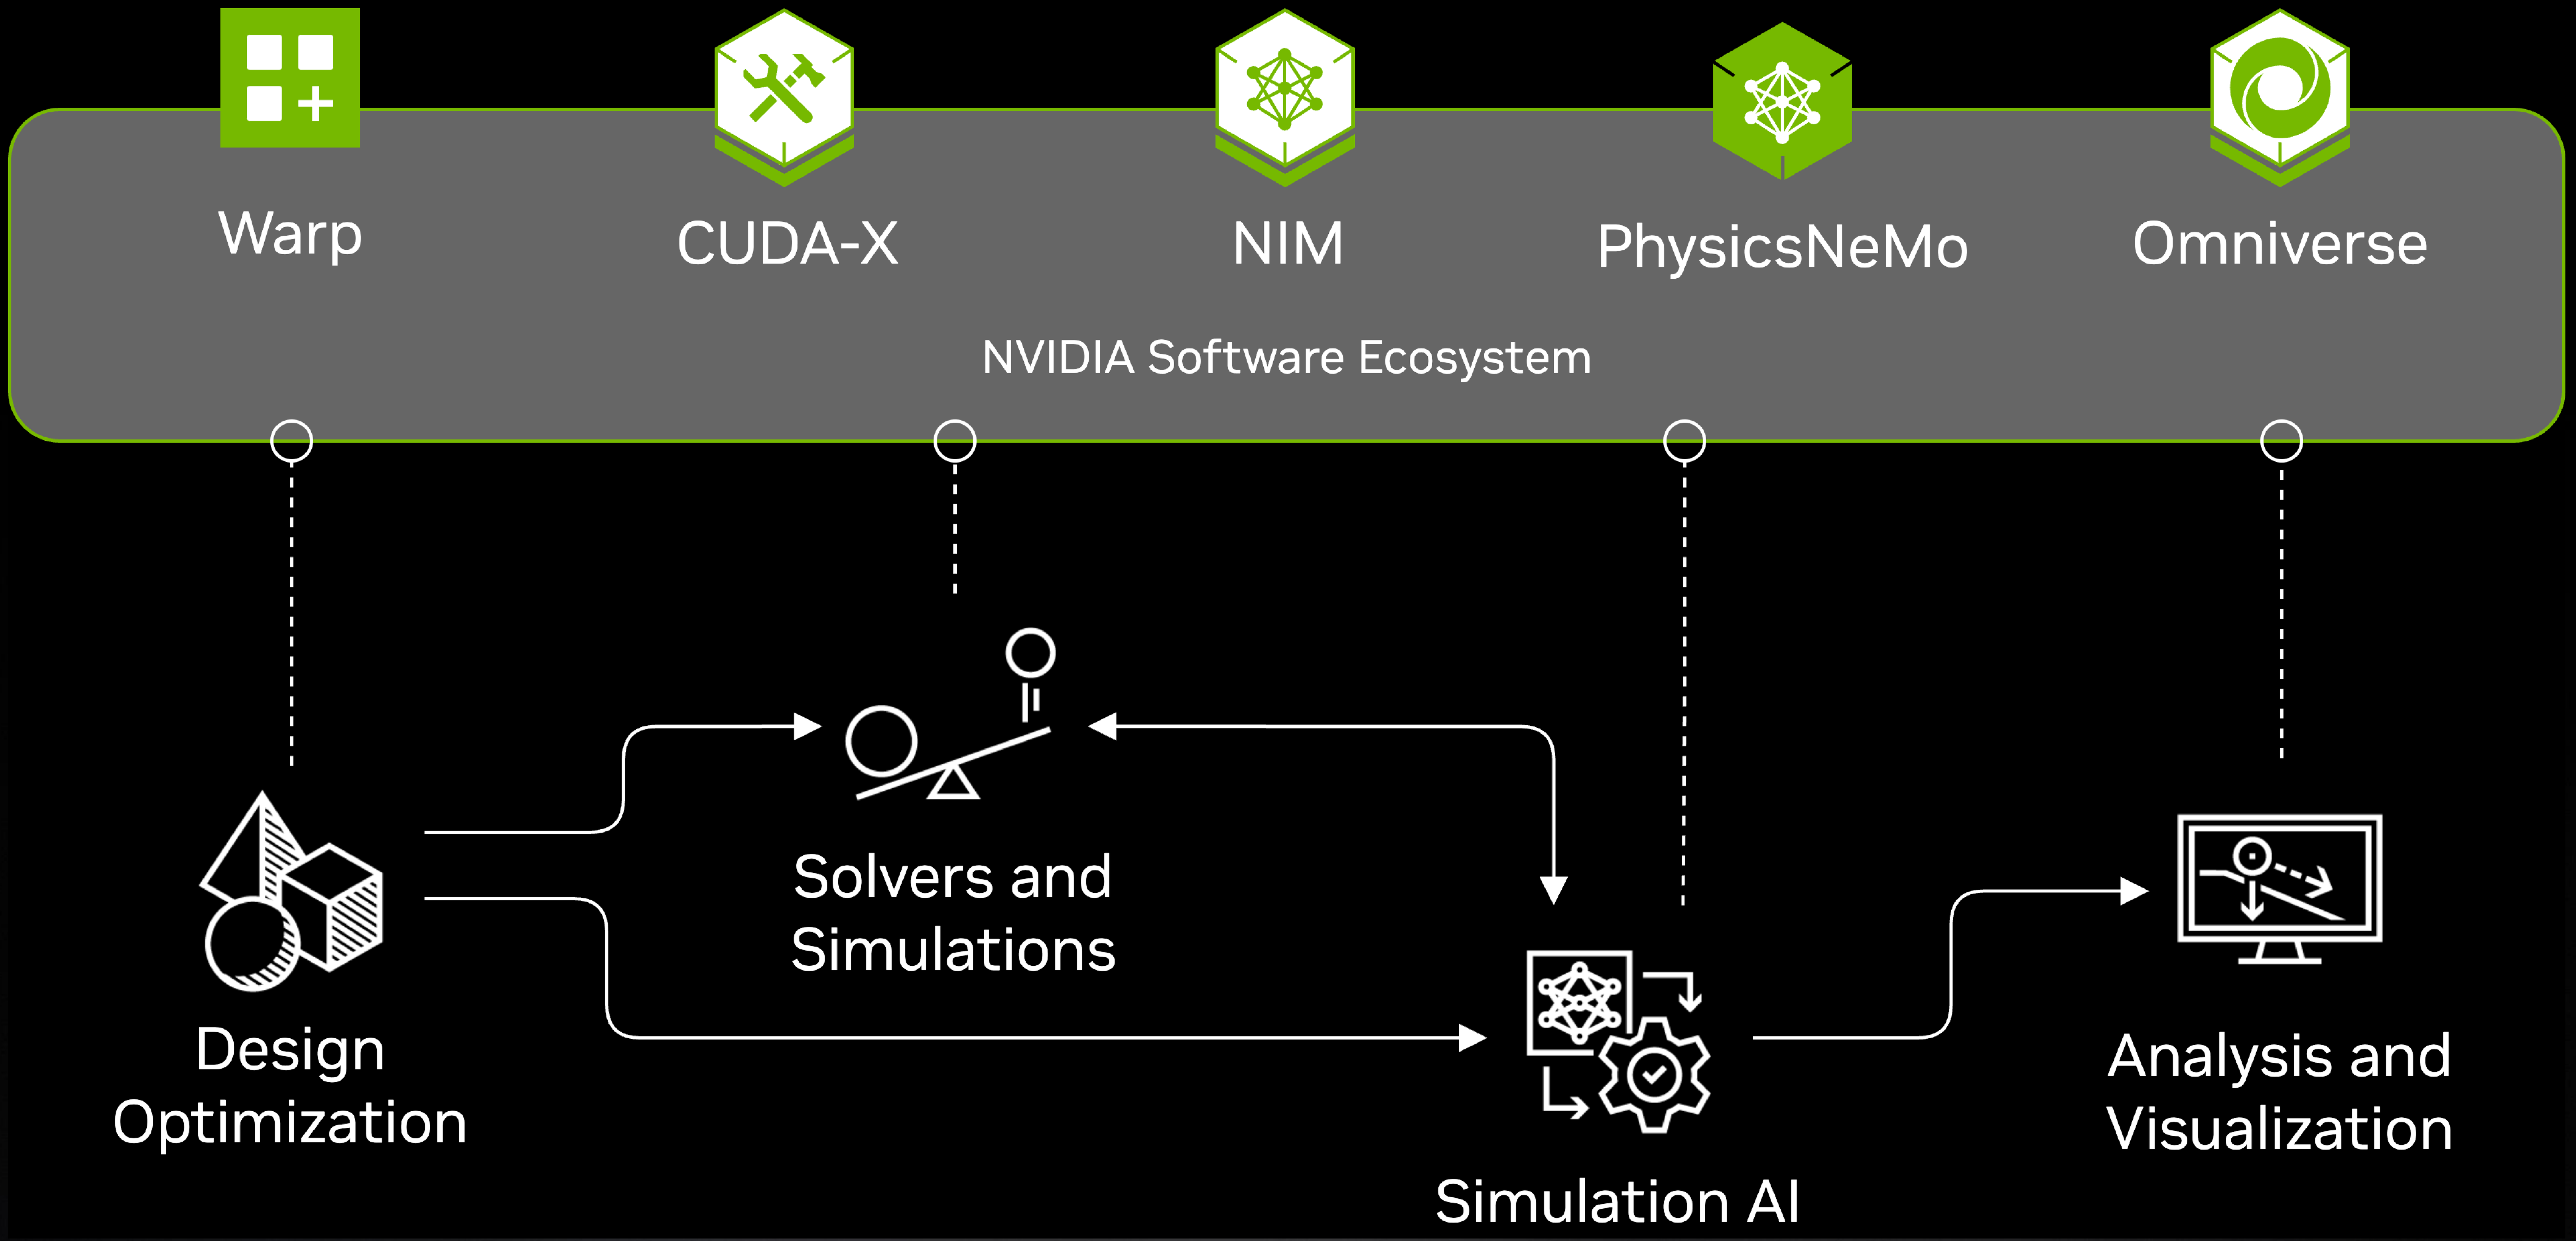

NVIDIA Warp

Purpose-built framework for accelerated simulation and AI

Accelerated Python for physical computing

- Warp is a Python DSL and framework for writing GPU-accelerated and differentiable kernels
- Kernel-based programming is often a more natural fit for routines found in simulation and geometry processing

Create scalable physical simulations and AI training pipelines

- Multi-GPU acceleration
- Integrated with PhysicsNeMo, Newton, Omniverse, and CUDA-X
- Native simulation data structure & algorithms and FEM module

Broad adoption across the CAE and Robotics ecosystem

- Autodesk, Amazon, Google/DeepMind, Siemens
- Applicable to every stage of the CAE/EDA workflow: **design → simulation → analysis & AI training**
- Applicable across robotics workflows and 3-computer system: **simulation → training → Inference on edge**

Open Source

- Links:
  - Repo: https://github.com/NVIDIA/warp
  - Docs: https://nvidia.github.io/warp/

```
pip install warp-lang
```

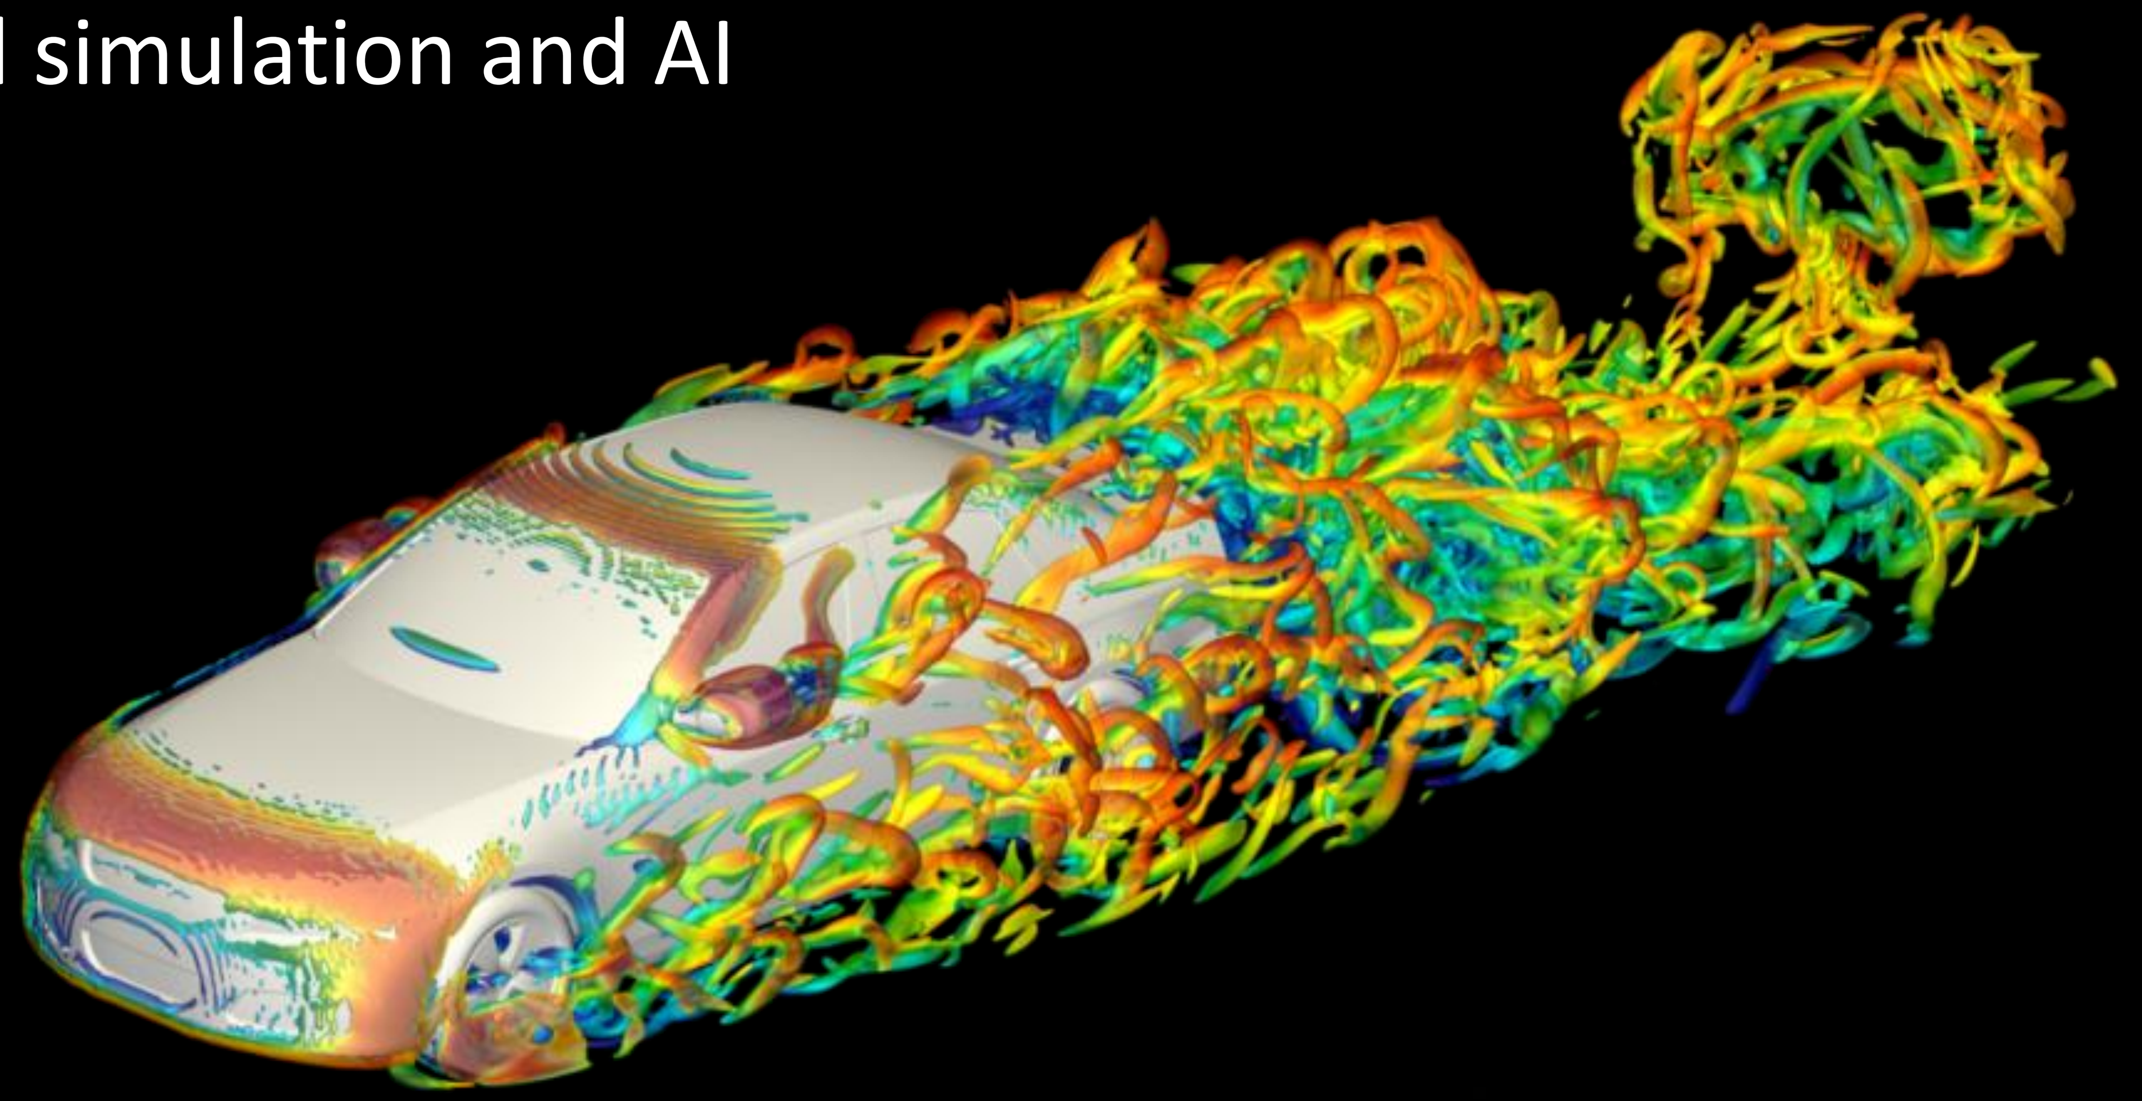

Distributed and Differentiable Fluid Dynamics with Accelerated Python

MuJoCo-Warp & Newton

NVIDIA CUDA-X Math Libraries

Accelerated Numerical Solvers

NVIDIA Warp

Accelerated physical computing

- A developer framework for high-performance simulation, rendering, and data processing
- Enables physics-informed machine-learning pipelines

Create scalable physical simulations and AI training pipelines

- Multi-GPU acceleration
- Integrated into Modulus, Omniverse, and CUDA-X
- Accelerated rendering and NanoVDB integration

Broad adoption across the CAE ecosystem

- Autodesk, Amazon, Google/Deepmind, Siemens
- Applicable to every stage of the CAE/EDA workflow: design → simulation → analysis & AI training

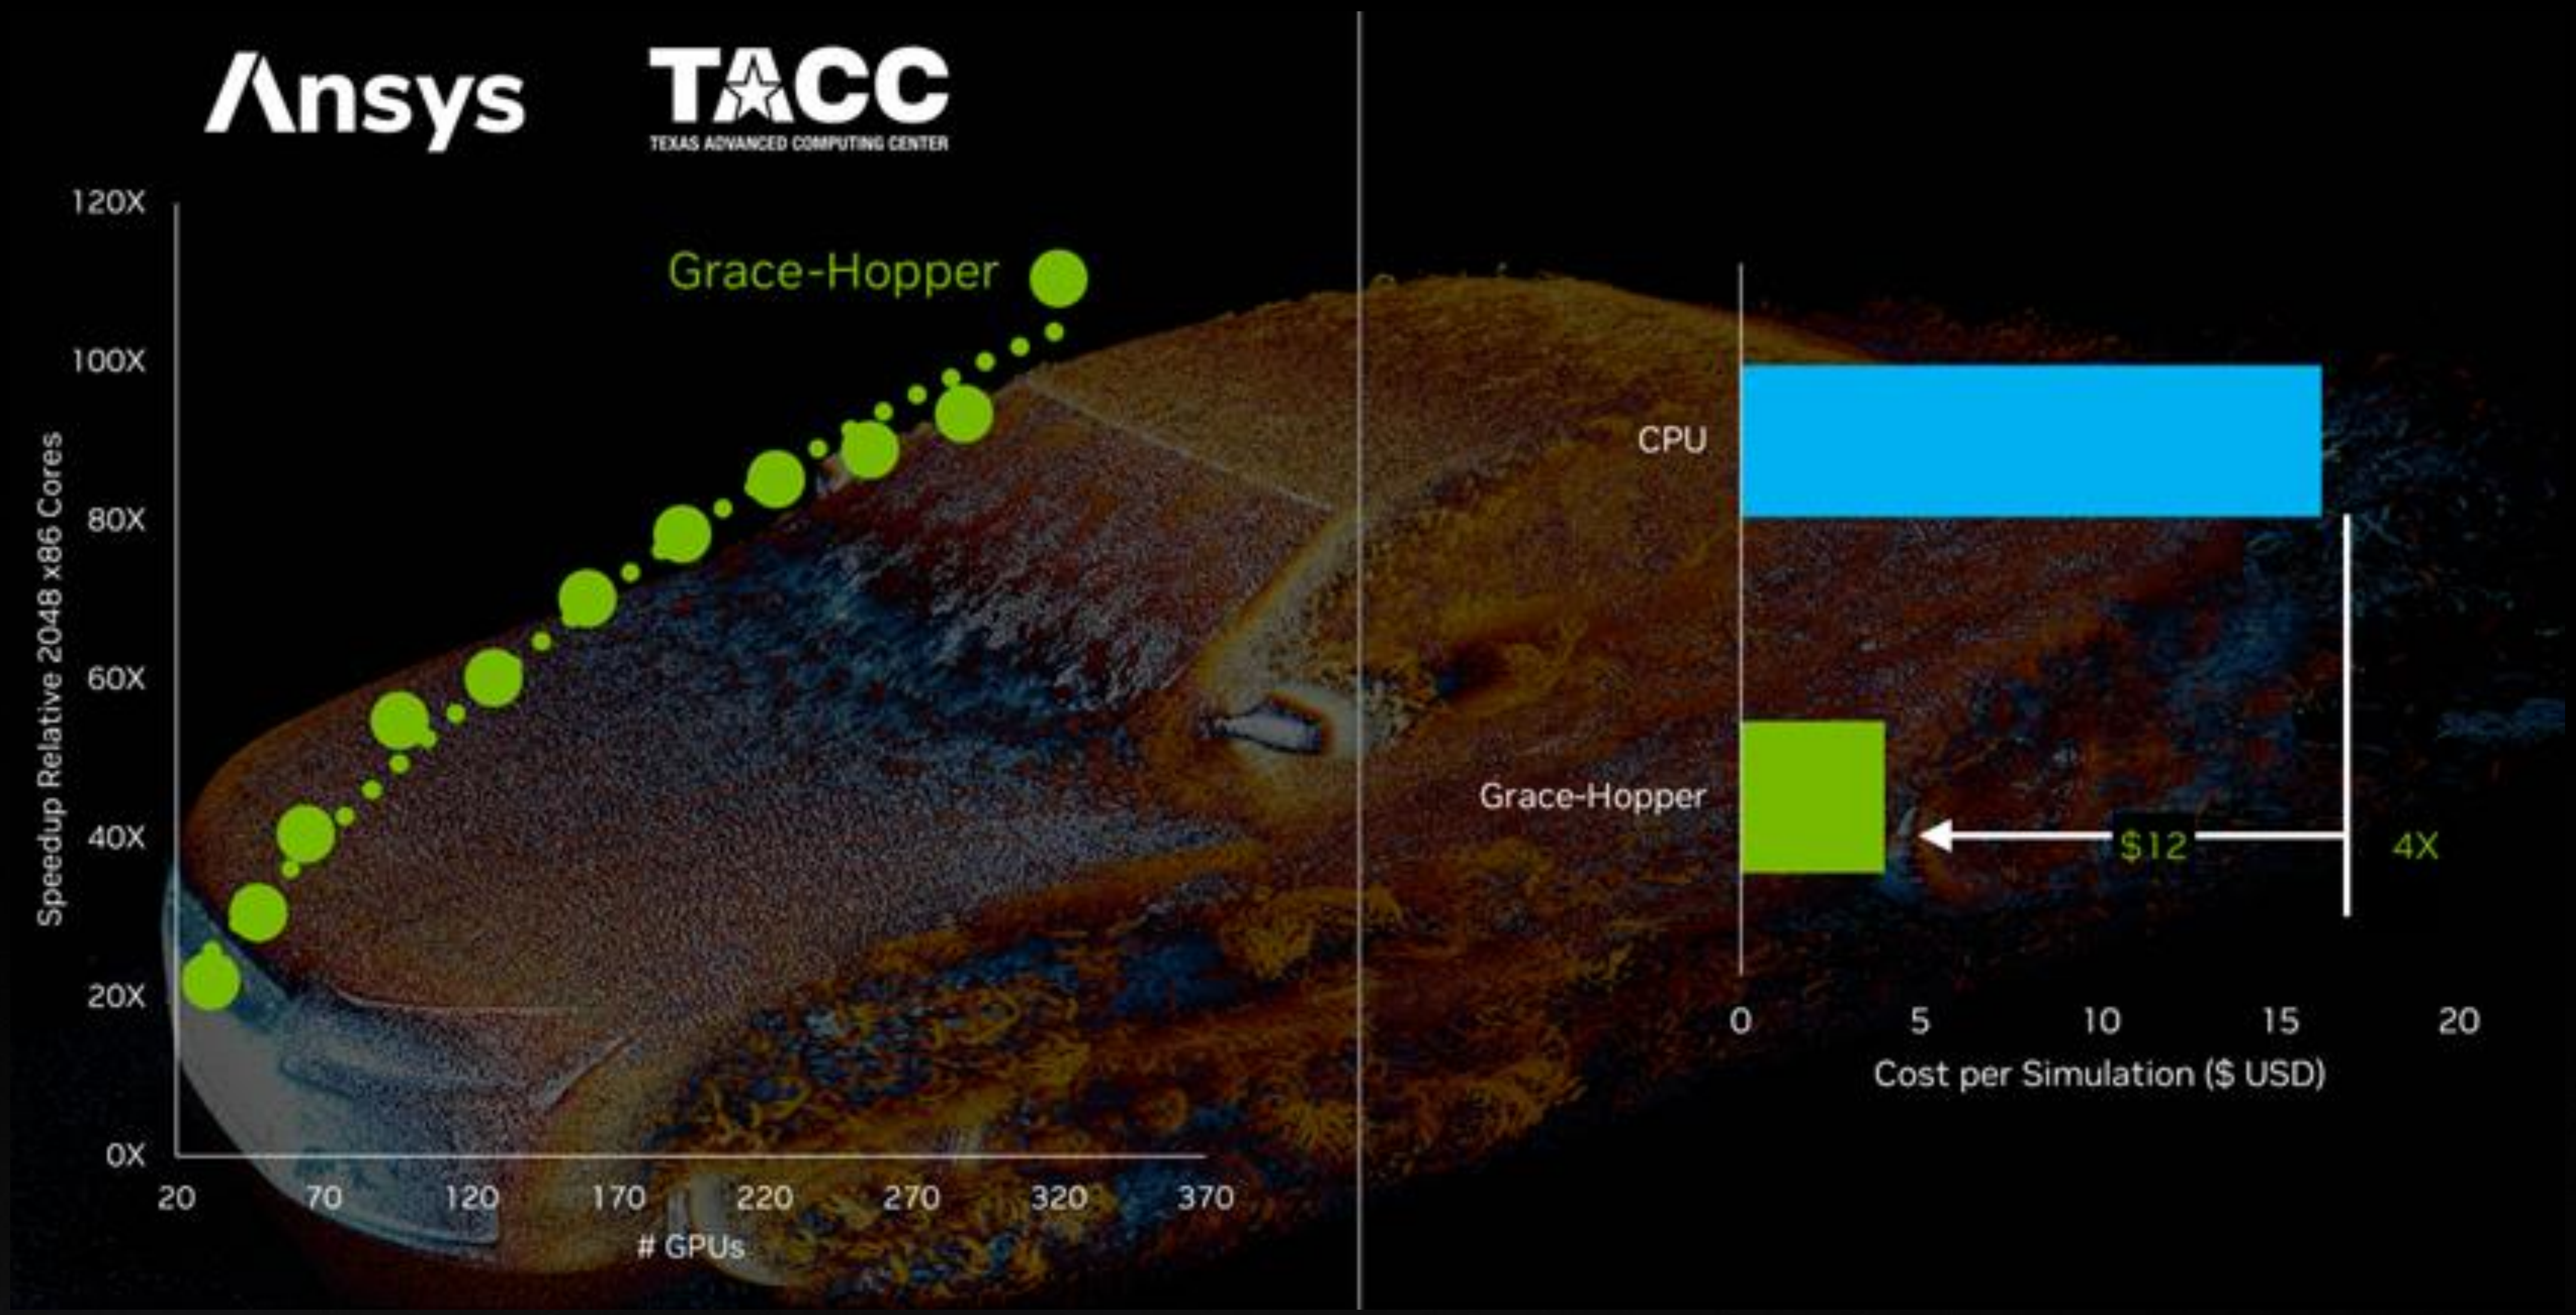

Simulations are a critical for engineering design

Simulations via numerical methods are computationally expensive

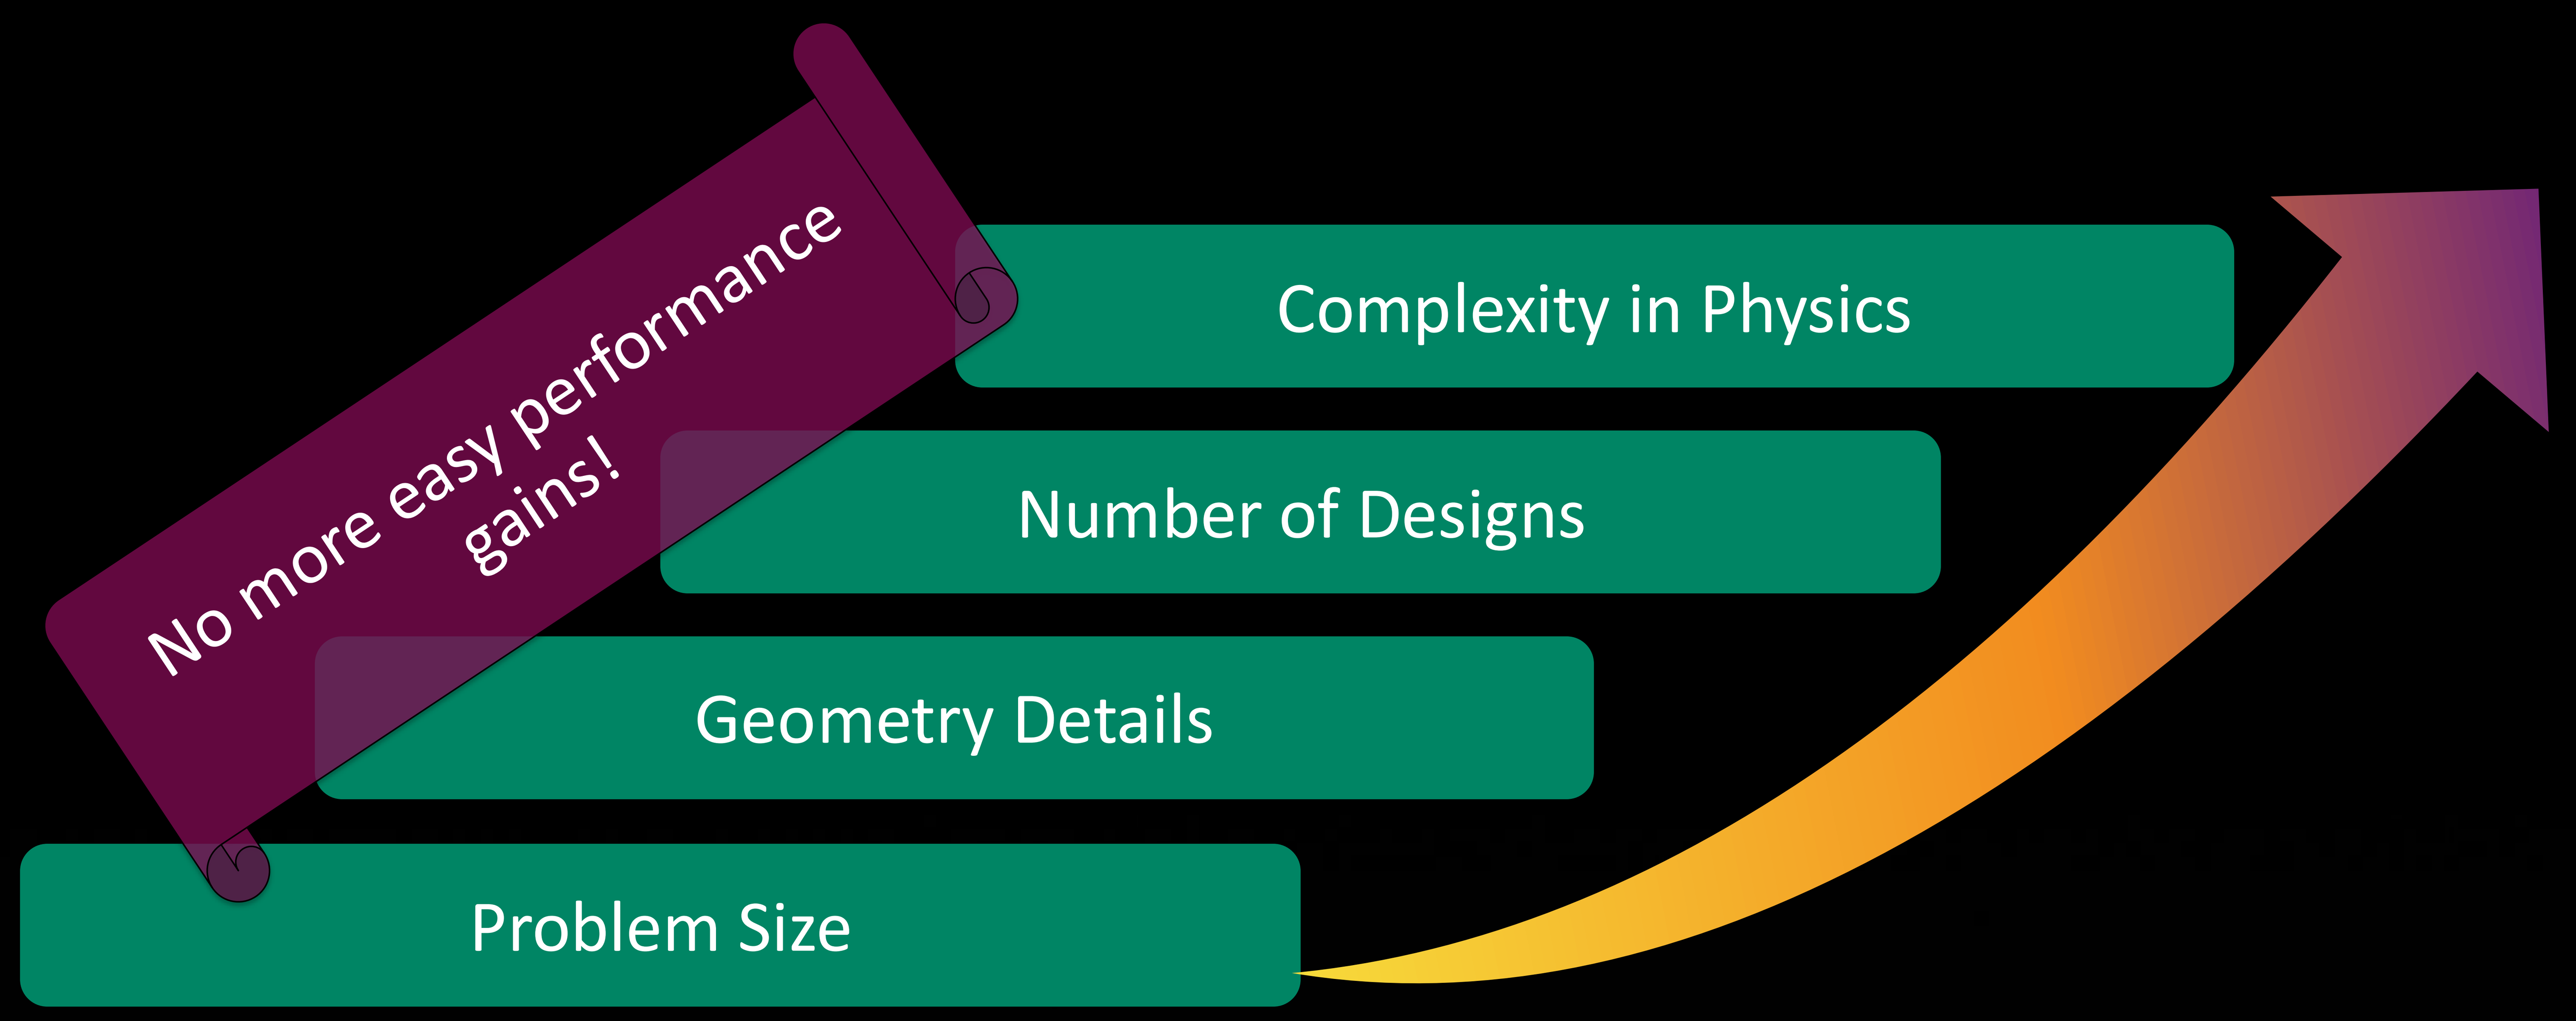

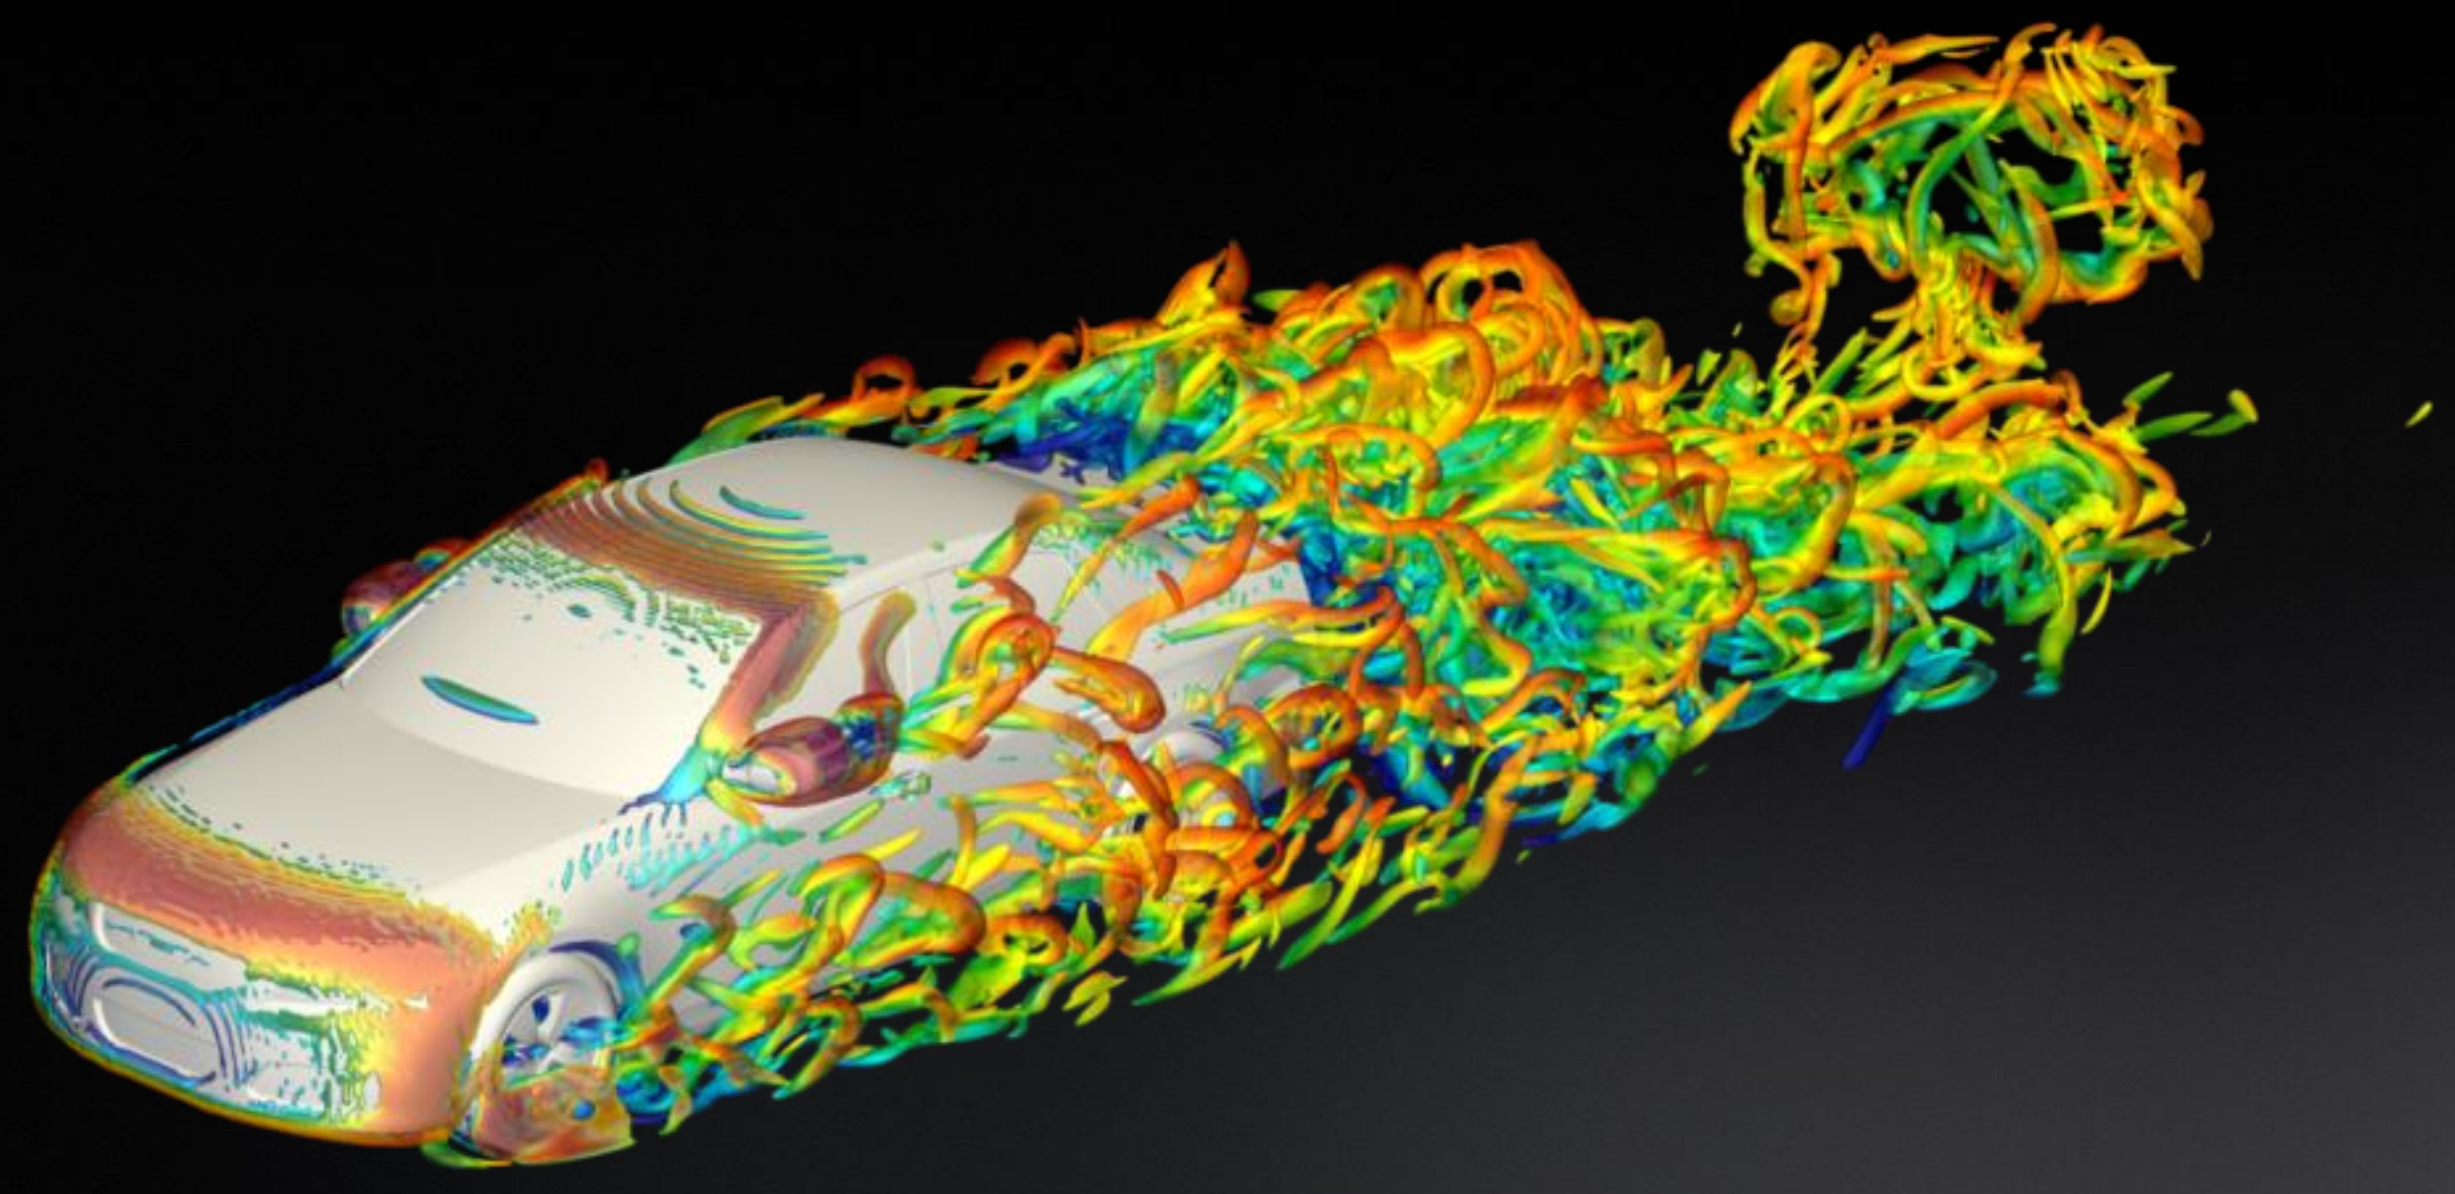

- Require ability to ingest data coming from sensors or prior simulations

Using AI for simulations?

- AI has already disrupted the way we think of computation in other domains and mapping to AI unleashes parallelism
- Doing once vs repetitive – learn once and infer over and over

- Near-real time emulation
- Enable high fidelity simulations
- Representative of the high dimensional geometry and parameter design space

Training: NVIDIA PhysicsNeMo

Open-Source Platform for Developing Physics-Based Machine Learning

Training Neural Networks using both Data and Governing Equations

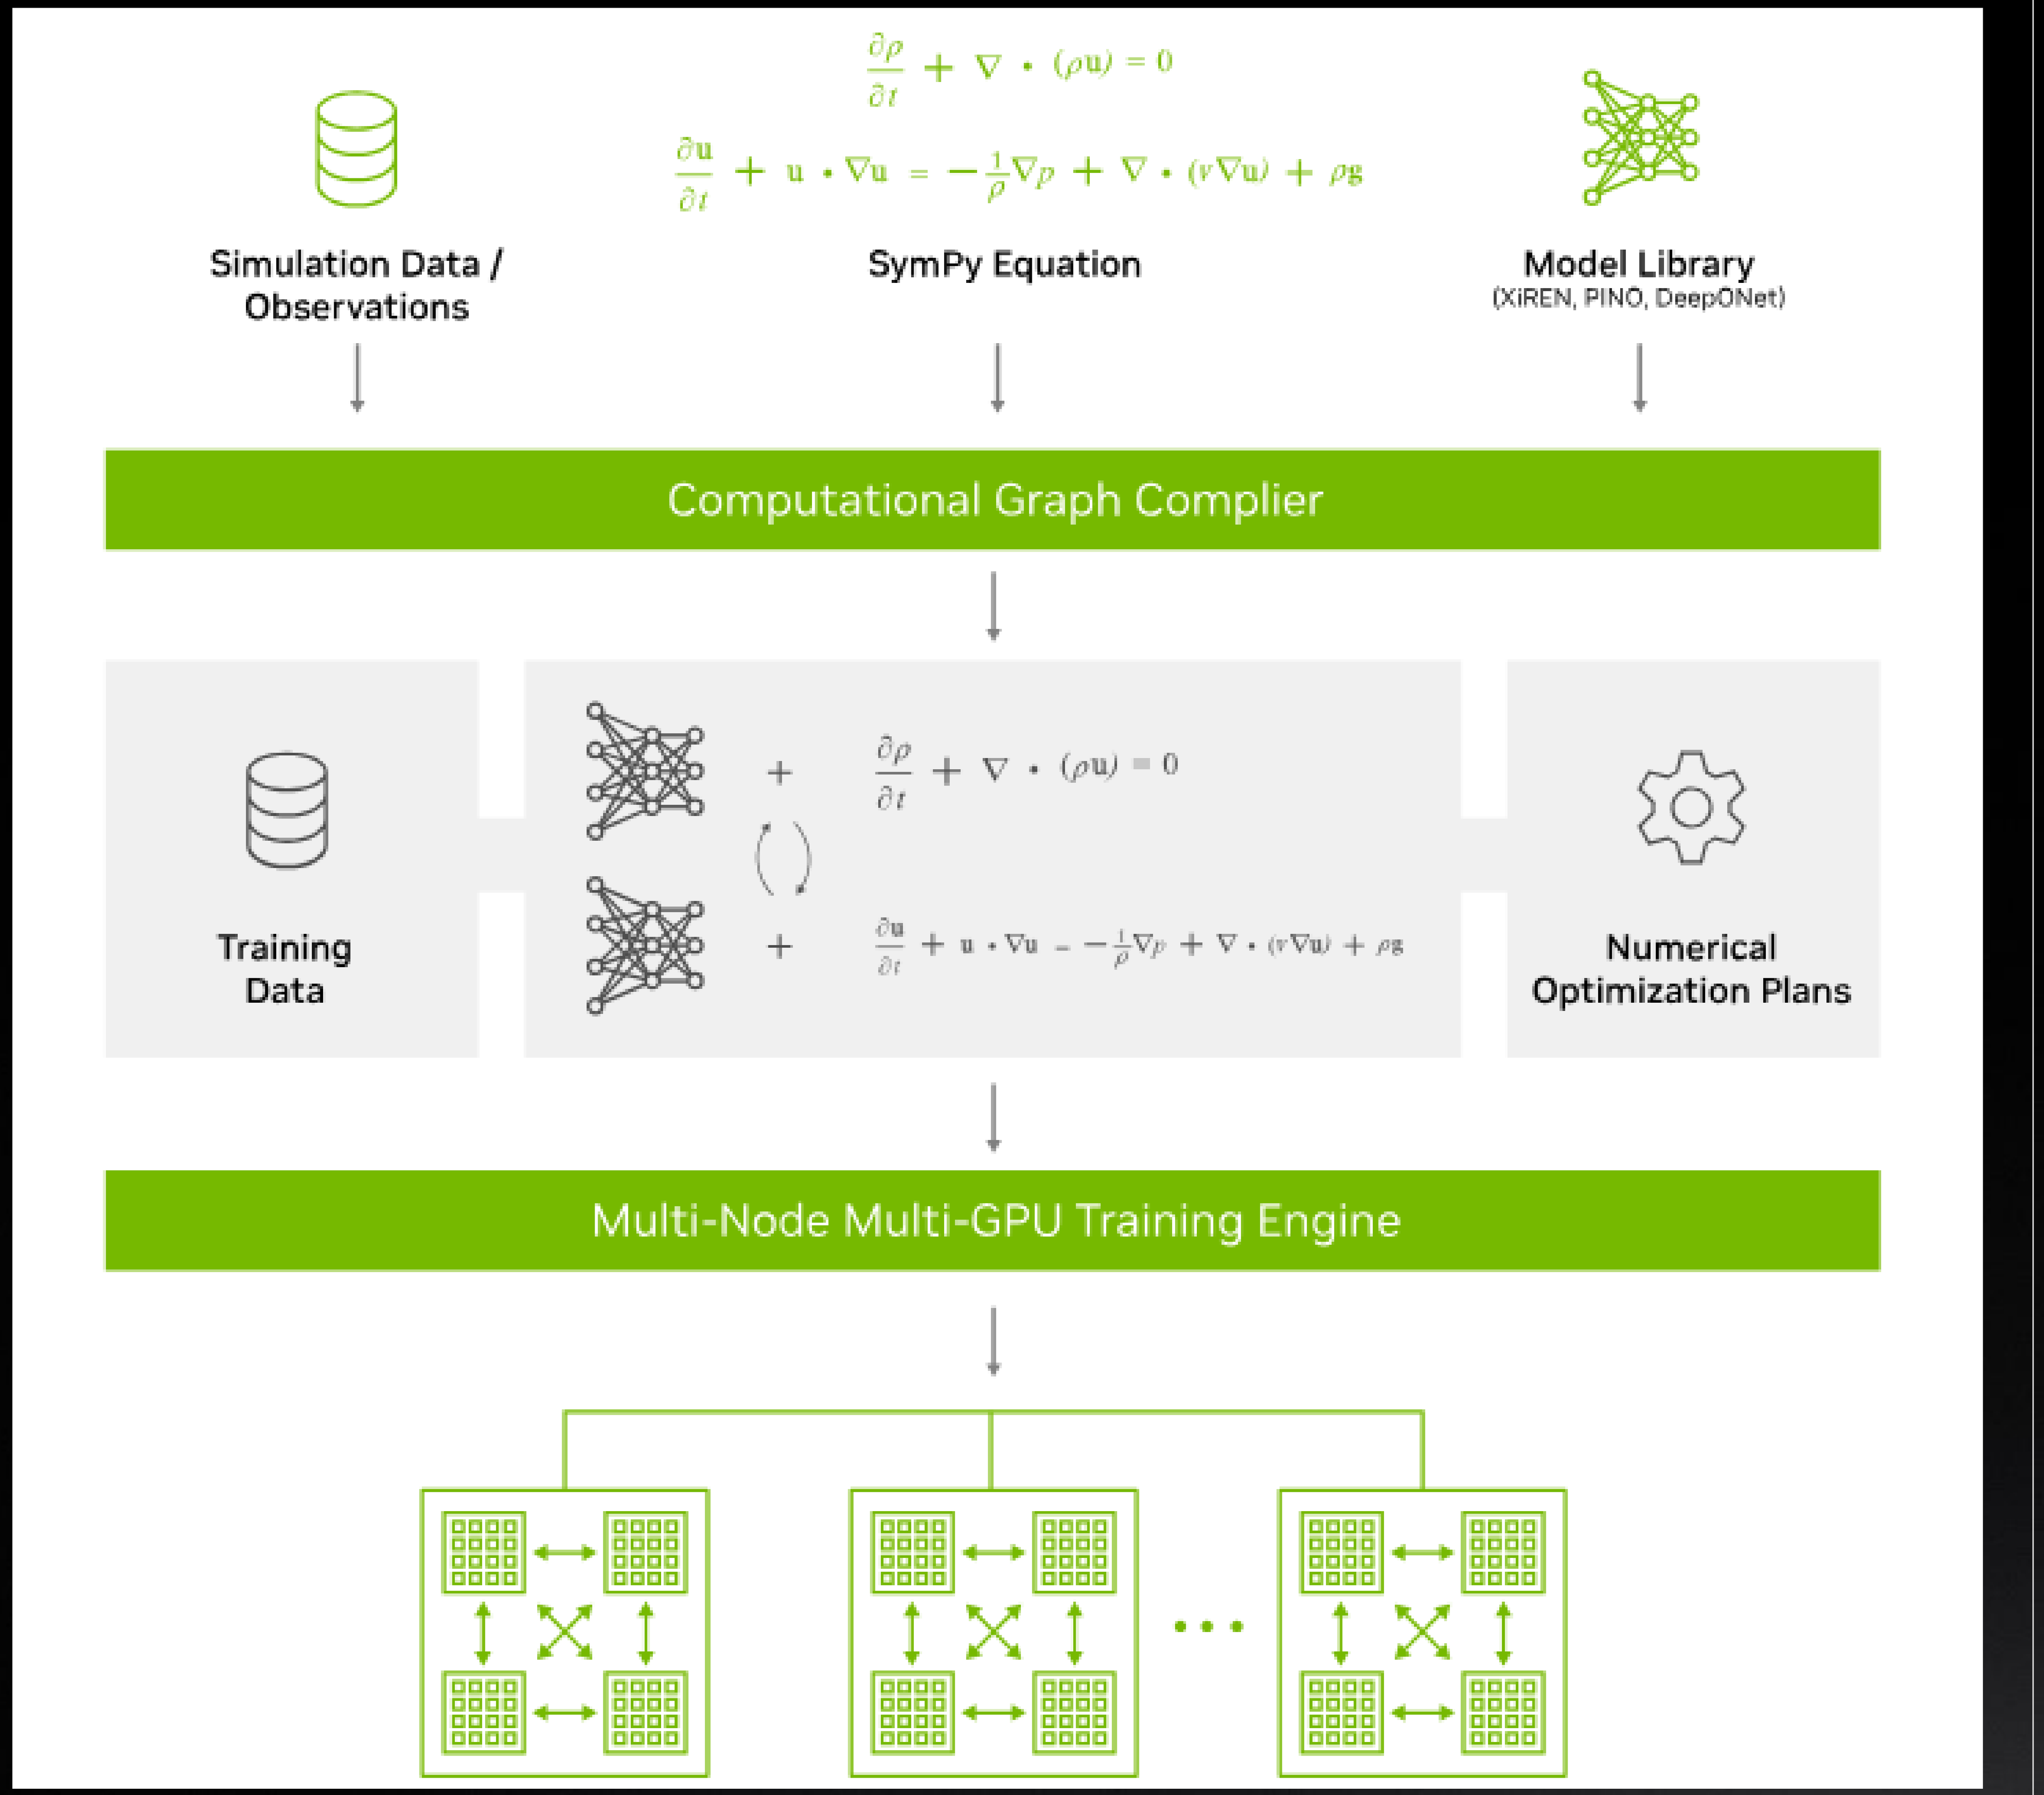

https://github.com/NVIDIA/physicsnemo

Advancing Scientific Discovery With PhysicsNeMo

- Open-source Python toolkit for physics-driven ML
- State-of-the-art architectures and pre-trained weights
- Efficient data loading and preprocessing components
- Easily scalable to multi-GPU, multi-node infrastructure

Renewable Energy
Siemens Gamesa: 4000X Faster wind turbine wake optimization

Climate Change
45,000X Faster extreme weather prediction with AI Weather Models

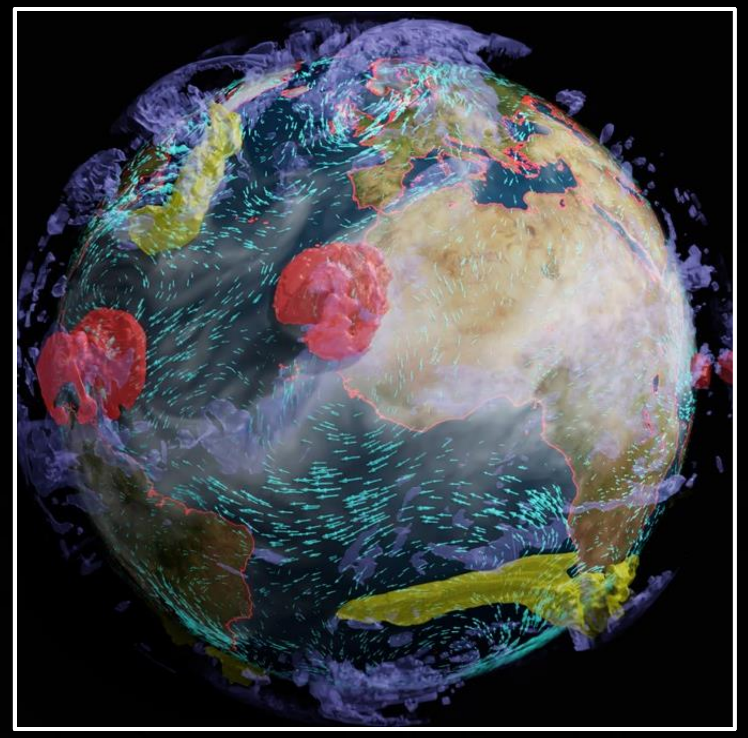

Healthcare
High-fidelity results faster for blood flow in inter-cranial aneurysm

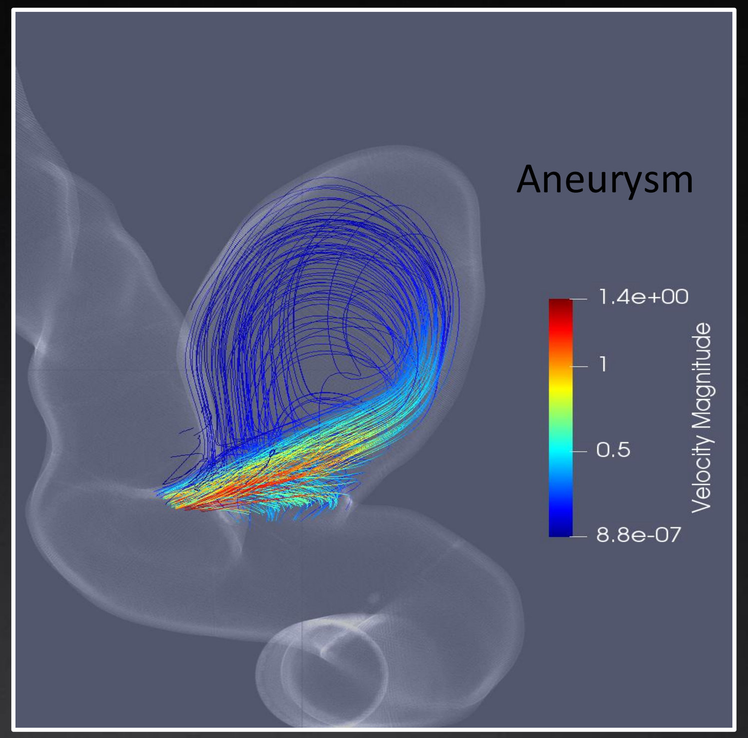

Digital Twins
Kinetic Vision: Design optimization using parameterized models

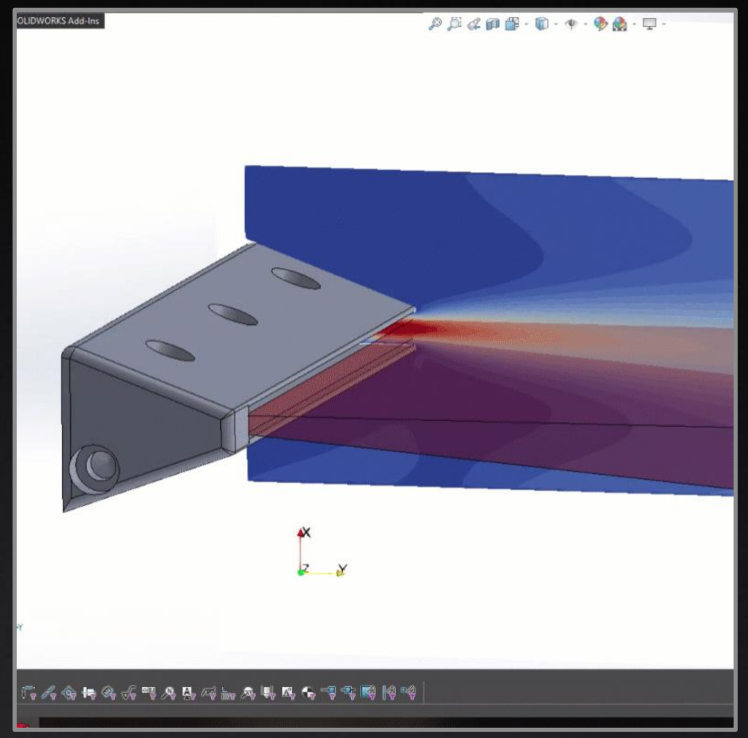

Industrial HPC
NETL: 10,000X Faster build of high-fidelity surrogate models

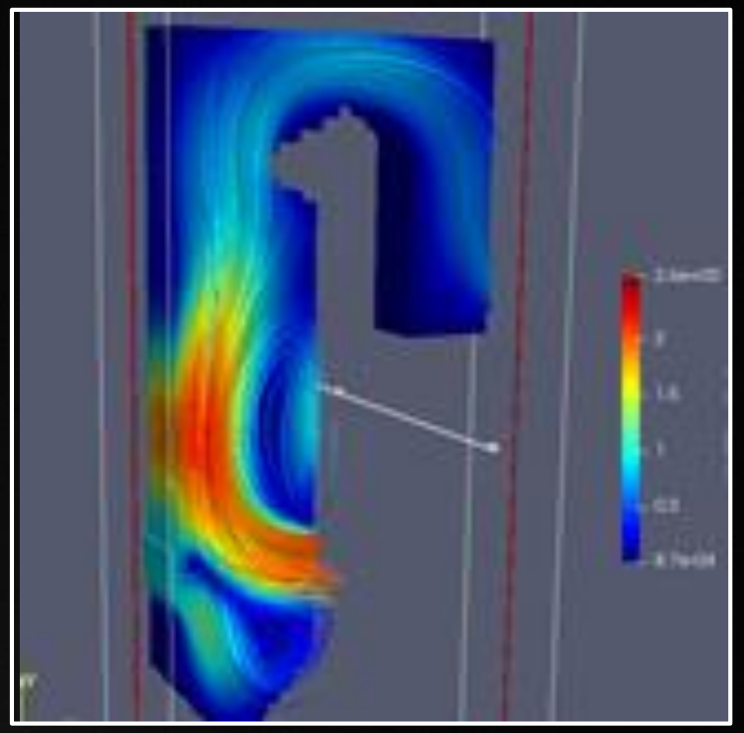

NVIDIA PhysicsNeMo

NVIDIA’s AI framework for developing Physics-AI models

Developing AI models for engineering and science applications

- Tools to develop solutions that obey first principles / domain knowledge
- Performant AI stack for real-world problem scale
- Model architectures and training pipelines tuned for CAE to accelerate adoption of AI

Unlocking accelerated simulations with AI

- AI models can run a simulation **1000x** faster than traditional numerical solvers
- Design cycles reduced to **seconds** from **hours**
- Enabling more simulations for better designs.

Optimal and Scalable training pipelines

- Memory optimized training pipelines and model architectures/layers
- Scale to multi-node systems out of the box – data and model parallel
- Reference AI enhanced sample applications

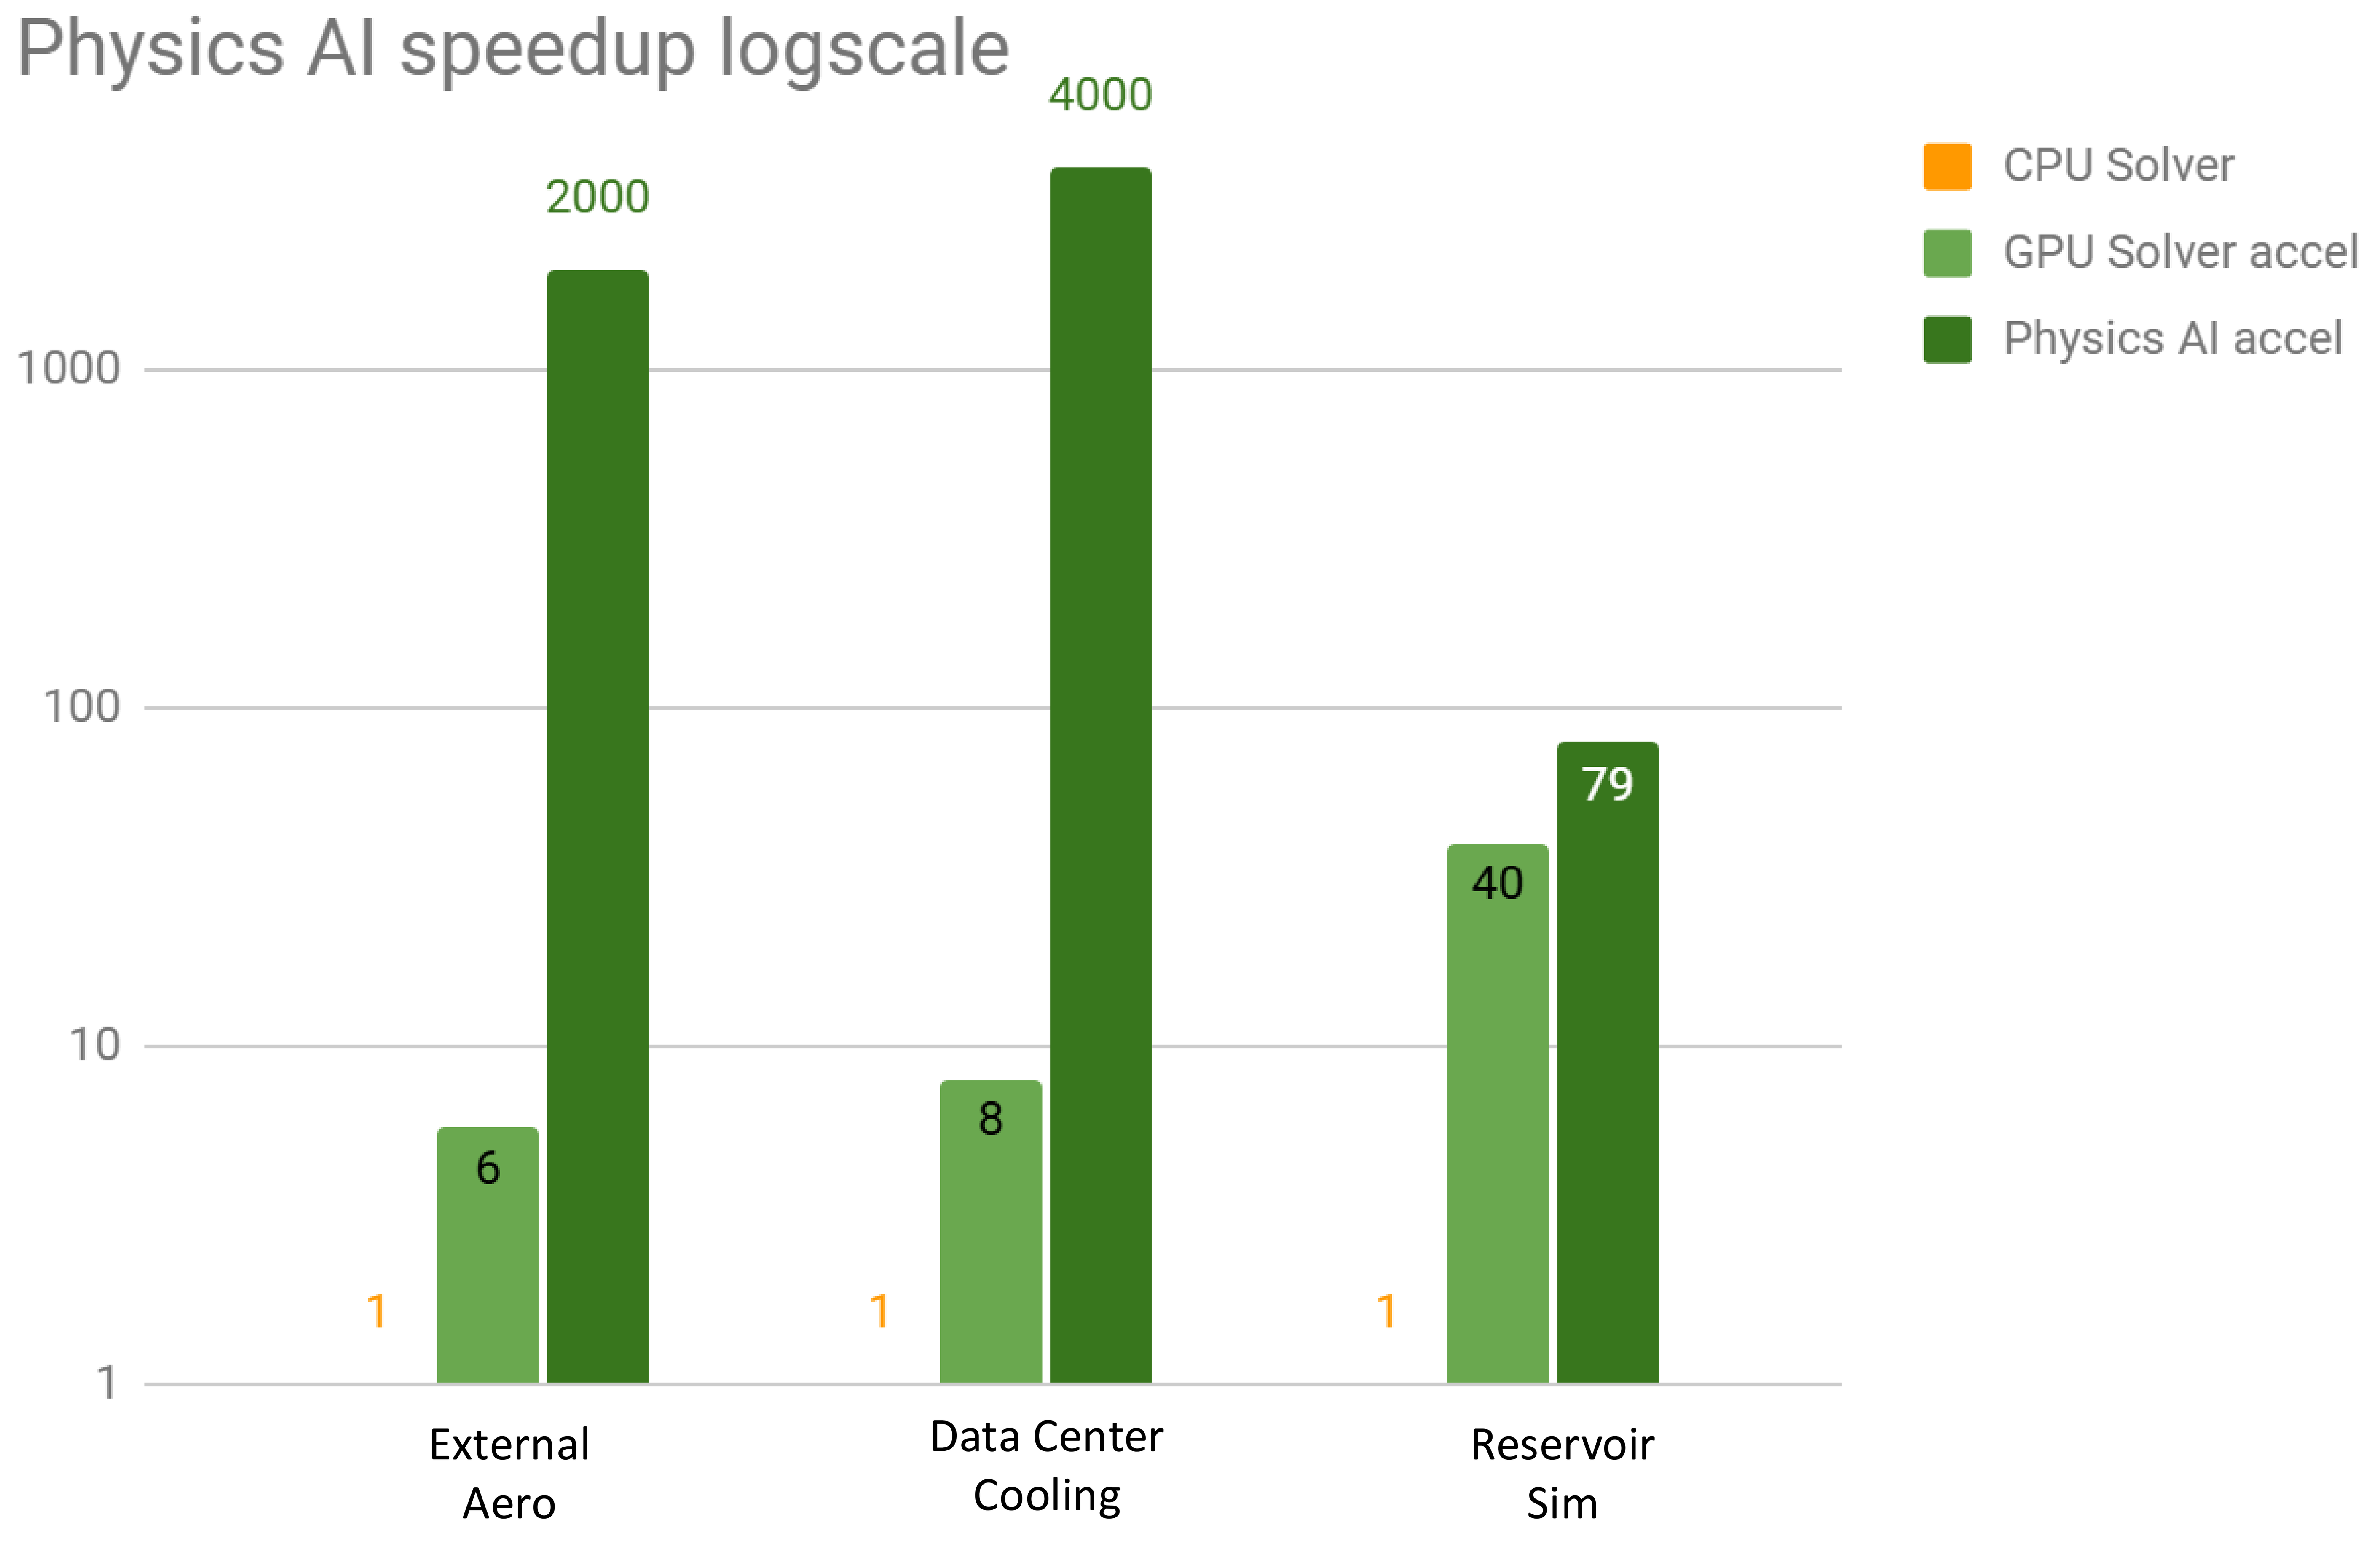

NVIDIA PhysicsNeMo

NVIDIA's AI framework for developing Physics-AI models

- Open-Source platform for Physics-ML innovation and development
- Stack for building vertical specific applications
- Platform for training and inference pipeline

PhysicsNeMo Framework

NIMs | Inference pipelines

Model architectures | Interfaces | Training recipes

Geometric module | Physics AI module | Distributed module (Model, Data and Domain parallel)

DALI | GPU Optimized PyTorch | Warp

https://github.com/NVIDIA/PhysicsNeMo

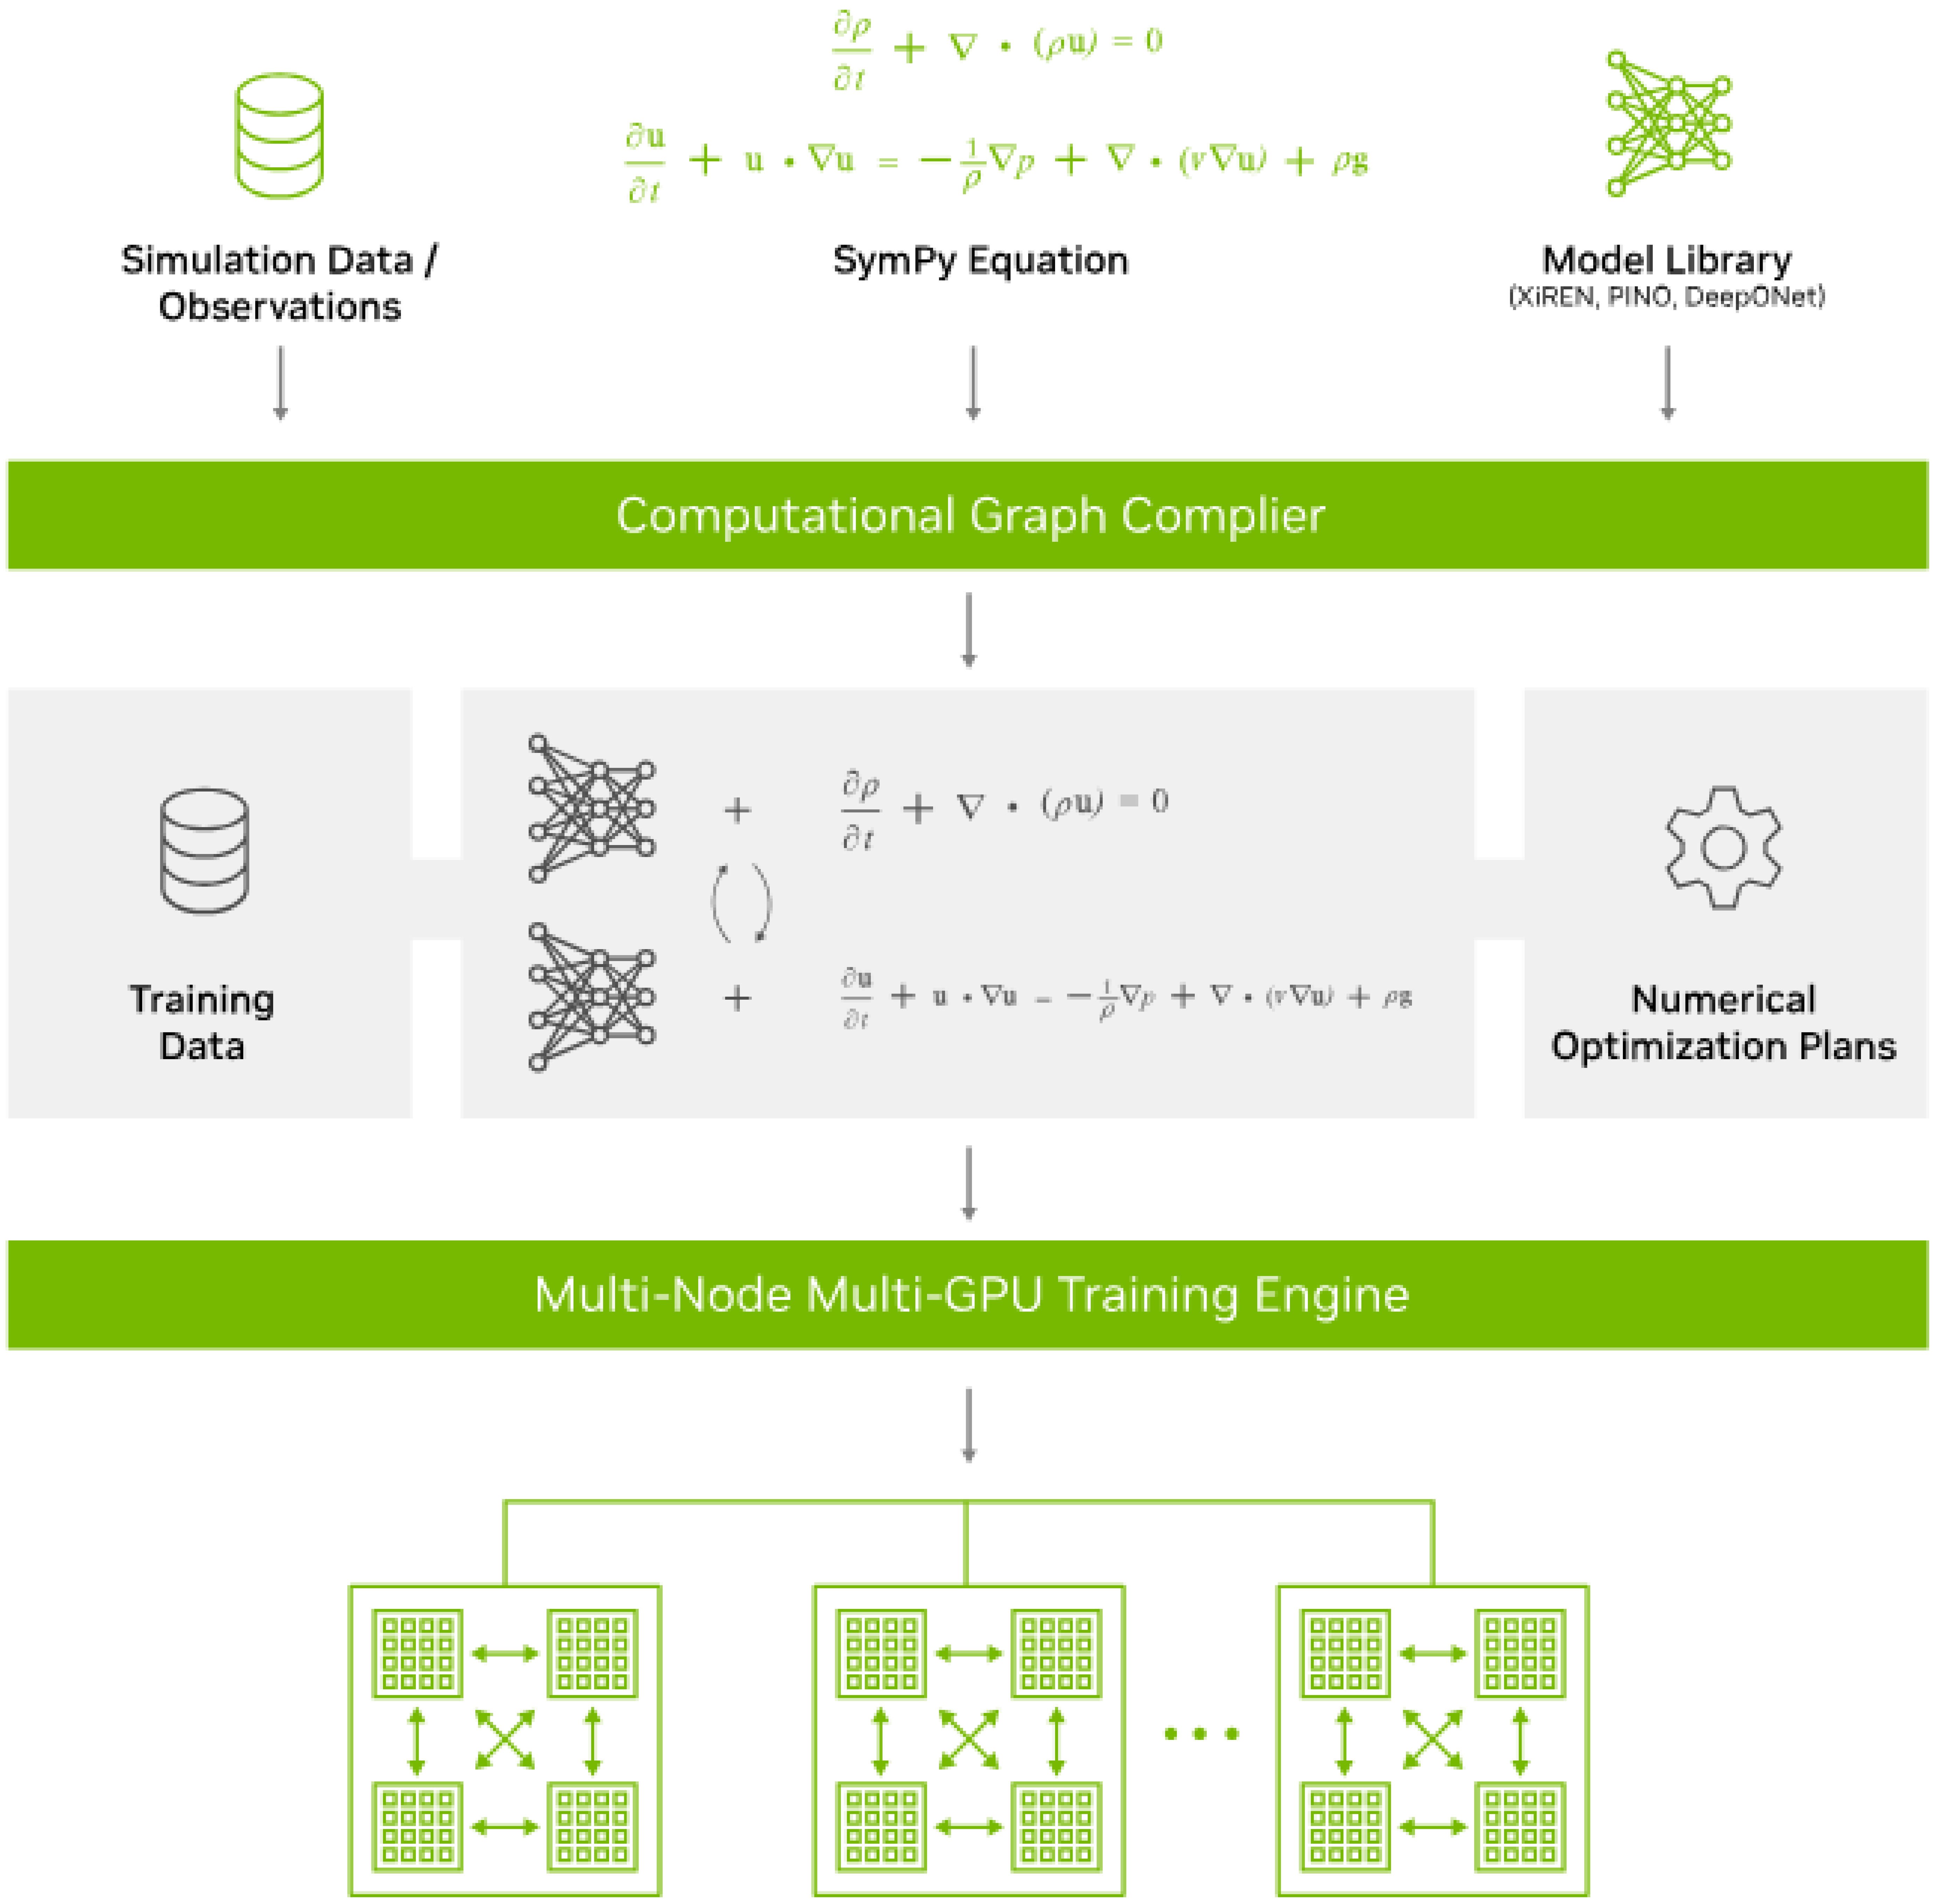

Open-Source AI Toolkit for Physics-based ML

Tailored for developing Physics AI models

The Problem

- Limited generalizability of off-the-shelf surrogate models
- Expensive, limited datasets for training
- Computer vision approaches insufficient due to convergence issues & spectral bias
- Need to satisfy governing principles for coupled PDEs and downstream applications

PhysicsNeMo

- Built-in Reference Samples: No more starting from scratch
- Modular Architecture: Abstracted ML layers as building blocks
  - Point/block conv, spectral, graph, recurrence, attention, ect.
- Supported AI Architectures
  - Graph Neural Networks (GNNs)
  - Neural Operators
  - Diffusion Models
  - Physics-Informed Neural Networks (PINNs)

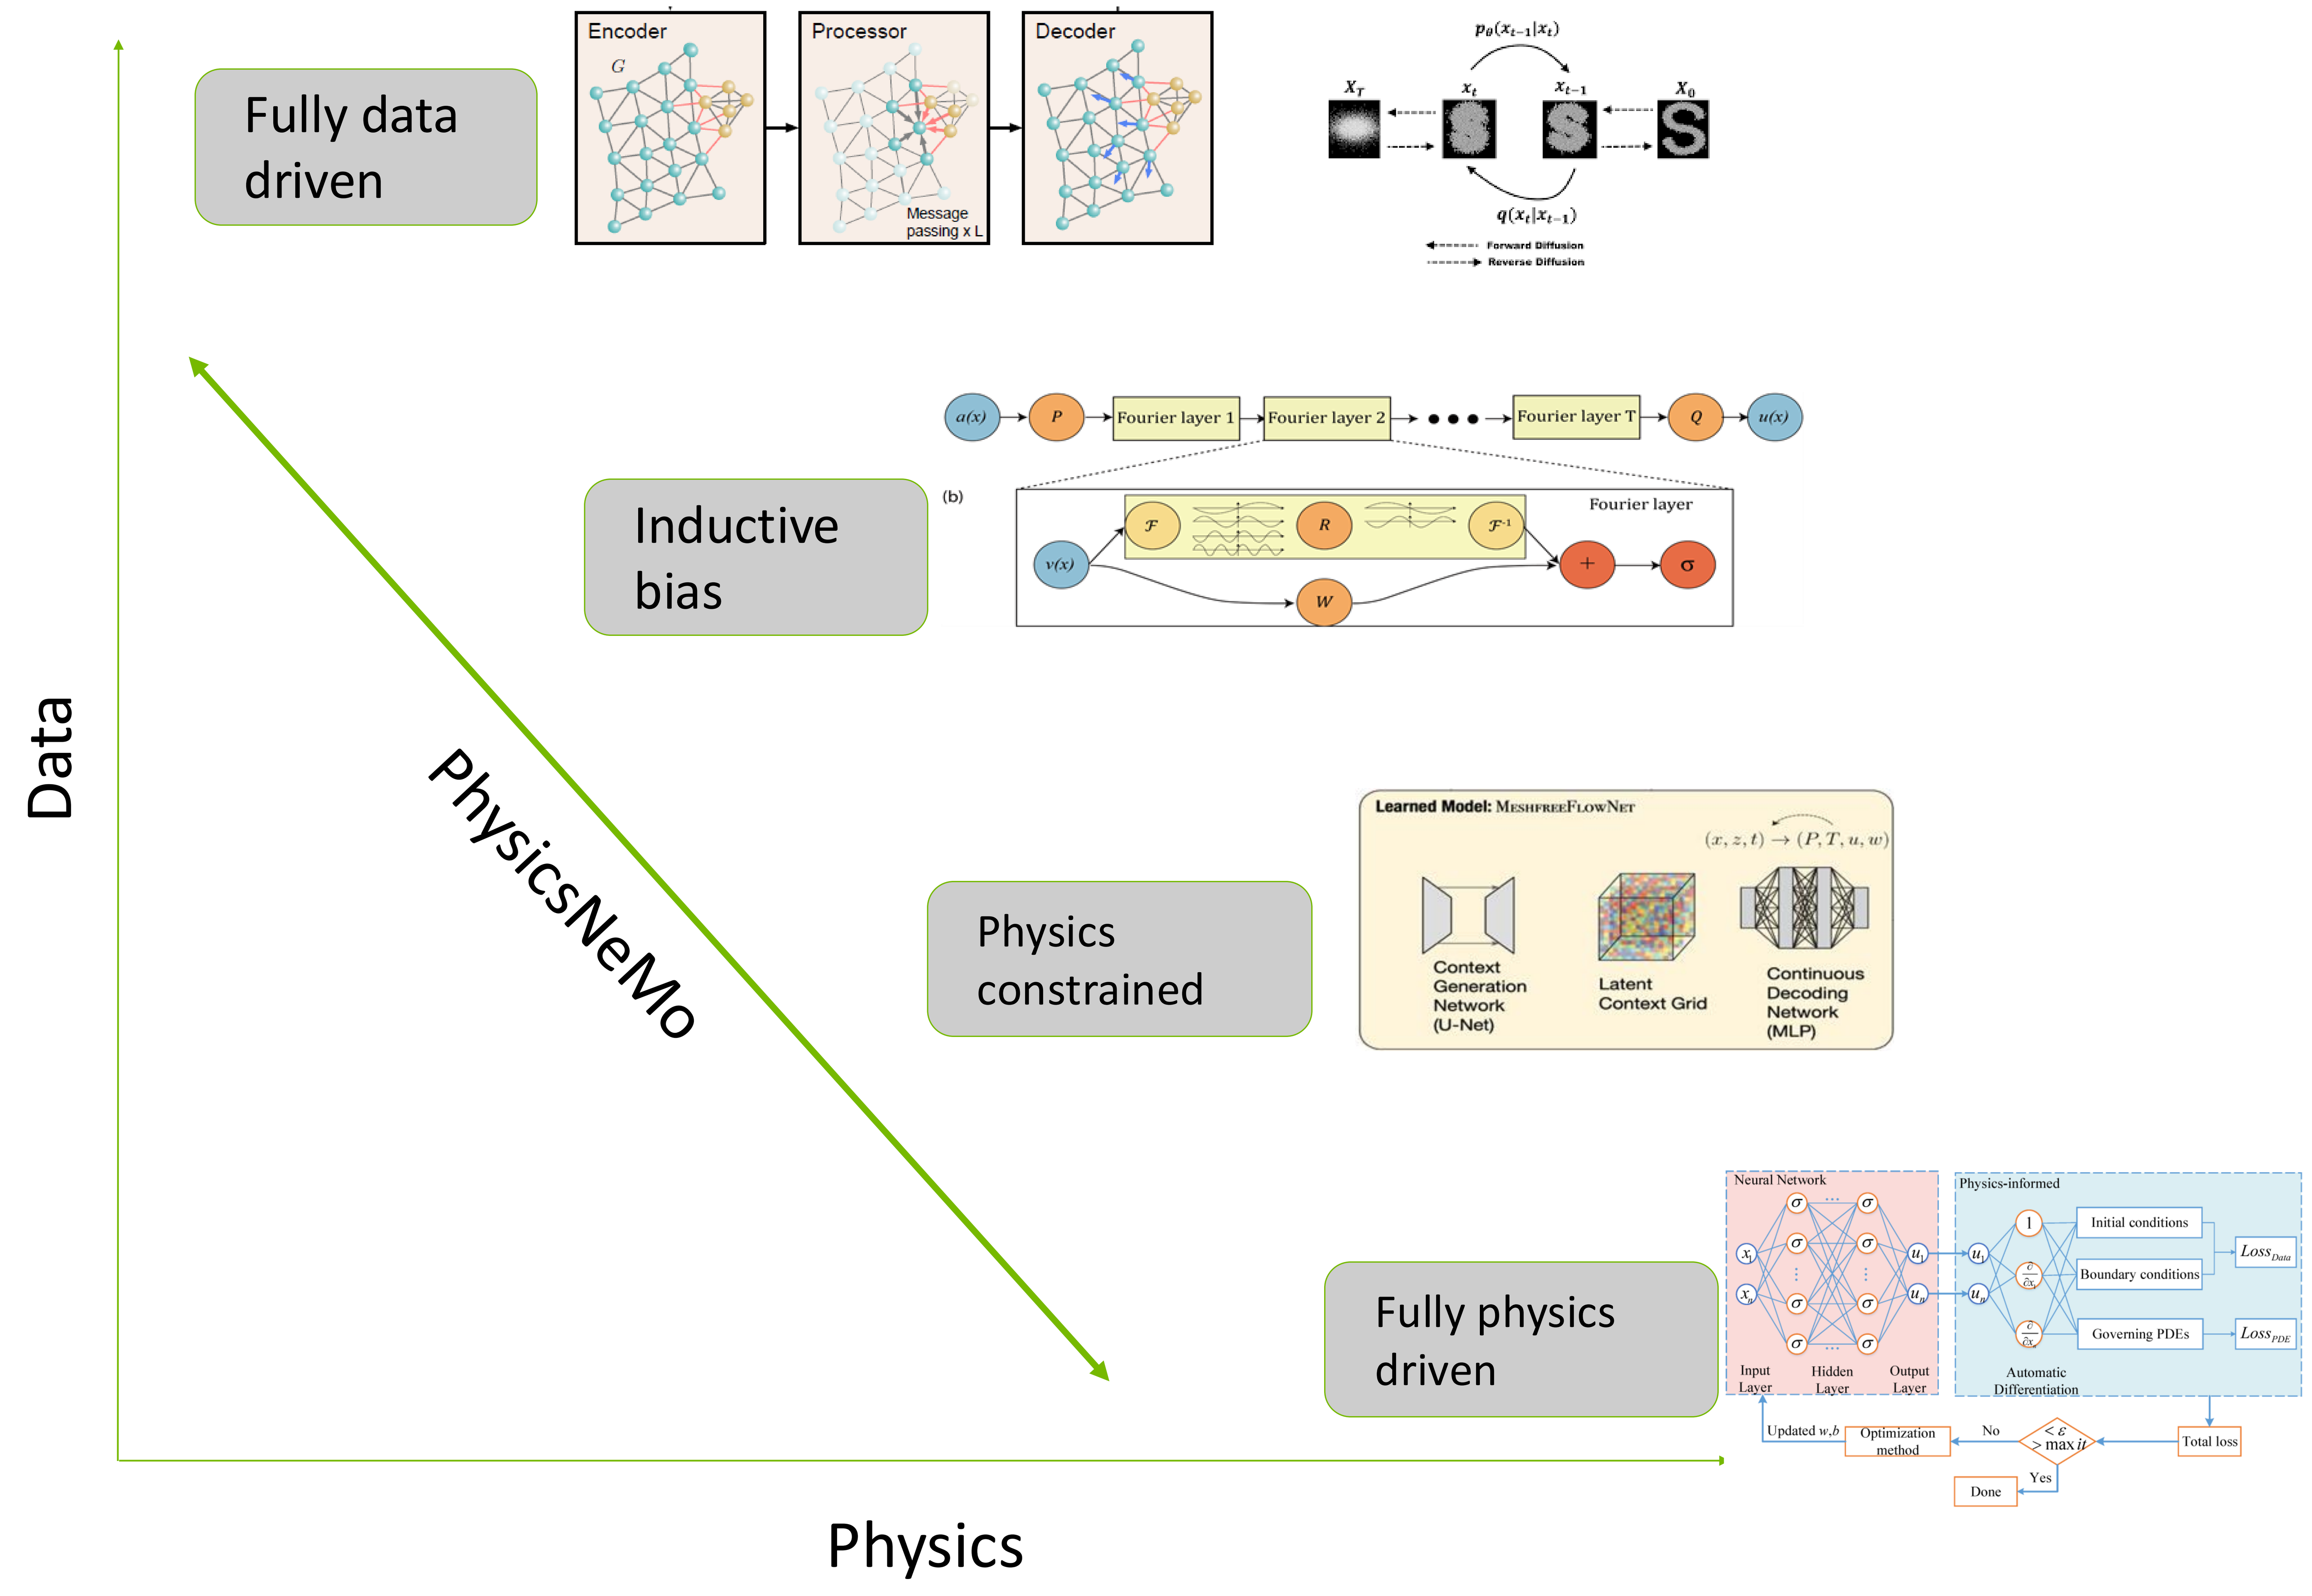

Optimized model architectures and ML layers

FNO Performance benchmarking vs PyTorch

6X more params

Model (# of Parameters)	Memory utilization(Gb) / Time taken(sec) per epoch			
	Batch size			
	1	2	3	4
PhysicsNeMo FNO (2.36m)	10.84/81.94	13.12/90.84	OOM	OOM
PhysicsNeMo FNO (1.33m)	10.62/64.1	12.75/74.1	14.04/72.9	OOO
PhysicsNeMo FNO (0.5m)	9.89/43	11.09/51	12.01/49.1	13.17/47.1
Base FNO (699k)	12.44/101.8	OOM	OOM	OOM
Base FNO (393k)	11.62/75.26	14.41/83.15	OOM	OOM

- Key takeaways:
  - PhysicsNeMo implemented FNO can accommodate a larger batch size for the same number of parameters. The training time per epoch is also better as compared to the Base FNO implementation.
  - PhysicsNeMo FNO can accommodate 6x larger number of model parameters for the same memory utilization.
- Low level kernels (FFT kernels, Vectorized and elementwise kernels etc.) evaluation is optimized in PhysicsNeMo implemented FNO .
- Minimal tensor copying resulting in better memory utilization.

Optimized model architectures and ML layers

MeshGraphNet Performance benchmarking

Model (Mesh size)	Number of GPUs				
	1	2	4	8	16
Base MGN (0.5 mil)	✅	x	x	x	x
PhysicsNeMo MGN (0.5 mil)	✅	✅	✅	✅	✅
PhysicsNeMo MGN (1 mil)	x	✅	✅	✅	✅
PhysicsNeMo MGN (2 mil)	x	x	x	✅	✅
PhysicsNeMo MGN (5 mil)	x	x	x	x	✅

- Key takeaways:
  - Base MeshGraphNet implementation is only data parallel. In PhysicsNeMo, an optimized, graph parallel GPU implementation is provided.
  - The graph-parallel MeshGraphNet implementation in PhysicsNeMo scales to multiple nodes. The distributed message passing is optimized for memory and performance.
  - Gradient checkpointing, fused activation and low-level network improvements for improvement in memory utilization.
- Minimal tensor copying resulting in better memory utilization.

Multi-level parallelism for enterprise engineering scale solutions

Domain, model, and data parallelisms.

- Parallelism in AI applications has several dimensions :
  - **Data** (or **batch**) parallelism distributes a minibatch over several GPUs. Ideal when the size of the data and number of model parameters are modest.
  - **Model** parallelism distributes model weights (and corresponding optimizer states) over GPUs. Useful as the number of trainable parameters grows.
  - **Pipeline** parallelism distributes entire layers over GPUs, connecting the output of one GPU as the input to another GPU.
- In scientific AI training, the driver of memory utilization is often the extremely high resolution data.
- Generalized, **domain parallelization** techniques to enable generic, sharded computation.
- Enabling multi-level parallelism to compose **domain, model, and data parallelisms**.

Increasing Data resolution ...

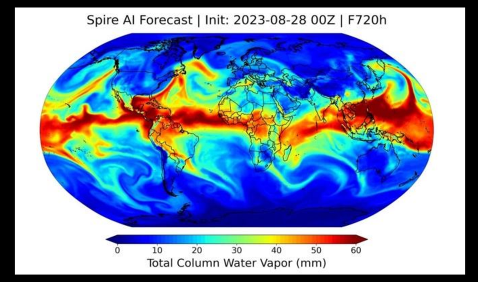

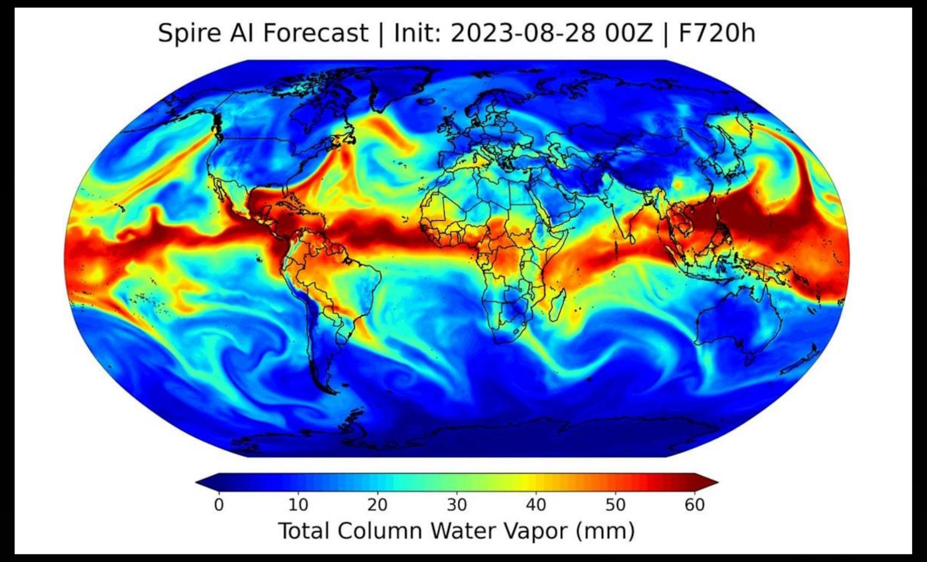

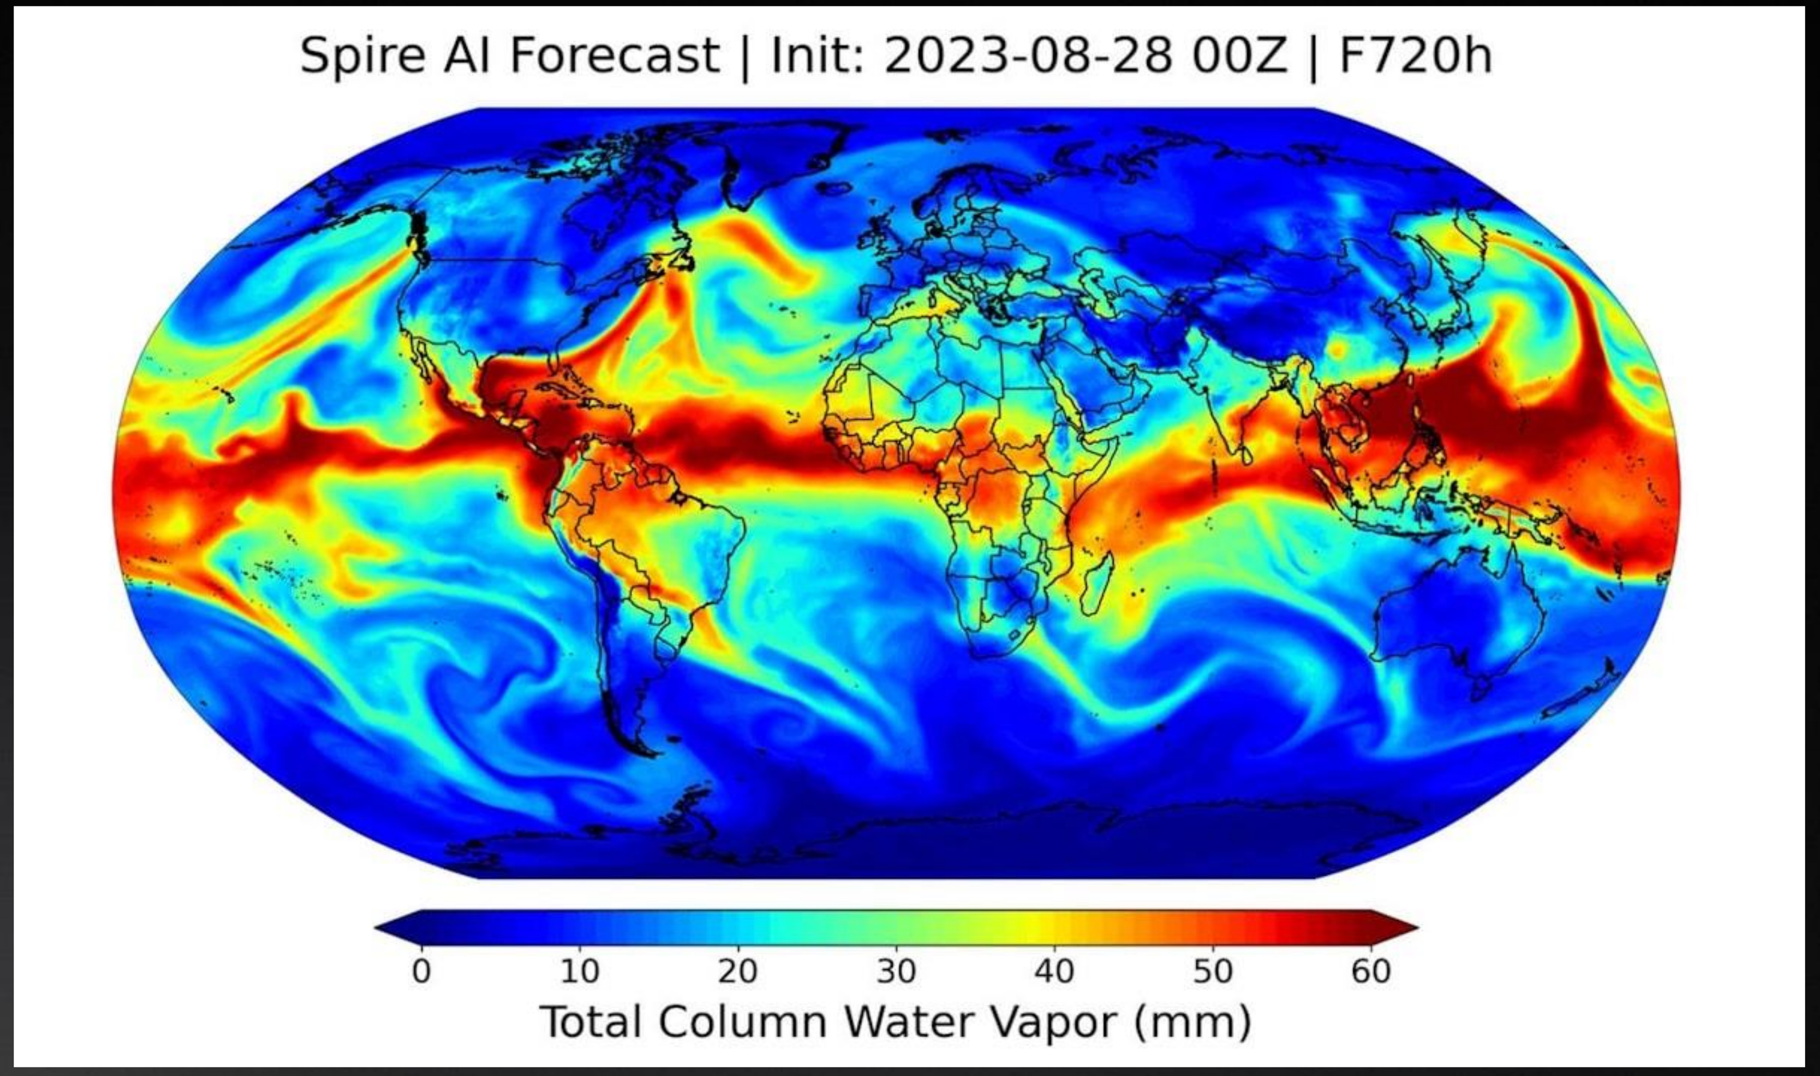

... leads to exponential growth of GPU memory and compute usage.

Multi-level parallelism for enterprise engineering scale solutions

PhysicsNeMo ShardTensor

- ShardTensor is a PhysicsNemo utility for building domain-parallel applications.
- ShardTensor:
  - Combines the data, metadata, and device orchestration concepts into one object.
  - Interoperates with Pytorch's FSDP framework to enable multilevel parallelism
  - Leverages Pytorch support for a variety of operations (tensor ops, reshaping, reductions) with extensions in PhysicsNeMo to enable critical operations (ex: Convolutions, Attention, GroupNorm, etc).

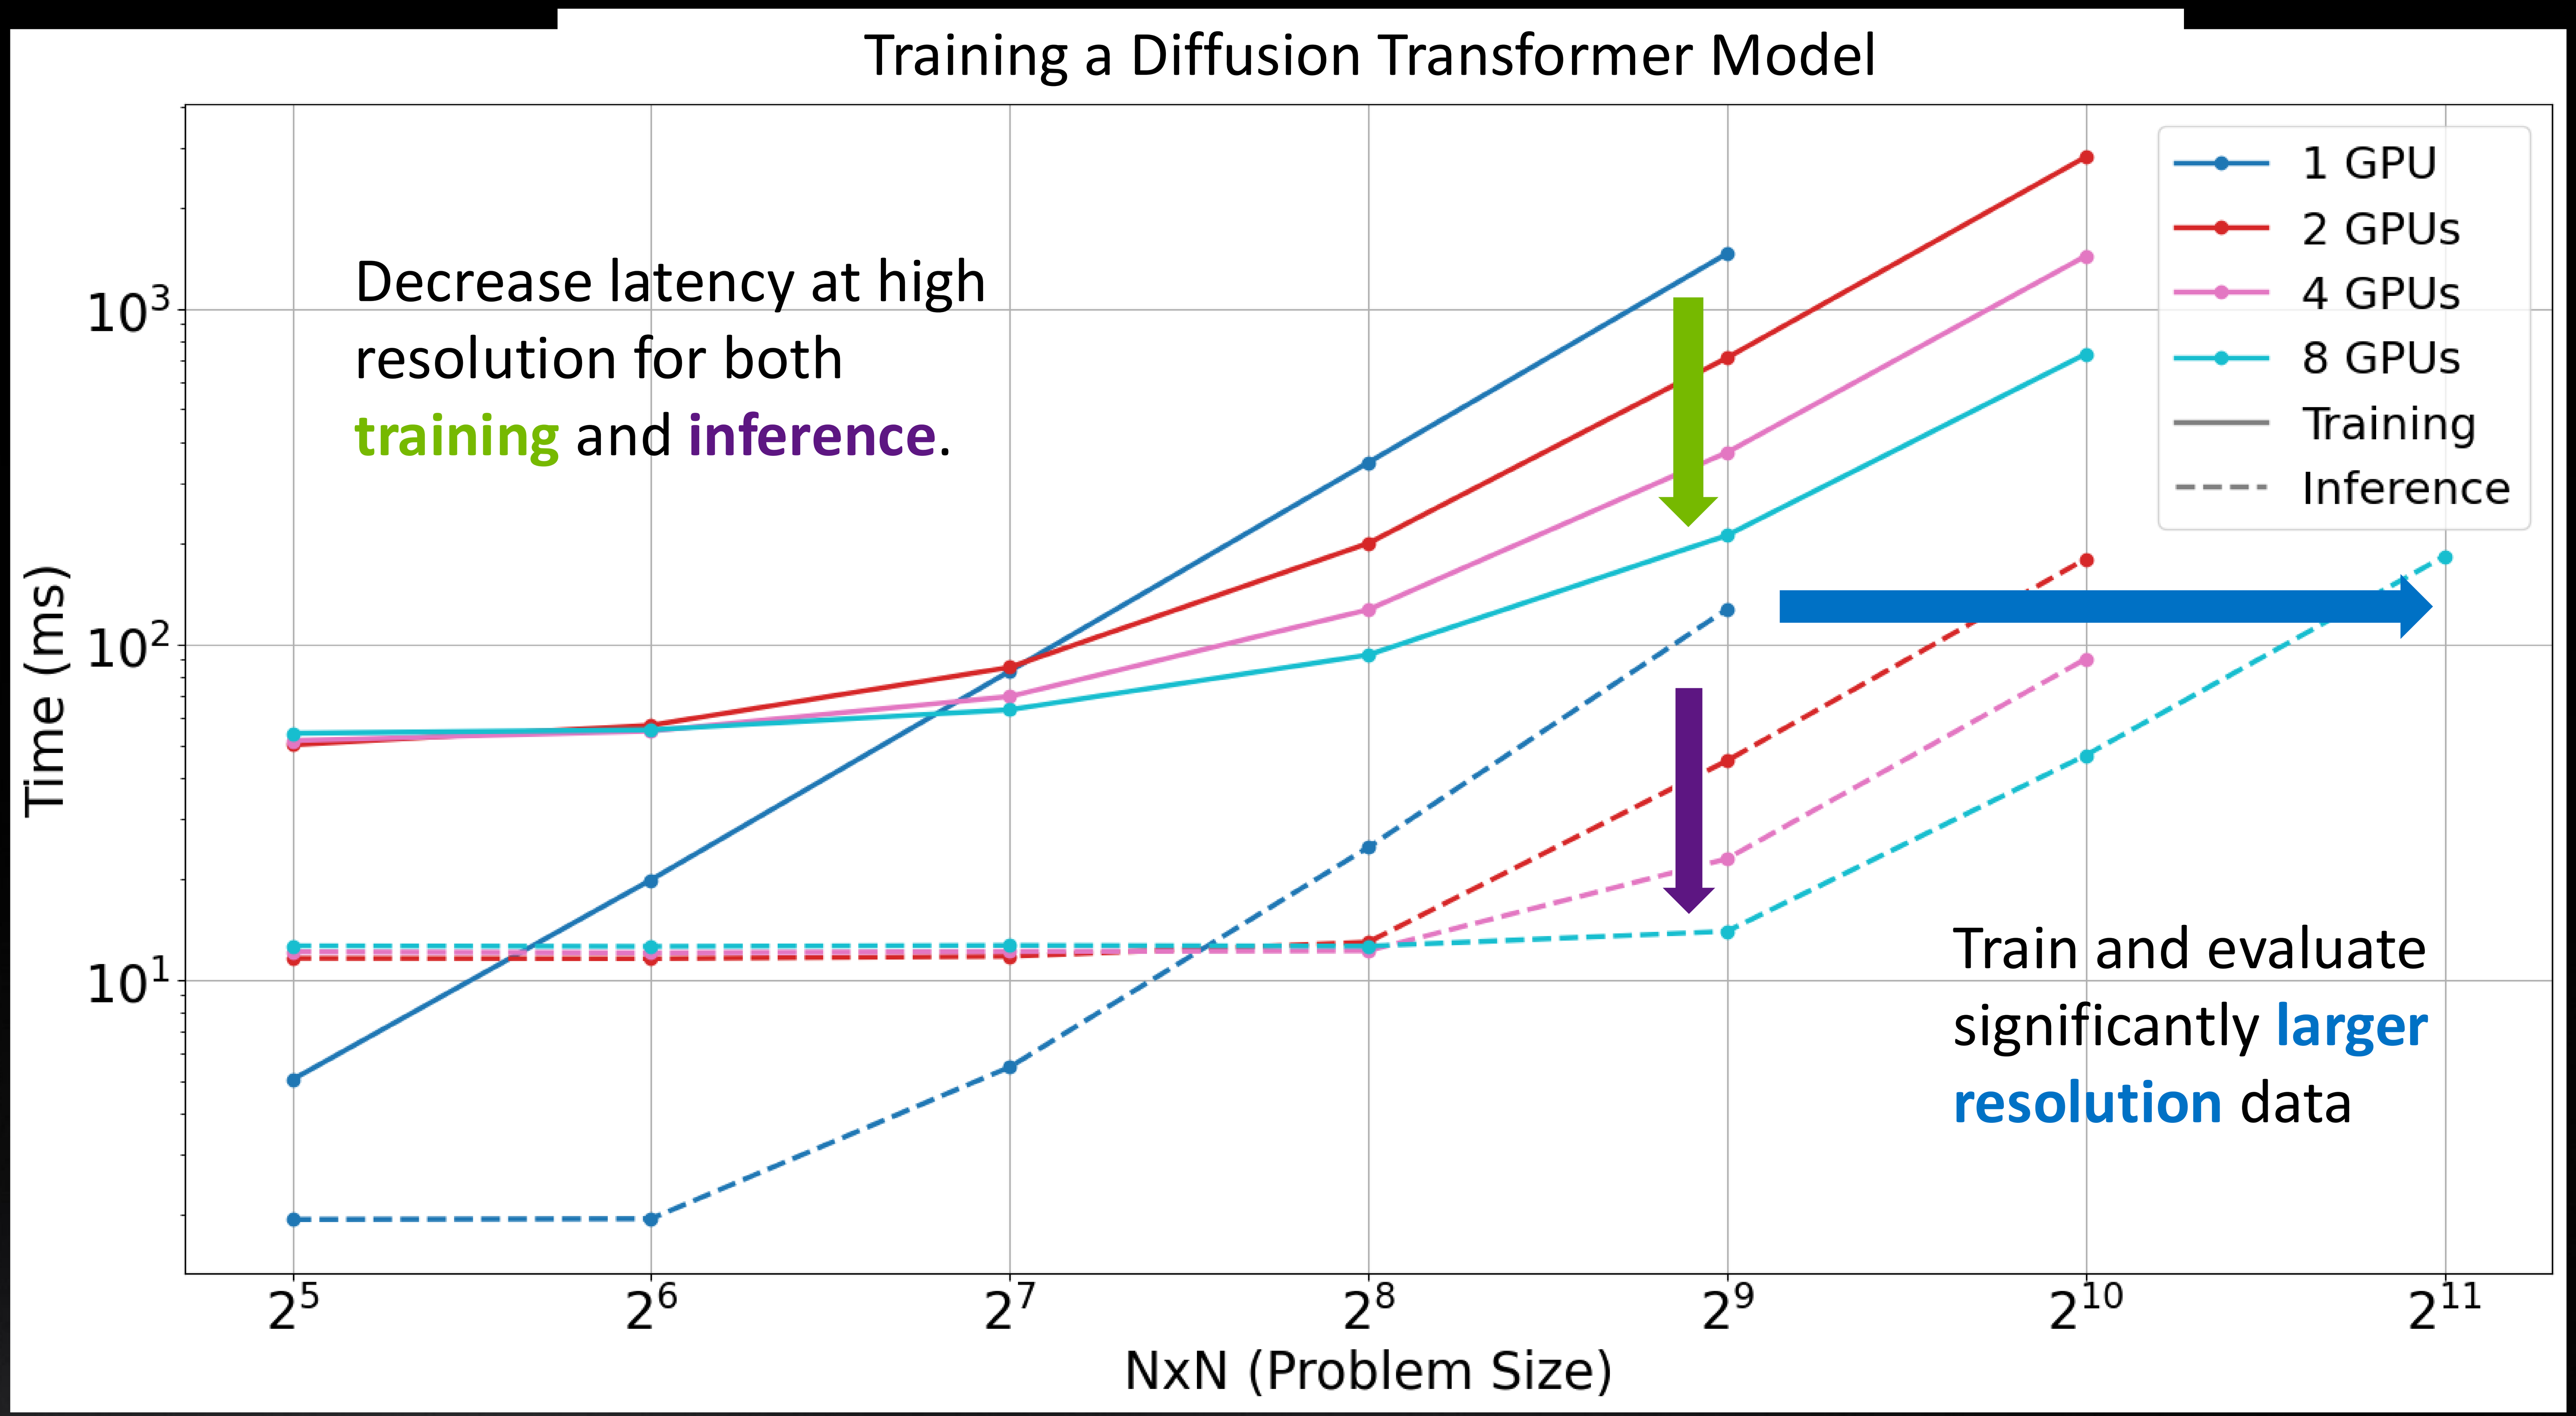

NVIDIA

Surrogates for Aerodynamic Prediction

Goal : Accurate and efficient predictions of aerodynamic quantities (surface pressure, wall shear stress, volume fields)

X-MeshGraphNet - Multi-scale Graph Neural Network for Physics Simulation

- Scales effectively to large meshes
- Multi-Scale Graphs: Combines coarse and fine-resolution point clouds to capture local and long-range interactions
- Accurately predicts surface pressure and wall shear stresses
- Overcomes scalability and mesh preprocessing bottlenecks in traditional GNNs

DoMINO - Decomposable Multi-scale Iterative Neural Operator

- Point Cloud-Based: Uses local geometric features to predict flow fields on discrete surface and volume points
- Mesh-Independent: Trained on one mesh, transferable across discretisation without re-training
- Predicts both surface and volume flow fields, scales to large engineering simulations

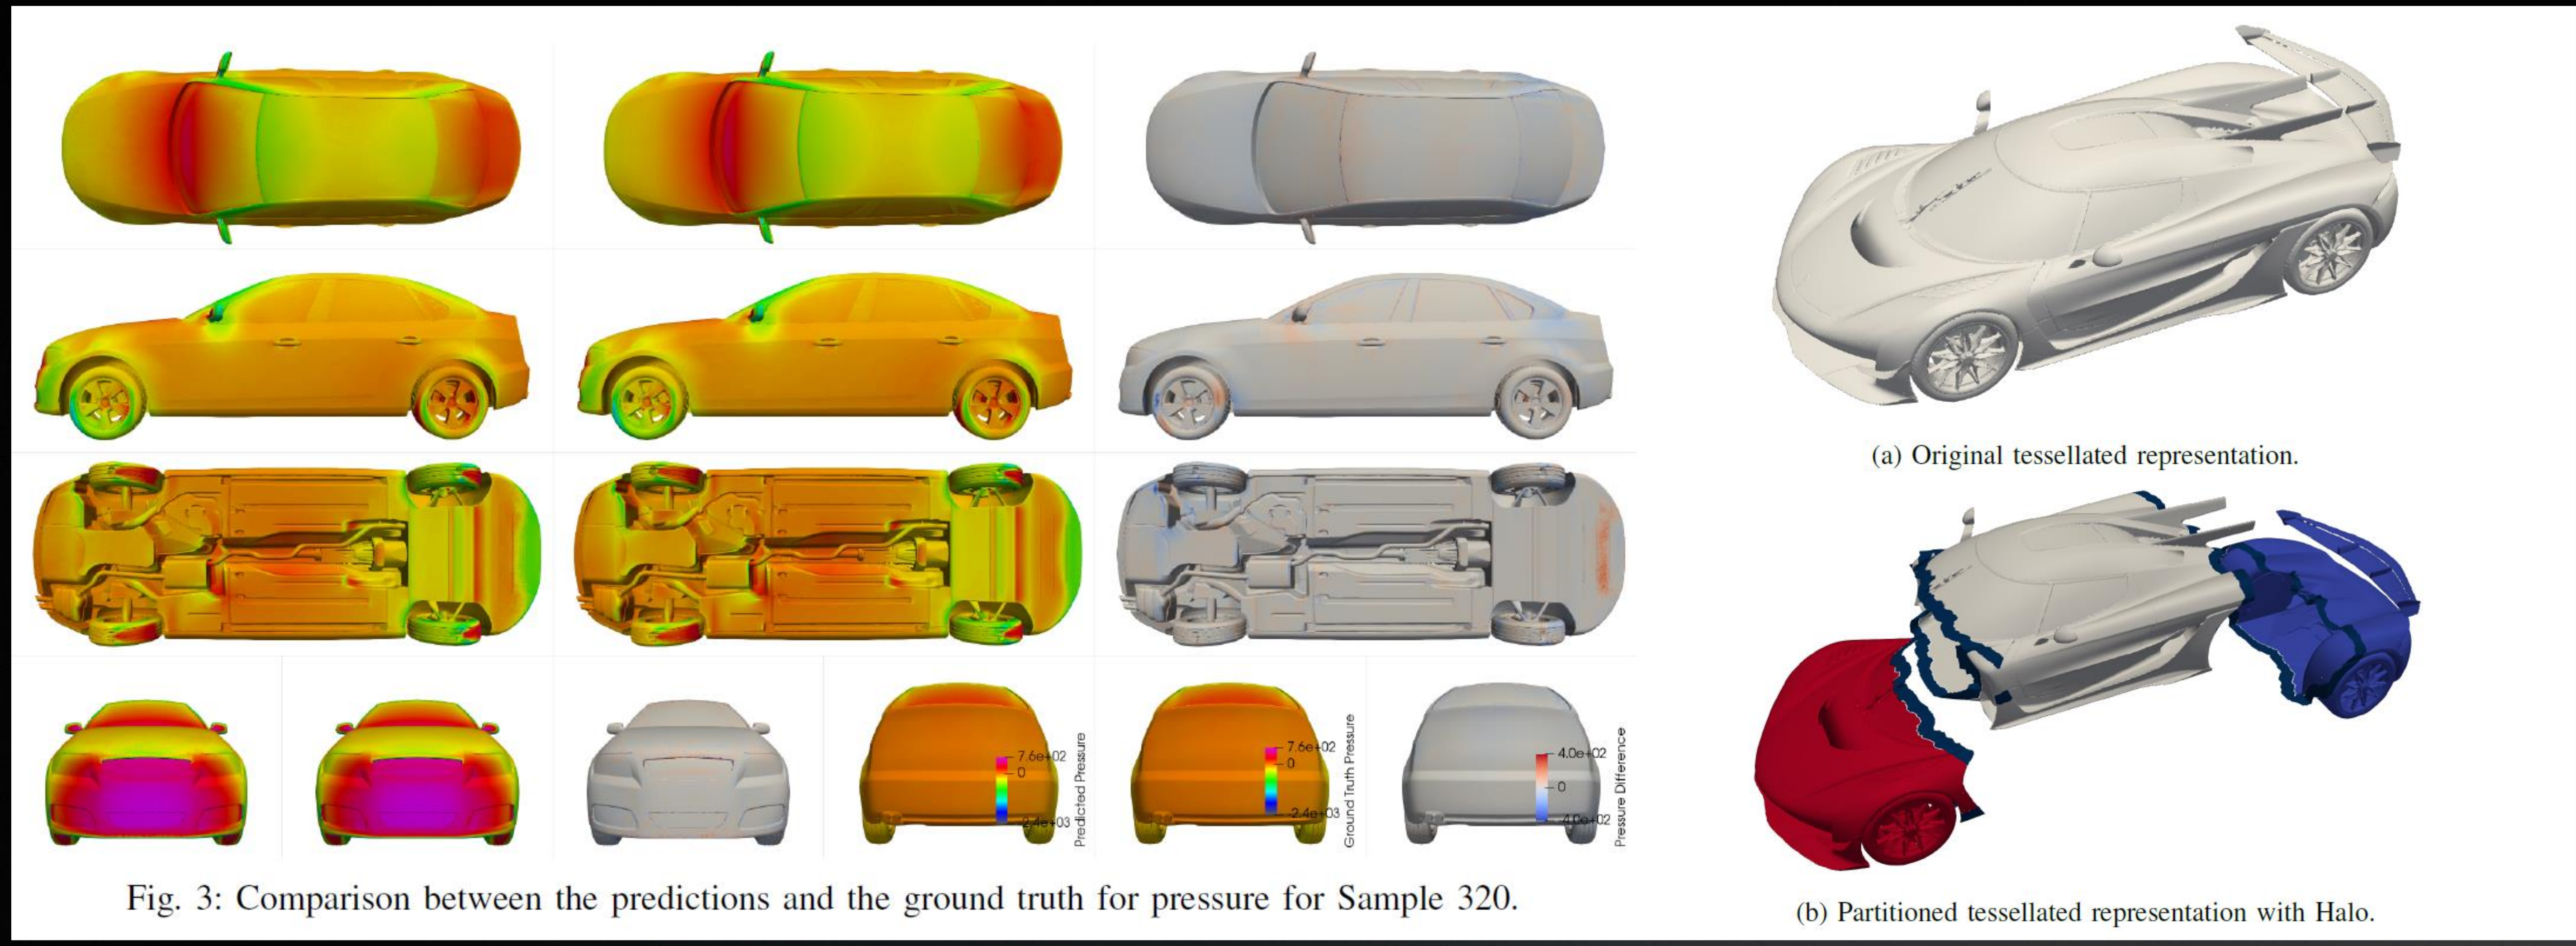

X-MeshGraphNet - https://arxiv.org/abs/2411.17164

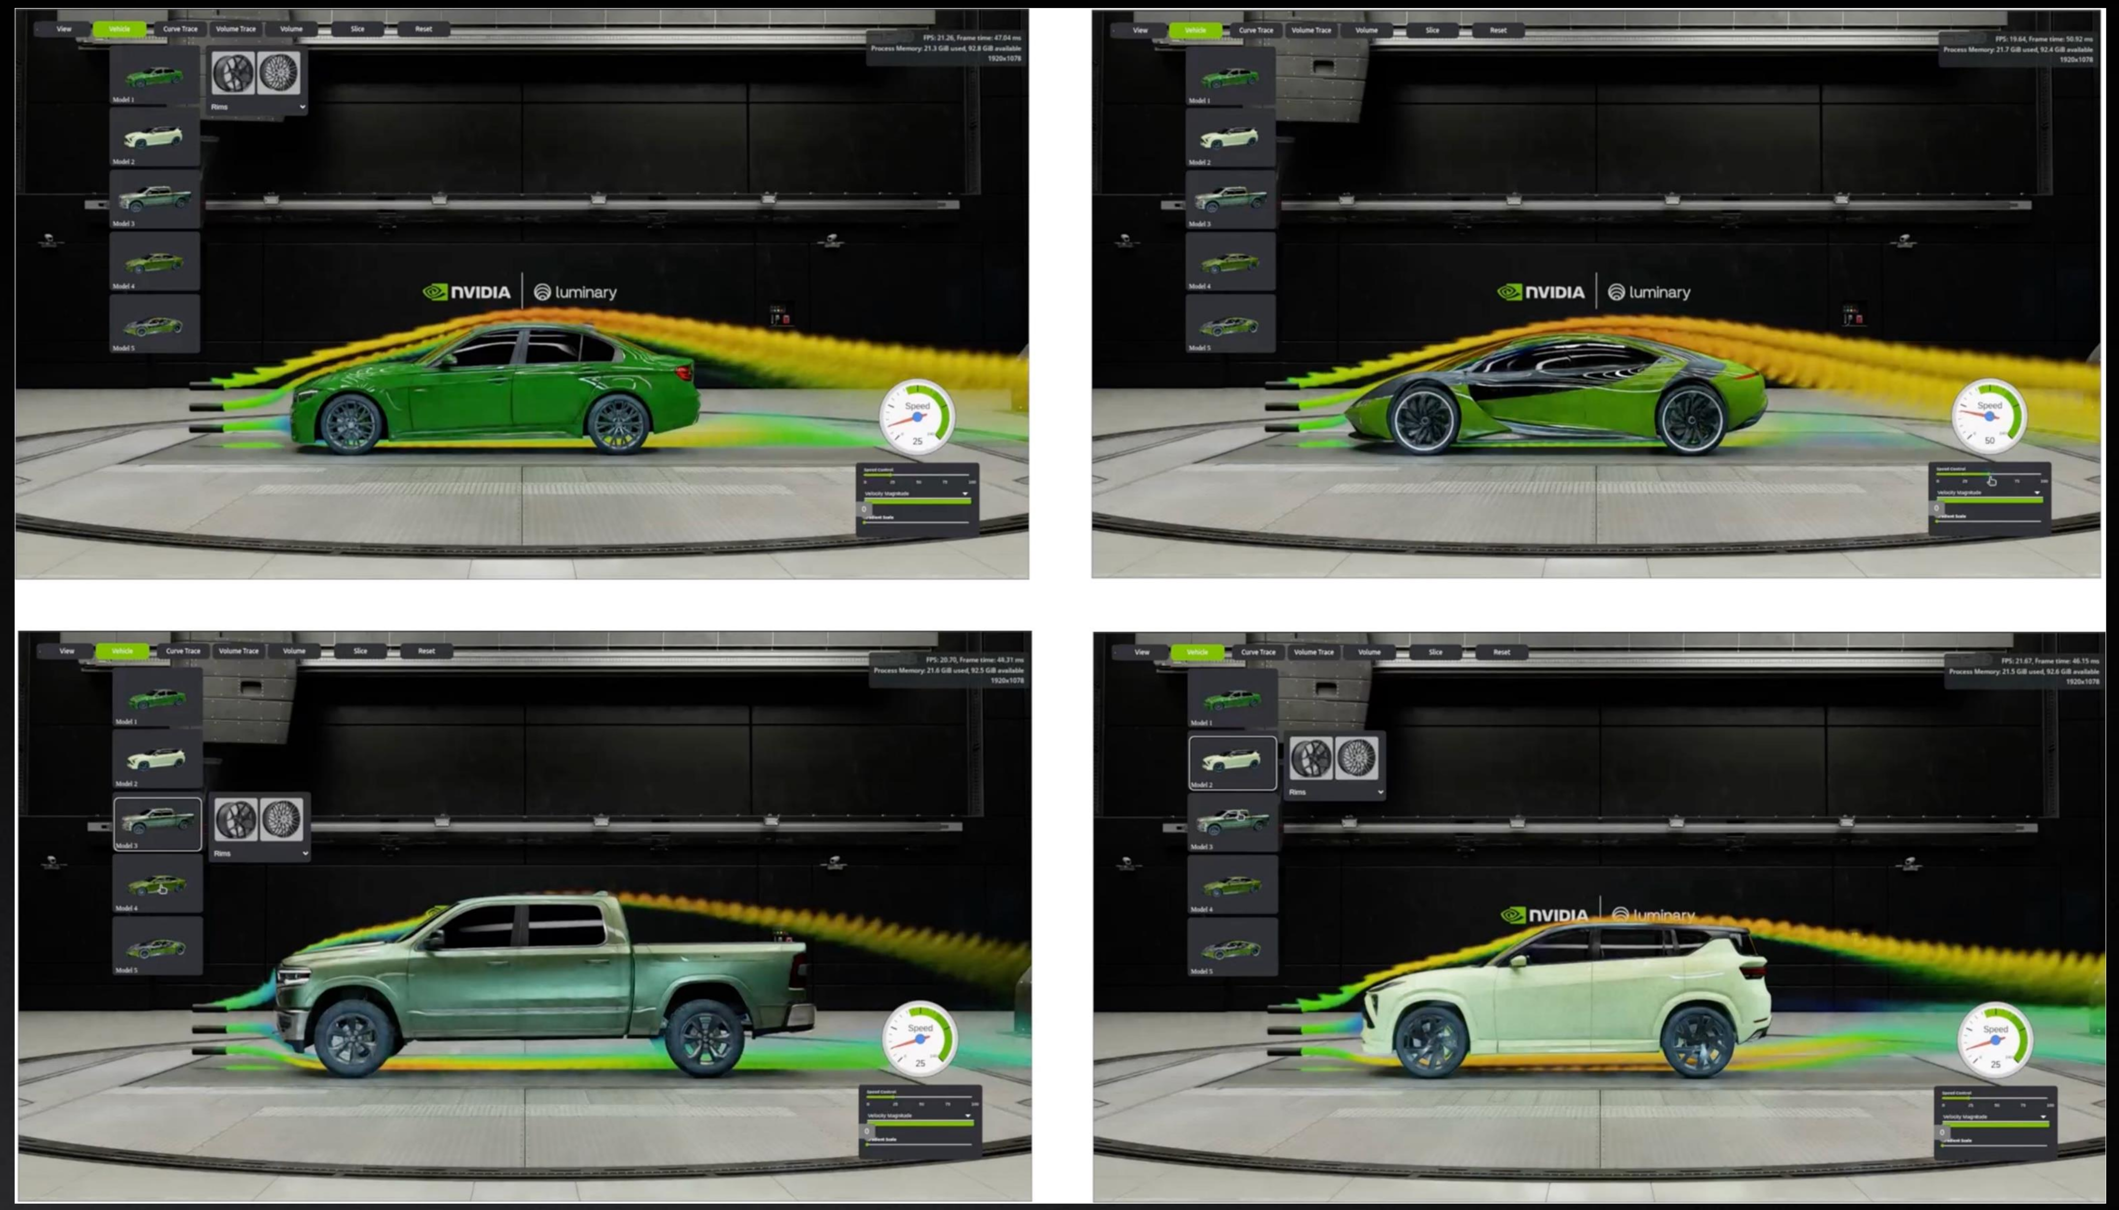

Demo - https://build.nvidia.com/nvidia/digital-twins-for-fluid-simulation

DoMINO - https://arxiv.org/abs/2501.13350

DoMINO

Point cloud based neural operator for scalability, accuracy and generalizability

- DoMINO stands for Decomposable Multiscale Iterative Neural Operator
- Key features:
  - **Neural operator:** Predicts point-wise volume and surface fields, infinitely scalable
  - **Decomposable:** Learns local geometry representations in sub-regions to improve solution accuracy
  - **Multi-scale:** Learns multi-scale point kernels to capture fine- and coarse-scale geometry features
  - **Iterative:** Facilitates long-range interaction by propagating geometry features into computational domain
  - **Basis in traditional numerical methods:** Builds dynamic computational stencils to learn non-linear basis functions
  - **Only STLs required at inference (no surface or volume meshes):** Insensitive to spatial structure and density of point cloud at inference
- DoMINO artifacts:
  - Preprint https://arxiv.org/pdf/2501.13350
  - Source code: https://github.com/NVIDIA/physicsnemo/blob/domino/physicsnemo/models/domino/model.py

DOMINO MODEL ARCHITECTURE

Overview

Three components:

- **Multi-res geometry NN:** Transforms STLs to structured representations using multi-scale point convolutional kernels
- **Local geometry representation:** Extracts local geometry encodings in multi-res subdomains from global representations
- **Aggregation NN:** Dynamically constructs a finite-volume stencil and approximates solution on the cell center conditioned on the local geometry representation

AI SURROGATE FOR DESIGN SPACE OPTIMIZATION

Real Time Digital Wind Tunnel for road vehicle aerodynamics

Use Case

- Enabling real-time design exploration through AI-accelerated virtual wind tunnel simulations
- Allowing engineers to iterate on complex aerodynamic geometries with immediate feedback

Challenges

- Traditional CFD simulations (RANS, LES) require significant computational resources and time — ranging from hours to months
- Limited number of design iterations due to the high cost and duration of conventional solvers

Solution

- AI surrogate model trained using solver-generated data to emulate flow physics at high fidelity and speed
- DoMINO architecture capable of handling varying geometries

Outcome

- ~1,200x speed-up in design iteration time
- Full real-time feedback loop for geometry modification and aerodynamic evaluation
- Model can be finetuned on new geometries to transfer to new domains, requiring smaller datasets

https://github.com/NVIDIA/modulus/tree/main/examples/cfd/external_aerodynamics/domino

NVIDIA

WORKFLOW DETAILS

DoMINO AI model

- Captures high-fidelity volume and surface flow fields on large meshes
- Local geometry representation
- Basis in traditional numerical methods
- No mesh required at inference, non-uniform point cloud can be sampled

DoMINO model was trained on ~1500 OpenFOAM simulations (20-50 mil meshes), comprising 5 vehicle classes (Sedans, SUVs, Pickups, Vans, Hatchbacks) for flow speeds ranging between 45 and 135 mph.

- Inputs to model:
  - Geometry STL (triangulated surface mesh), Inlet velocity, sampled point on volume and surface
- Outputs of model
  - Surface fields: pressure, wall-shear vector and engineering metrics drag, lift forces.
  - Volume fields: pressure, velocity, turbulent viscosity and kinetic energy etc.

Compute details:

- Data generation (~8 hrs per case on 64 CPUs and ~2 hrs on 8 H100s with partially accelerated GPU solver)
- Training on 2 H100 nodes and took about 4 days
- NIM optimized to run on a single GPU with H100, A100 and L40.
  - Takes about ~4 seconds end-to-end to evaluate 0.5 million points (includes time taken for sending-receiving inference requests)

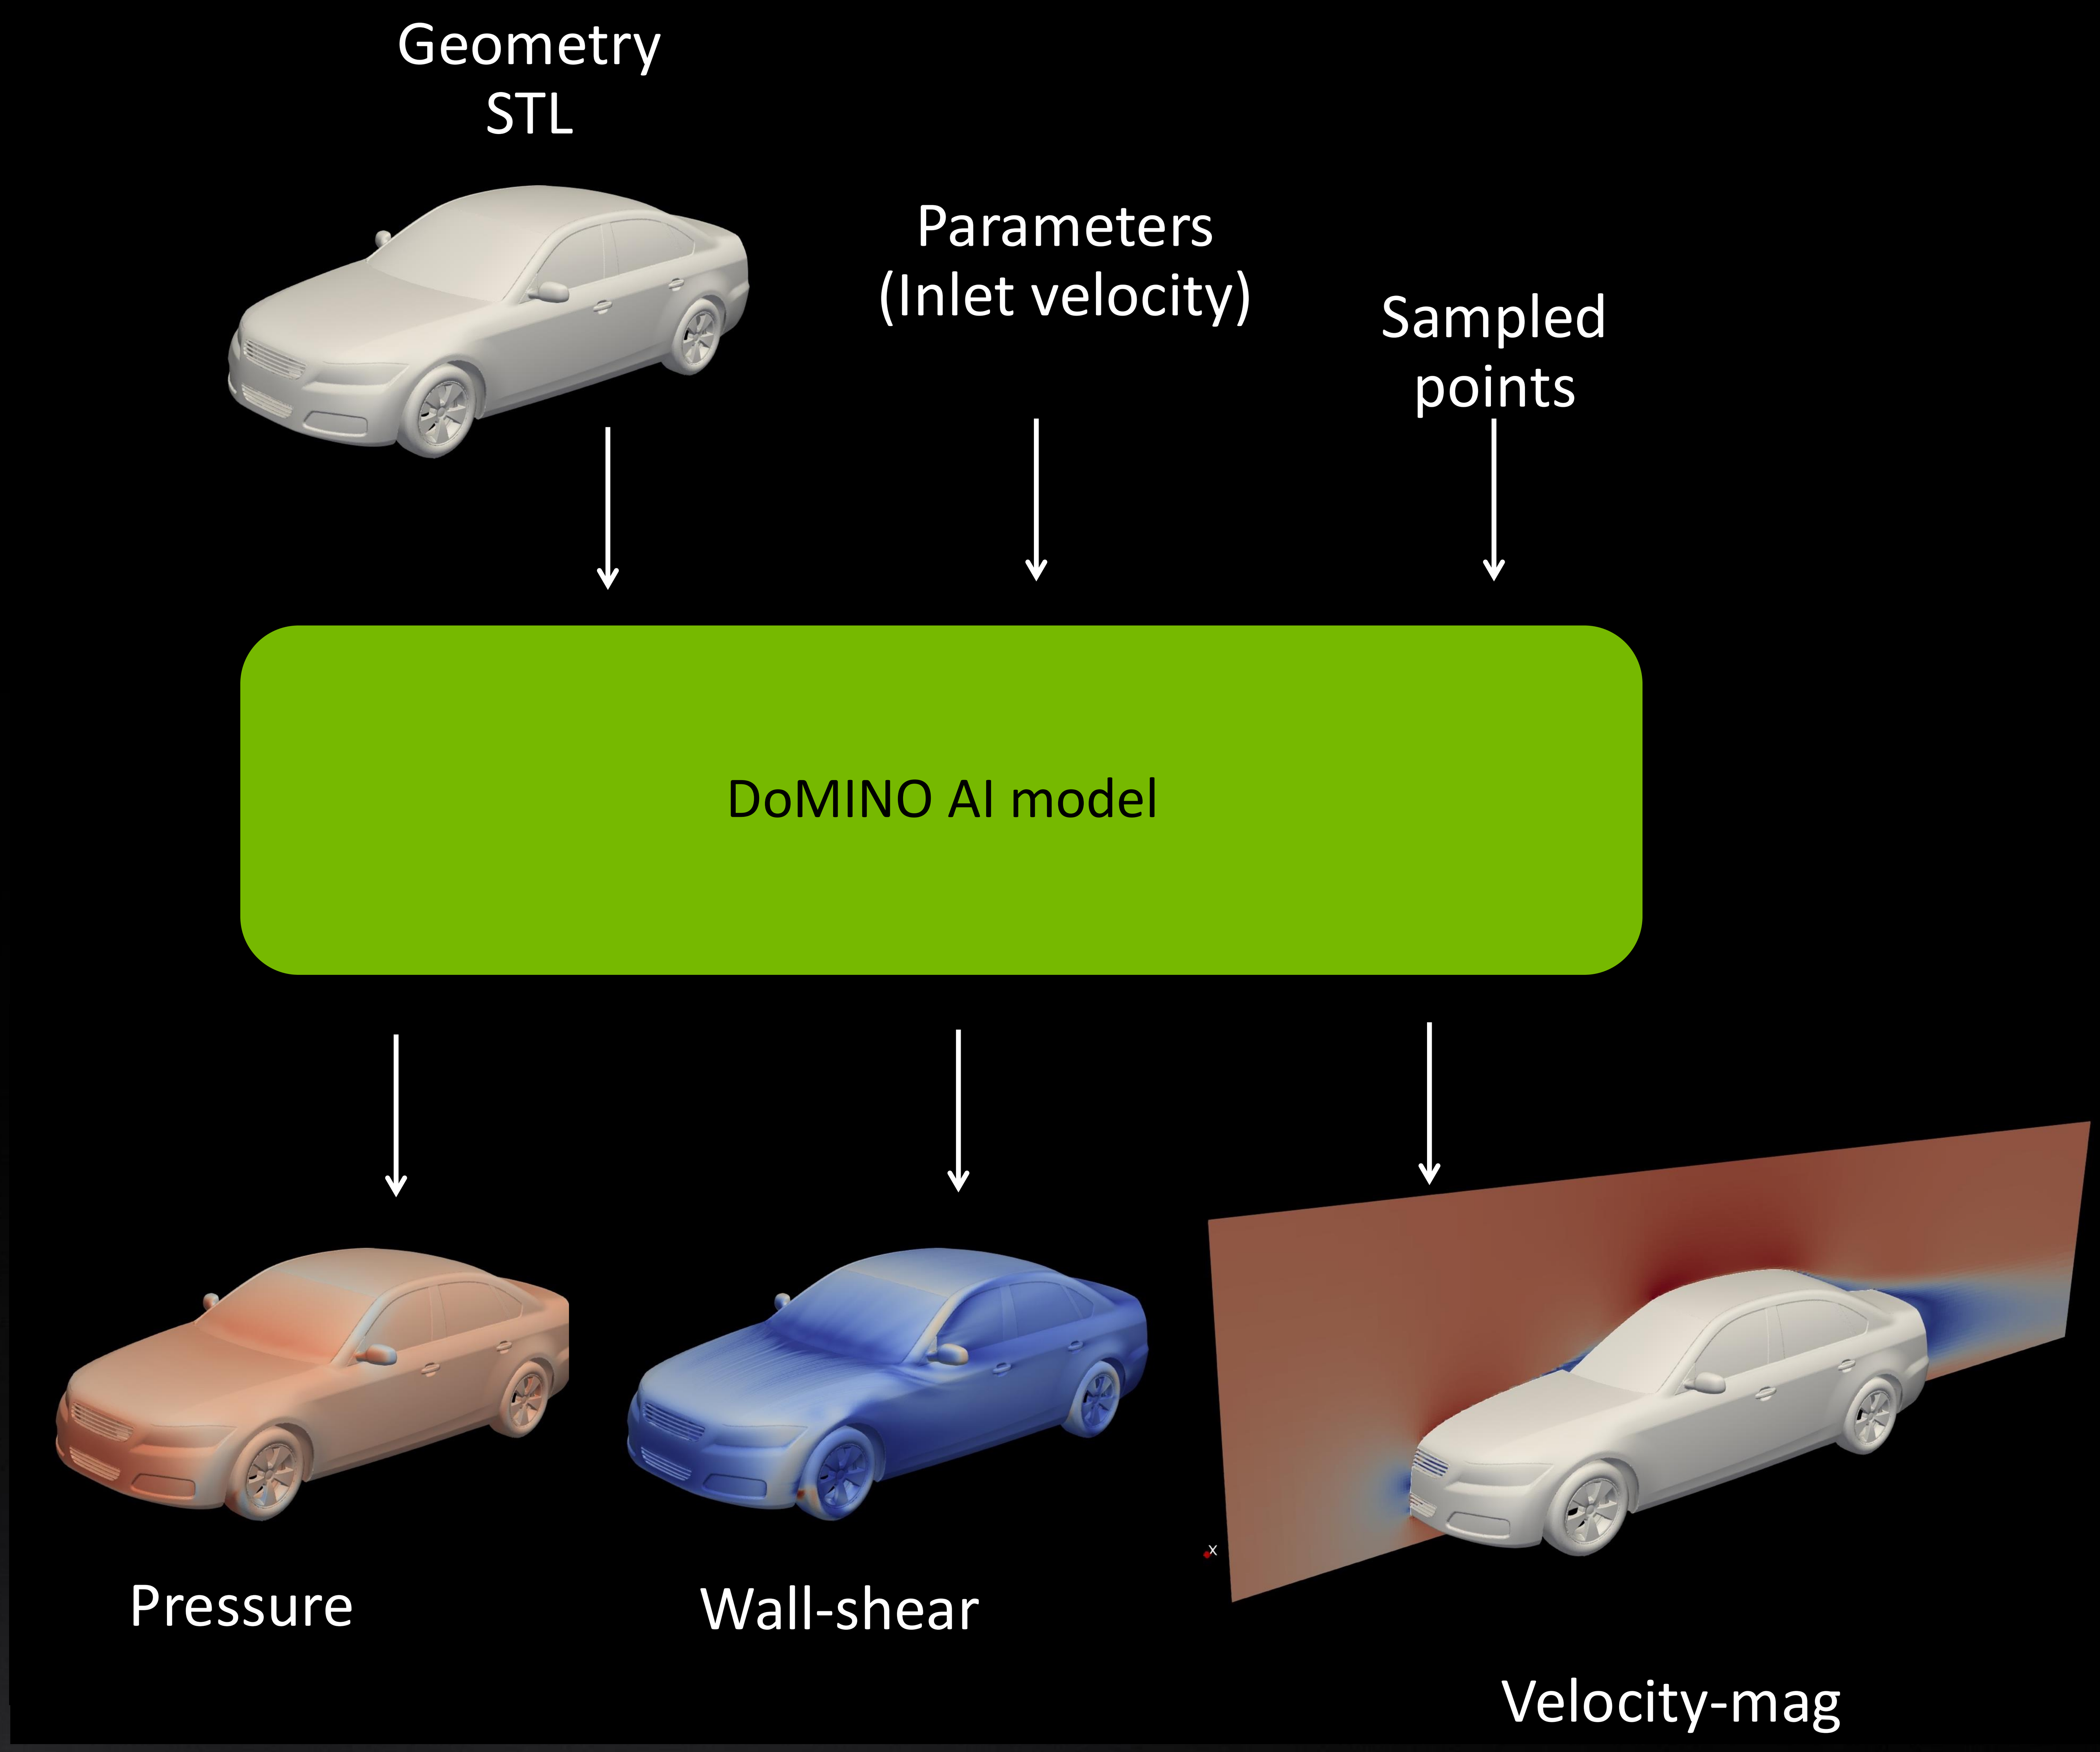

External Aerodynamics benchmarking

Utilities and workflows for benchmarking ML models for external aero

- Developed utilities and workflows for benchmarking to analyze ML models in a consistent and transparent manner using CAE specific metrics
  - Users train their ML models with the OSS DrivAerML dataset (https://caemldatasets.org/drivaerml/) with a specified train and test split
  - Trained models are used to predict on test set and written back into the VTPs and VTUs provided with the dataset
  - Use our benchmarking utilities with these files to generate CAE specific results (few examples showed below)
    - L_2 error metrics, drag force R^2 coefficients and design trends, comparisons of surface and volume contours, centerline plot on surface, line plots in different volume regions such as wake, underbody etc. for different field variables

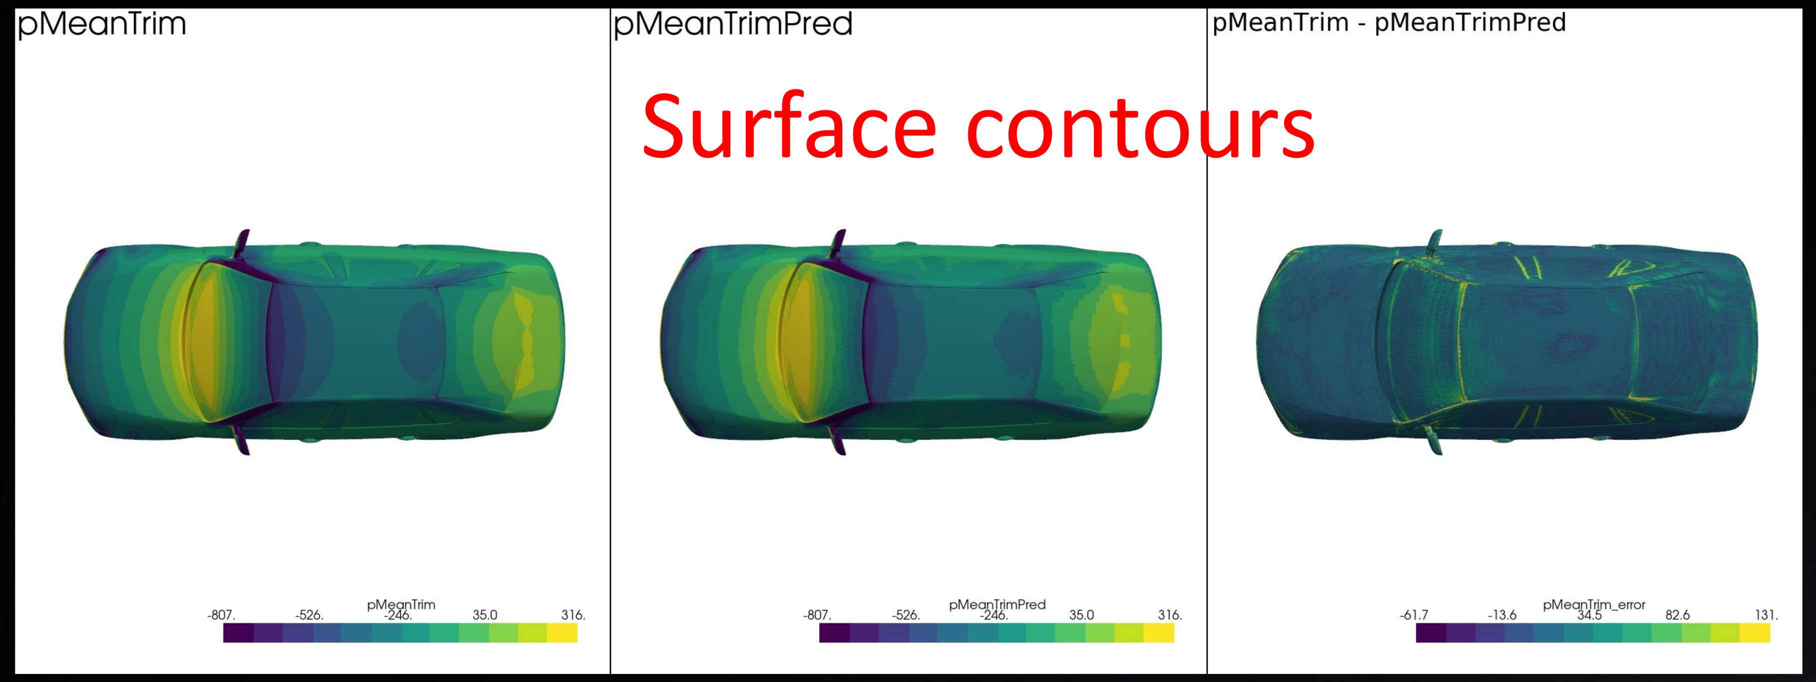

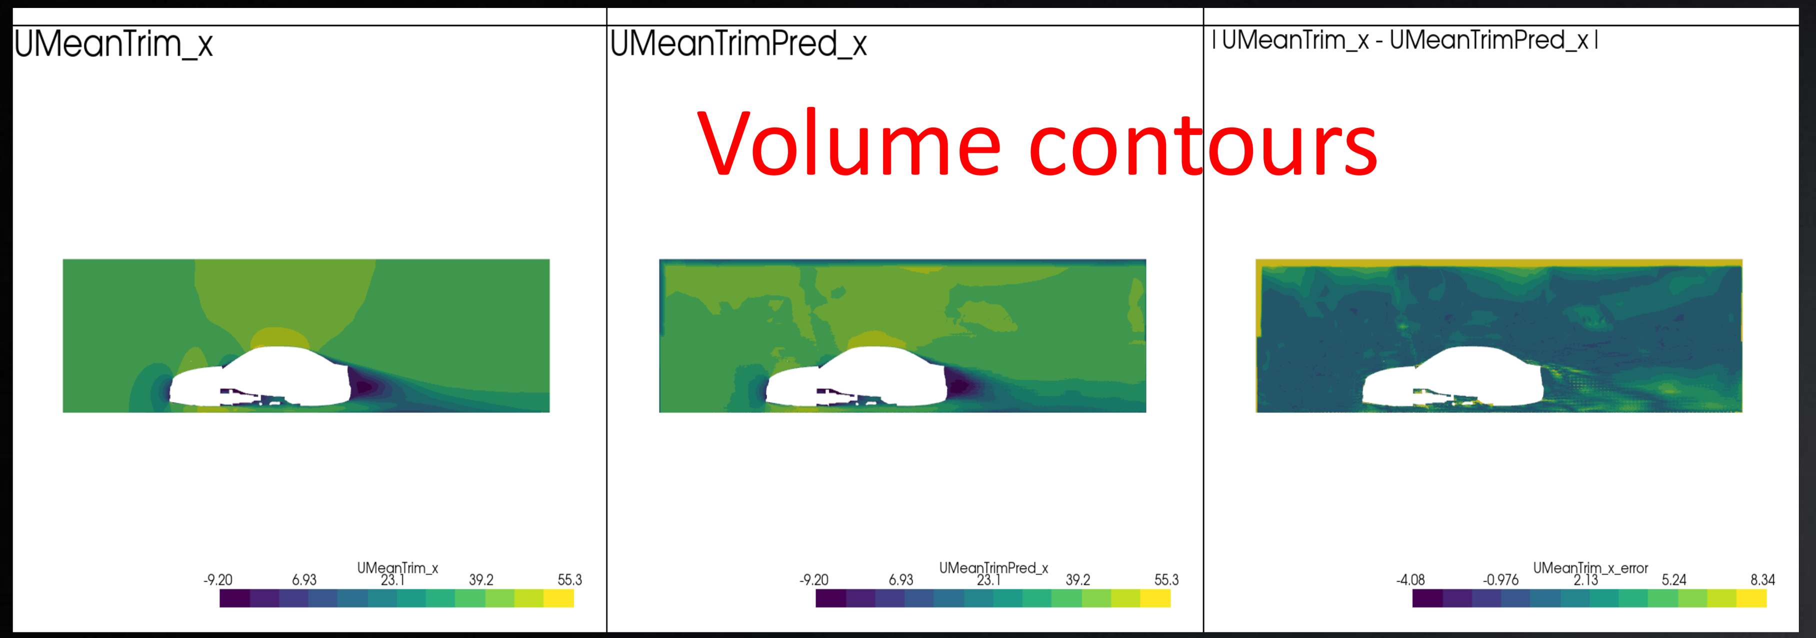

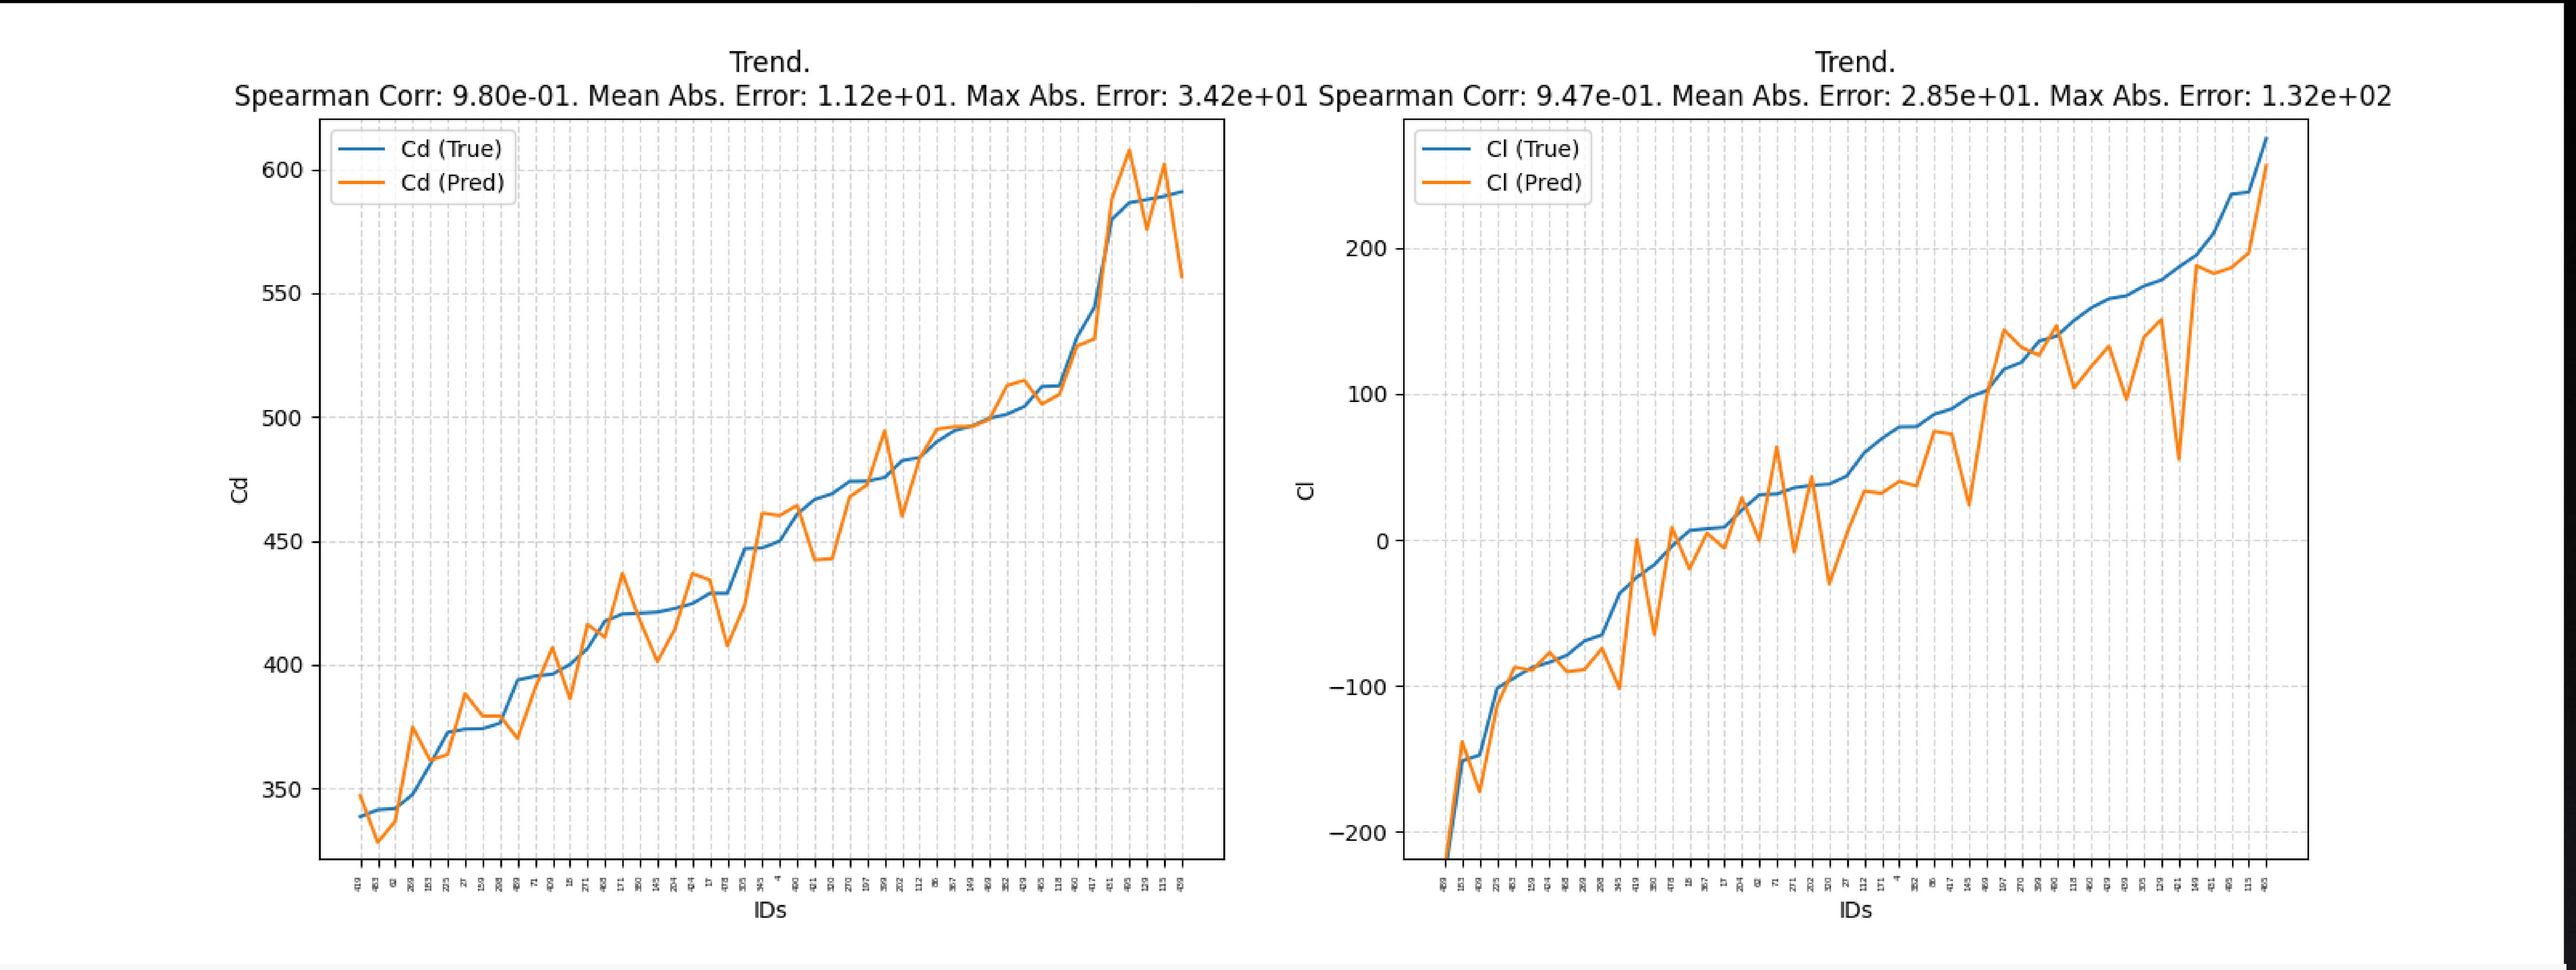

Drag and lift force design trends

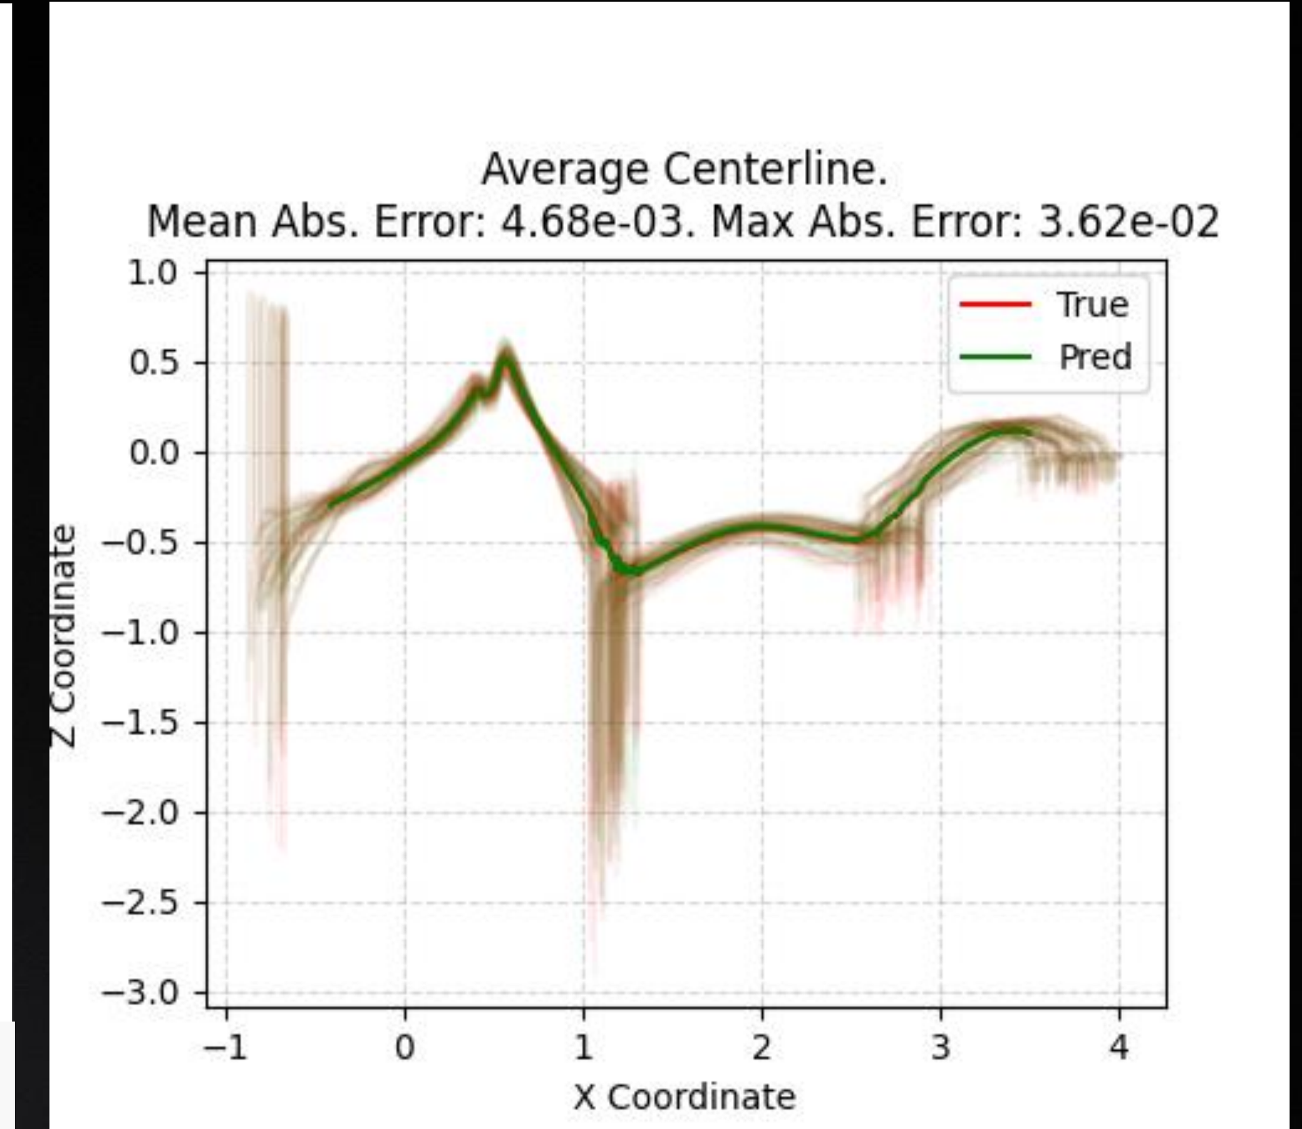

Centerline surface

NVIDIA Omniverse

Development platform for building digital twins

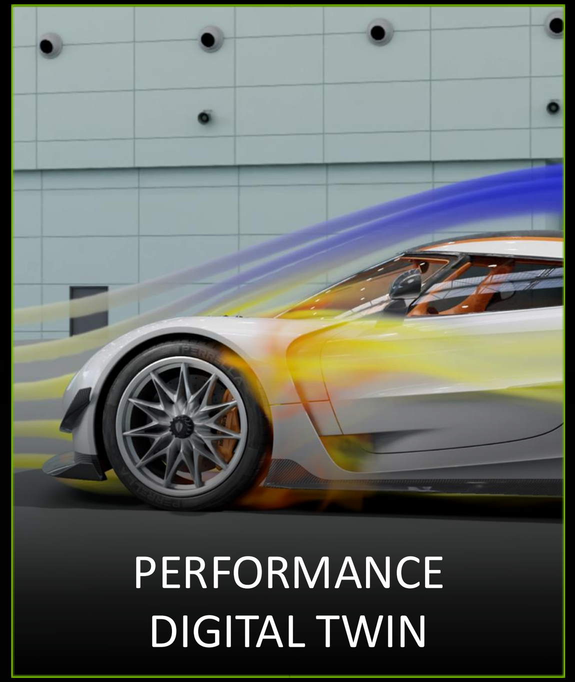

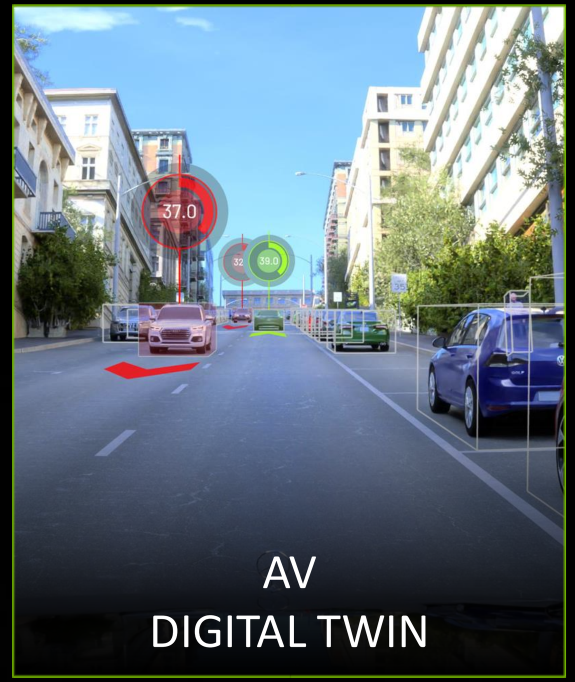

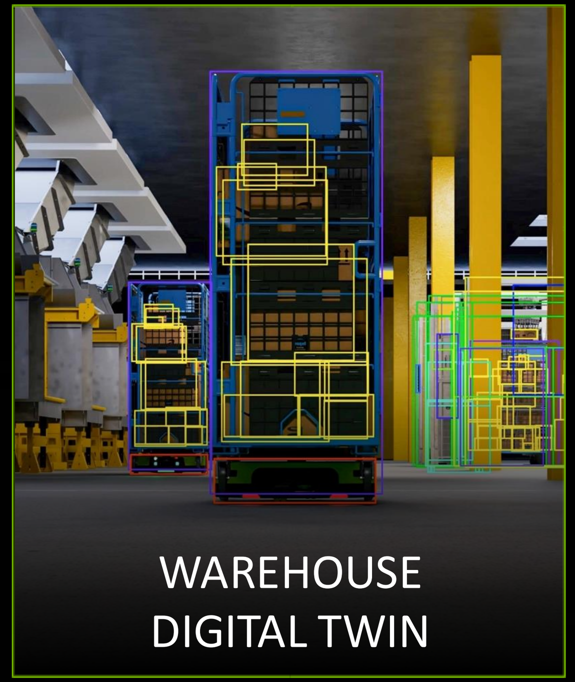

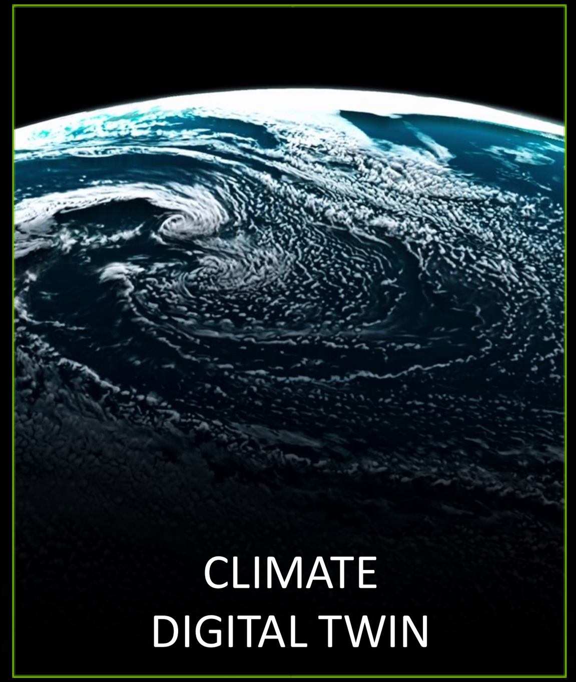

NVIDIA
Omniverse

NVIDIA Omniverse Turbocharges Self-Driving Car Development

NVIDIA Omniverse Cloud APIs deliver large-scale, high-fidelity sensor simulation, paving the path to autonomous driving.
By bringing together a rich ecosystem of simulation tools, applications, and sensors, these APIs let developers safely explore the wide variety of real-world scenarios autonomous systems will encounter. This enables vehicles to drive millions of miles in a wide range of simulated scenarios, so they hit the road running safely.

Foxconn's Robotic Factory Ecosystem Runs on NVIDIA

Foxconn, one of the world's largest makers of electronics, uses Omniverse to build their robotic factories. This lets them orchestrate robots running on NVIDIA Isaac™ to build NVIDIA AI supercomputers, which in turn train Foxconn's robots.

Omniverse Digital Twin

Real Factory

Amazon Robotics Builds Digital Twins of Warehouses in NVIDIA Omniverse

Amazon has over 200 robotics facilities that handle millions of packages each day. Using NVIDIA Omniverse and Isaac Sim, Amazon Robotics is building AI- enabled digital twins of its warehouses to better optimize warehouse design and flow, and train more intelligent robotic solutions.

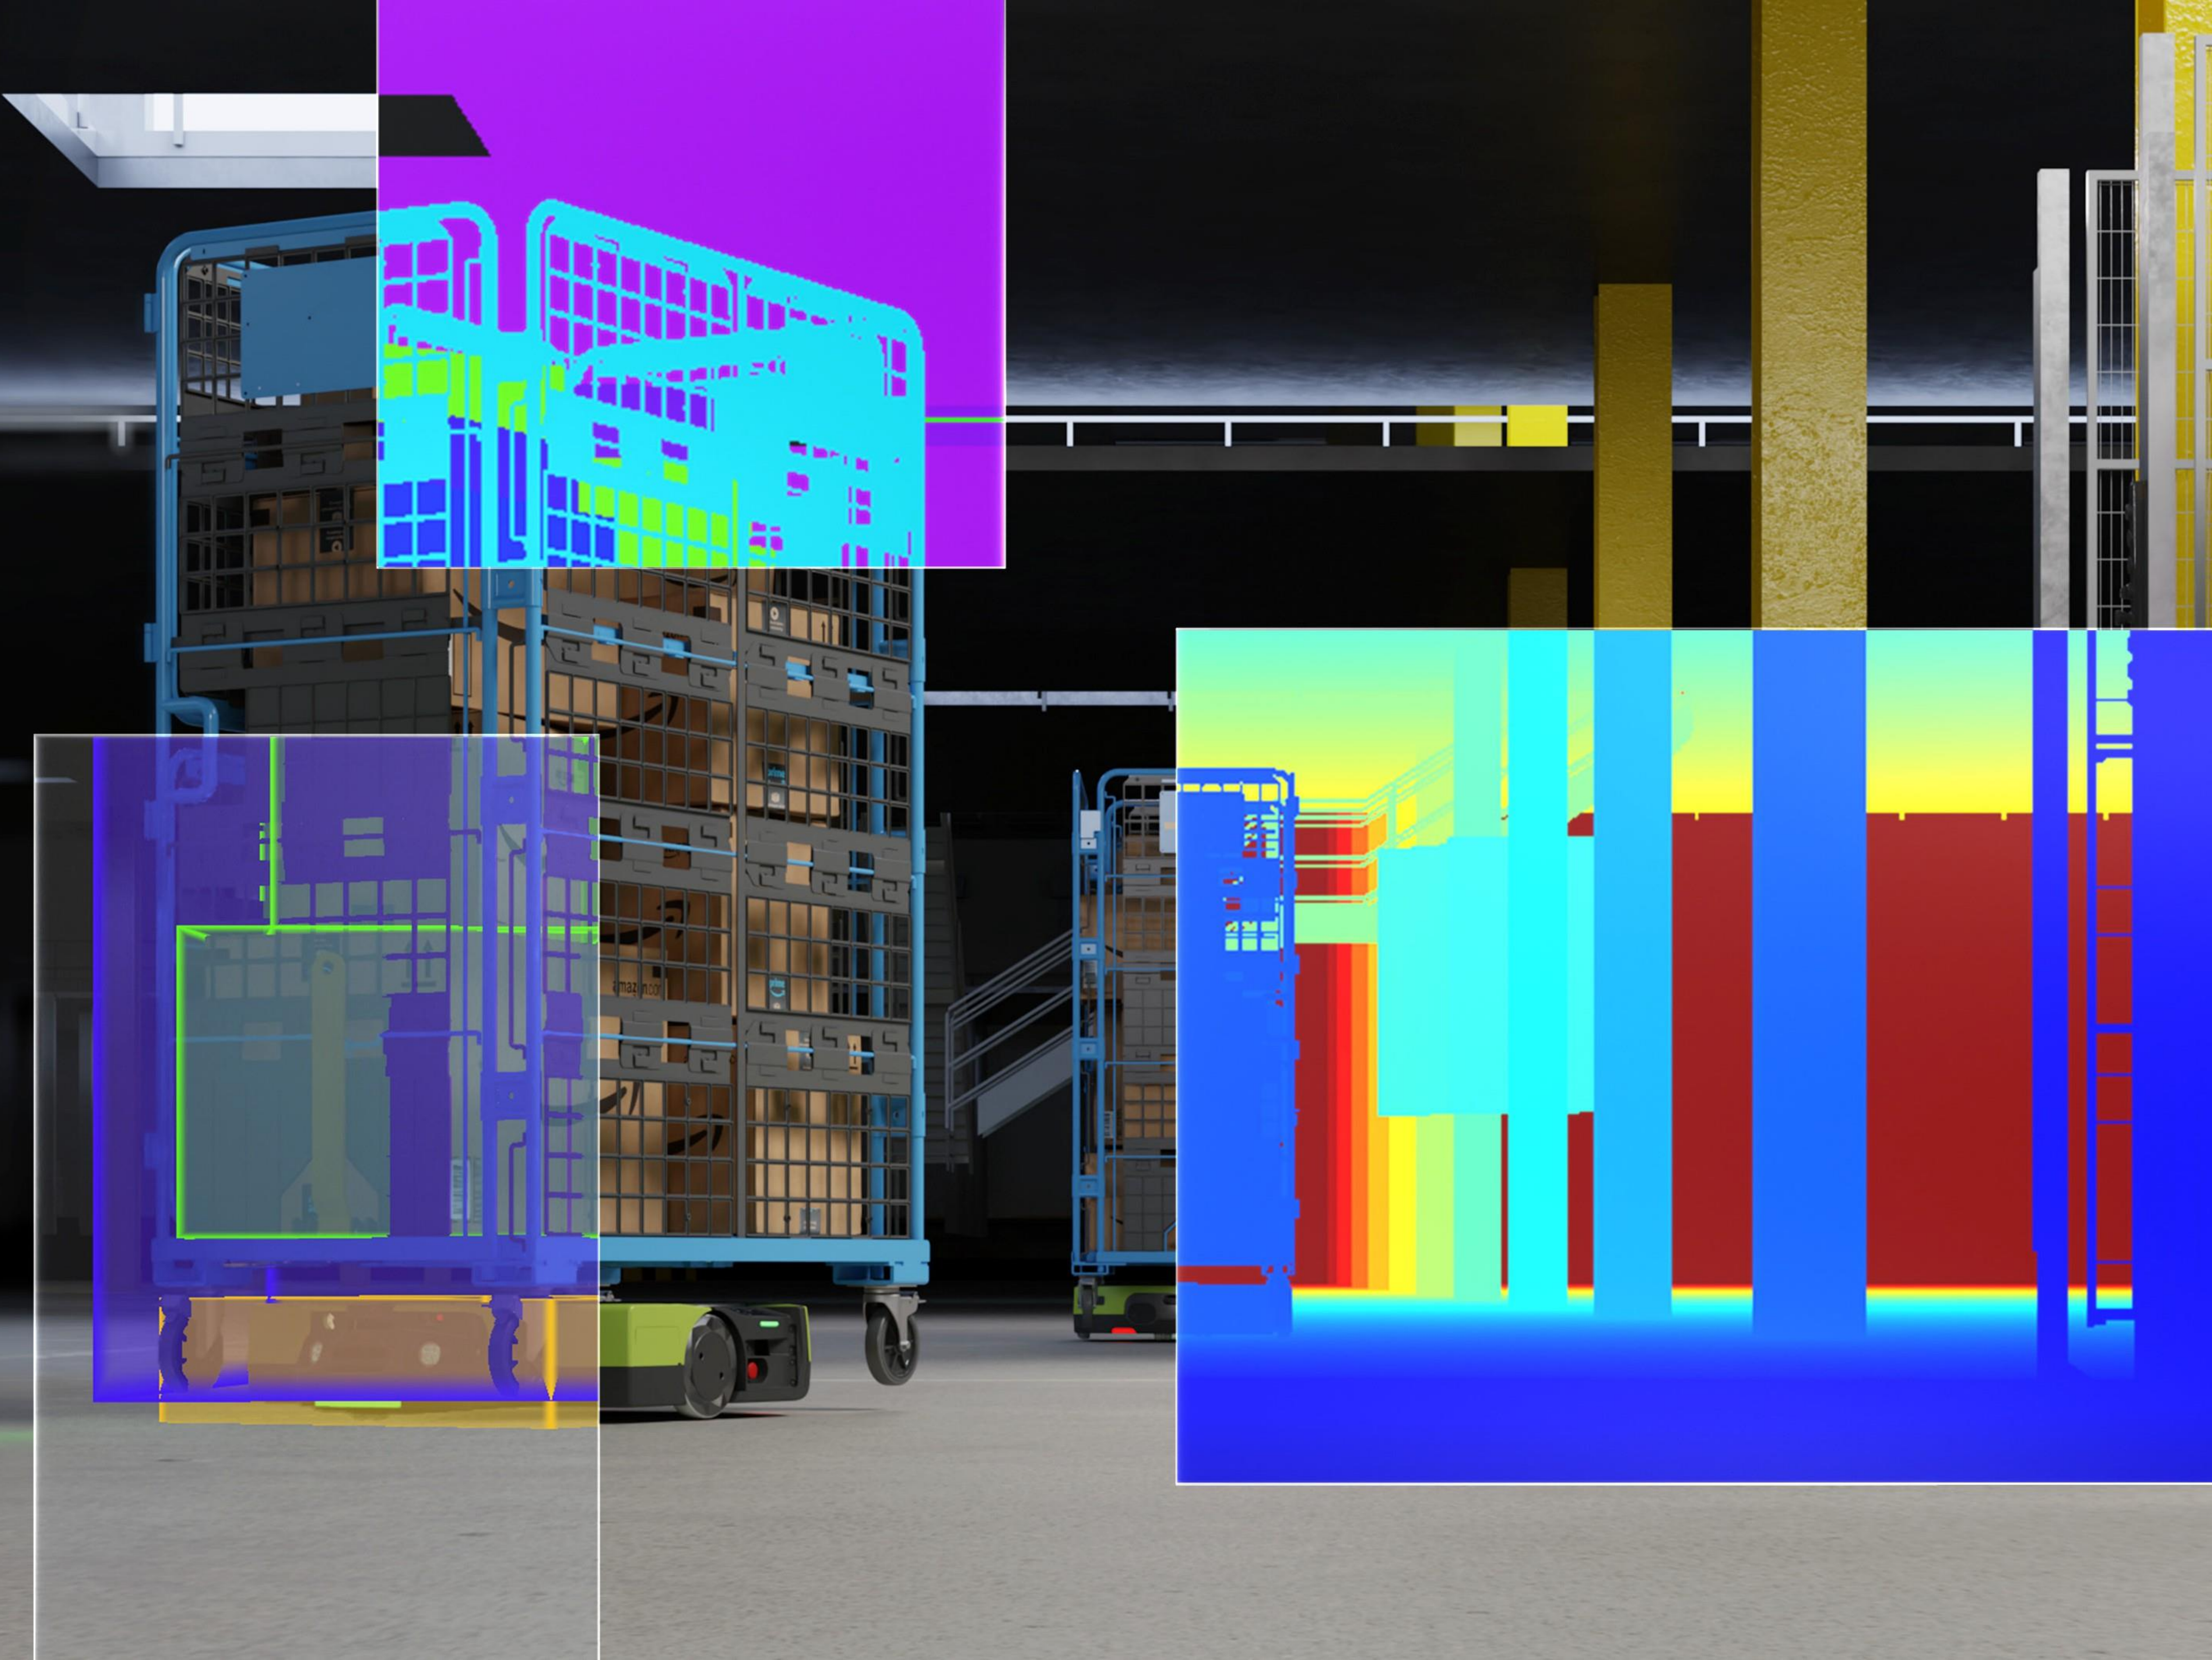

Amazon Robotics

Digital Twin: Actionable Results an Actionable Time

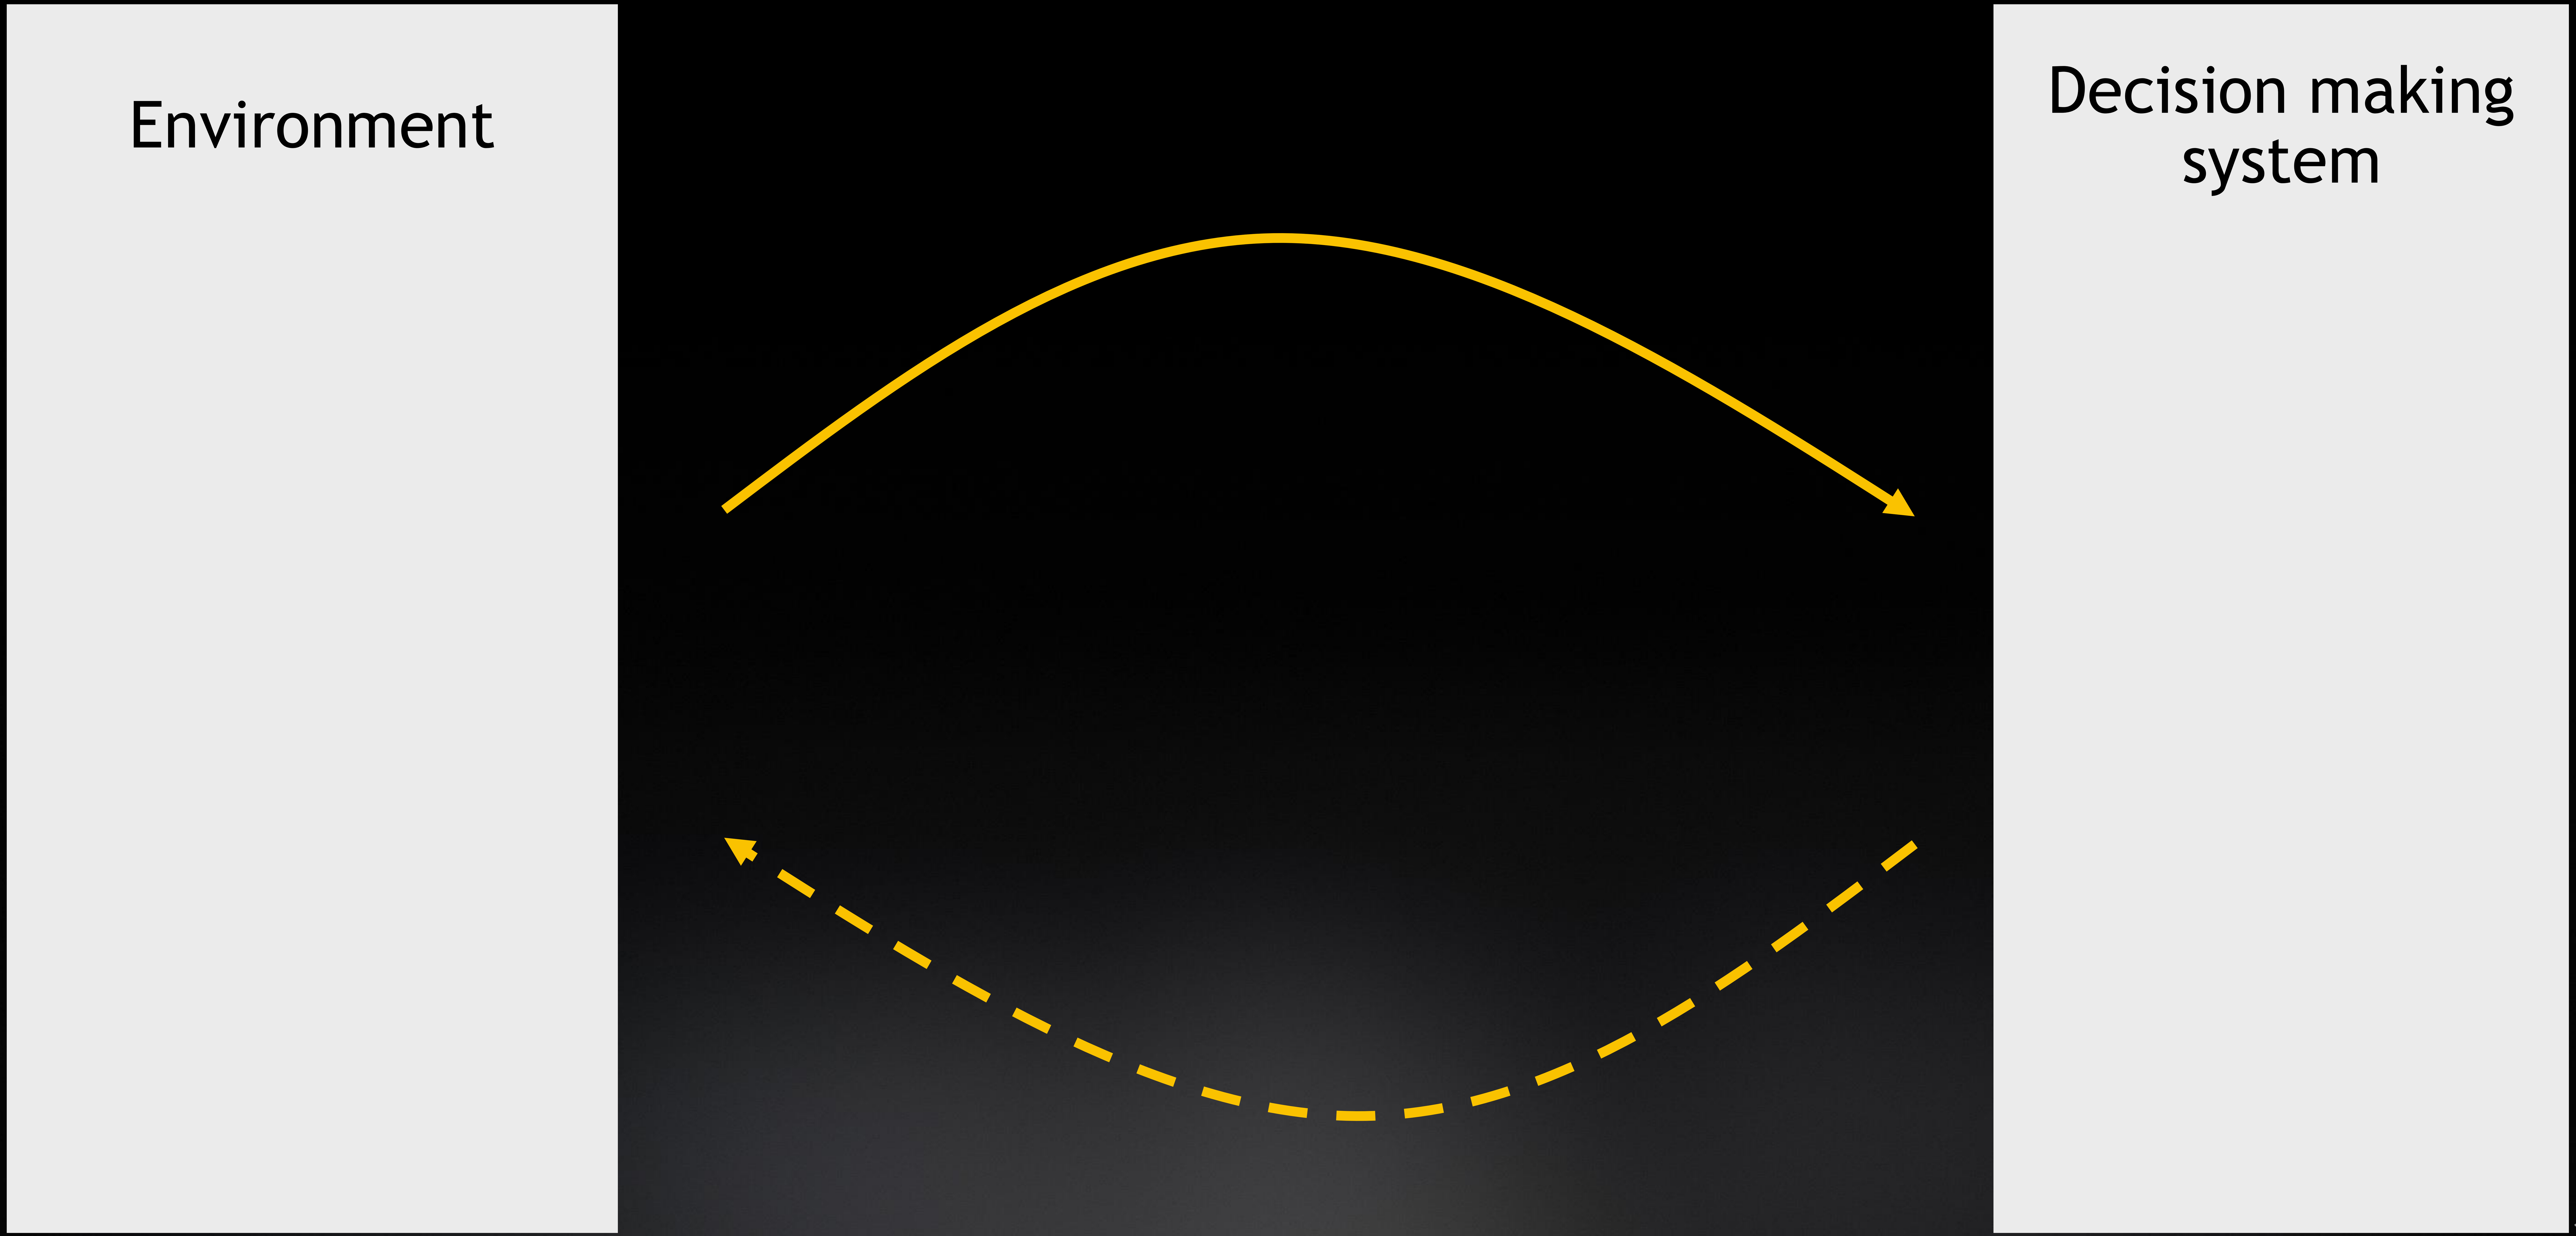

Keeping the twin in sync with Reality

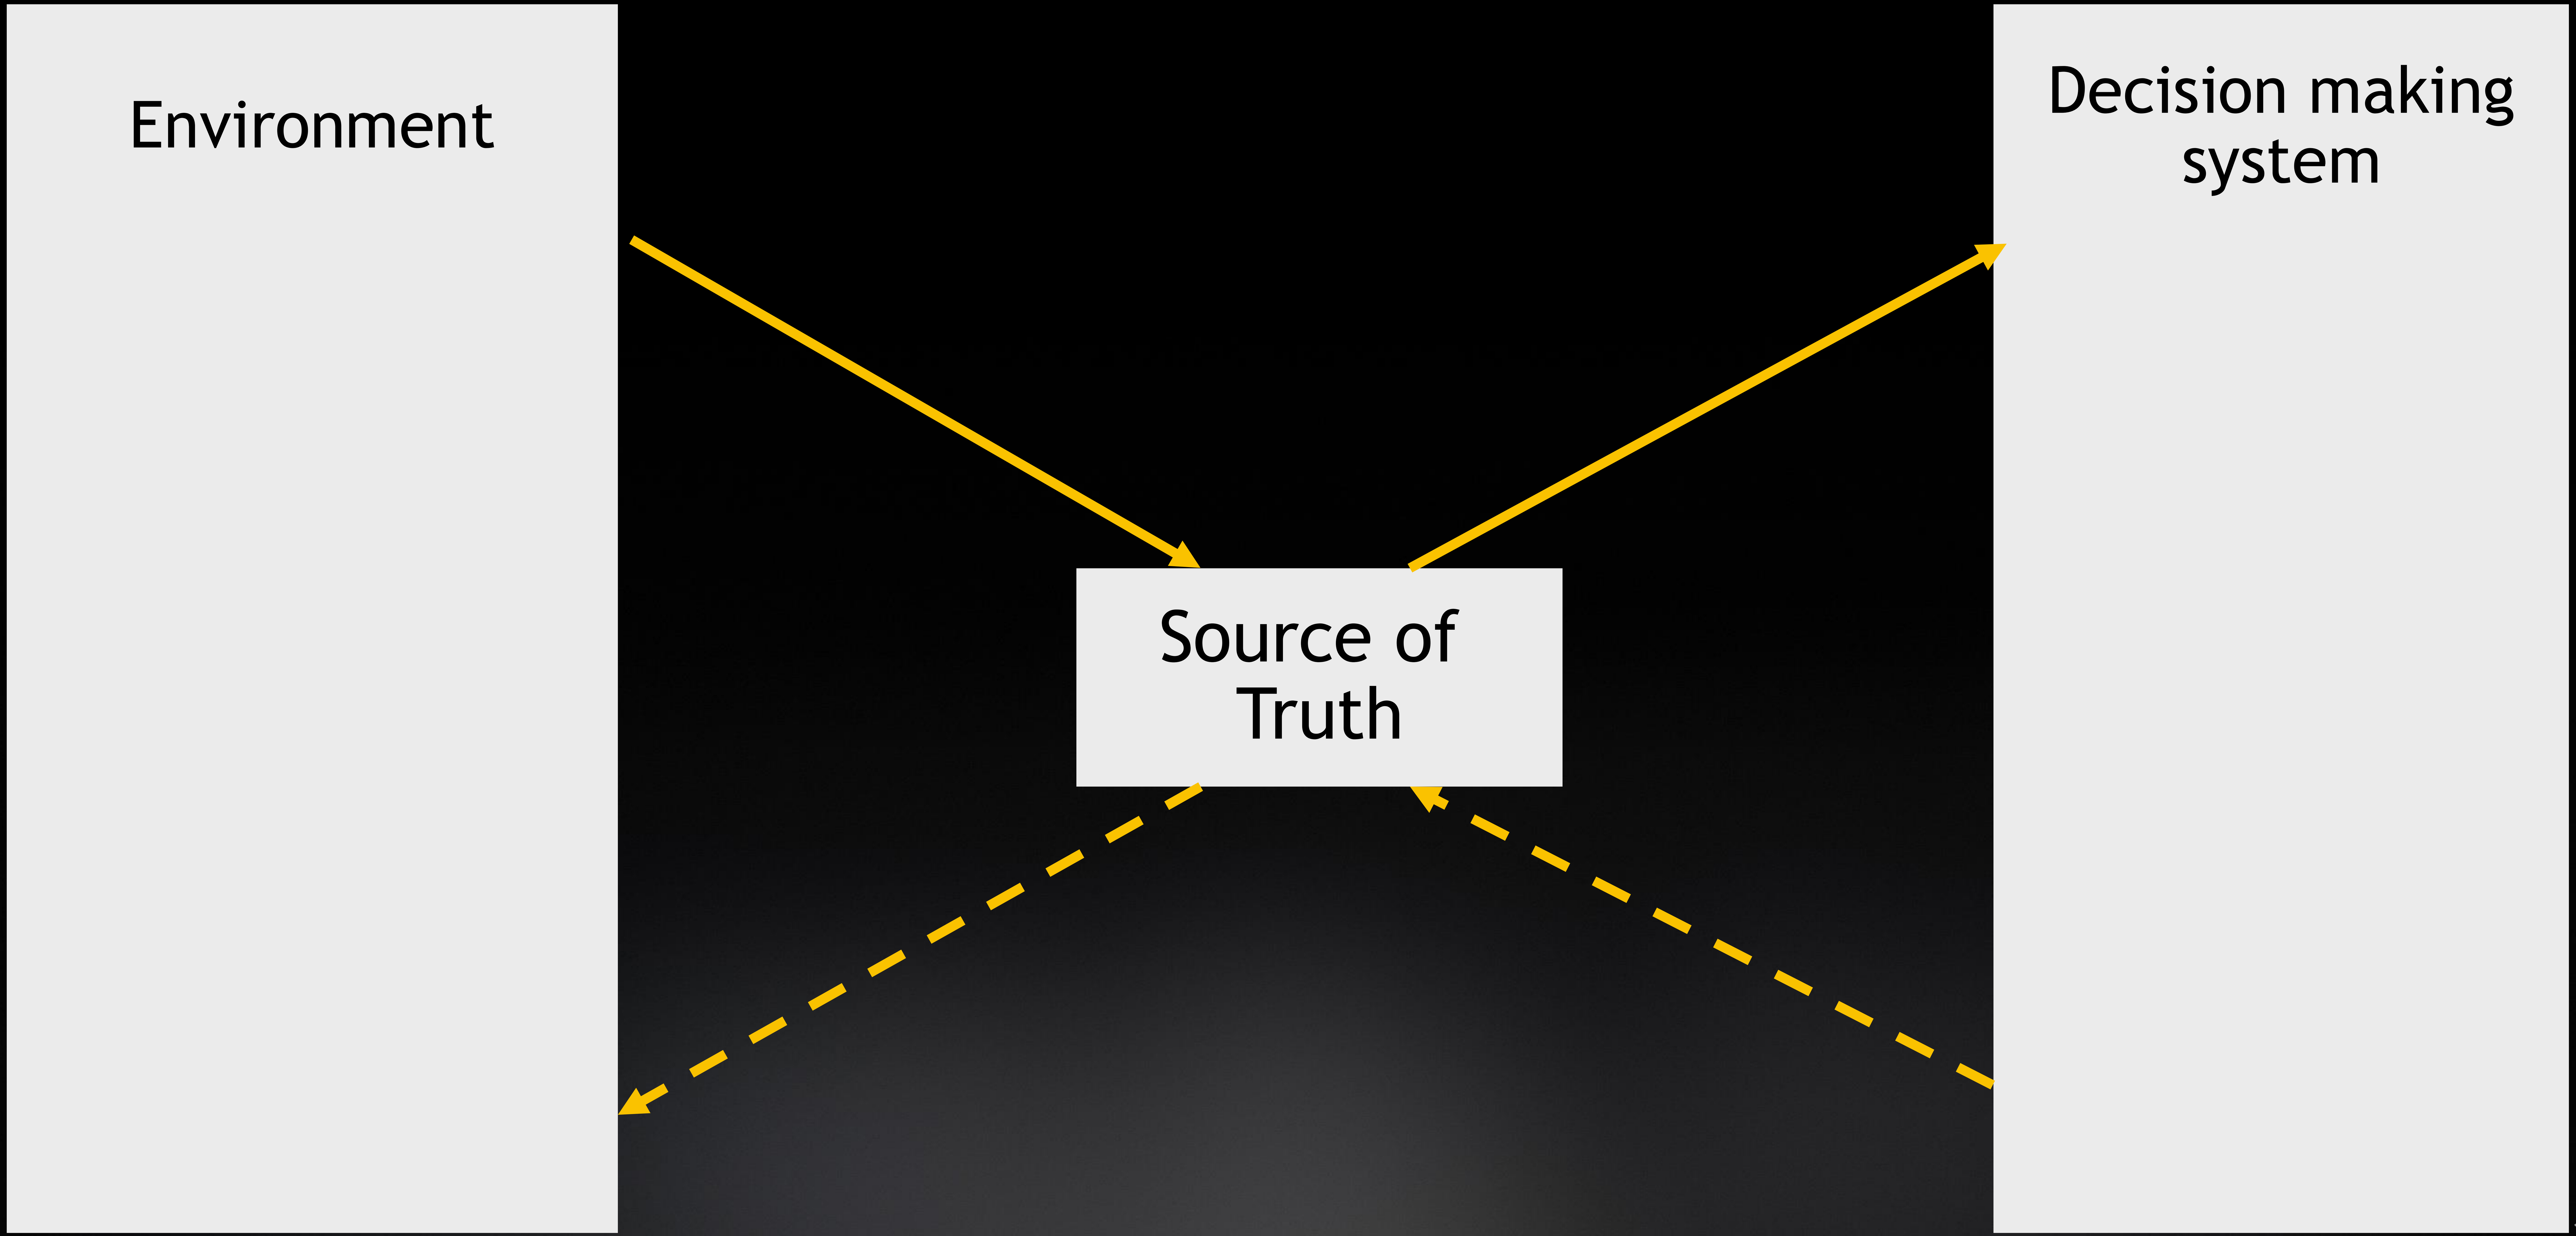

Digital twin at realistic complexity

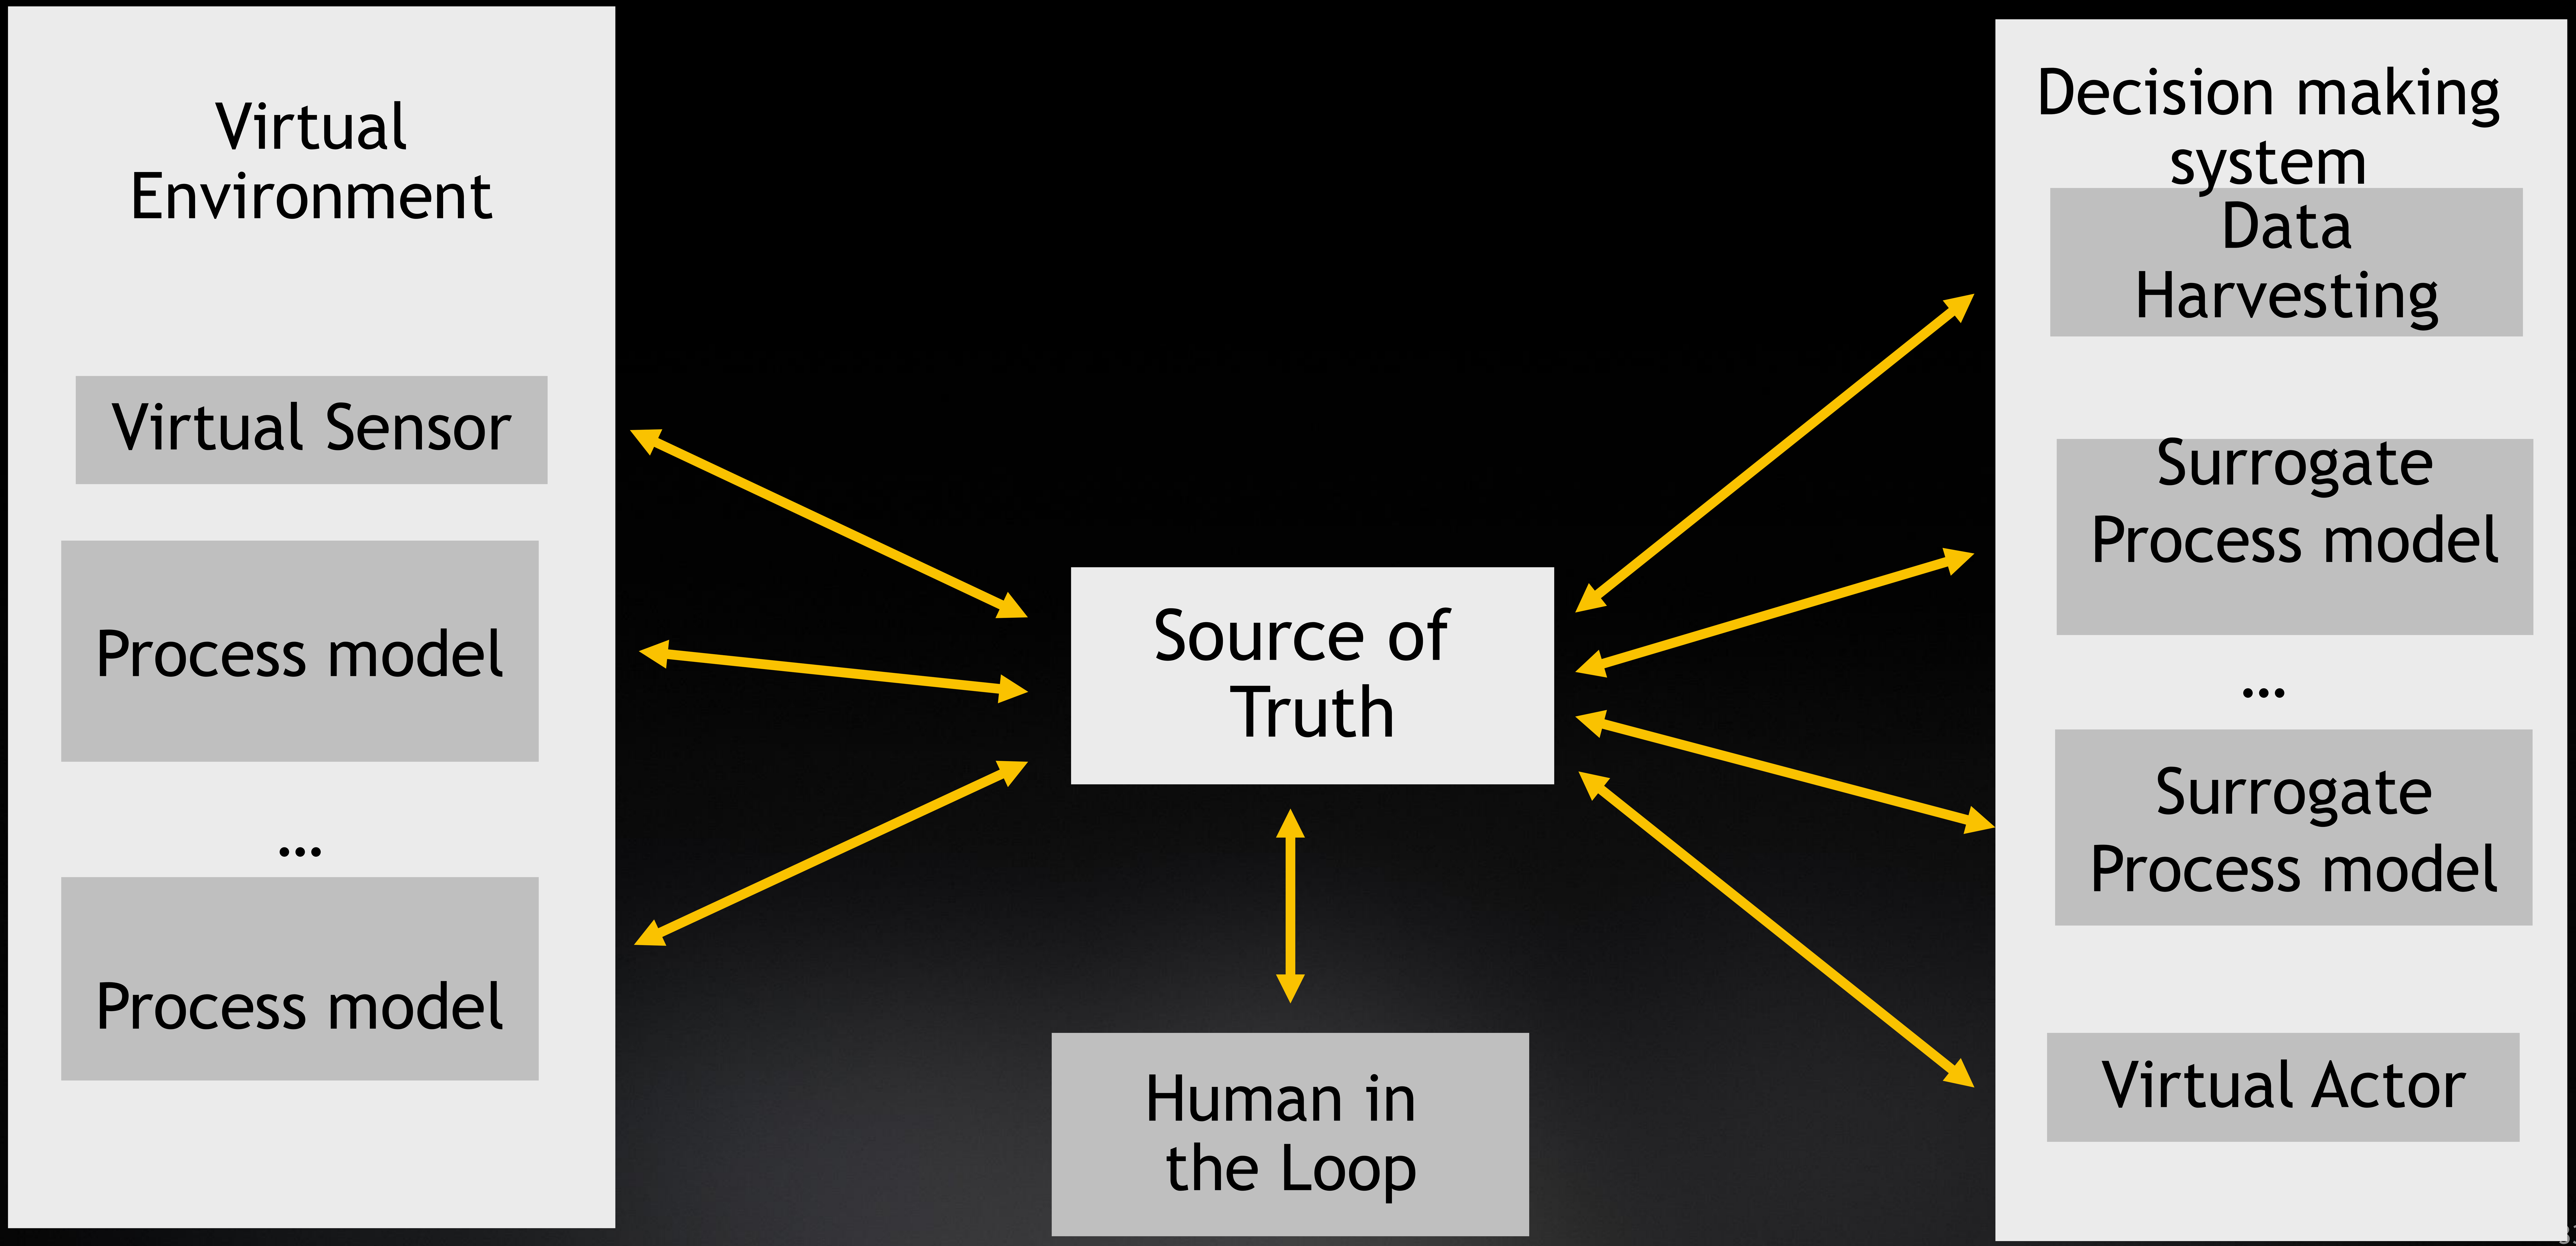

Omniverse: Platform for Building digital Twins

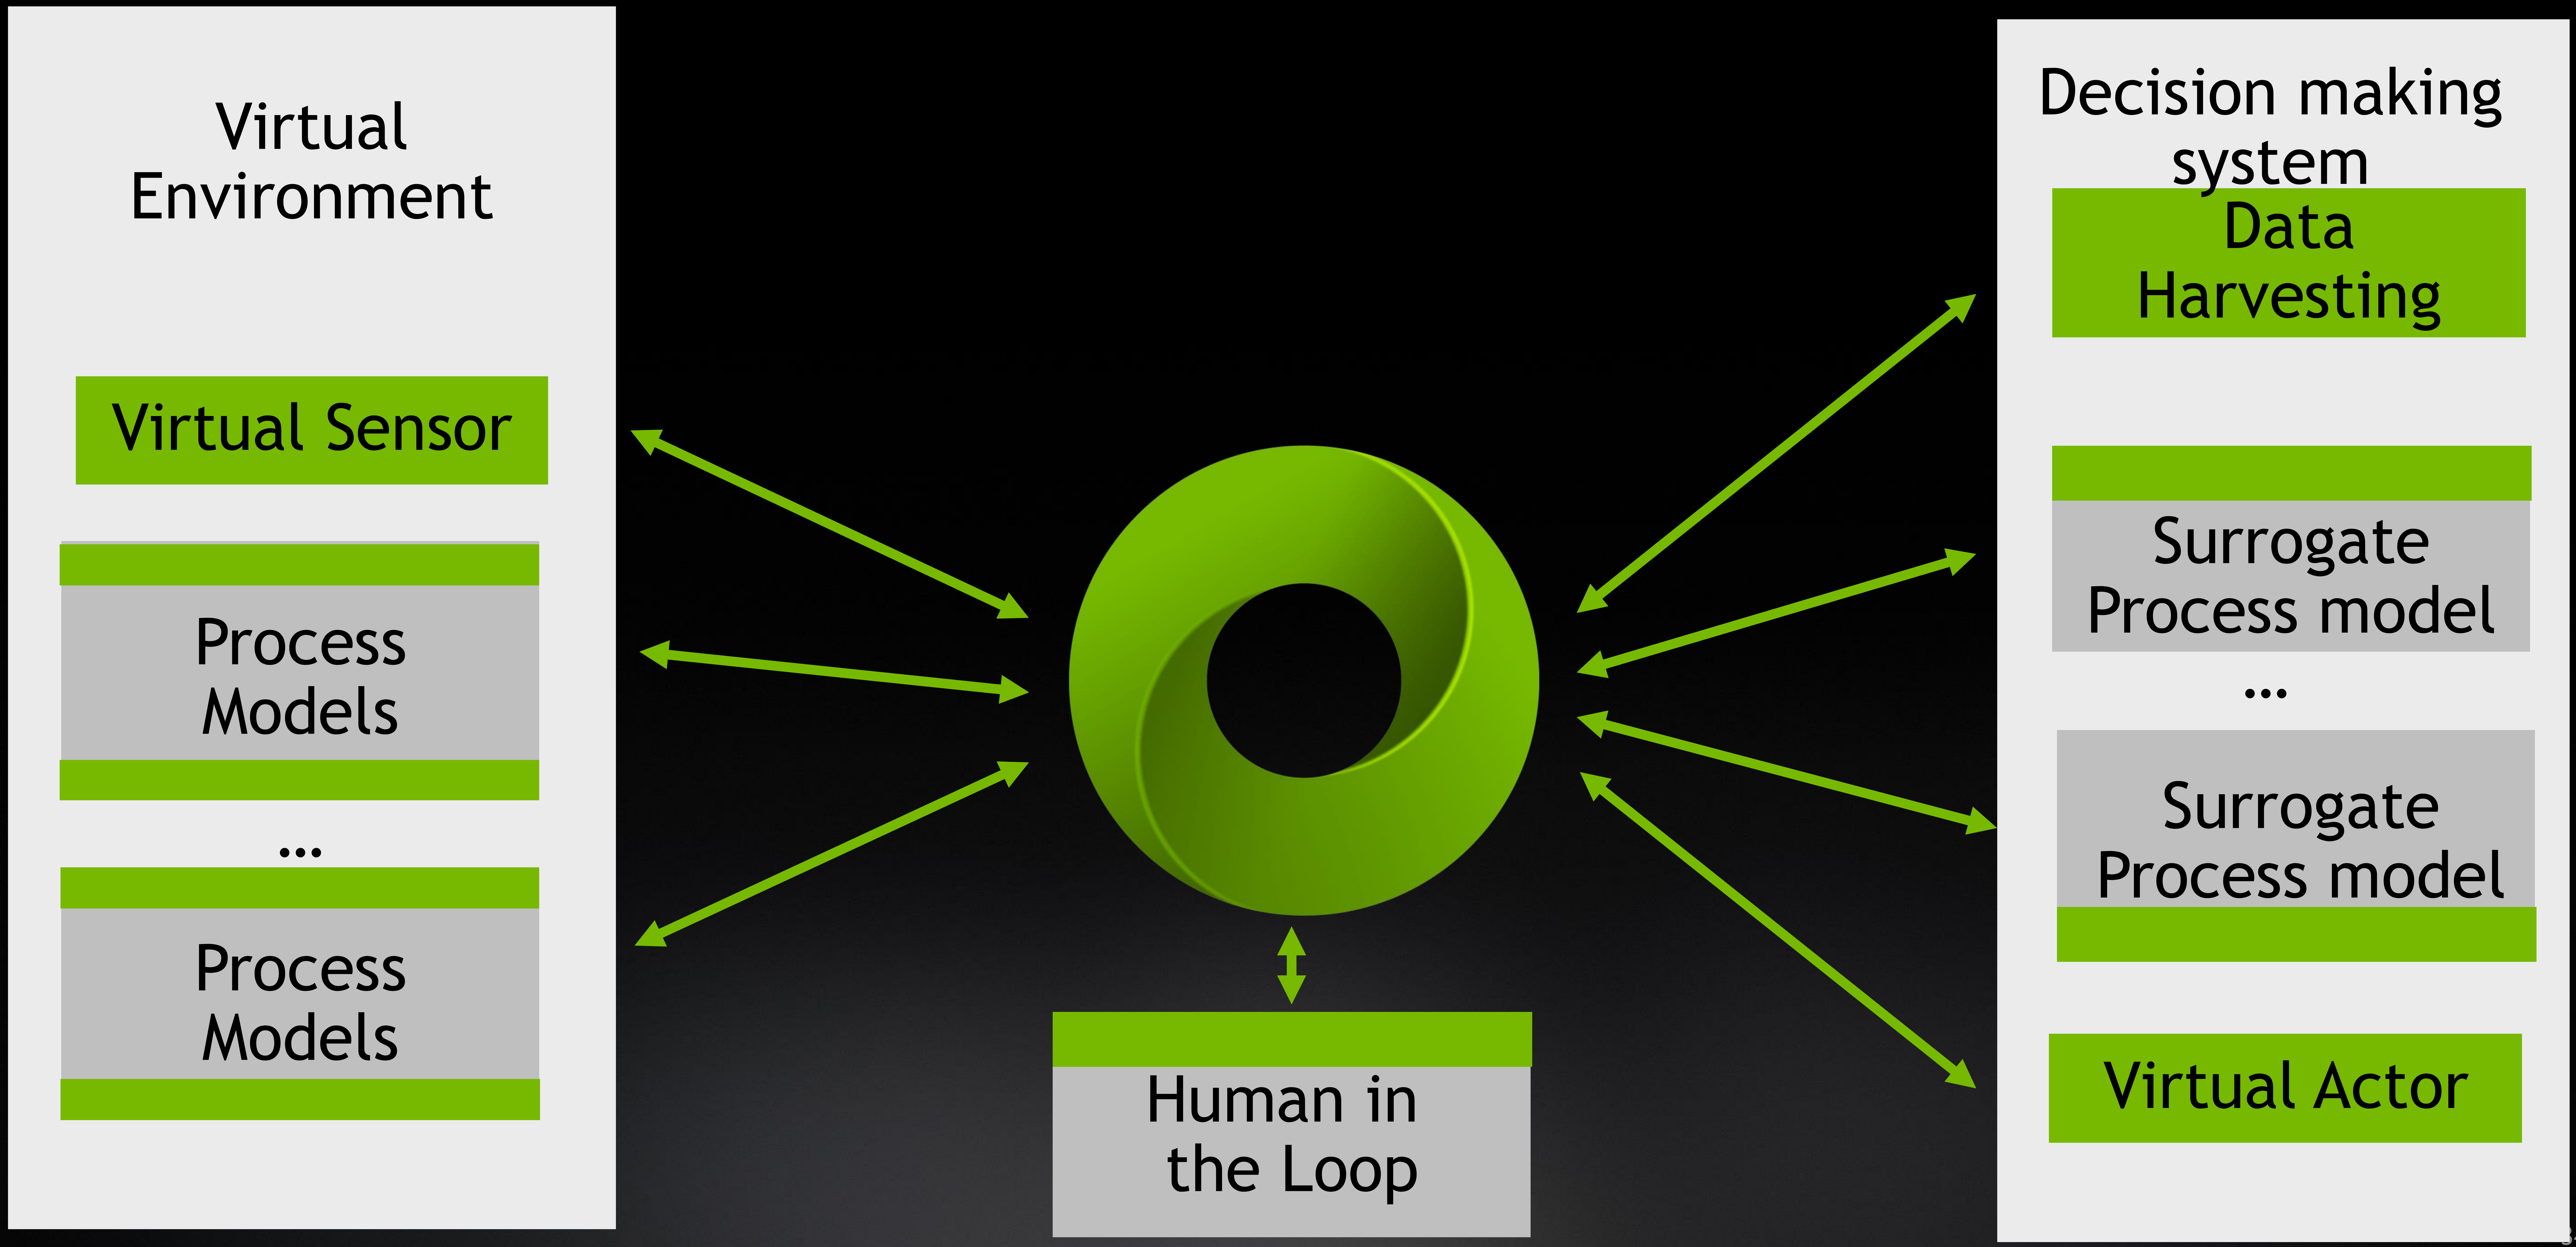

Advanced Tools and technologies

Foundational Platform Components

NUCLEUS

Source of truth

CONNECT

Coupling

KIT

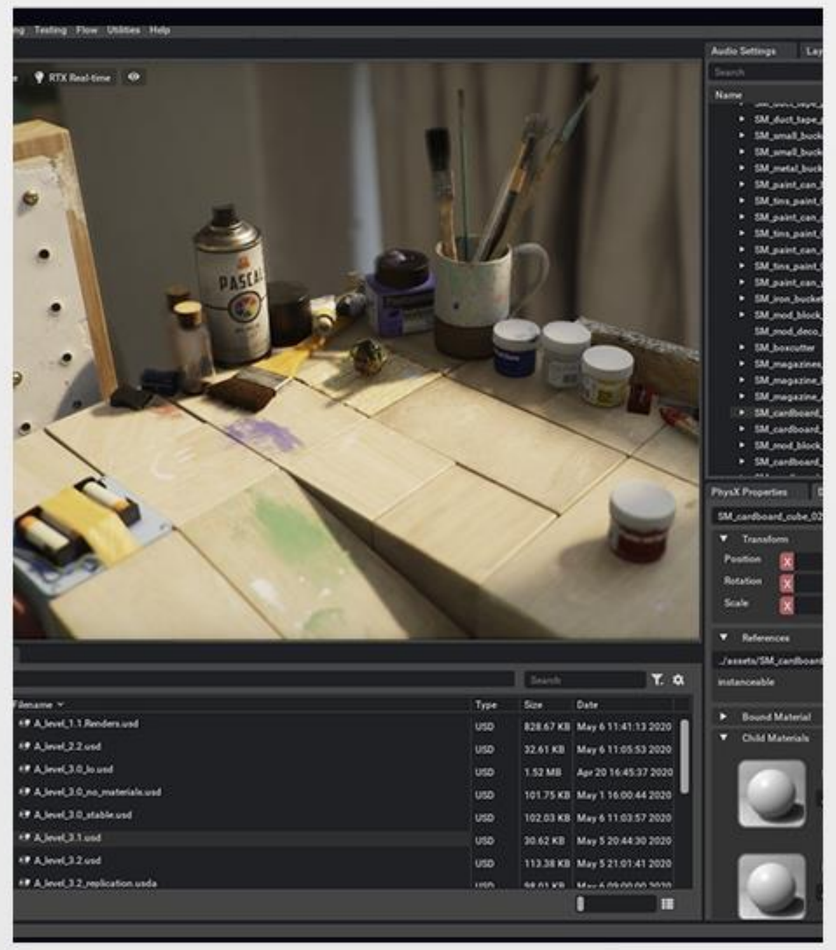

Application API
User experience

SIMULATION

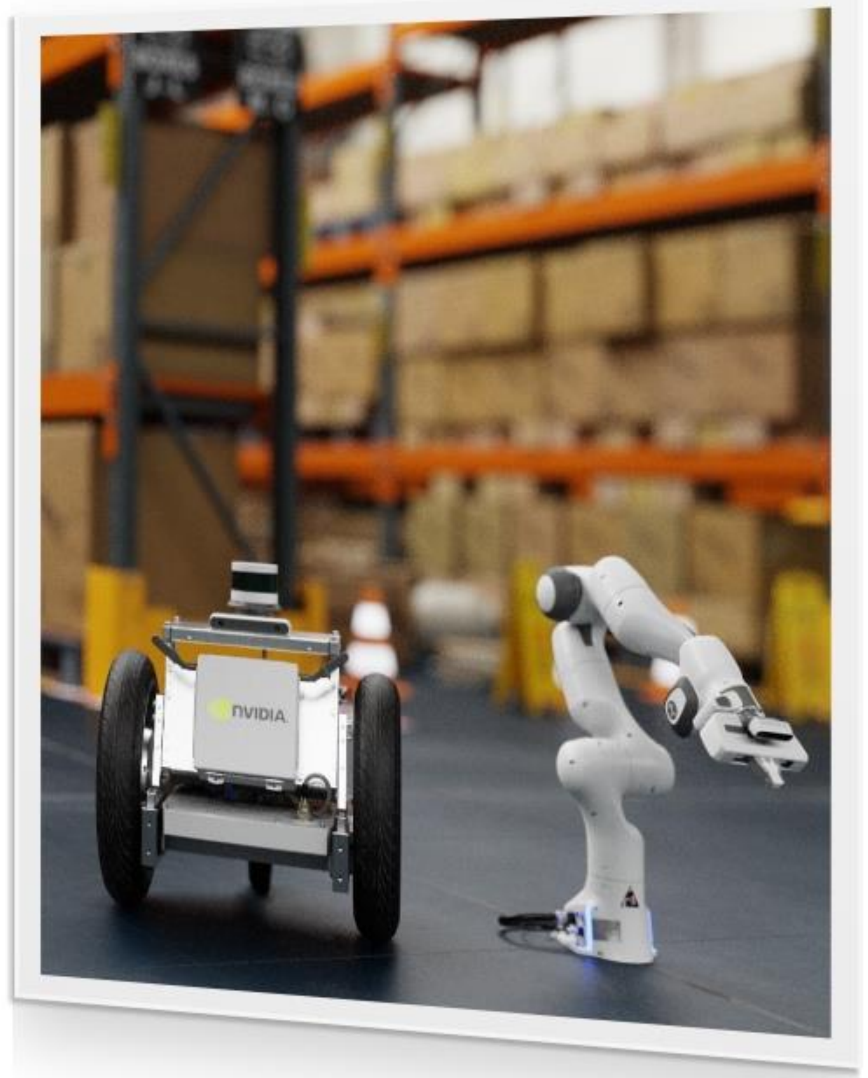

Virtual Actor

RTX RENDERER

Virtual Sensor

PhysicsNeMo Surrogate Models for Transient CFD Initialization

- Many high-fidelity automotive aerodynamics simulations require transient CFD:
  - LES/DDES turbulence modeling
  - Massively separated flows, base drag prediction
- Traditional initializations either:
  - Slow (steady RANS) → 2-40 hours additional compute
  - Inaccurate (uniform / potential flow) → much slower transient convergence
- ML-Enhanced workflow:
  - **~2x faster convergence** vs. uniform/potential flow
  - **Uses existing solvers** for the subsequent transient solve – no need to redo costly validation studies
  - **Fast initialization** (1-10 minutes)

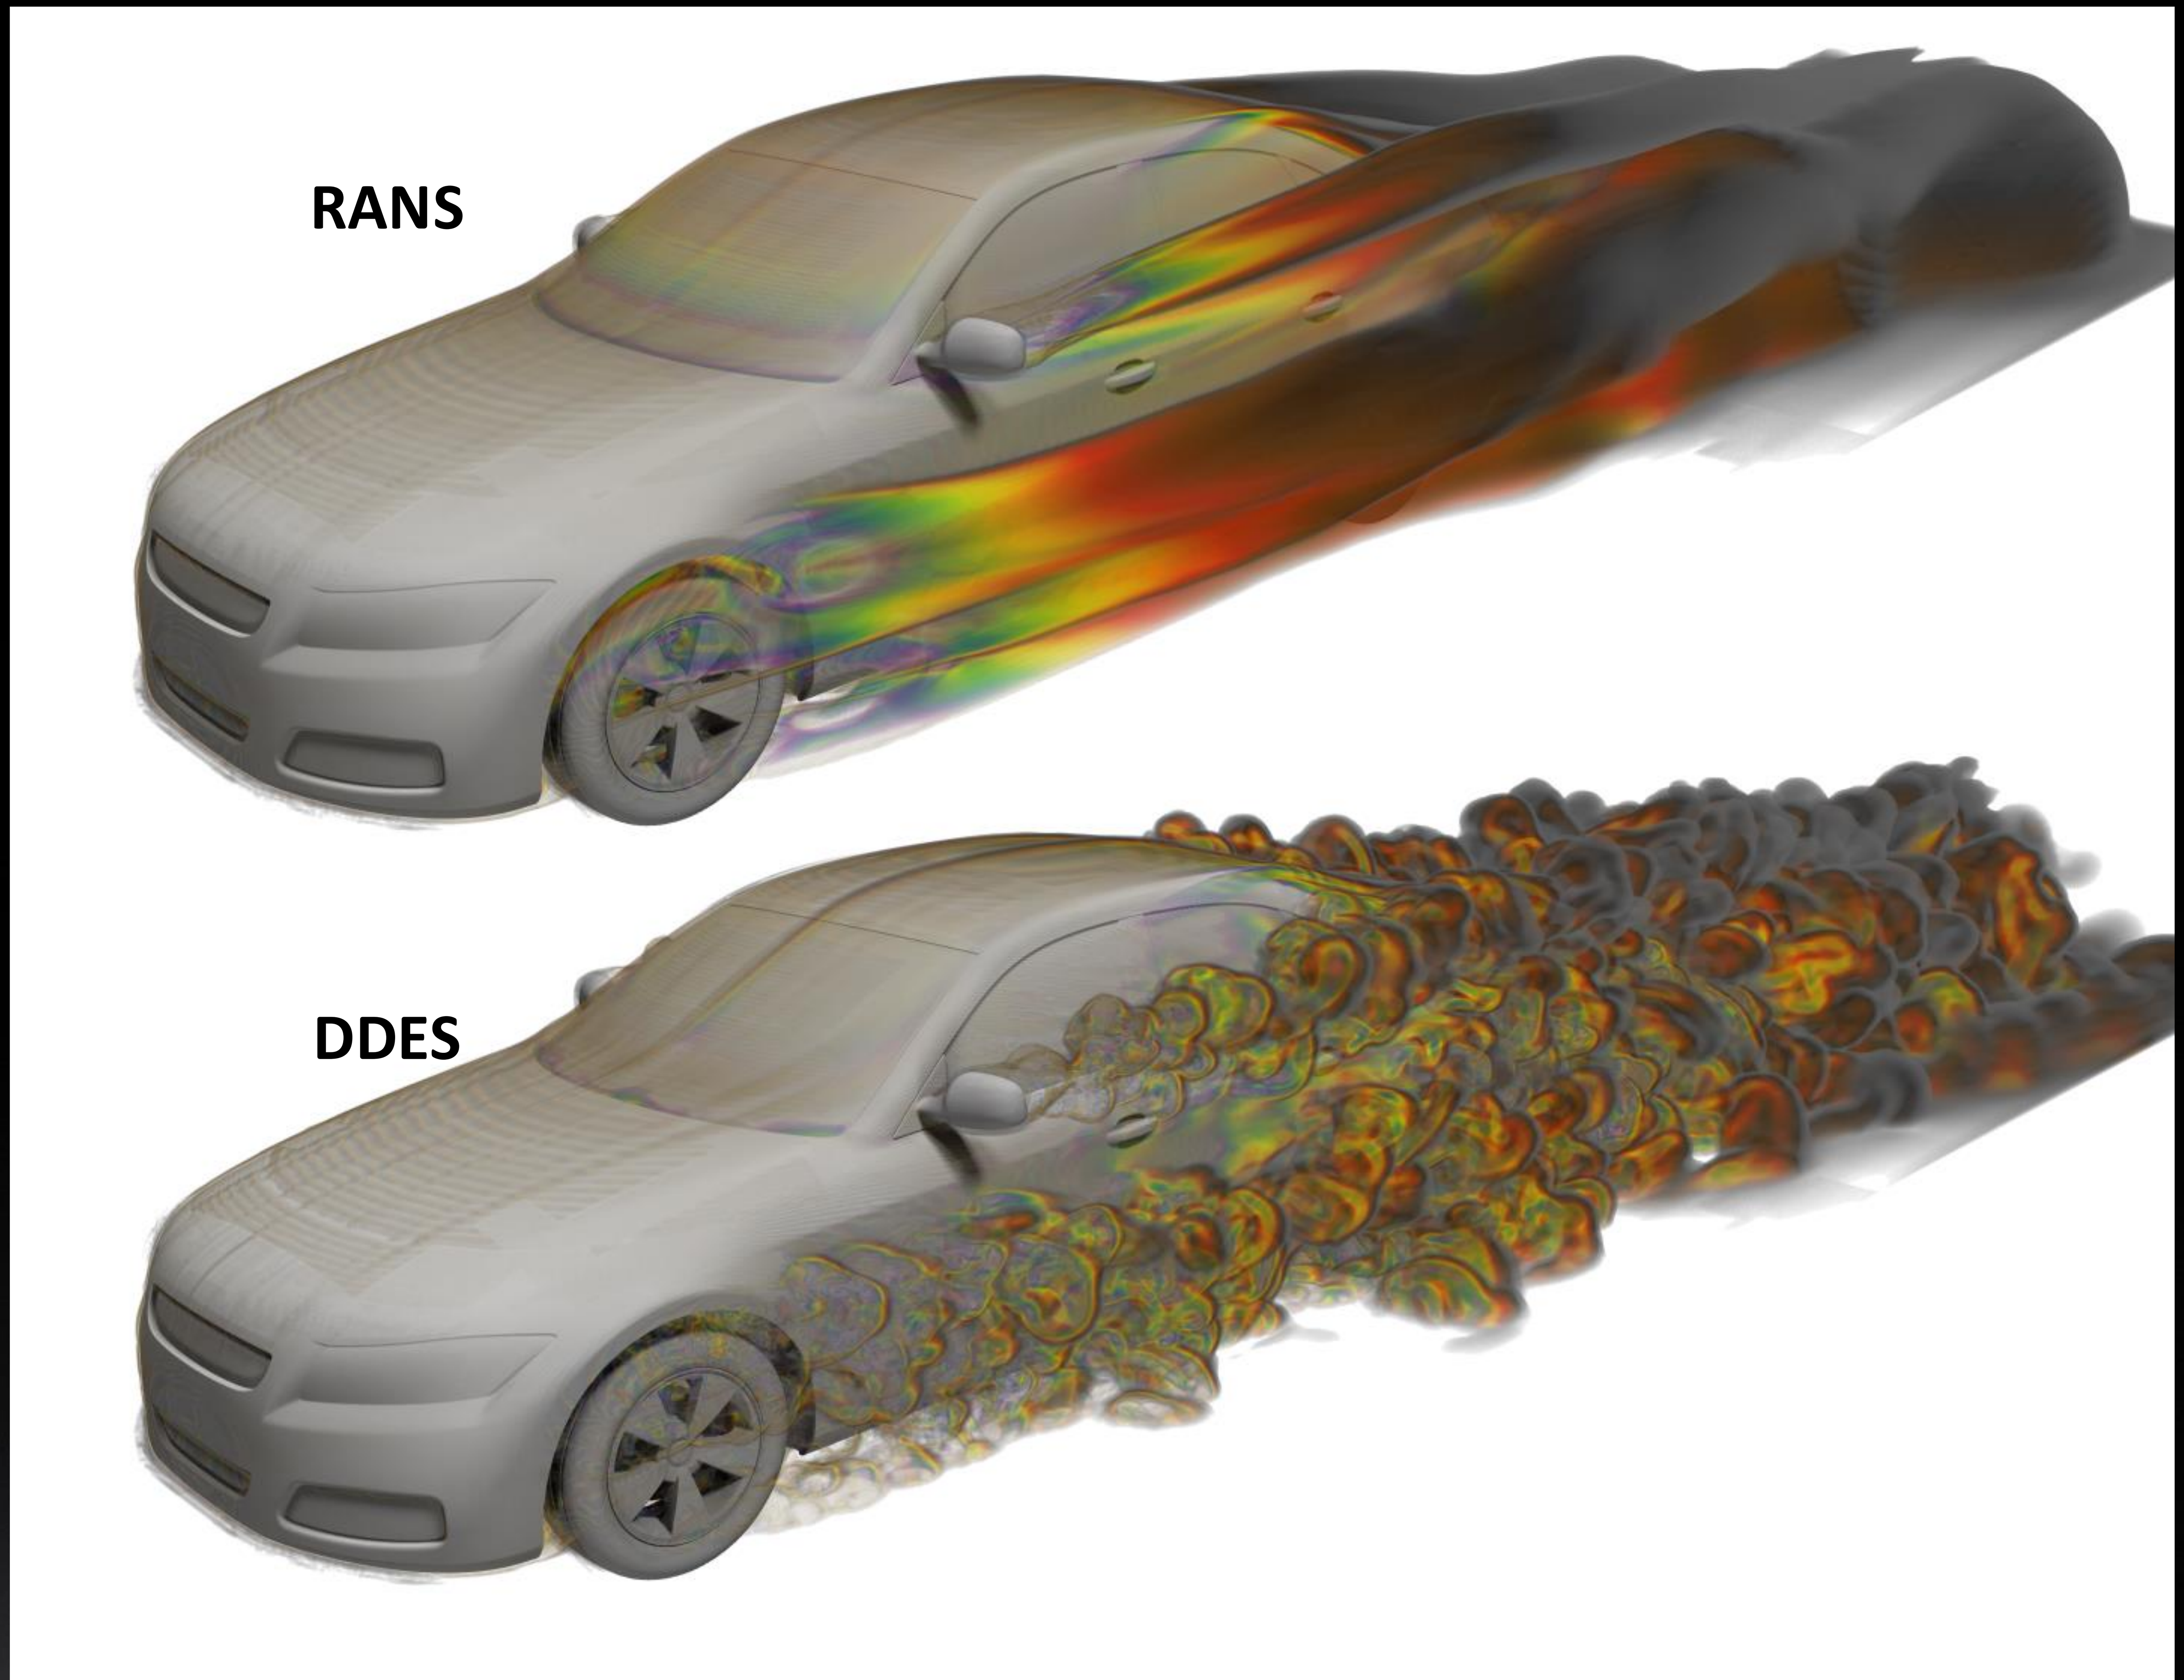

Q-criterion isosurfaces on DrivAer geometry with RANS (top) and DDES (bottom).
Figure reproduced from DrMesh using Star-CCM+, CC-BY-SA.

PhysicsNeMo Surrogate Models for Transient CFD Initialization

Problem Setup

- Various initialization strategies tested:
  - Traditional, fast: uniform flow, potential flow
  - Traditional, slow: steady RANS, DDES snapshot
  - ML-based: DoMINO via NVIDIA PhysicsNeMo
    - Various strategies used to extend prediction range from the near-field ML domain to the larger CFD domain

Traditional Potential Flow Initialization

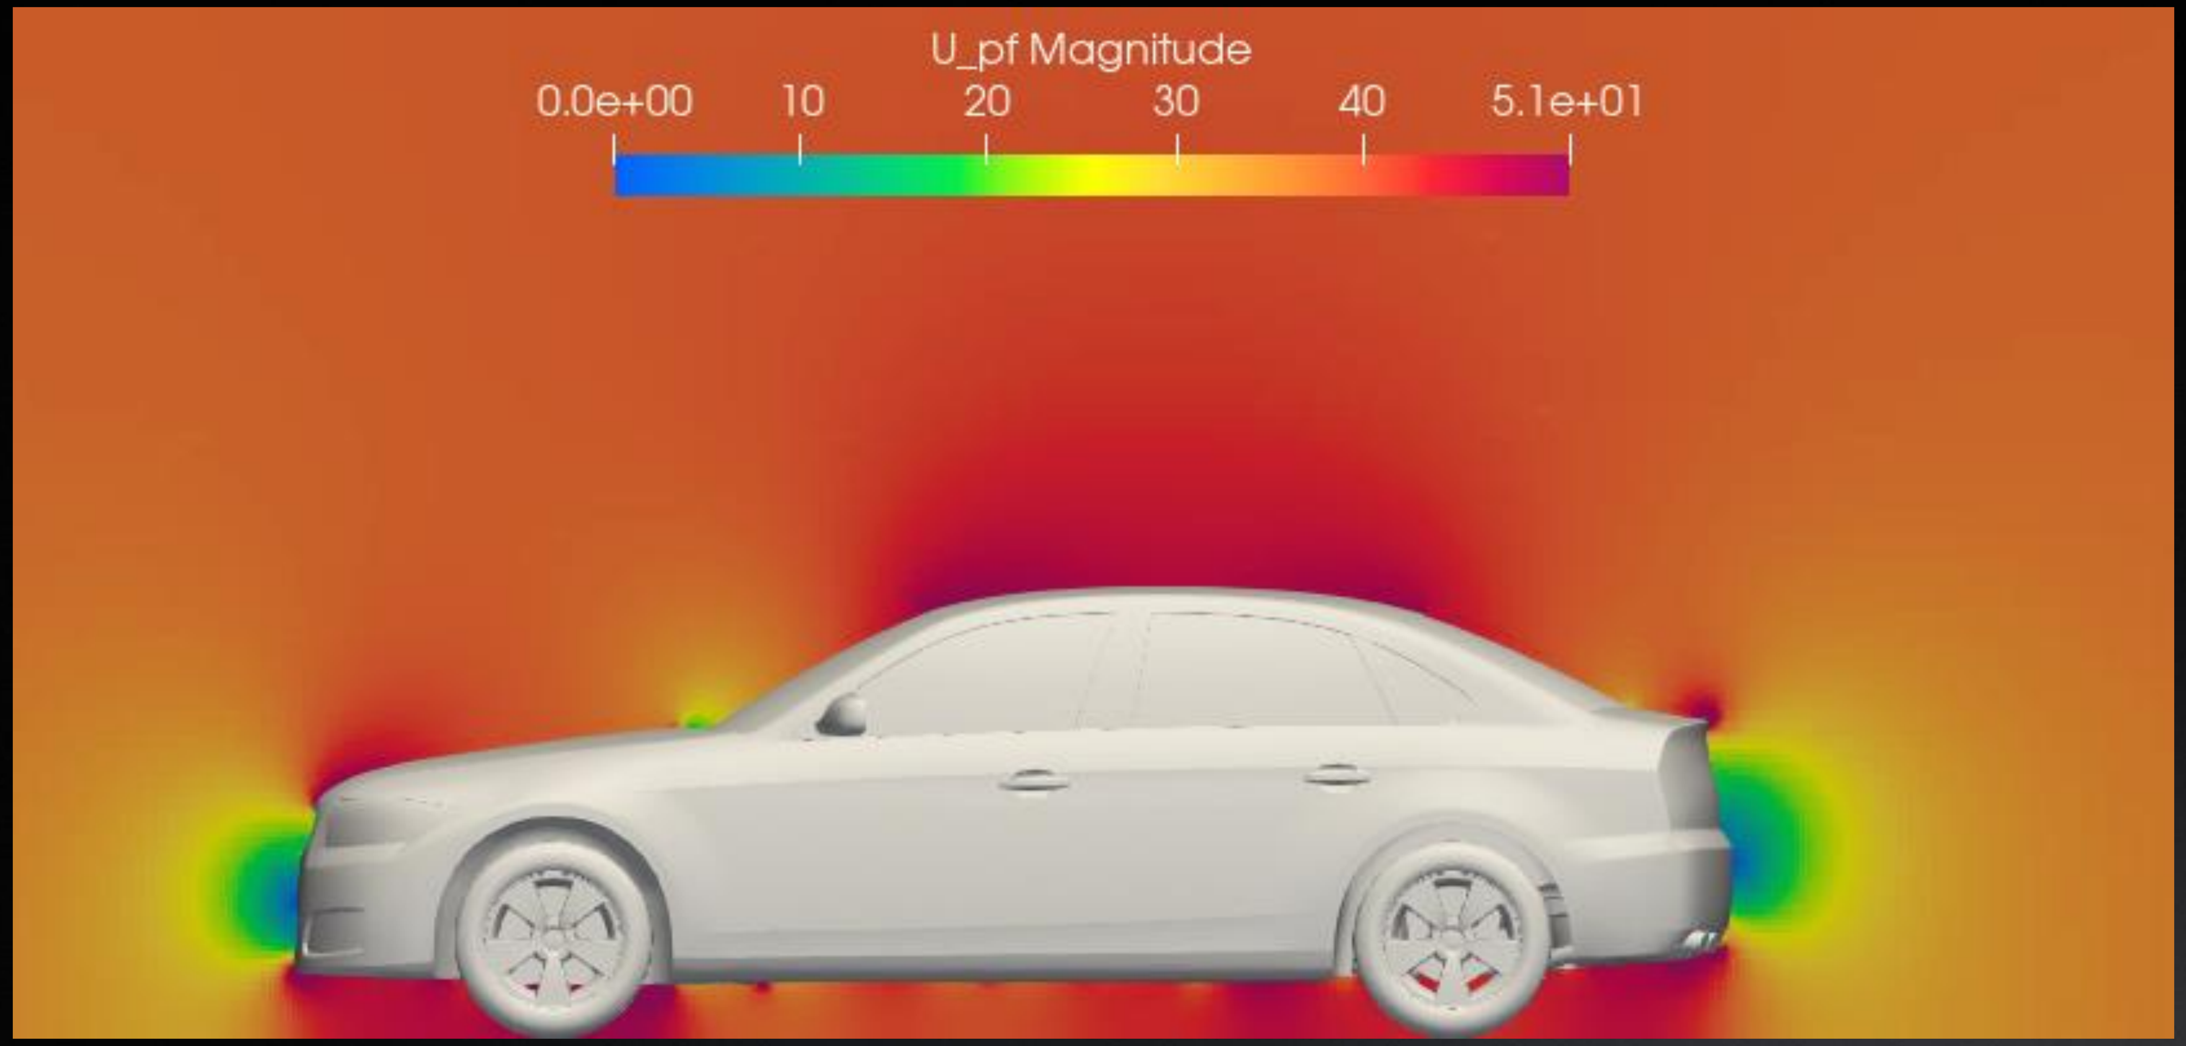

DoMINO-based Initialization

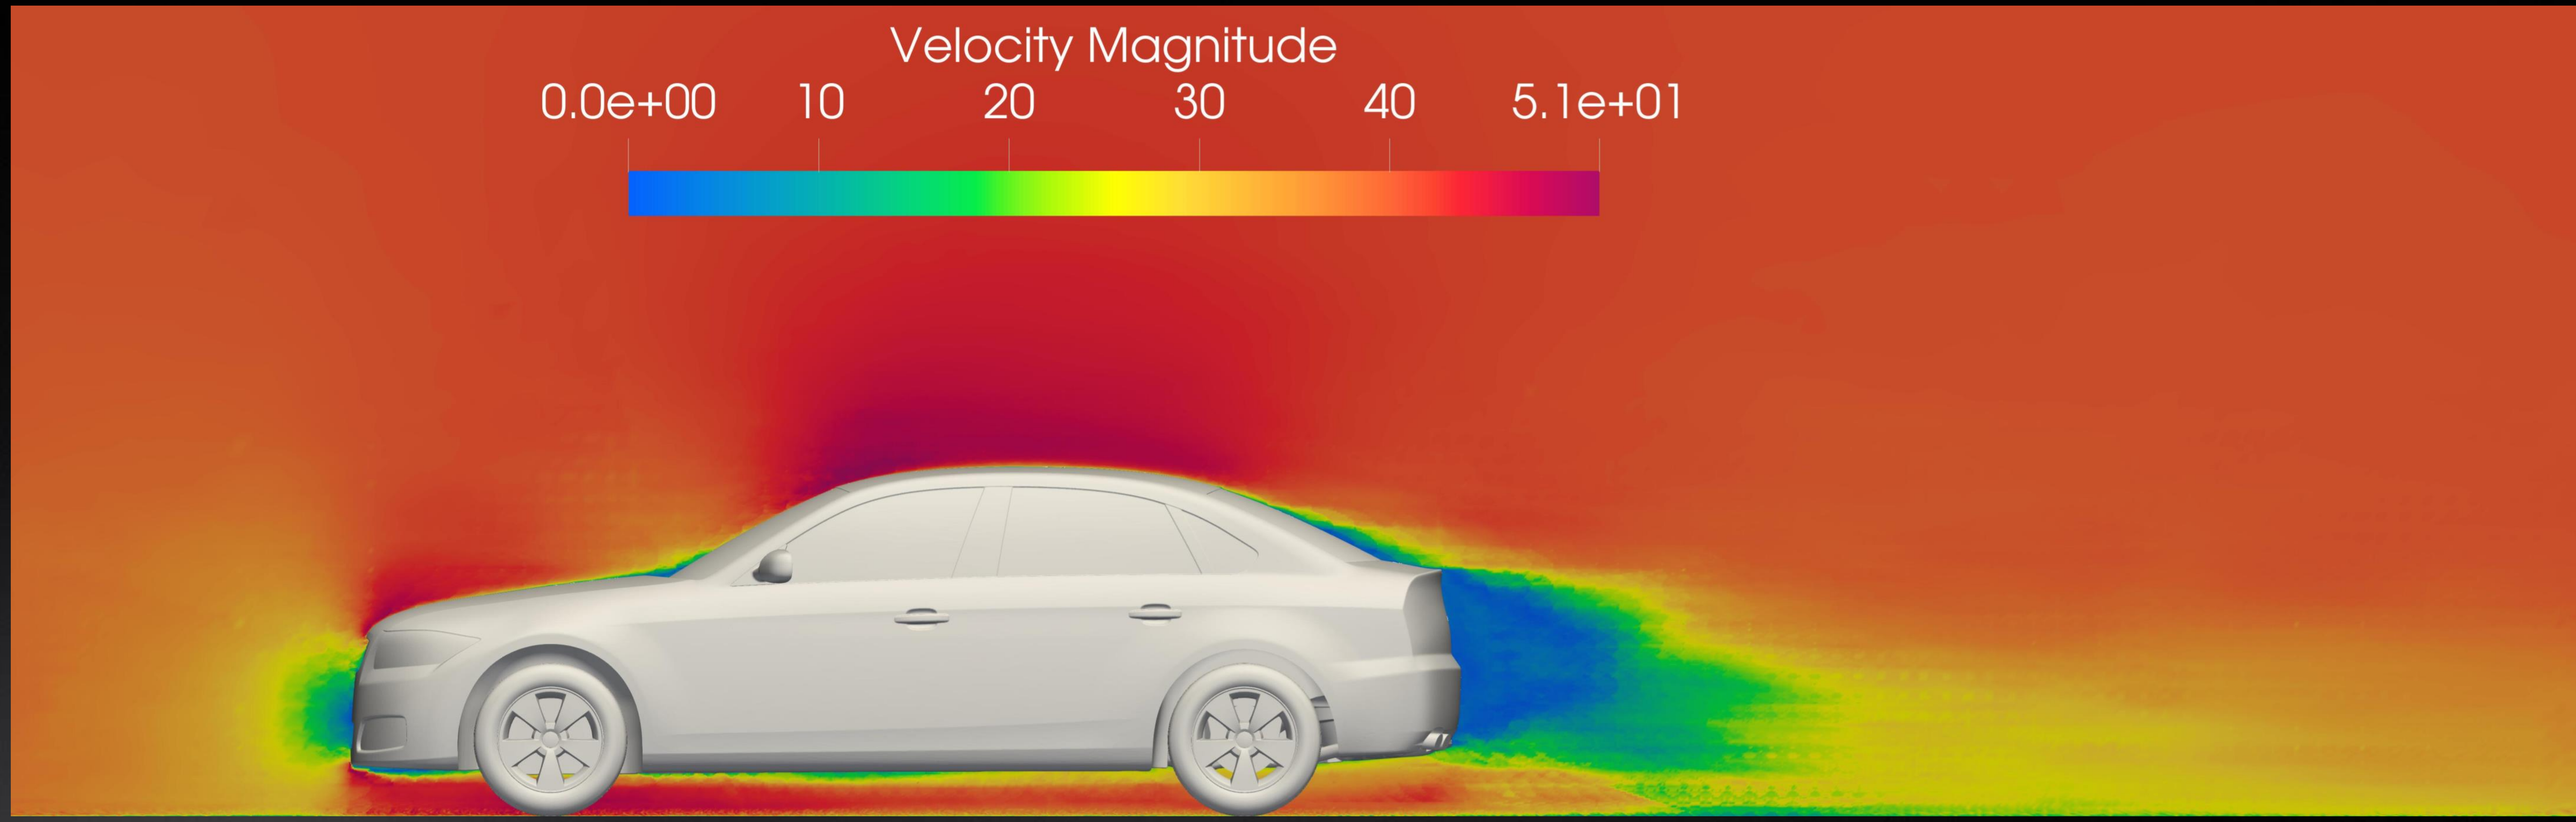

Paper: “Accelerating Transient CFD through ML-based Flow Initialization”, ArXiv 2025

PhysicsNeMo Surrogate Models for Transient CFD Initialization

Problem Setup

- Geometry & Mesh:
  - Sample from public DrivAerML automotive CFD dataset
  - 16.7M cells, hex-dominant, via SnappyHexMesh
- ML Model for Initialization:
  - DoMINO architecture via NVIDIA PhysicsNeMo
  - Training data:
    - DriveSim dataset, an in-house automotive CFD dataset that includes sedans, pickup trucks, hatchbacks, etc.
    - Notably: does not include any DrivAerML samples – the geometry for this case is **out-of-distribution**.
- Transient CFD Solver:
  - OpenFOAM
  - Incompressible, URANS with k-ω SST turbulence
  - 39 m/s freestream, mixed ground boundary conditions

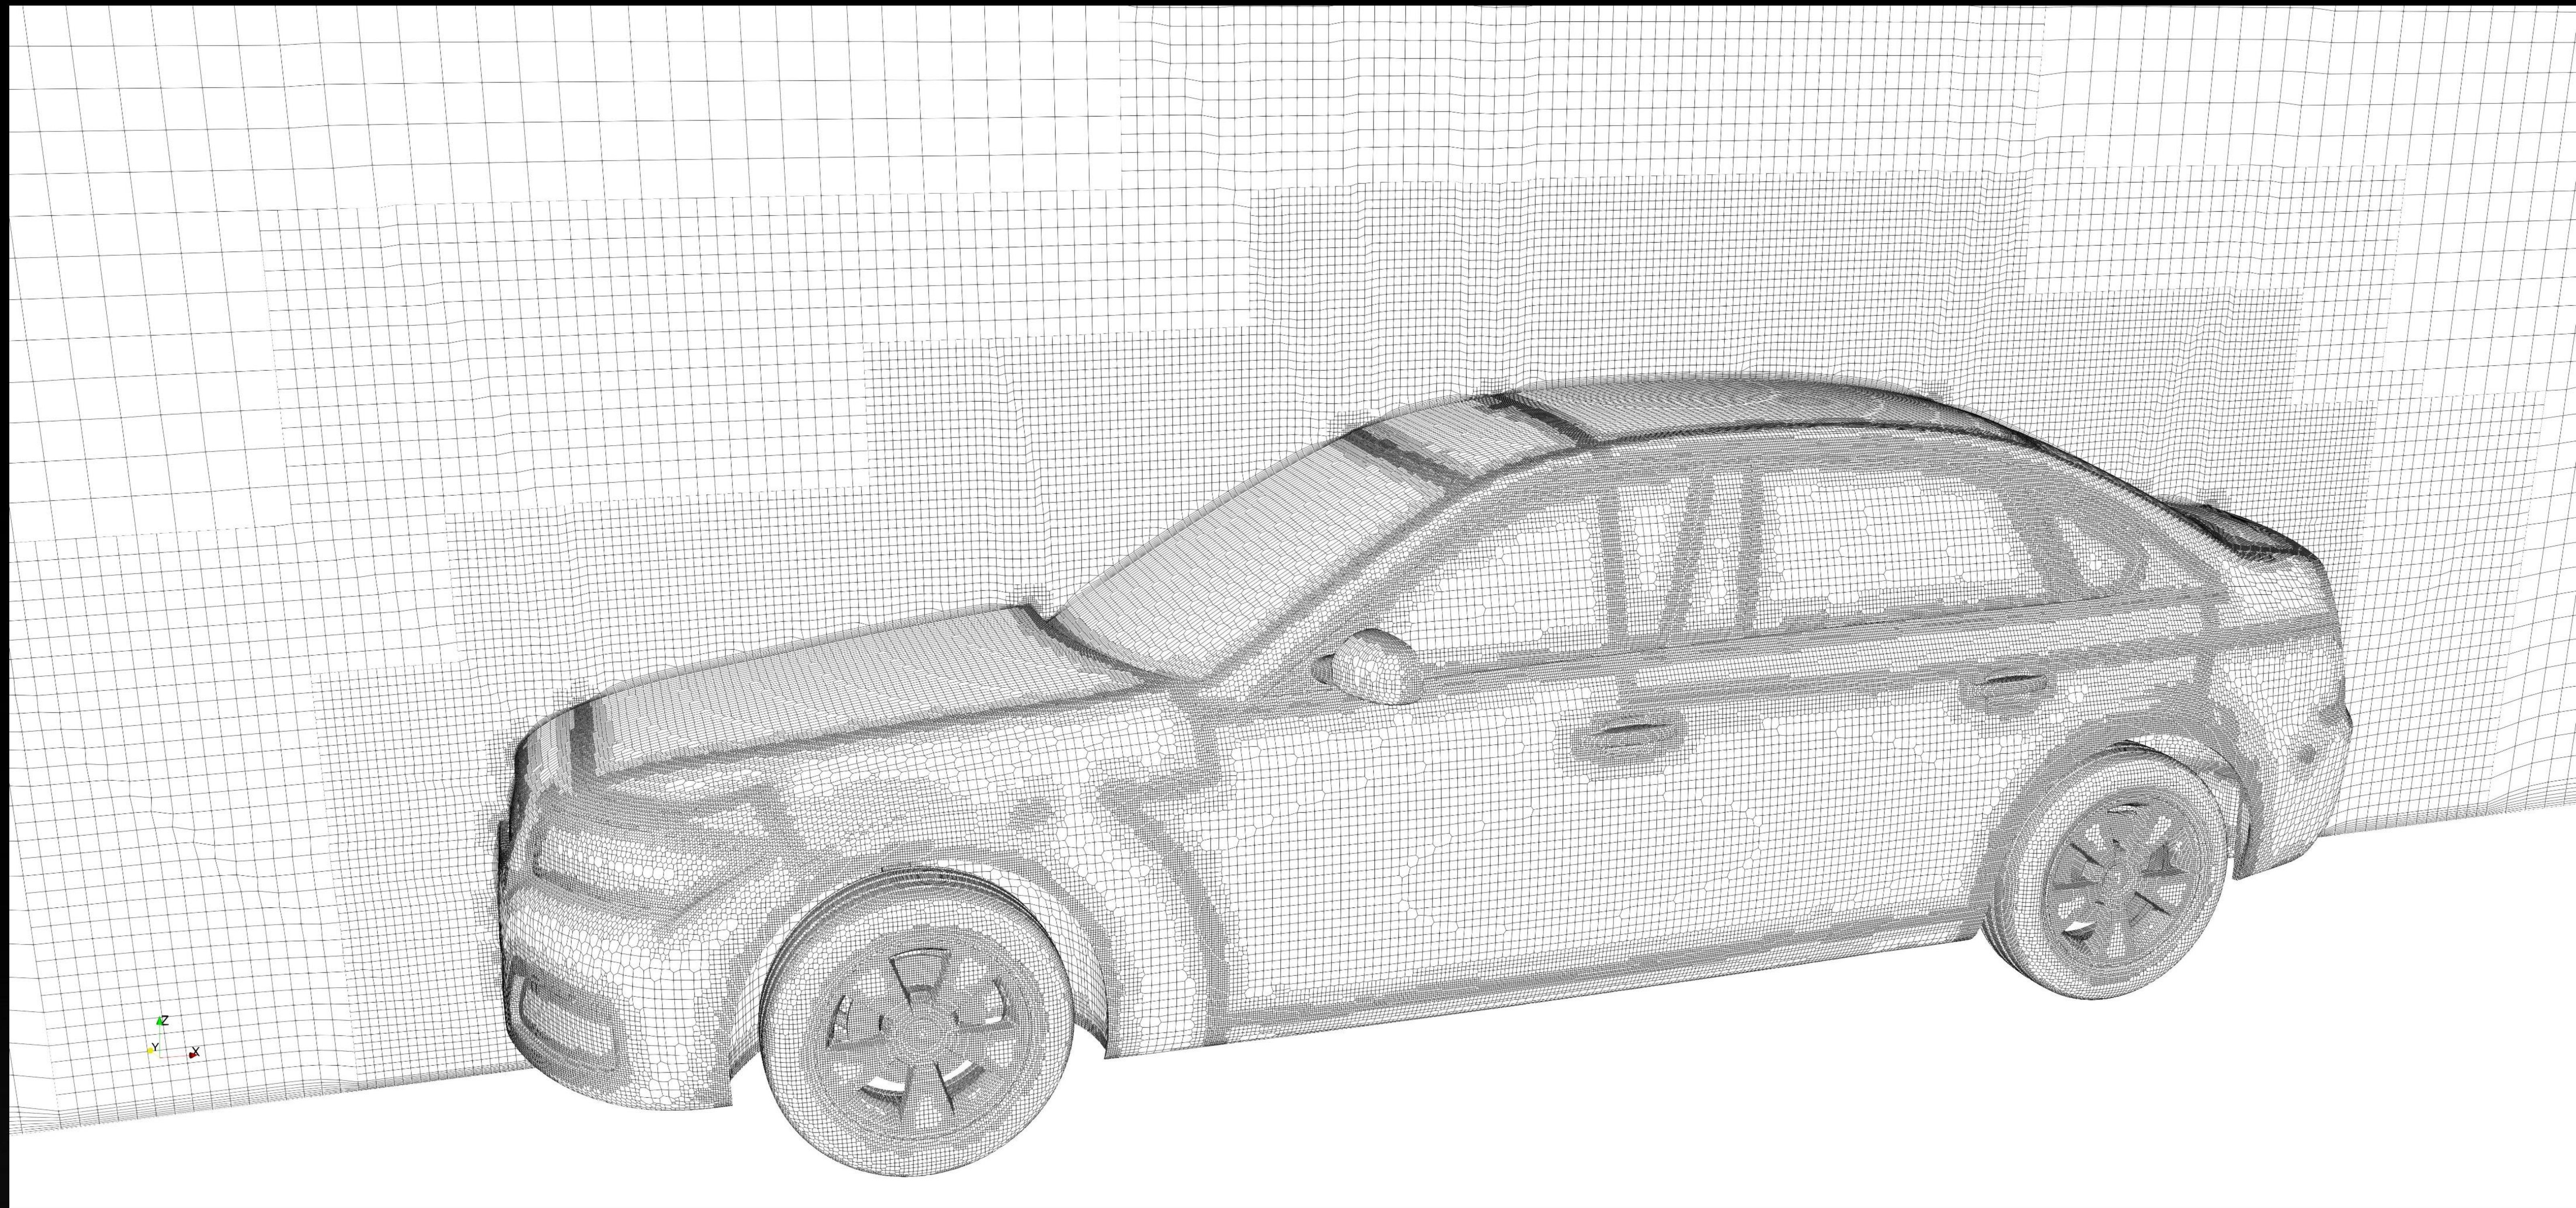

Paper: "Accelerating Transient CFD through ML-based Flow Initialization", ArXiv 2025

PhysicsNeMo Surrogate Models for Transient CFD Initialization

Key Results

- Developed two strategies (DoMINO + Uniform, DoMINO + Potential) that yield substantial speedups over traditional methods.
  - ~2x faster convergence of the subsequent transient solve, relative to traditional initialization methods with comparable cost
- With LES instead of URANS, wall-clock speedup becomes even more important

	Initialization Strategy	Initialization wall-clock runtime (hours)	Time Required for Transient Convergence: Physical simulation time (sec.)	Time Required for Transient Convergence: Wall-clock runtime (hours)
Traditional	**Uniform Flow**	Instant	0.7642	19.5
	Potential Flow	0.18	0.8668	22.1
	Steady RANS	2.4	0.1852	4.7
	DDES Flow	40†	0.5050	12.9
ML-based	**DoMINO + Uniform**	0.02	0.5540	14.1
	DoMINO + IDW	0.03	1.4198	36.2
	DoMINO + Potential (hybrid)	0.21	0.3146	8.0

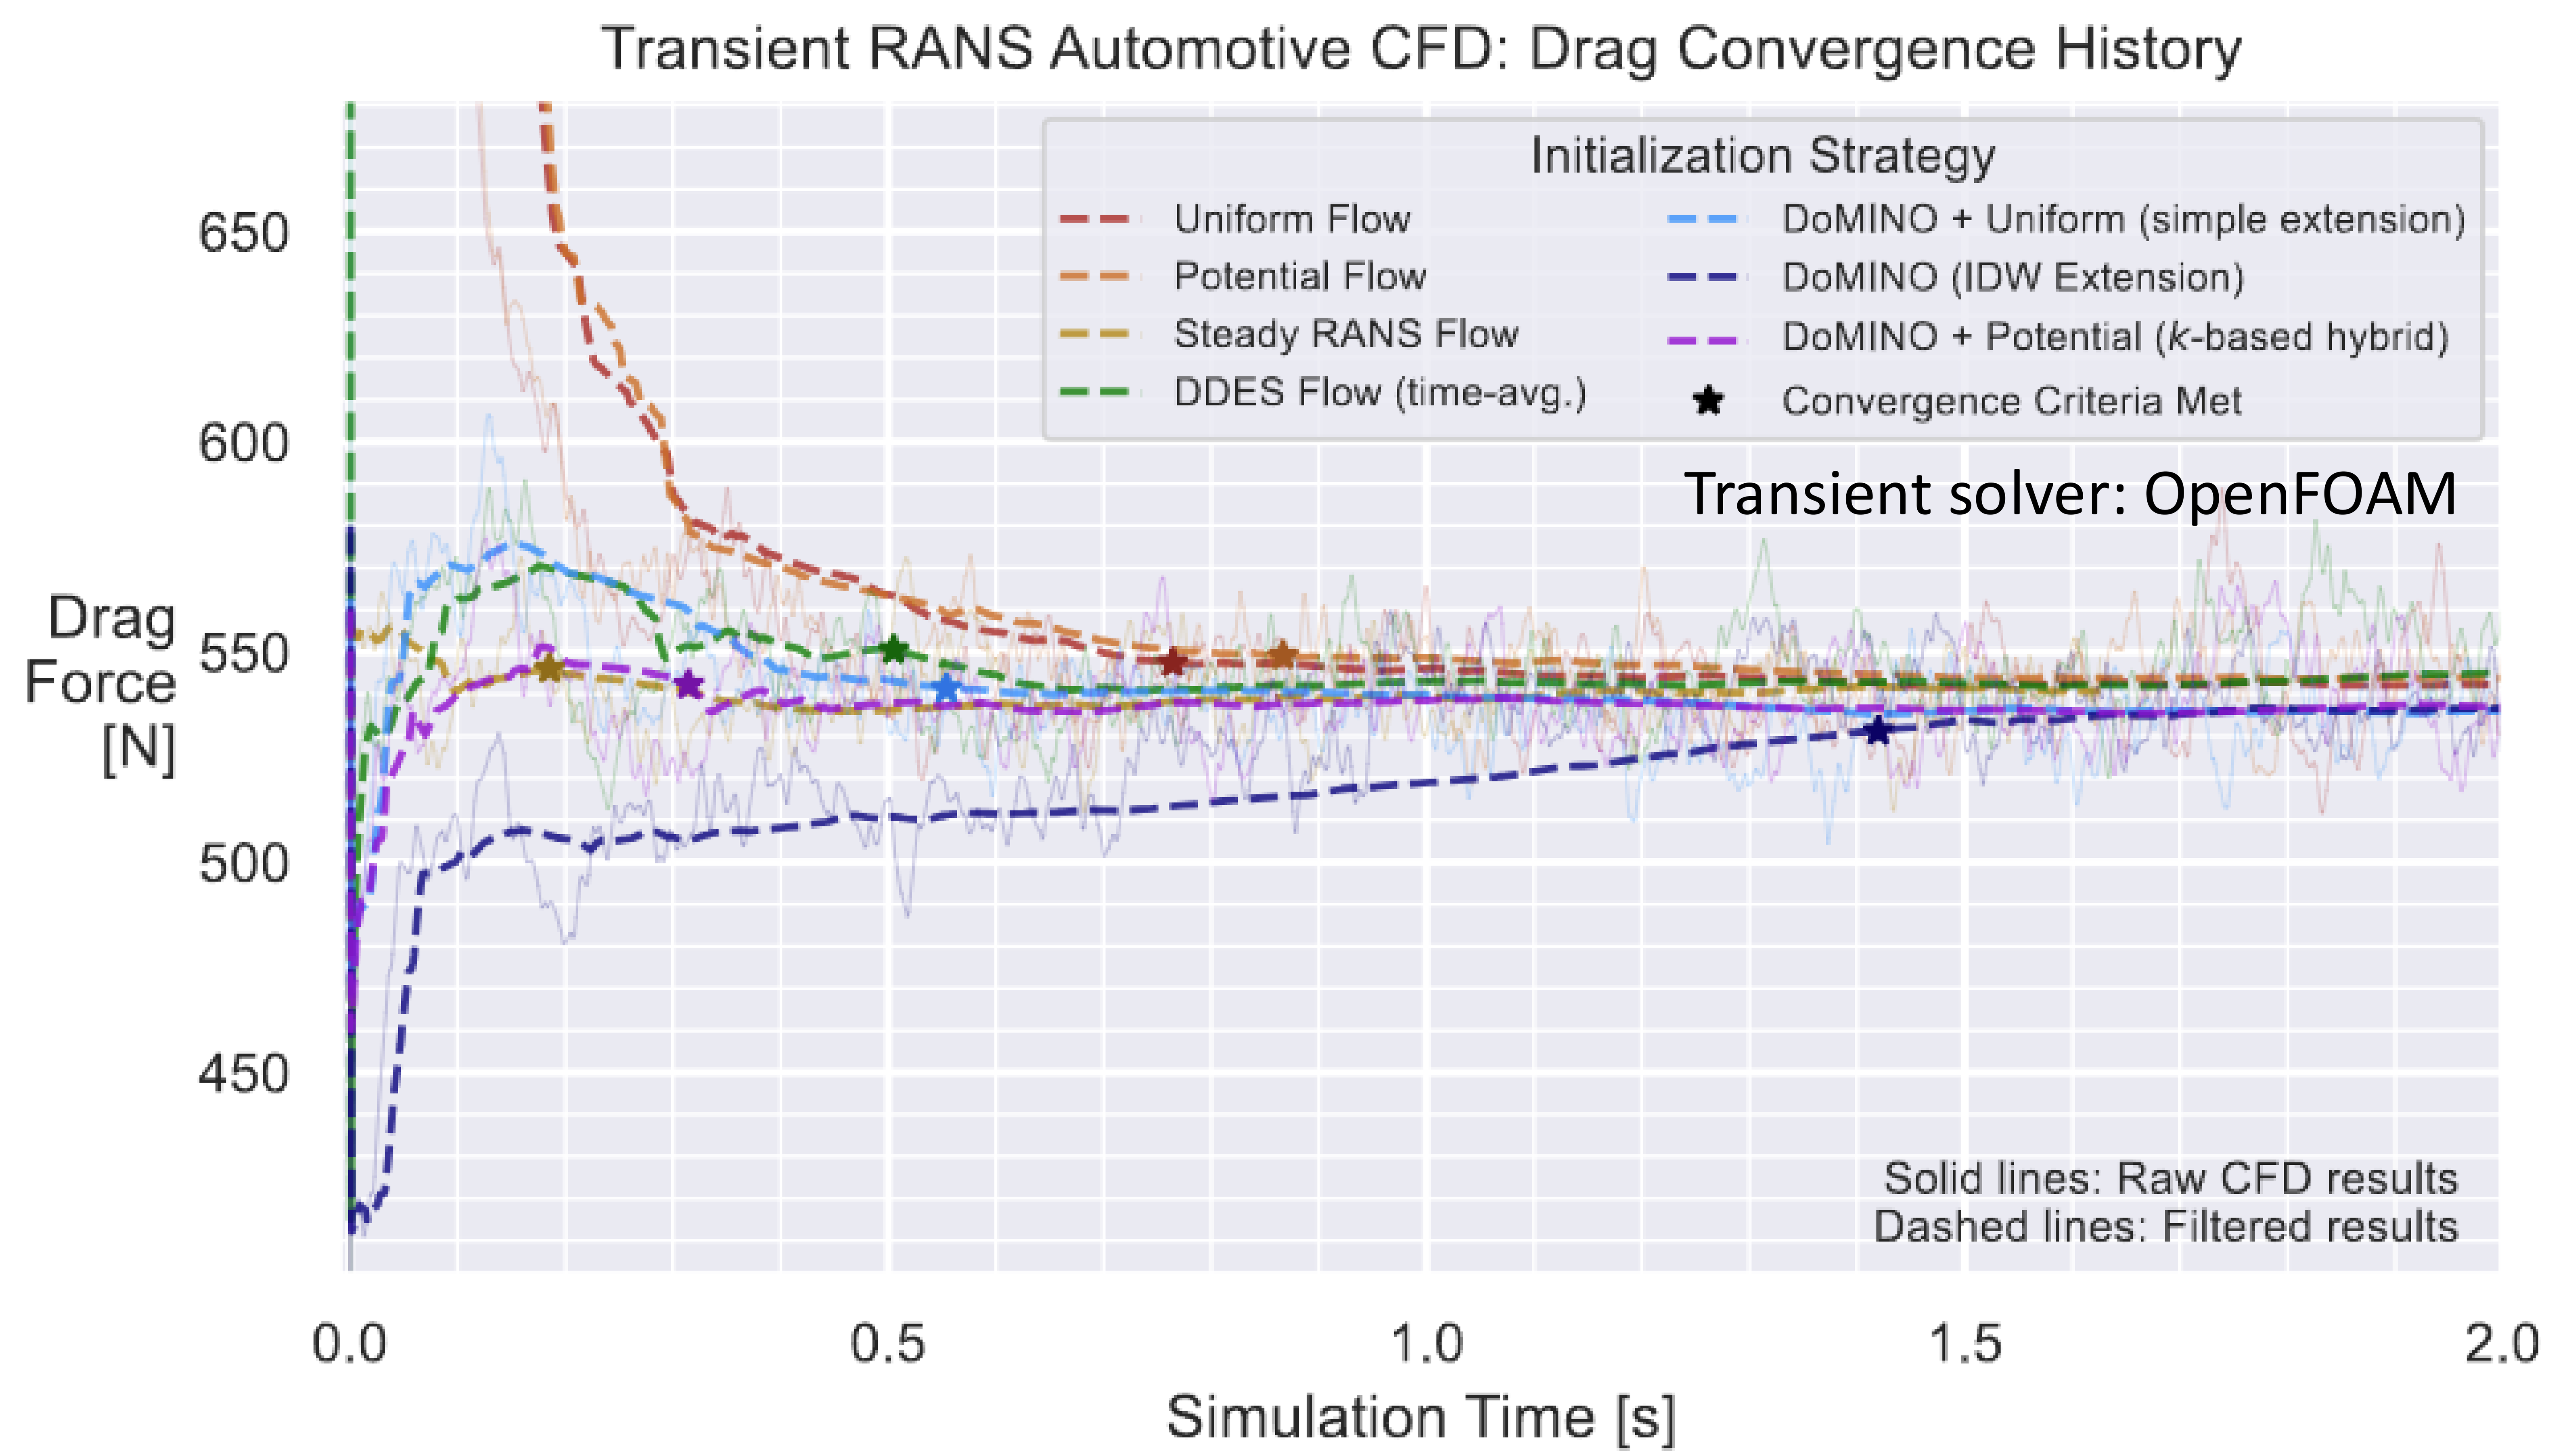

Paper: “Accelerating Transient CFD through ML-based Flow Initialization”, ArXiv, 2025

Wind Turbine Wake Optimization — Siemens Gamesa

Enabling high resolution simulation

Use Case

- Developing optimal engineering wake models to optimize wind farm layouts
- Simulating the effect that a turbine might have on another when placed in close proximity

Challenges

- Generating high-fidelity simulation data from Reynolds-averaged Navier-Stokes (RANS) or Large Eddy Simulations (LES) can take over a month to run, even on a 100-CPU cluster.

Solution

- NVIDIA Omniverse and PhysicsNeMo enable accurate, high-fidelity simulations of the wake of the turbines, using low-resolution simulations as inputs and applying super resolution using AI.

Outcome

- **~4,000x** speedup for high-fidelity simulation
- Optimizing wind farm layouts in real-time increases overall production while reducing loads and operating costs.

Demo

Earth 2

A digital twin of our planet

Weather Prediction is Integral to Modern Society

Simulations and forecasts drive planning and decision-making

Agriculture & Forestry

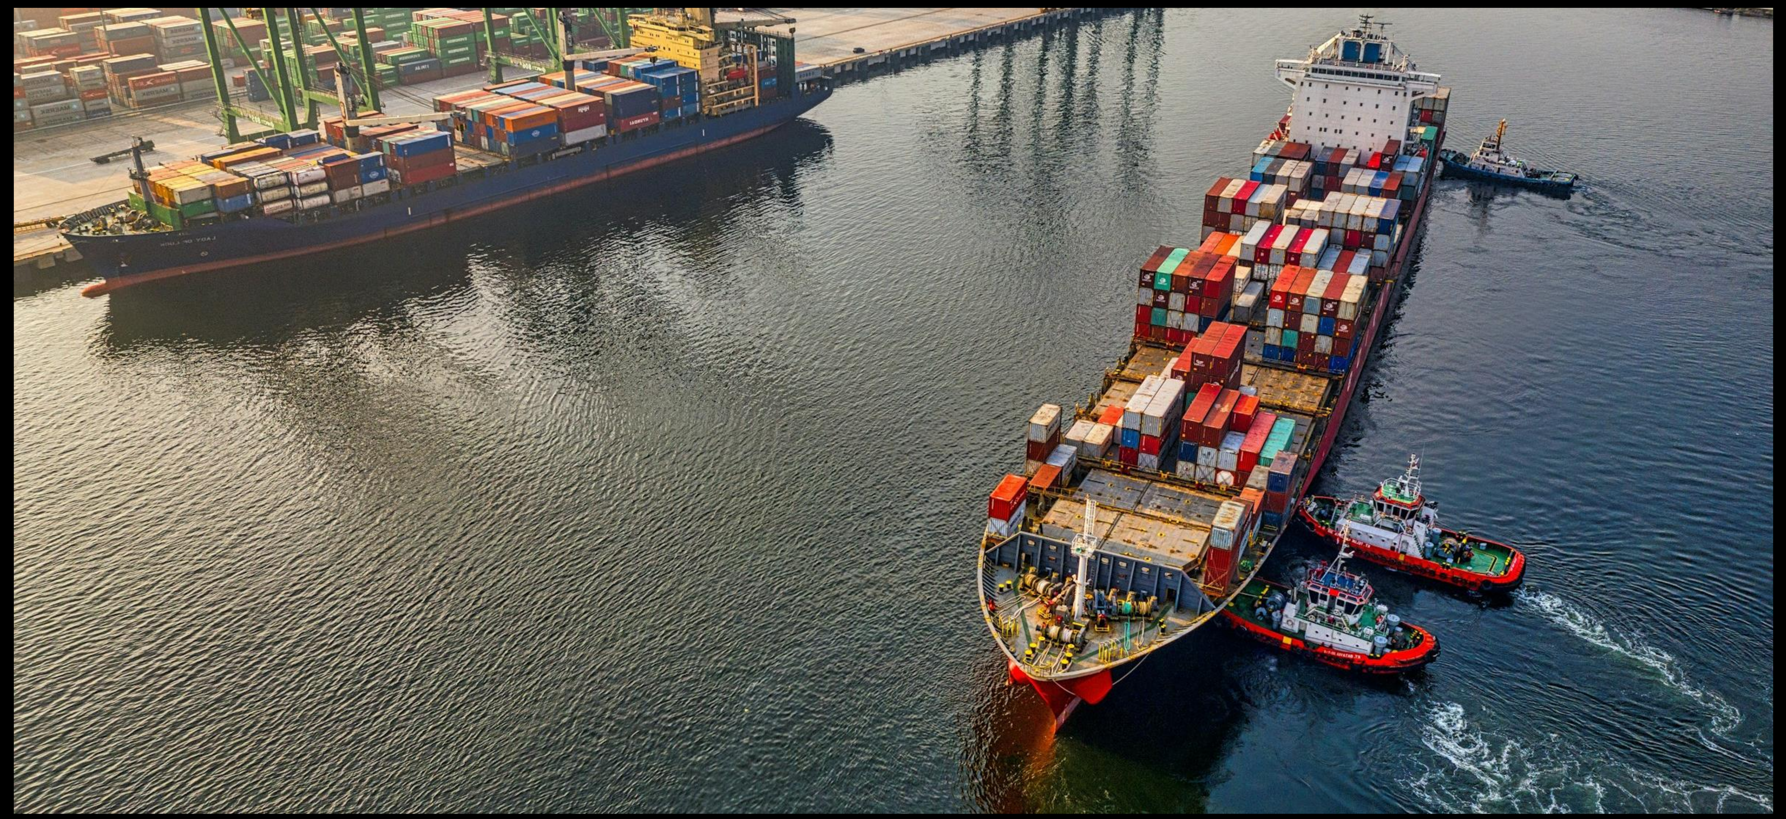

Transport & Logistics

Energy

Risk management

Land development

Recreation

Extreme Weather Events

Extreme weather events have become more frequent and more severe

Ahr Valley Flood:

https://commons.wikimedia.org/wiki/File:Hochwasser_in_Altenahr_Altenburg.jpg

Hurricane Ian

https://commons.wikimedia.org/w/index.php?curid=123606695

Pakistan Floods

https://commons.wikimedia.org/wiki/File:Flood_in_Pakistan_2022.png

Our future climate will be very different from the past

Global Mean Surface Temperatures (0°C – 6°C)

Will the Horn of Africa struggle with unending drought?

Will the hurricane season in North America intensify?

Will billions in South Asia suffer a failed summer monsoon?

Will southeastern Australia burn worse than in previous years?

Will Europe be submerged under incessant rain and heavy flooding?

Advantages of AI Weather Models

Compelling skill, resource requirements, and accessibility

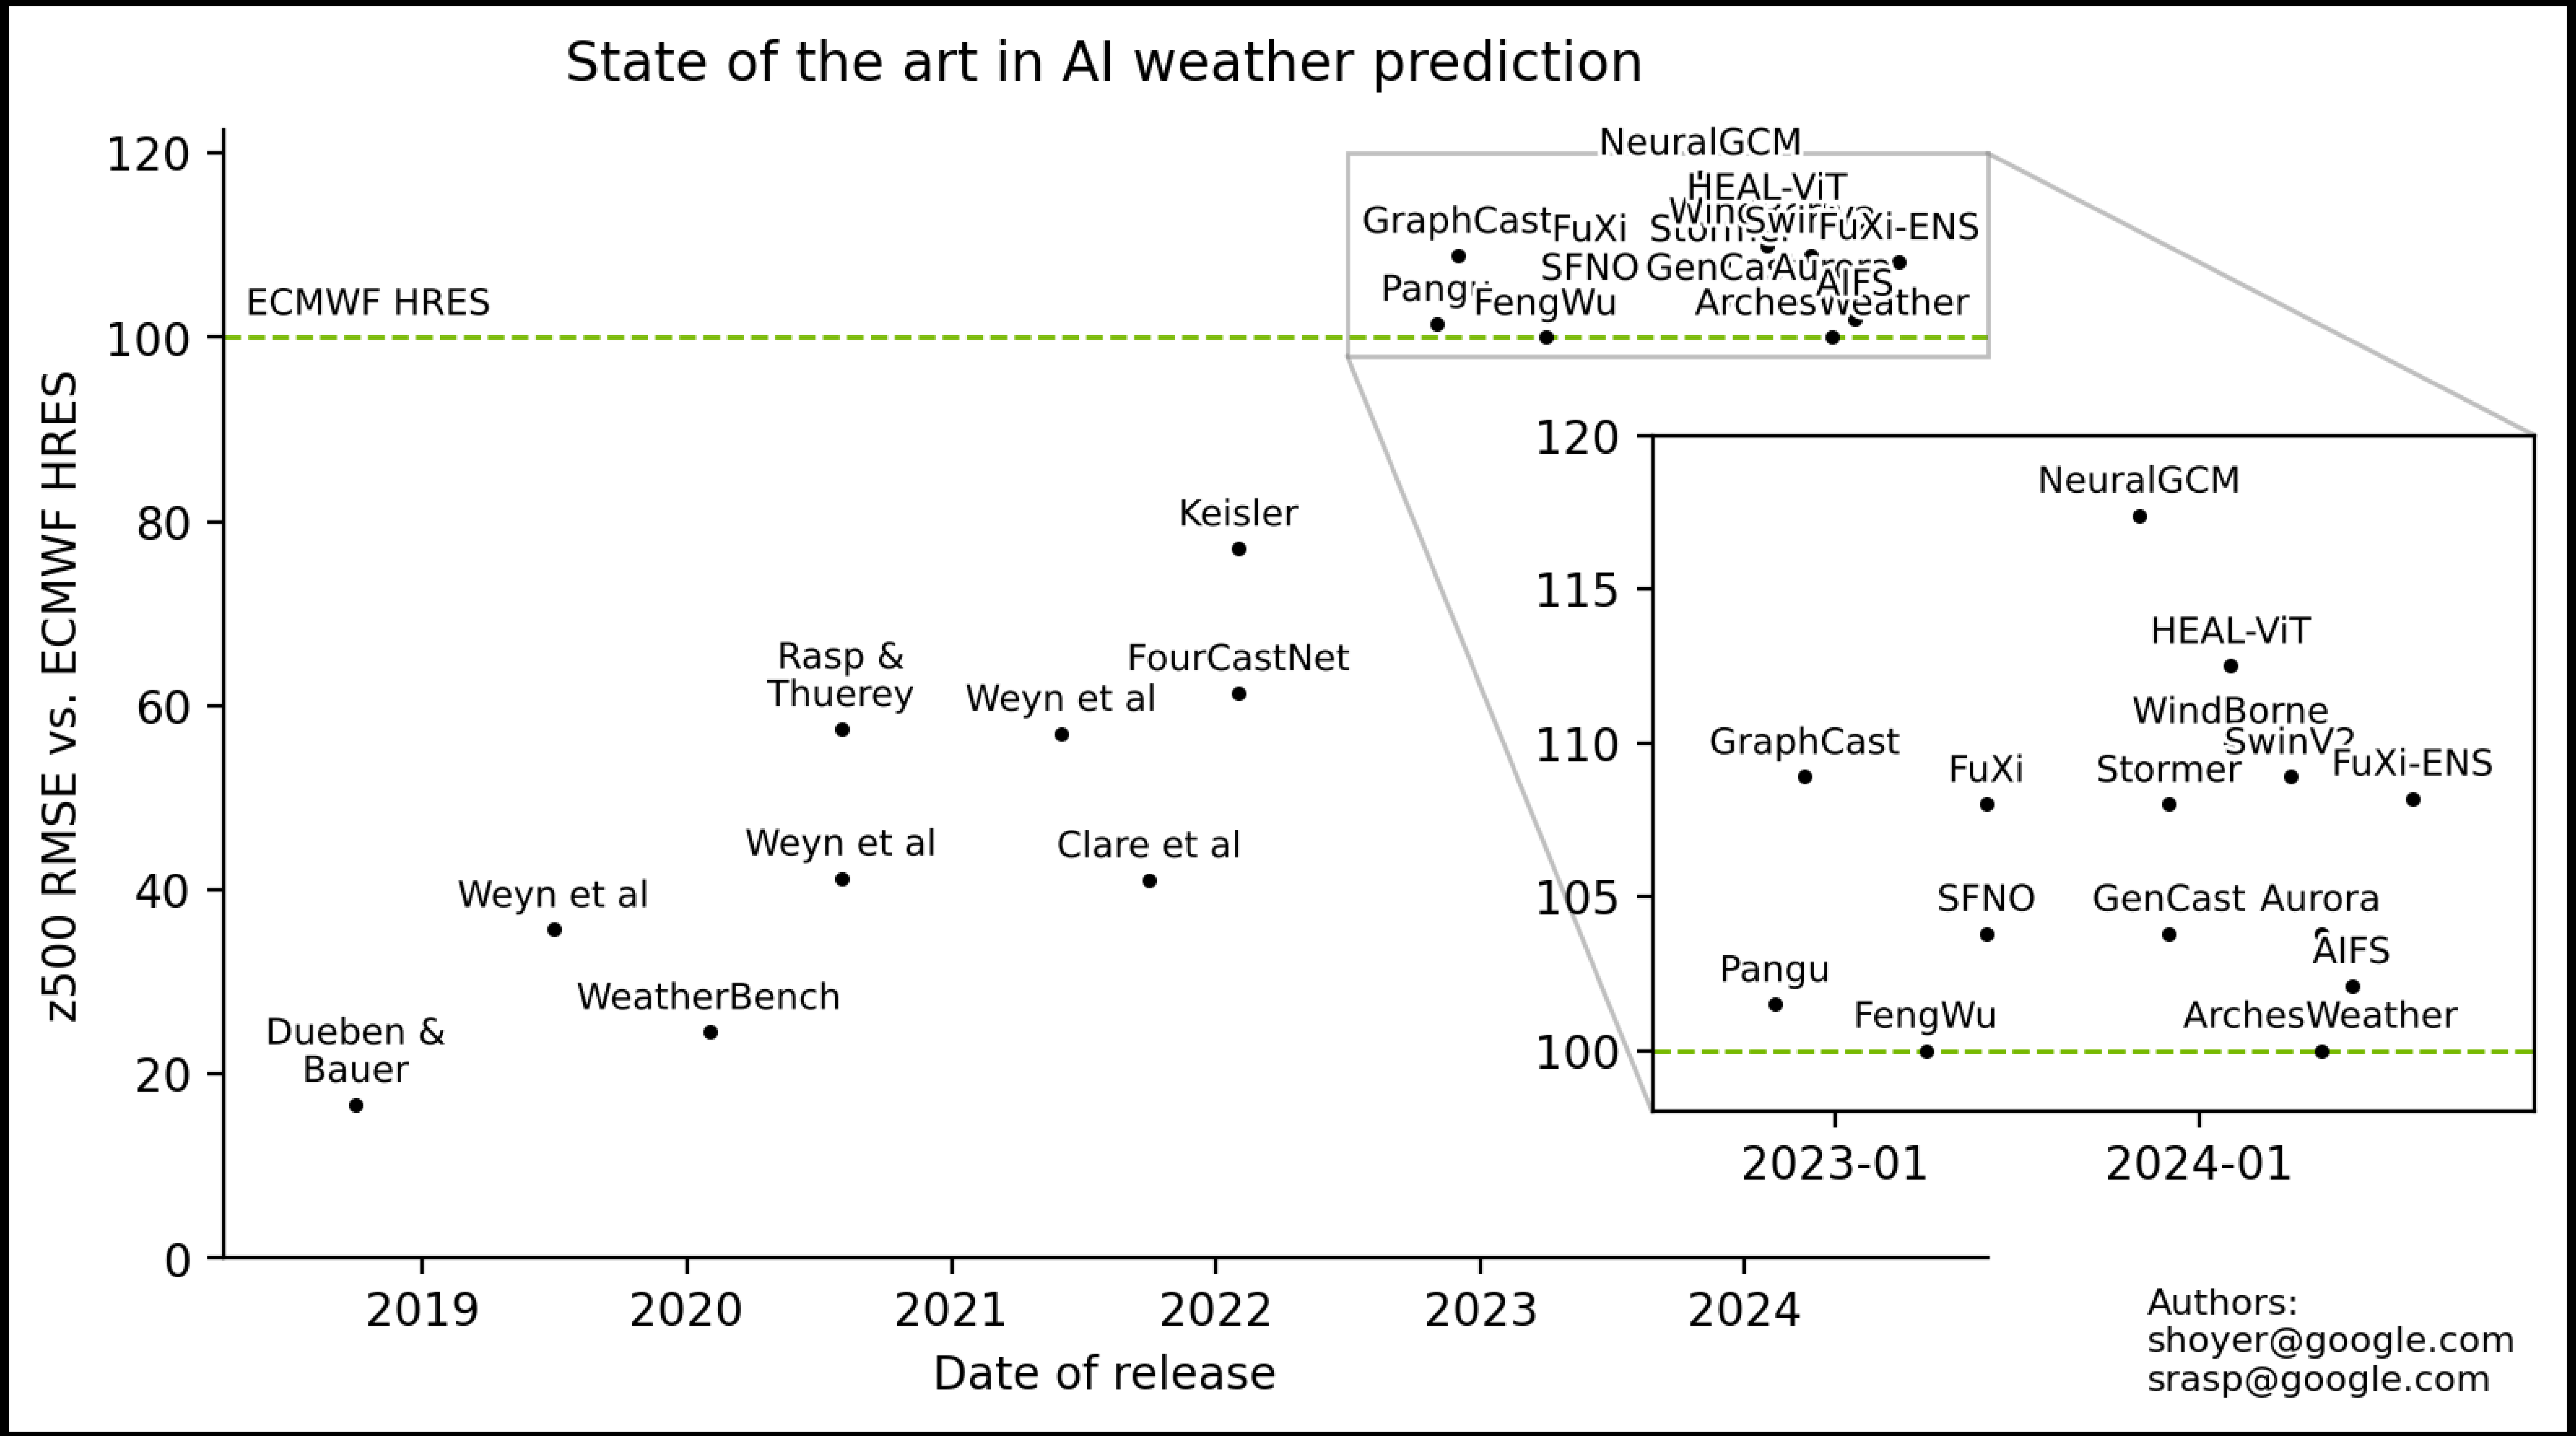

Rasp, Stephan (2024). AI-Weather SotA vs. Time.
https://doi.org/10.6084/m9.figshare.28083515.v1

High skill

Less hardware

Faster inference

More accessible

AI Weather Models in comparison

	Numerical WP	AIWP
hardware requirement	cluster	$\geqslant$ 1 GPU
time to produce 10-day forecast	hours	seconds
achievable ensemble sizes	O(100)	$\gg$1,000
customizability	low	high
user level	modelling expert	basic technical and domain knowledge

AI Could Side-Step Moore's Law With Implications for Weather Forecasts

Current PDE solvers will take 40 years to achieve meter-scale resolutions needed for local planning

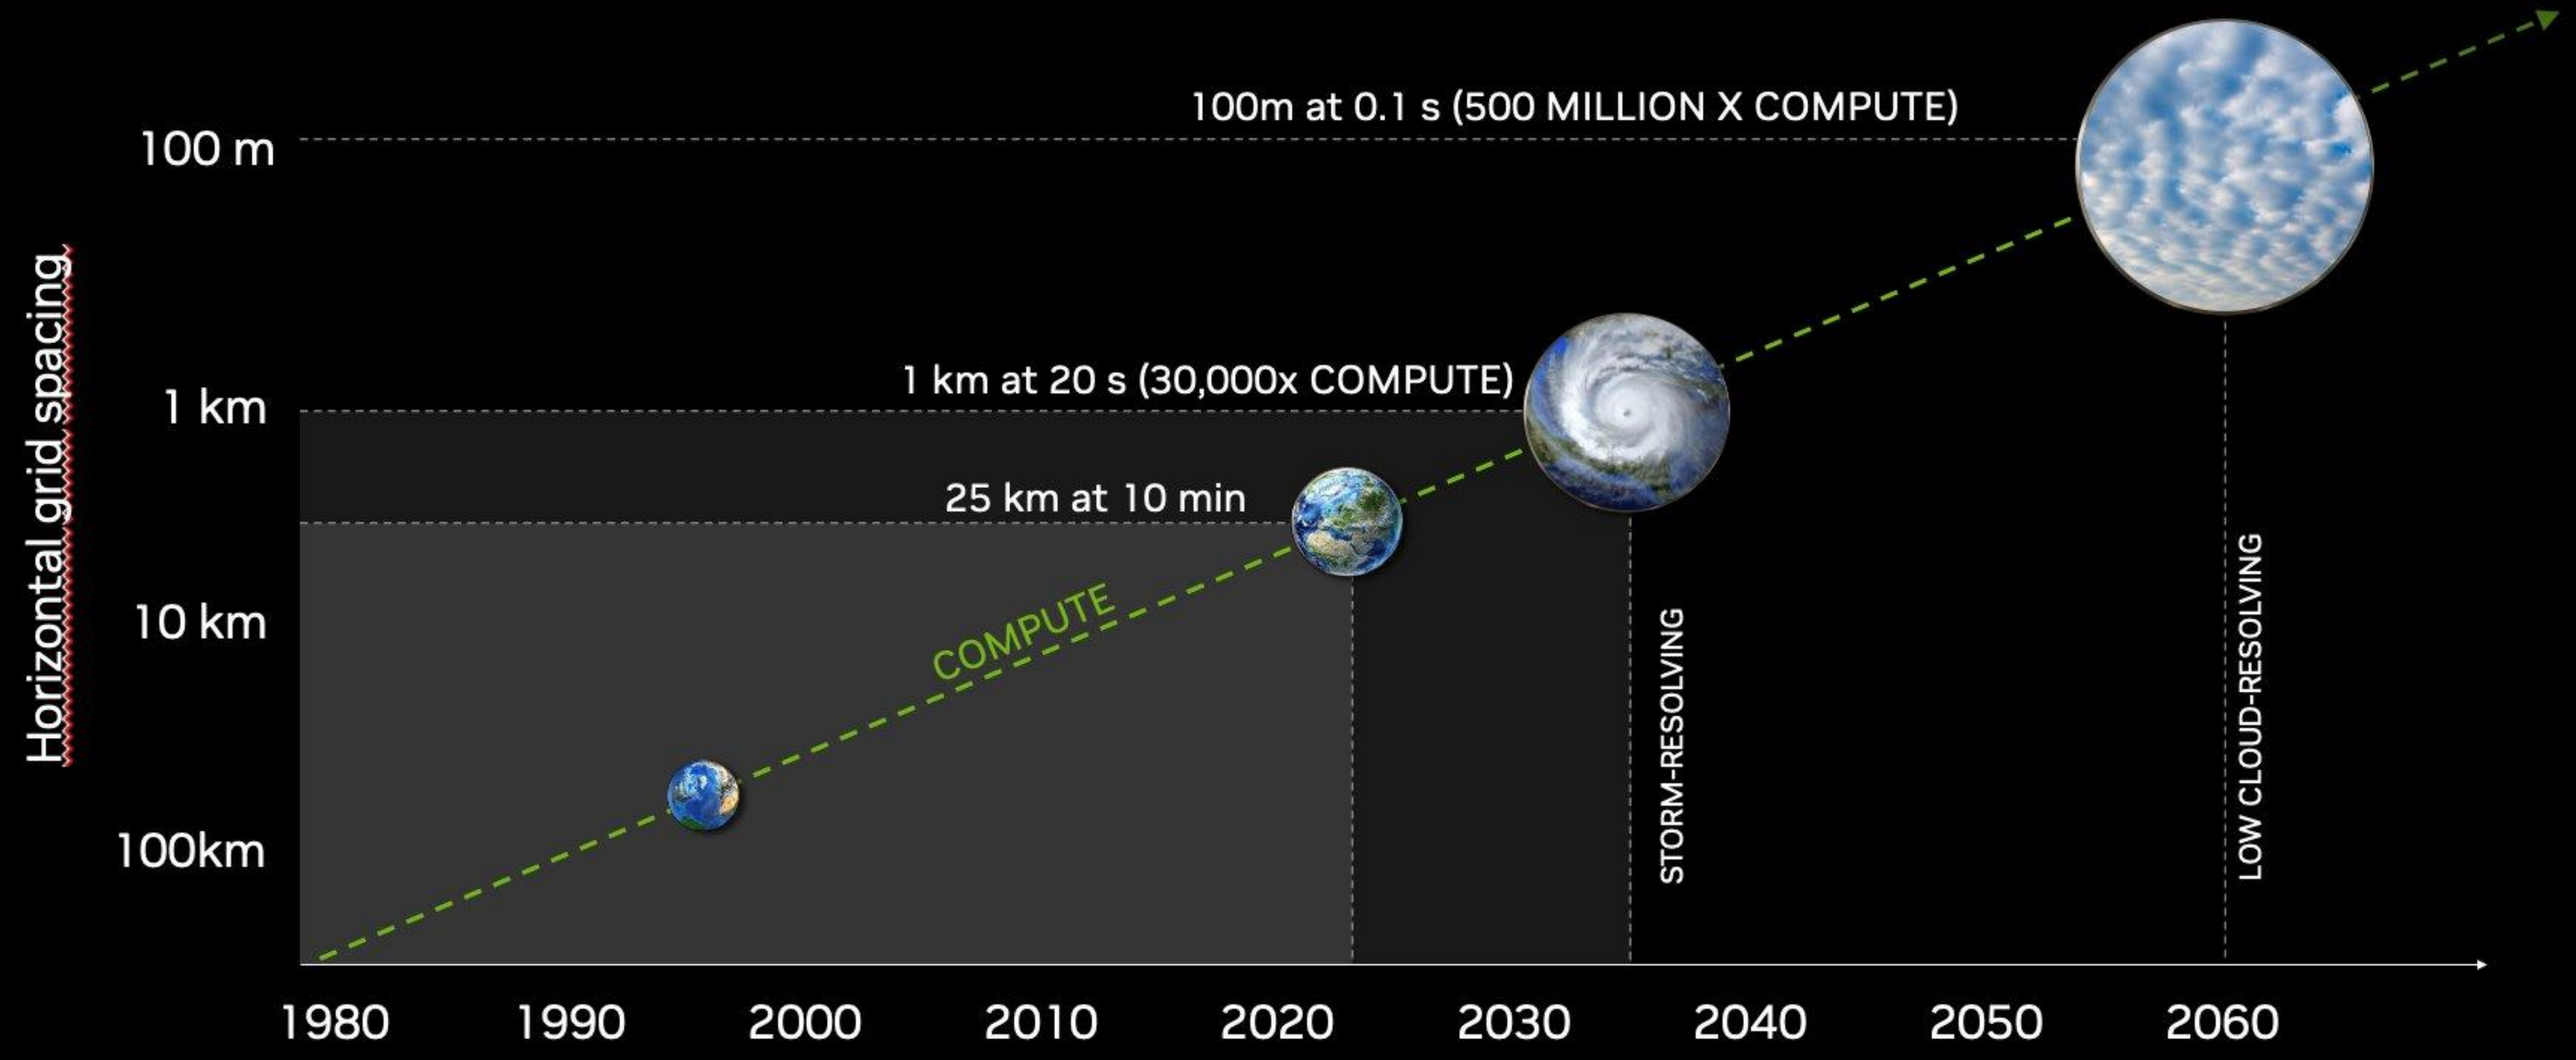

Can Breakthroughs in AI for Atmospheric Simulation Unlock Bigger Ensembles & Higher Resolution?

Which Requirements must Simulations fulfil to Predict Severe Weather Events?

High Resolution

Details like land-sea breeze, topography, or small-scale physics have a huge impact on the atmosphere, e.g. the track of a hurricane. High-resolution simulations are required to capture such small features.

Massive Number of Forecasts

Extreme events like floods are rare. Predicting rare extremes with high confidence requires a huge set (ensemble) of forecasts (~10,000 forecasts).

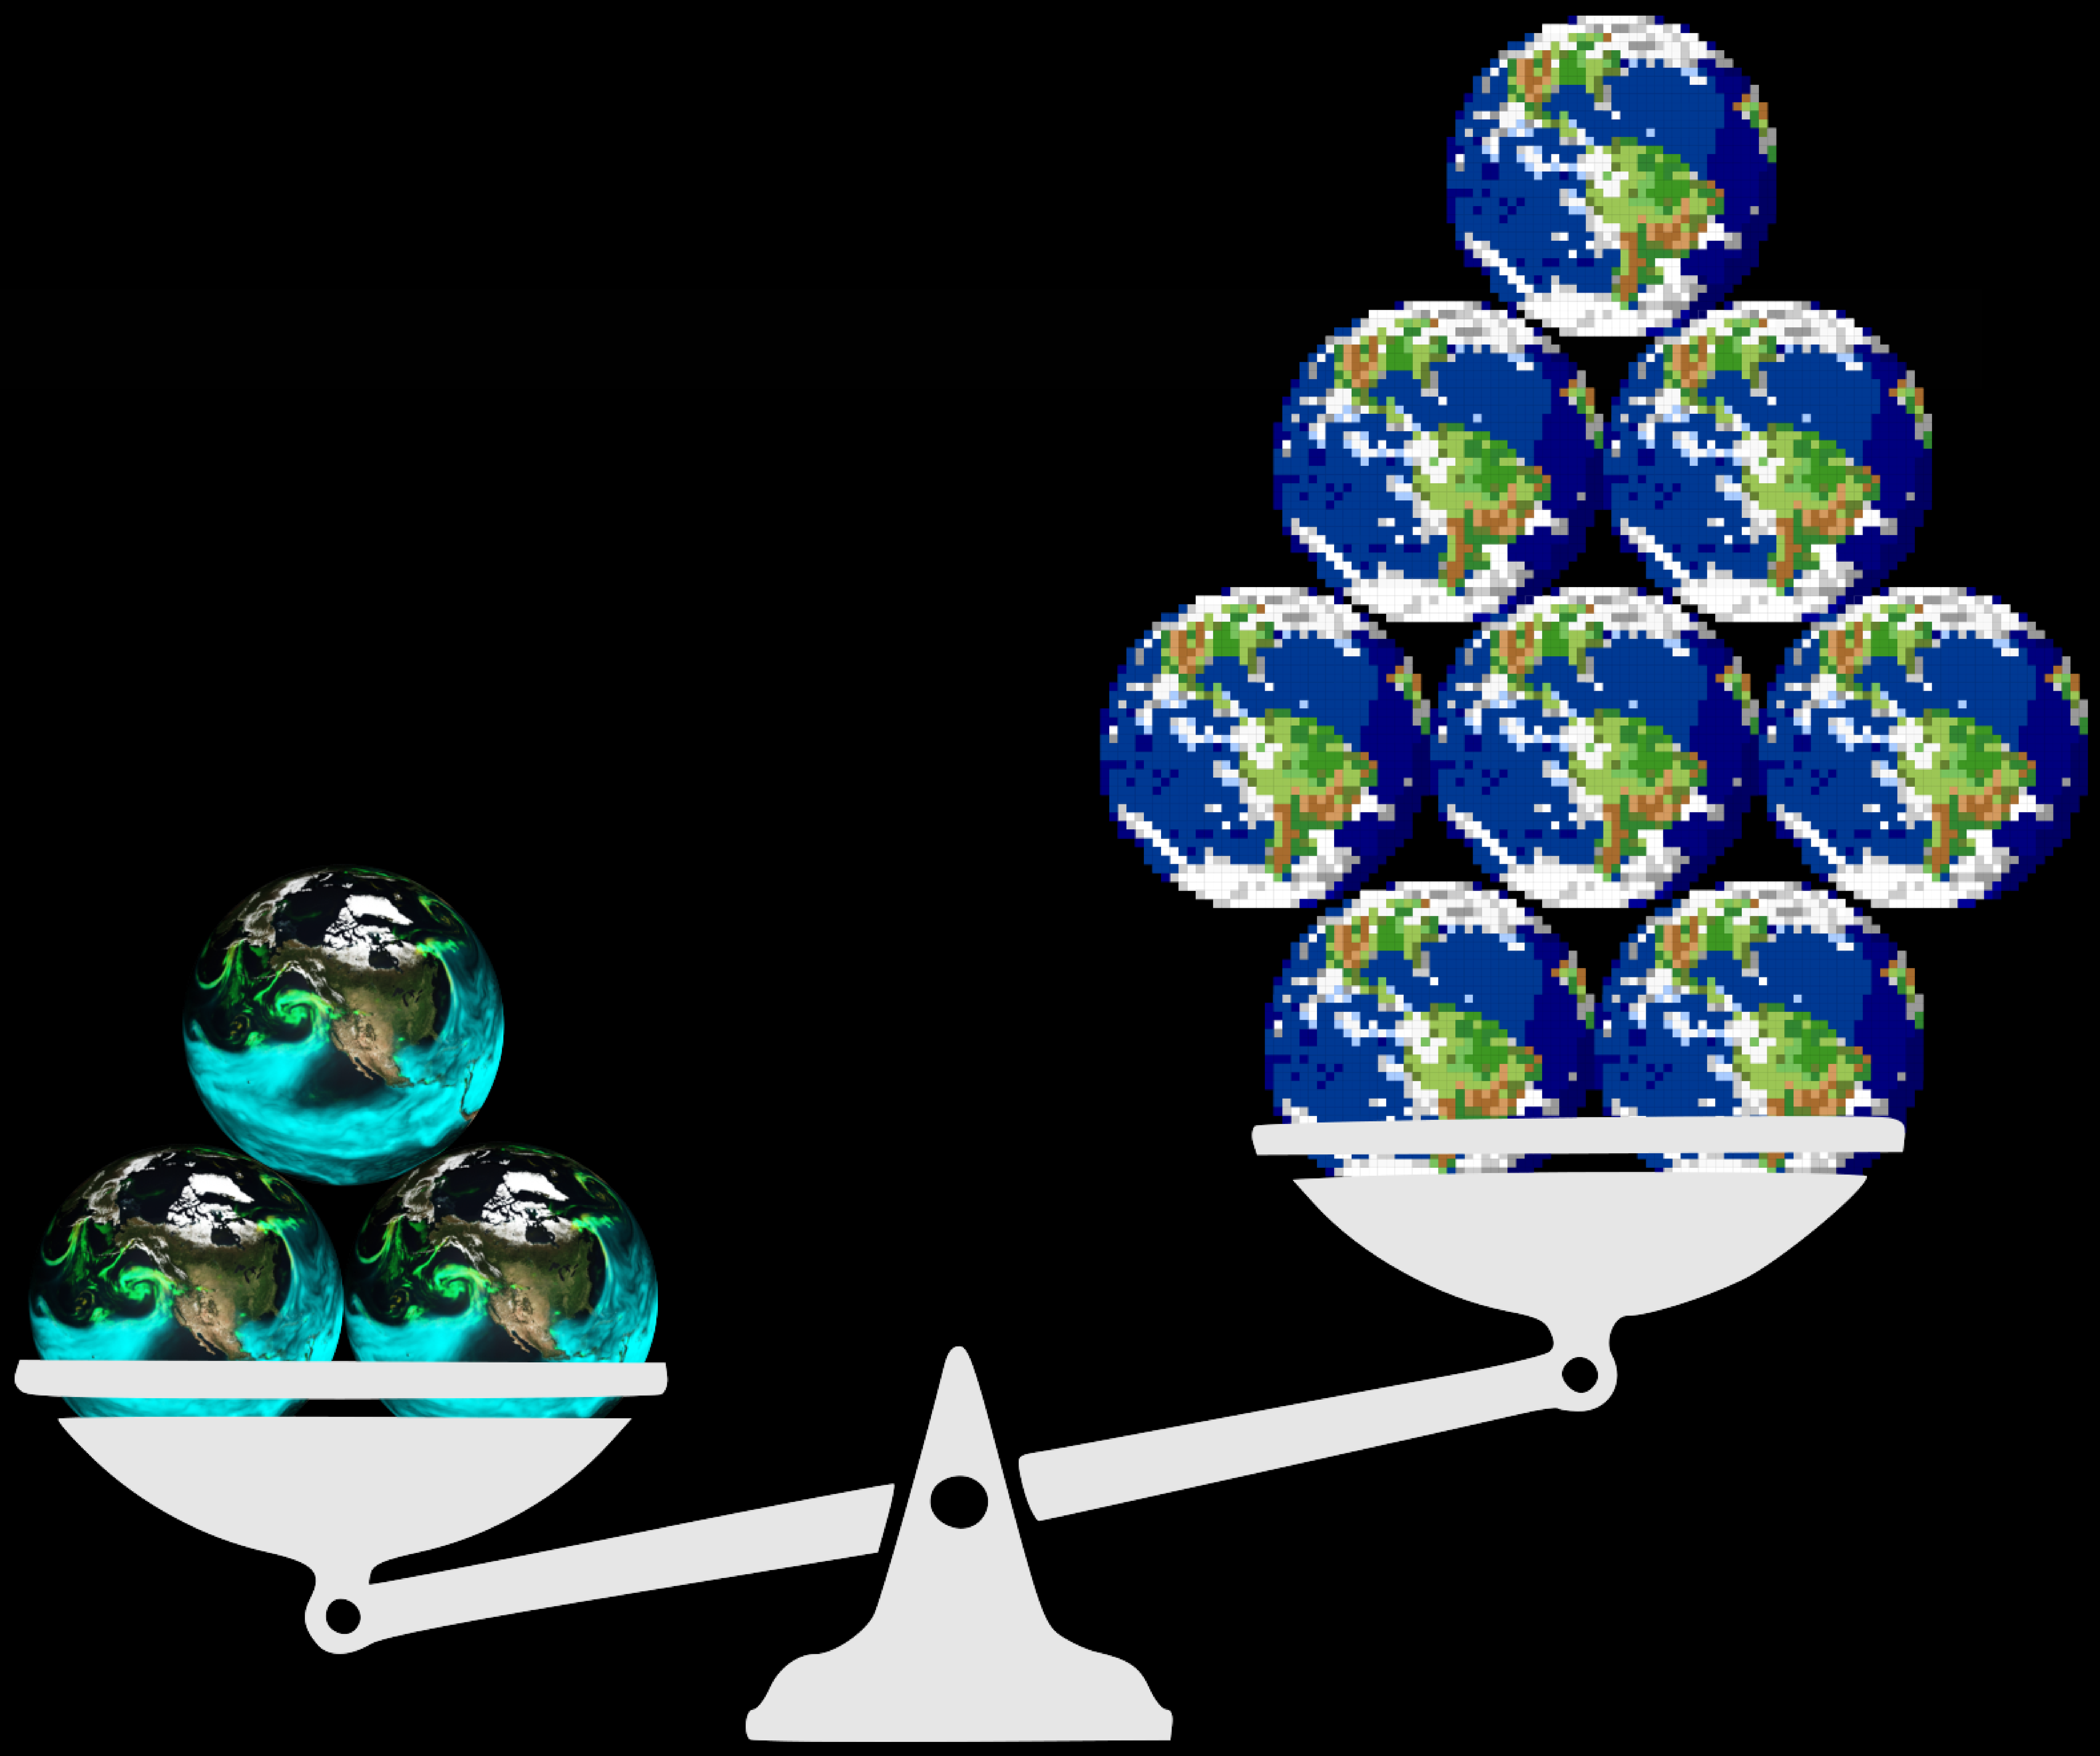

Under computational constraints, the number of forecasts must be balanced against their resolution.

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



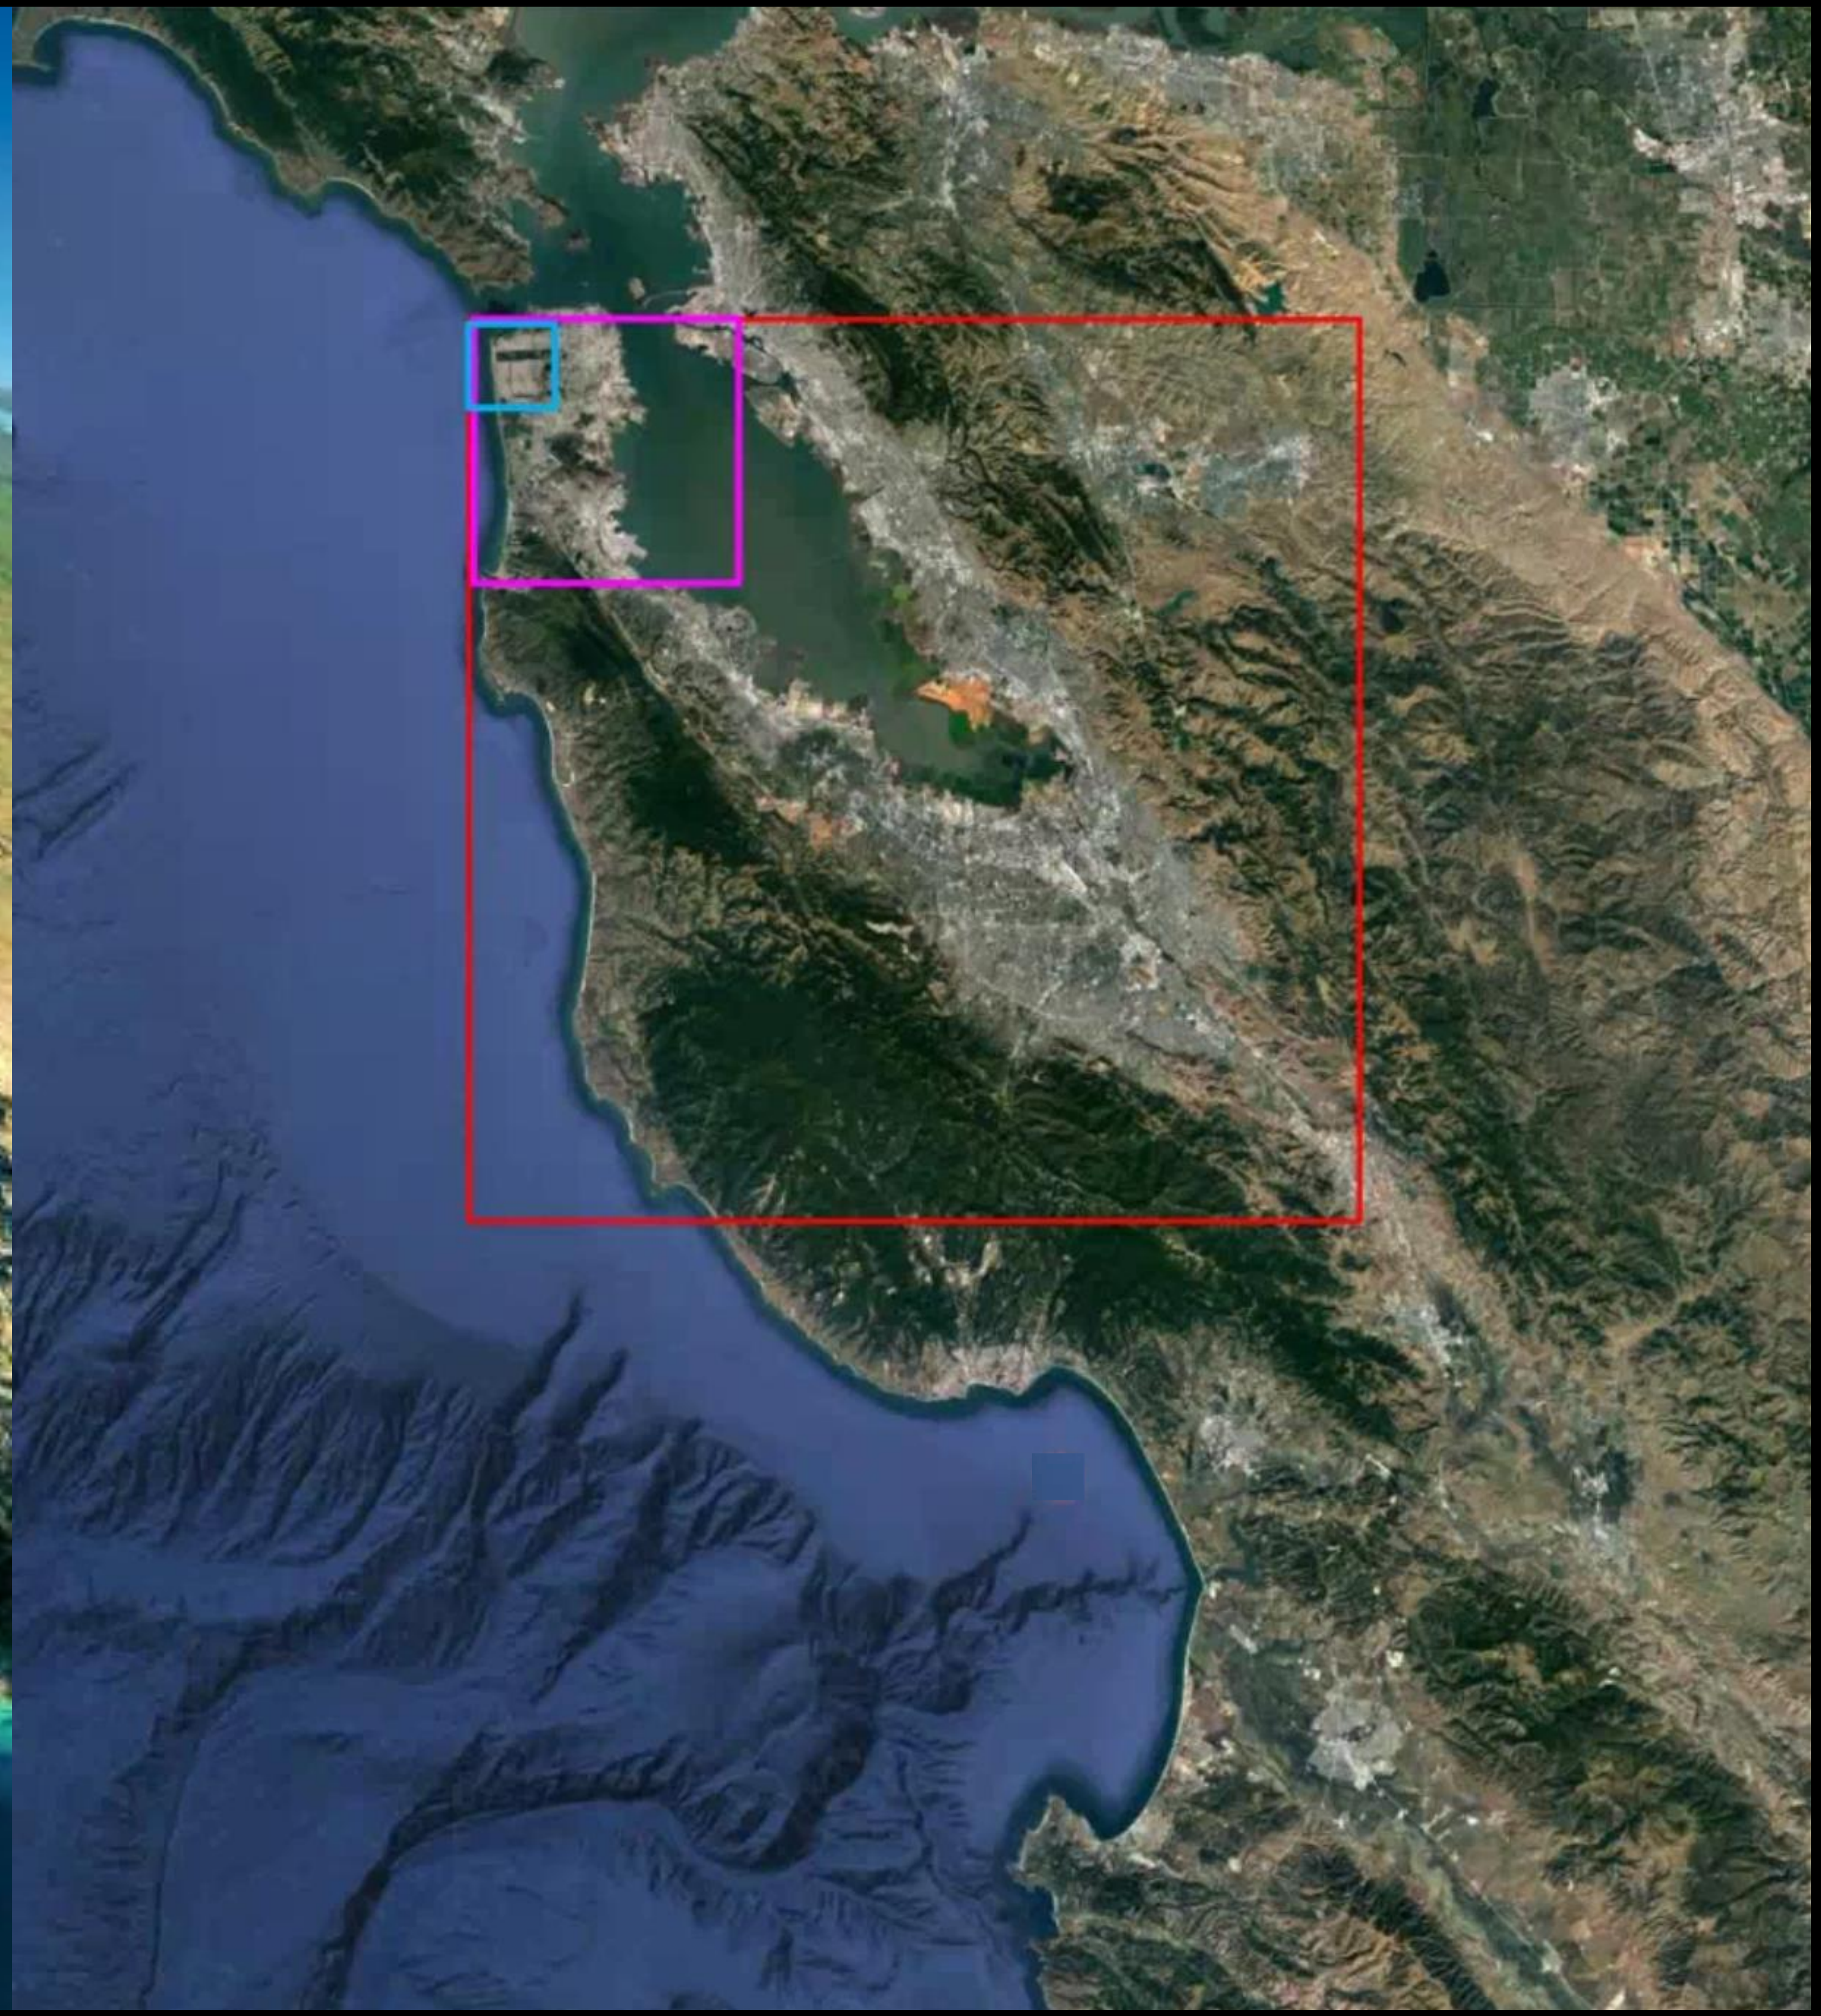

… Answer Questions about Climate Change's Impacts

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

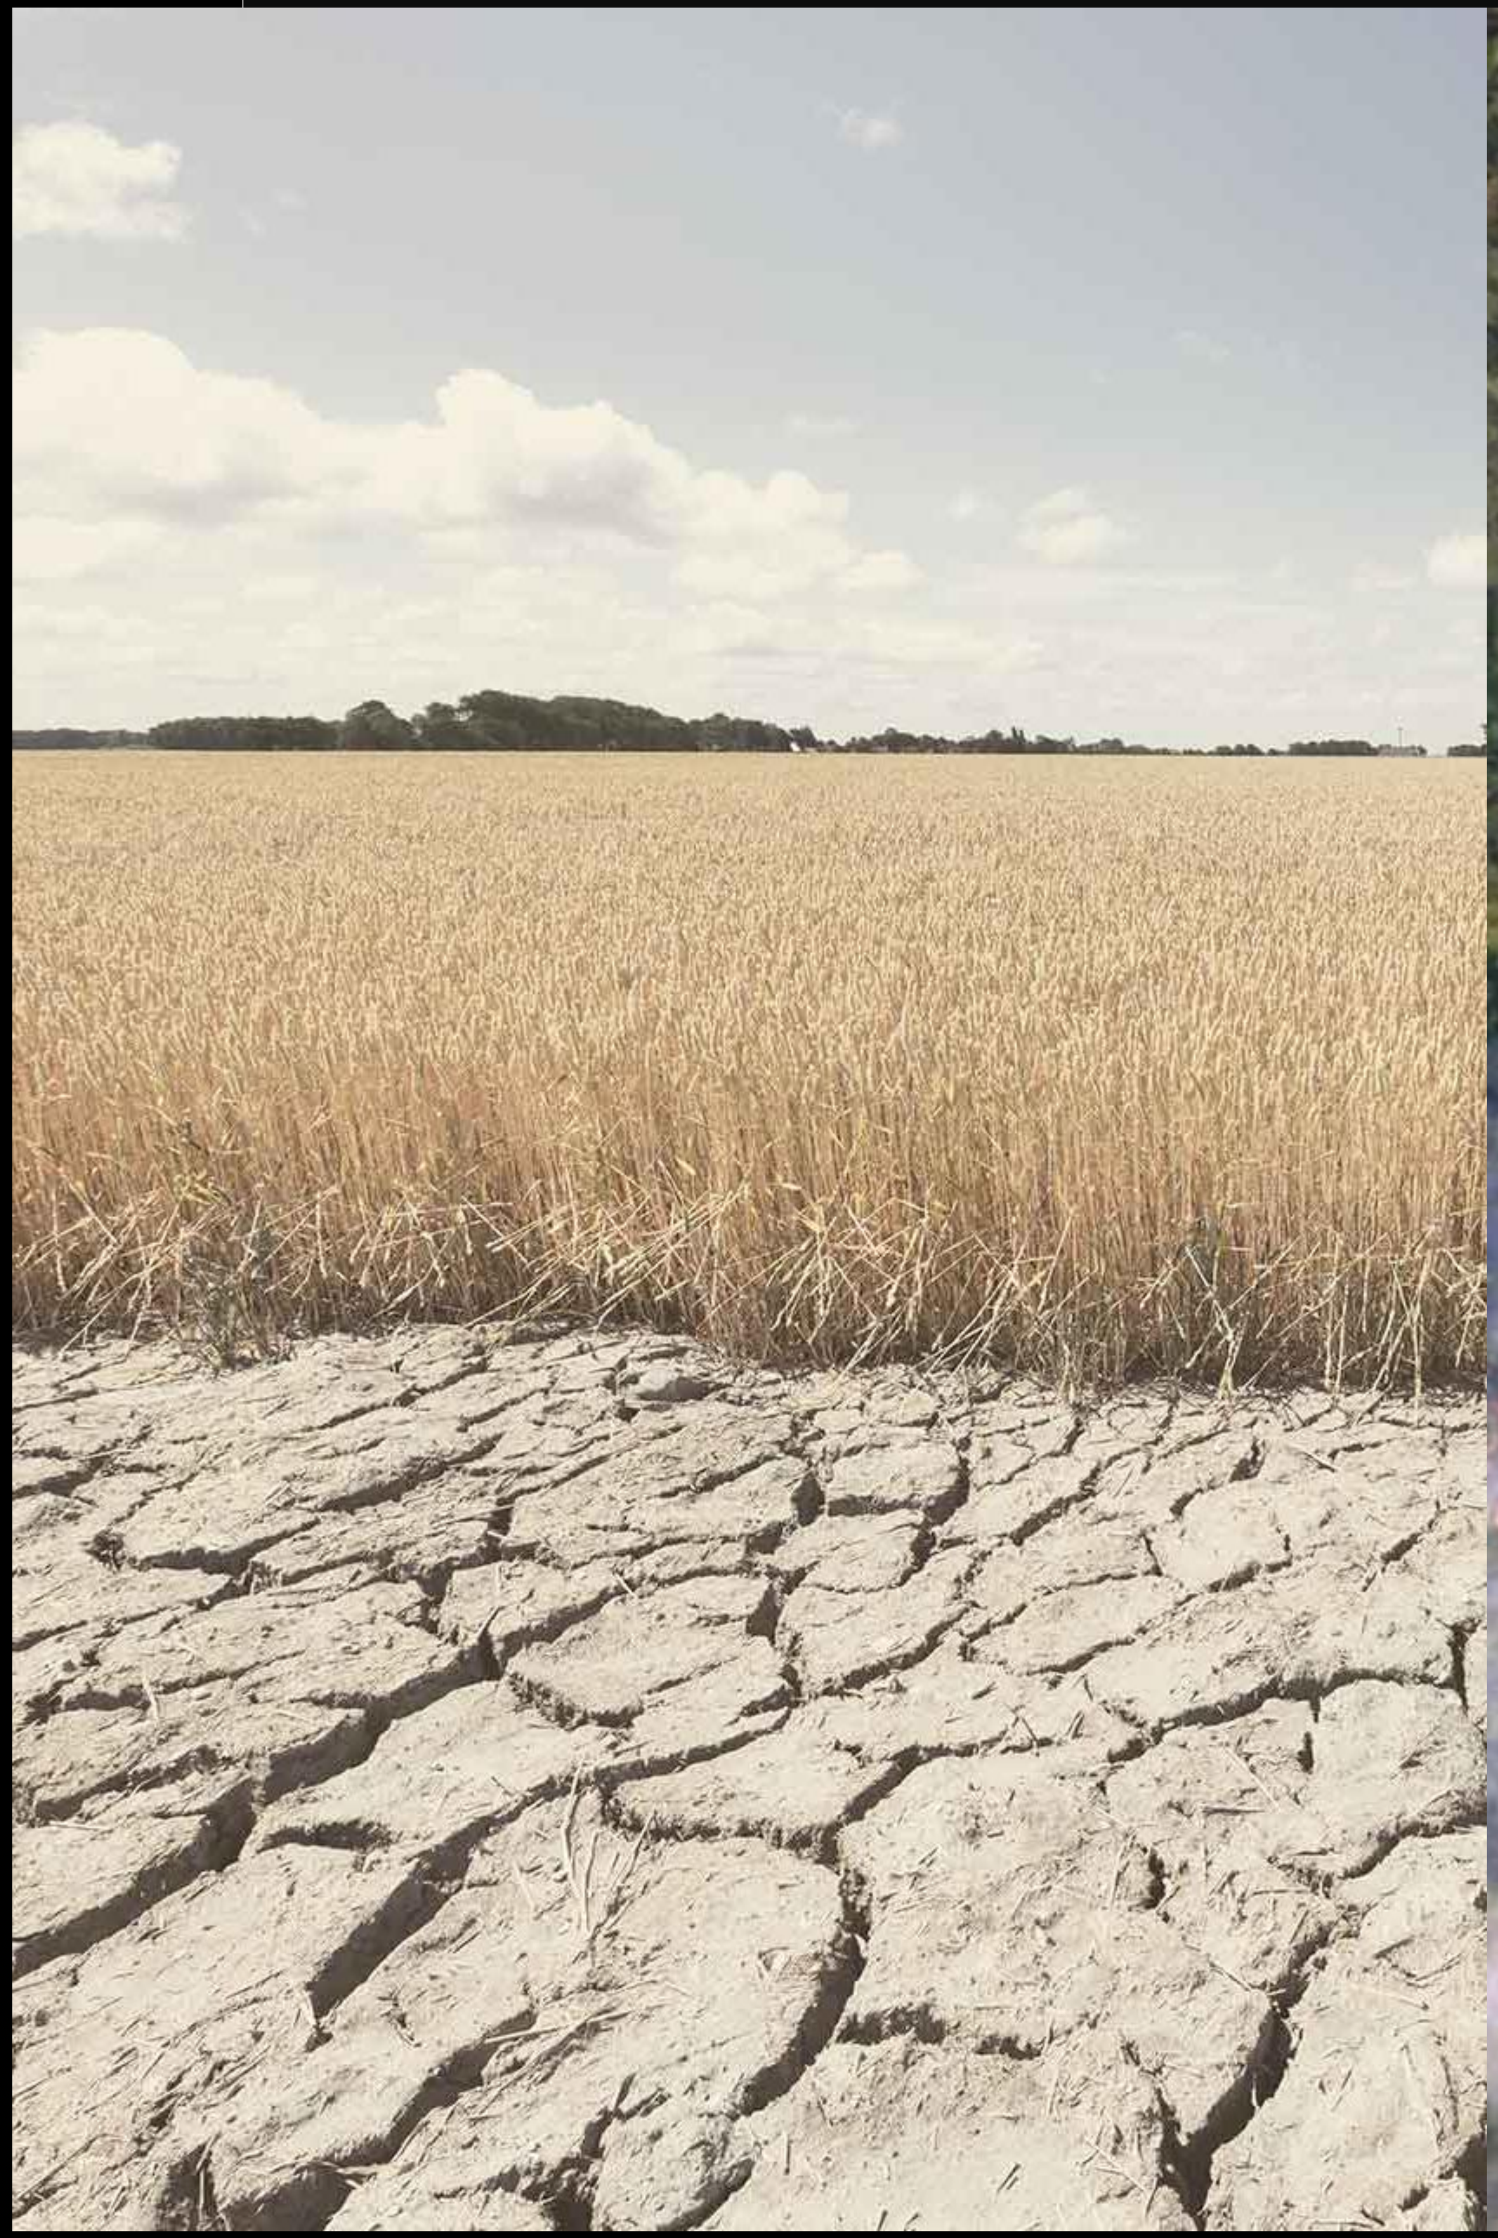

To Explore Consequences of Actions

And Optimize for Desired Outcomes

Research Areas

Mission: Seed New AI Tech Across Earth System Simulation & Informatics Stack

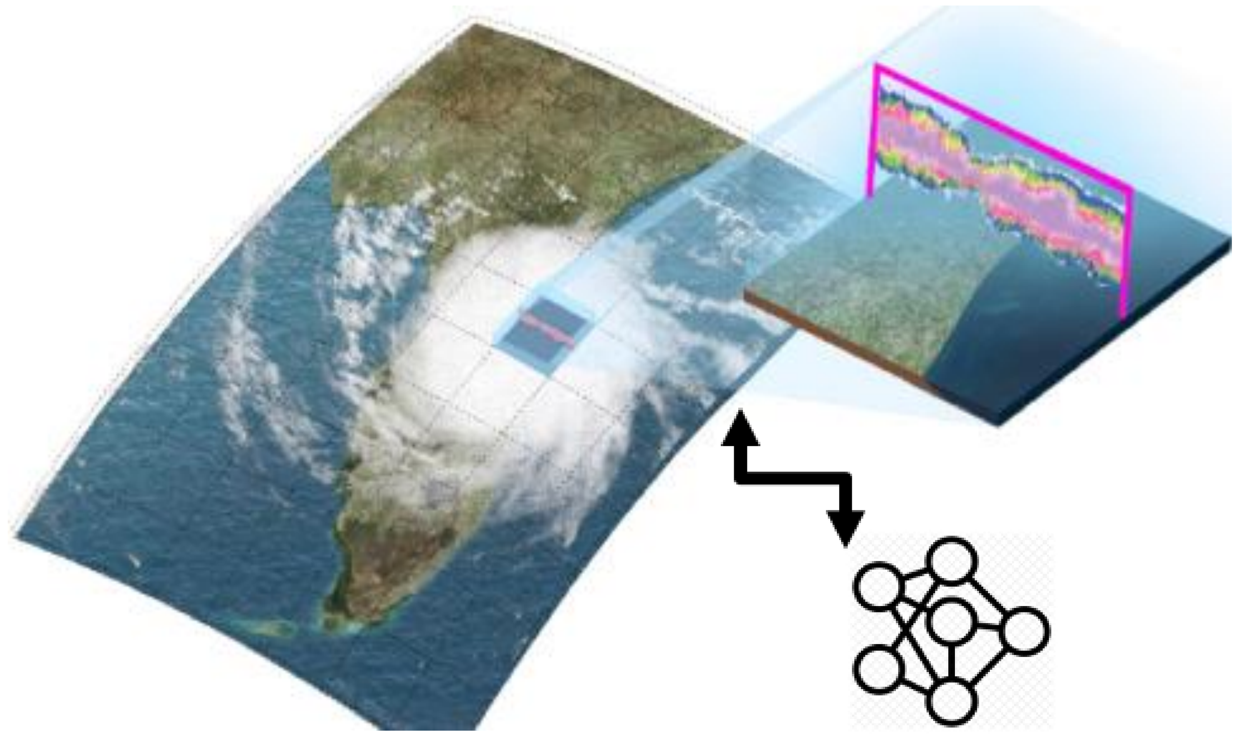

Hybrid Physics-ML Climate Simulation

Accelerating multi-scale predictions to capture storm physics

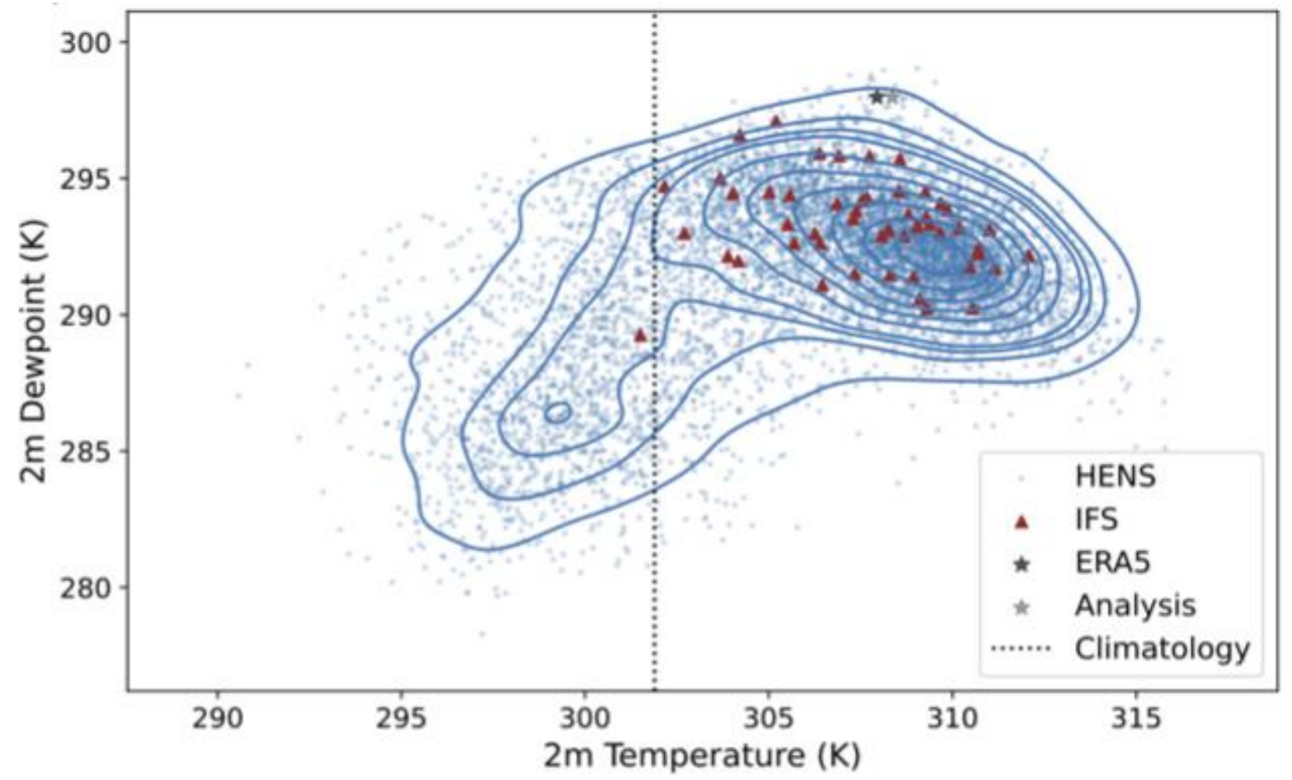

Global AI Weather Prediction

Probabilistic Calibration for Tail Risk Use Cases

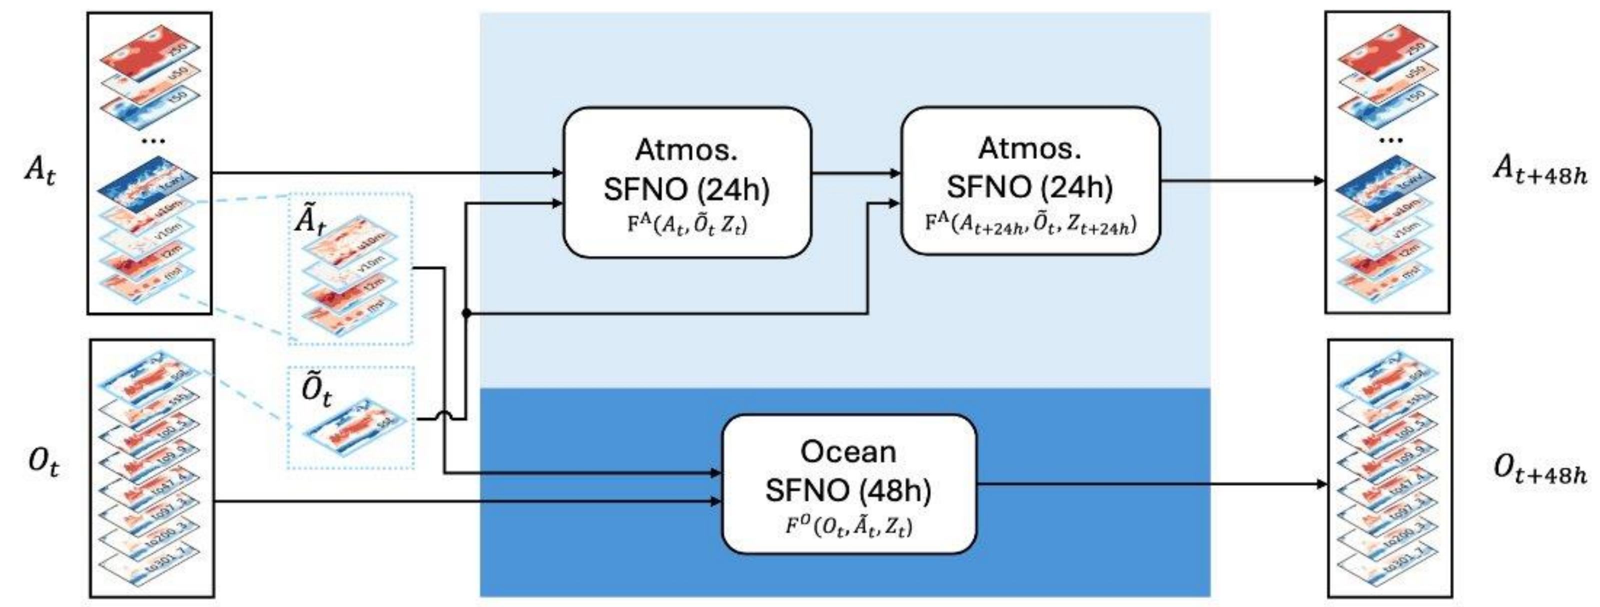

Ocean-Coupled Earth System Prediction

Learning Subseasonal, Seasonal & Climate Physics

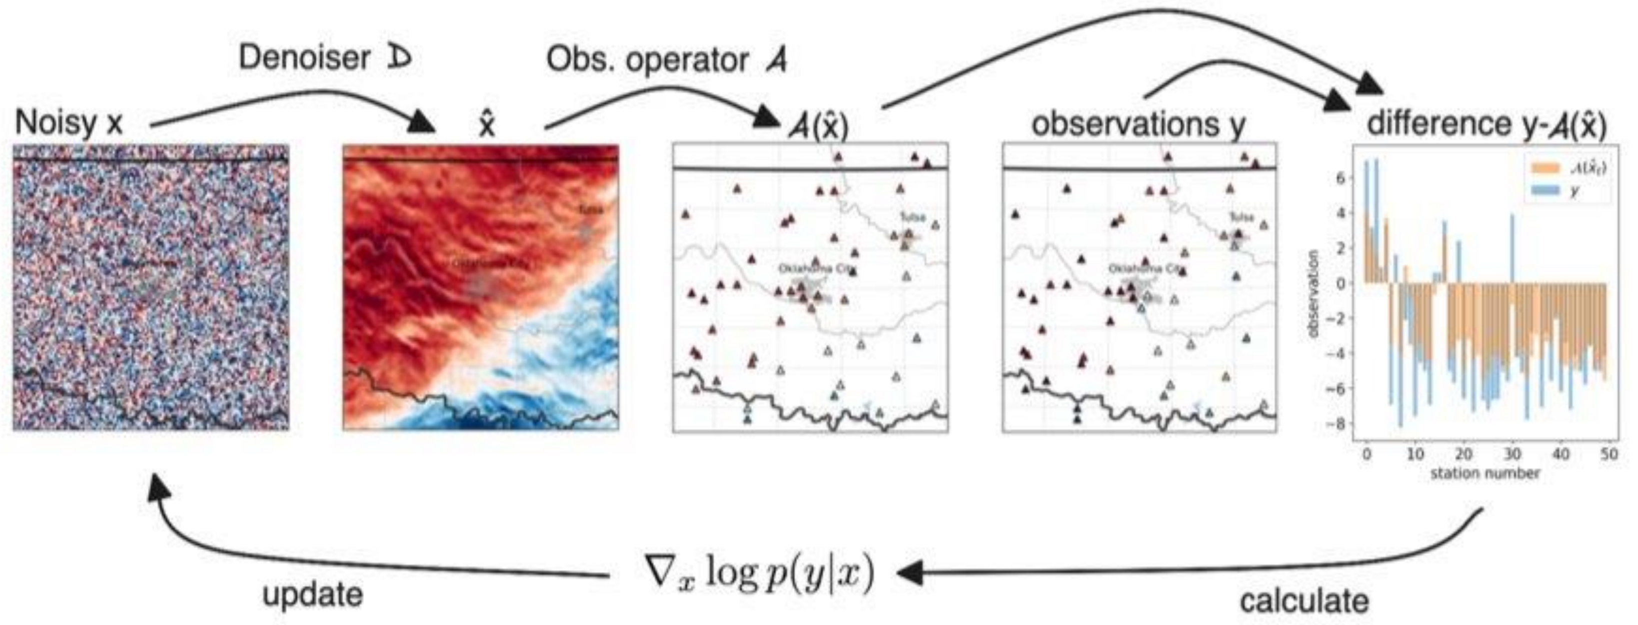

Generative Data Fusion

In-filling & multi-modal atmospheric state estimation

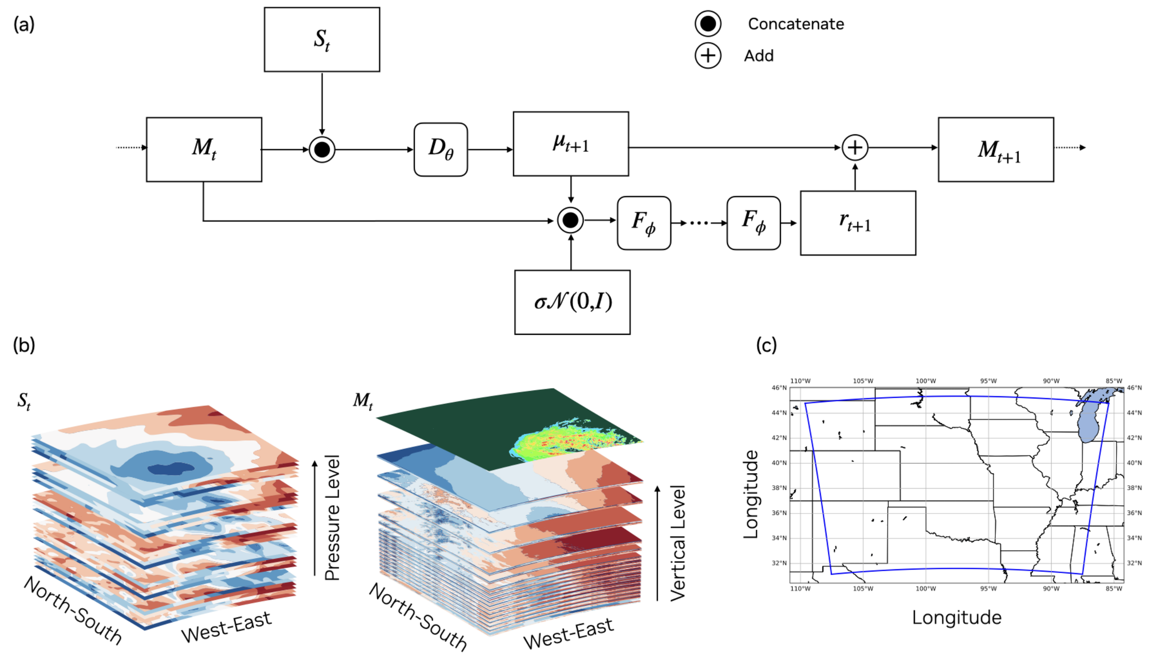

Km-Scale Weather Forecasting

Learning to Predict Fundamental Storm Physics

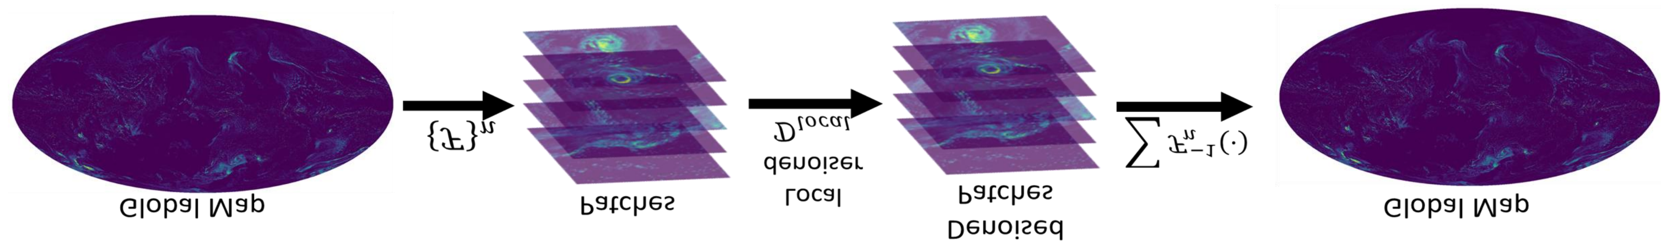

Climate Foundation Modeling

Planetary Diffusion Models For Generative Informatics

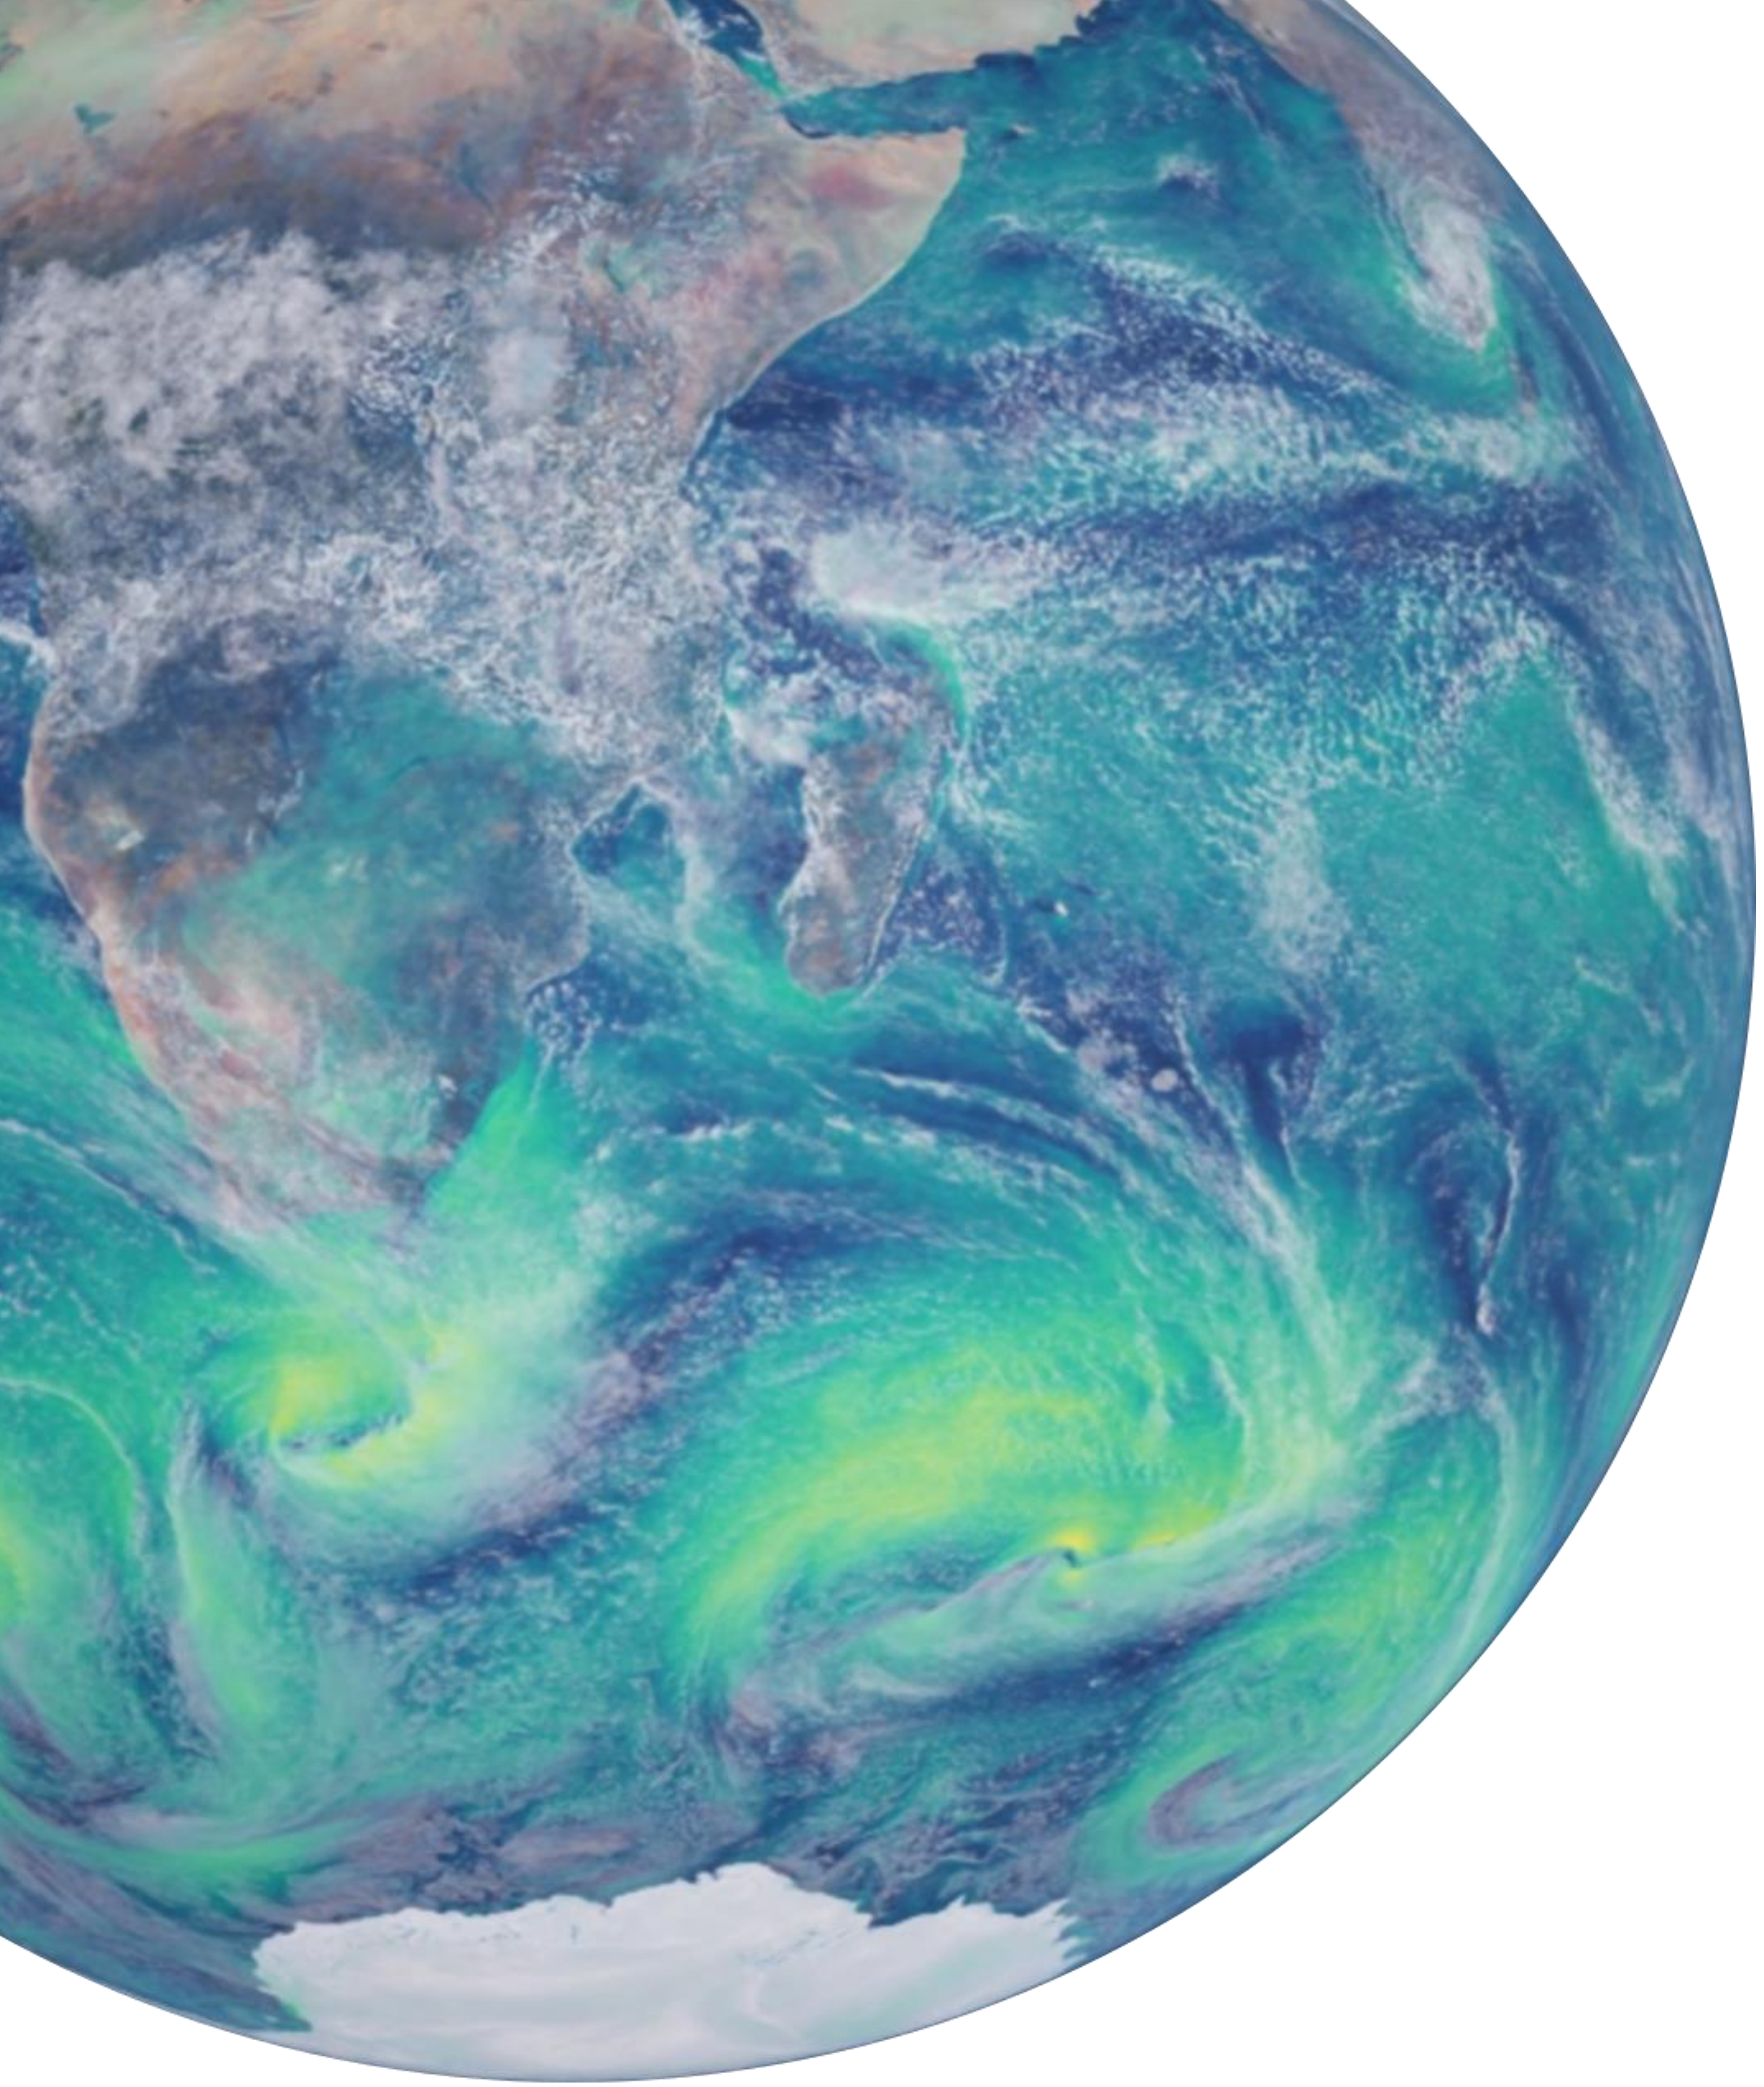

NVIDIA Earth-2 Libraries and Tools

PhysicsNeMo: Training framework for foundational scale models

- Scalable training framework for weather models
- SOTA architectures for weather applications optimized for GPUs
- Training pipelines to train on peta byte scale datasets like ERA5
- Scaling to multi-GPU and multi-node training
- Training recipes for:
  - CorrDiff: Downscaling model for custom region
  - StormCast, ReGen, ….

https://github.com/NVIDIA/PhysicsNeMo-launch/tree/main/examples/weather

NVIDIA Earth-2

Toolbox for understanding the impact of weather and climate

CUDA/ Warp

Simulation

Accelerated numerical simulation of weather and climate physics.

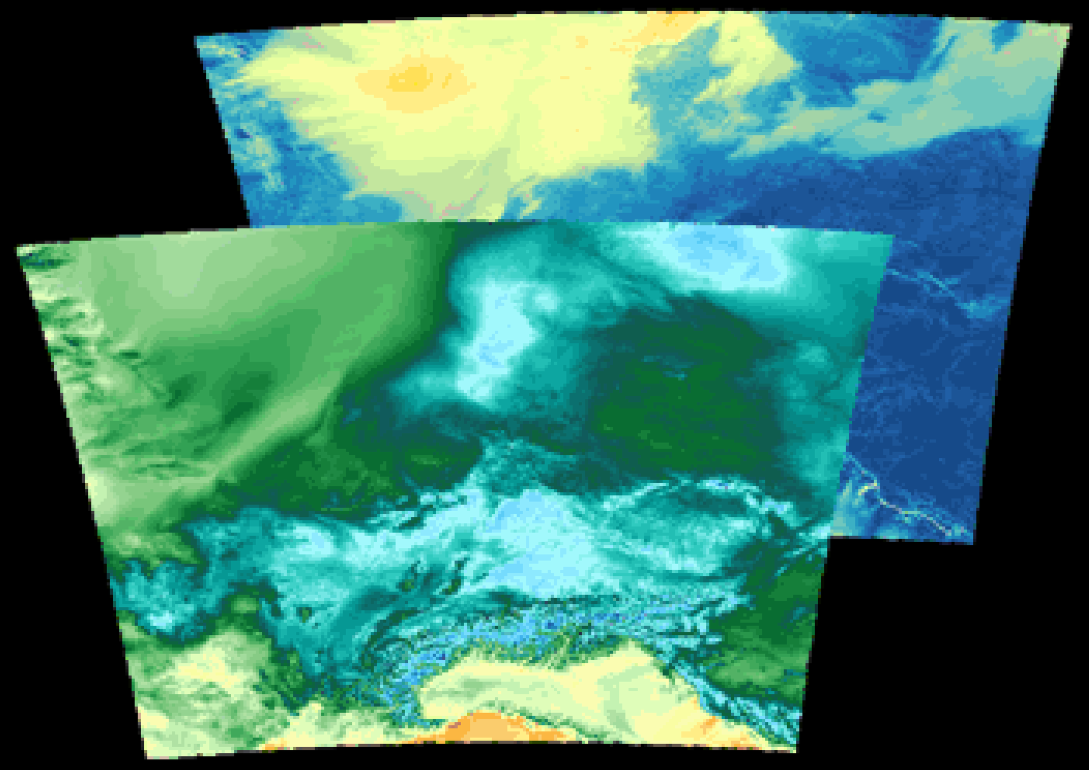

Physics Nemo

AI Training

Training of AI models for forecasting, downscaling, interpolation, and other applications.

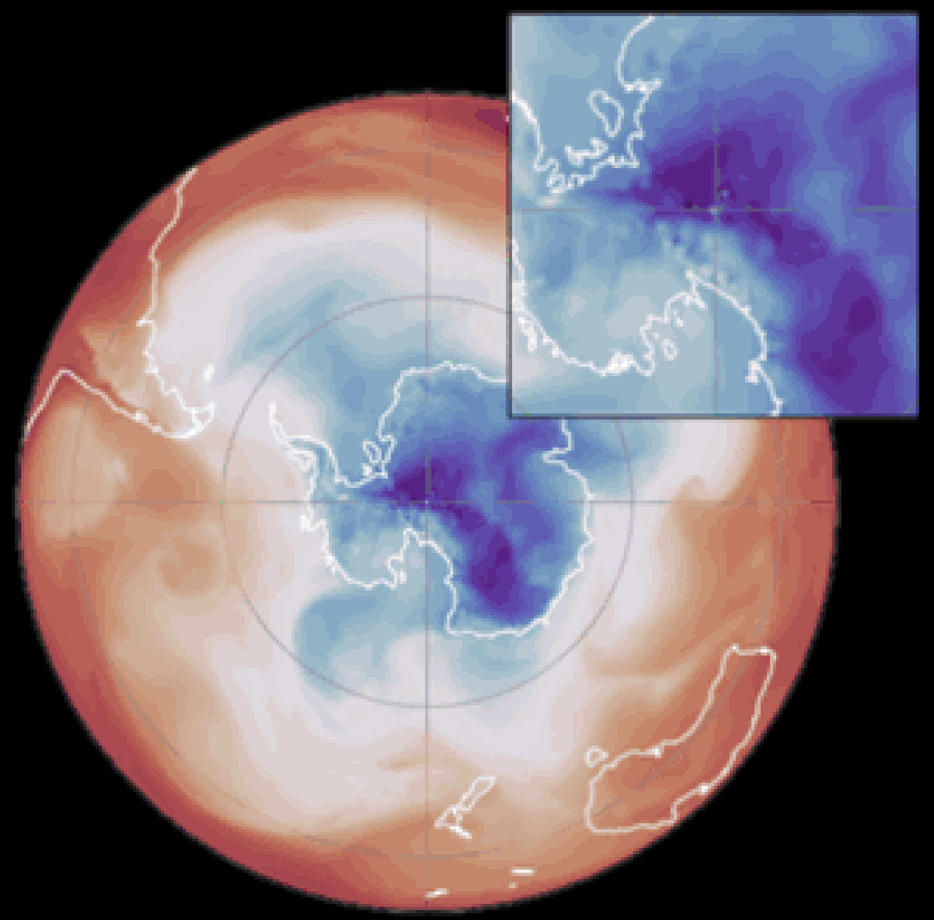

Earth 2 Studio

AI Inference

Efficient inference, connecting AI models, data sources, and downstream applications.

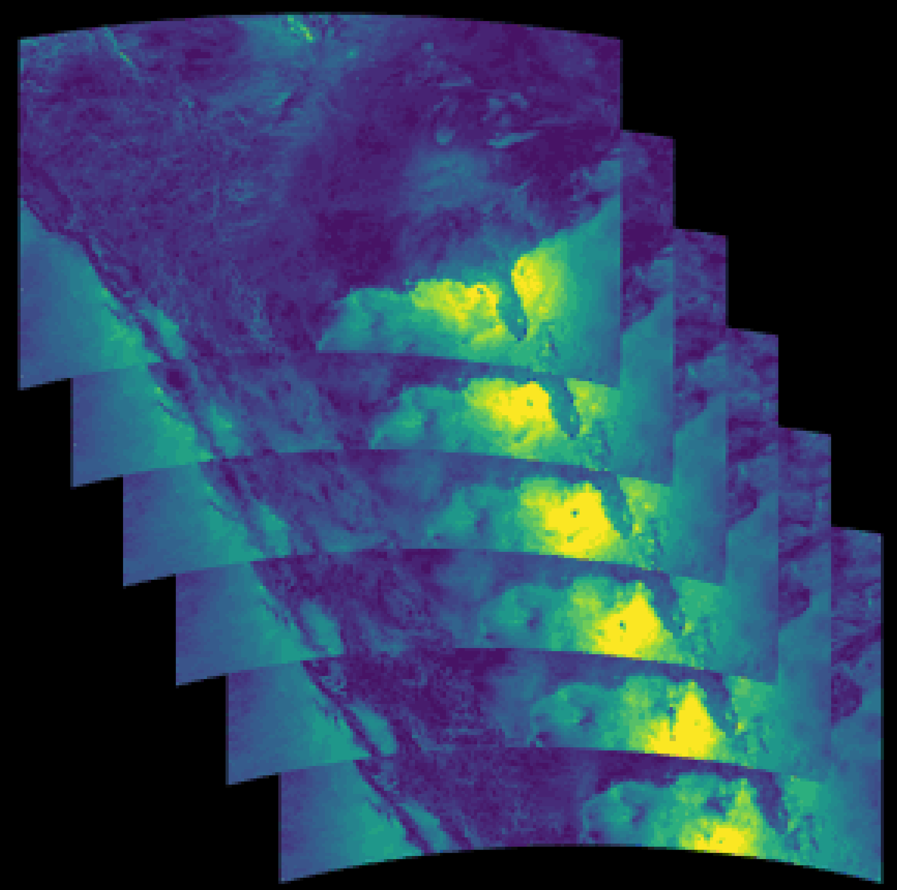

Omniverse

Virtual World

Interactive second Earth

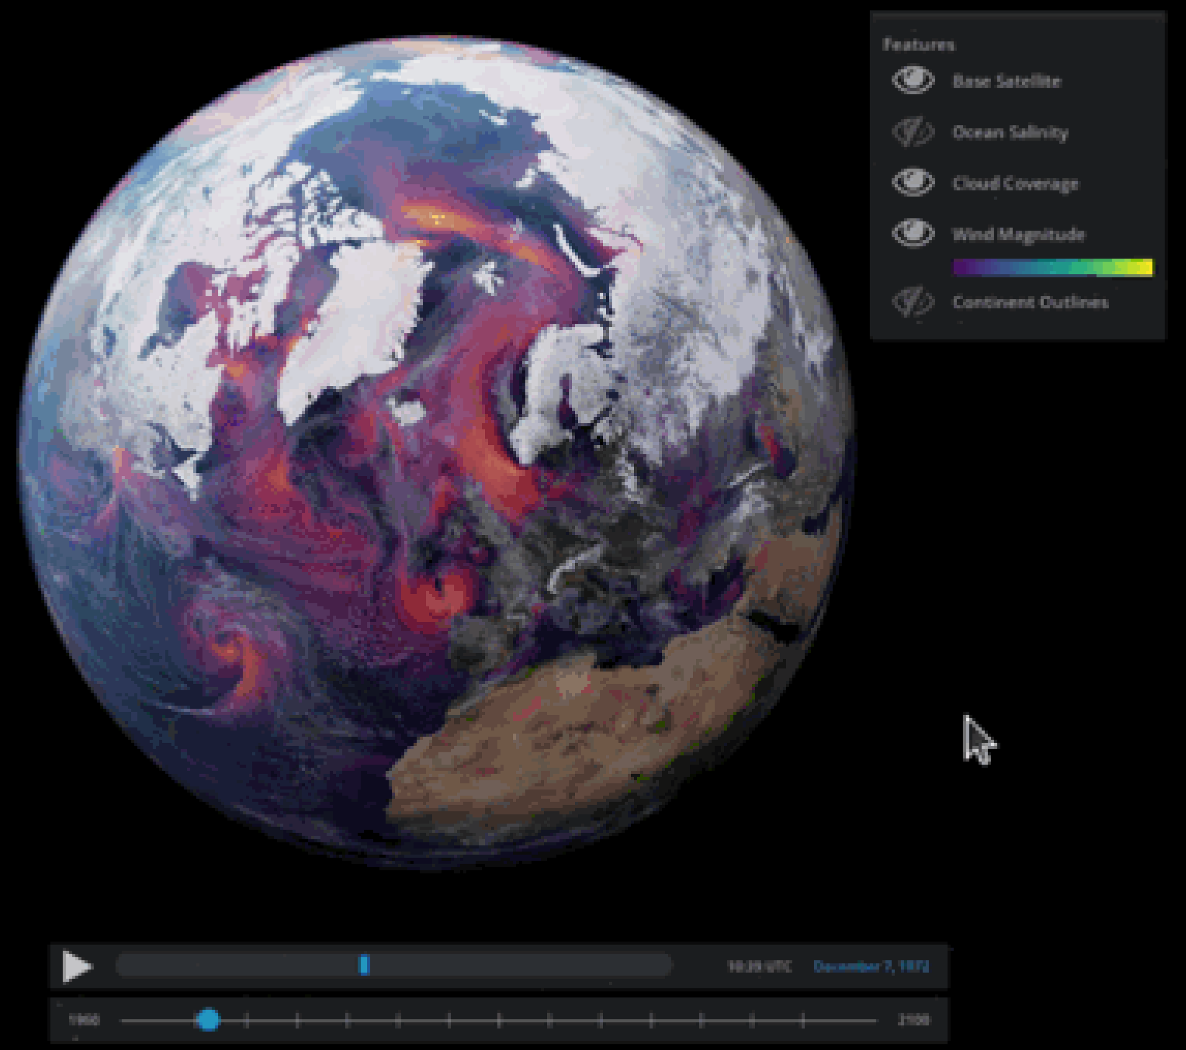

Data Federation

Two-Phase Life Cycle of AI Weather Models

Once trained, AI models make rapid predictions

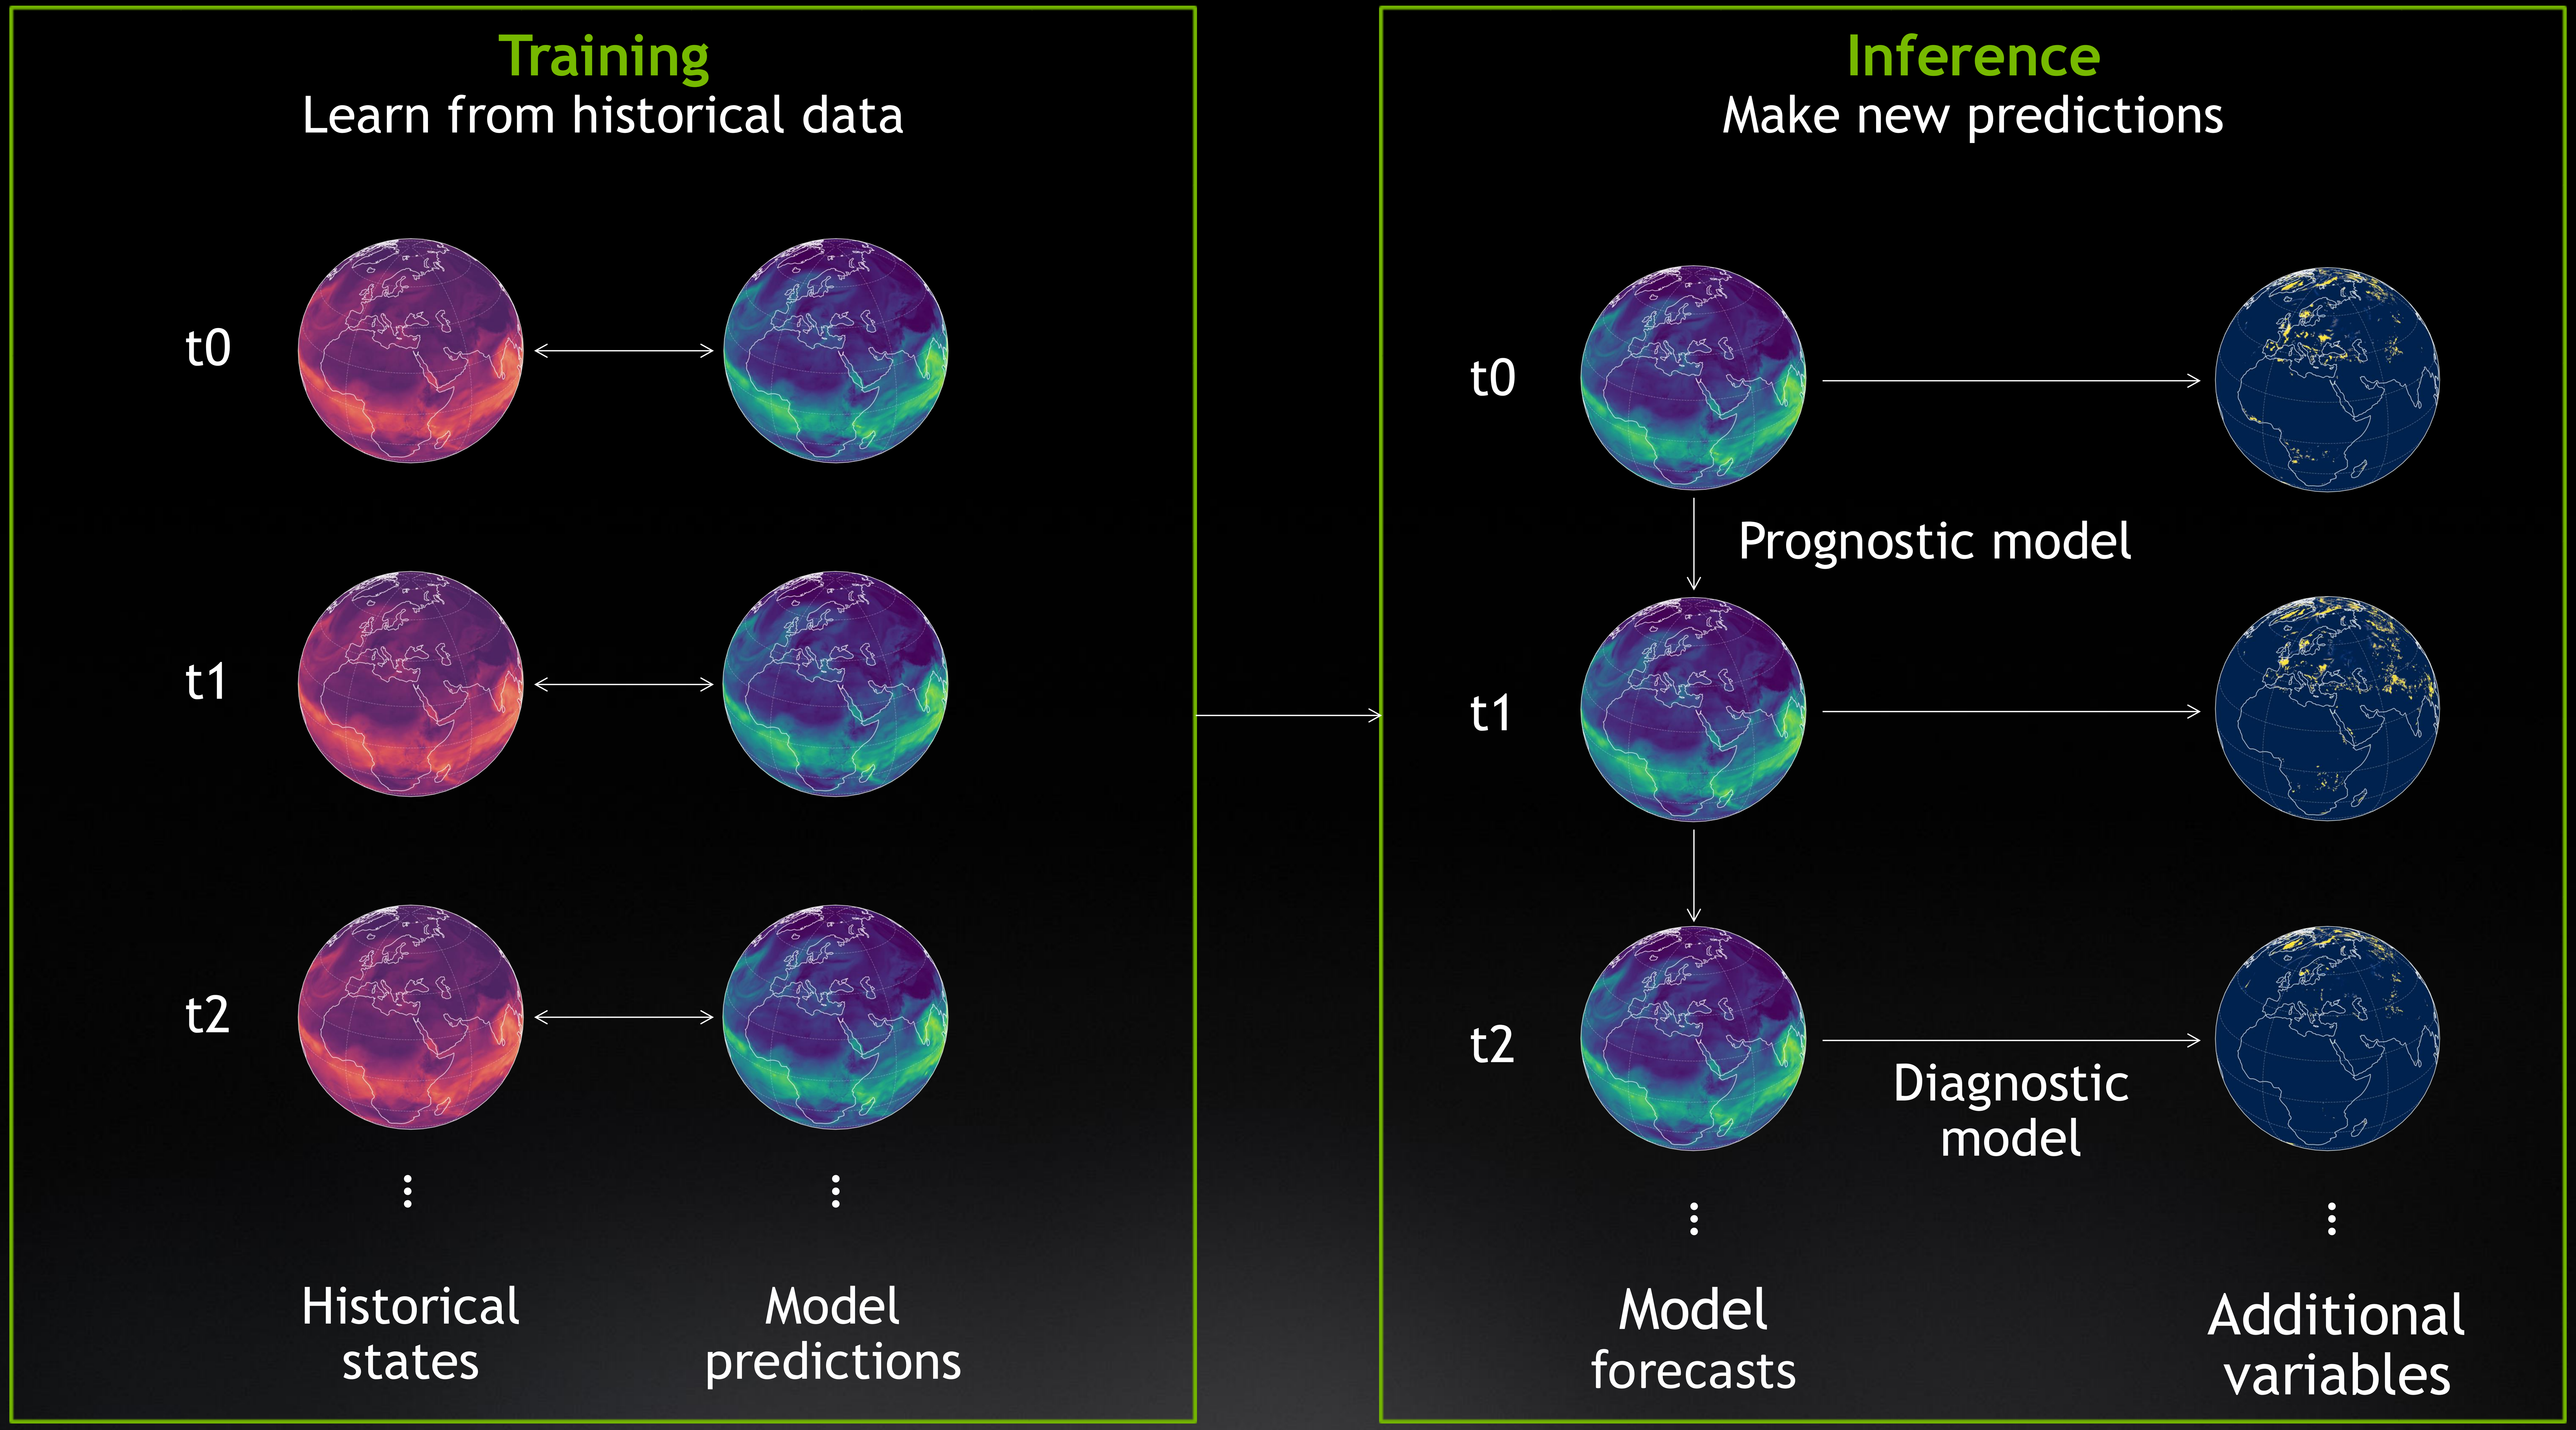

Inference: Earth2Studio

High-level API for rapid experimentation and deployment

- **Easy:** Open-source Python library for accessible AI weather model inference
- **Modular:** Combine data sources, perturbations, models, and IO utilities into workflows
- **Model agnostic:** Get immediately started with pretrained models or bring your own model

https://github.com/NVIDIA/earth2studio

Forecasting

- FourCastNet (AFNO/SFNO) — Publication
- FengWu (Transformer) — Publication
- FuXi (Transformer) — Publication
- PanguWeather (Transformer) — Publication
- Aurora (Transformer) — Publication
- DLWP (CNN) — Publication
- AIFS (GNN) — Publication
- GraphCast (GNN) — Publication
- Coming soon: DLWP HEALPix (CNN) — Publication

Downscaling

- CorrDiff (Diffusion) — Publication
- StormCast (Diffusion) — Publication

Diagnostics

- Precipitation (AFNO) — Publication
- Tropical cyclones (CNN) — Publication
- Atmospheric rivers (CNN) — Publication
- Temporal interpolation (AFNO) — Publication
- Solar irradiance (AFNO) — Publication
- More to come…

Self-paced Course

Applying AI Weather Models with NVIDIA Earth-2

Explore state-of-the-art AI weather prediction models and learn how to integrate them into custom workflows.

Self-paced courses are temporarily unavailable for purchase outside the USA as we transition to a new ecommerce system. We apologize for any inconvenience. **Free courses** remain available for enrollment.

About Course Objectives Stay Informed Contact Us **Enroll Now**

About this Course

This course is free for a limited time.

Weather forecasts are indispensable for planning and decision-making in the public and private sector, with weather affecting anything from supply chain resiliency to energy production. Traditional numerical weather prediction systems are difficult to operate and place heavy demands on time and compute resources. With the recent advances in AI weather modeling, non-expert practitioners are now enabled to run forecasts tuned to their own needs. This course explores the possibilities offered by state-of-the-art AI weather prediction models and teaches how to integrate them into custom workflows. In this course, students will learn how AI weather models are revolutionizing the approach to weather forecasting. They will gain hands-on experience running AI weather forecasts, validating model outputs, and explore how super-resolution AI models can make fine-grained predictions. After the course, students will be able to build their own custom AI weather pipelines.

Learning Objectives

Course Details

Duration: 03:00

Price: Free

Level: Technical - Beginner

Subject: Deep Learning

Language: English

Course Prerequisites:

- Basic familiarity with Python.
- Familiarity with Deep Learning beneficial but not required.

Prefer learning from an instructor?

Request a private workshop or view our public workshop

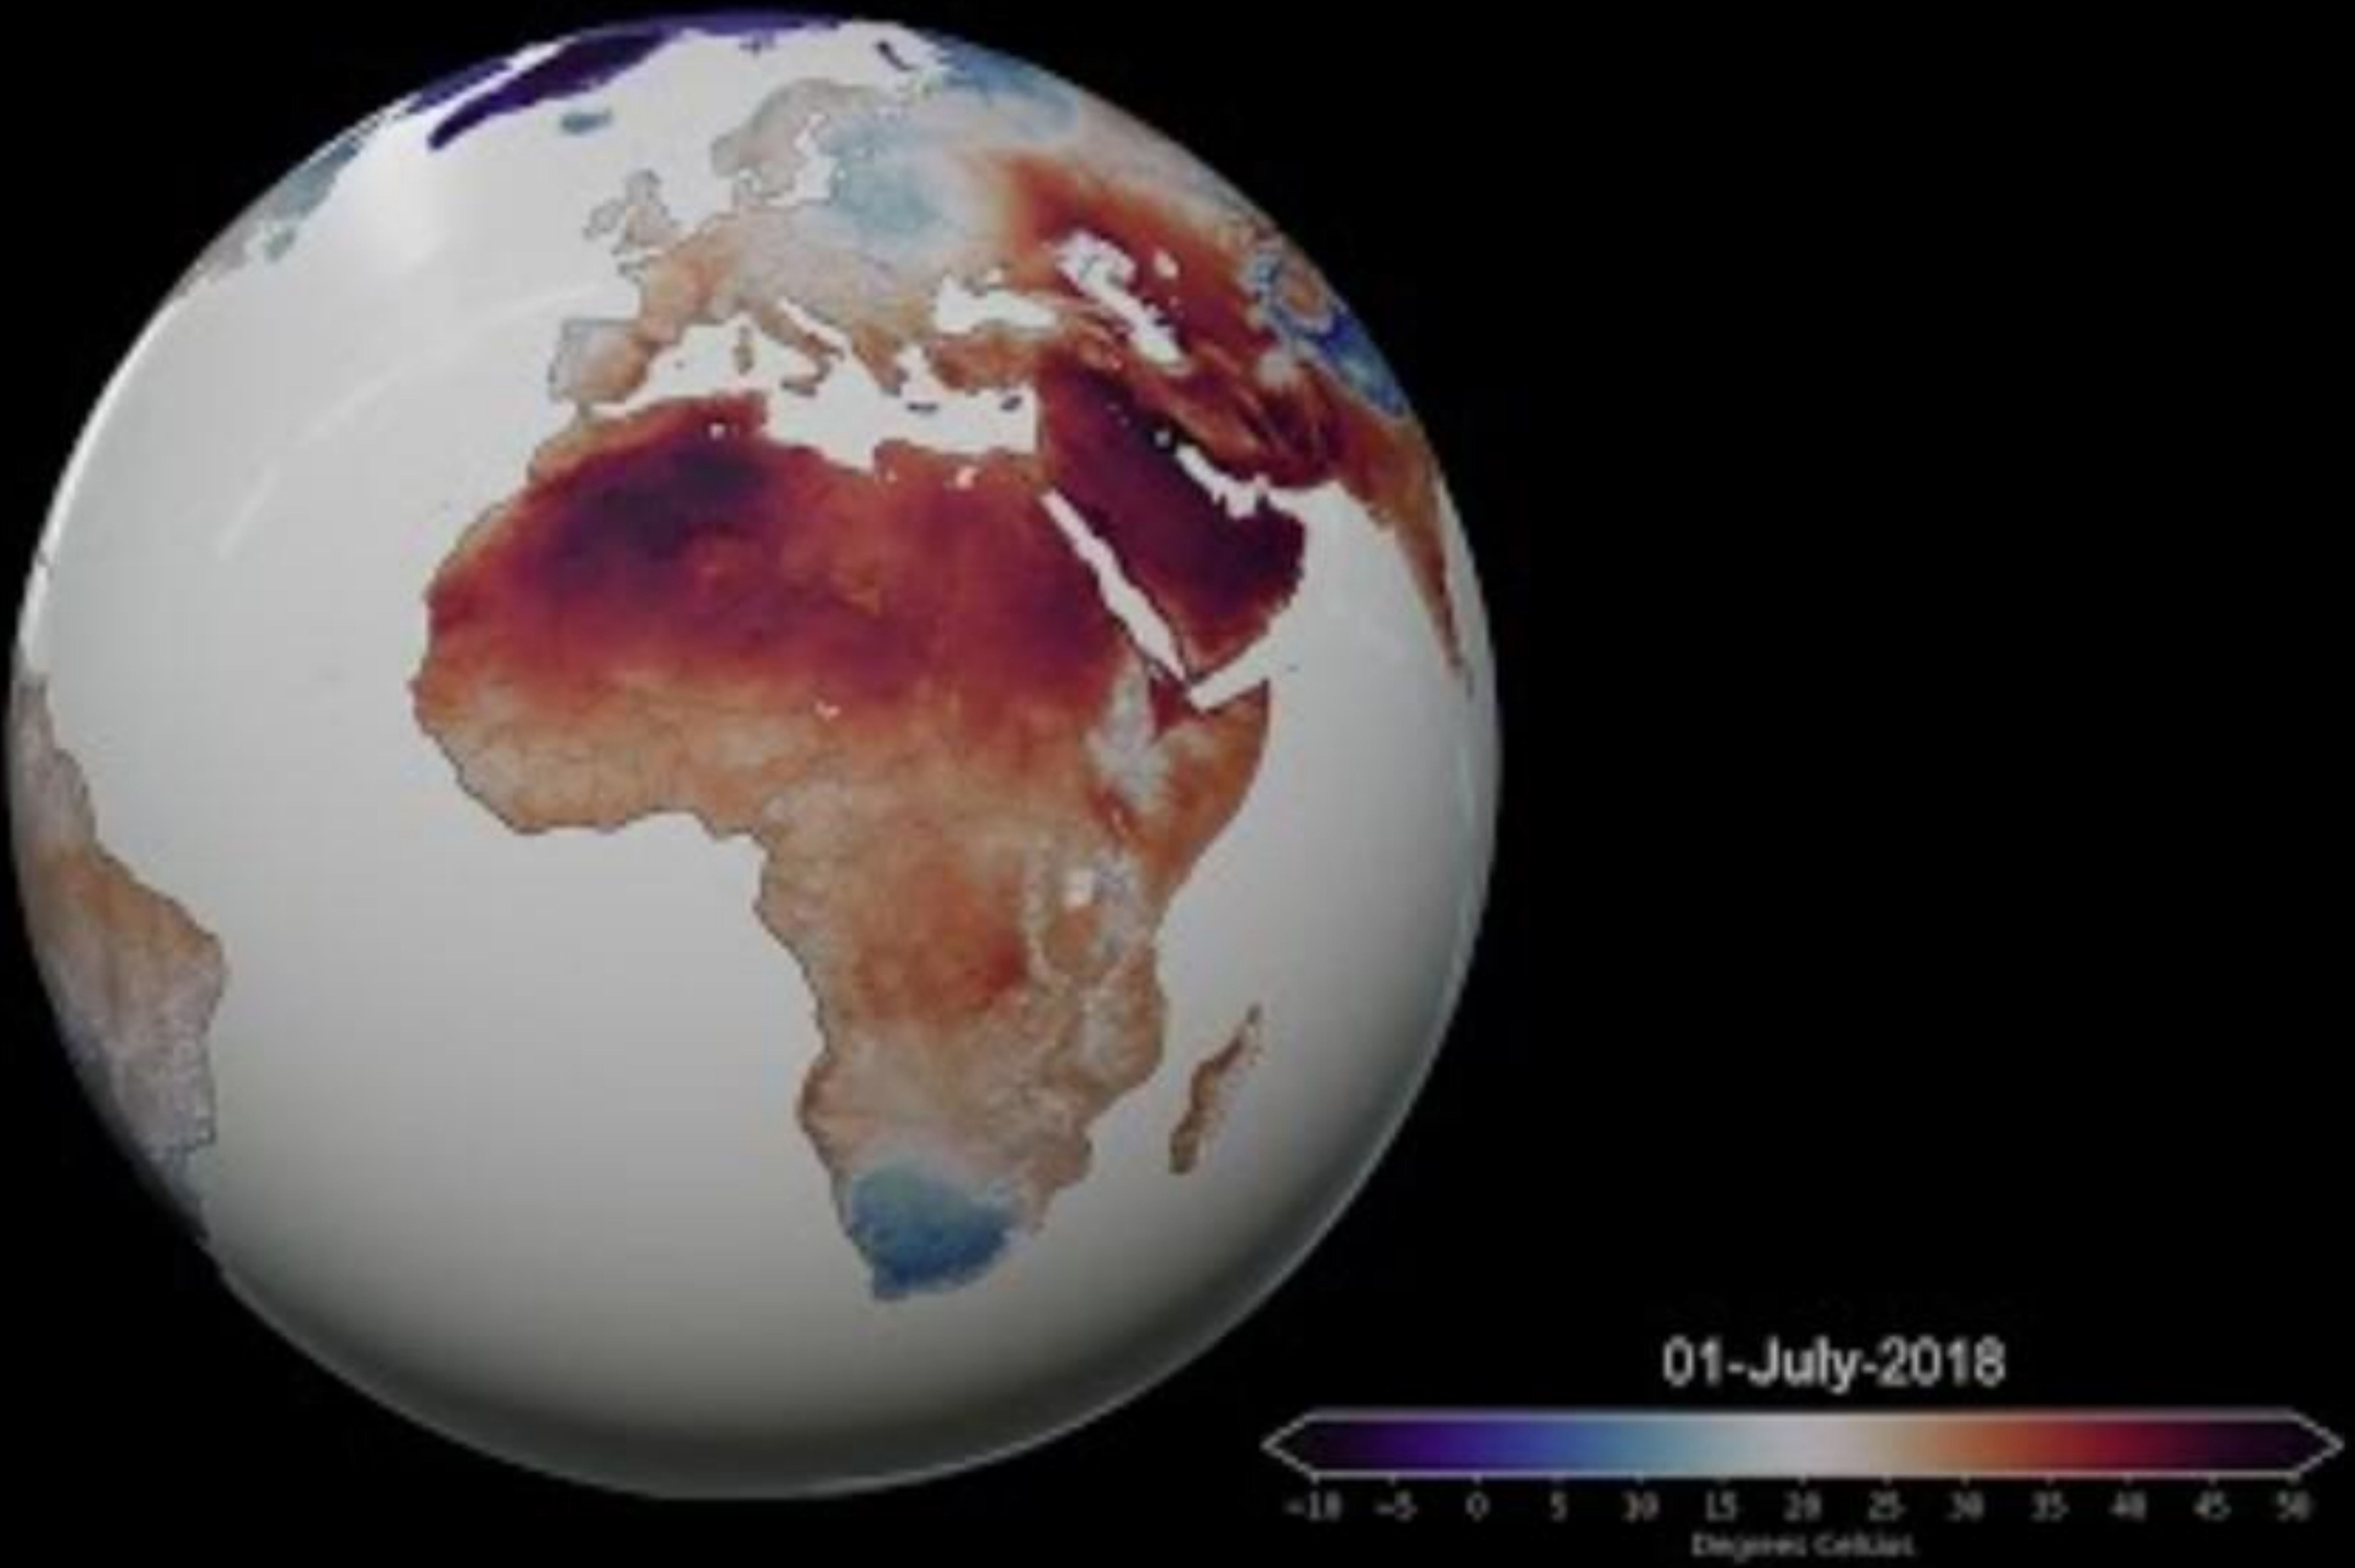
01-July-2018
-10 -5 0 5 10 15 20 25 30 35 40 45 50
Degrees Celsius

Produce an AI Weather Forecast with a few lines of code

https://github.com/NVIDIA/earth2studio

```python
from earth2studio.models.px import FCN
from earth2studio.data import GFS
from earth2studio.io import ZarrBackend
import earth2studio.run as run

# Load FourCastNet pretrained model
model = FCN.load_model(FCN.load_default_package())

# Create the data source
data = GFS()

# Create a Zarr IO Backend
io = ZarrBackend()

# Run 20 steps of inference
output_datastore = run.deterministic(["2024-01-01"], 20, model, data, io)
```

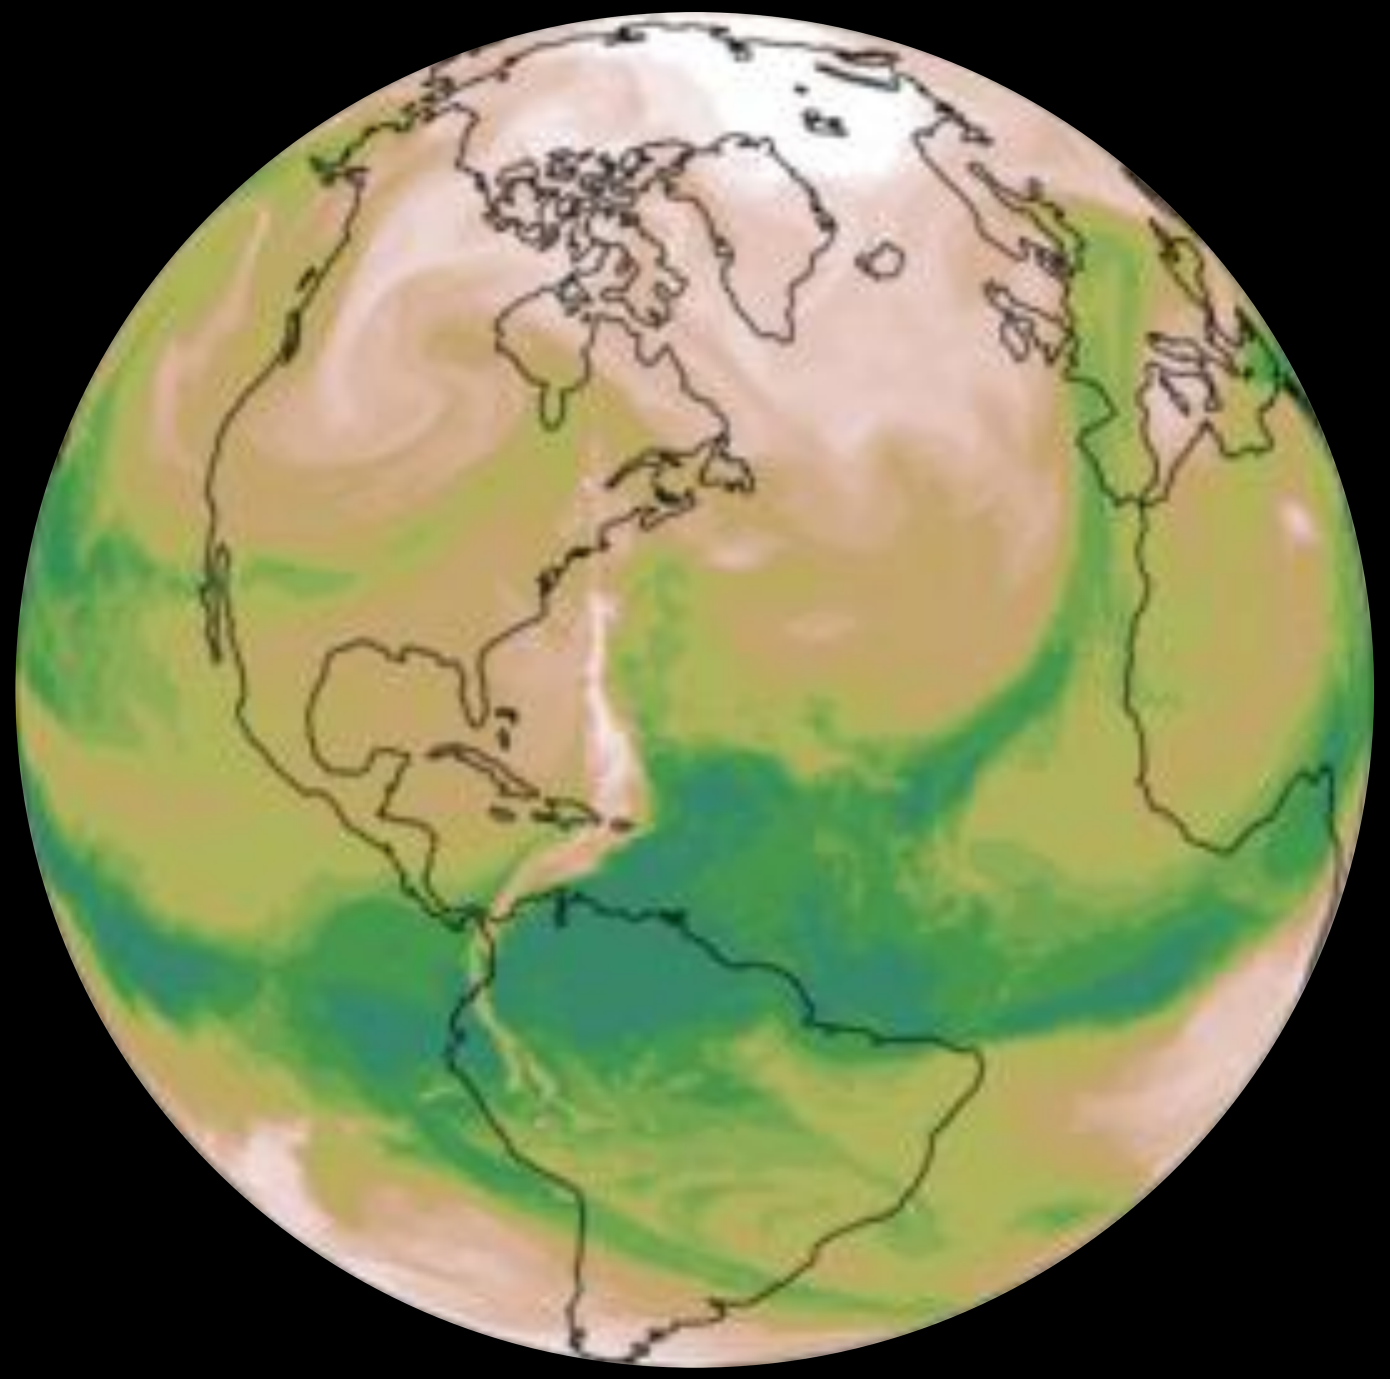

Total column water vapour field at last time step of the forecast

FourCastNet: Formulating Fourier Neural Operators on the Sphere

Accounting for Geometry: Equivariant treatment of spherical geometry overcomes instabilities

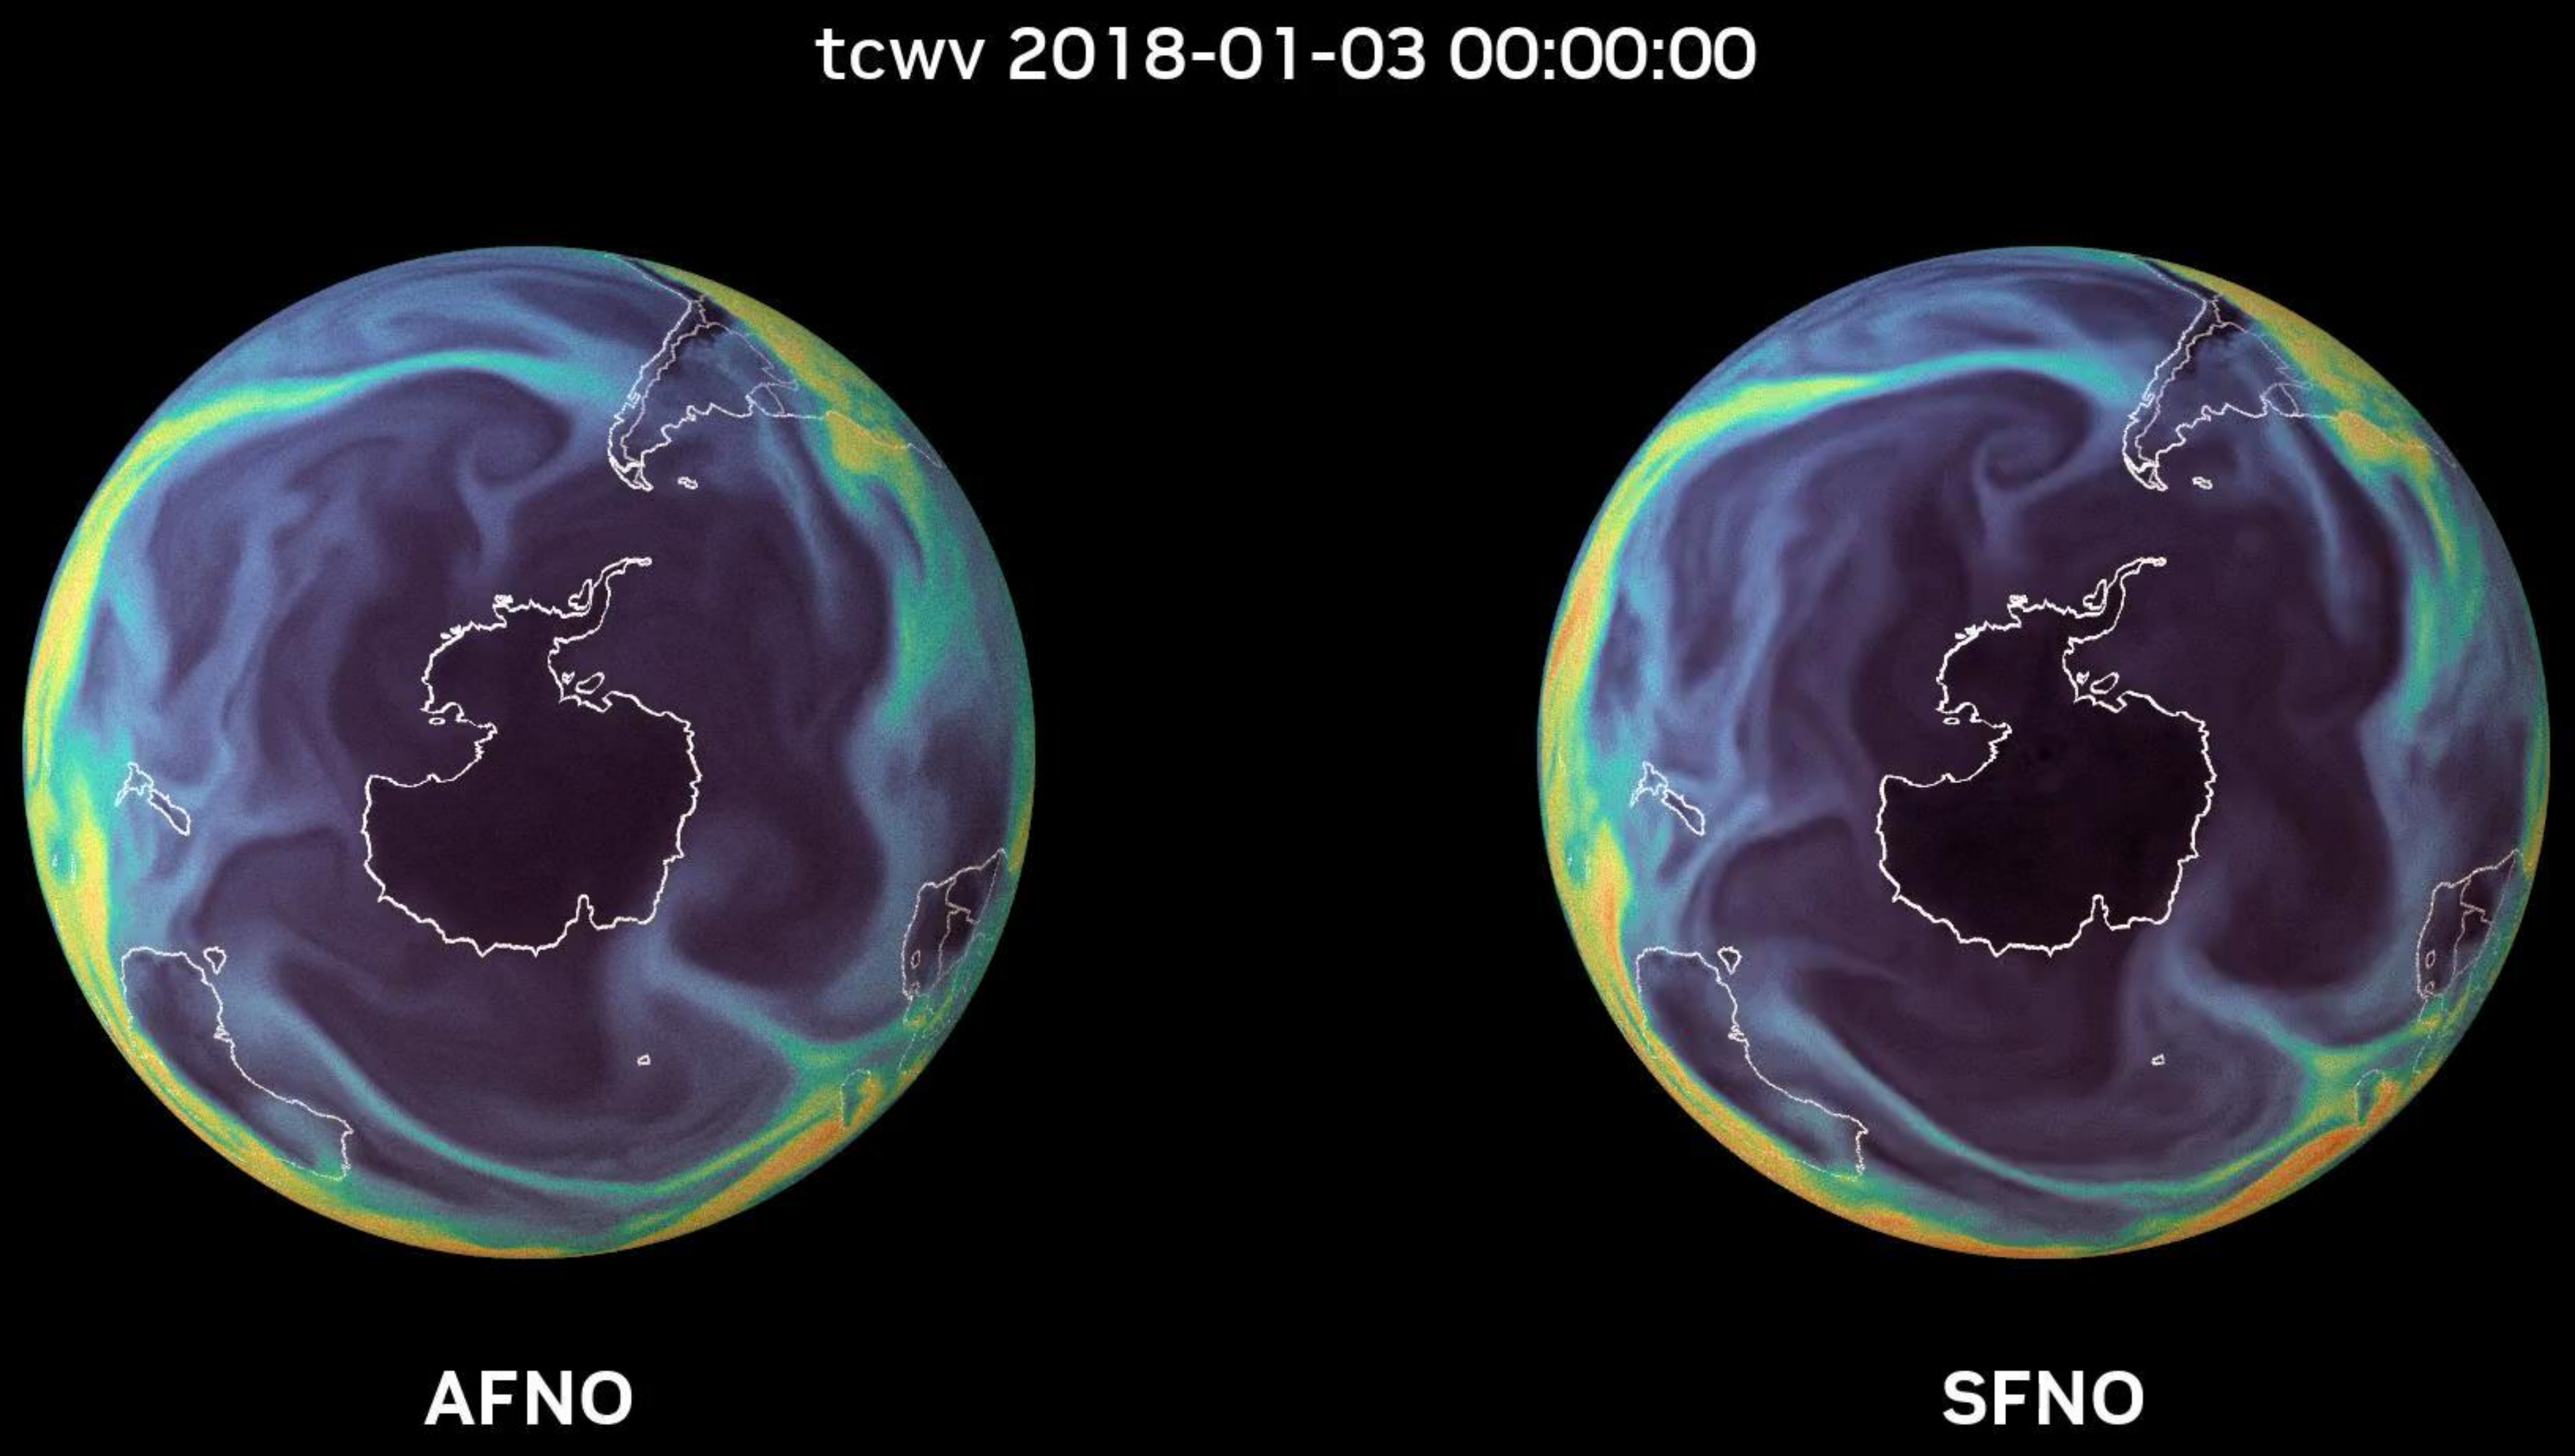

Visualization by Boris Bonev

harmonics

- Open-Source library under MIT license: https://github.com/NVIDIA/torch-harmonics
- Efficient calls for forward and inverse spherical harmonic transformations
- Autograd support as differential layers in PyTorch
- Support for distributed computation across several GPUs

FourCastNet v3

Bonev et al (2025) - A geometric approach to probabilistic machine-learning weather forecasting at scale

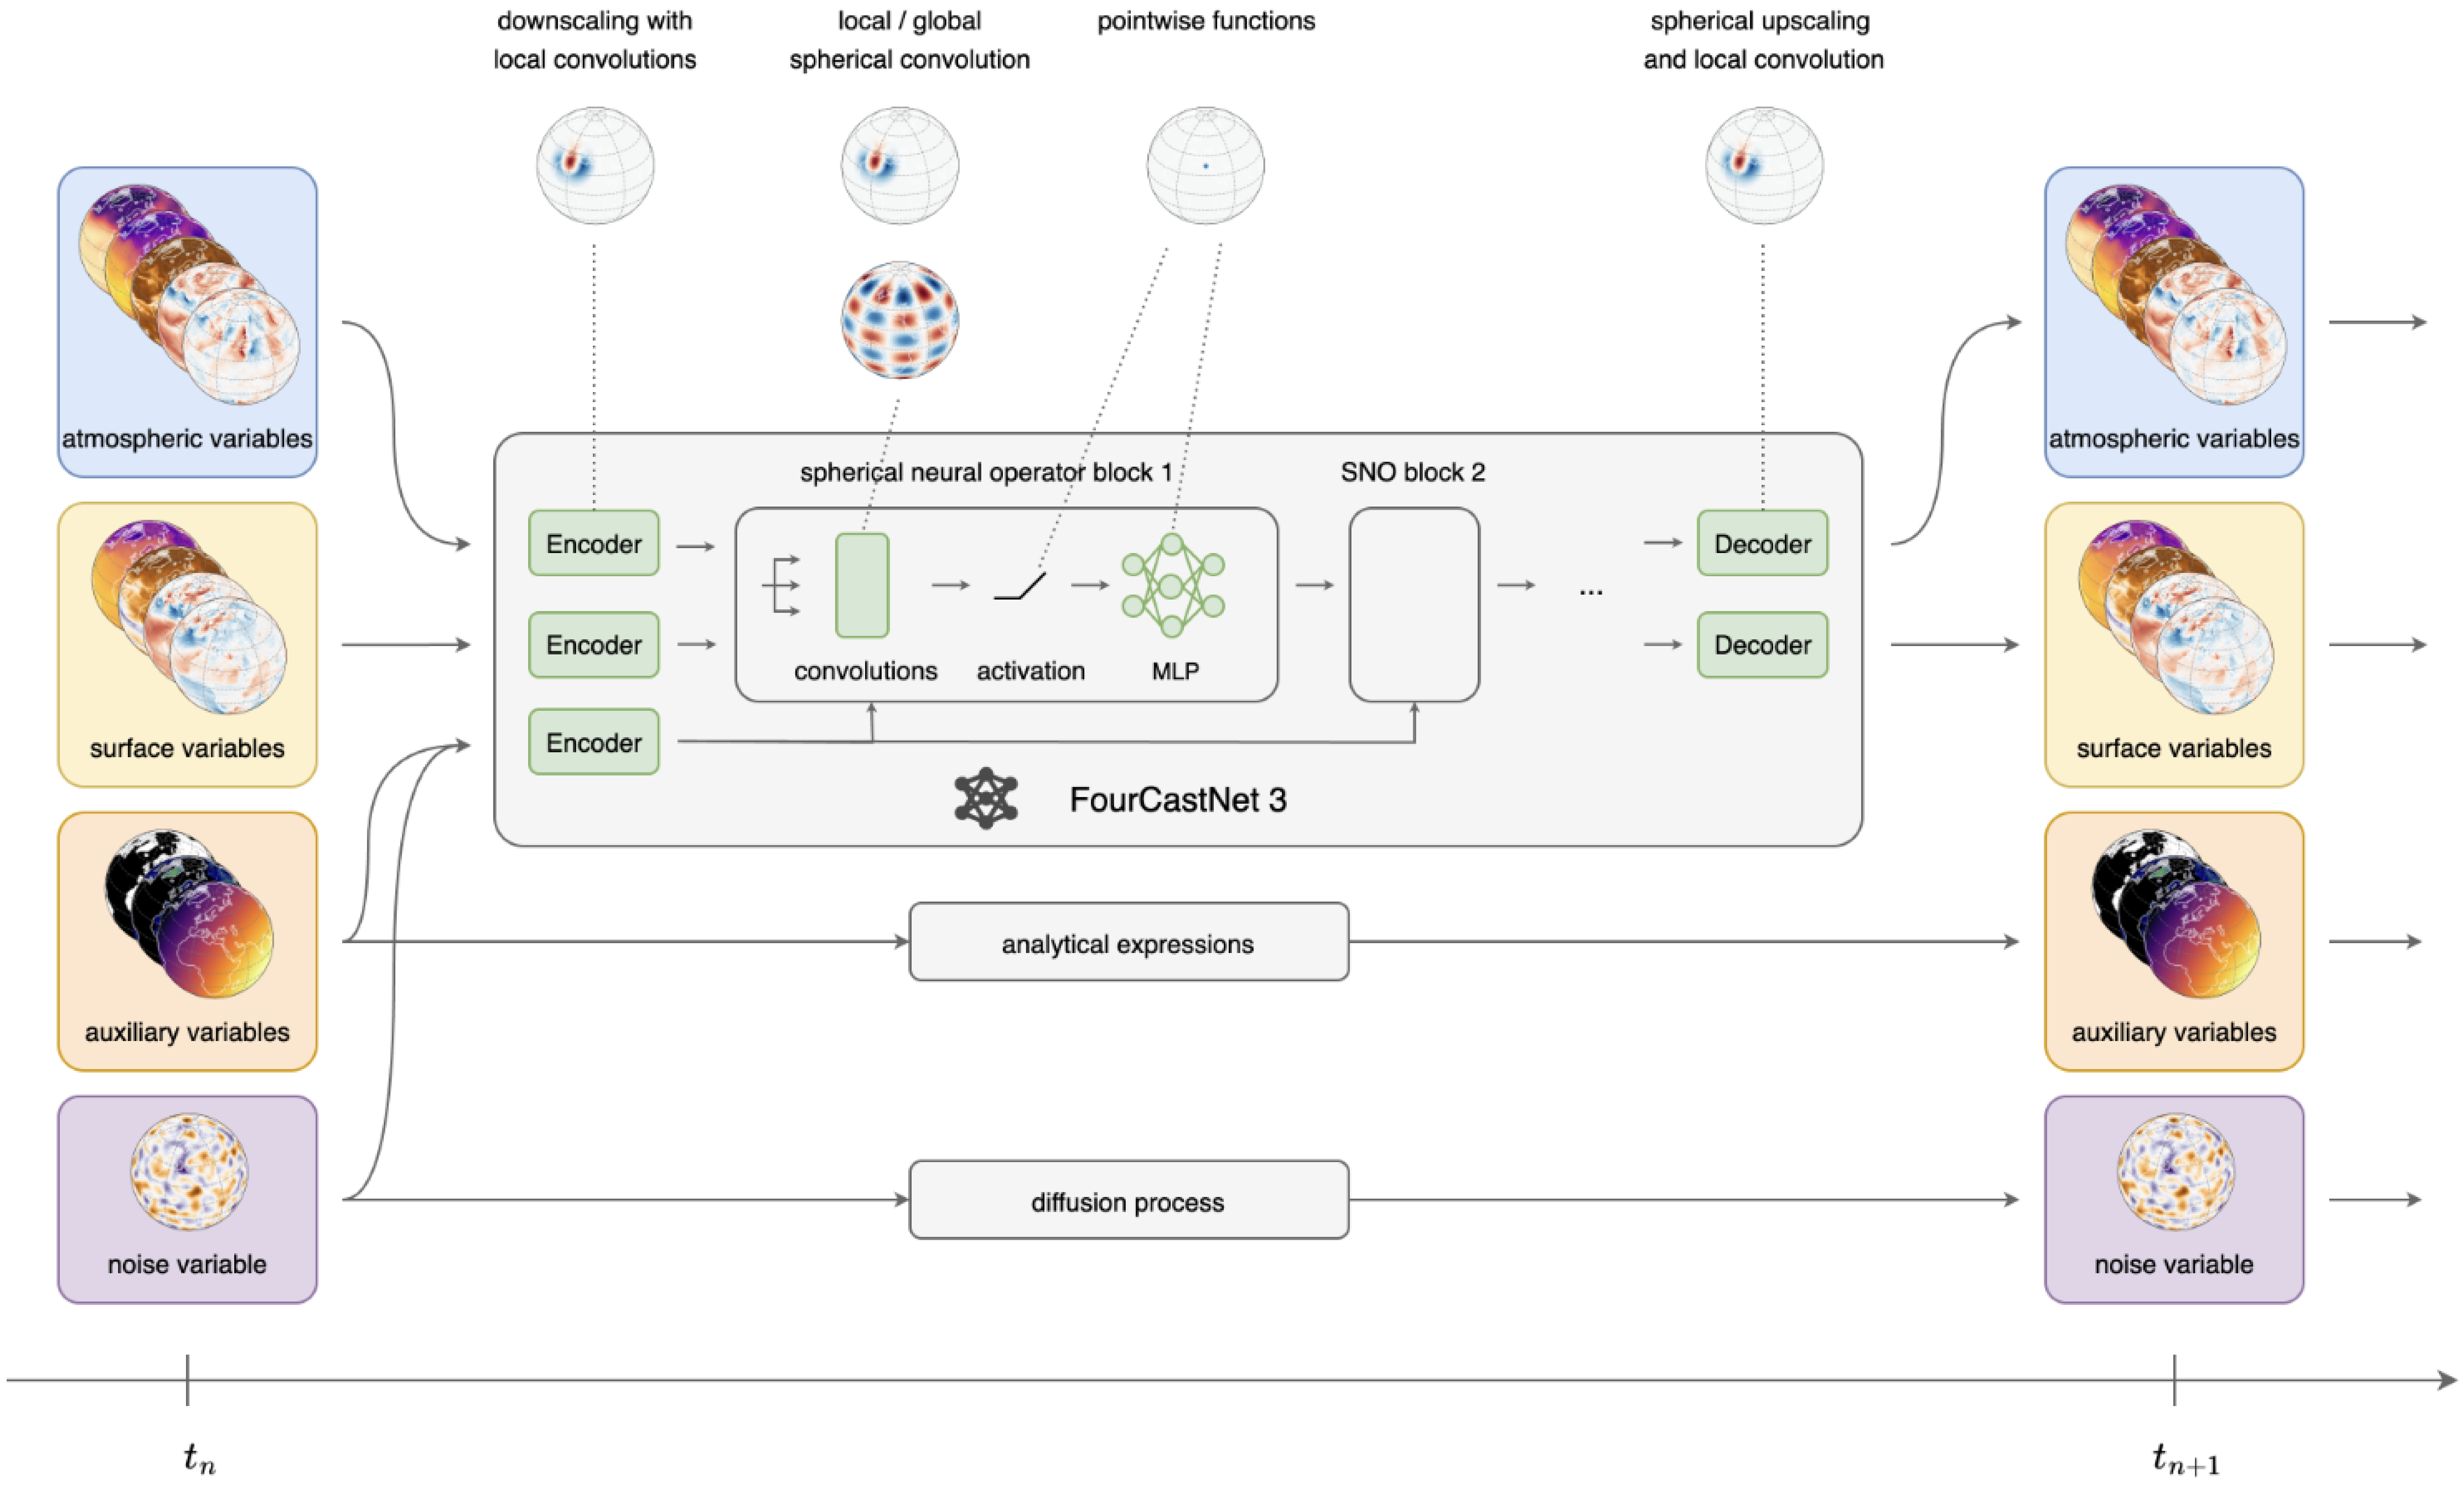

FourCastNet v3

Bonev et al (2025) - A geometric approach to probabilistic machine-learning weather forecasting at scale

Combines Local and Global Convolutions to capture large scales efficiently while capturing small scale processes

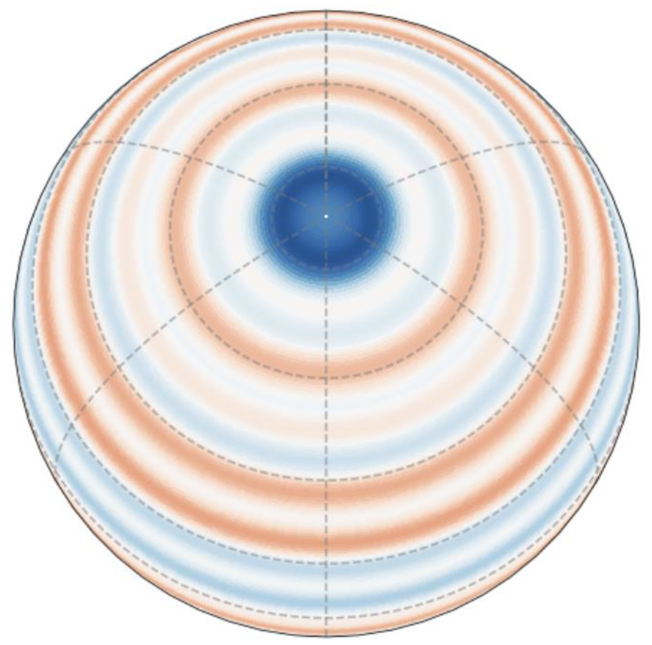

(a) spectral convolution filter

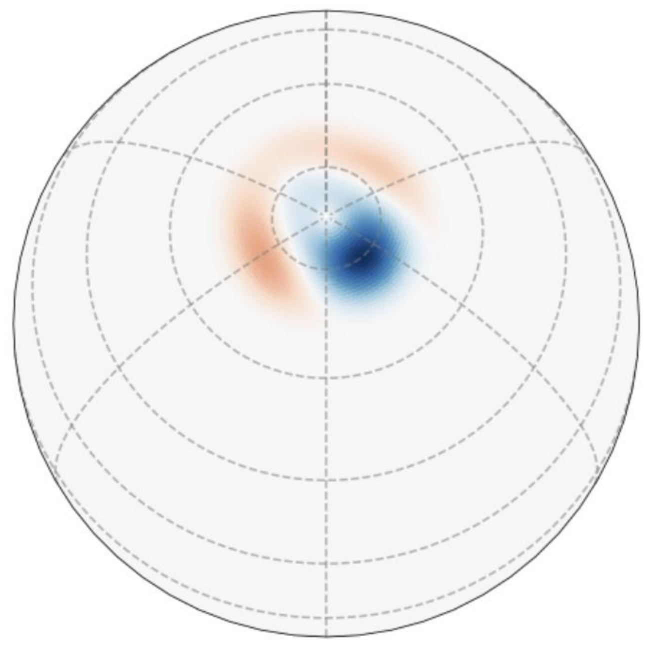

(b) local convolution filter

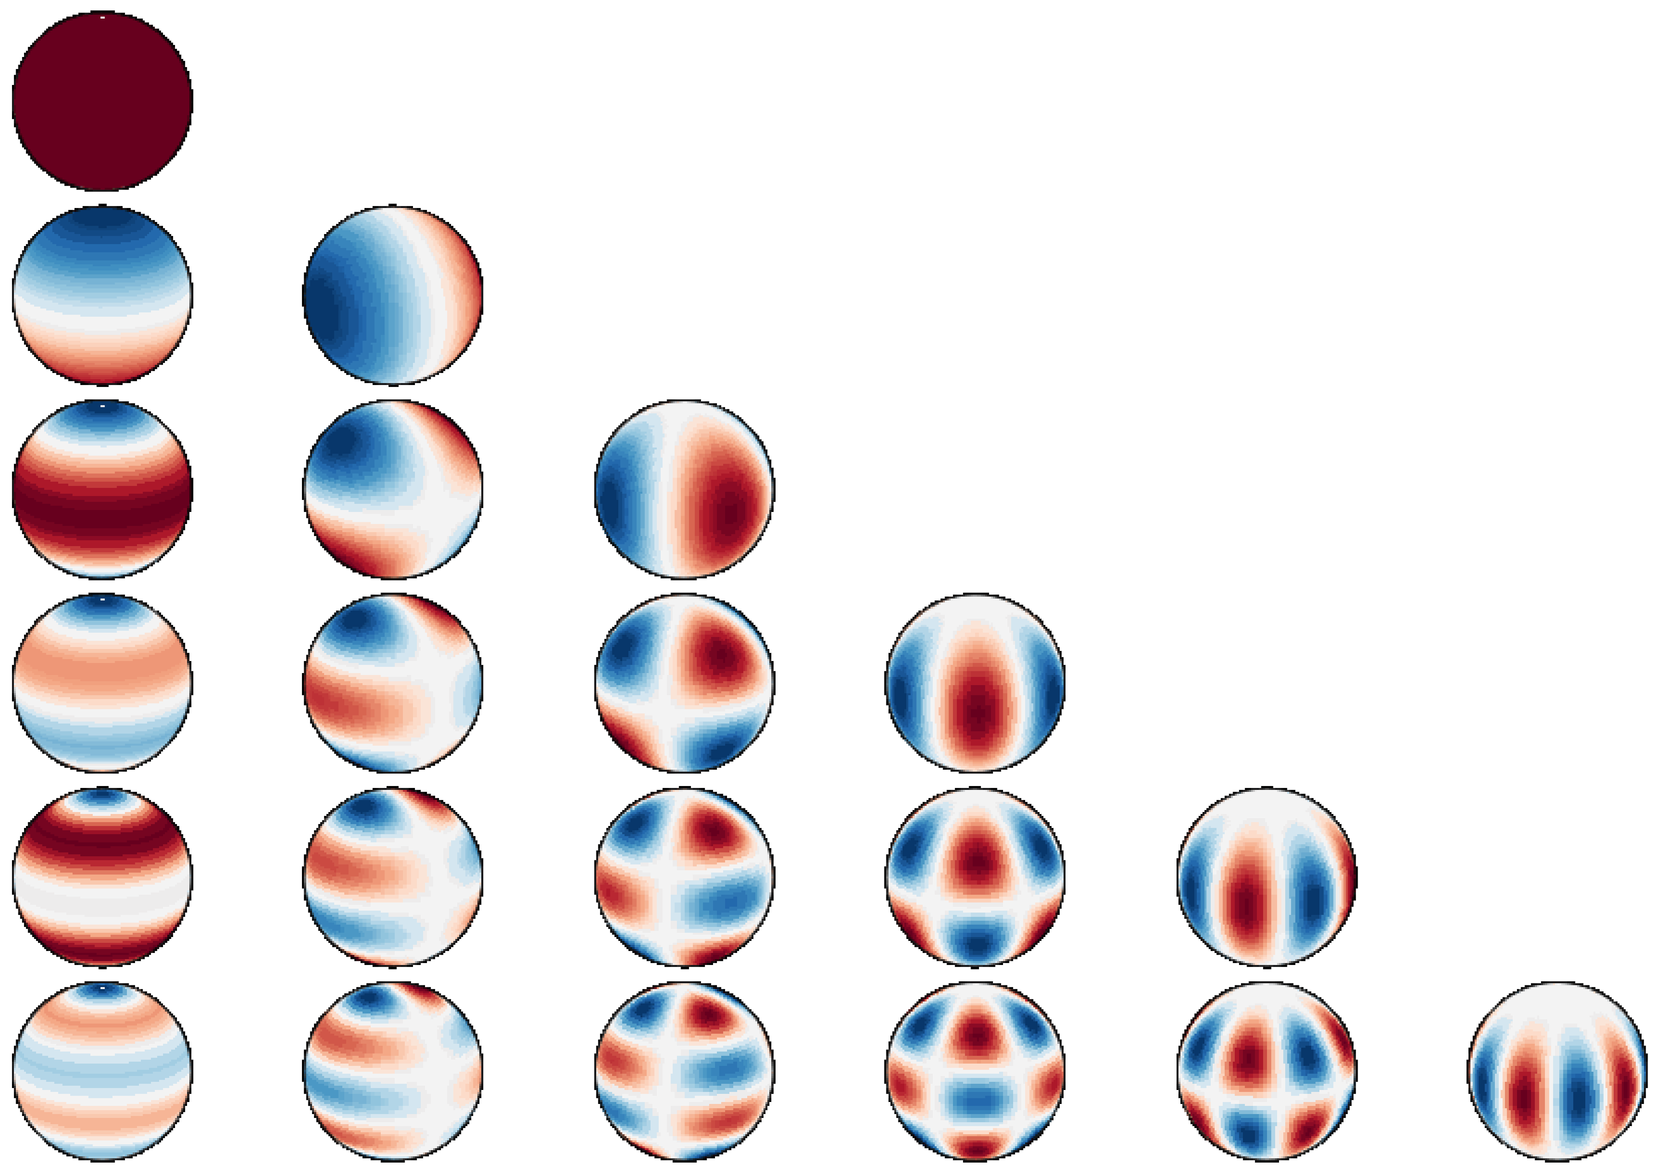

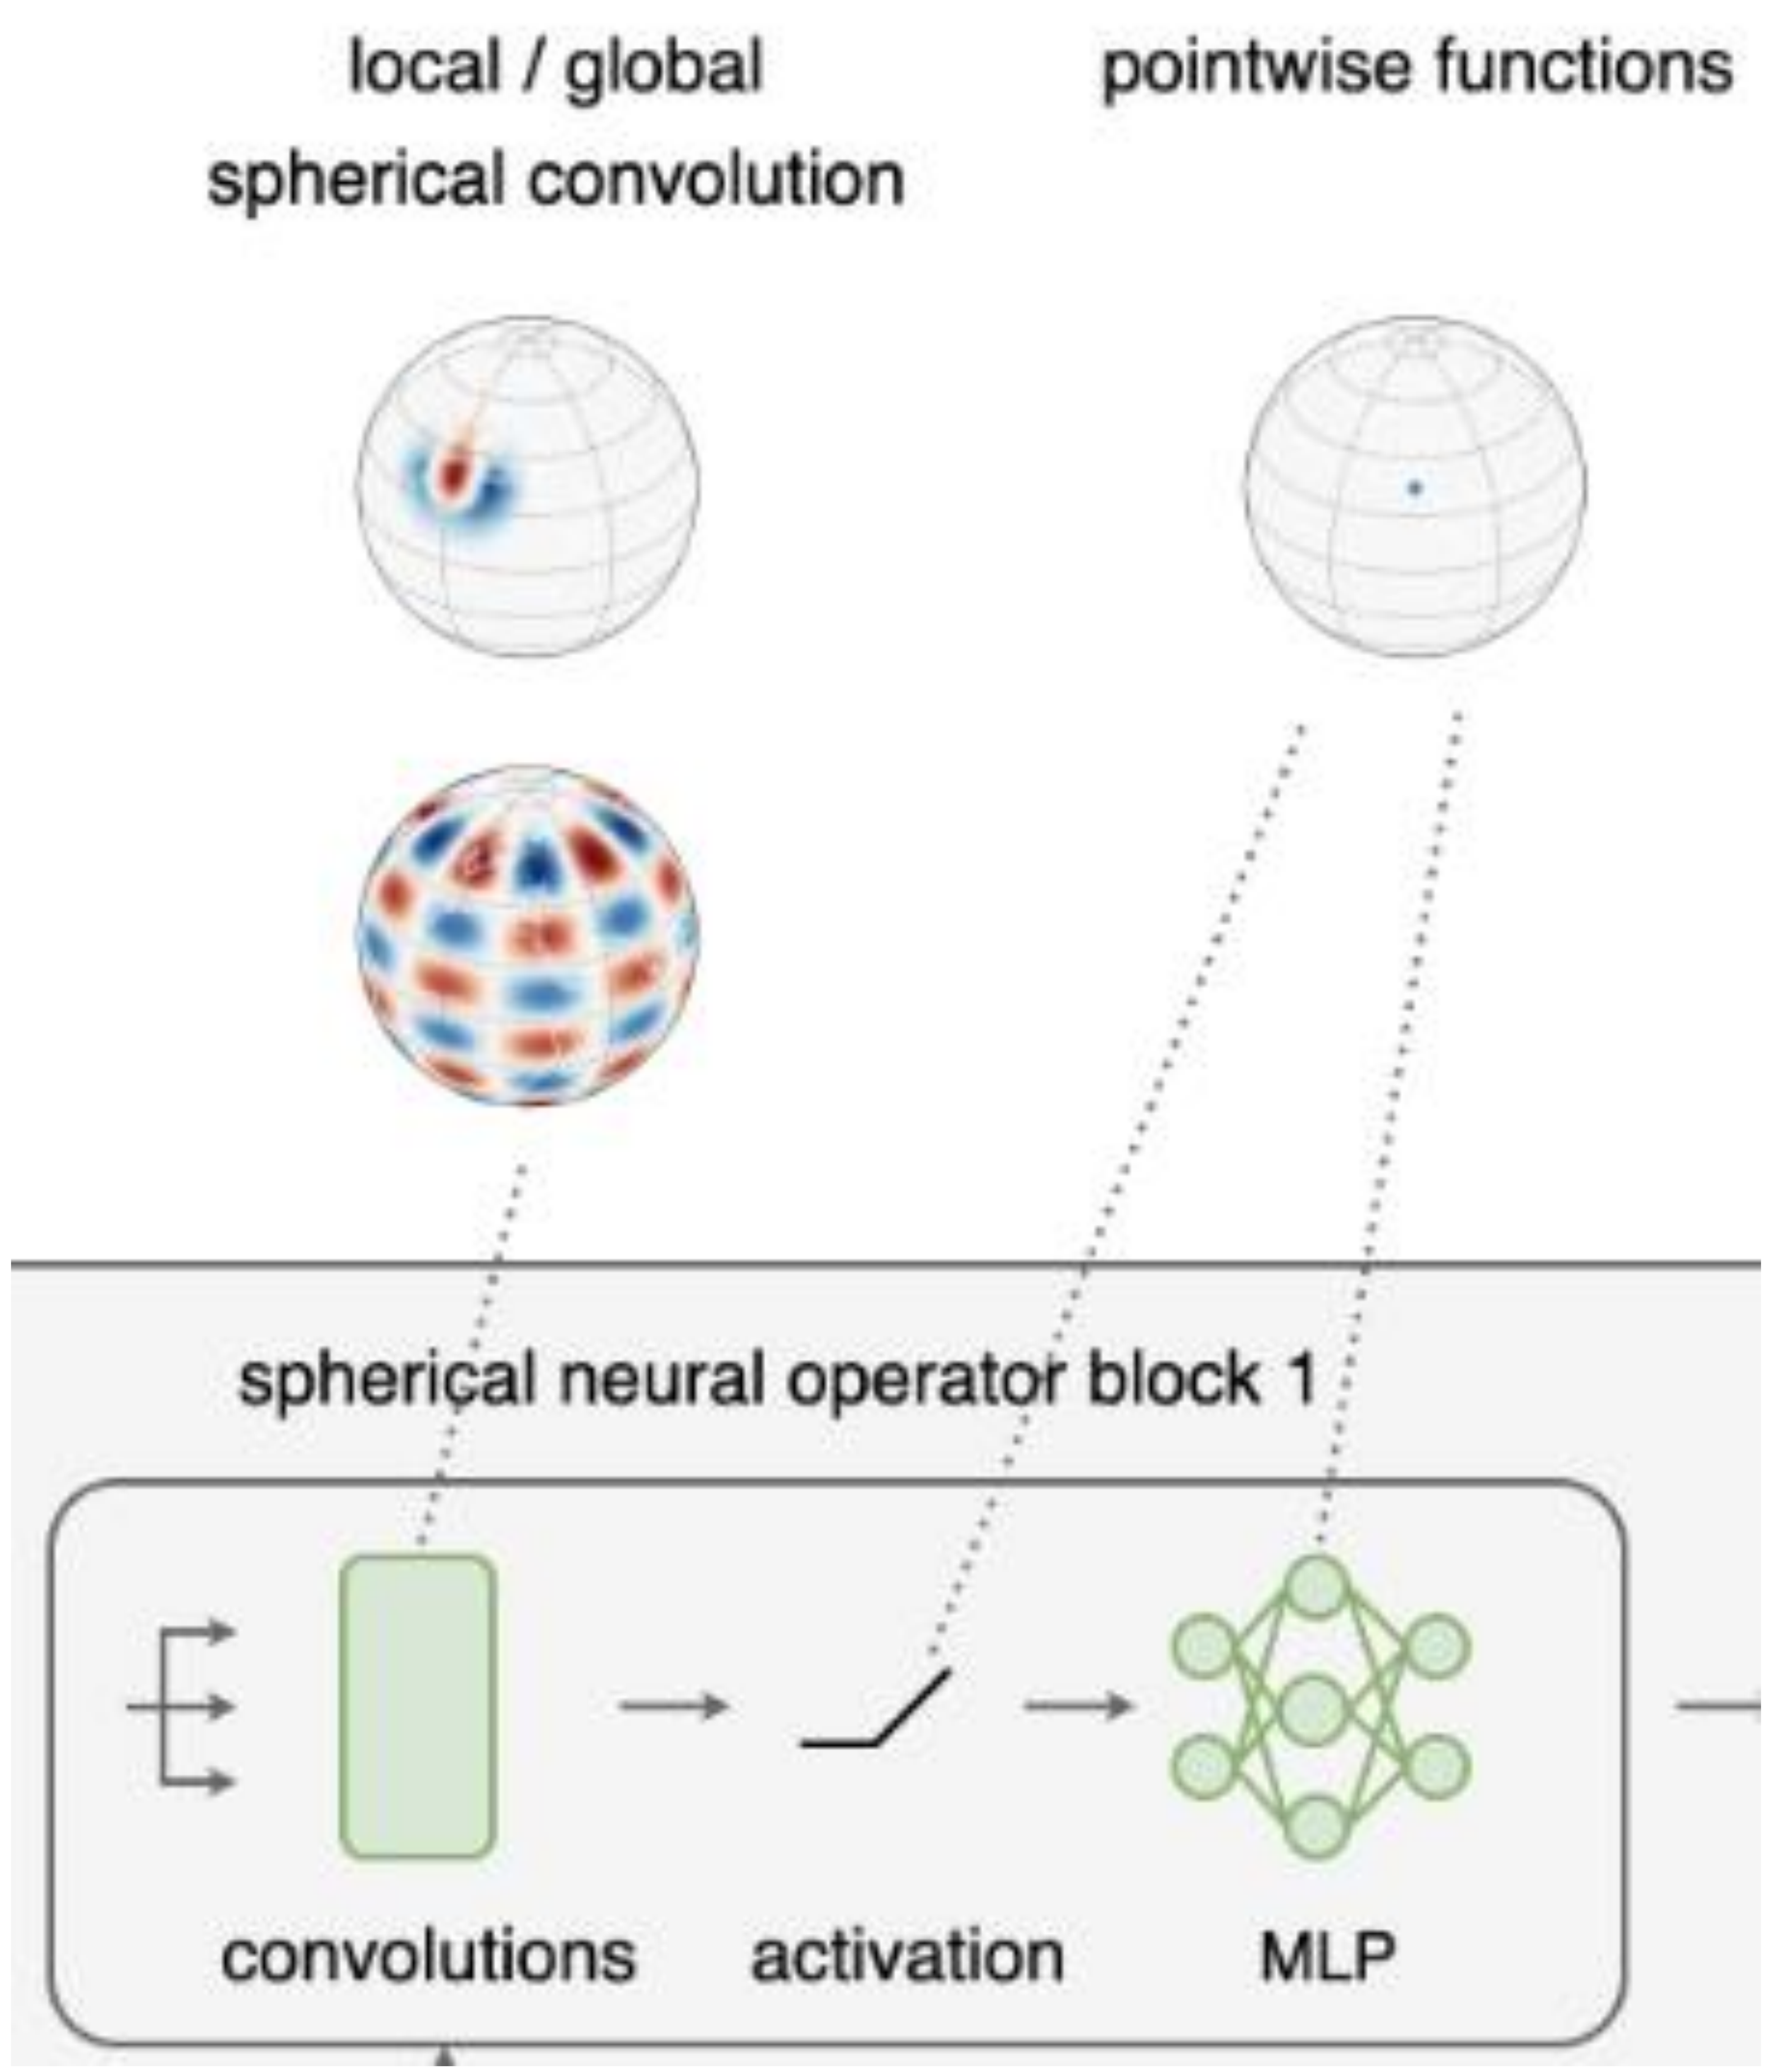

FourCastNet v3

Bonev et al (2025) - A geometric approach to probabilistic machine-learning weather forecasting at scale

Domain parallelism to scale to 1024 GPUS

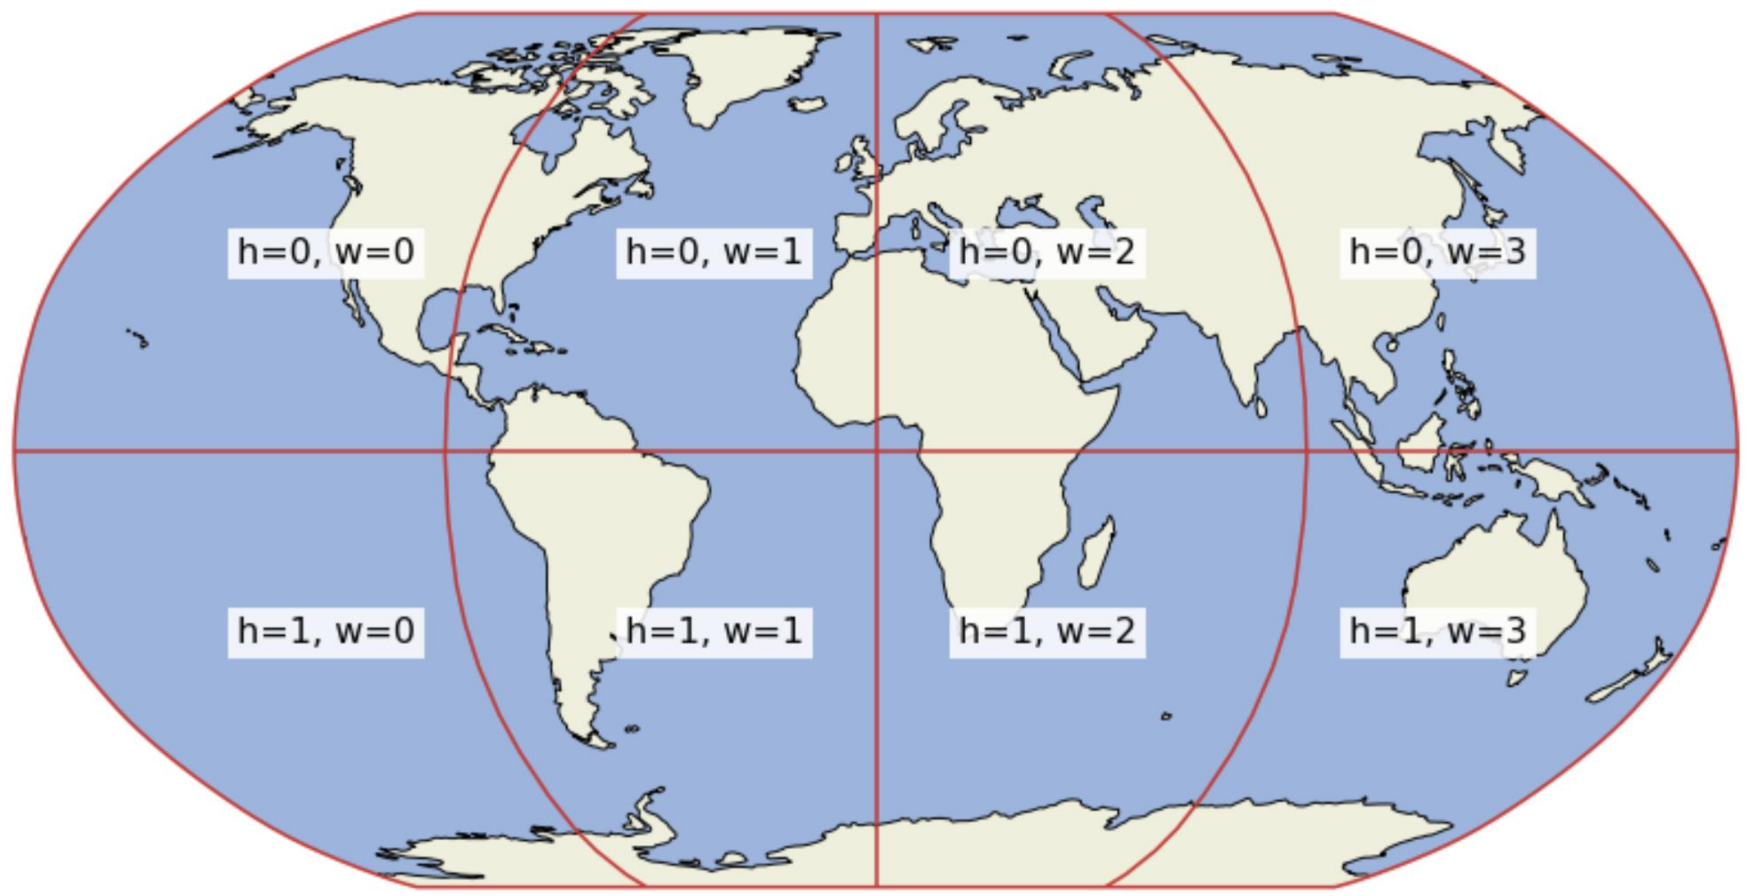

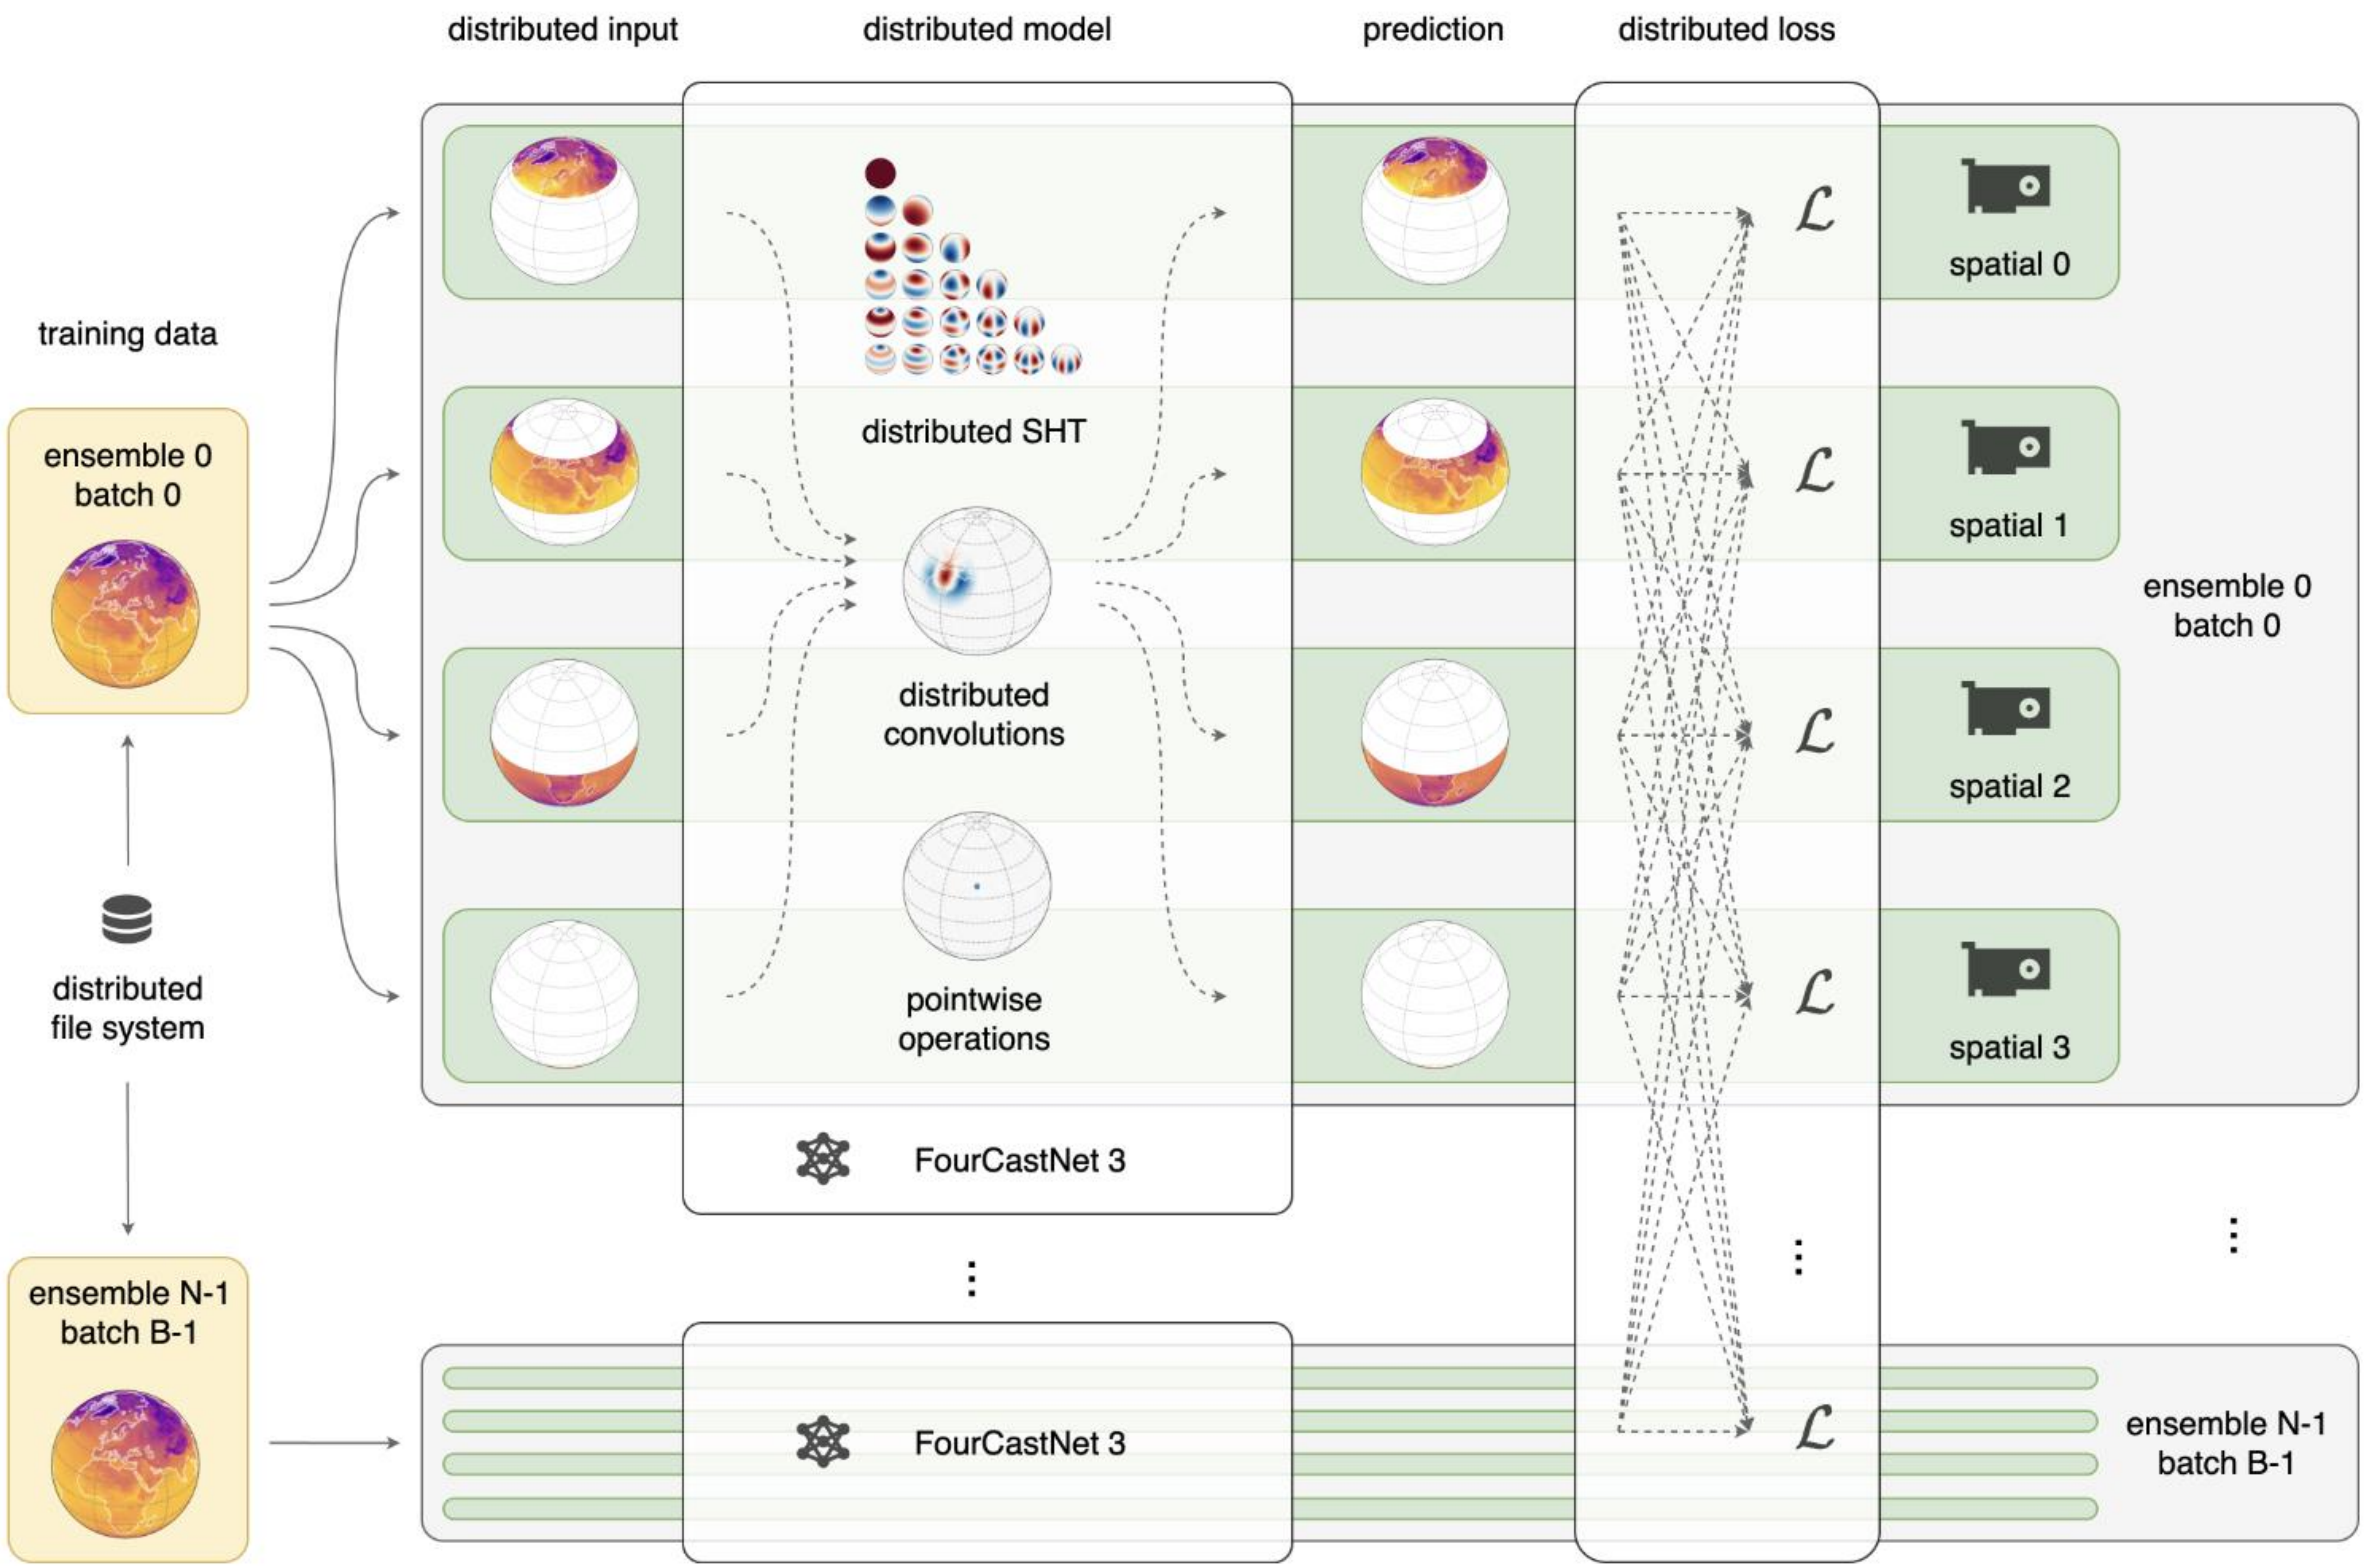

FourCastNet v3

Bonev et al (2025) - A geometric approach to probabilistic machine-learning weather forecasting at scale

Training Objective CRPS

$$\mathrm{CRPS}(F_{\mathrm{ens}}, u^*) = \frac{1}{N_{\mathrm{ens}}} \sum_{e=1}^{N_{\mathrm{ens}}} |u_e - u^*| - \frac{1}{2N_{\mathrm{ens}}^2} \sum_{e=1}^{N_{\mathrm{ens}}} \sum_{i=1}^{N_{\mathrm{ens}}} |u_e - u_i|$$

Composite Training Objective:

$L = \mathrm{CRPS}_{\mathrm{spatial}} + \mathrm{CRPS}_{\mathrm{spectral}}$

Retains realistic spectra, even at extended lead times of up to 60 days

60-day global forecast at 0.25°, 6-hourly resolution in under 4 minute, 60x faster than diffusion

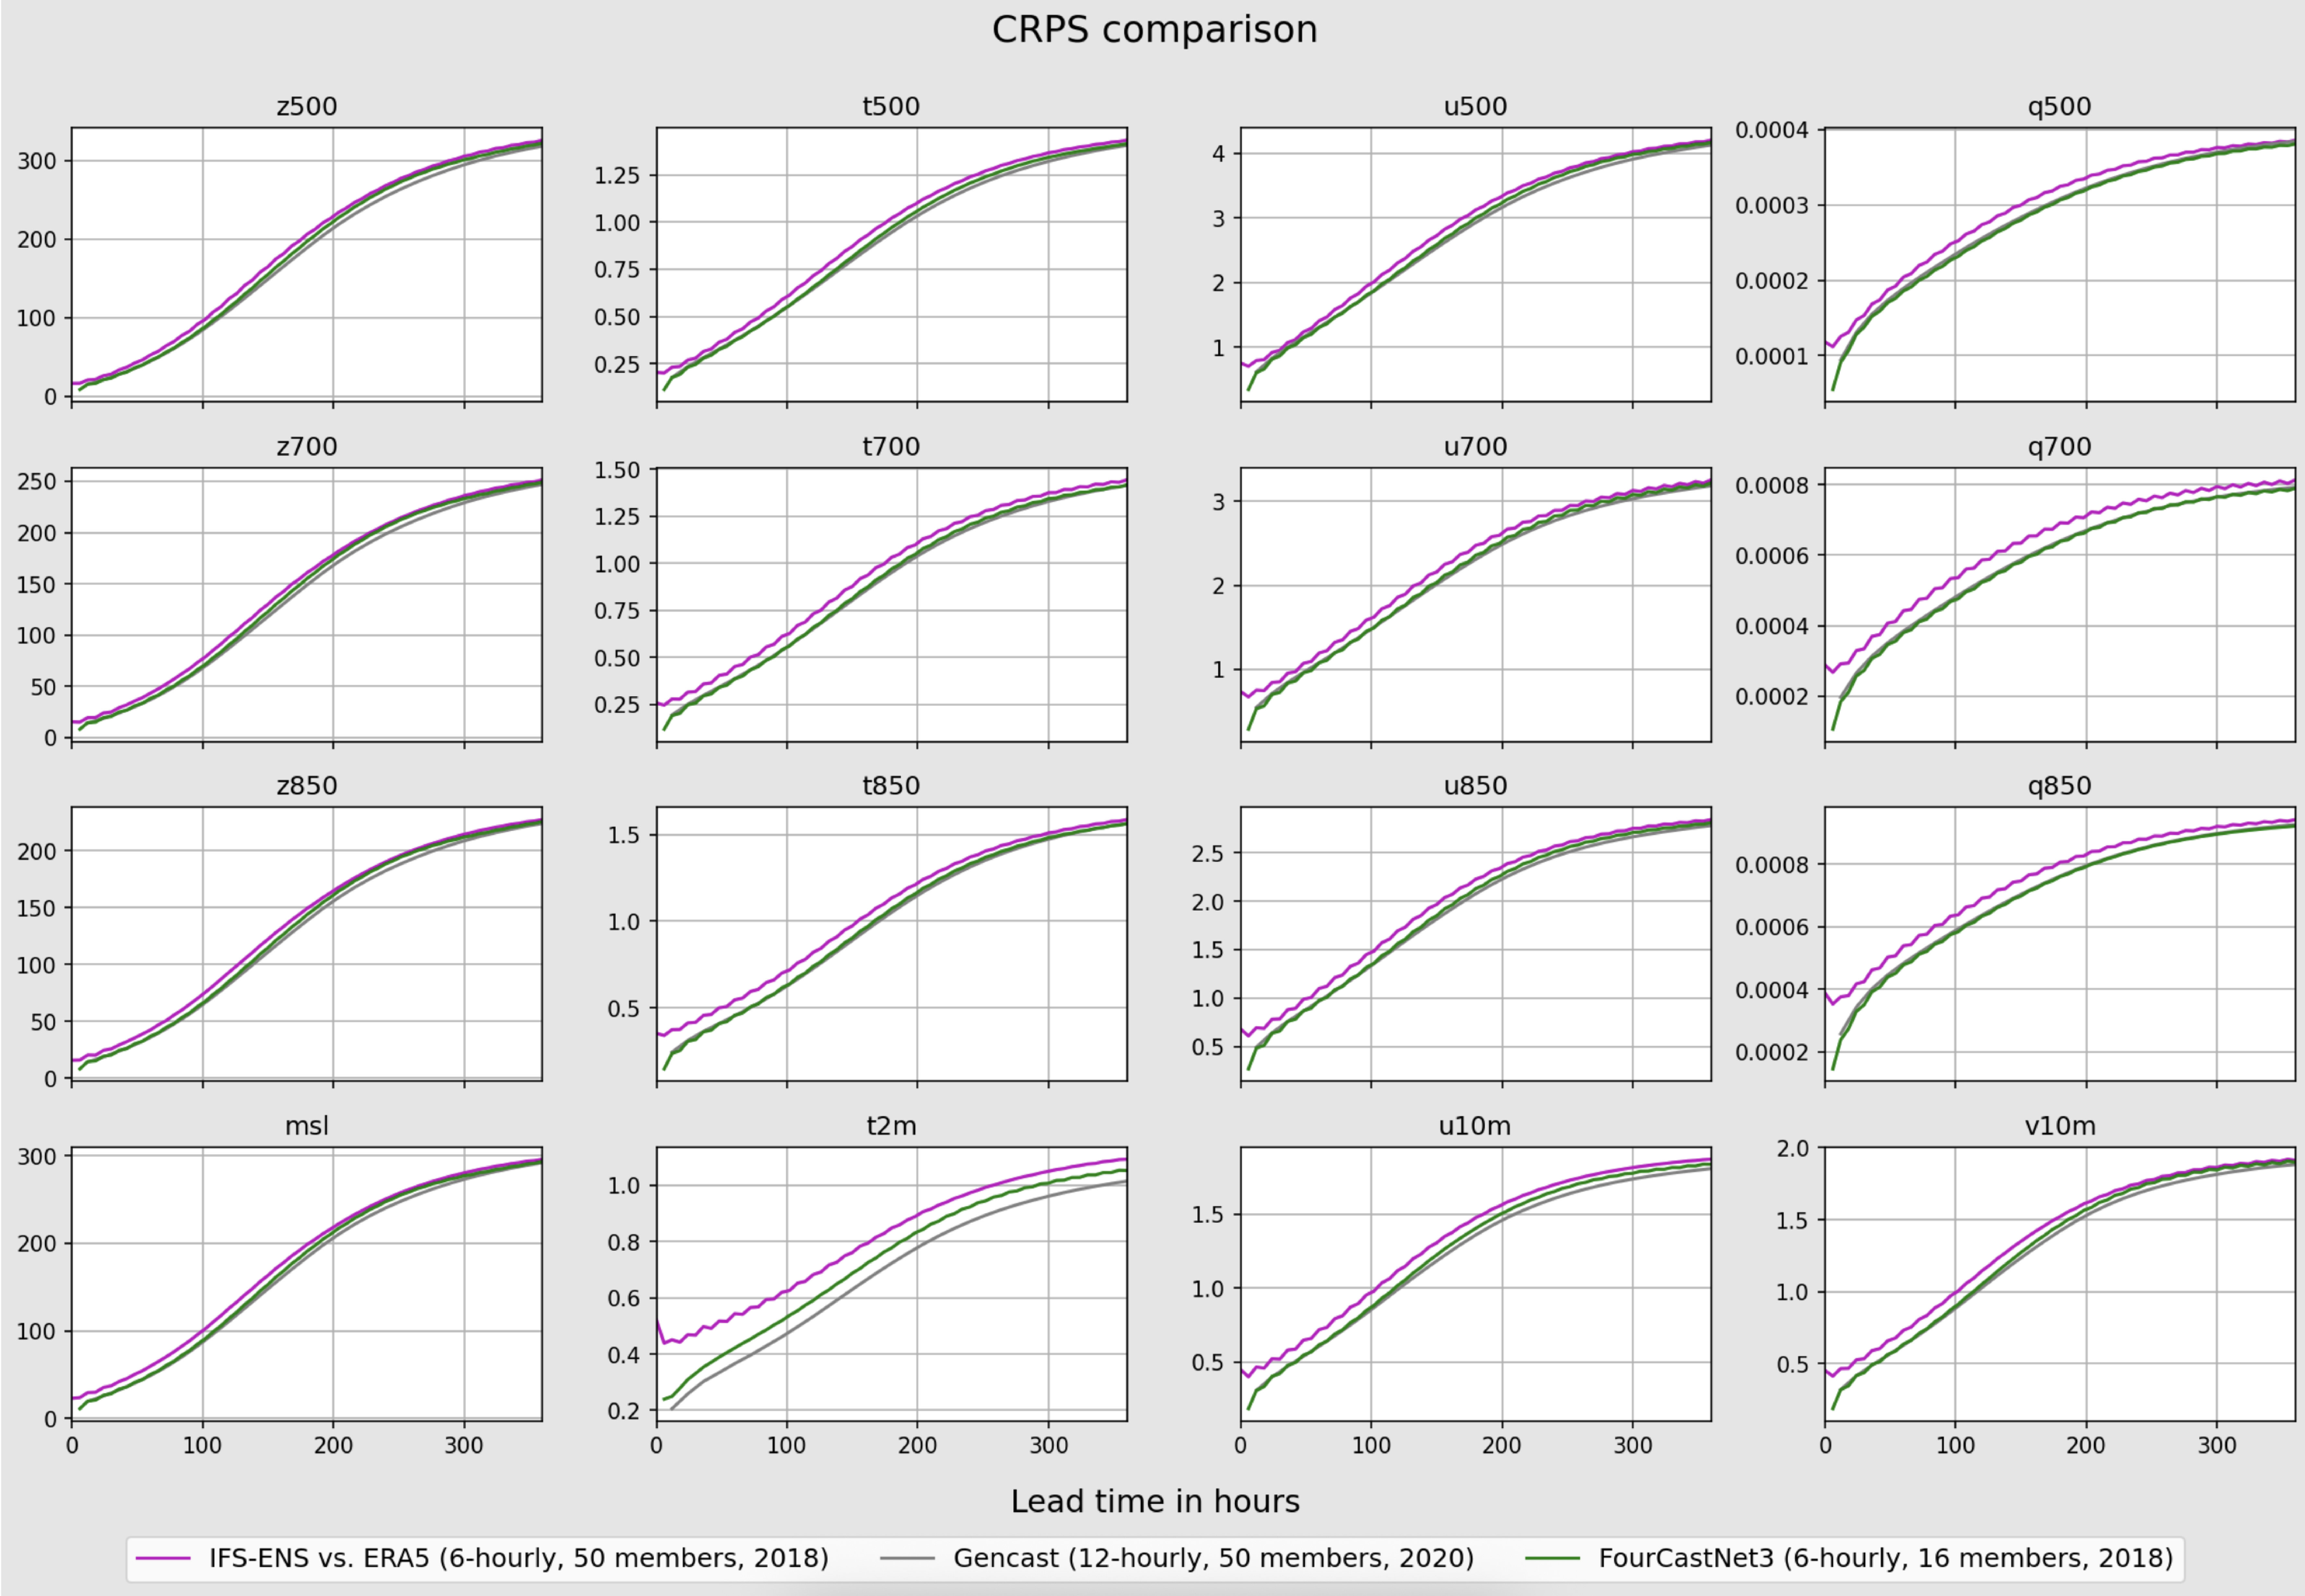

FourCastNet v3

Bonev et al (2025) - A geometric approach to probabilistic machine-learning weather forecasting at scale

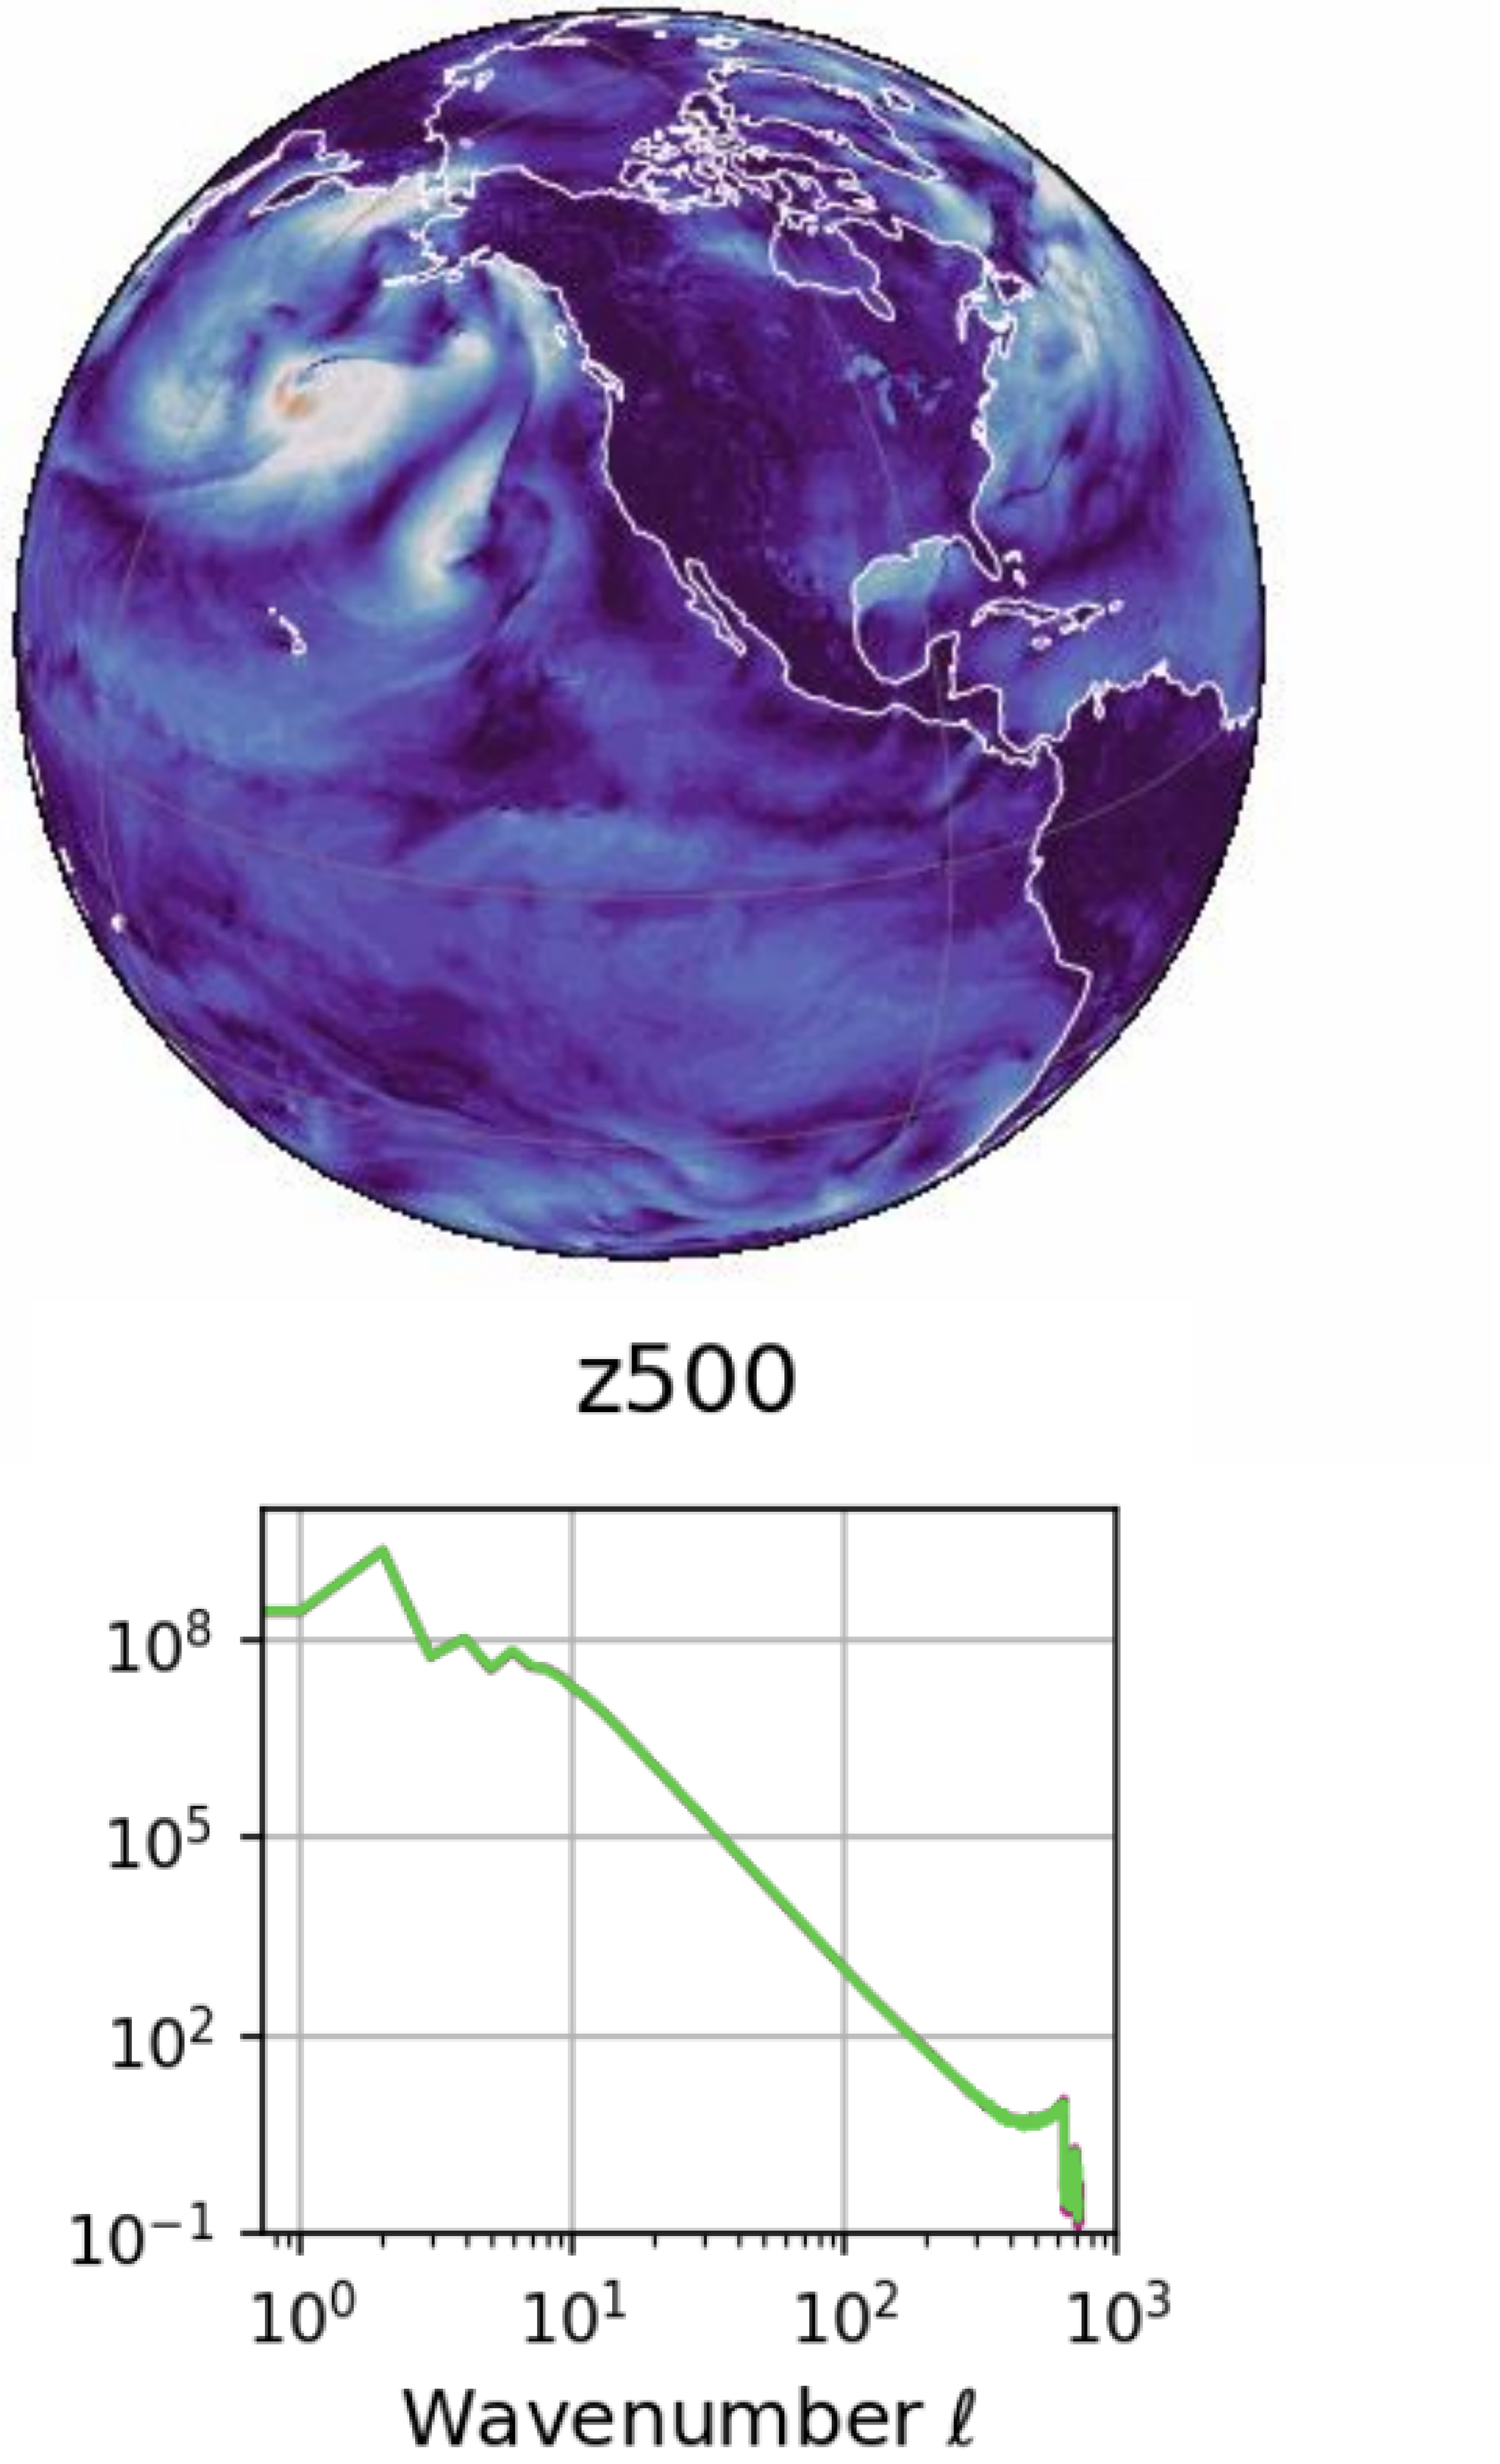

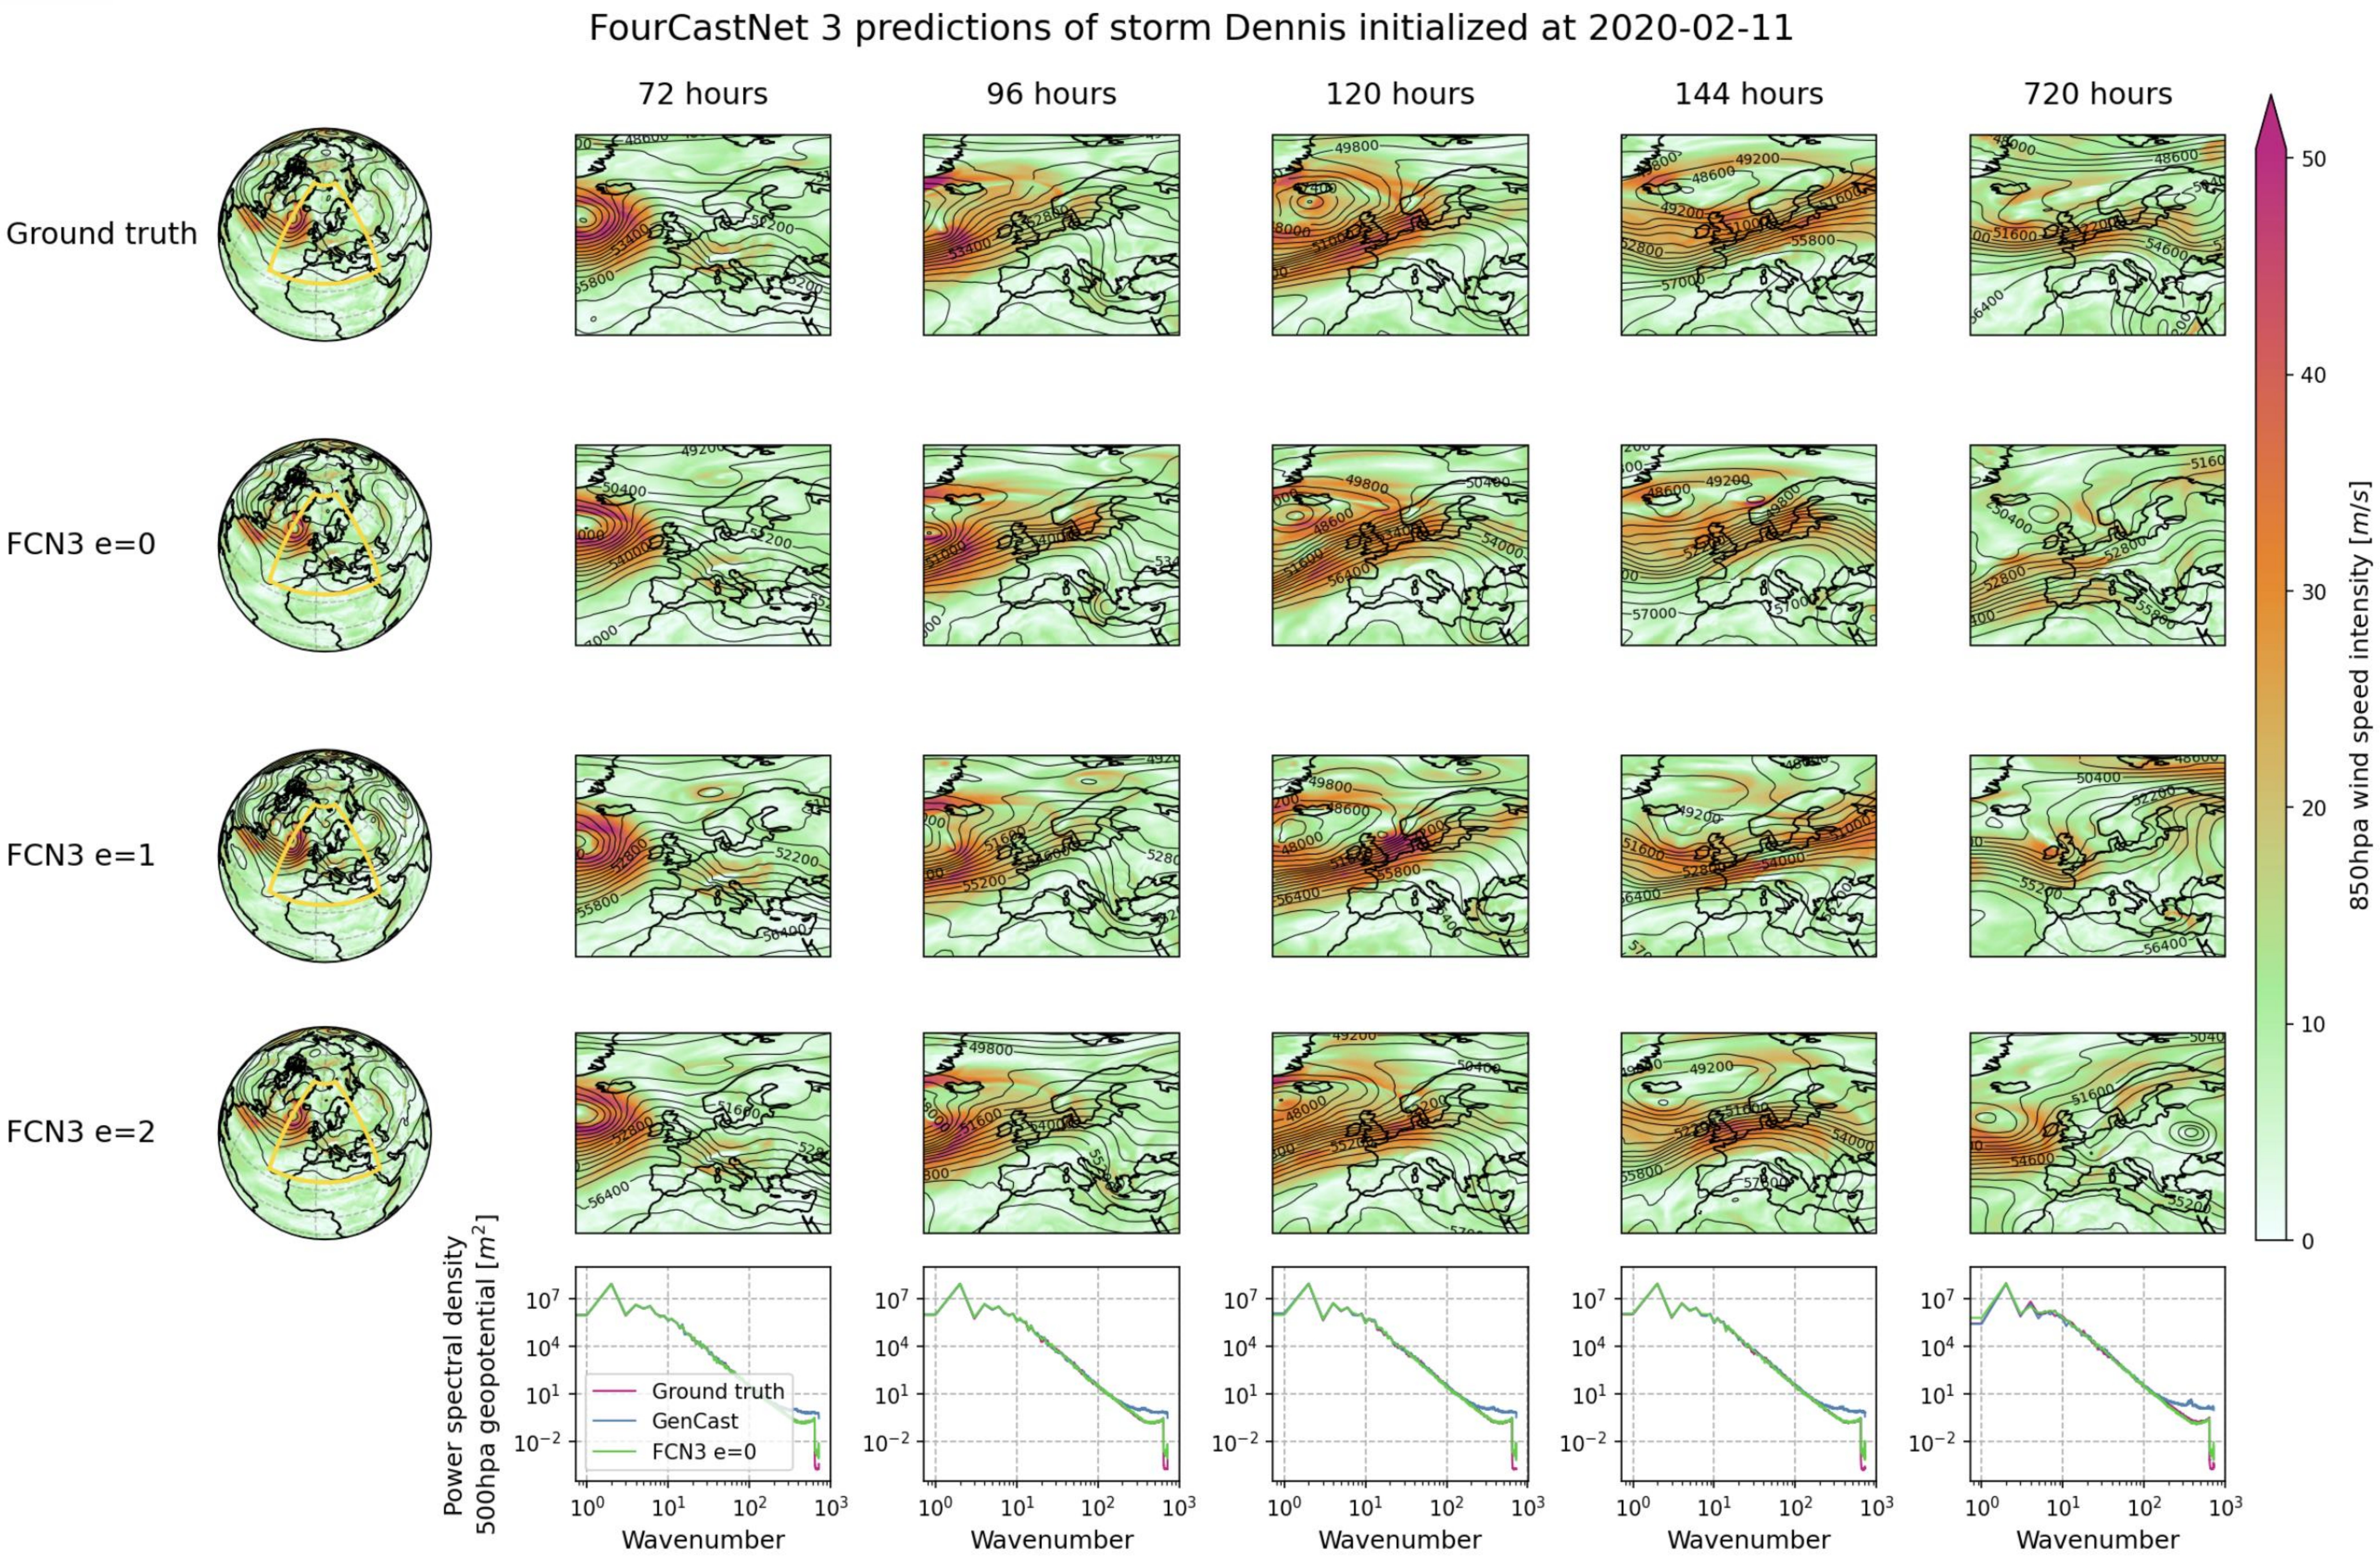

Huge Ensembles for Accurate Statistics

Probabilistic applications benefit from large event sets

- Modern numerical systems produce 50-100 ensemble members
- AI pipelines can be easily **scaled to 10,000+ ensemble members**
  - Improve statistics on extreme weather events like hurricanes
  - Predict extremely rare events like once-in-a-century rainfalls
  - Capture coincident extremes like high humidity during a heat wave
- **Long stable rollouts** enable creation of event sets for risk modeling

FourCastNet validated to produce ensembles for low-likelihood, high-impact extreme weather events

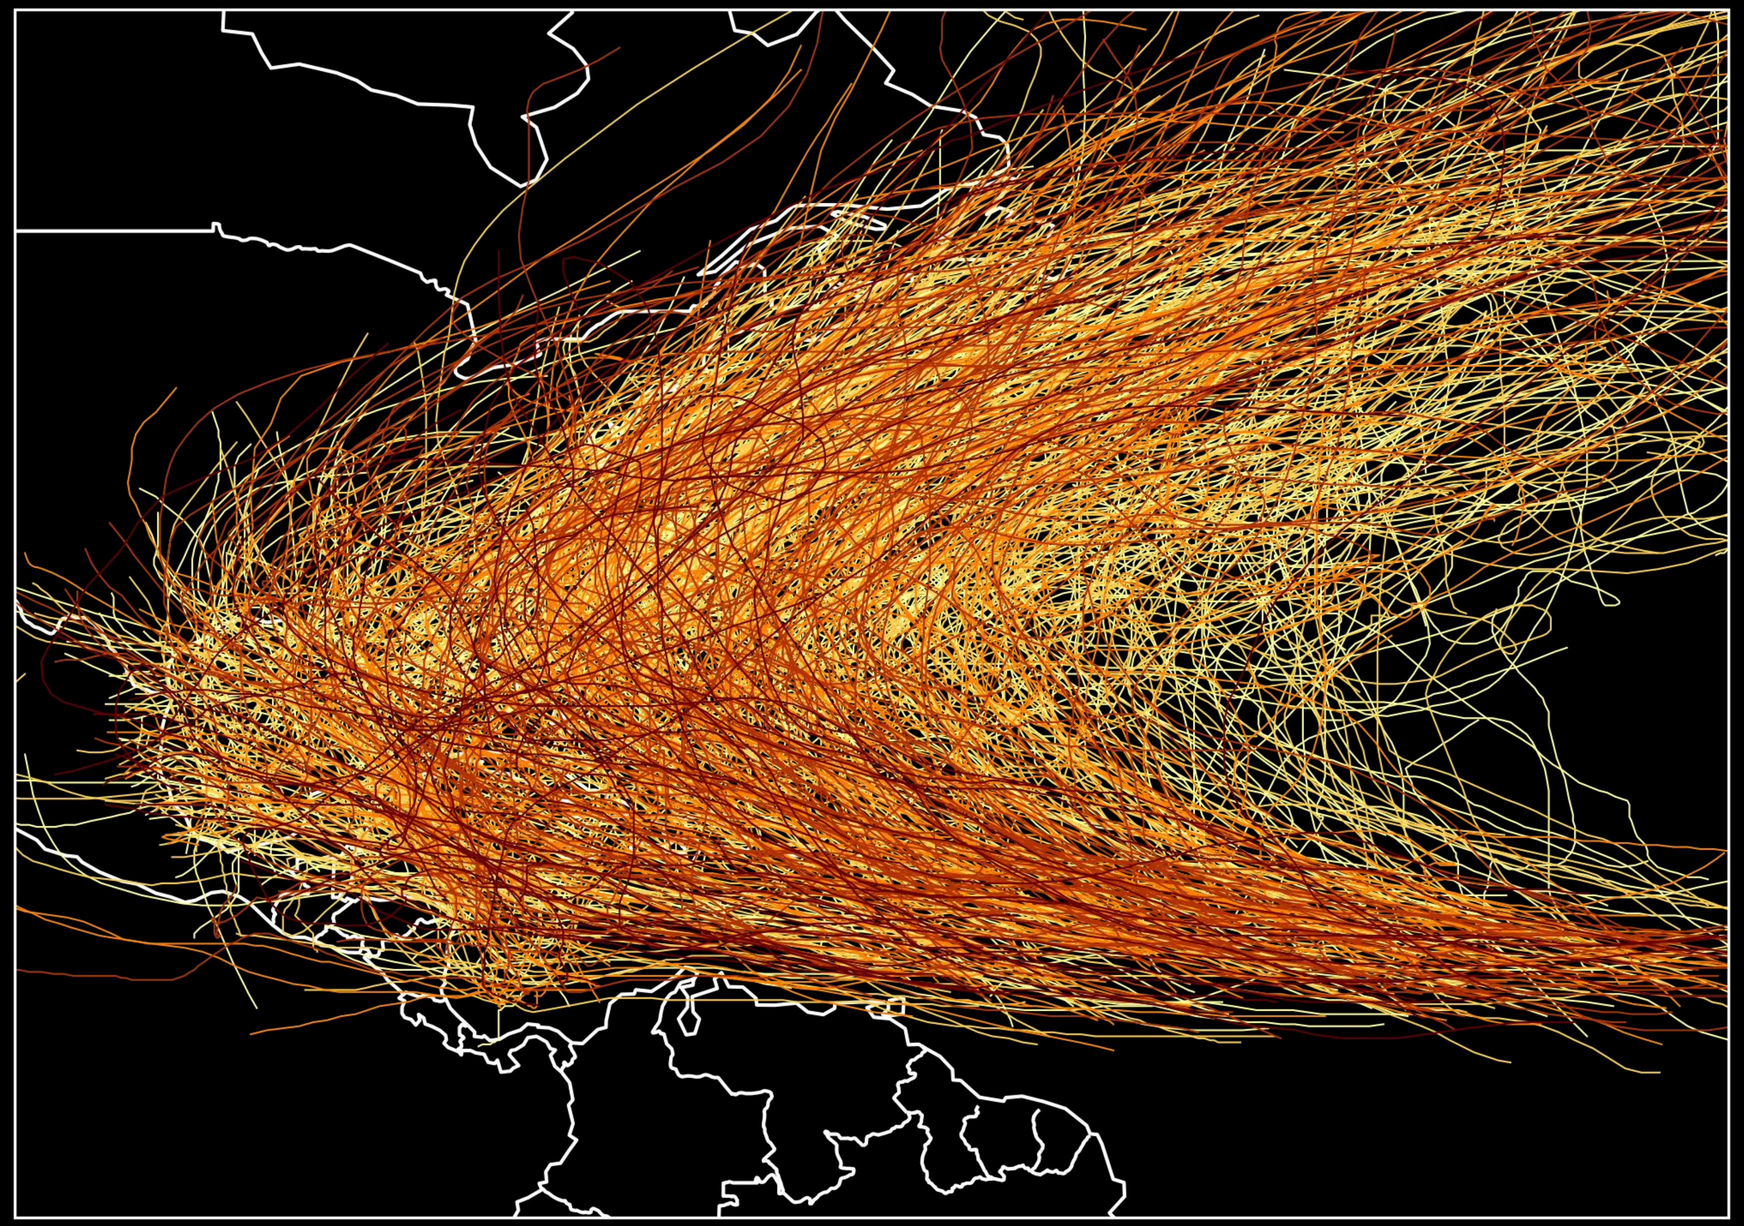

Storm tracks over the Northern Atlantic

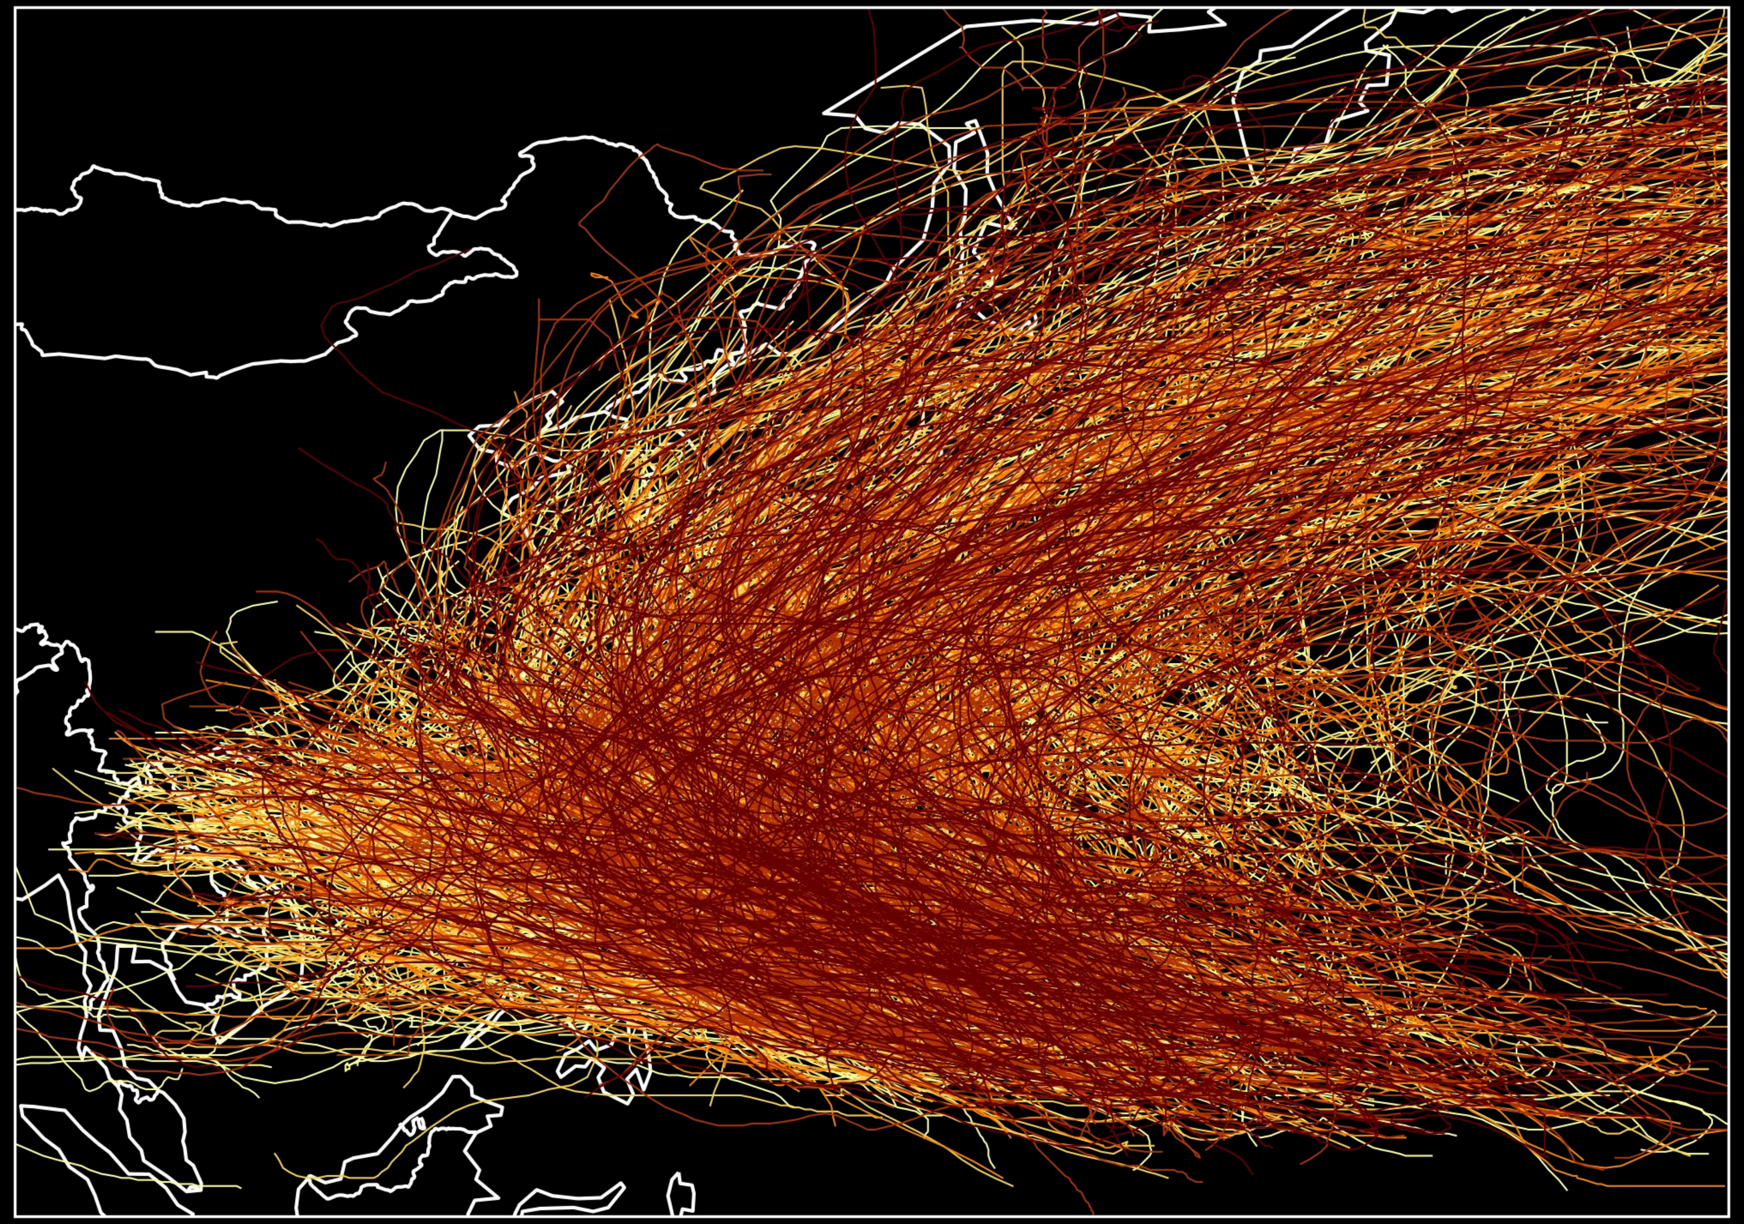

Storm tracks over the Western Pacific

Modelling Extremes

AI-Driven Hurricane Risk Assessment – Partnered with AXA

Modern numerical systems produce 50-100 ensemble members

AI pipelines can be easily scaled to 10,000+ ensemble members

Hurricanes represent low-probability but high-impact extreme weather events

AI enables the generation of ensemble forecasts for storm tracks, demonstrated on **Hurricane Helene**

This approach provides valuable insights into the frequency, intensity, and risk associated with rare natural disasters and predict extremely rare events like once-in-a-century events

Facilitates faster, more accurate, and proactive decision-making for disaster preparedness and response

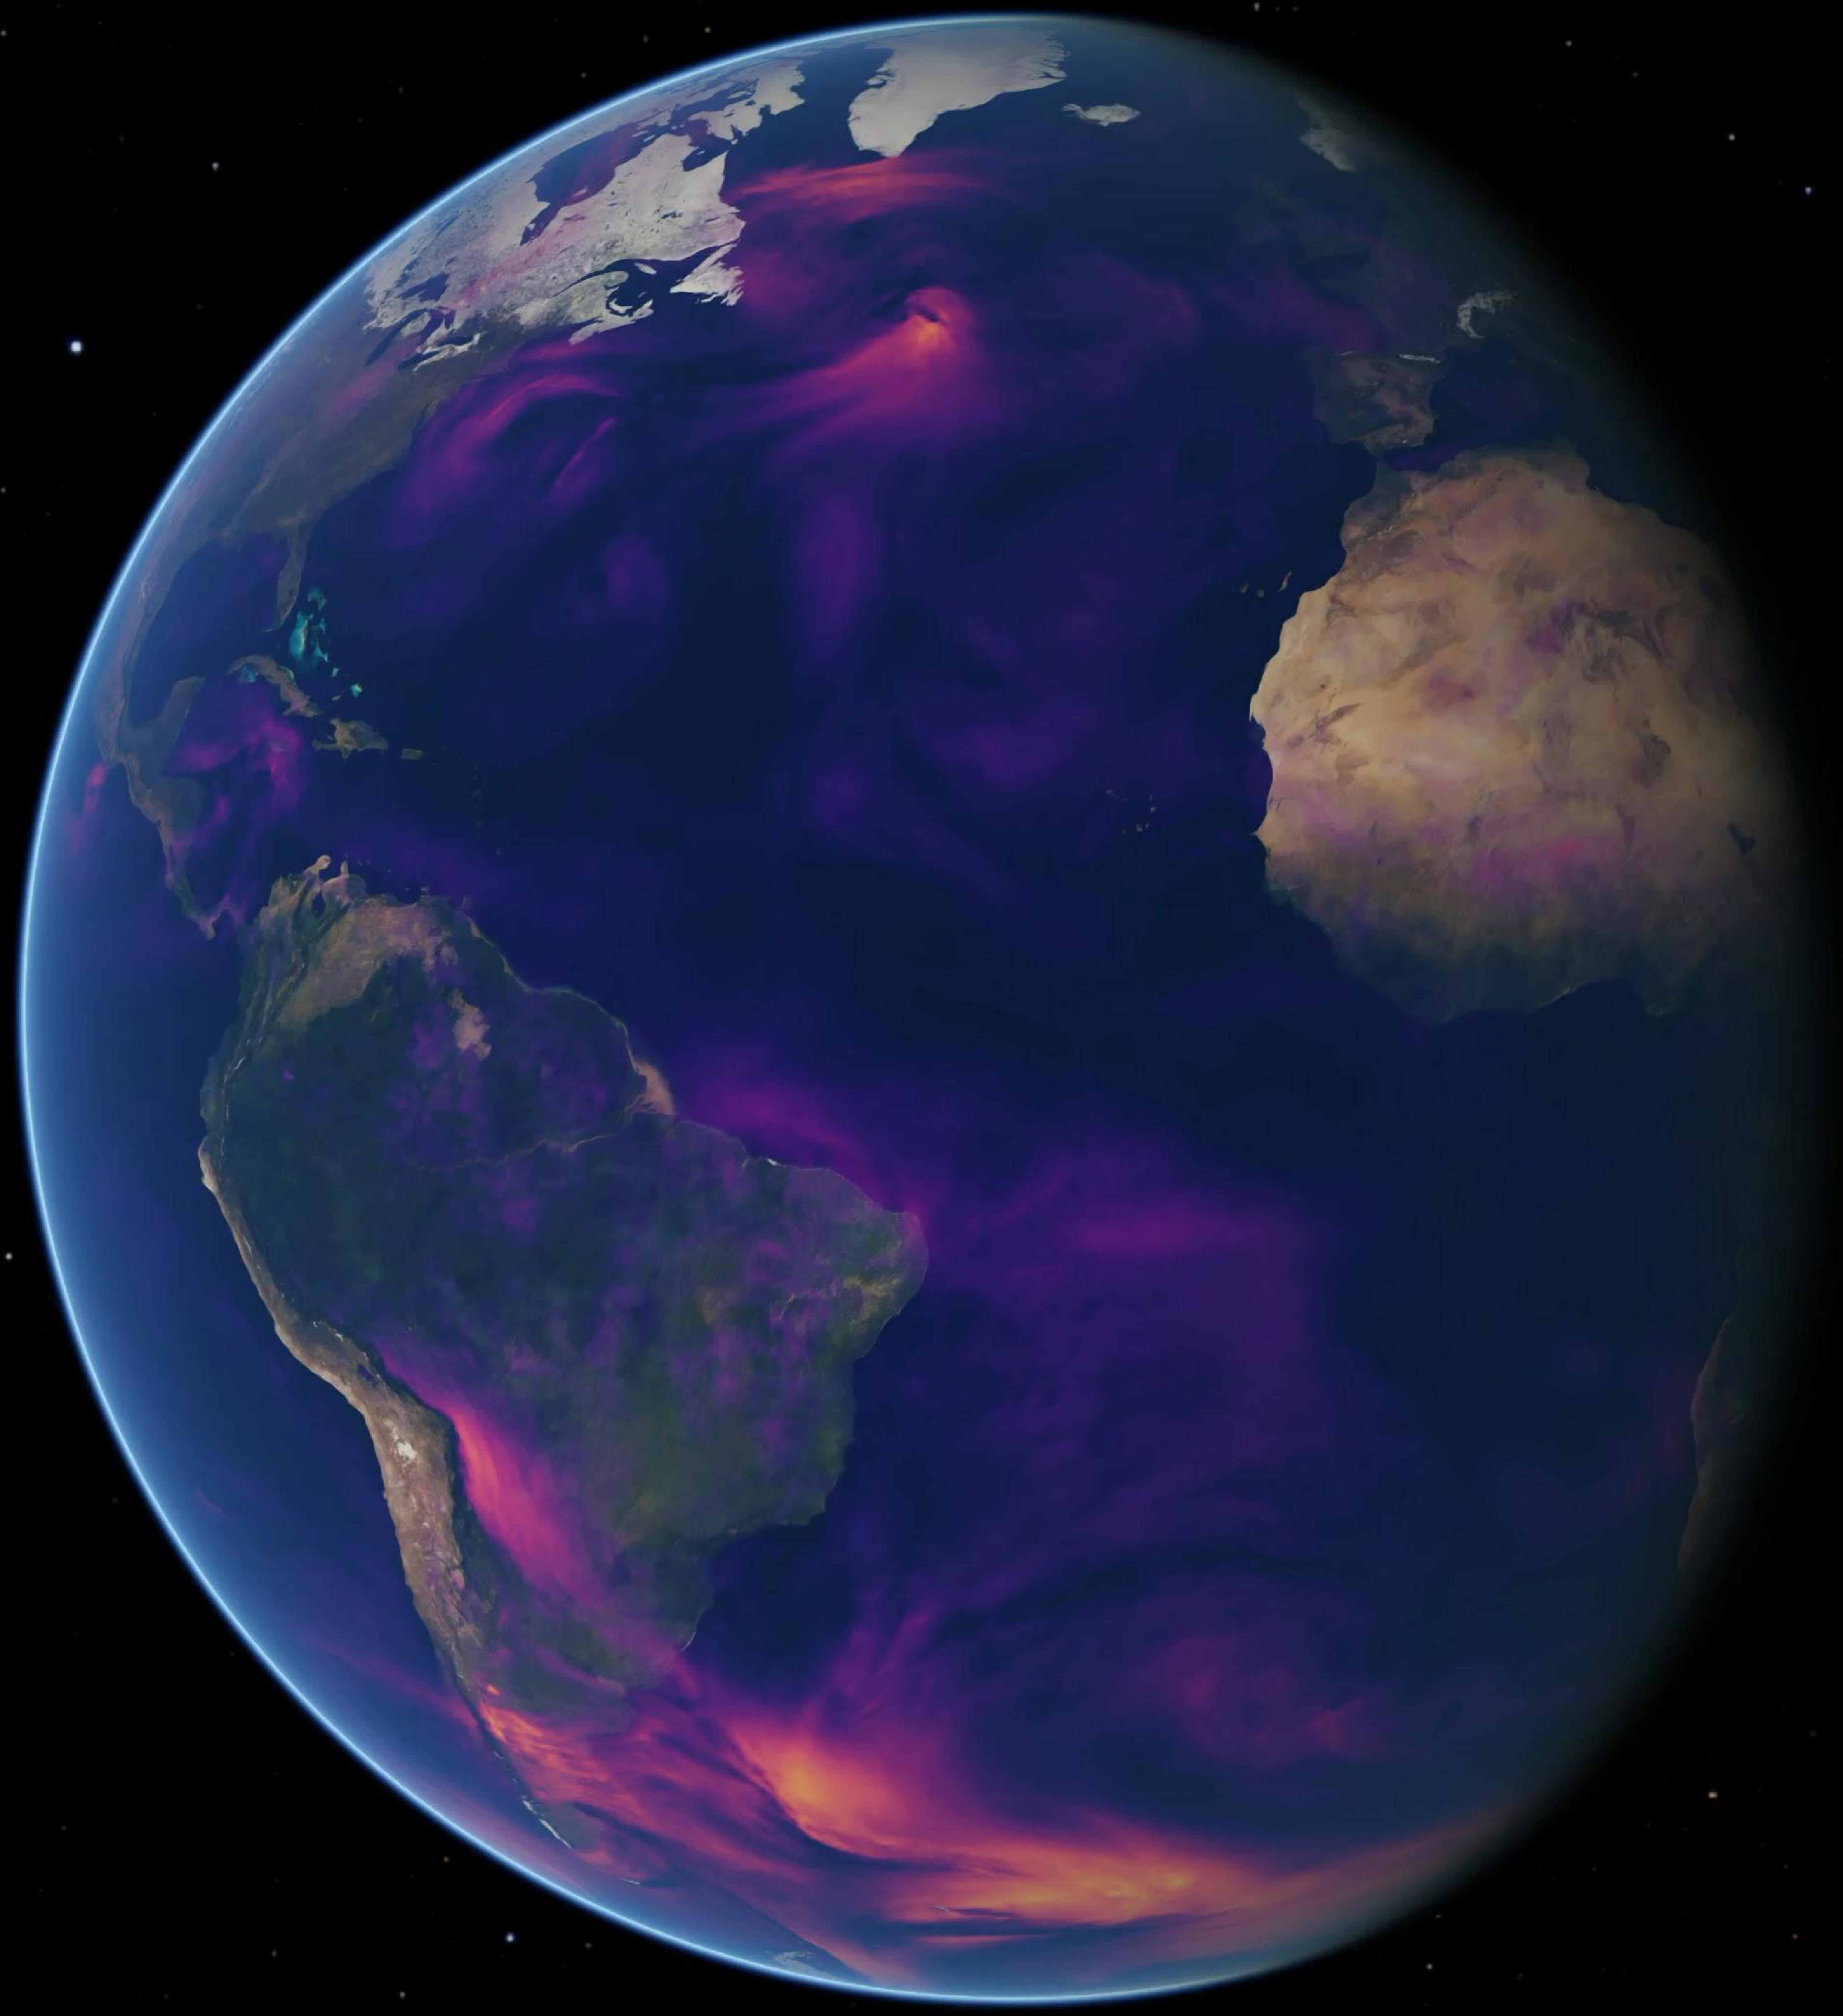

https://developer.nvidia.com/blog/spotlight-axa-explores-ai-driven-hurricane-risk-assessment-with-nvidia-earth-2/

The Mesoscale: An AI Weather Forecasting Frontier

Macroscale Weather

- Global datasets | 10-30 km resolution.
- Large phenomena – 100s to 1000s of km (e.g. cyclones)
- Negligible vertical acceleration of air; hydrostatic balance
- AI forecasts more skillful than physics models.

Microscale Weather

- National datasets | 1-5km resolution.
- Small phenomena: Thunderstorms, convection complexes.
- Hydrostatic balance not assumed. Buoyancy, stochasticity.
- **Potential of AI unknown**

Atmospheric River
~1000 km

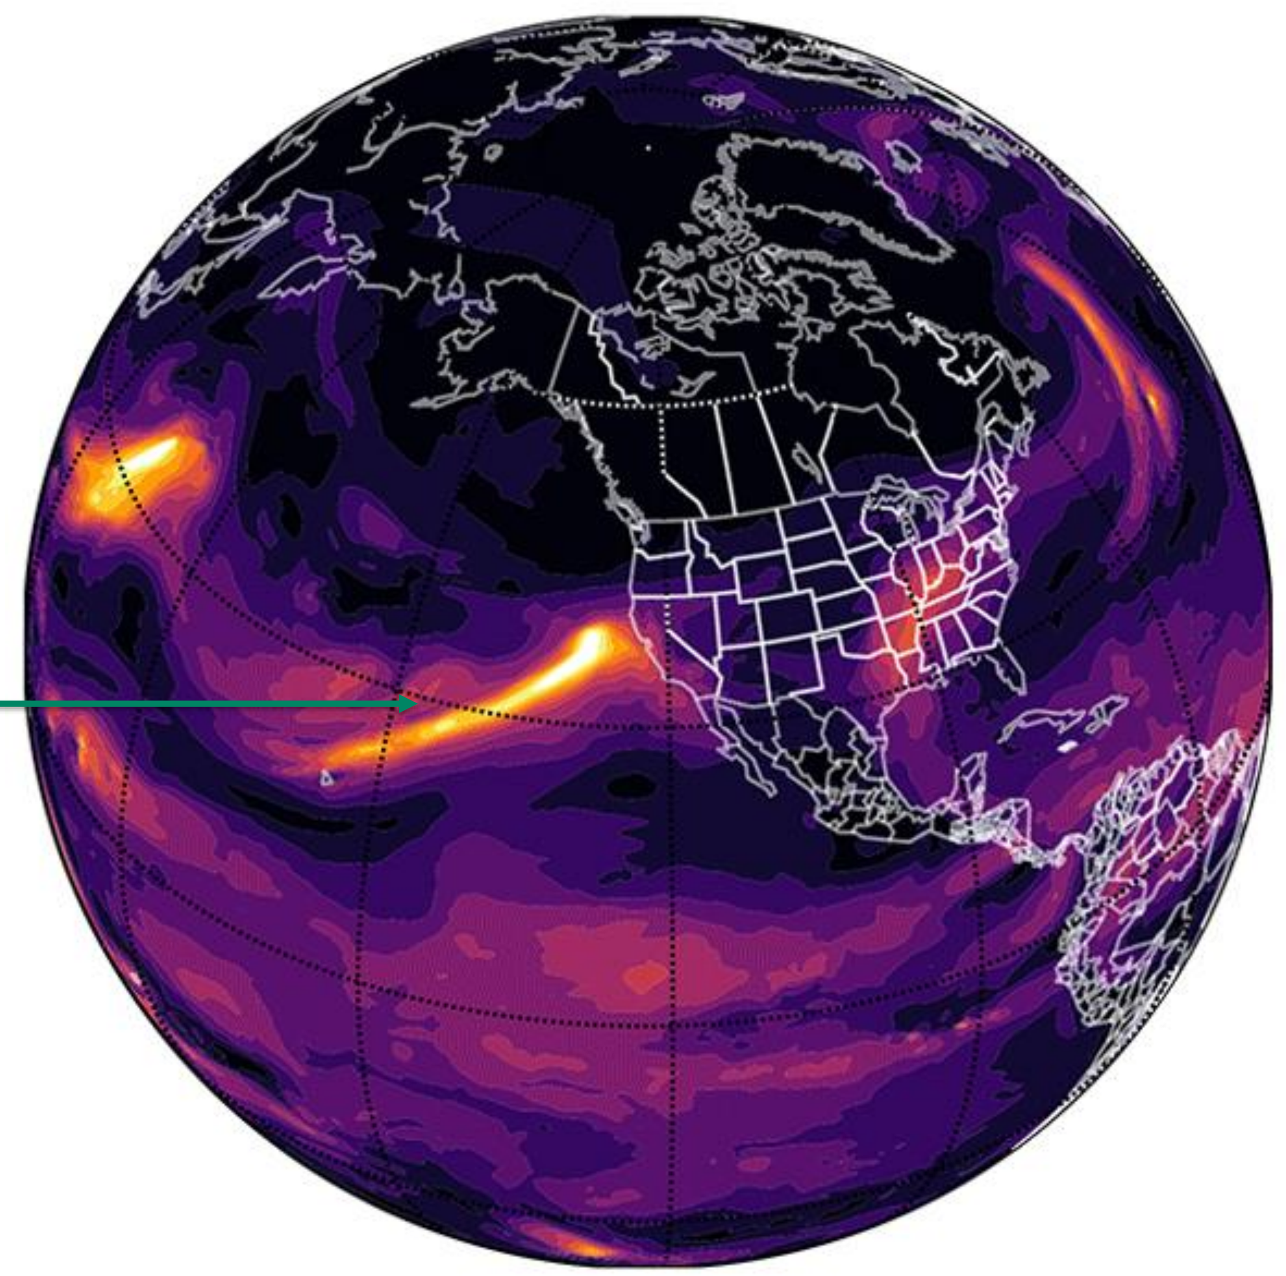

Organized Storm
Complex
~ 10 km

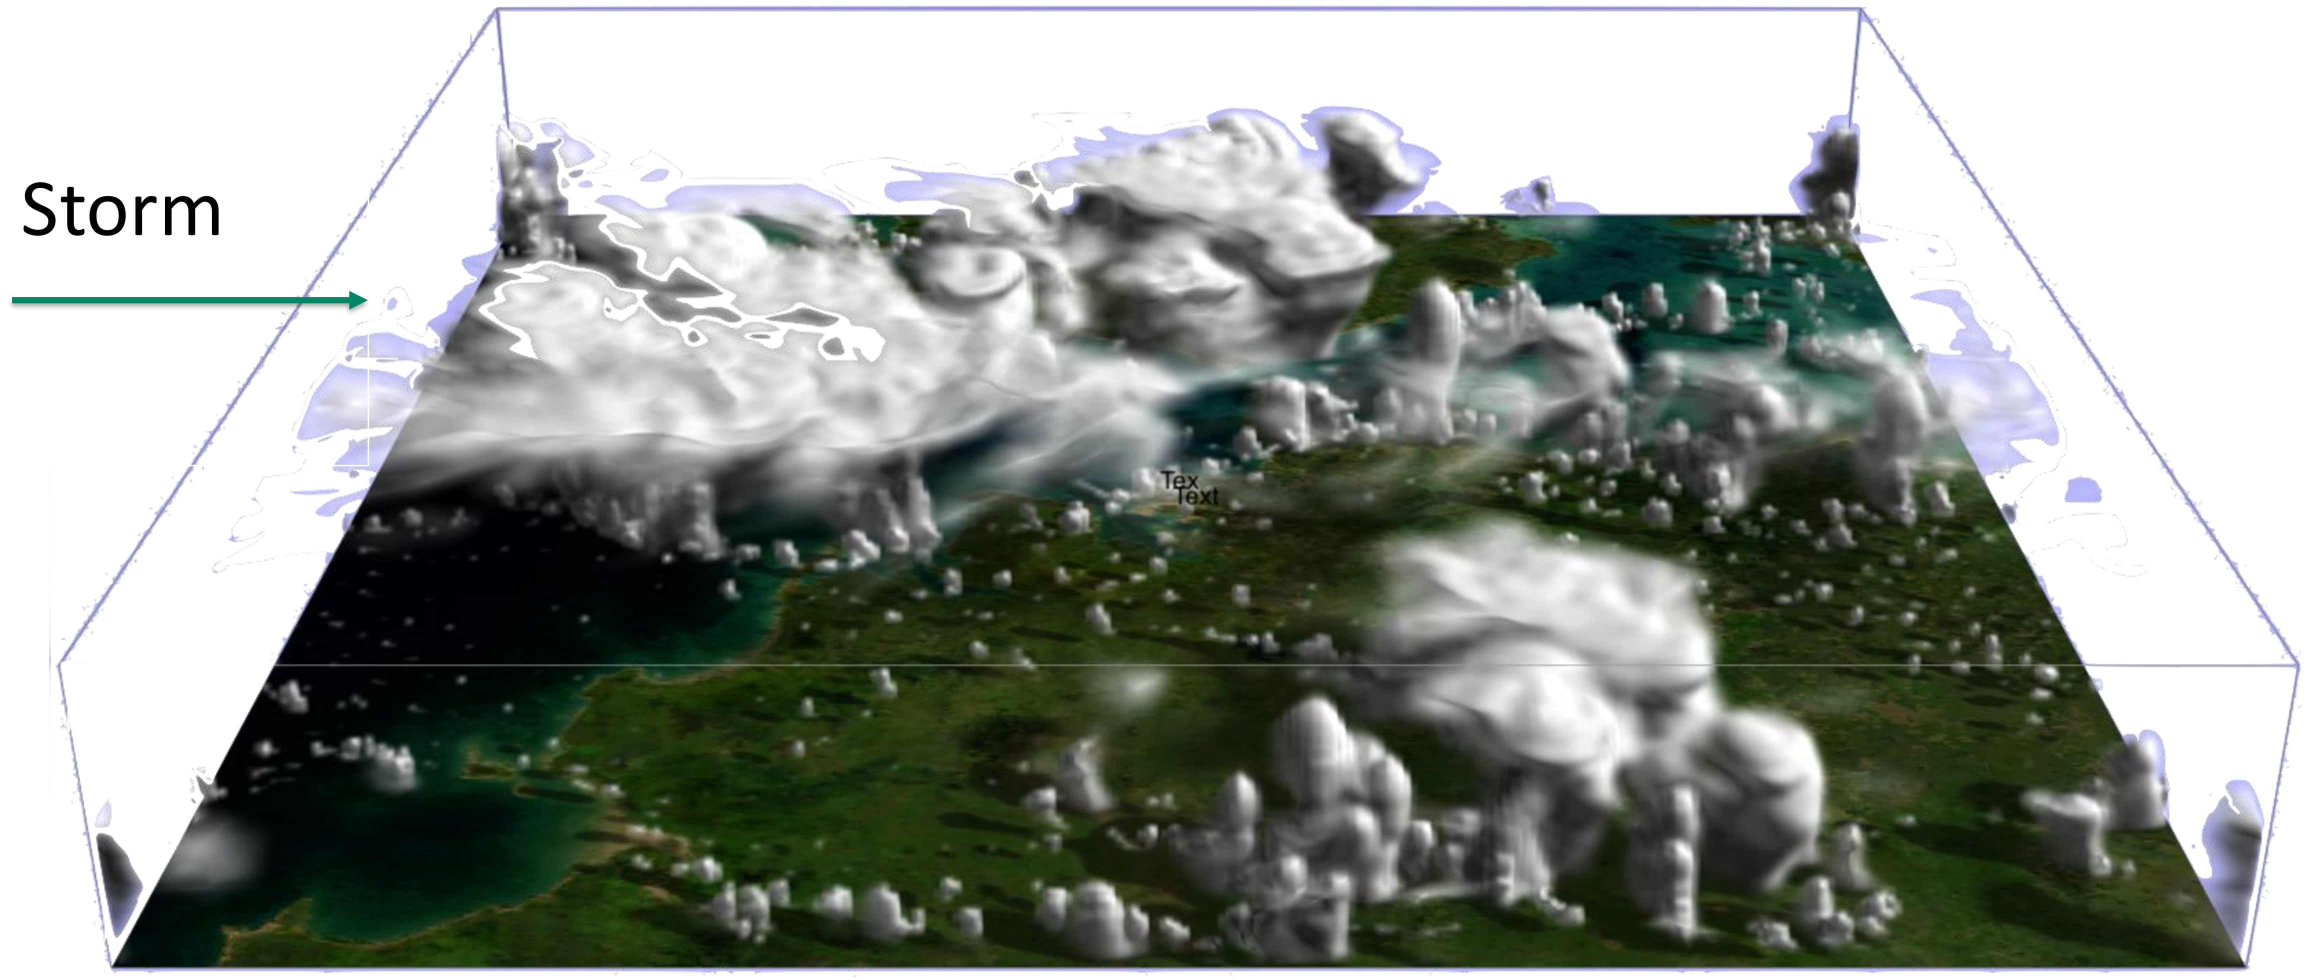

Efficient Downscaling with CorrDiff

Super fast super-resolution with generative diffusion

- AI models provide **physically realistic** representations of **small-scale weather**
- Trained to go from low resolution to high resolution (e.g., 25 km ERA5 to km-scale WRF)
- Stochasticity of generative models allows generating **ensemble of downscaled realizations**
- Inference is **orders of magnitude faster** than numerical (dynamical) downscaling

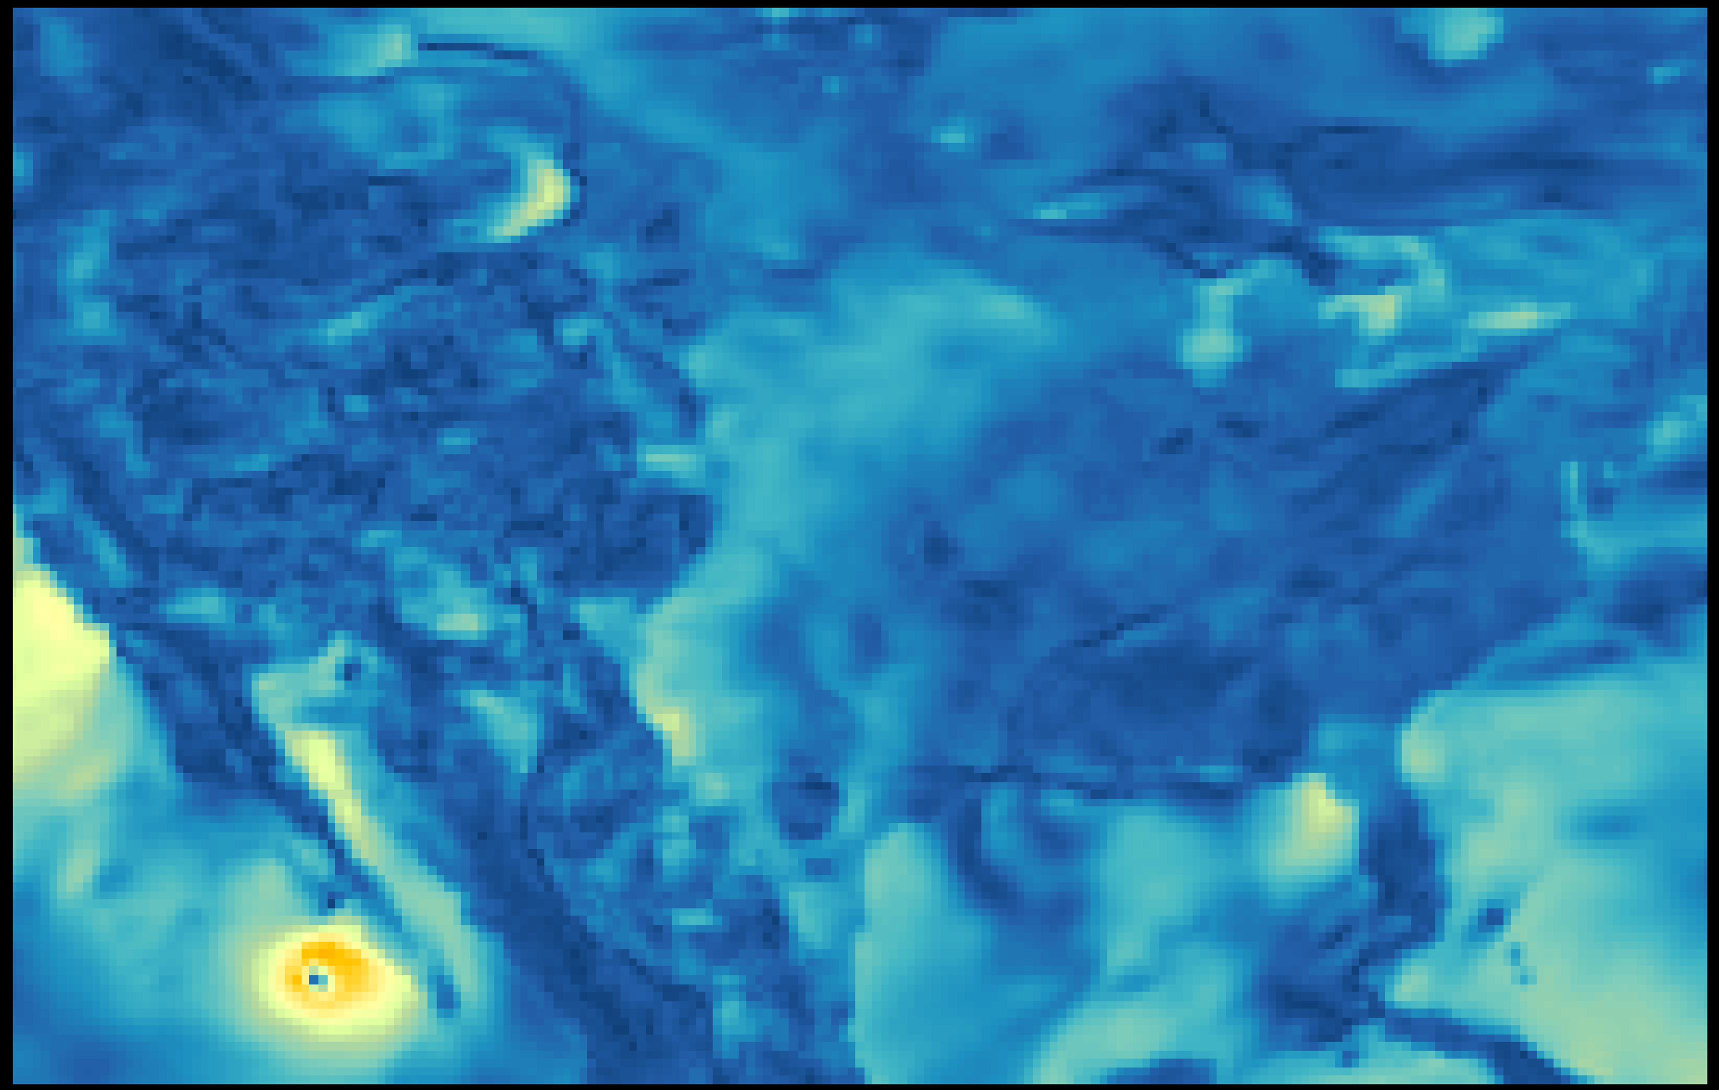

ca. 25 km per pixel

CorrDiff →

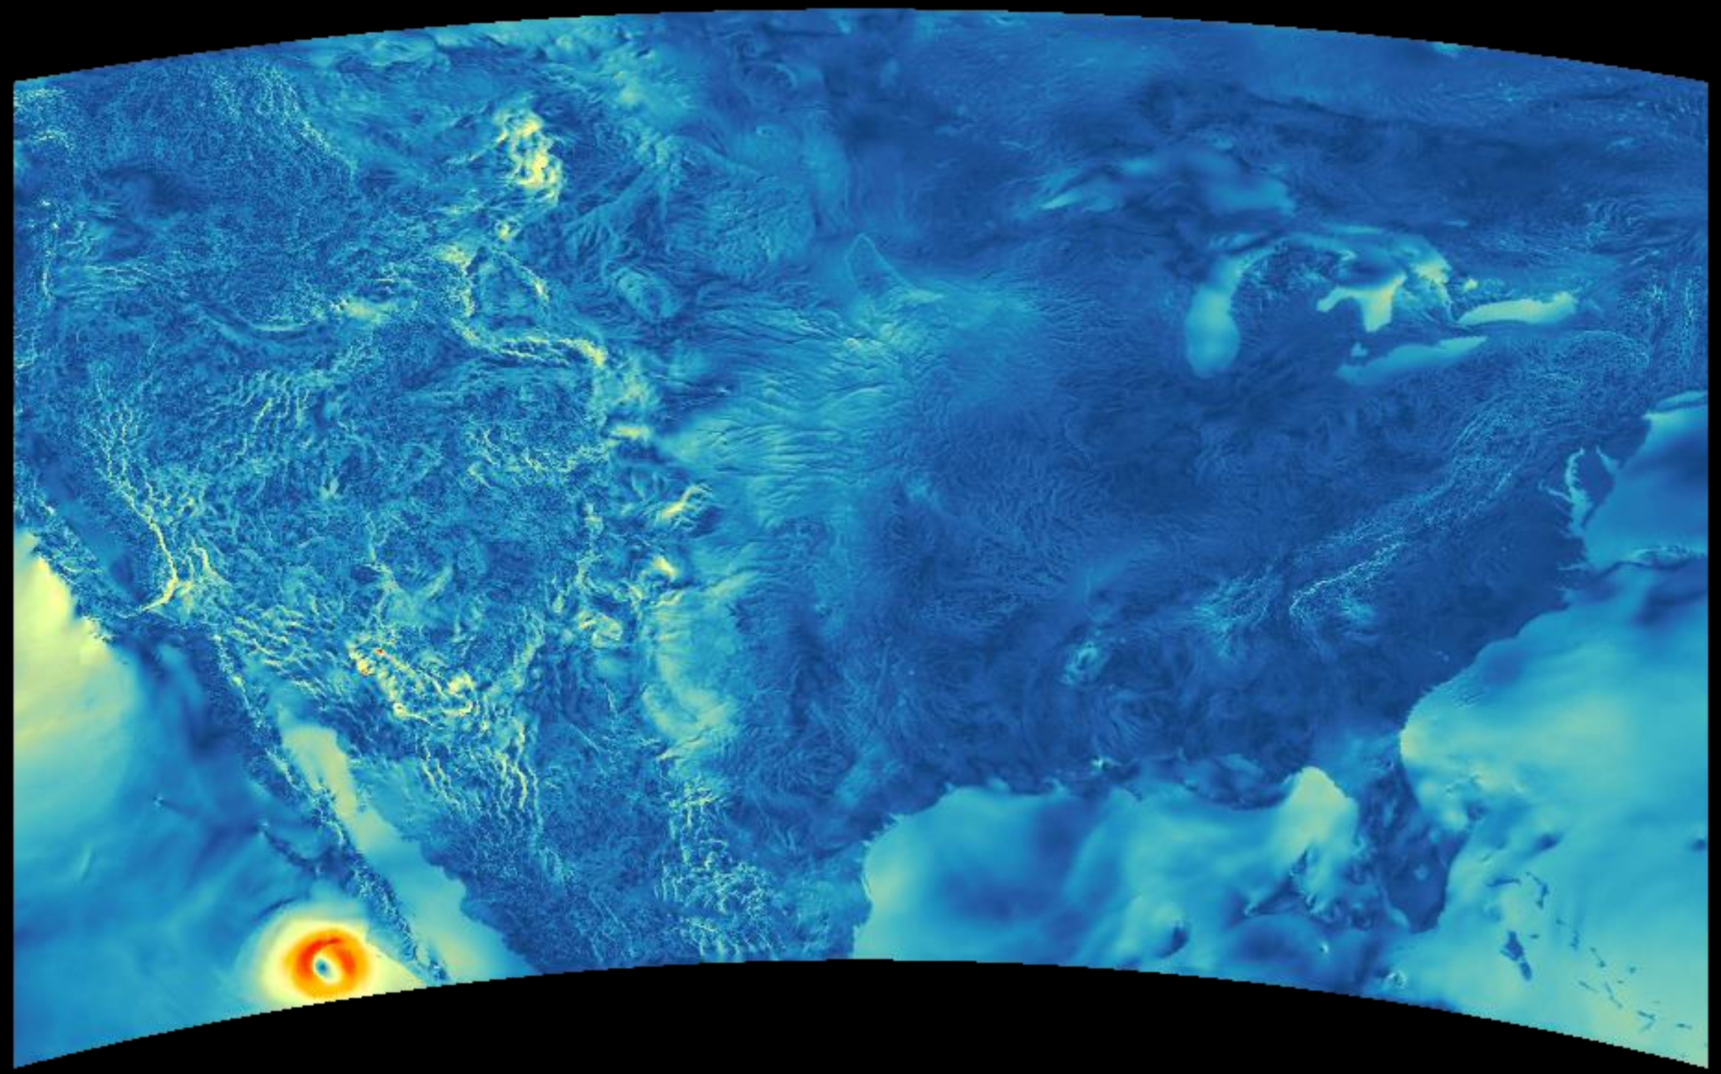

ca. 3 km per pixel

https://build.nvidia.com/nvidia/corrdiff

CorrDiff and StormCast Combine Regression and Diffusion

Two-step downscaling specifically developed for weather data

- **Regression model** predicts mean of high-resolution weather variables
- Mean is stochastically corrected through **diffusion model**
- Two-step process is thought to help bridge the significant distribution shift between input and output
- **CorrDiff** conducts **super-resolution**, **StormCast** makes **forecasts at km scale**

Corrdiff: https://arxiv.org/abs/2309.15214.

StormCast: https://arxiv.org/abs/2408.10958v1.

Downscaling with Diffusion Models

Iterative noise removal for resolving the fine scales of data

Application	Prompt
CorrDiff	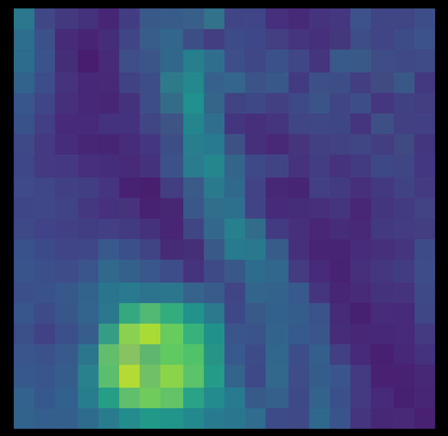
Super-resolution	
Image generation	“Photo of a tree”

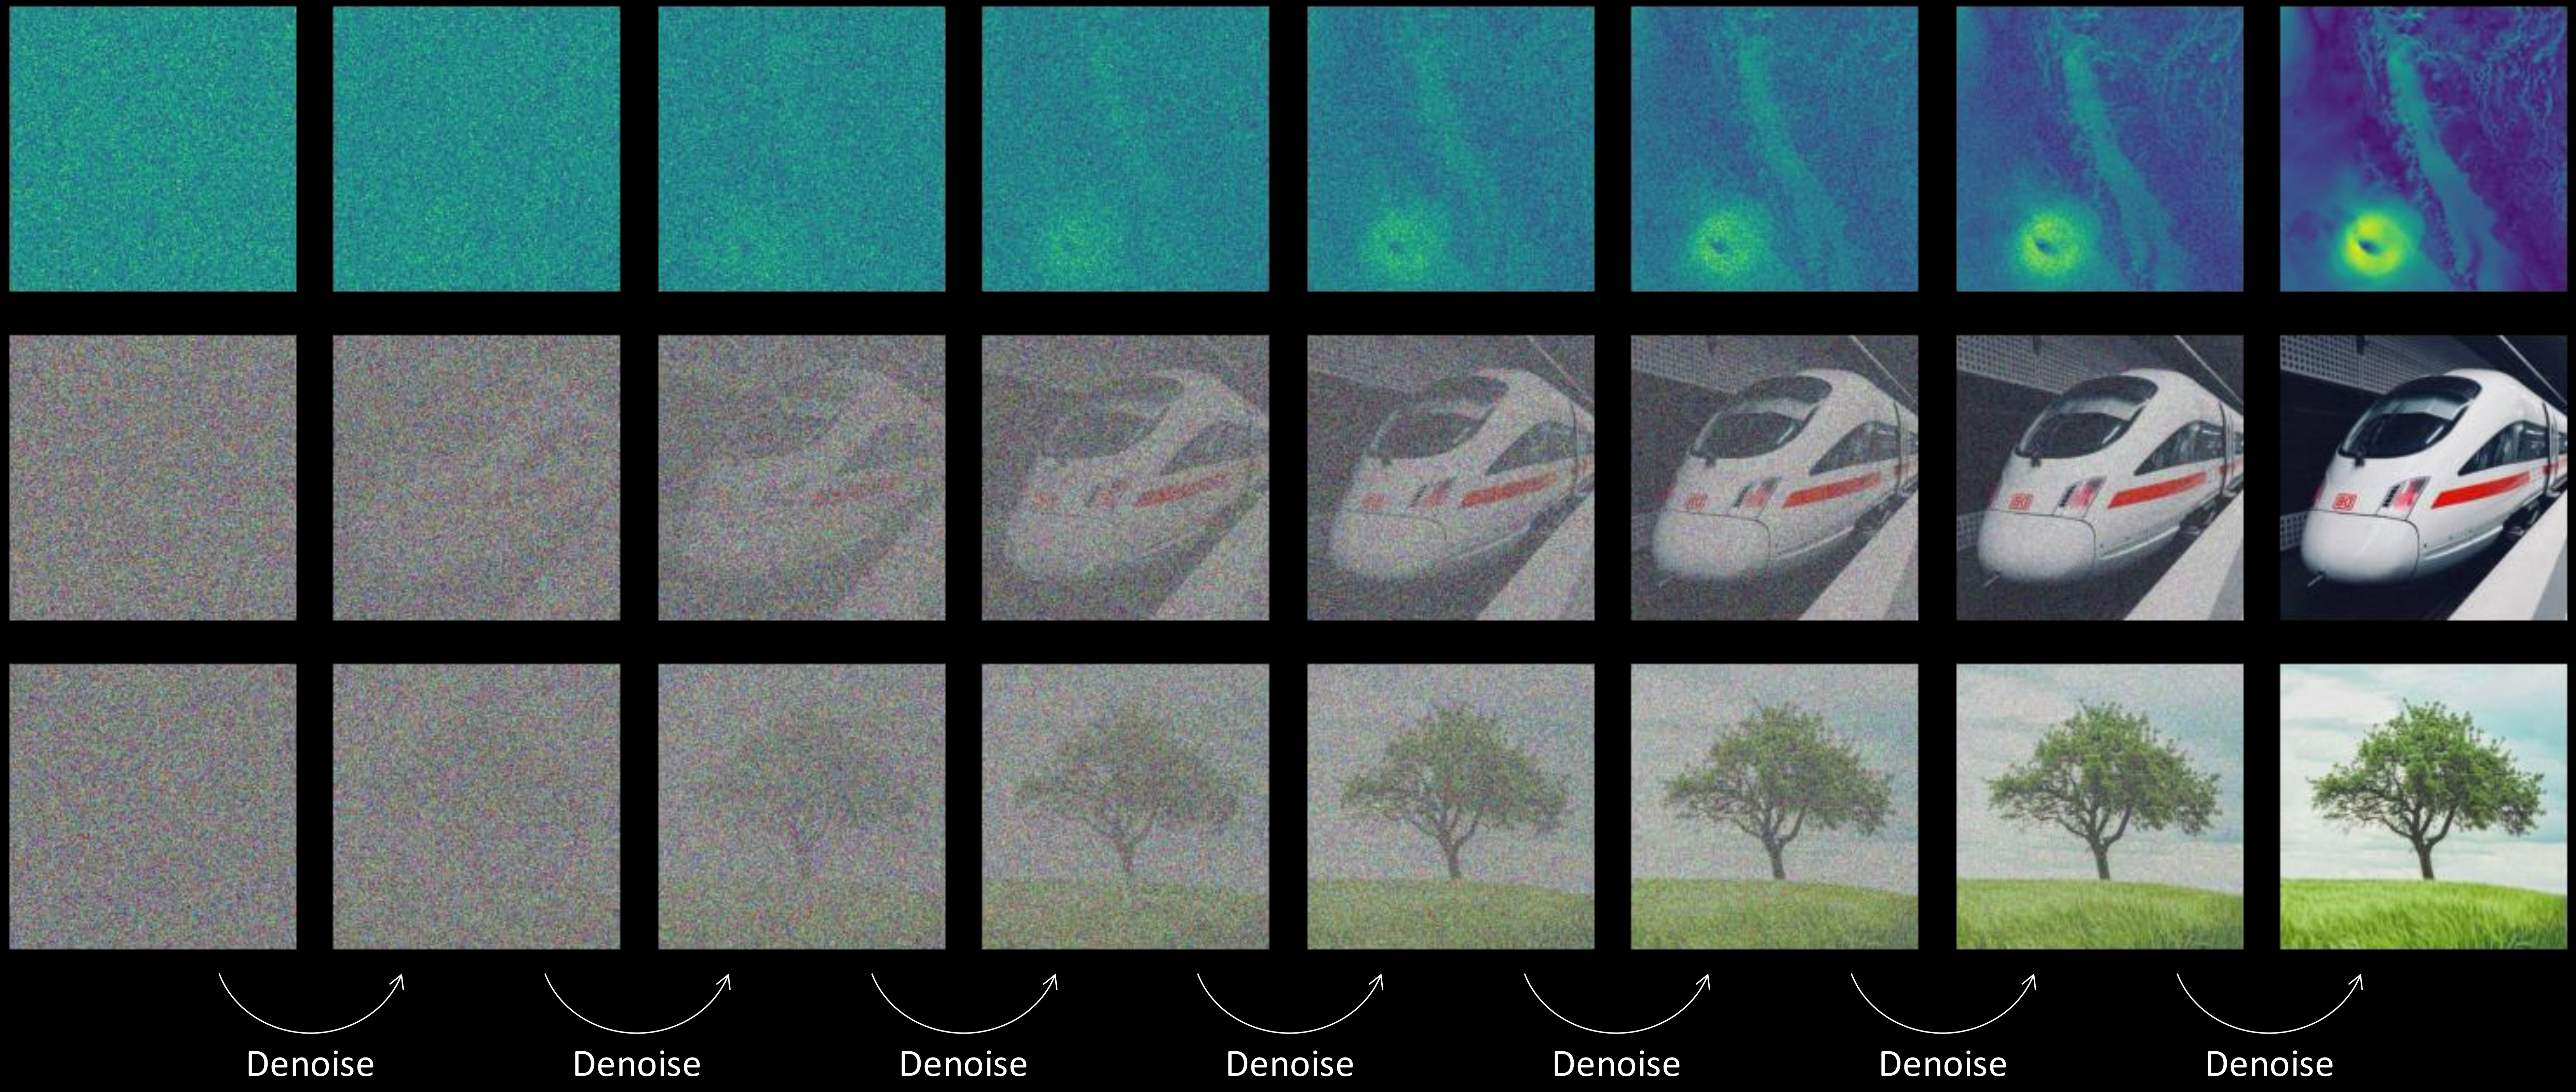

CorrDiff Downscaling Pipeline

Convert low-resolution to high-resolution forecast

Path to Bias Correction & Super-Resolution

Diversity across samples to account for downscaling uncertainty

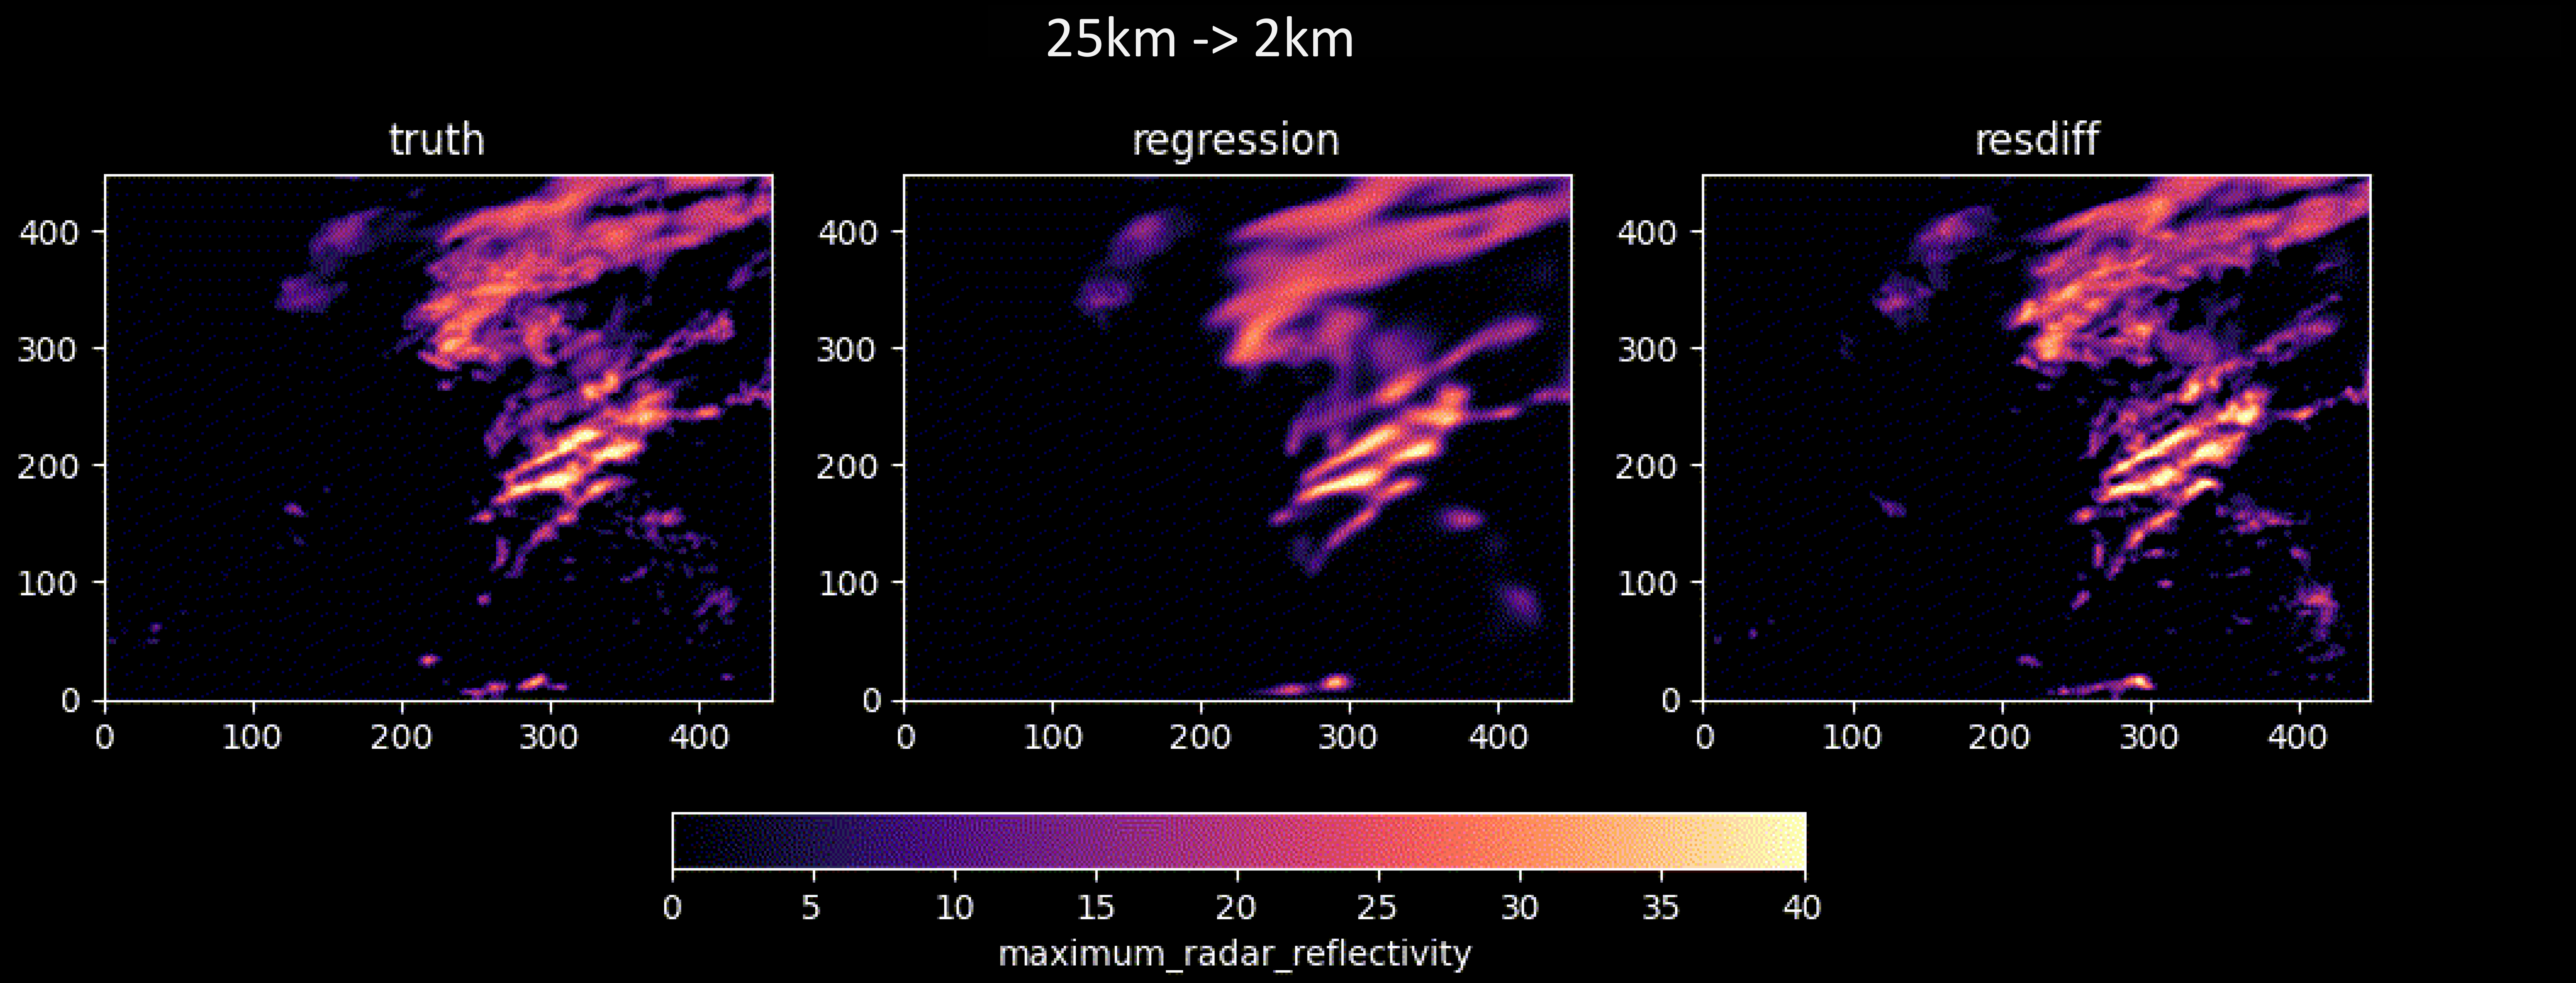

Mardani et al. 2023

StormCast - Motivation

Existing km-scale national AI weather models:

- Do not surpass physical such models on metrics of ensemble skill.
- Struggle to learn from sparse hourly data closest to observations.
- Can emulate physical models given subhourly output:

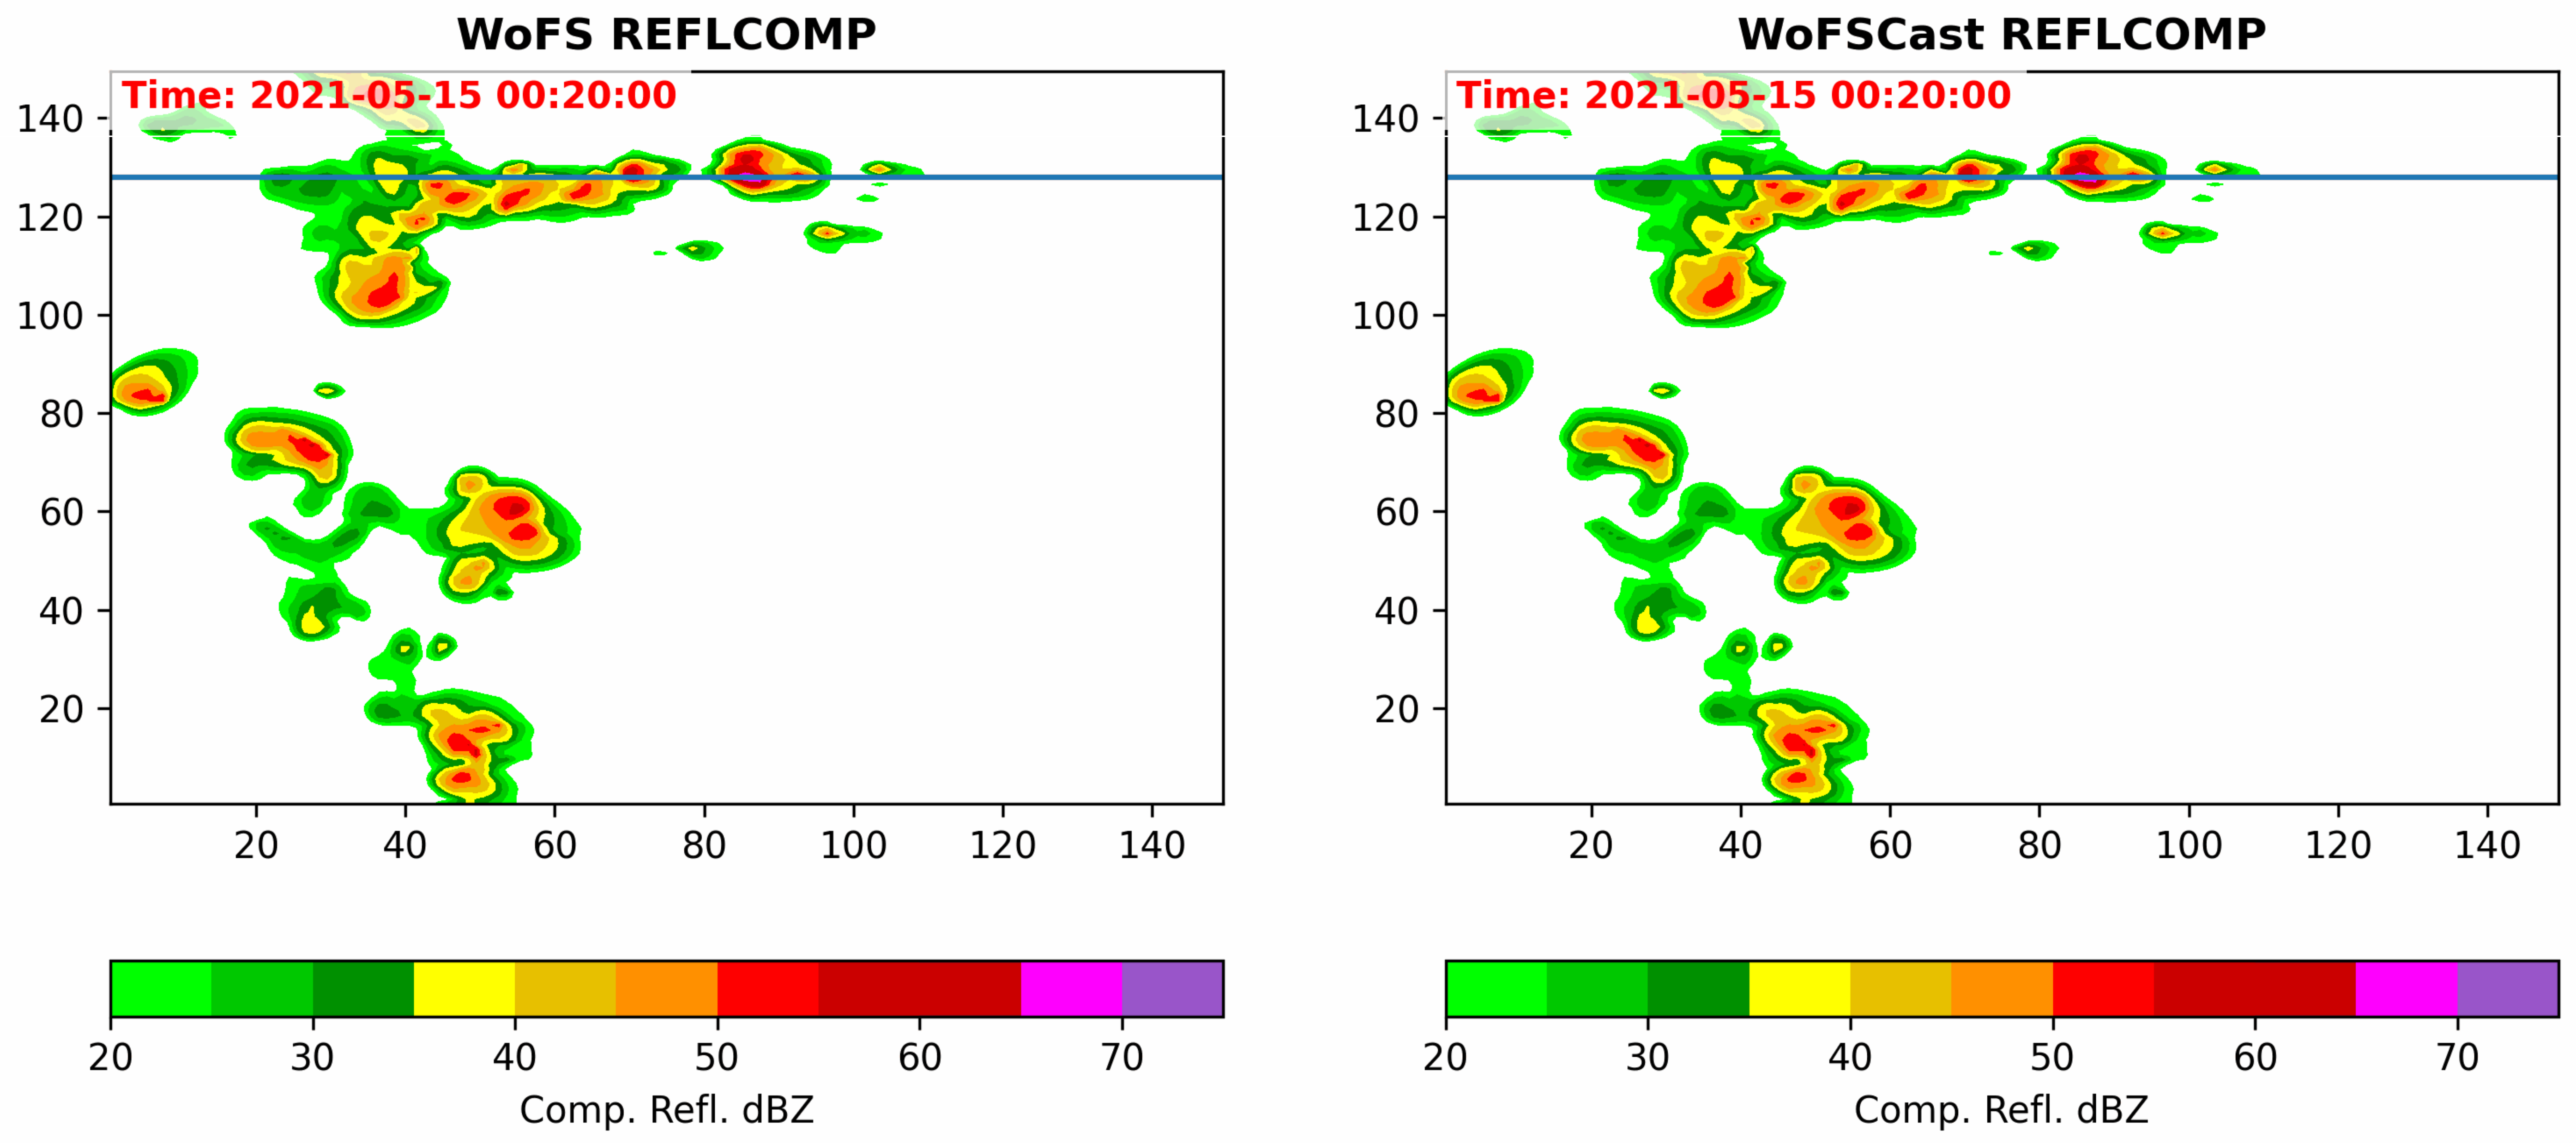

→

Multi-Scale Stochastic Architecture

- Cope with a sparse hourly time step
- Aligned with radar data fusion frequency.
- Generative methods to handle chaos.
- Condition on synoptic prior prediction.

Multi-Scale Inference Design

Analogous to the forcing of the US High Resolution Rapid Refresh (HRRR) National Weather Model

Coarse-resolution: ML model is made aware of a prior, 25-km global forecast.

Fine-resolution: Is initialized with radar-assimilating HRRR analysis...

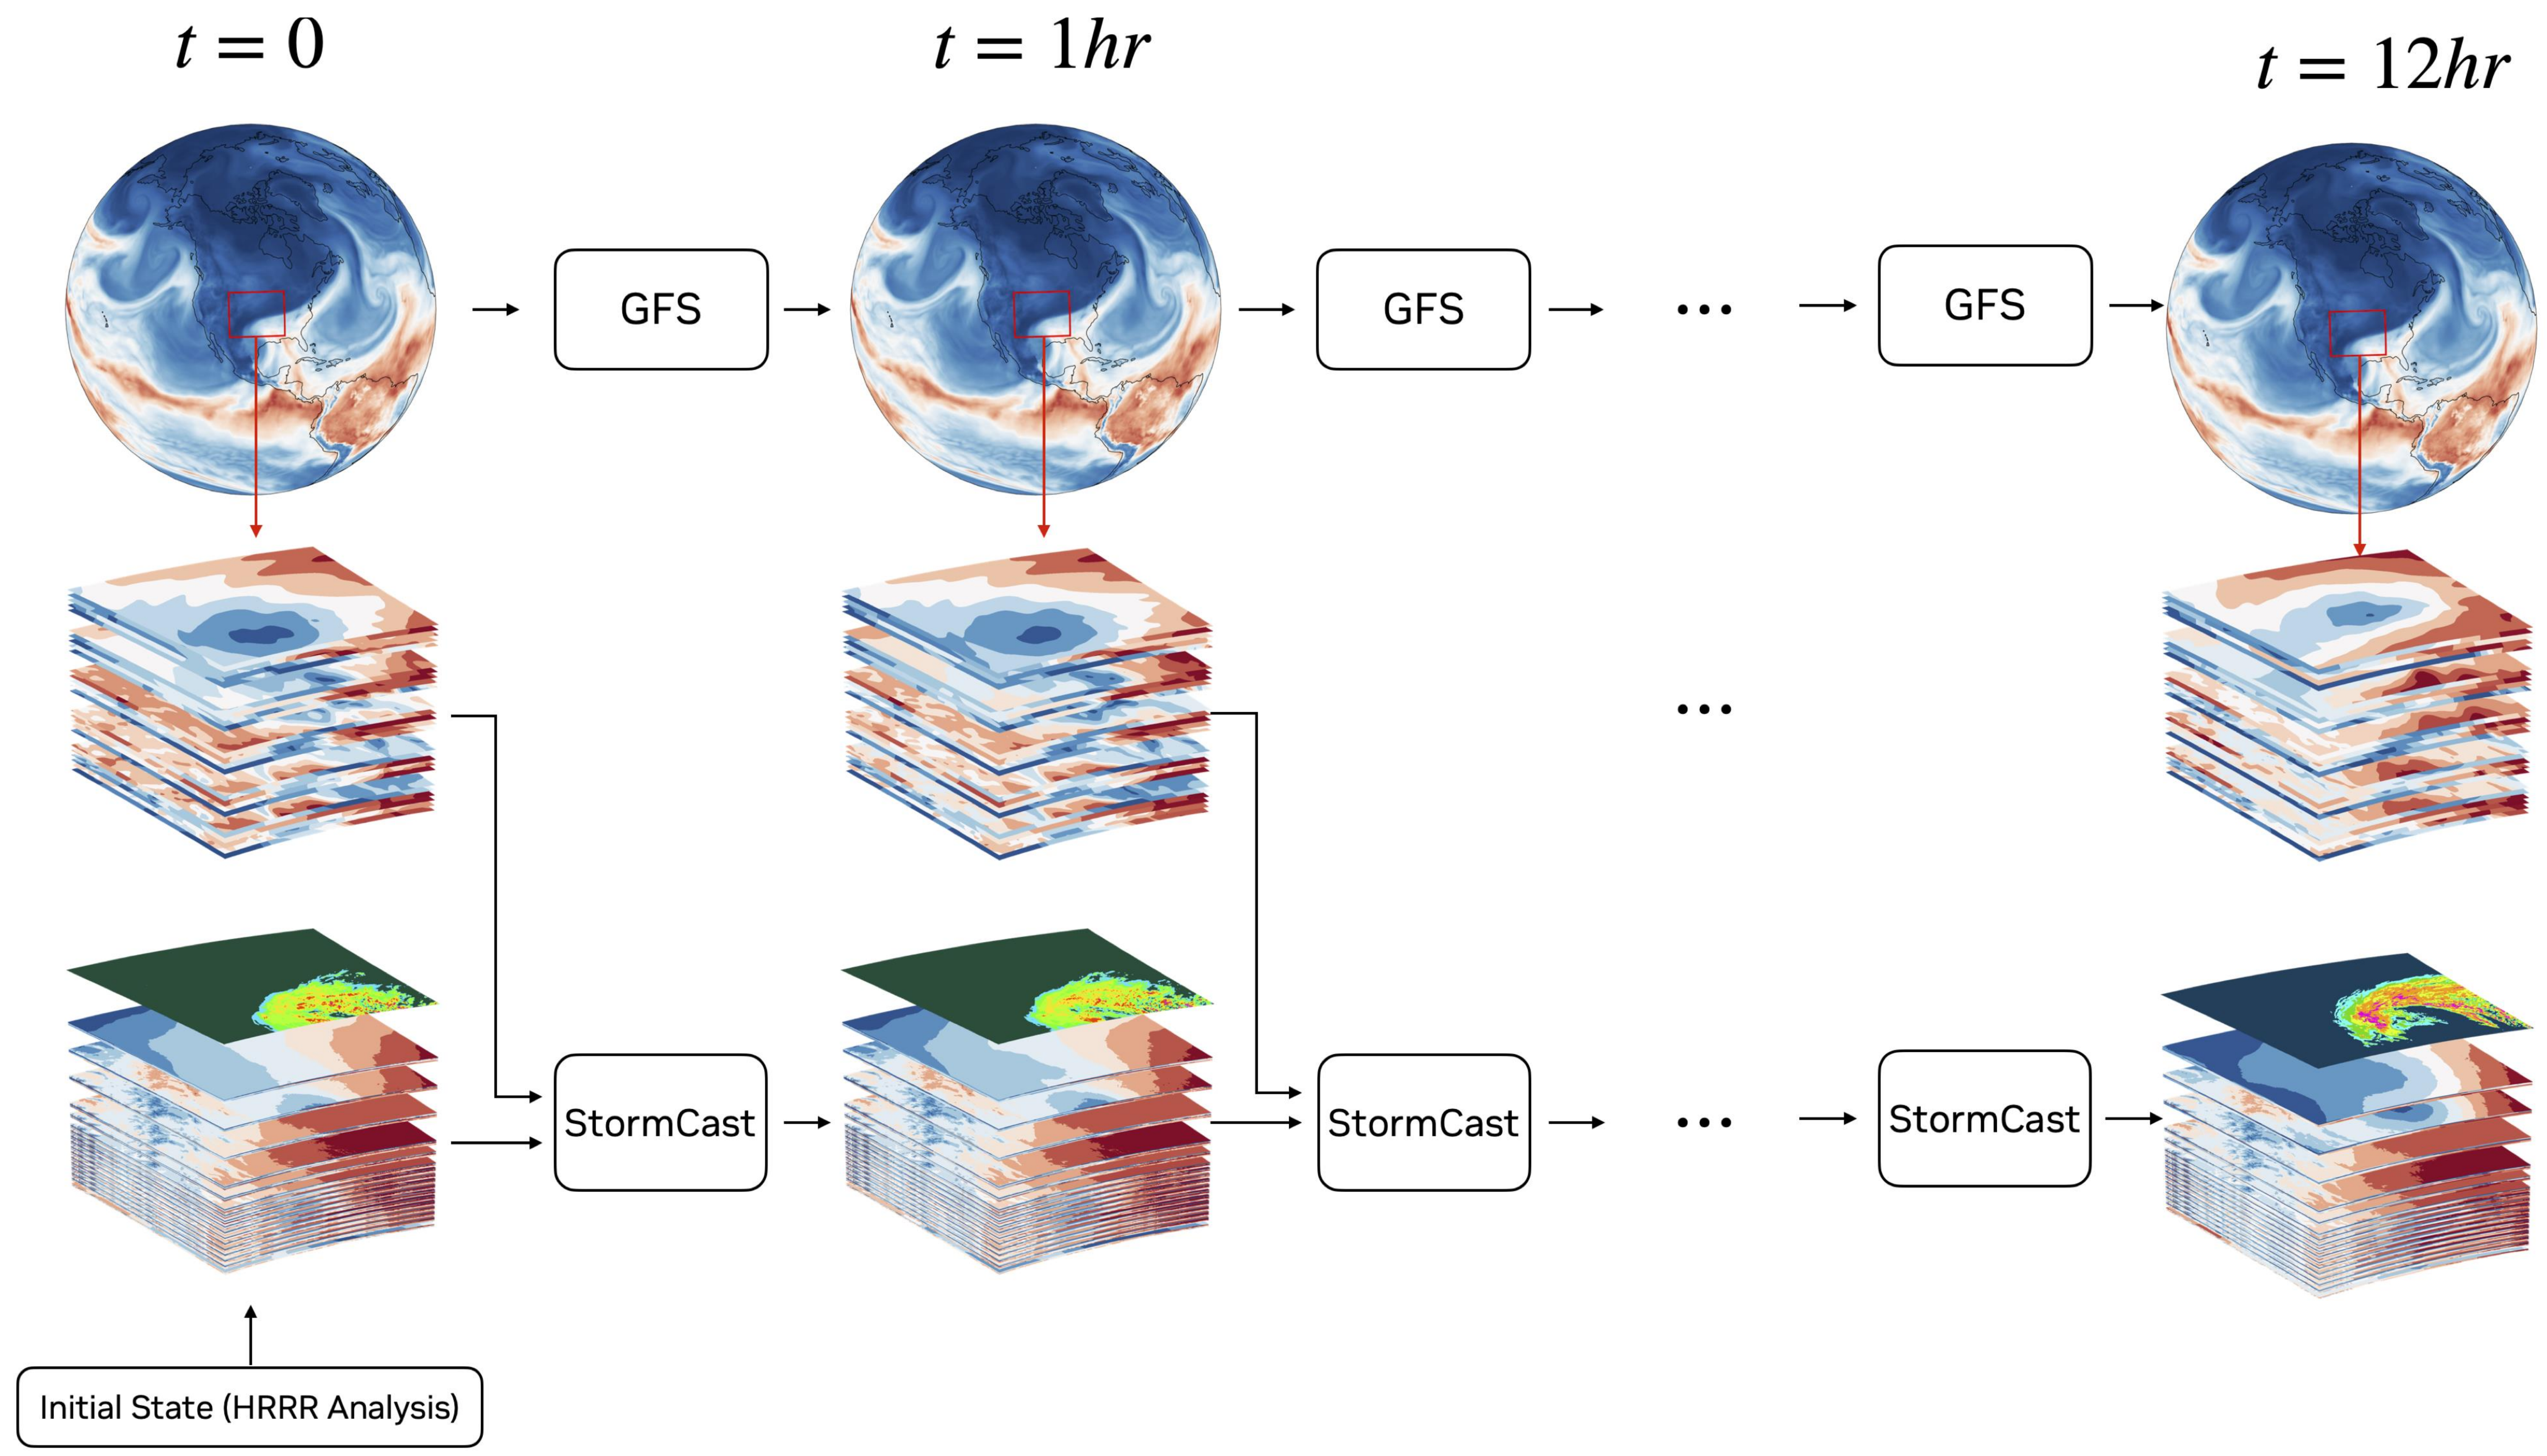

StormCast Leverages Residual Trick like CorrDiff

NVIDIA's Generative Super-Resolution Architecture Powered by Residual Diffusion

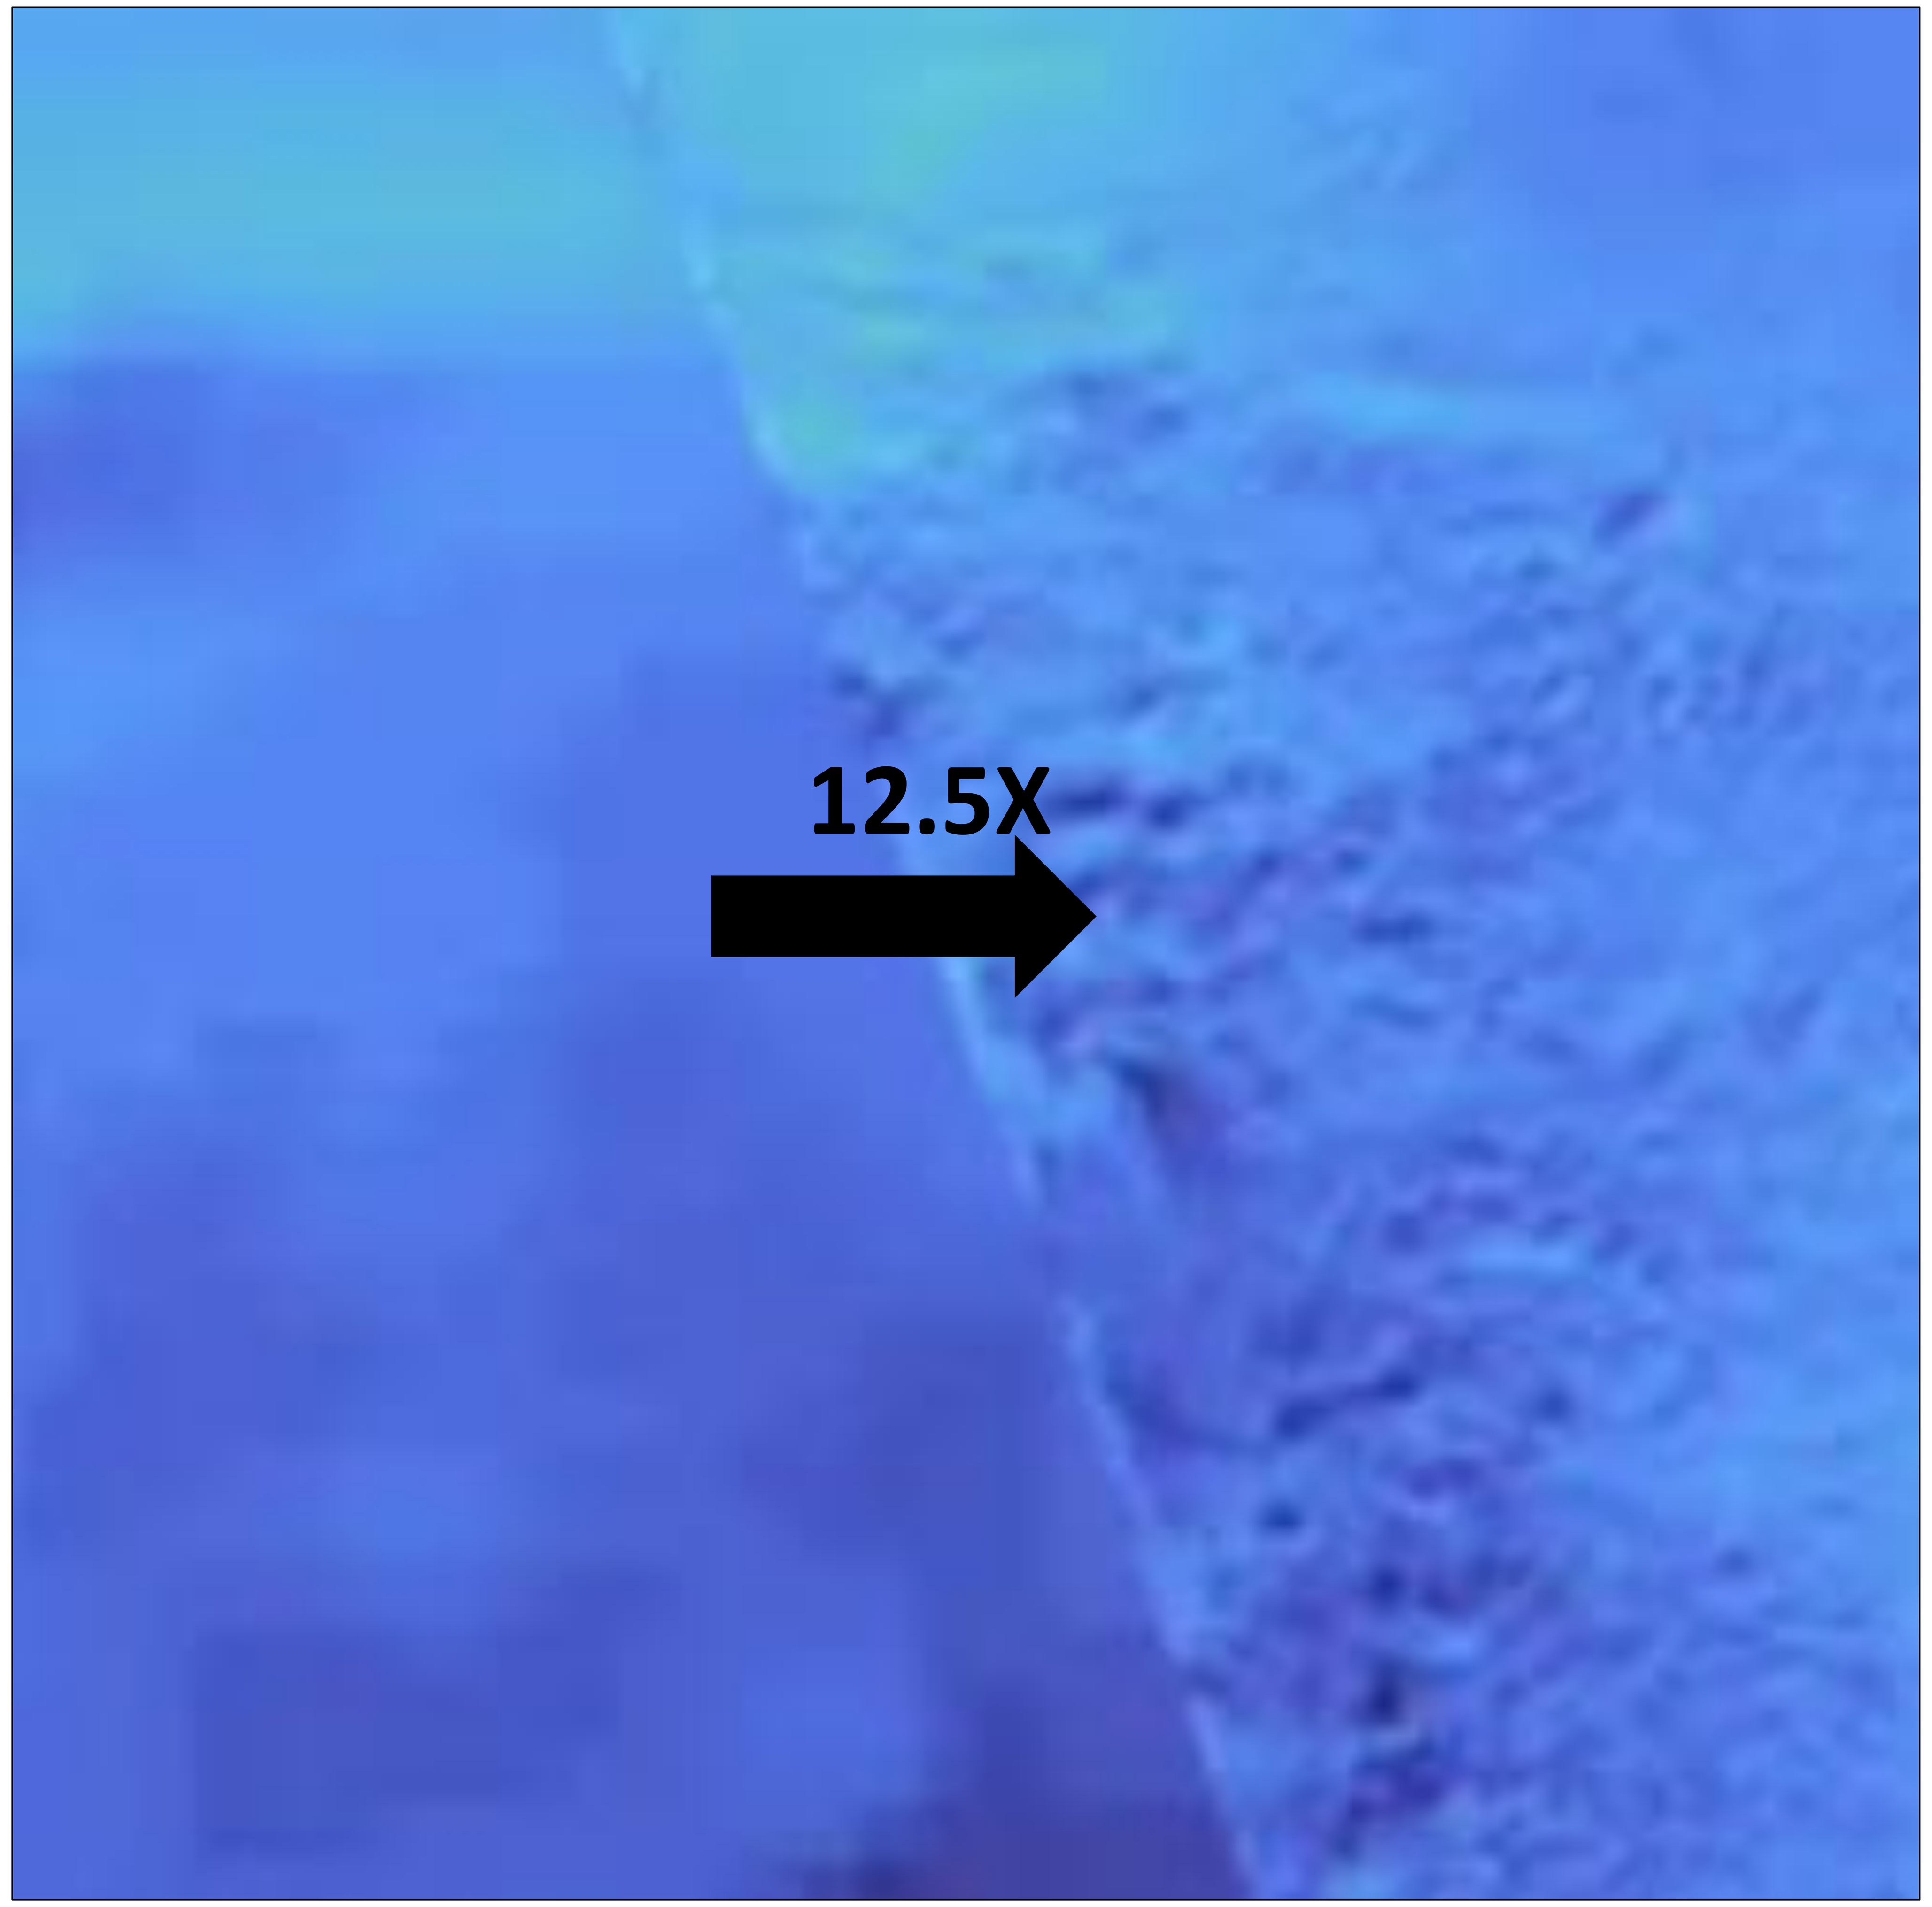

CorrDiff: Super-Resolution & Channel Synthesis

communications earth & environment

Article

https://doi.org/10.1038/s43247-025-02042-5

Residual corrective diffusion modeling for km-scale atmospheric downscaling

Check for updates

Morteza Mardani[1,3], Noah Brenowitz[1,3], Yair Cohen[1,3], Jaideep Pathak[1], Chieh-Yu Chen[1], Cheng-Chin Liu[2], Arash Vahdat[1], Mohammad Amin Nabian[1], Tao Ge[1], Akshay Subramaniam[1], Karthik Kashinath[1], Jan Kautz[1] & Mike Pritchard[1]

State of the art for weather and climate hazard prediction requires expensive km-scale numerical simulations. Here, a generative diffusion model is explored for downscaling global inputs to km-scale, as a cost-effective alternative. The model is trained to predict 2 km data from an operational regional weather model over Taiwan, conditioned on a 25 km reanalysis. To address the large resolution ratio, different physics and synthesize new channels, we employ a two-step approach. A deterministic model first predicts the mean, followed by a generative diffusion model that predicts the residual. The model exhibits encouraging deterministic and probabilistic skills, spectra and distributions that recover power law relationships in the target data. In case studies of coherent weather phenomena, it sharpens gradients in cold fronts and intensifies typhoons while synthesizing rainbands. Calibration of model uncertainty remains challenging. The prospect of unifying such methods with coarser global models implies a potential for global-to-regional machine learning simulation.

StormCast: Example Forecast. July 17, 2024

Evolution of Central US Radar Reflectivity – a Proxy for Precipitation

StormCast Ensemble Members

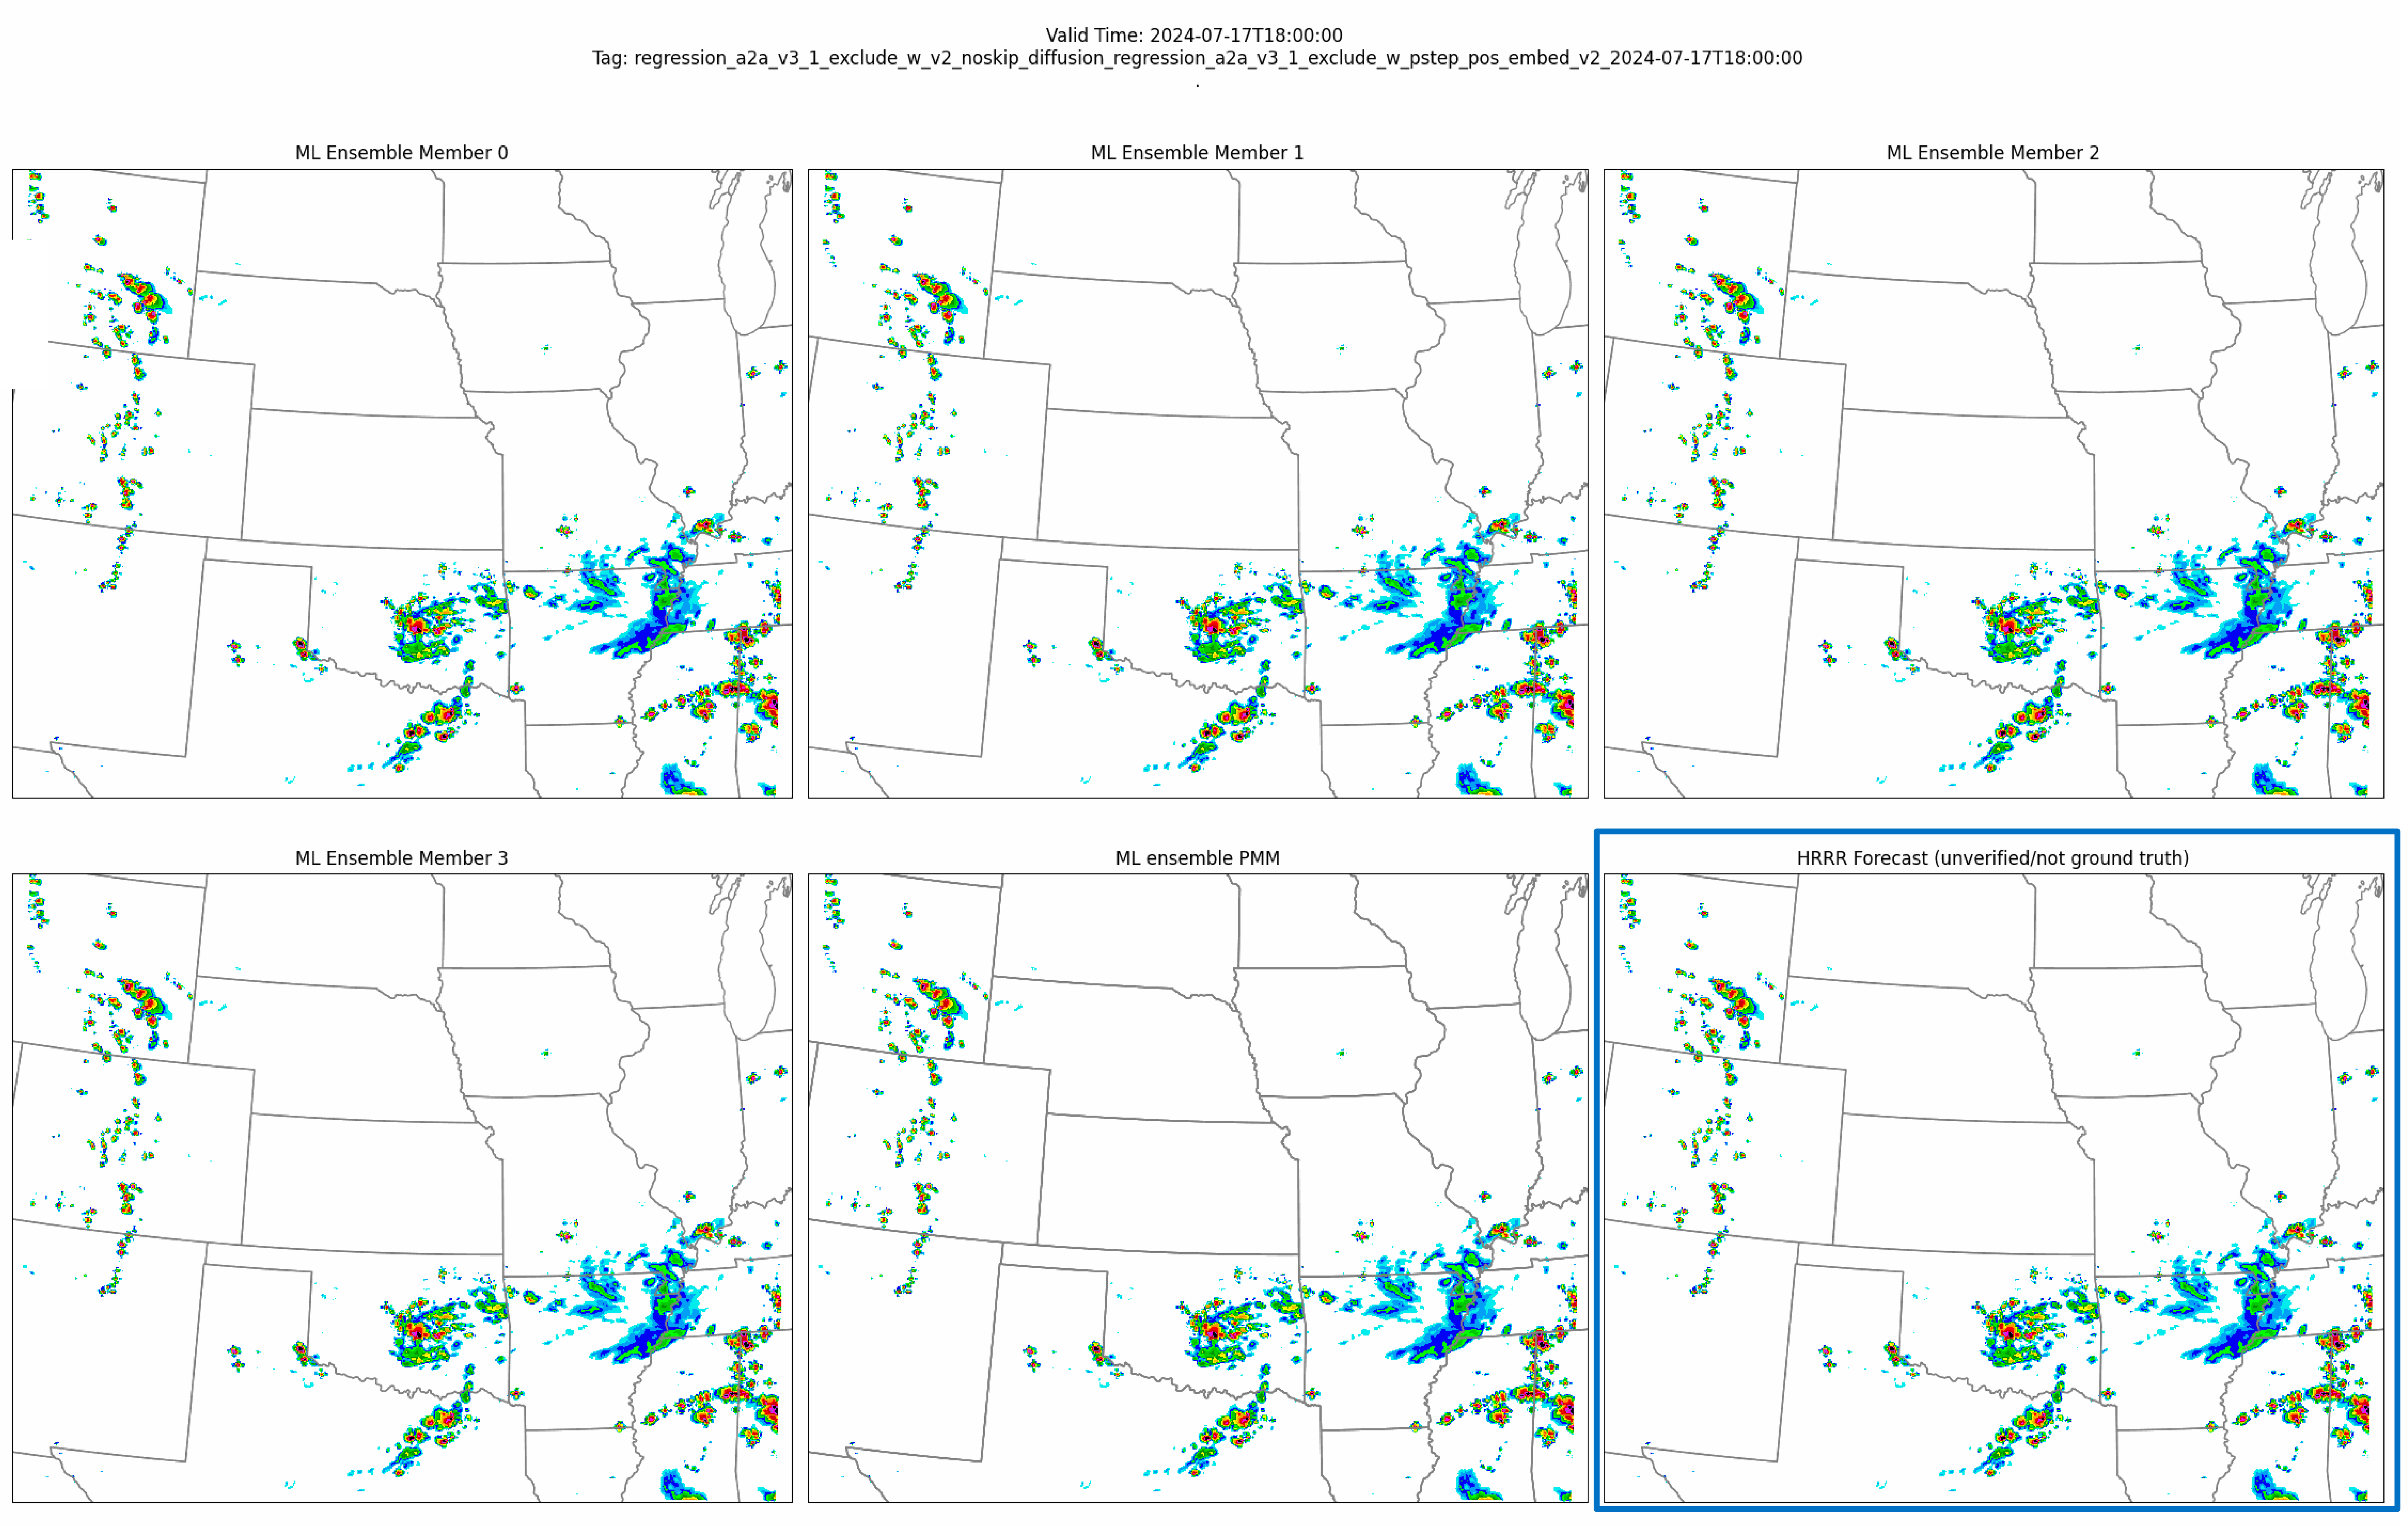

HRRR baseline

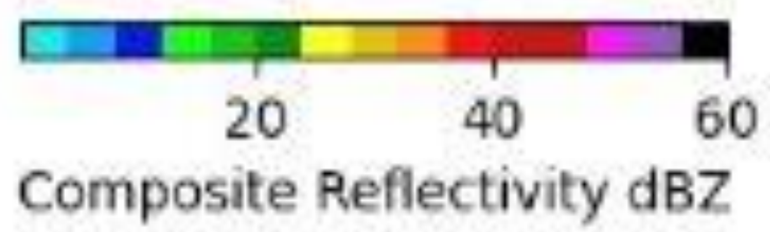

Skill Surpasses HRRR Physical Model

Using a 5-Member Ensemble Probability Matched Mean (PMM)

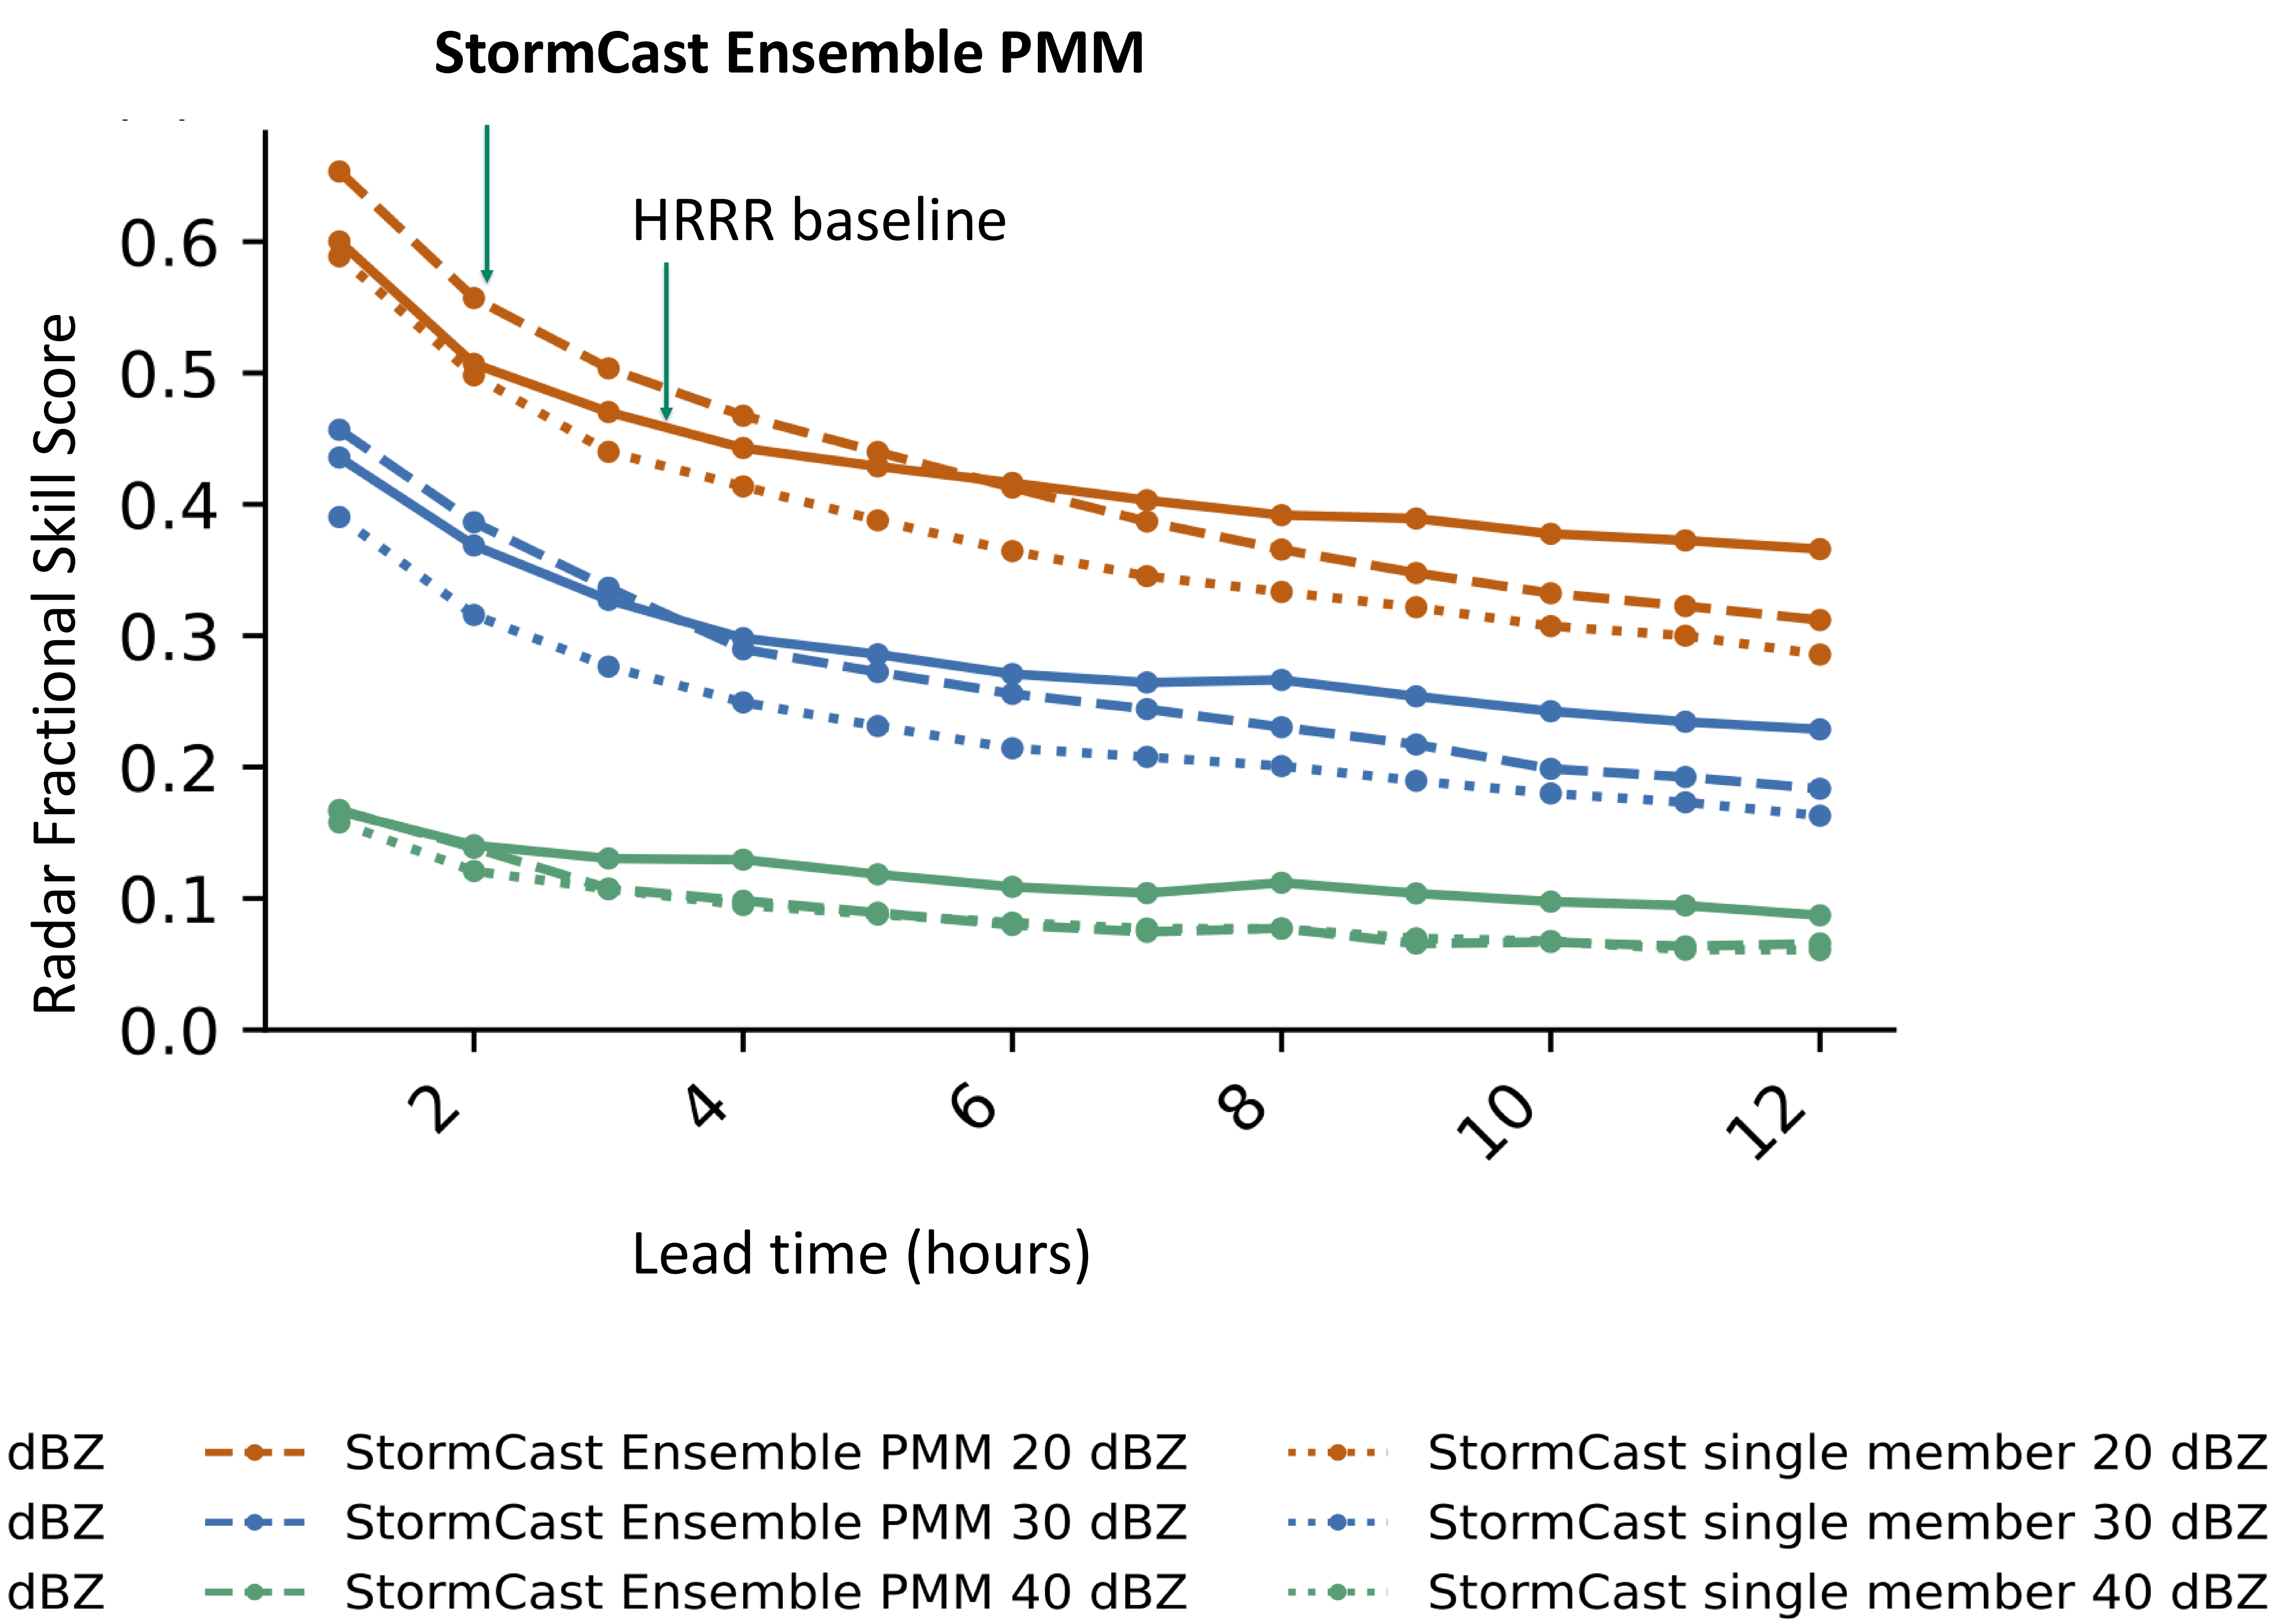

Building Level Resolution

Regional AI Weather Forecasting in the Taiwan

High-resolution regional forecasting helps mitigate **risks from severe rain, fog, dust storms, and heat** by protecting **infrastructure, transport, and public safety**.

Localised forecasts support **grid stability and renewable energy integration**.

Capturing winds at **building scale** enables hyper-local forecasting for **infrastructure protection, disaster response**, and **renewable energy integration**.

https://developer.nvidia.com/blog/spotlight-axa-explores-ai-driven-hurricane-risk-assessment-with-nvidia-earth-2/

Flood risk over and entire season

With JBA Risk

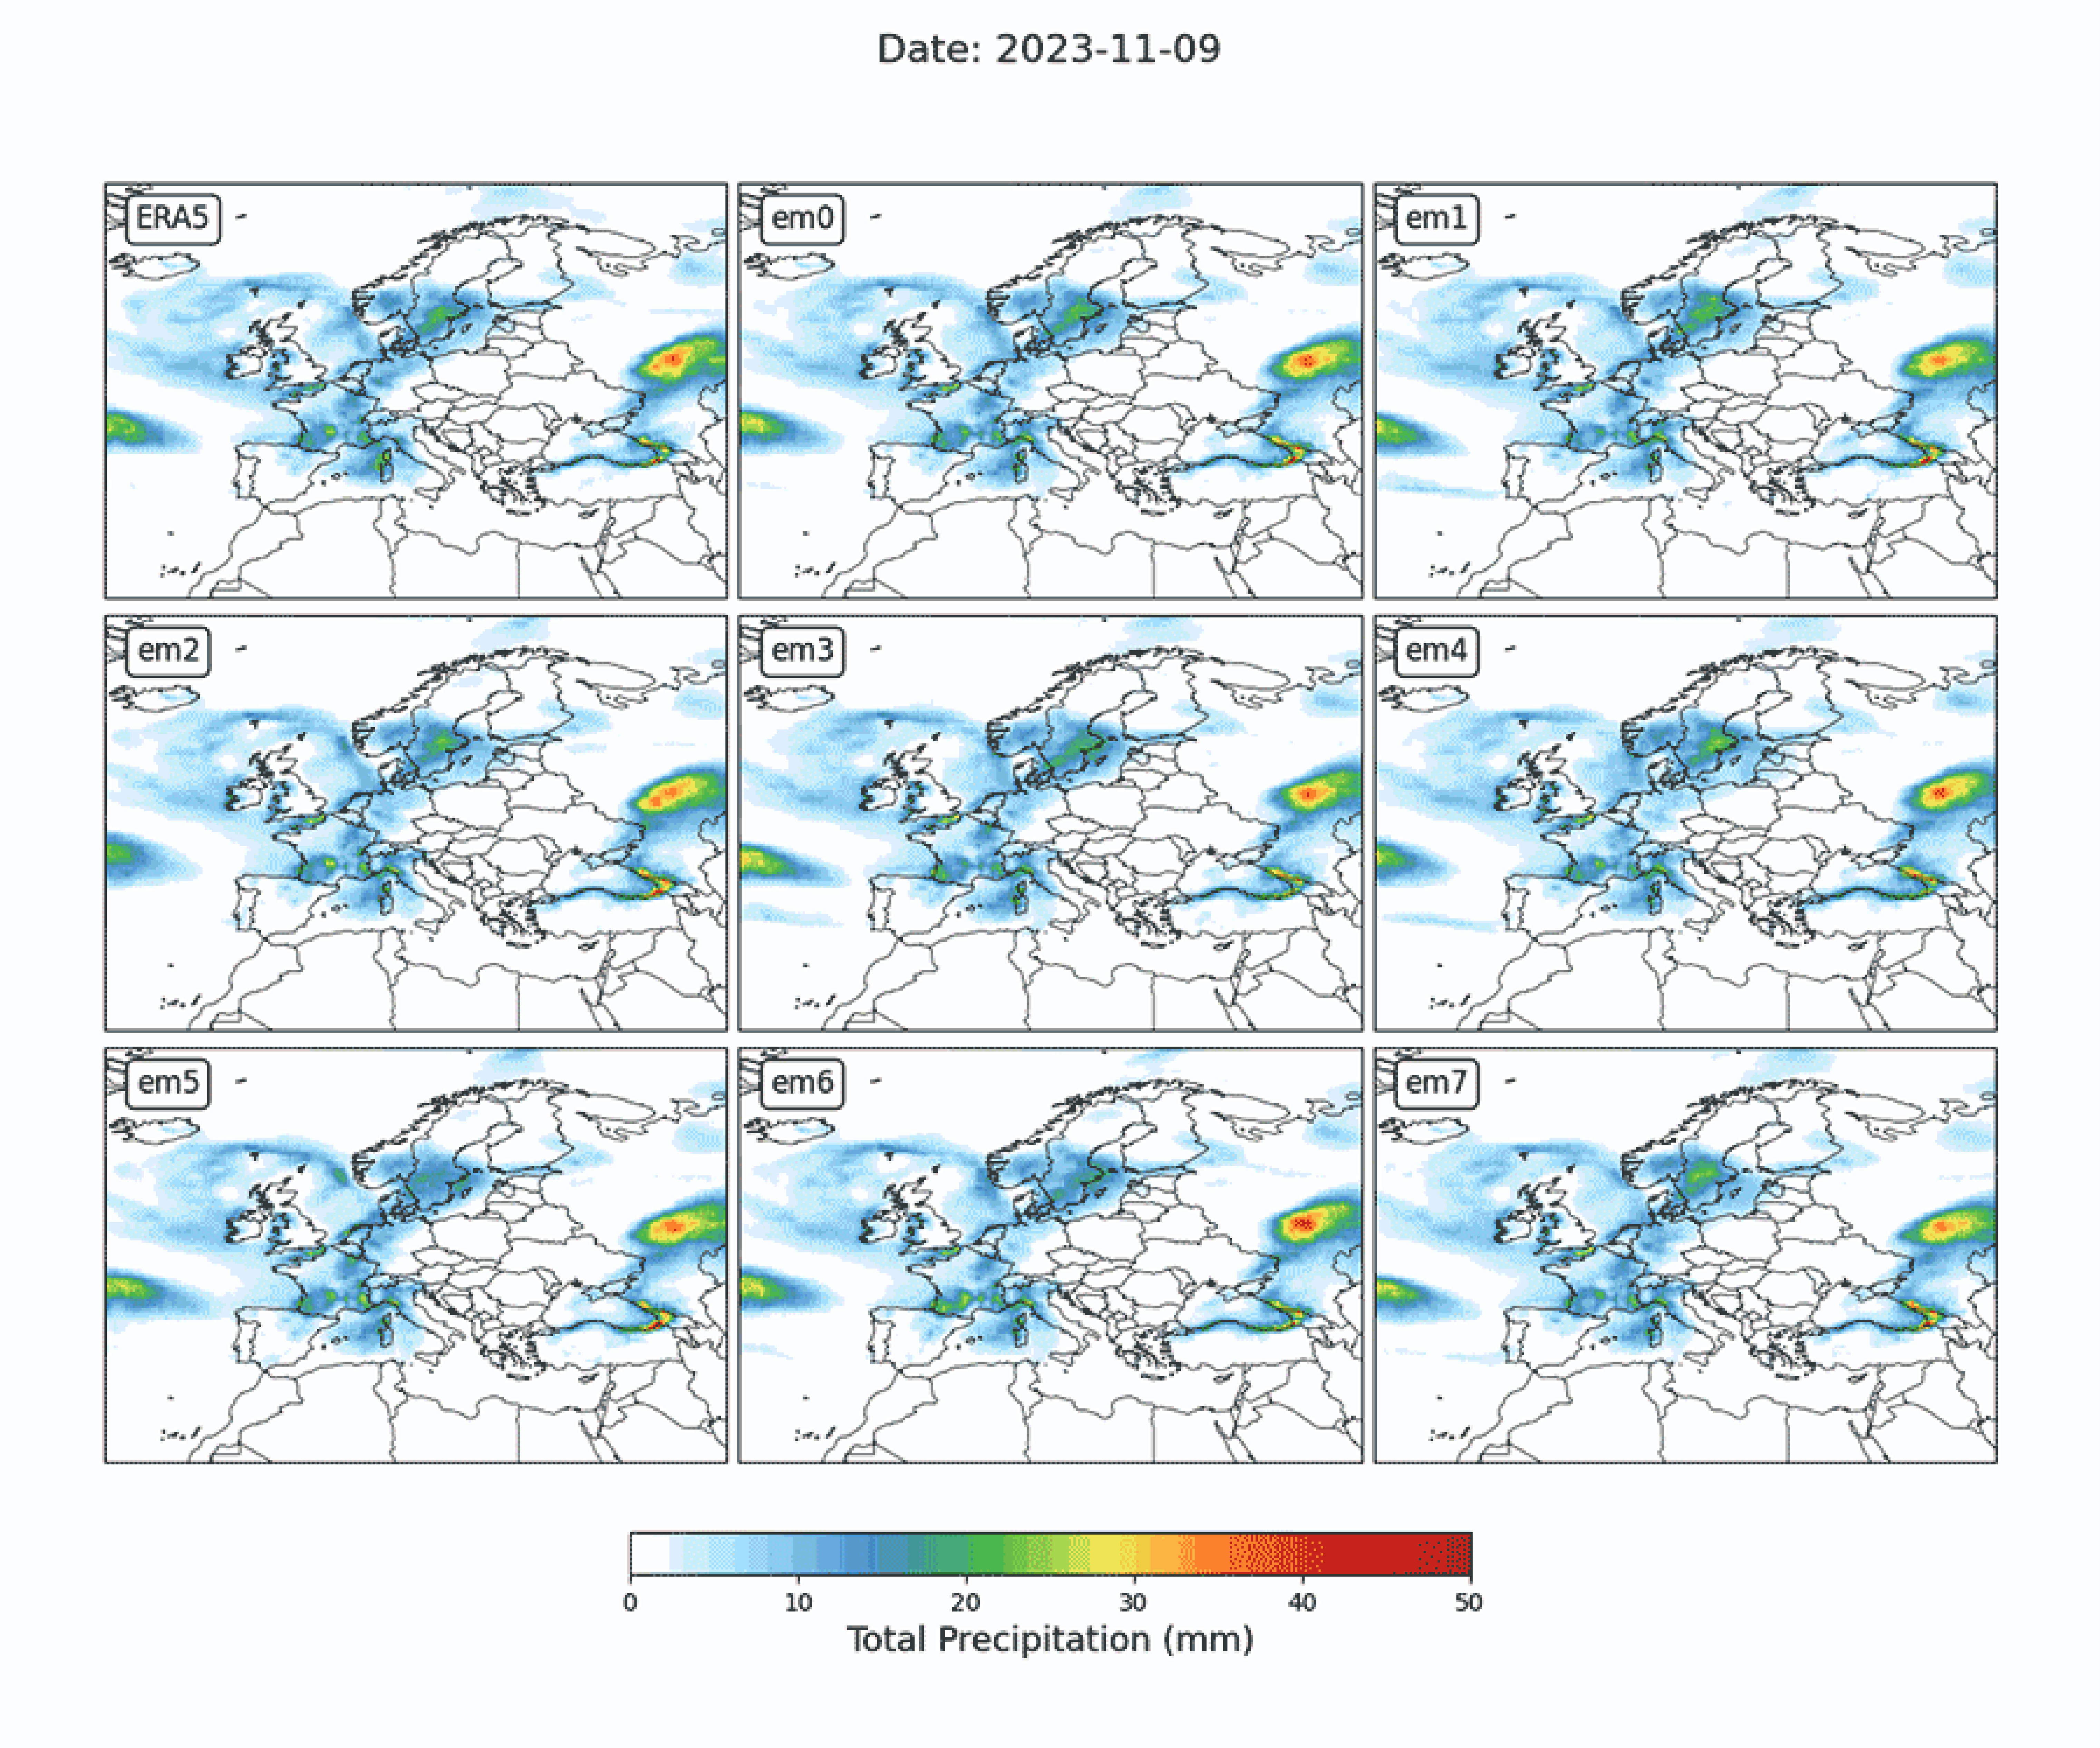

x 1008

Prediction of PV production with GCL

Solar radiation prediction
- 5,000x speedup
- 10,000x more energy efficient

PV power prediction
- 10% increase in accuracy
- 2 Bn ¥/year (~$300m)

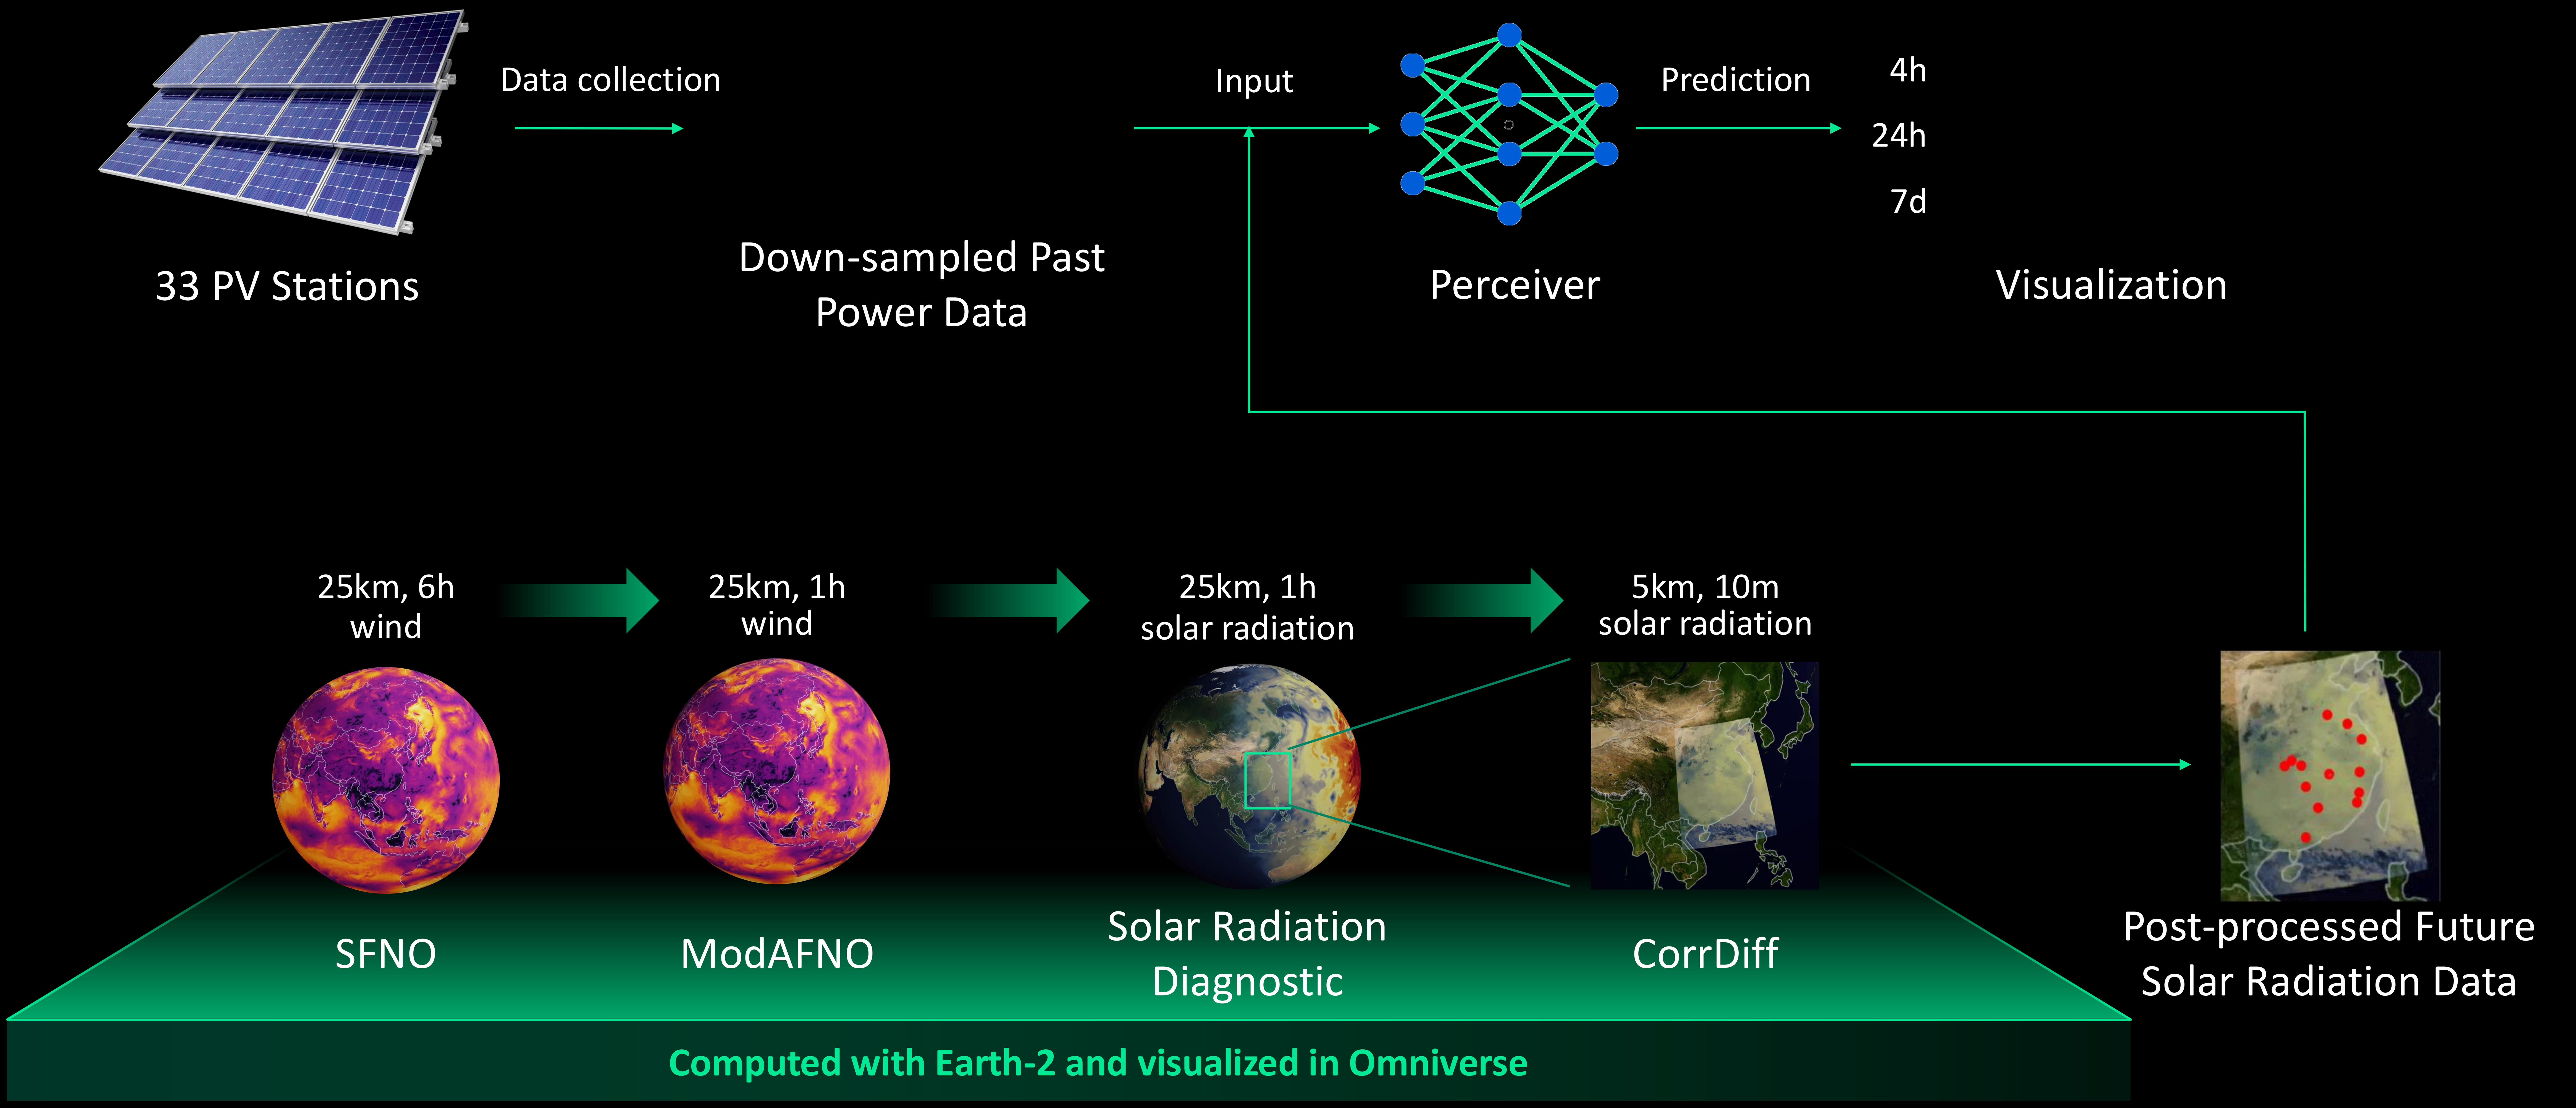

Load from Metadata File
/home/kpaul/Data/CMIP6_FC_UK/
.minikube
.ngc
.npm
.nv
.nvidia-omniverse
.nvm
.pki
.pytest_cache
.ssh
.vnc
.vscode
Data
2024_Greece_Fire
2024_GTC24_CorrDiff
2024_GTC24_ICON_BlueI
ClimateJust_UK
NFVI_UK
CMIP6_FC_UK
GridInfo
Desktop
Development
Documents
Downloads
Music
Notebooks
Pictures
CO
NO
SO2_SRF
TS
meta.json
File name:
Select
Cancel
Image Layers
Base Satellite
ClimateJust: NFVI (UK)
Misc
Continents Outline
Lights
Sun
Headlight
Ambient Light
Atmosphere
09:28
Mar 17, 2011
MANCHESTER 1824
The University of Manchester
NHS

NVIDIA

Climate Emulators Are Being Built On AI Weather Prediction Models

“ACE” by The Allen Institute for AI Builds off NVIDIA FourCastNet

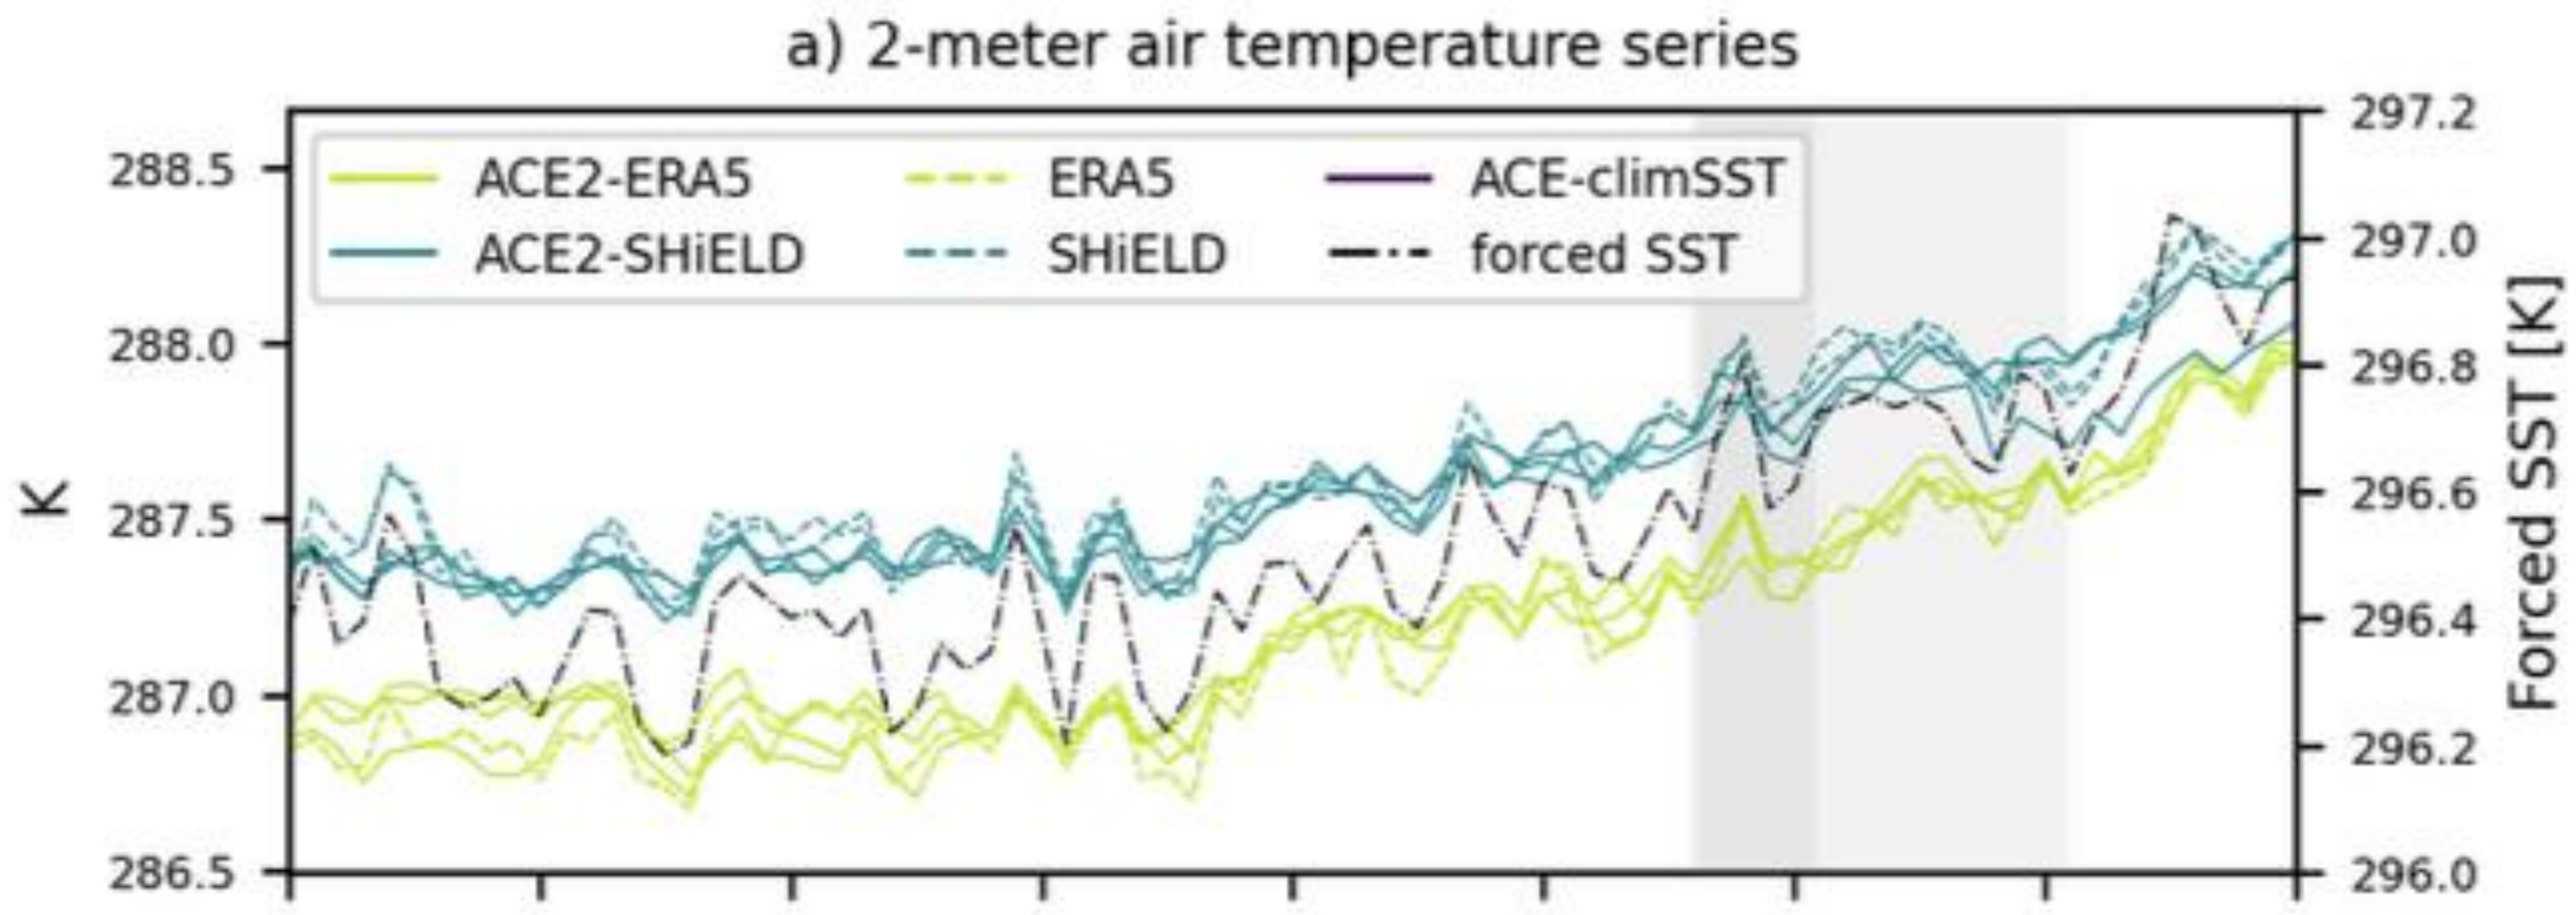

"ACE2": Watt-Meyer et al., 2024
ArXiv: 2411.11268

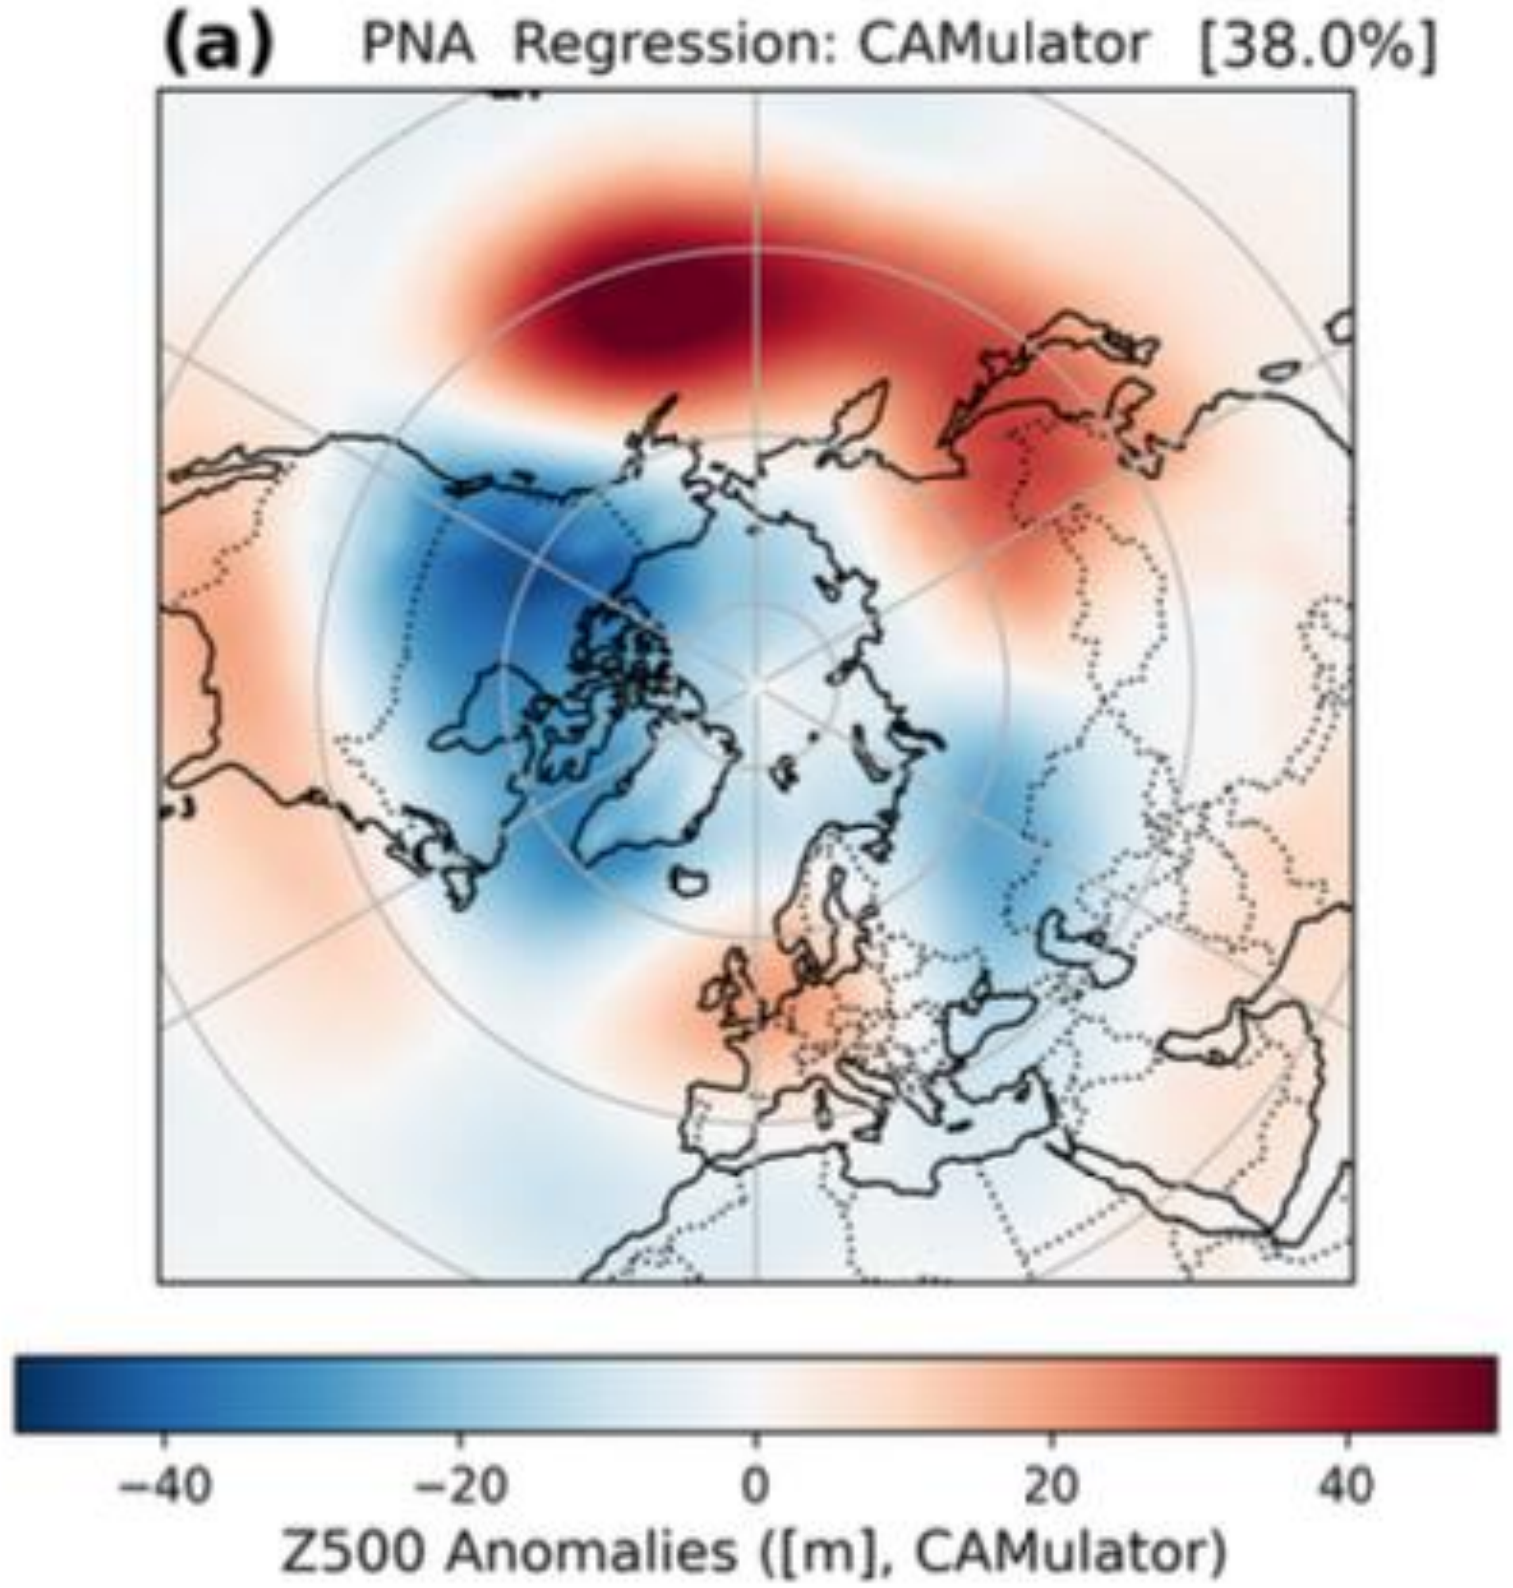

"CAMulator": Chapman et al., 2025
ArXiv: 2504.06007

A Frontier: Multi-Component AutoRegressive AI Earth System Modeling

Linking AI Earth System component models

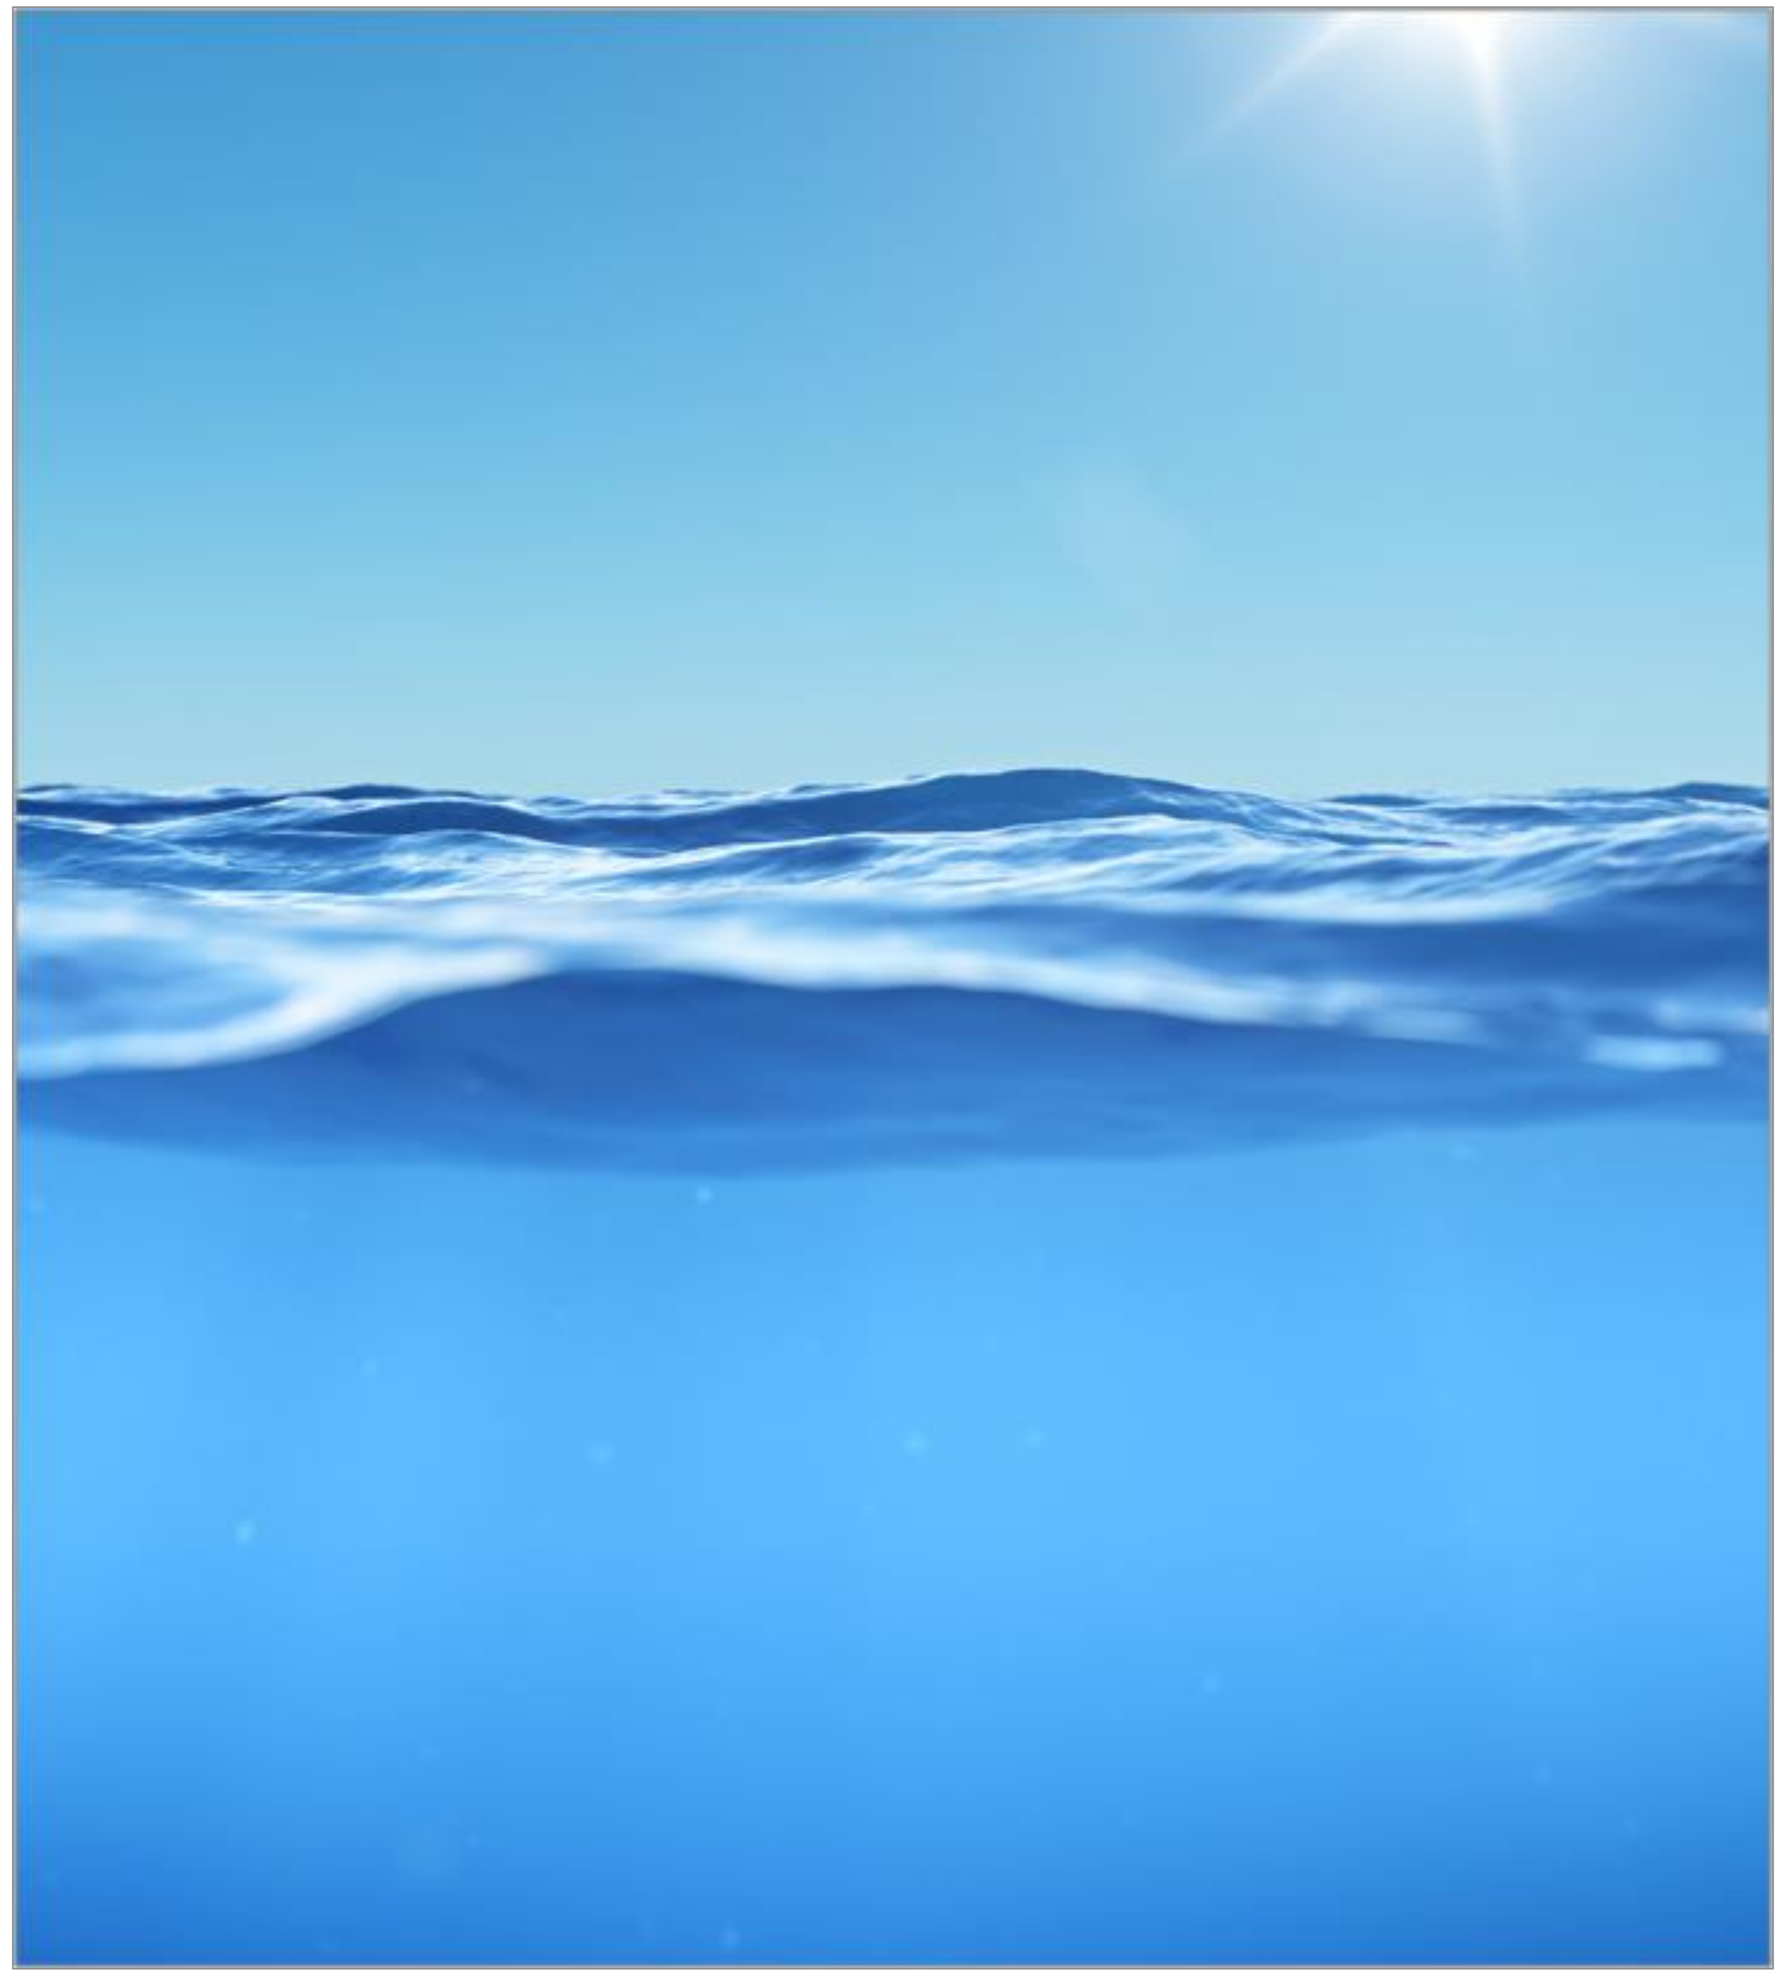

Ocean

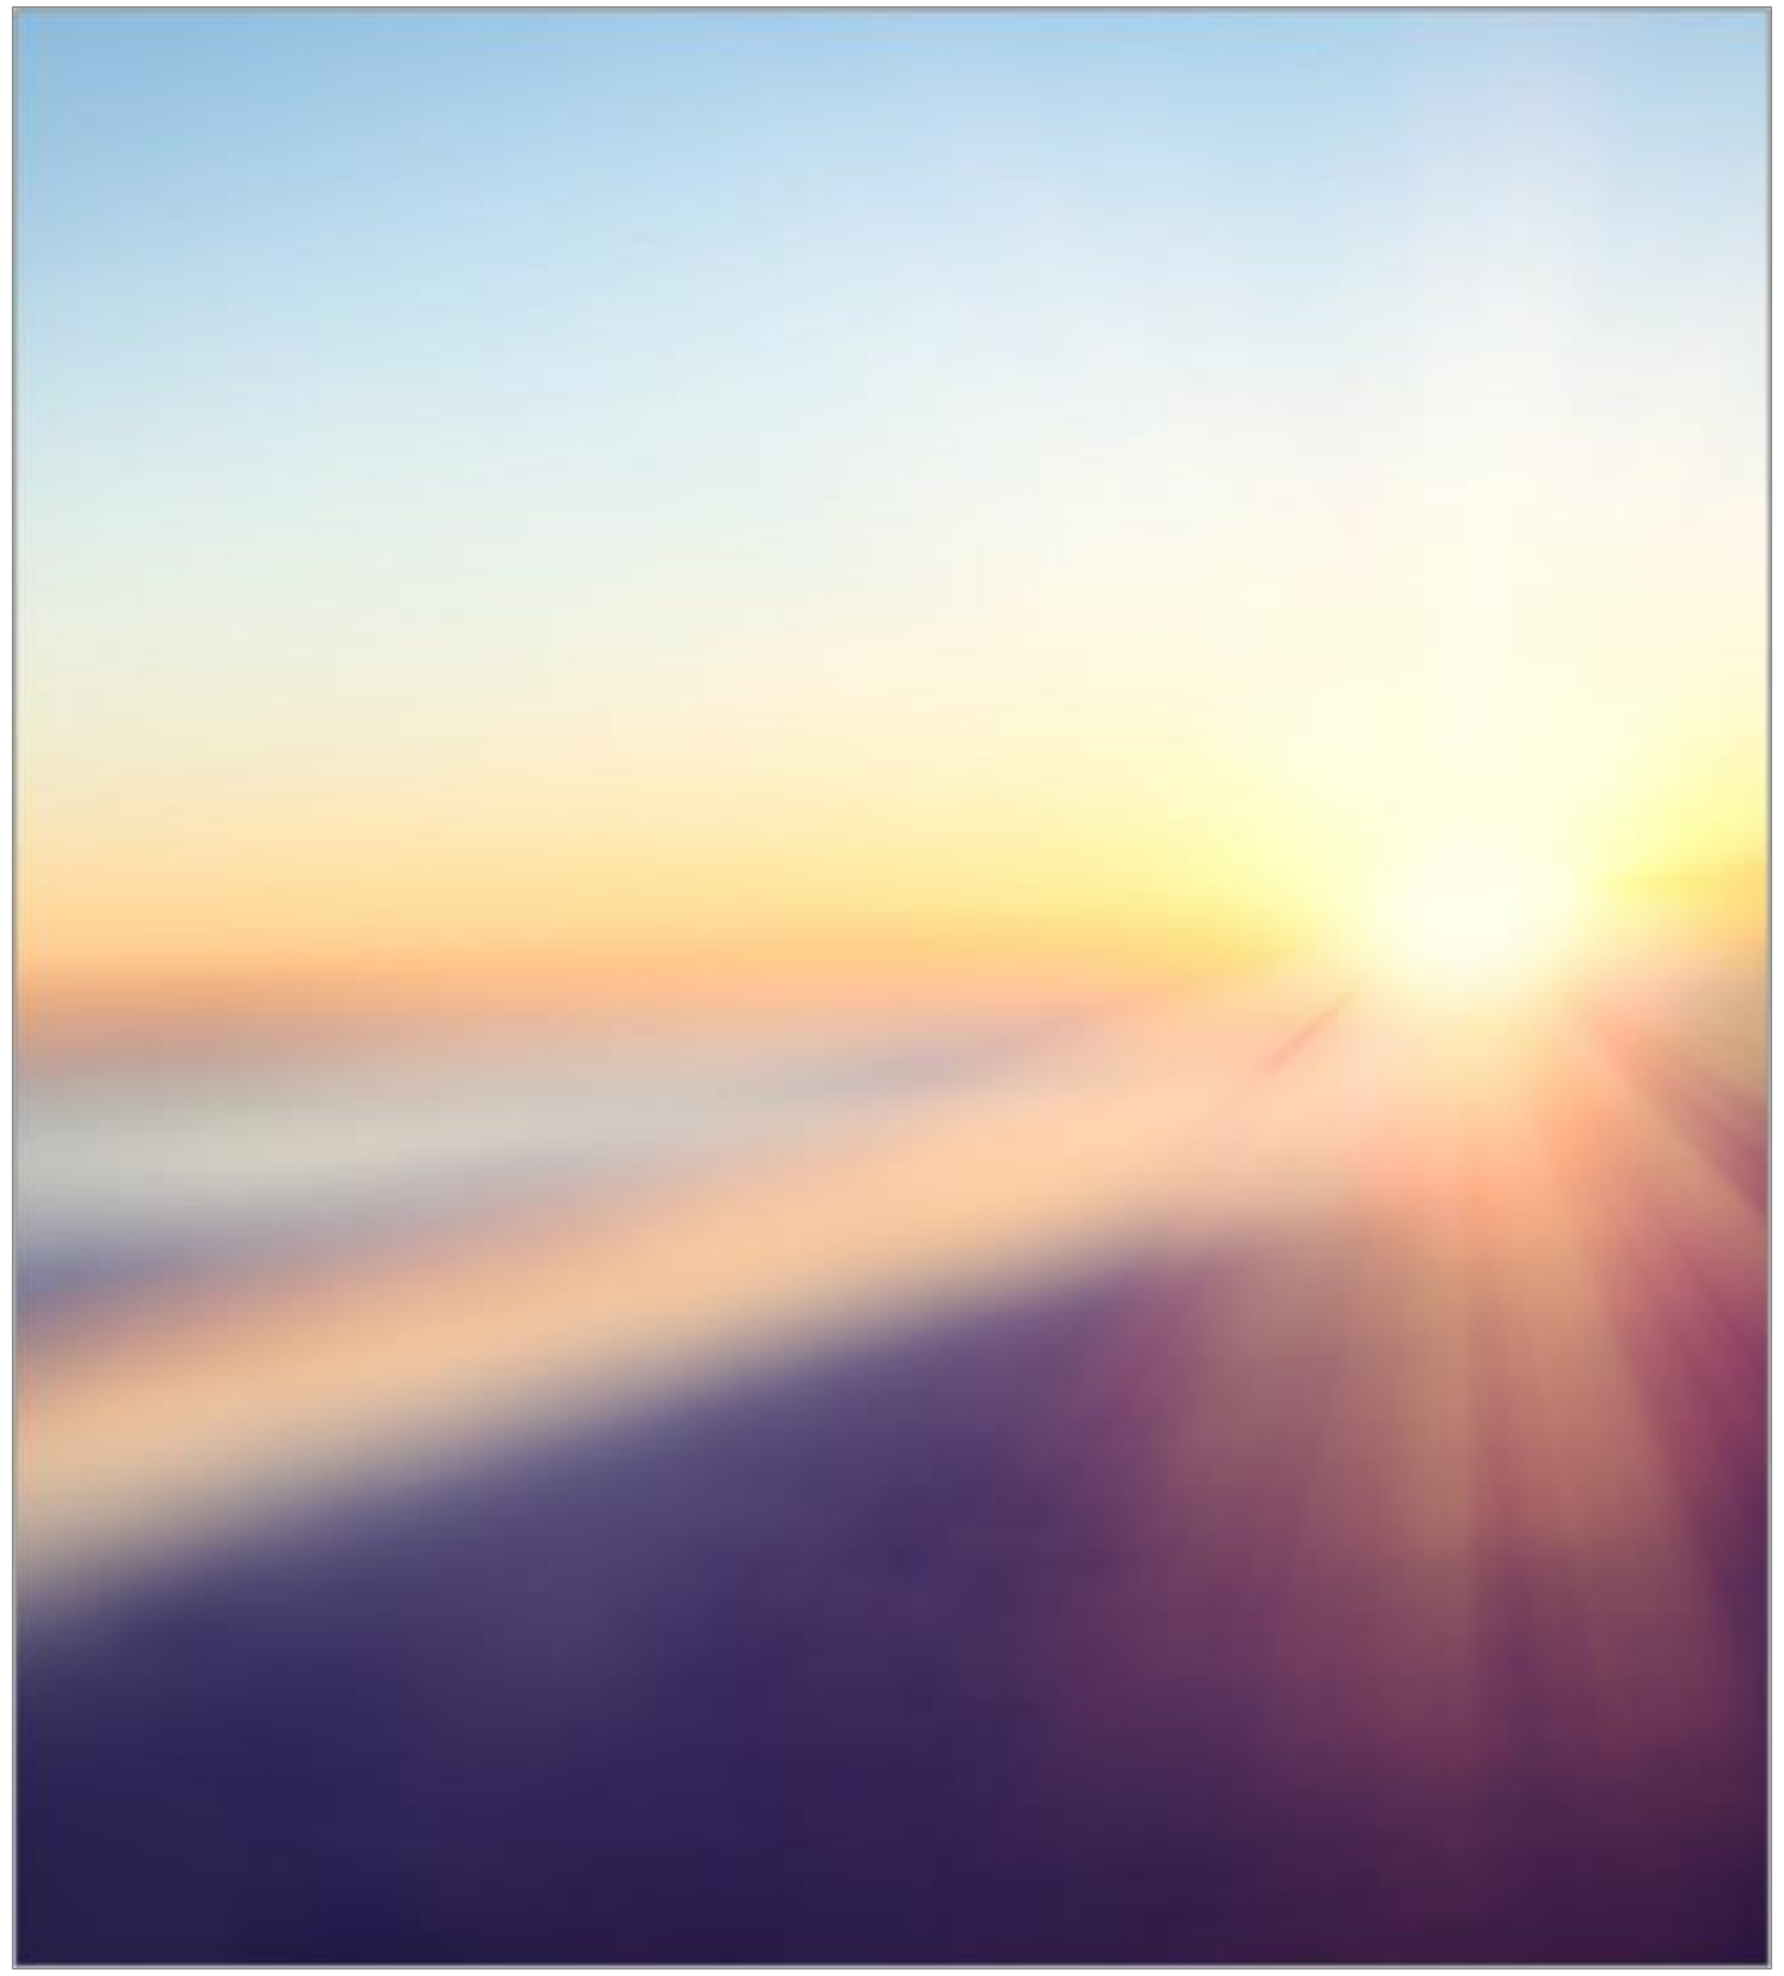

Atmosphere

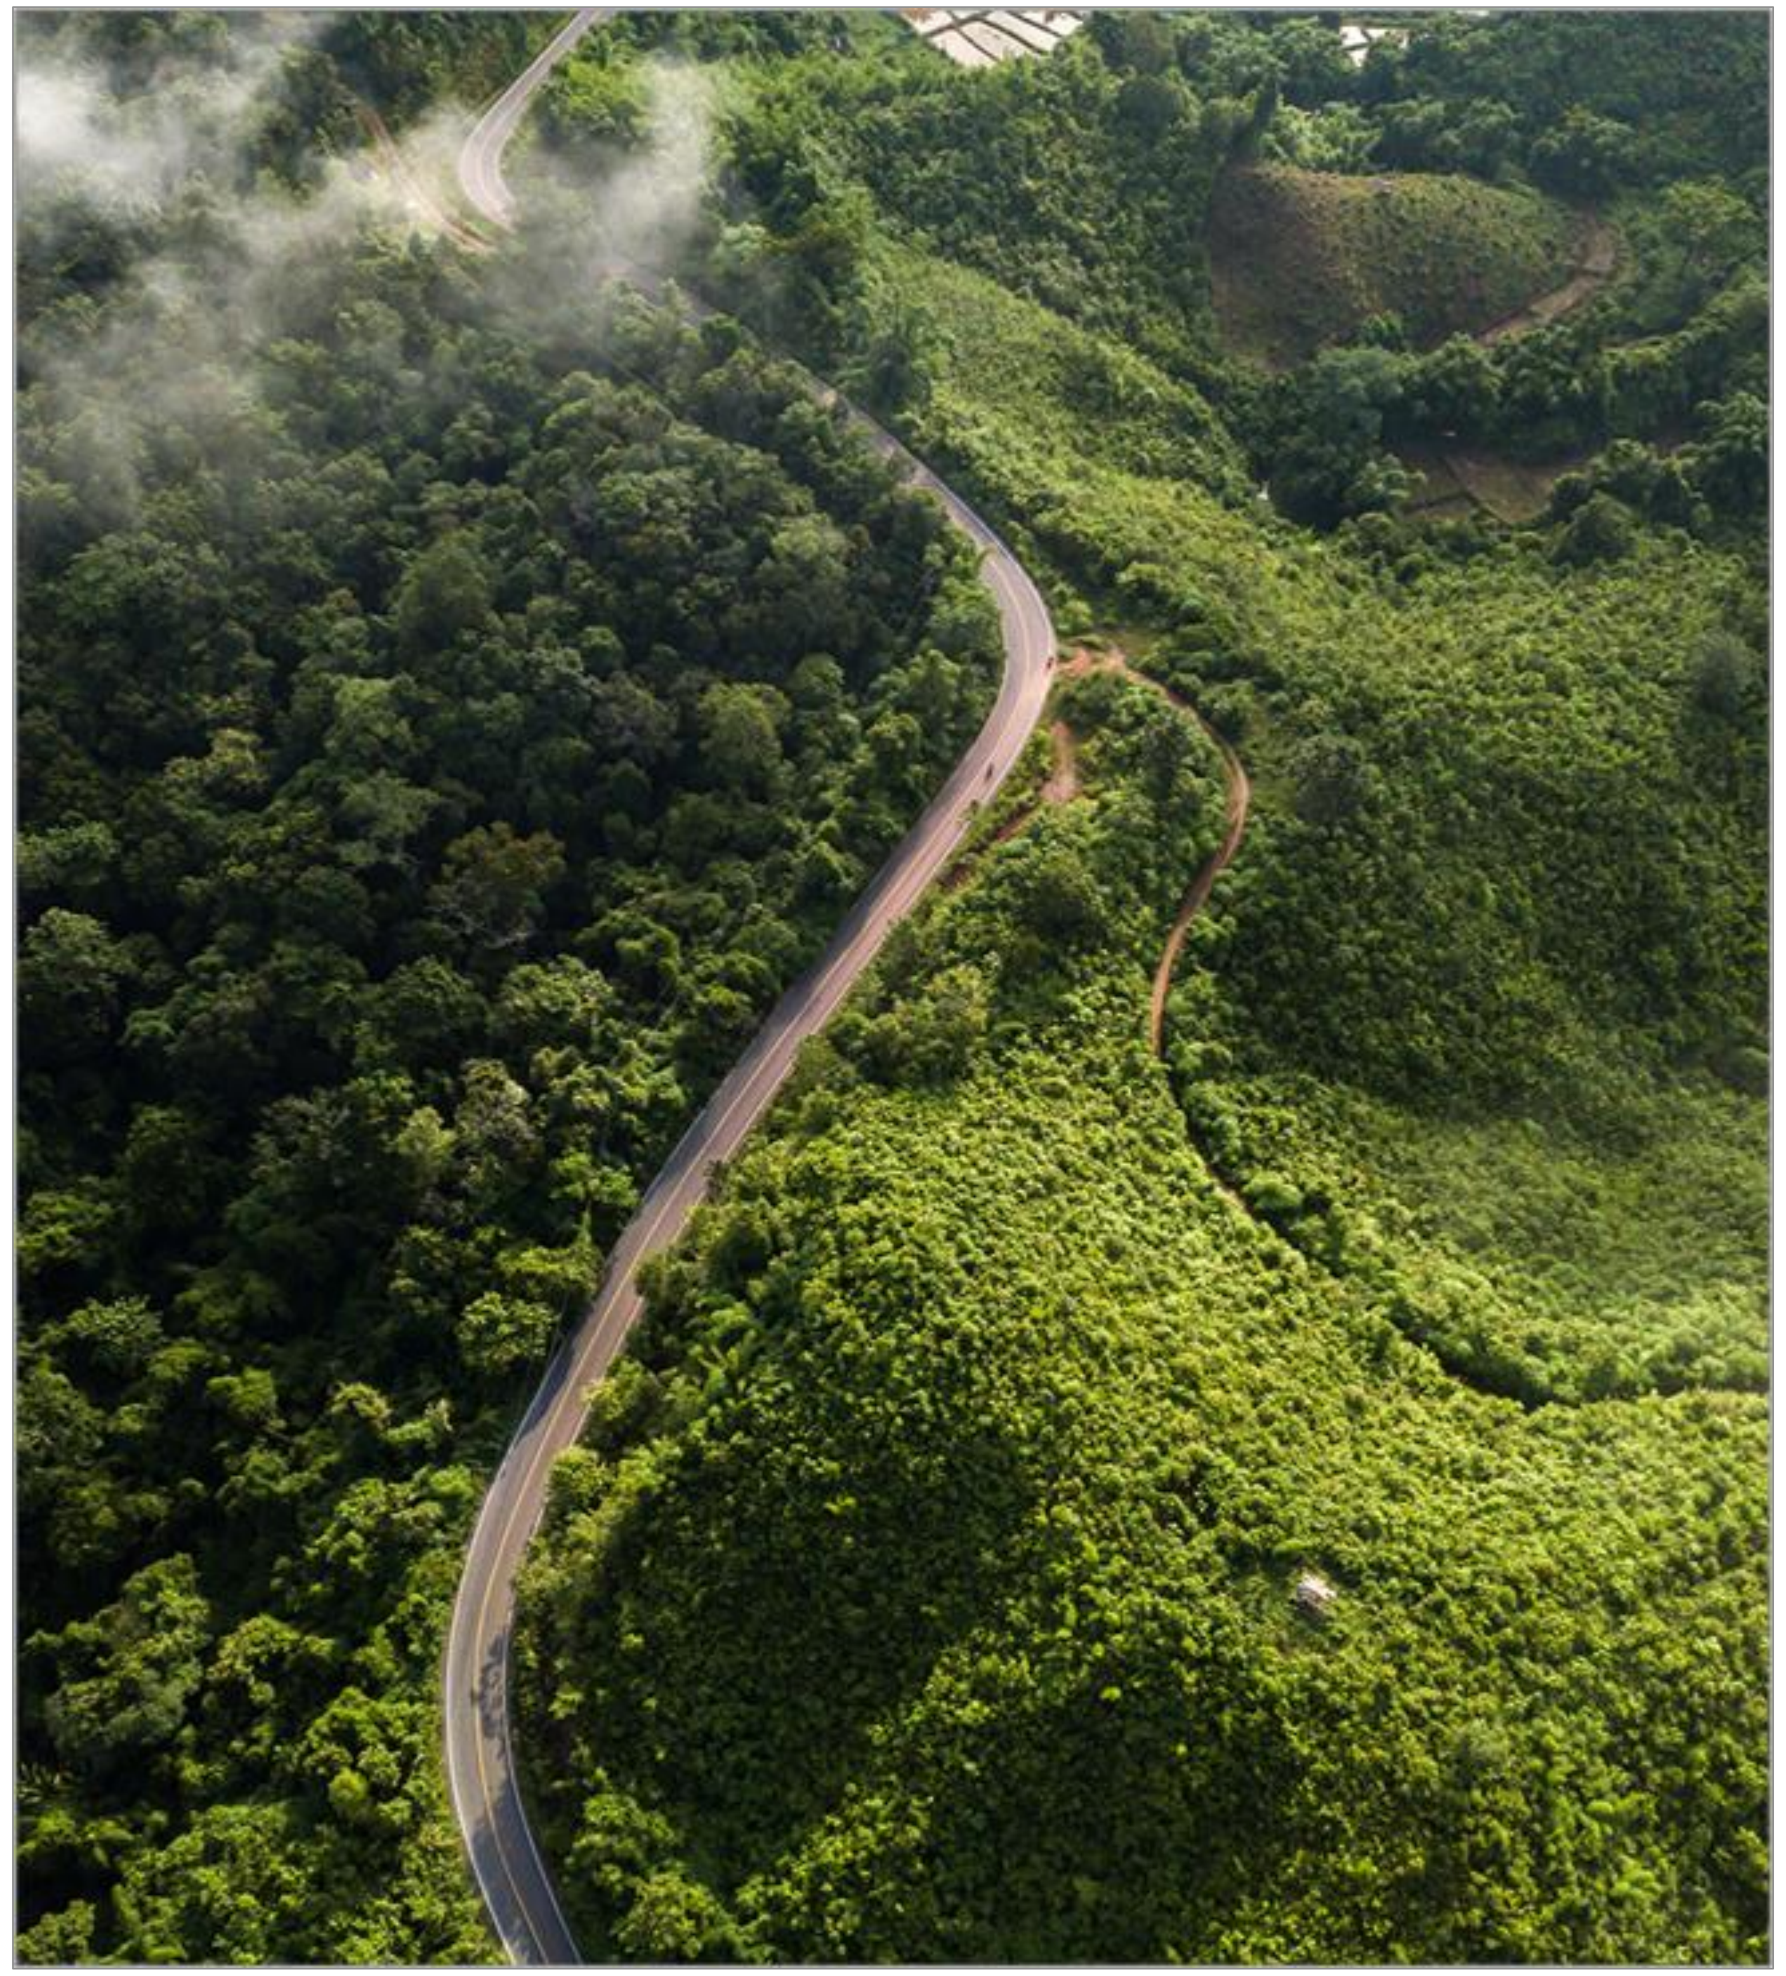

Land

Beyond Atmospheric Forecasting, Other Challenges in Weather and Climate

ReGen Generates Full States from Sparse Observations

Data source: NOAA weather stations. See Manshausen et al, arxiv.org/abs/2406.16947

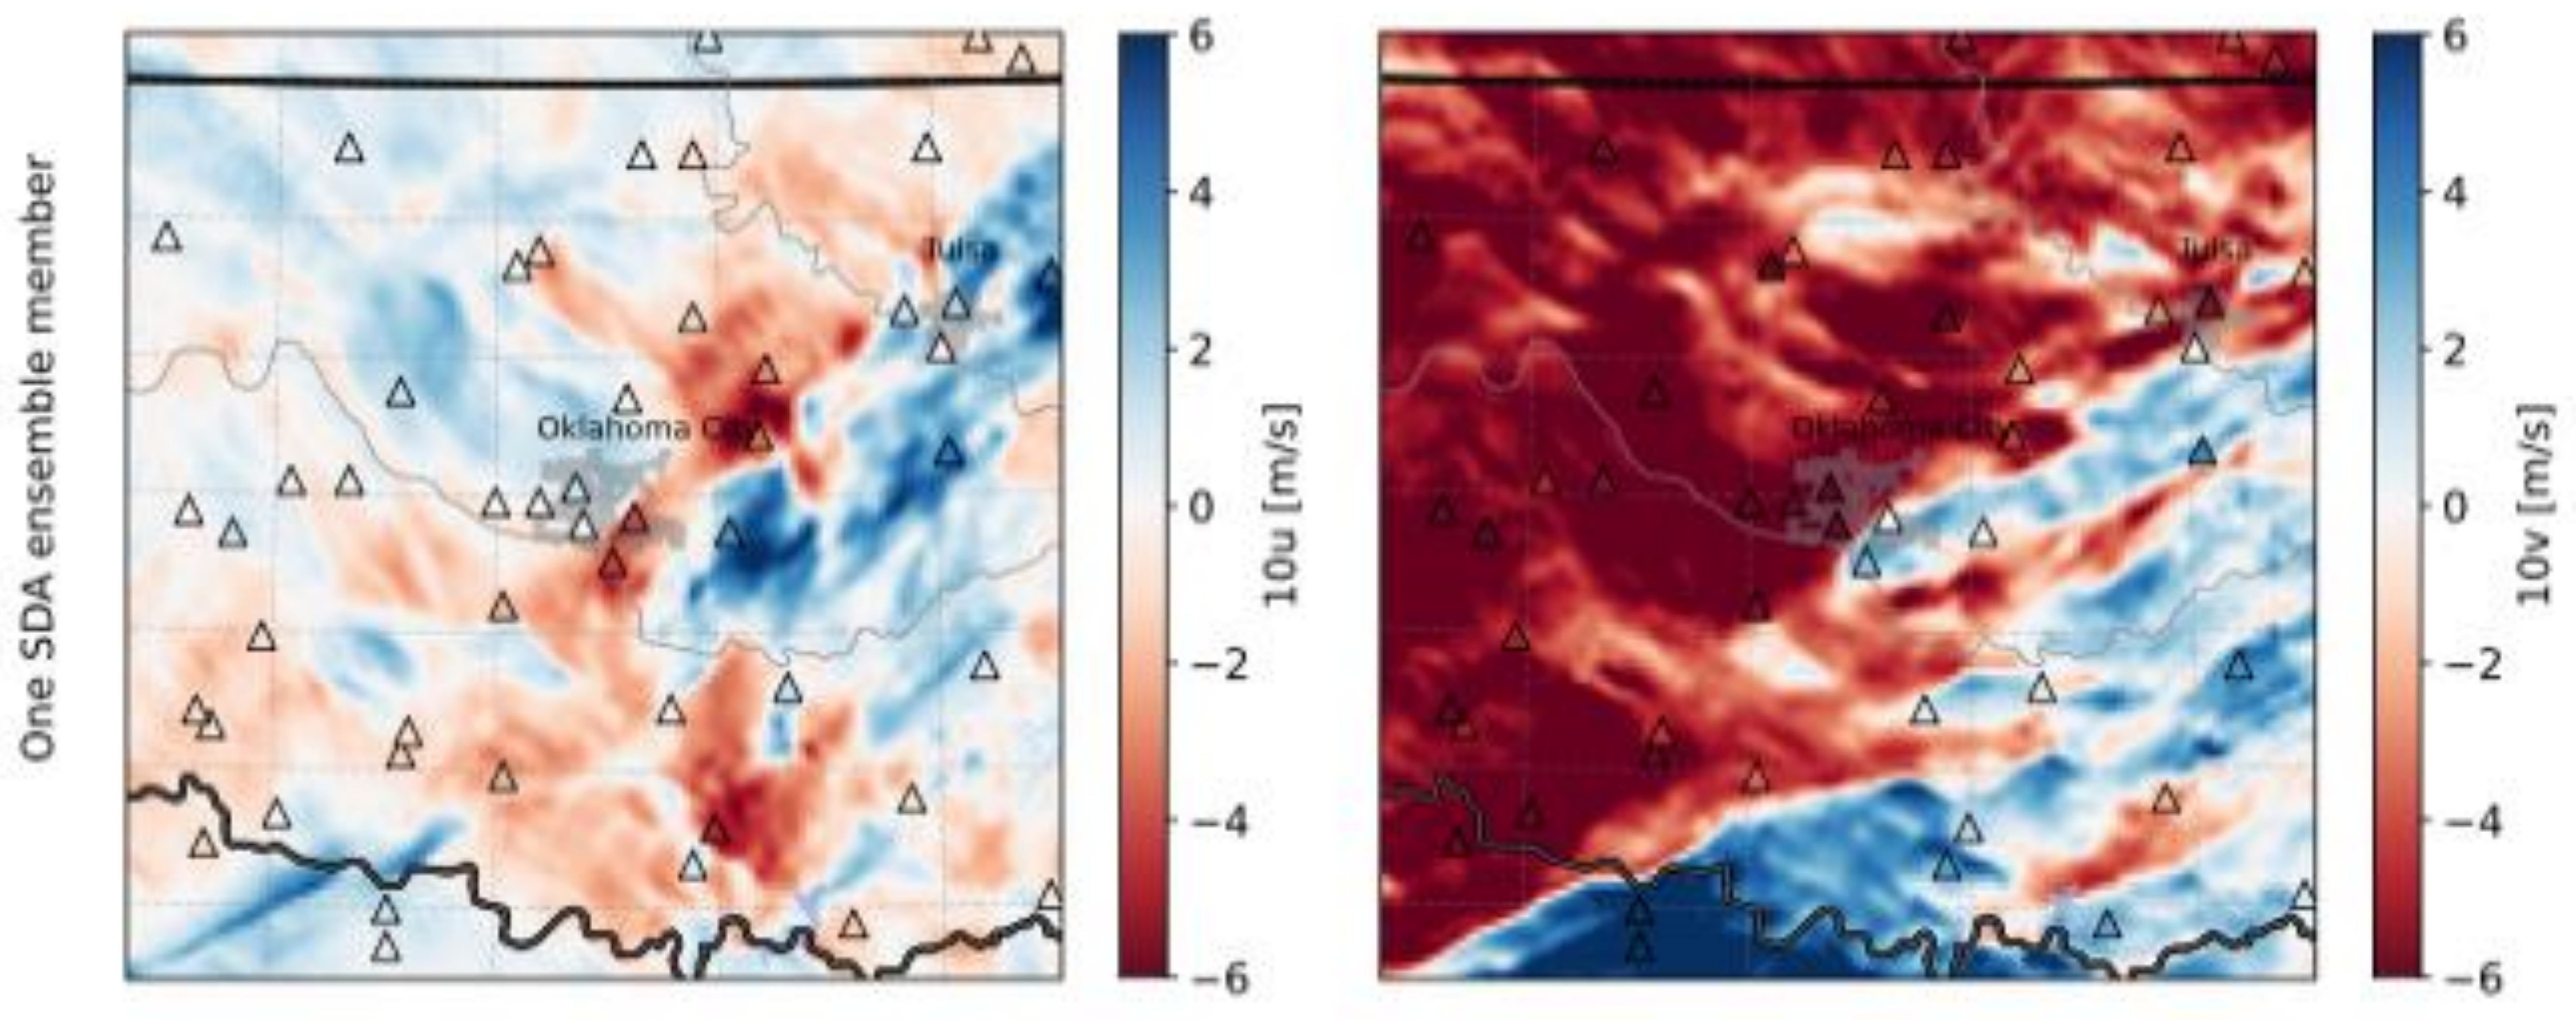

Spherical Fourier Neural Operators for Coupled Ocean-Atmosphere Prediction

"Ocean-Linked Atmosphere": SST, SSH & sub-surface theta. See Wang & Pathak et al, arXiv:2406.08632

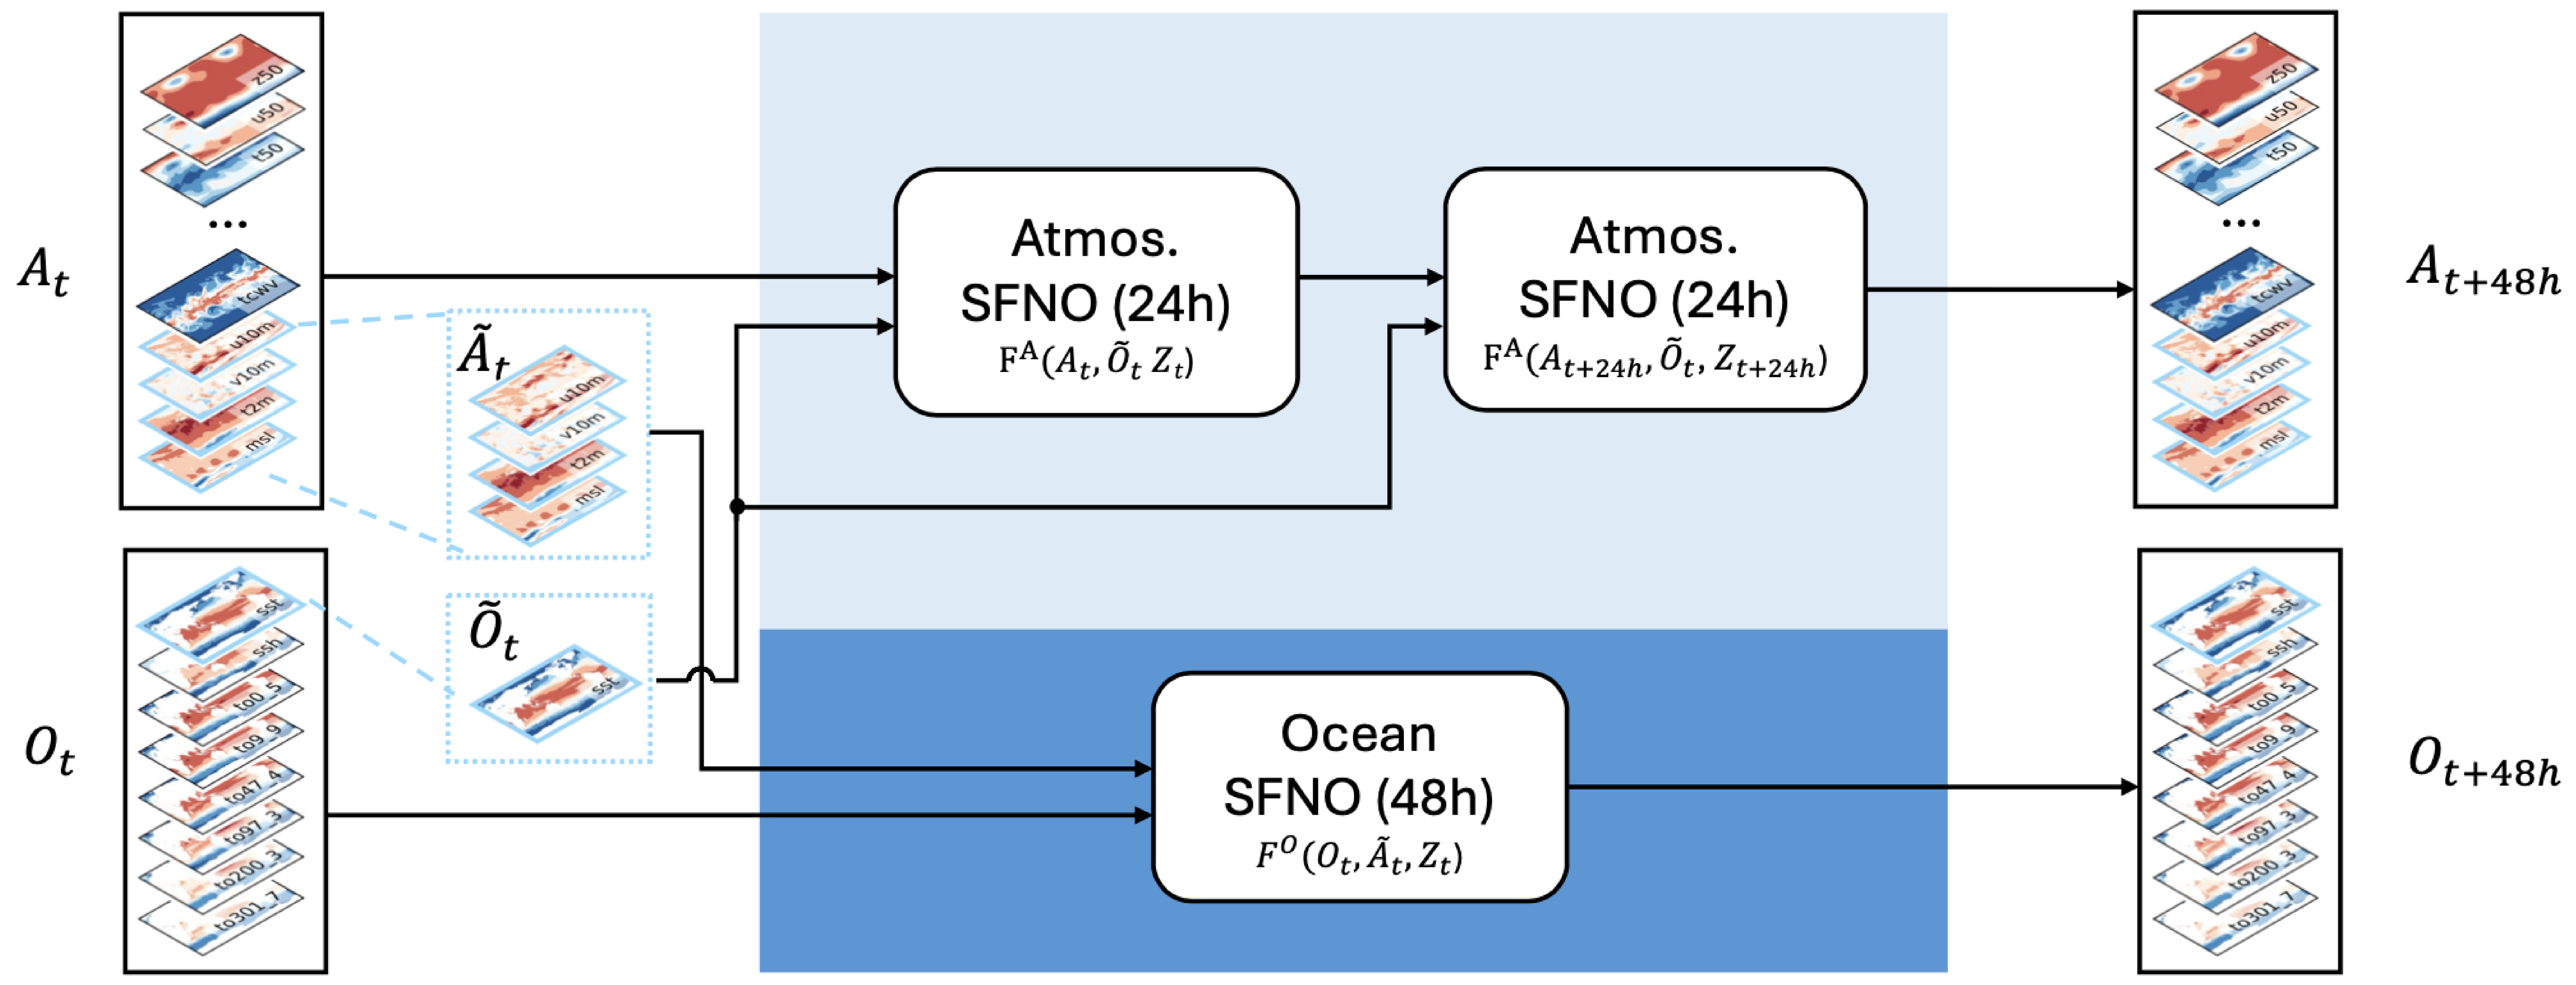

Atmosphere-Ocean Coupled Architectures Stable for Centuries Exist

Realistic midlatitude cyclones & seasonal modulation of tropical cyclogenesis 100-years into roll-out

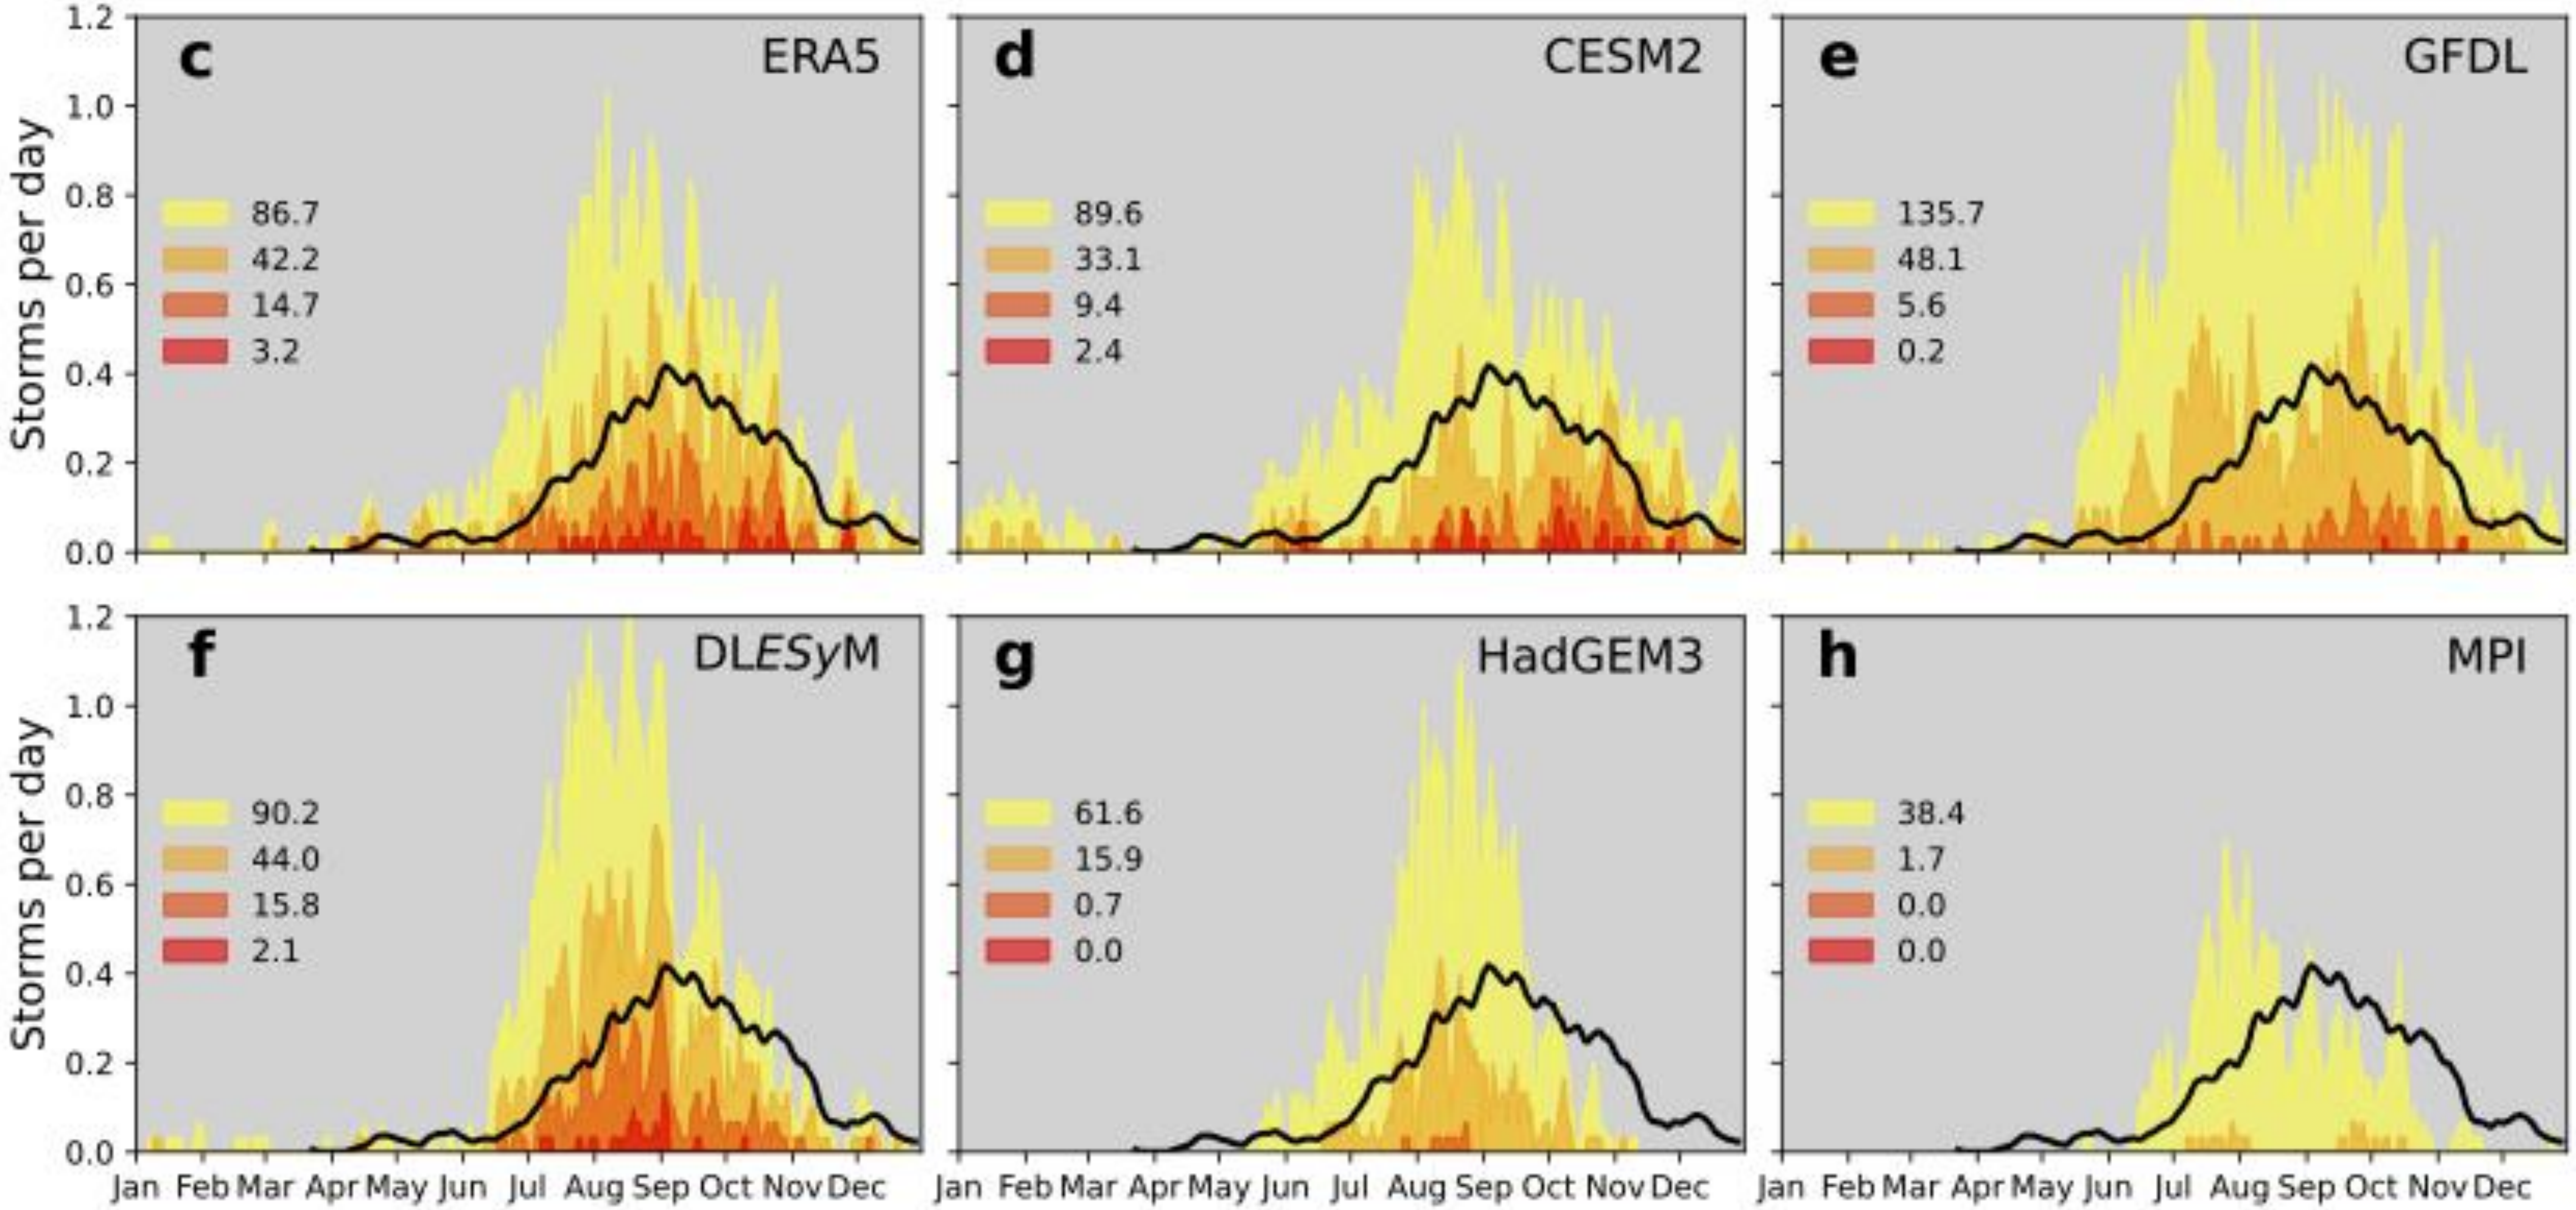

A Deep Learning Earth System Model for Stable and Efficient Simulation of the Current Climate

Nathaniel Cresswell-Clay, Bowen Liu, Dale Durran, Andy Liu, Zachary I. Espinosa, Raul Moreno, Matthias Karlbauer

https://arxiv.org/abs/2409.16247

UNIVERSITY of WASHINGTON

Aside: SubSeasonal-To-Seasonal (S2S) Ocean-Atmosphere Skill Is Here

DLESyM: Competitive with IFS ENS Physics Baseline

Lower is better

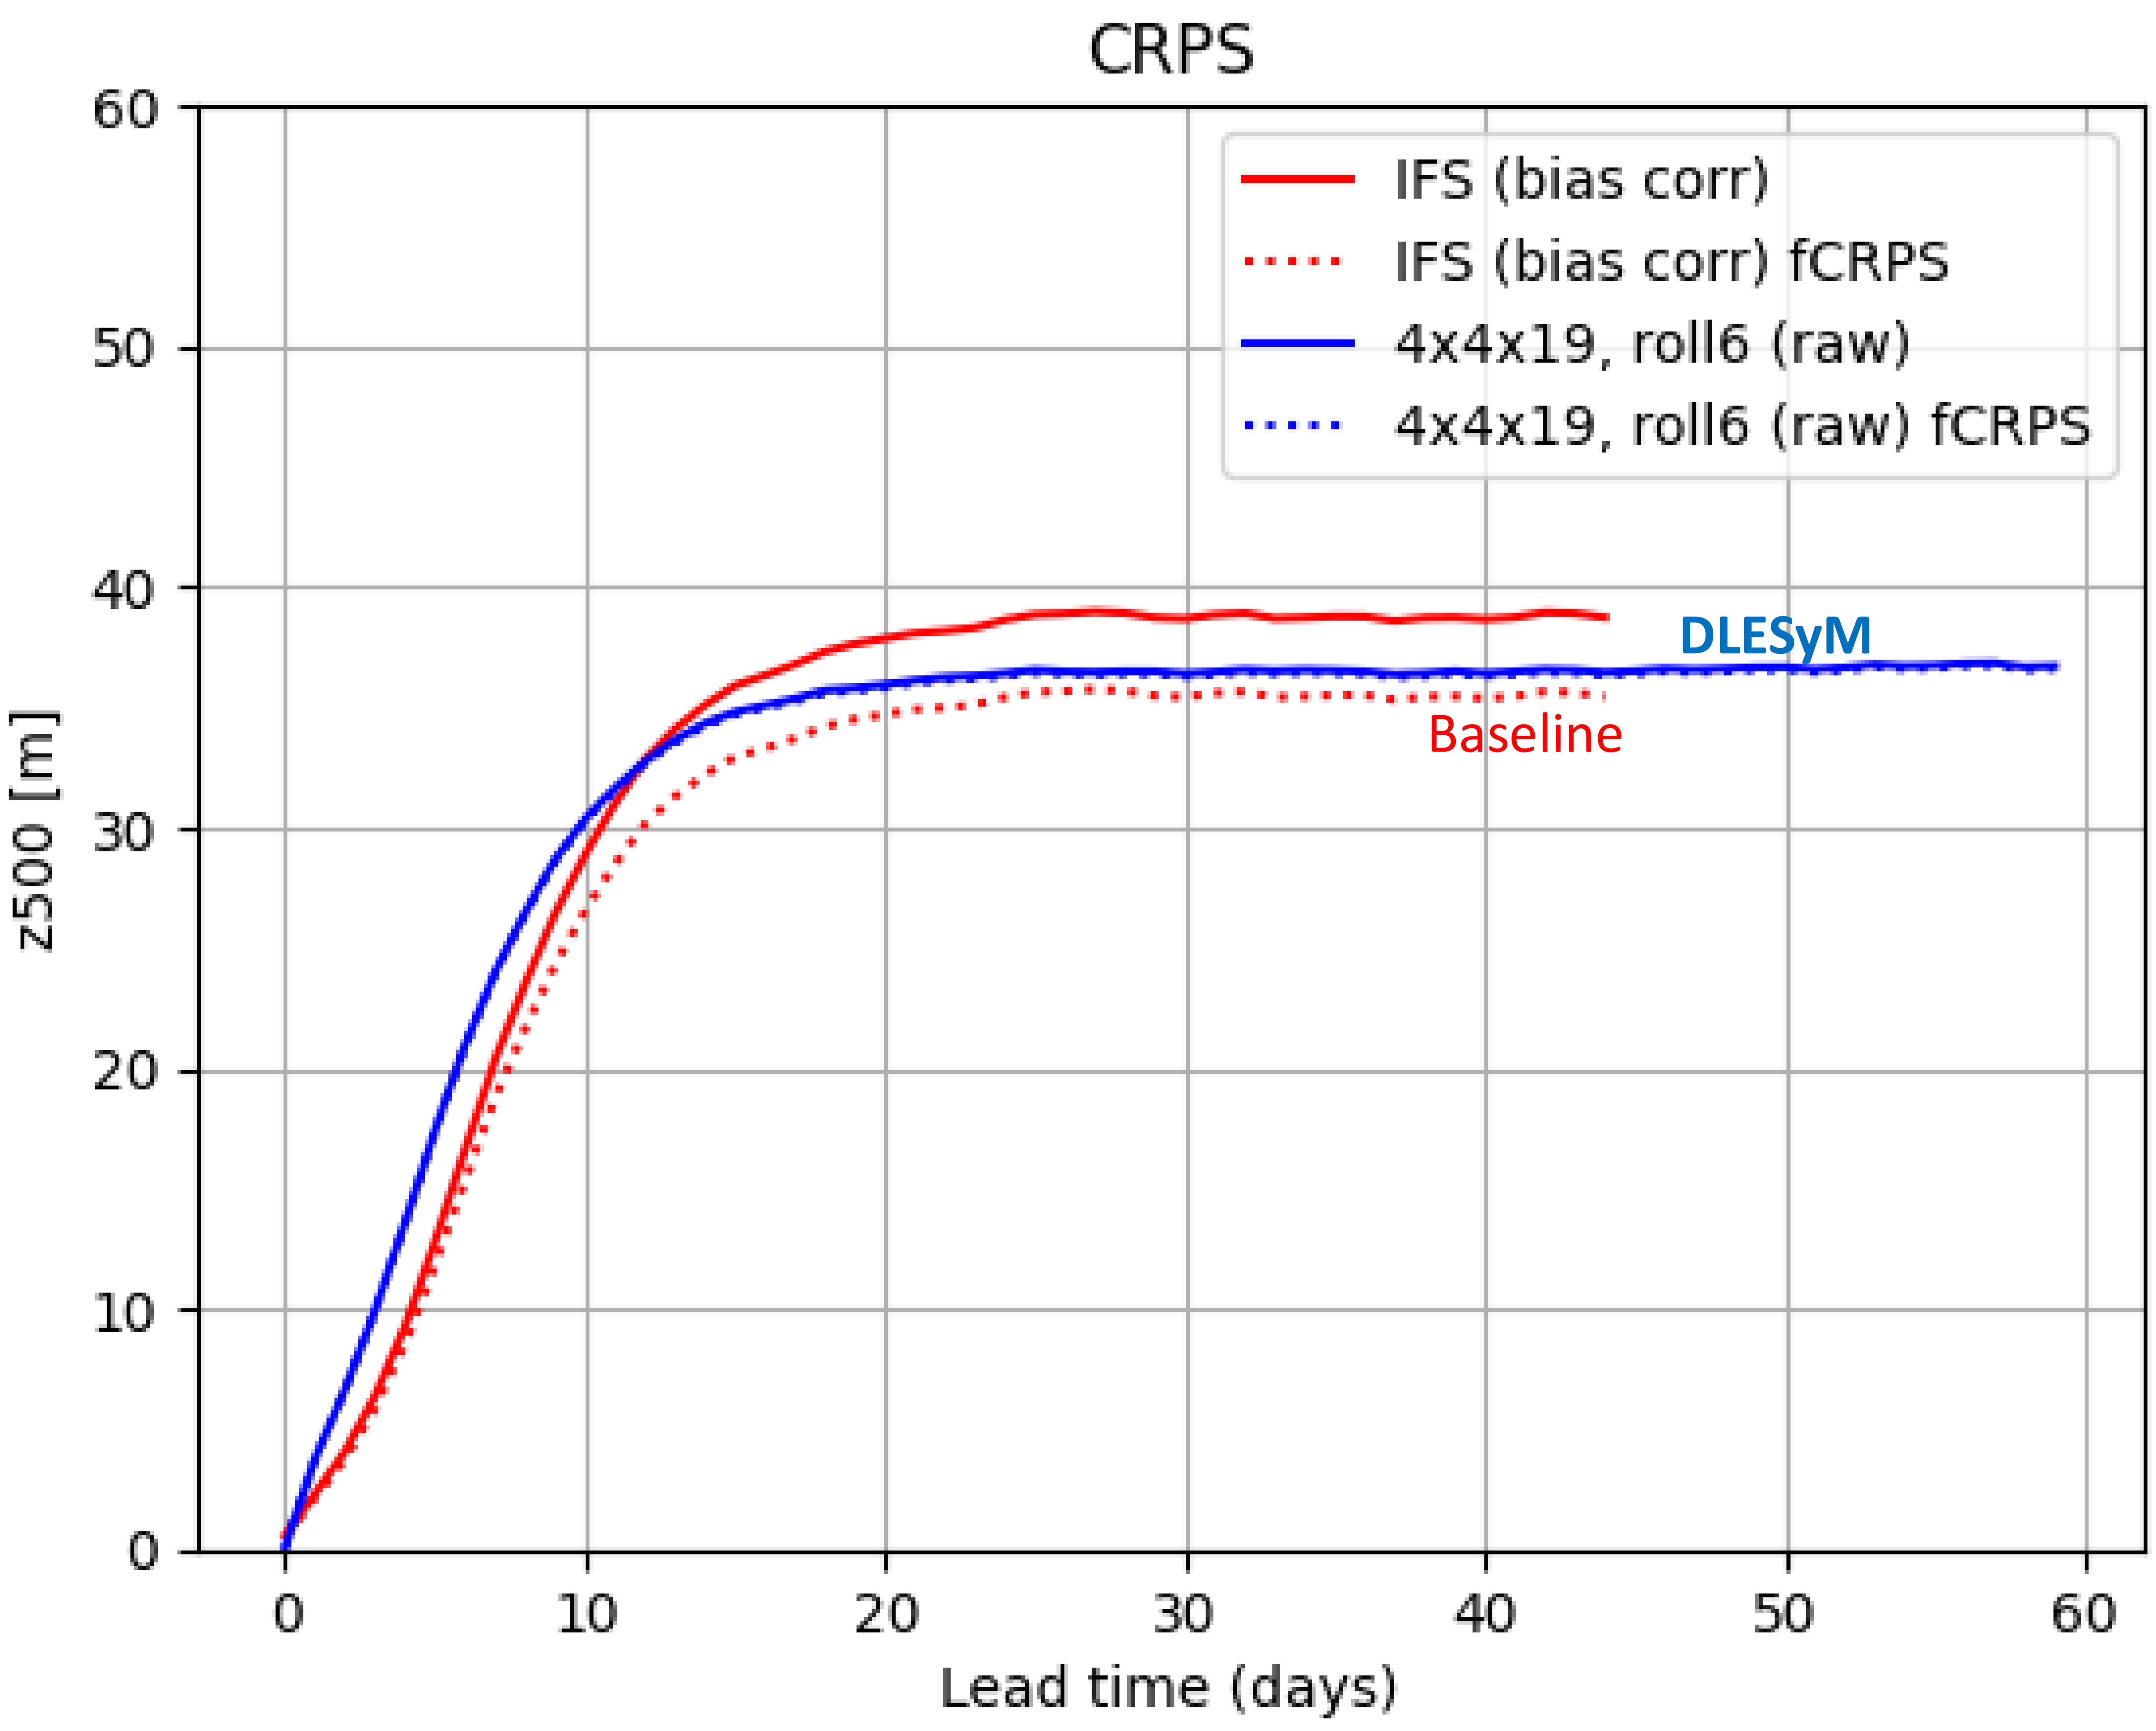

Climate in a Bottle

Problem: Intractable
High-Resolution Data

- Climate simulations generate **petabytes** of data
- Only major supercomputing centers can store / access it
- Interactive analysis is nearly impossible

Solution:
Fast AI data-generator

- Compress massive climate datasets into a small AI model
- Generate realistic km-scale climate data on demand

Climate emulators and extreme events

How emulators can help

- More samples
- Better samples
- Event attribution
- Enhance existing archives of climate data (CMIP):
    - Bias-correction
    - Downscaling/super-resolution
    - In-filling

The Pacific NW Heatwave of 2021, *Nature*

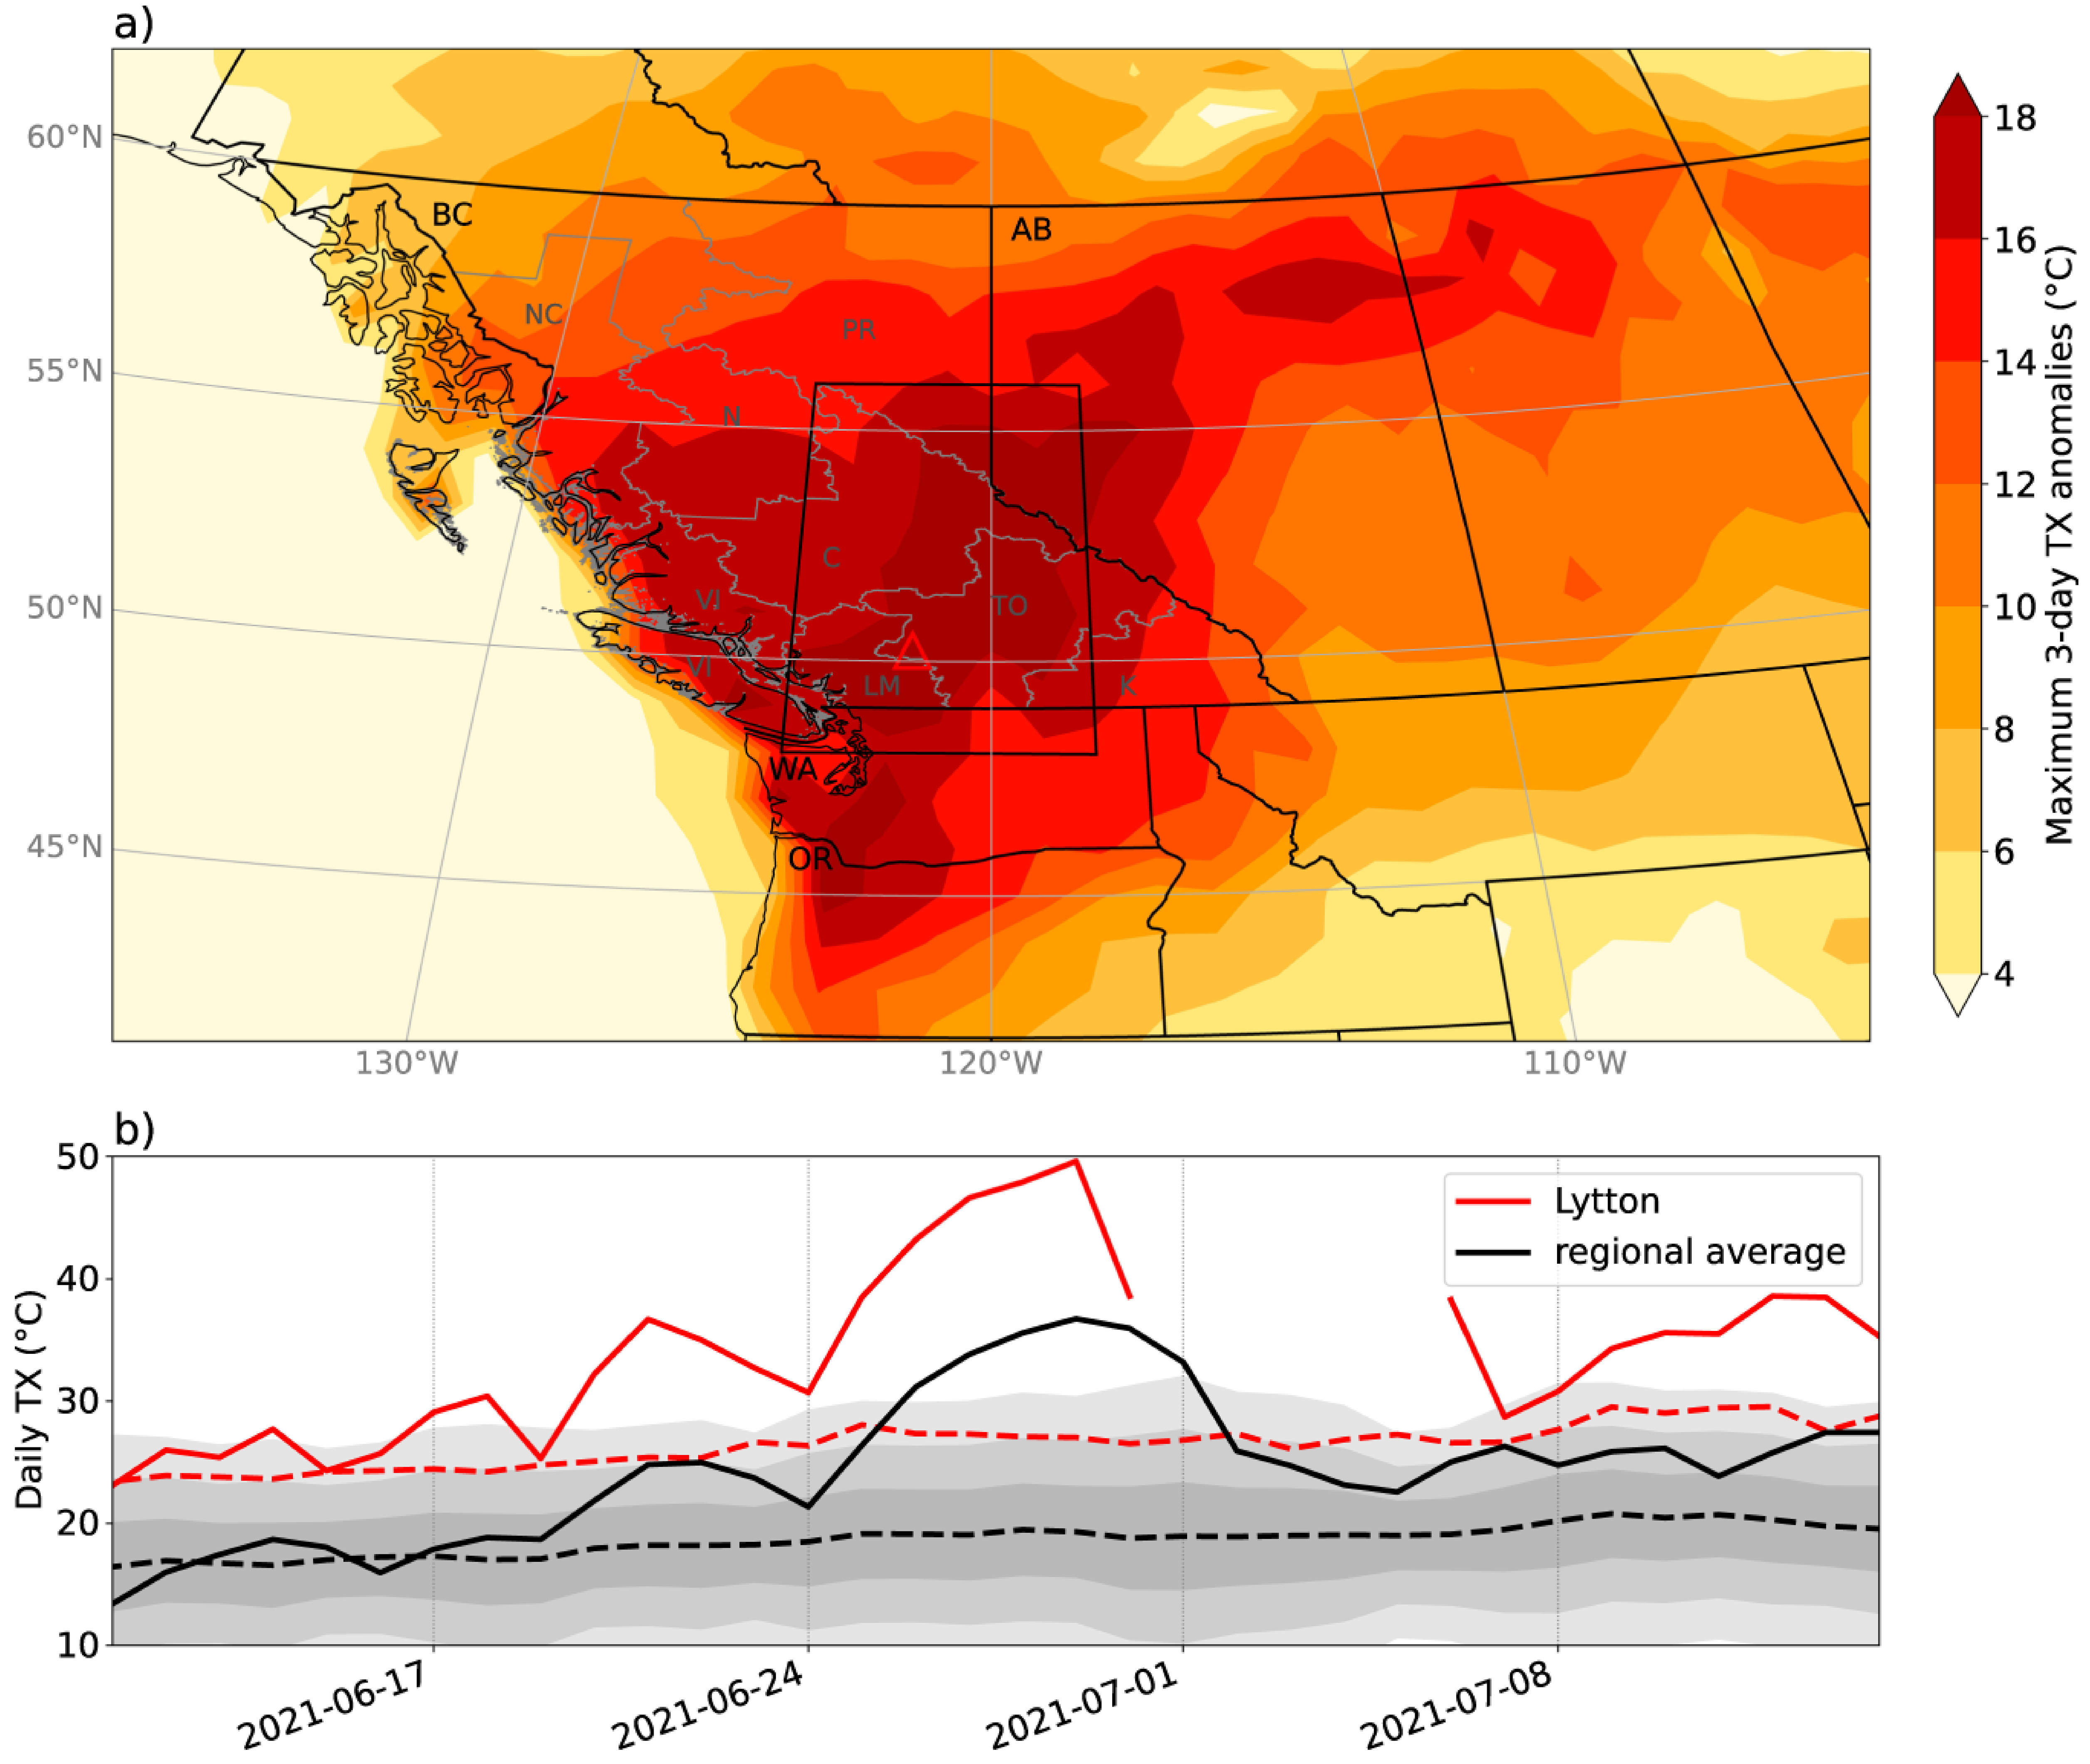

Sample Inflation With Autoregressive Models (Bill Collins' talk this aft)

Huge Ensembles Part I and II, Mahesh & Collins et. al. 2025, in press at *GMD*

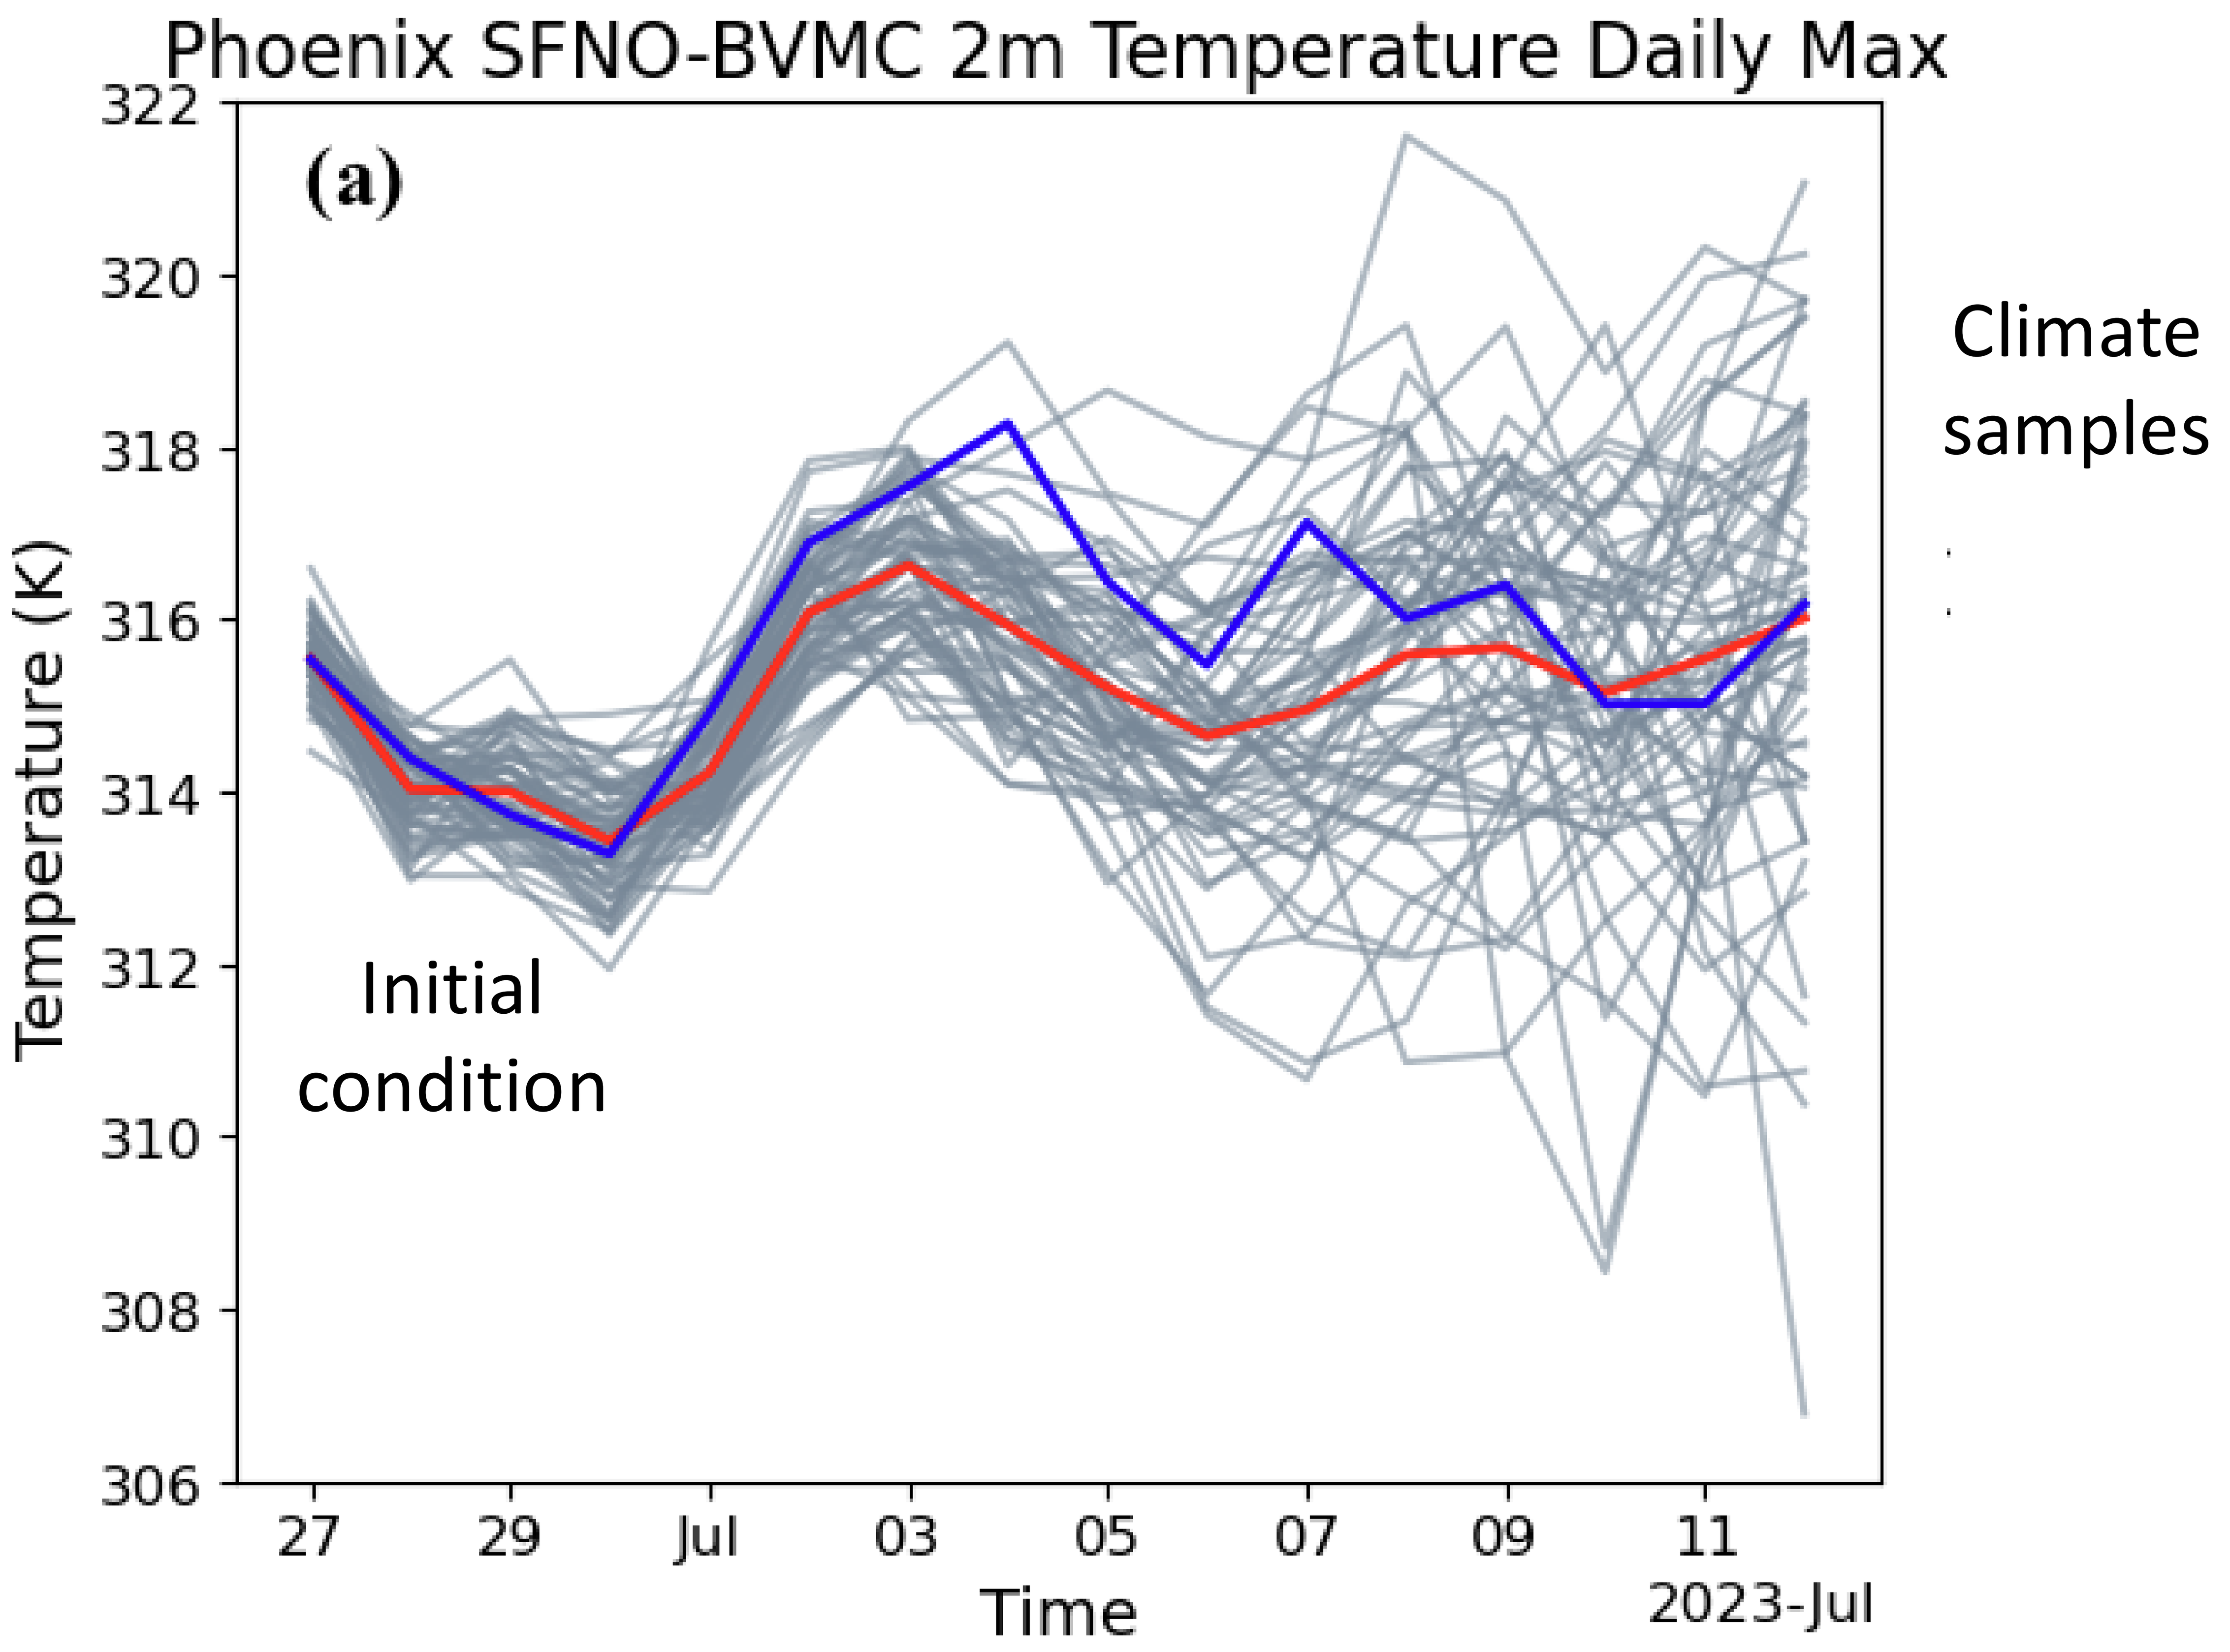

Kansas city heatwave at 10-day leads.

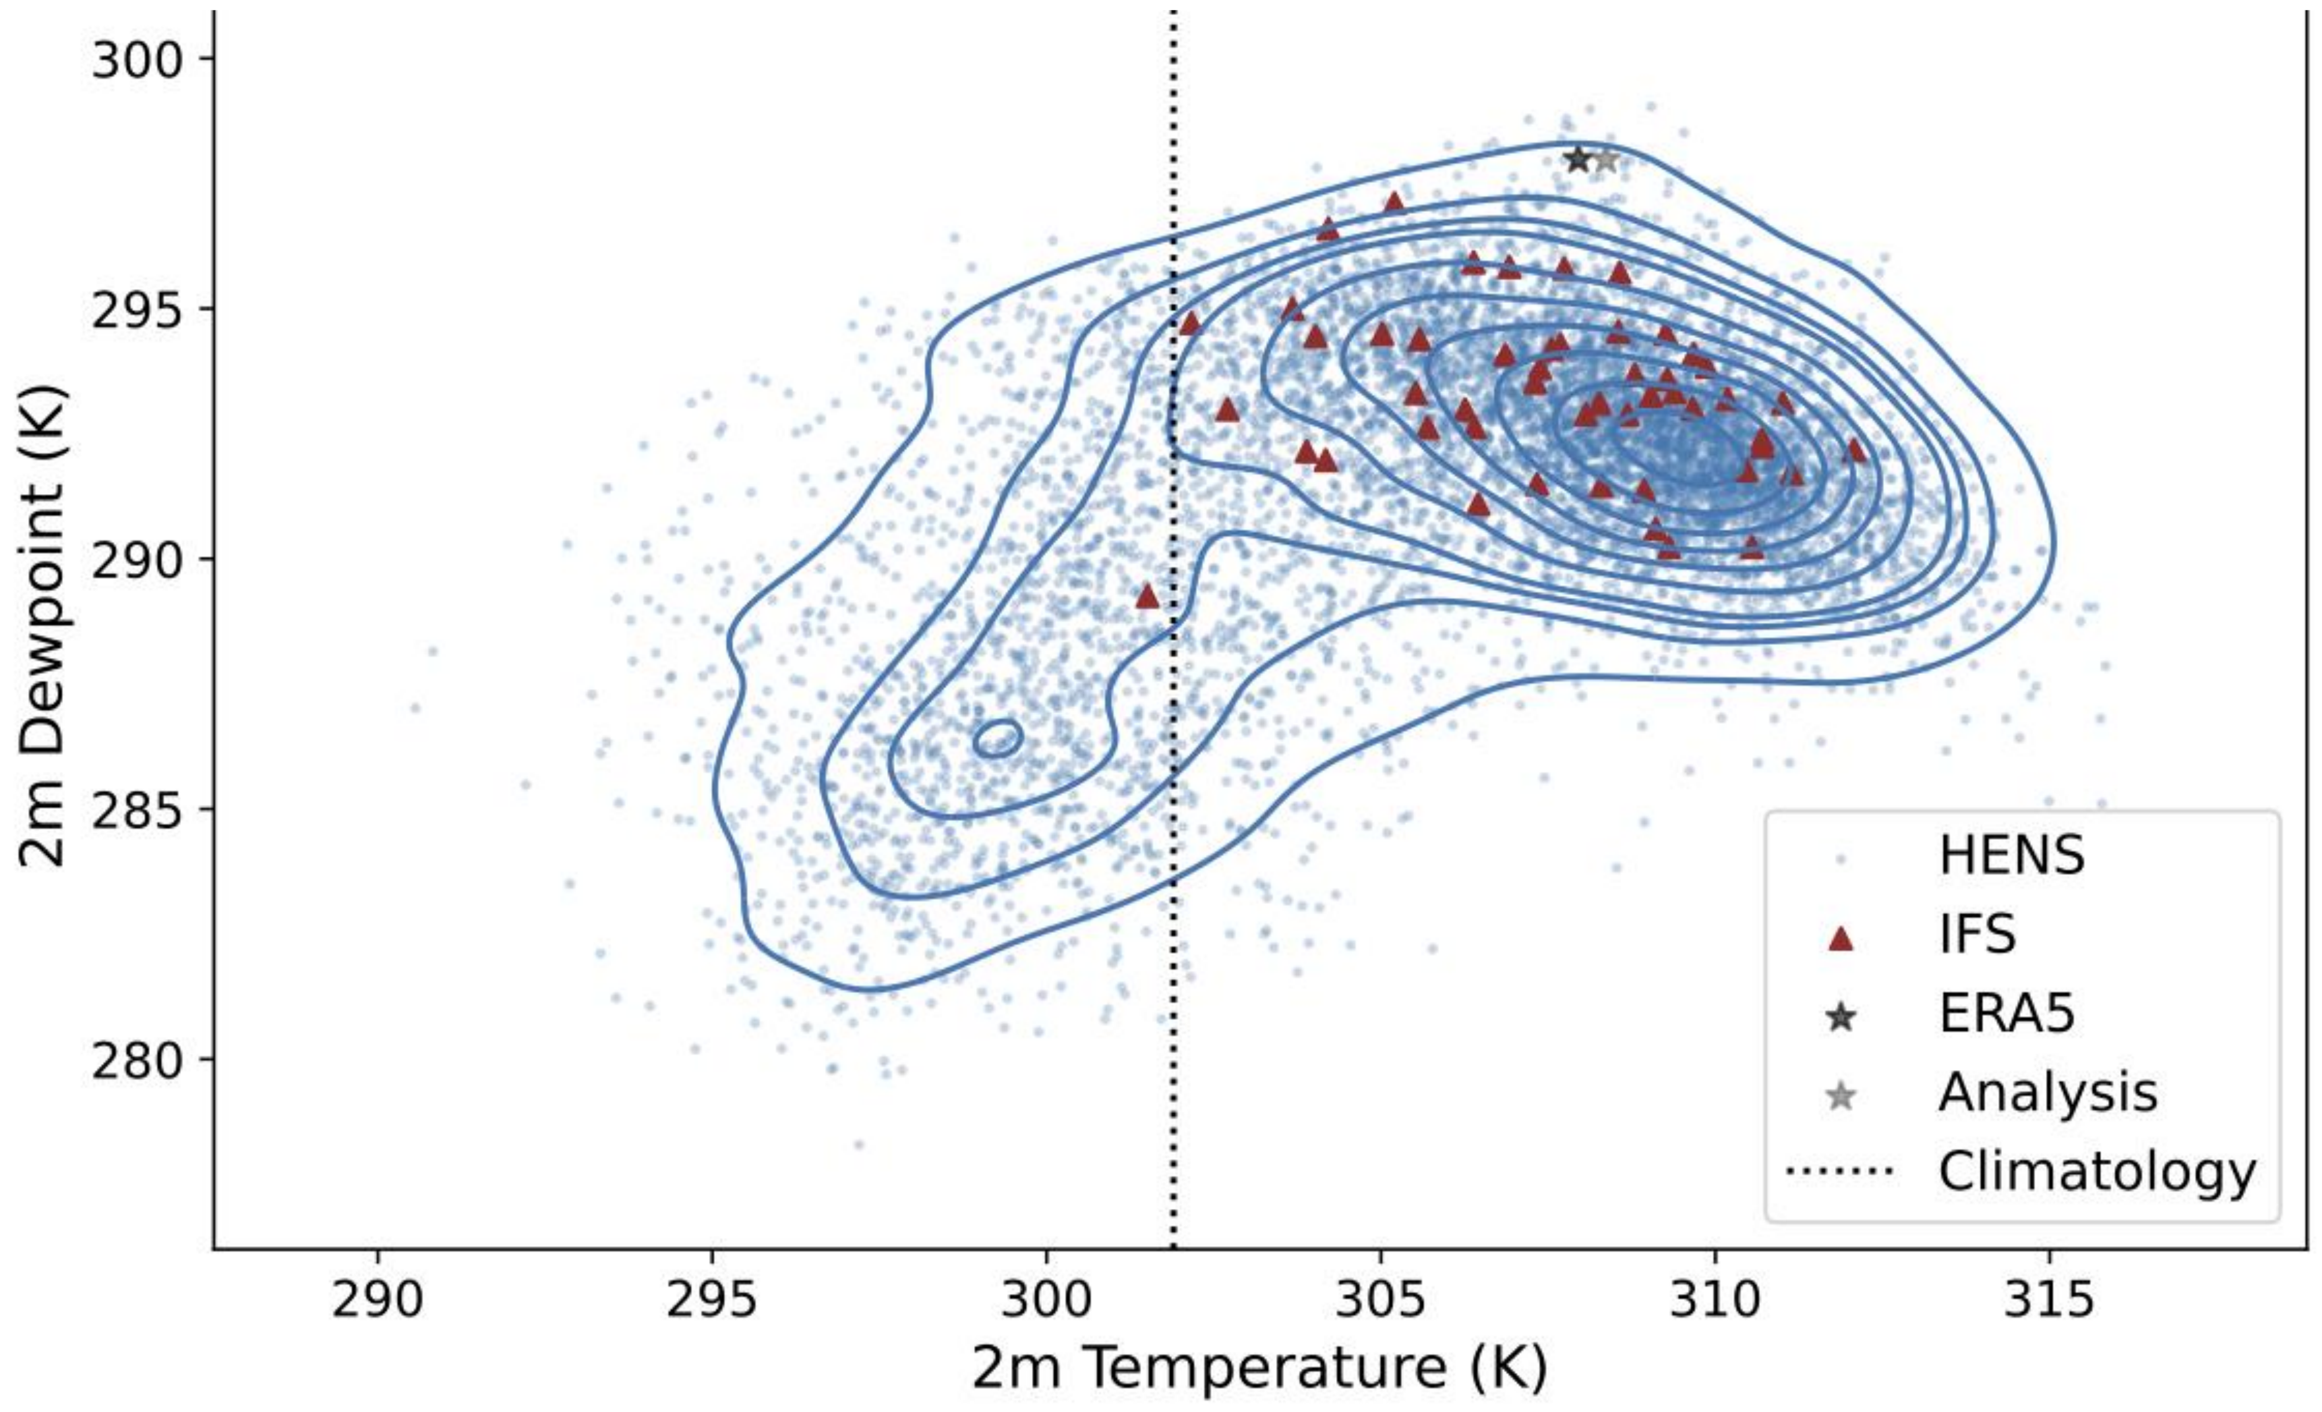

Inspiration: Scaling Diffusions for High-Resolution Image Synthesis

Source: https://imagen.research.google/; https://cascaded-diffusion.github.io/assets/cascaded_diffusion.pdf

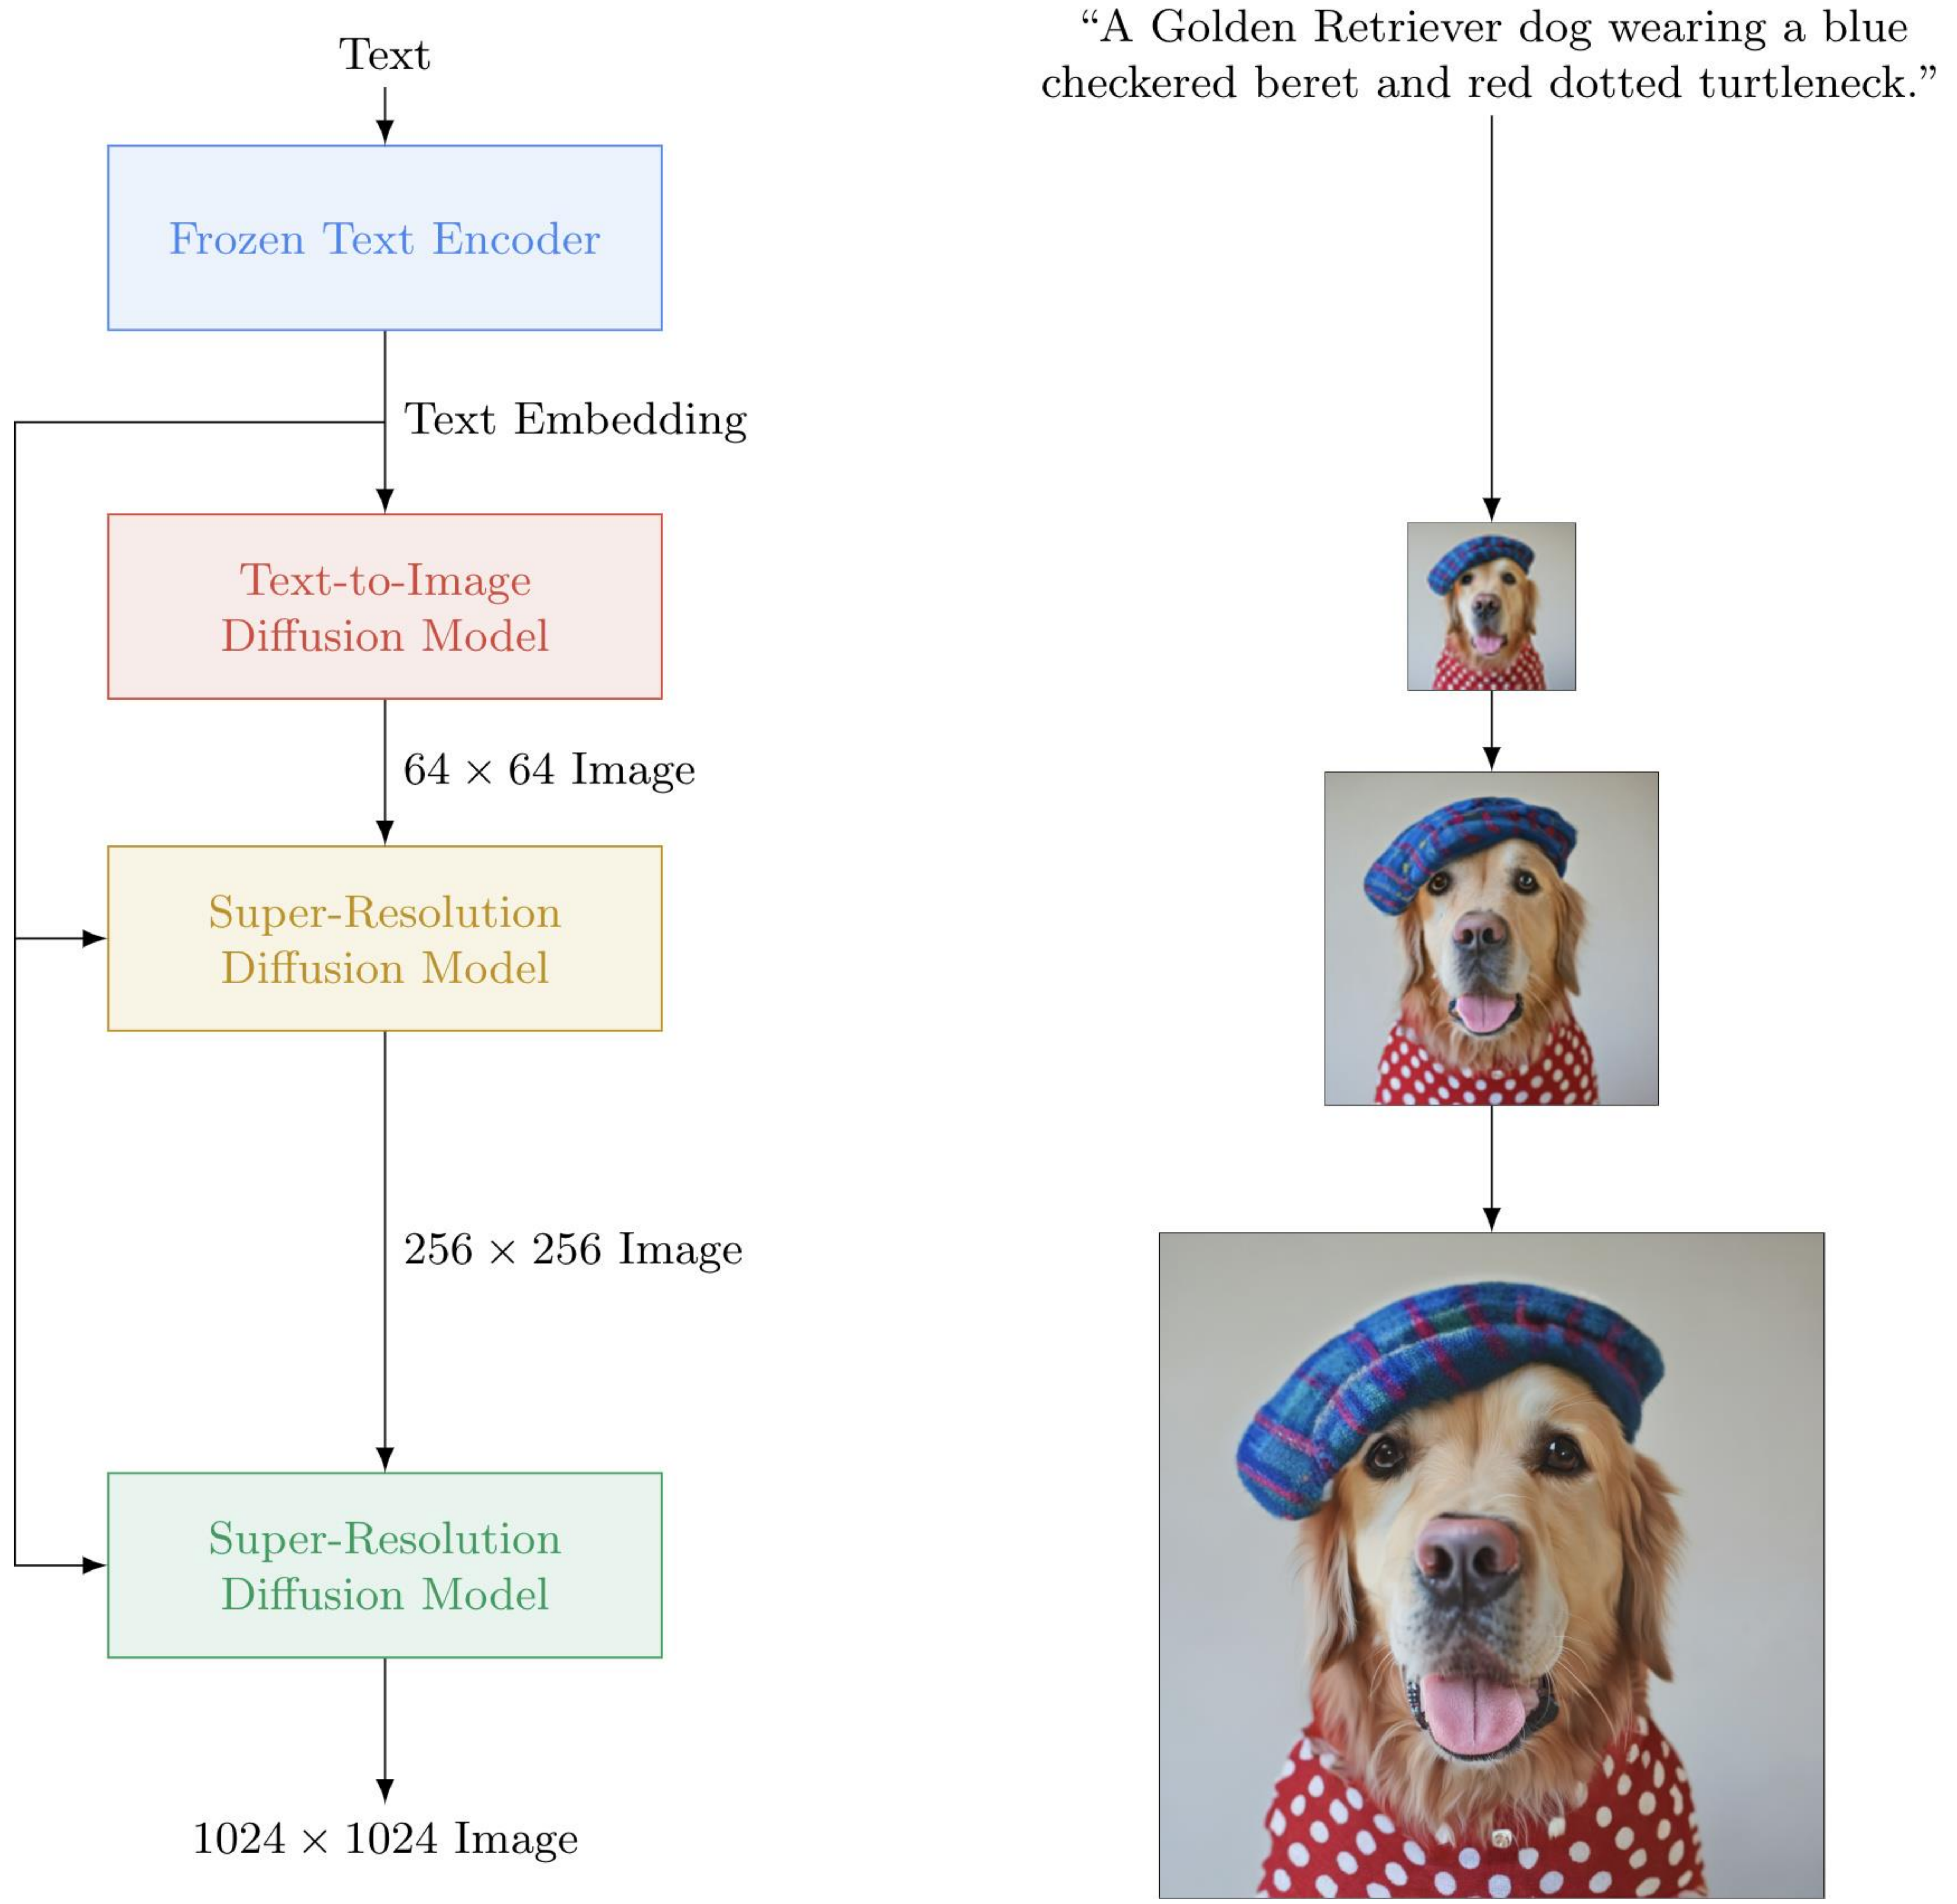

Inspiration: Users Can Control Diffusion Models

User draws sketch

Some generated samples

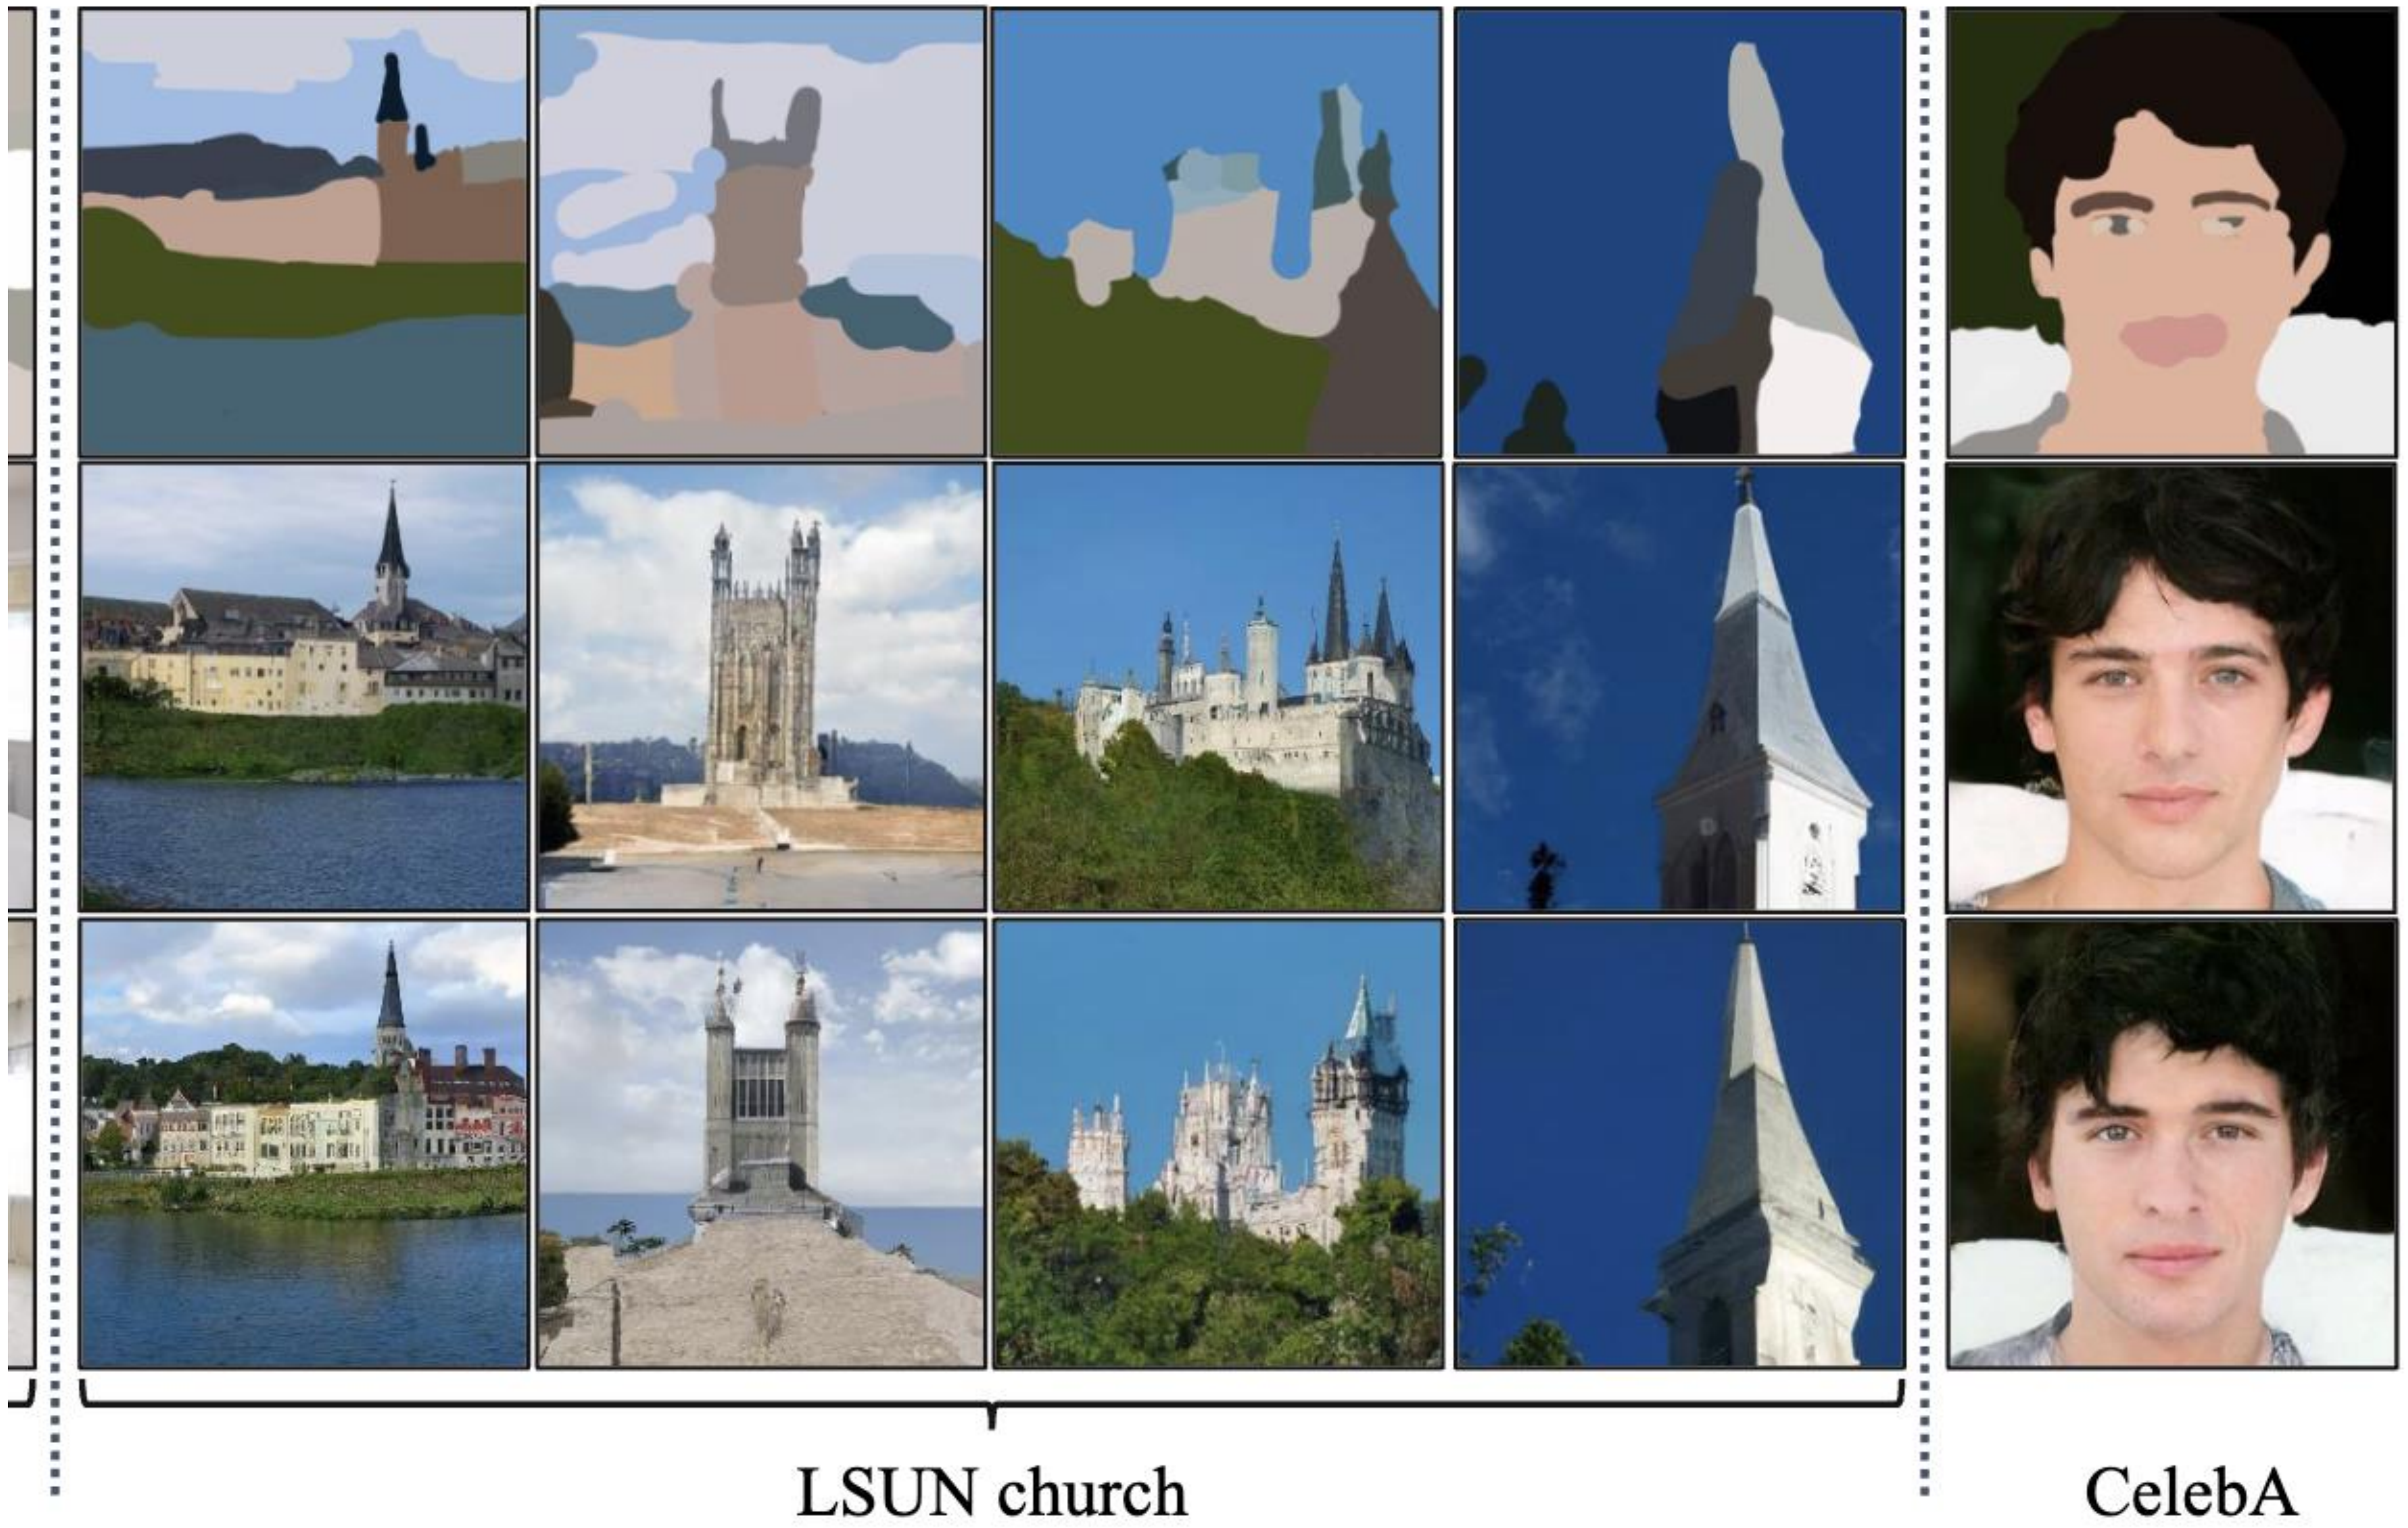

Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.-Y., & Ermon, S. (2021). SDEdit: Guided image synthesis and editing with stochastic differential equations. In *arXiv [cs.CV]*. arXiv. http://arxiv.org/abs/2108.01073

Steerable Climate Sampling Would Be Useful

"AI-on-Top" Should Enable Interactive Experiences

User guided scenario

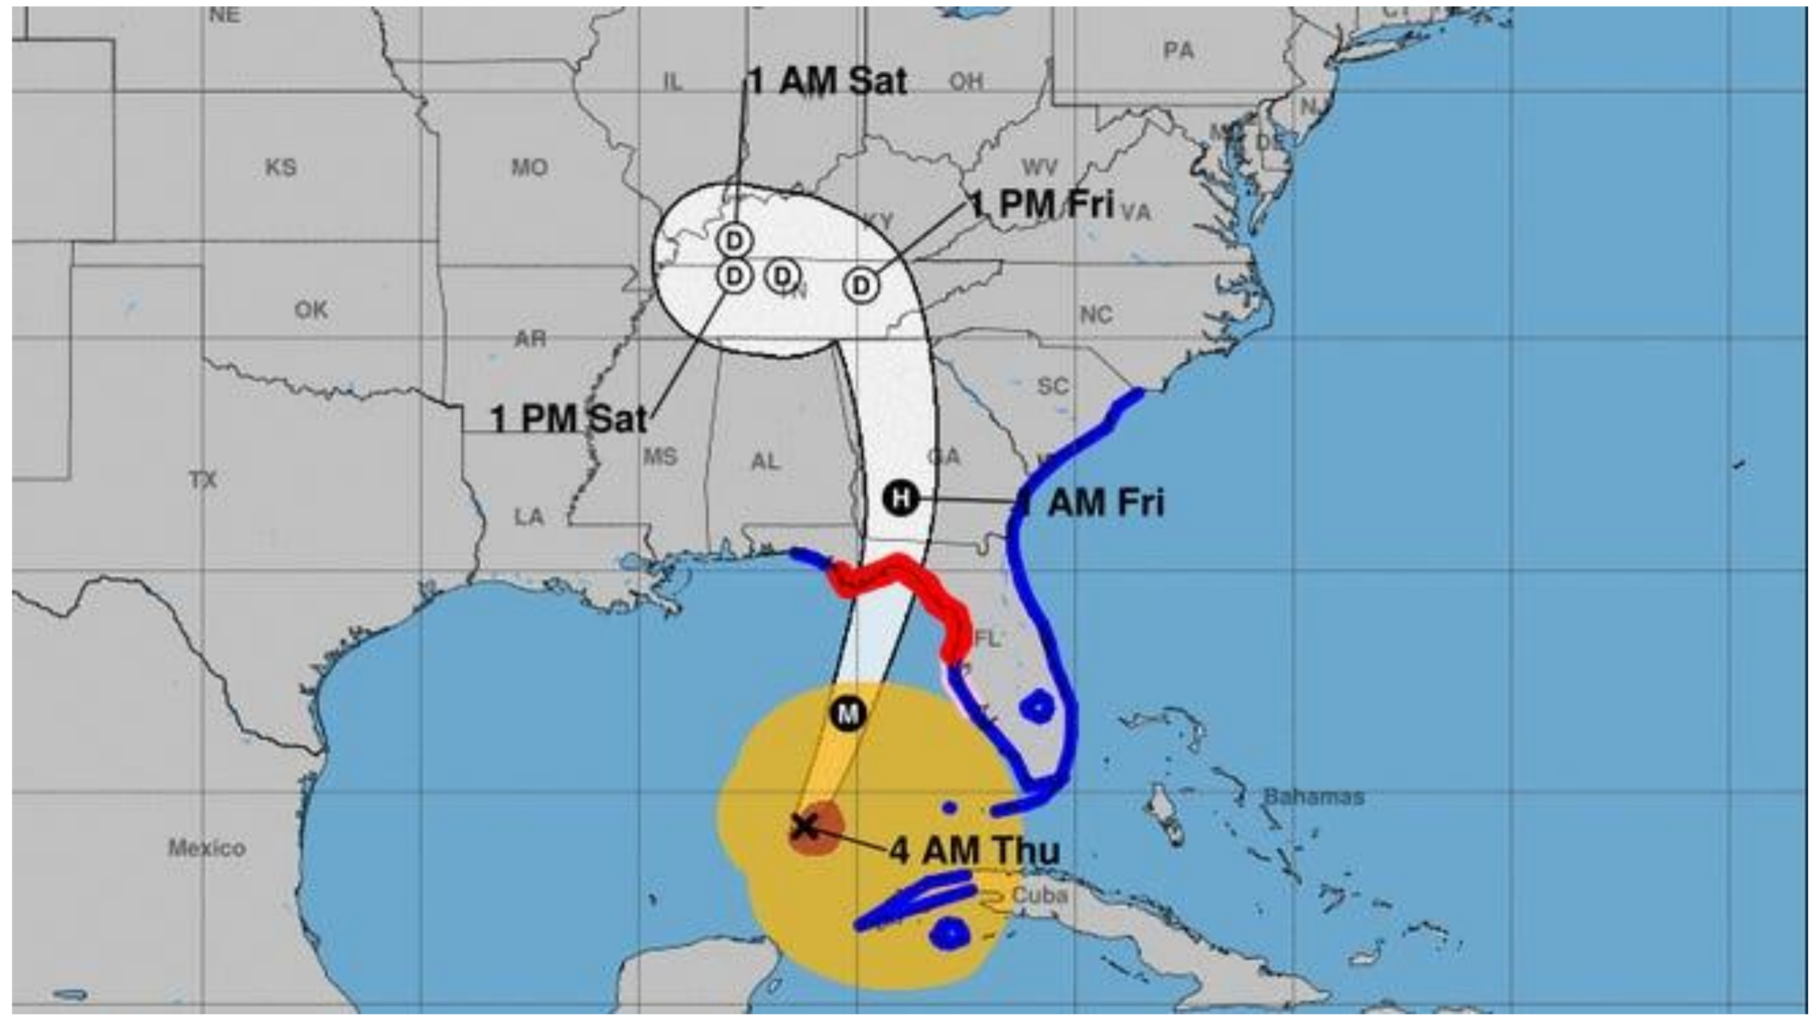

or

Prompt: Show me several realizations of a potential hurricane impacting the gulf coast given the current average surface conditions. Output precipitation and wind damage maps.

Multiple realizations

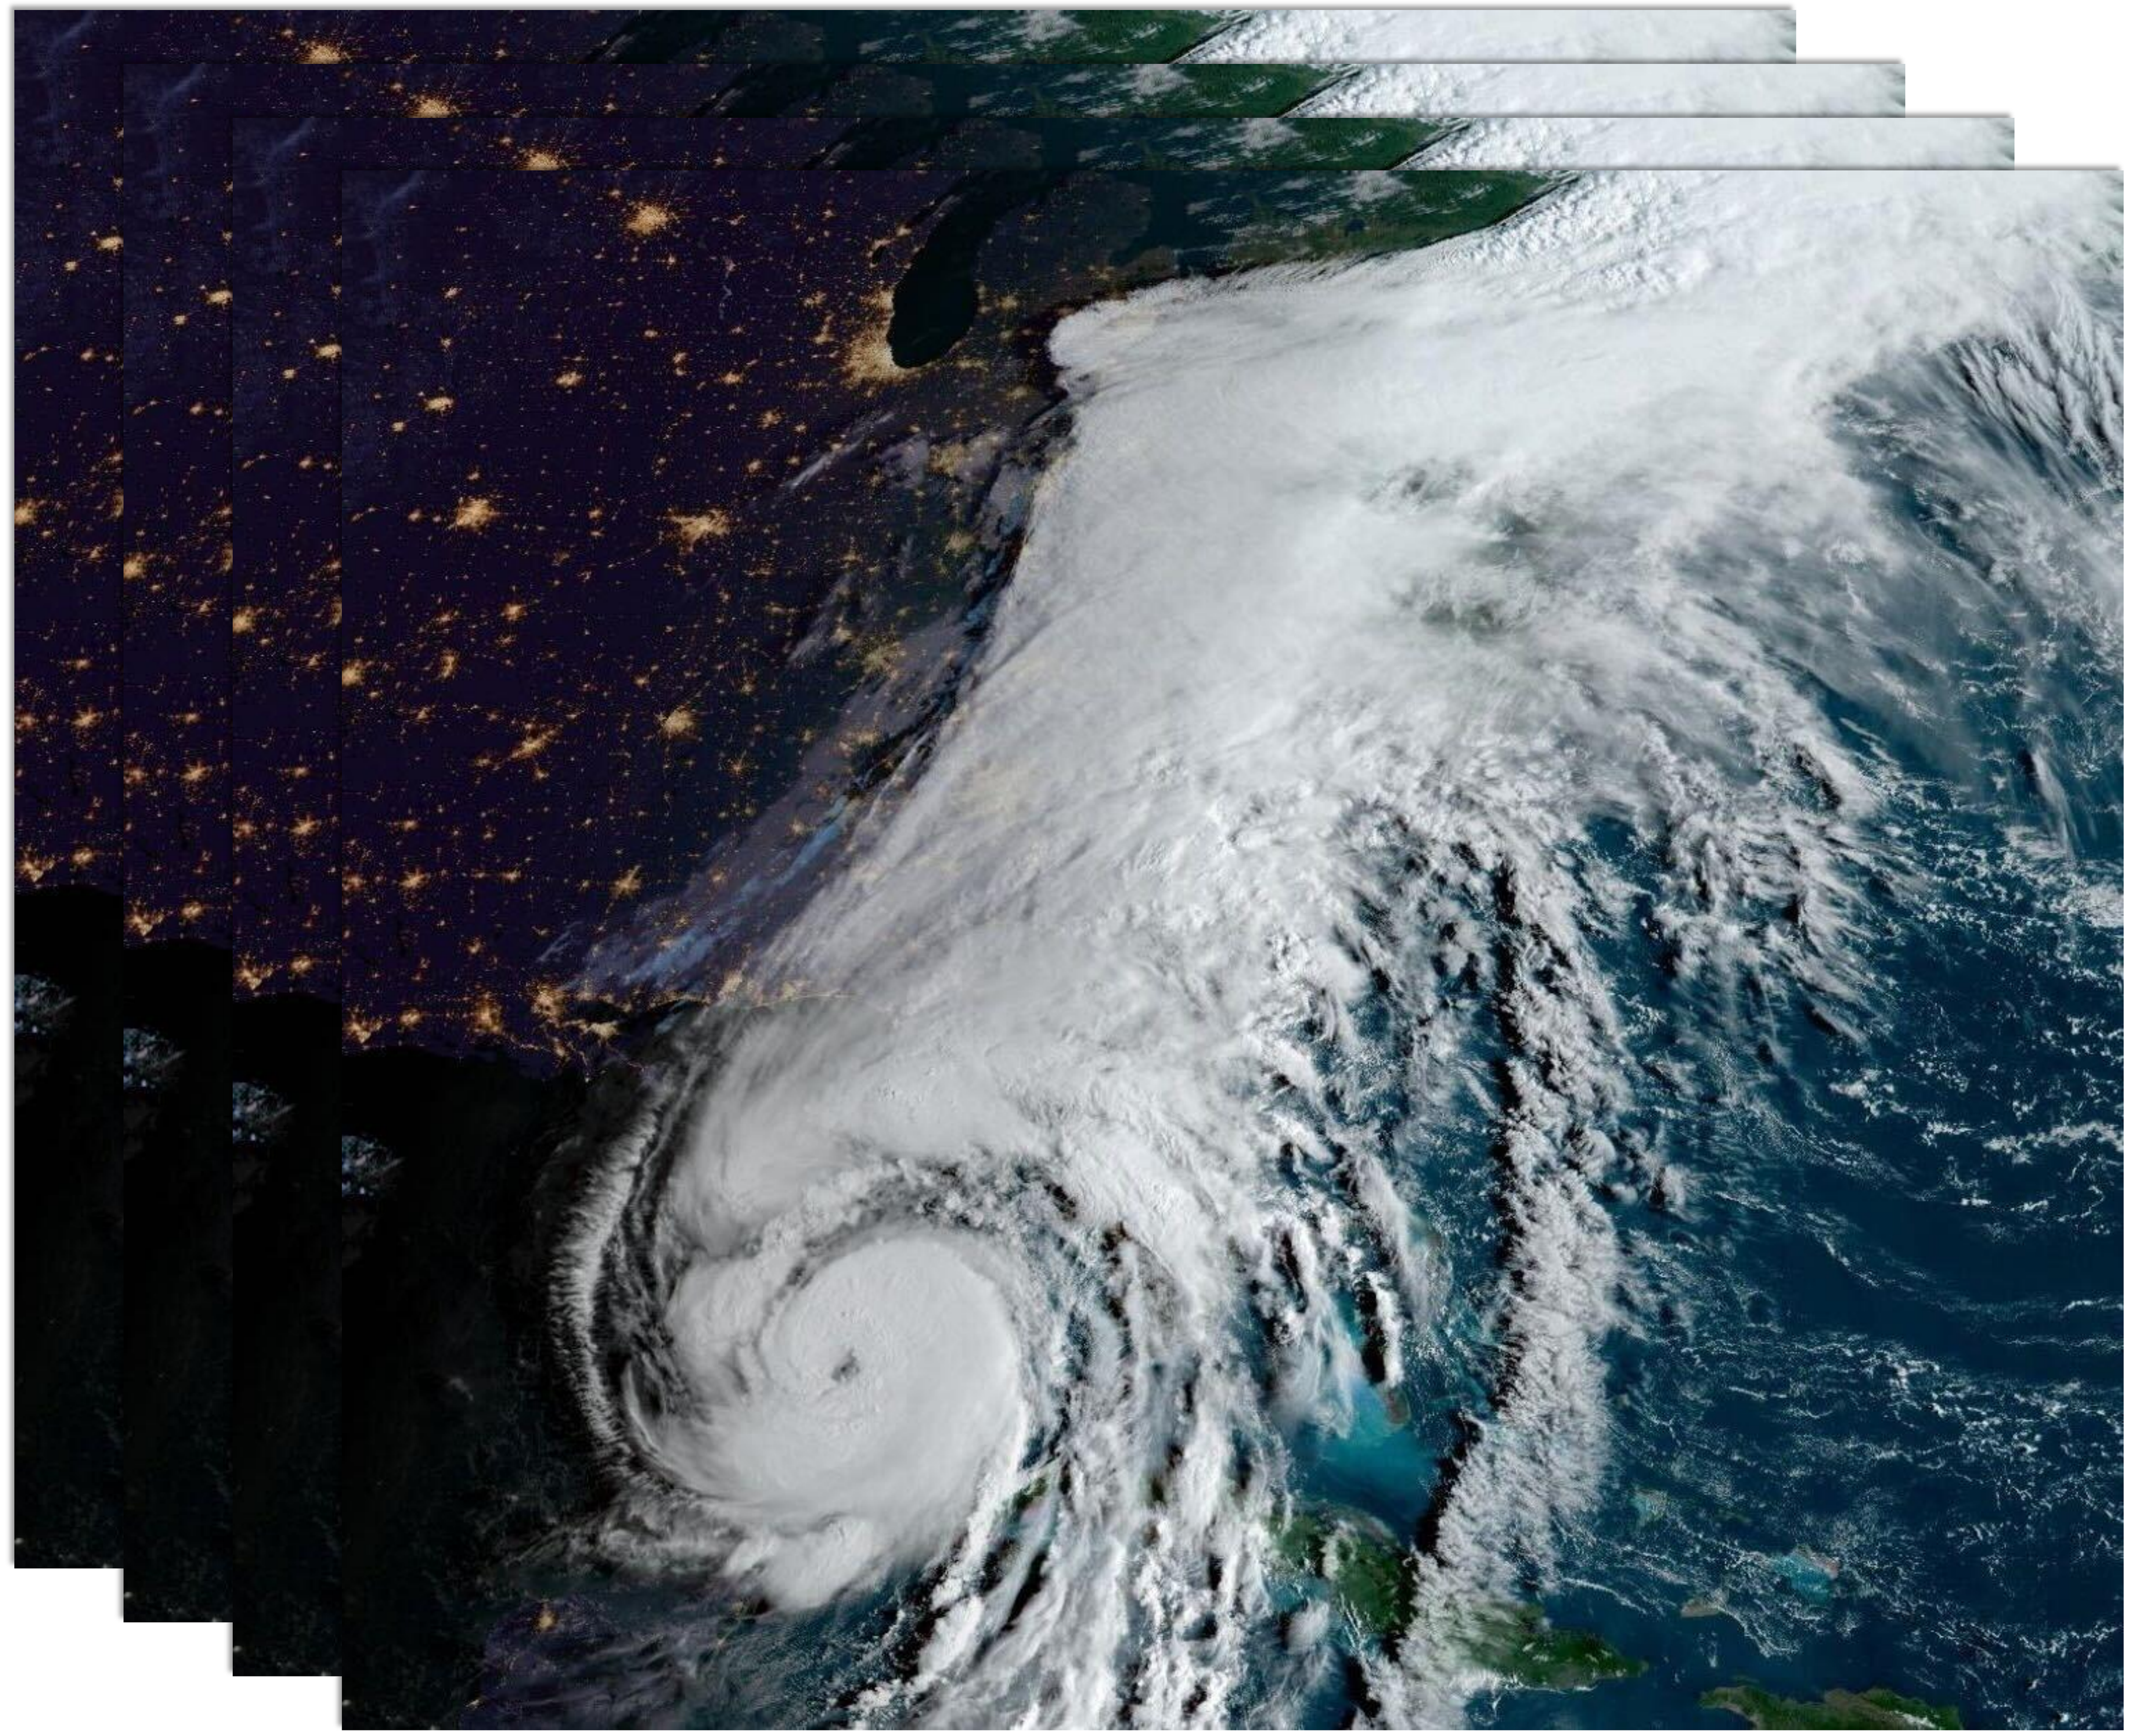

Hurricane Helene, Sept 26, *NOAA*

Two Ways to Generate Atmospheric States

Autoregression

x(n+1) ~ P(x(n)) for n = 1 to ∞

- + it's how physics works
- + it's how GCMs work
- + it's how graphcast/ace/dlwp/FCN/etc works
- + easy to train
- - deterministic requires many channels
- - ML has no guarantees for $n \to \infty$
  - No consistency + stability = accuracy result
  - Blurring
  - Hallucinations, drifts, blow-ups, etc
- Difficult to control.
  - Data assimilation is hard
  - How would you generate a realistic hurricane impacting a specific location?

Denoising diffusion models

x(n-1) = f(x(n), y) from n=T to 0

- + Limited number of steps
- + Grounded in theory of maximum likelihood and stochastic processes
- + Controllability: User control as one pleases (y)
- + Simple training
- - Complicated and expensive inference
- - More hyperparameters. Tuning noise schedules in addition to input/output normalization

Unifying Two Distinct Climate Data Modalities

Video Language Model → Reanalysis CMIP Model?

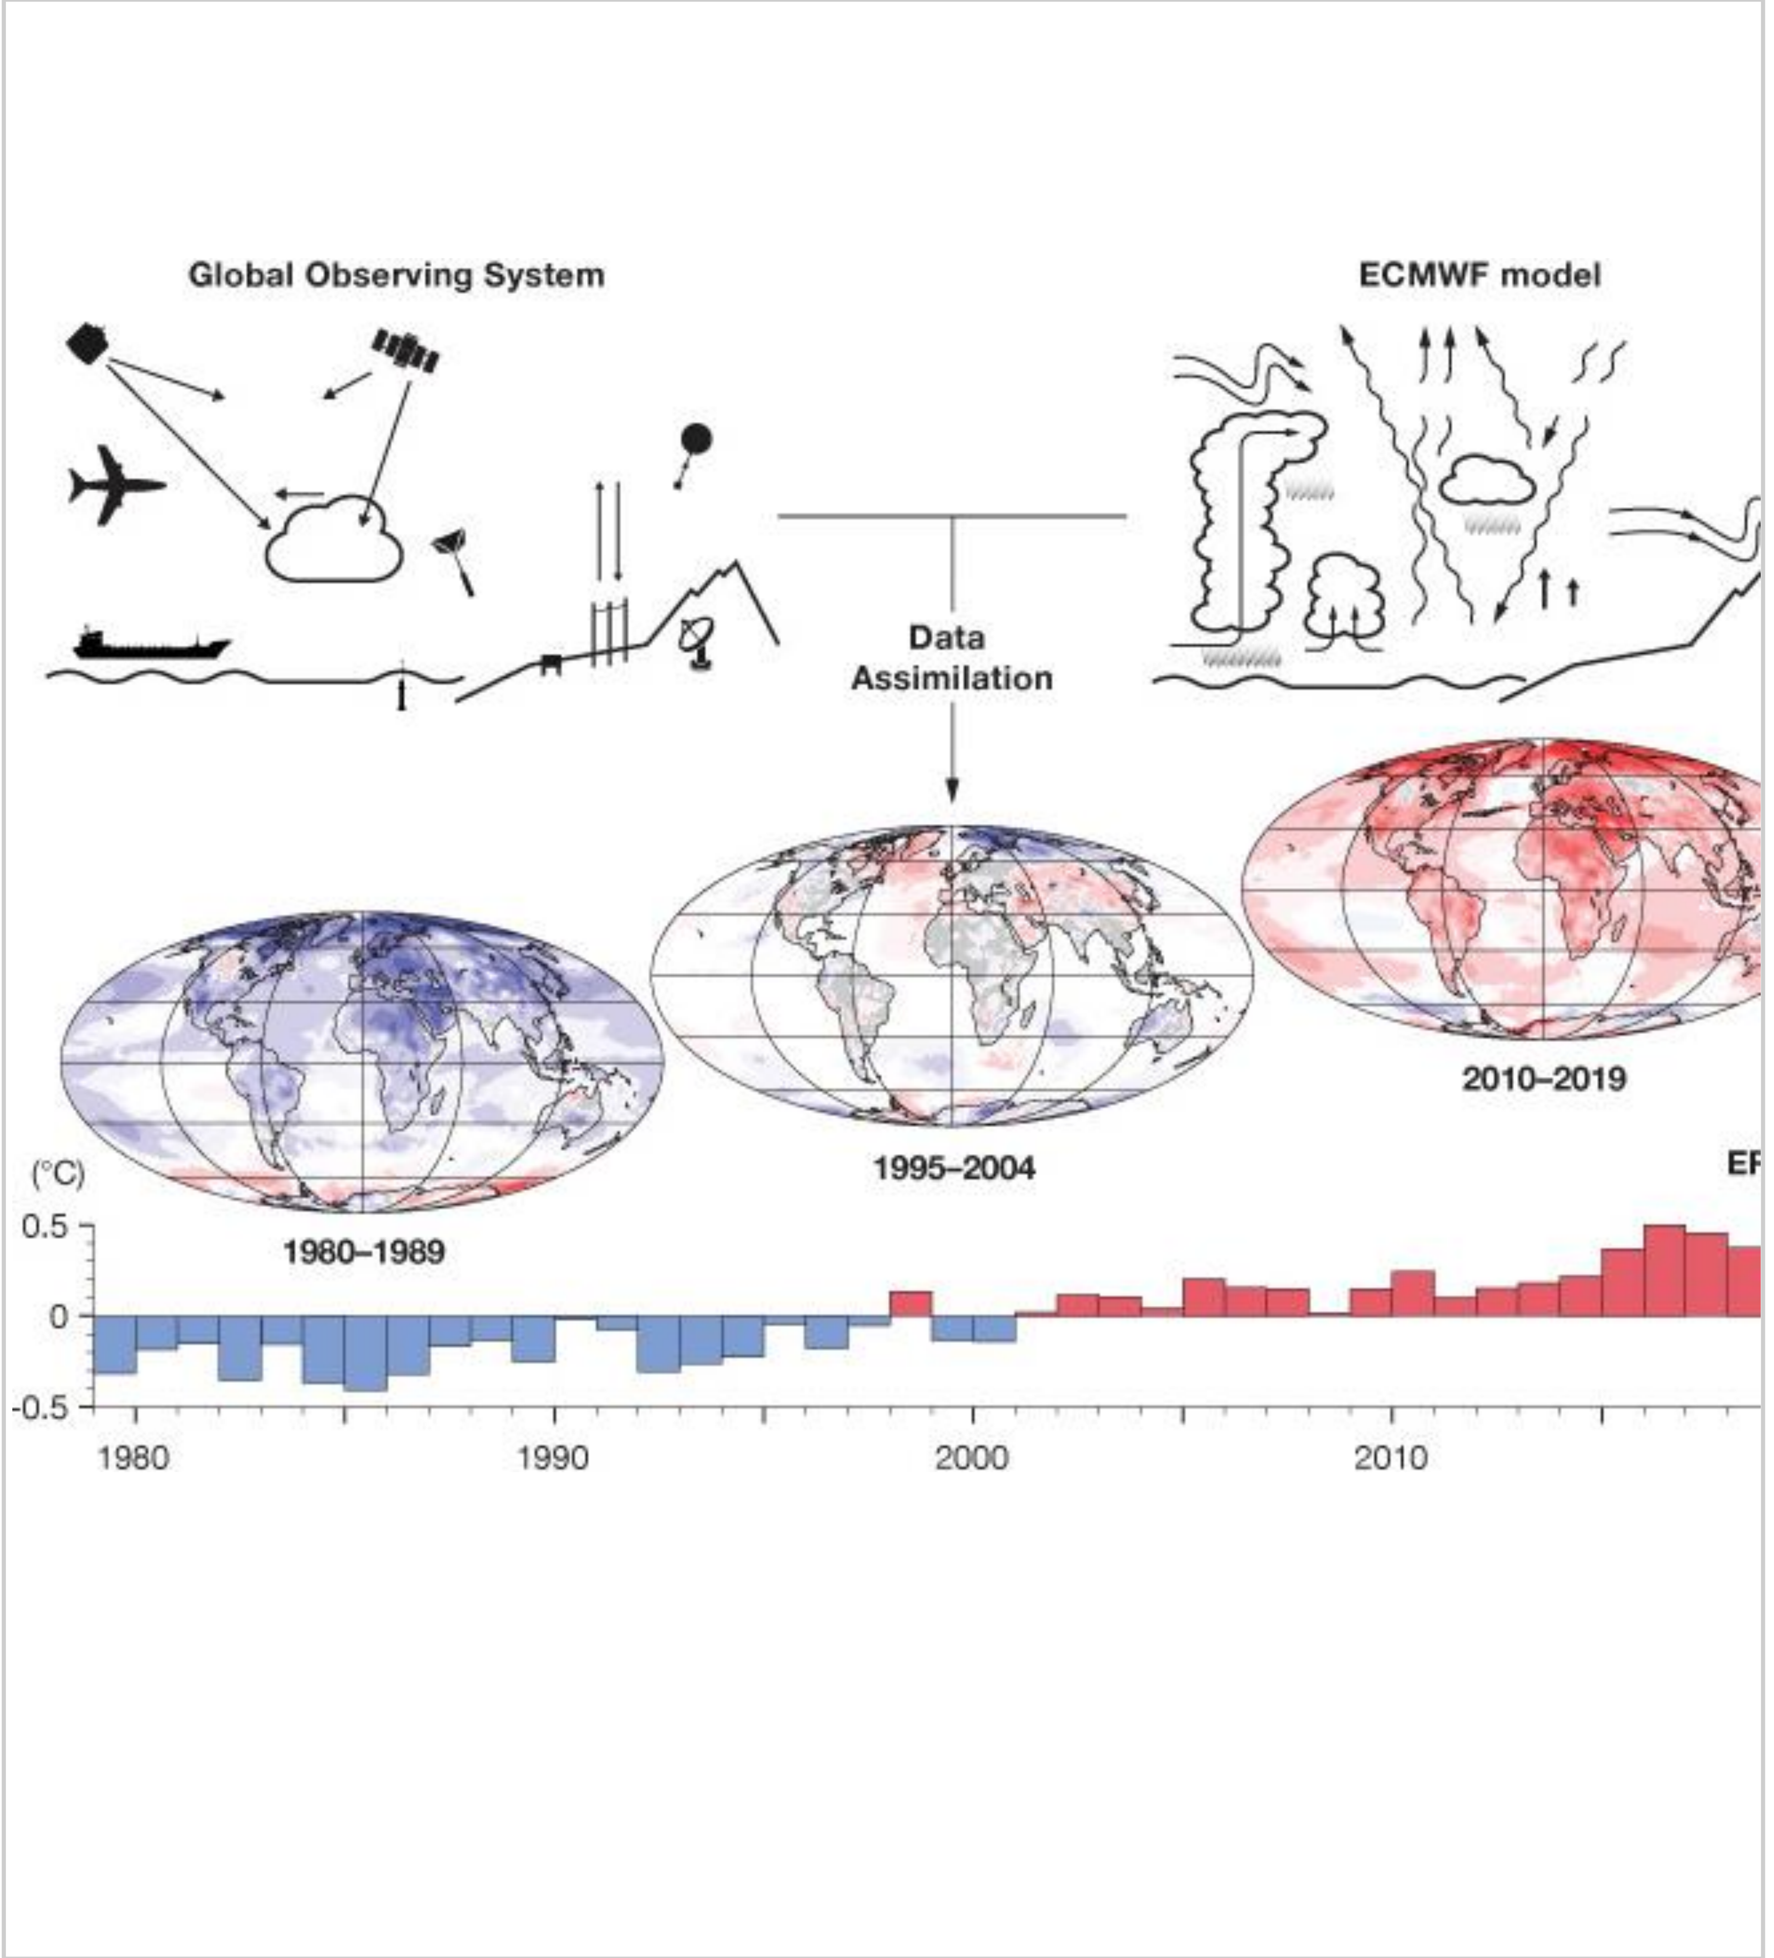

HPX64 Reanalysis (ERA5)

~ 50k pixels / channel / sample

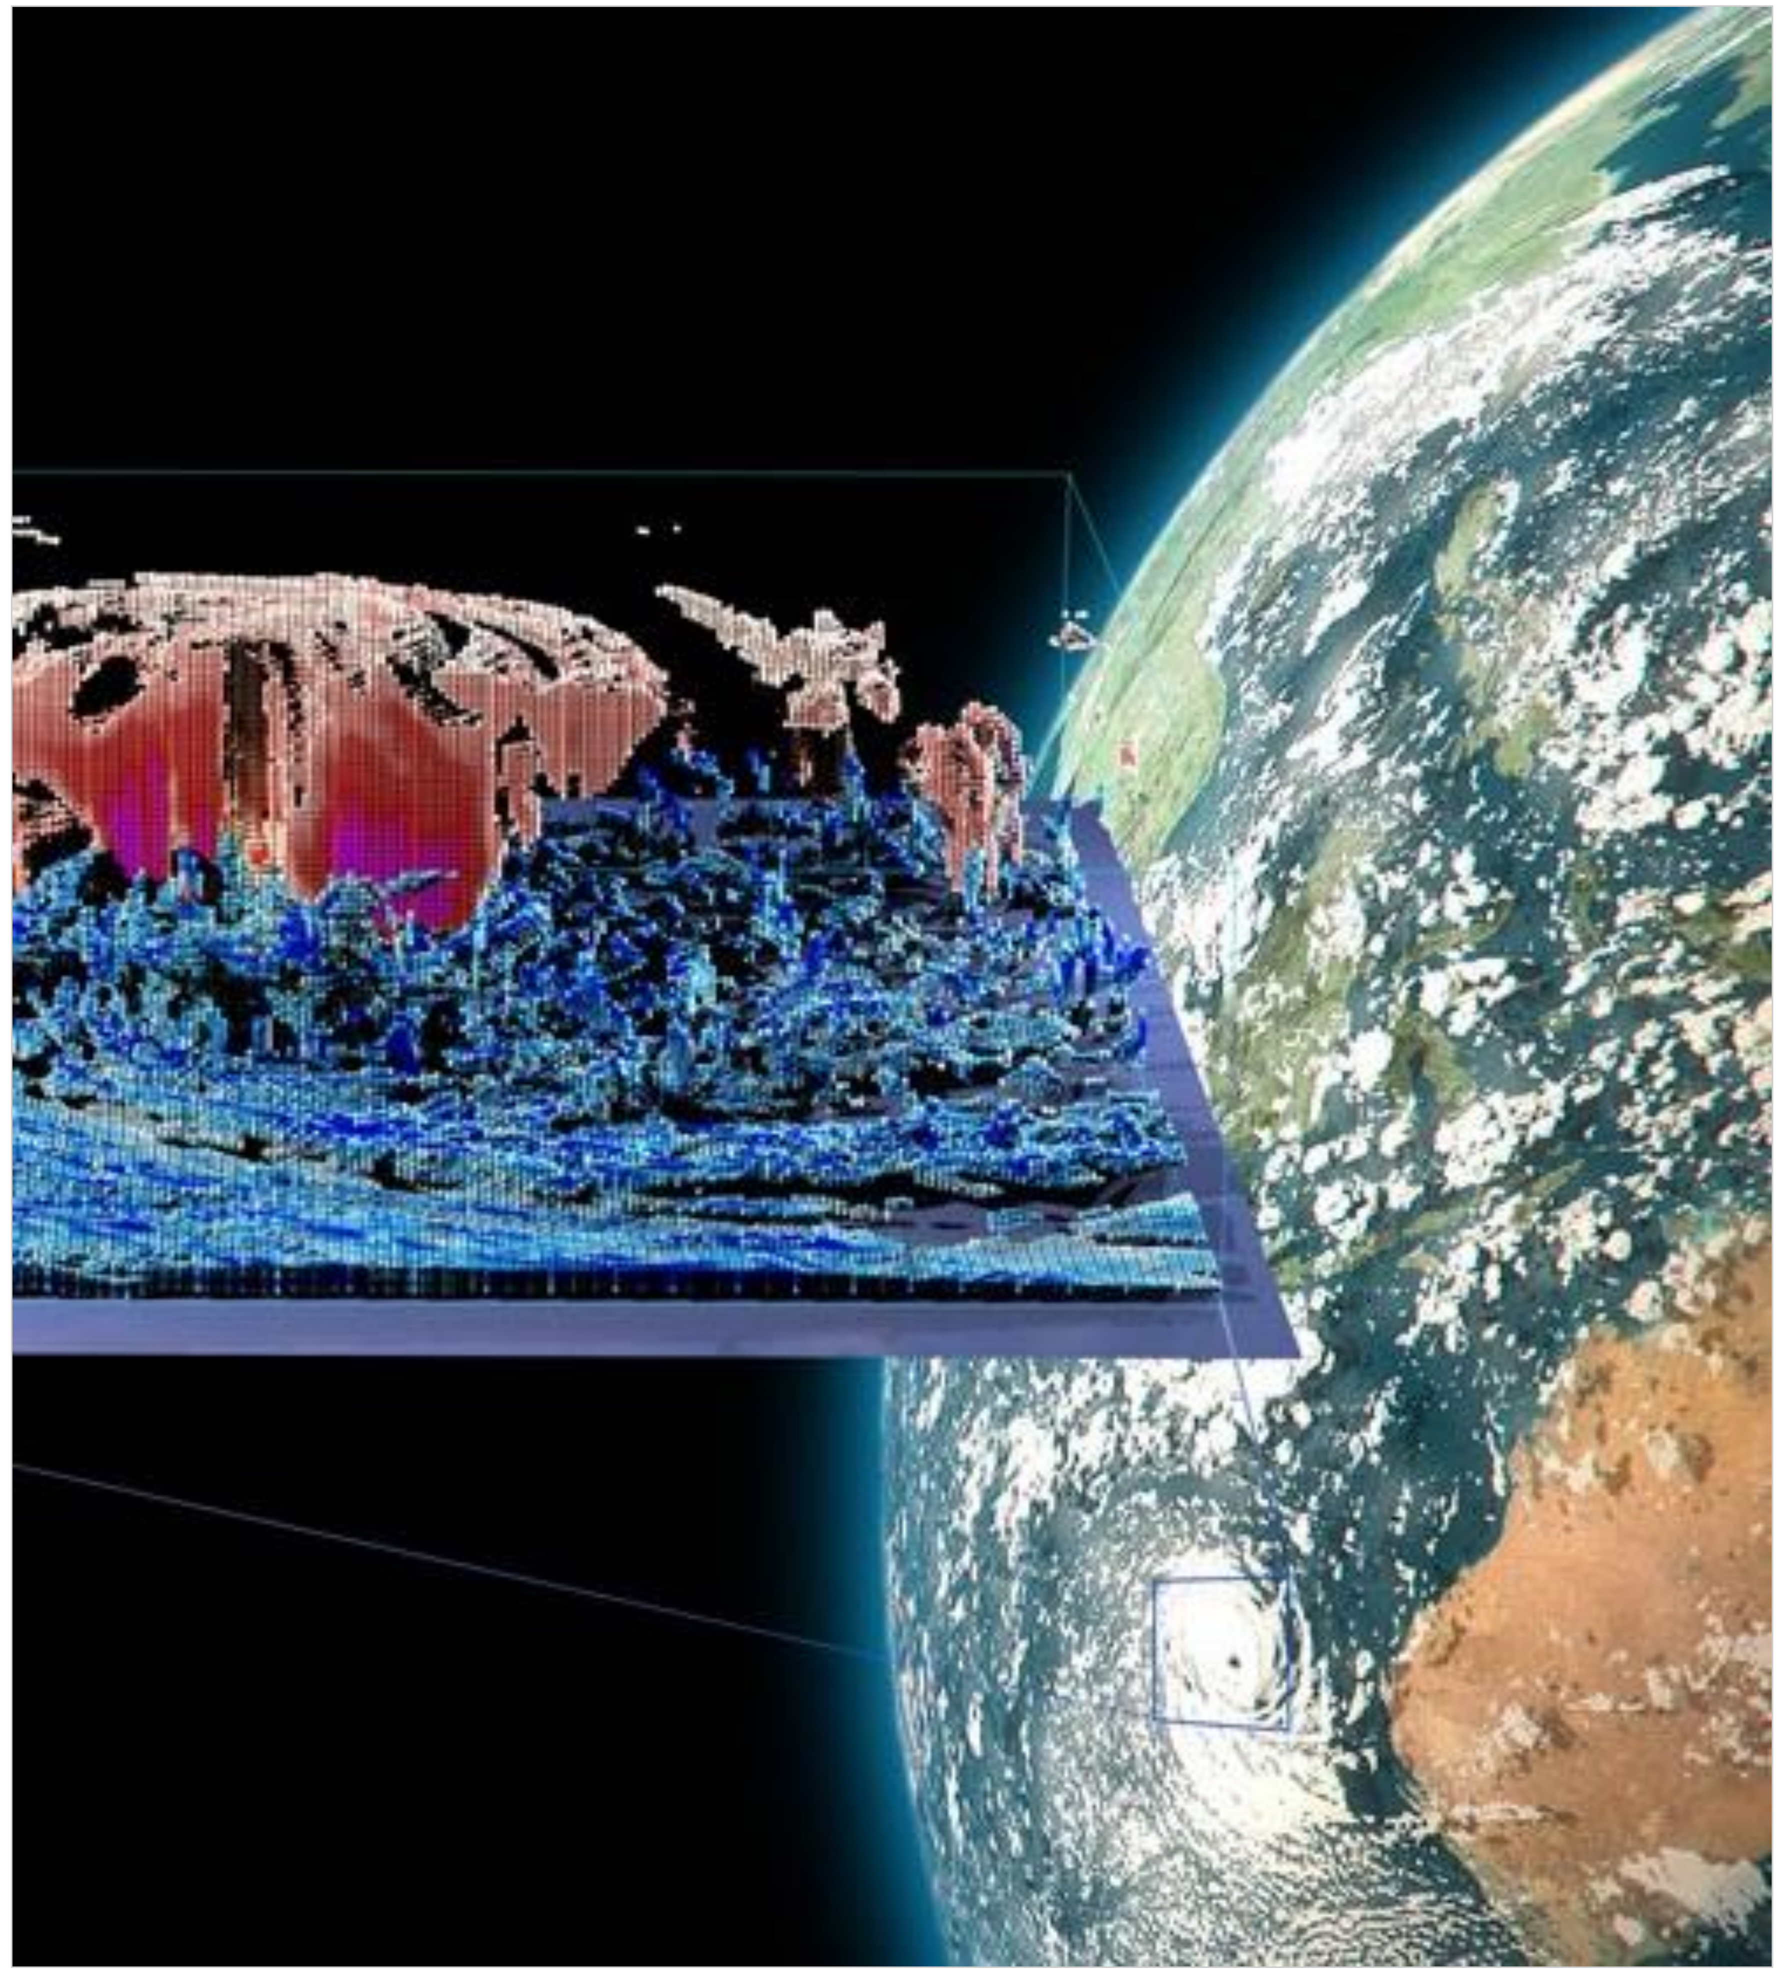

HPX1024 ICON

12.5M pixels / channel / sample

Component 1: Multi-Modal MacroScale HPX64 Generator

Capable of Generating Either the ICON or ERA5 Modality at HPX64 Resolution (~100-km, 50k pixels)

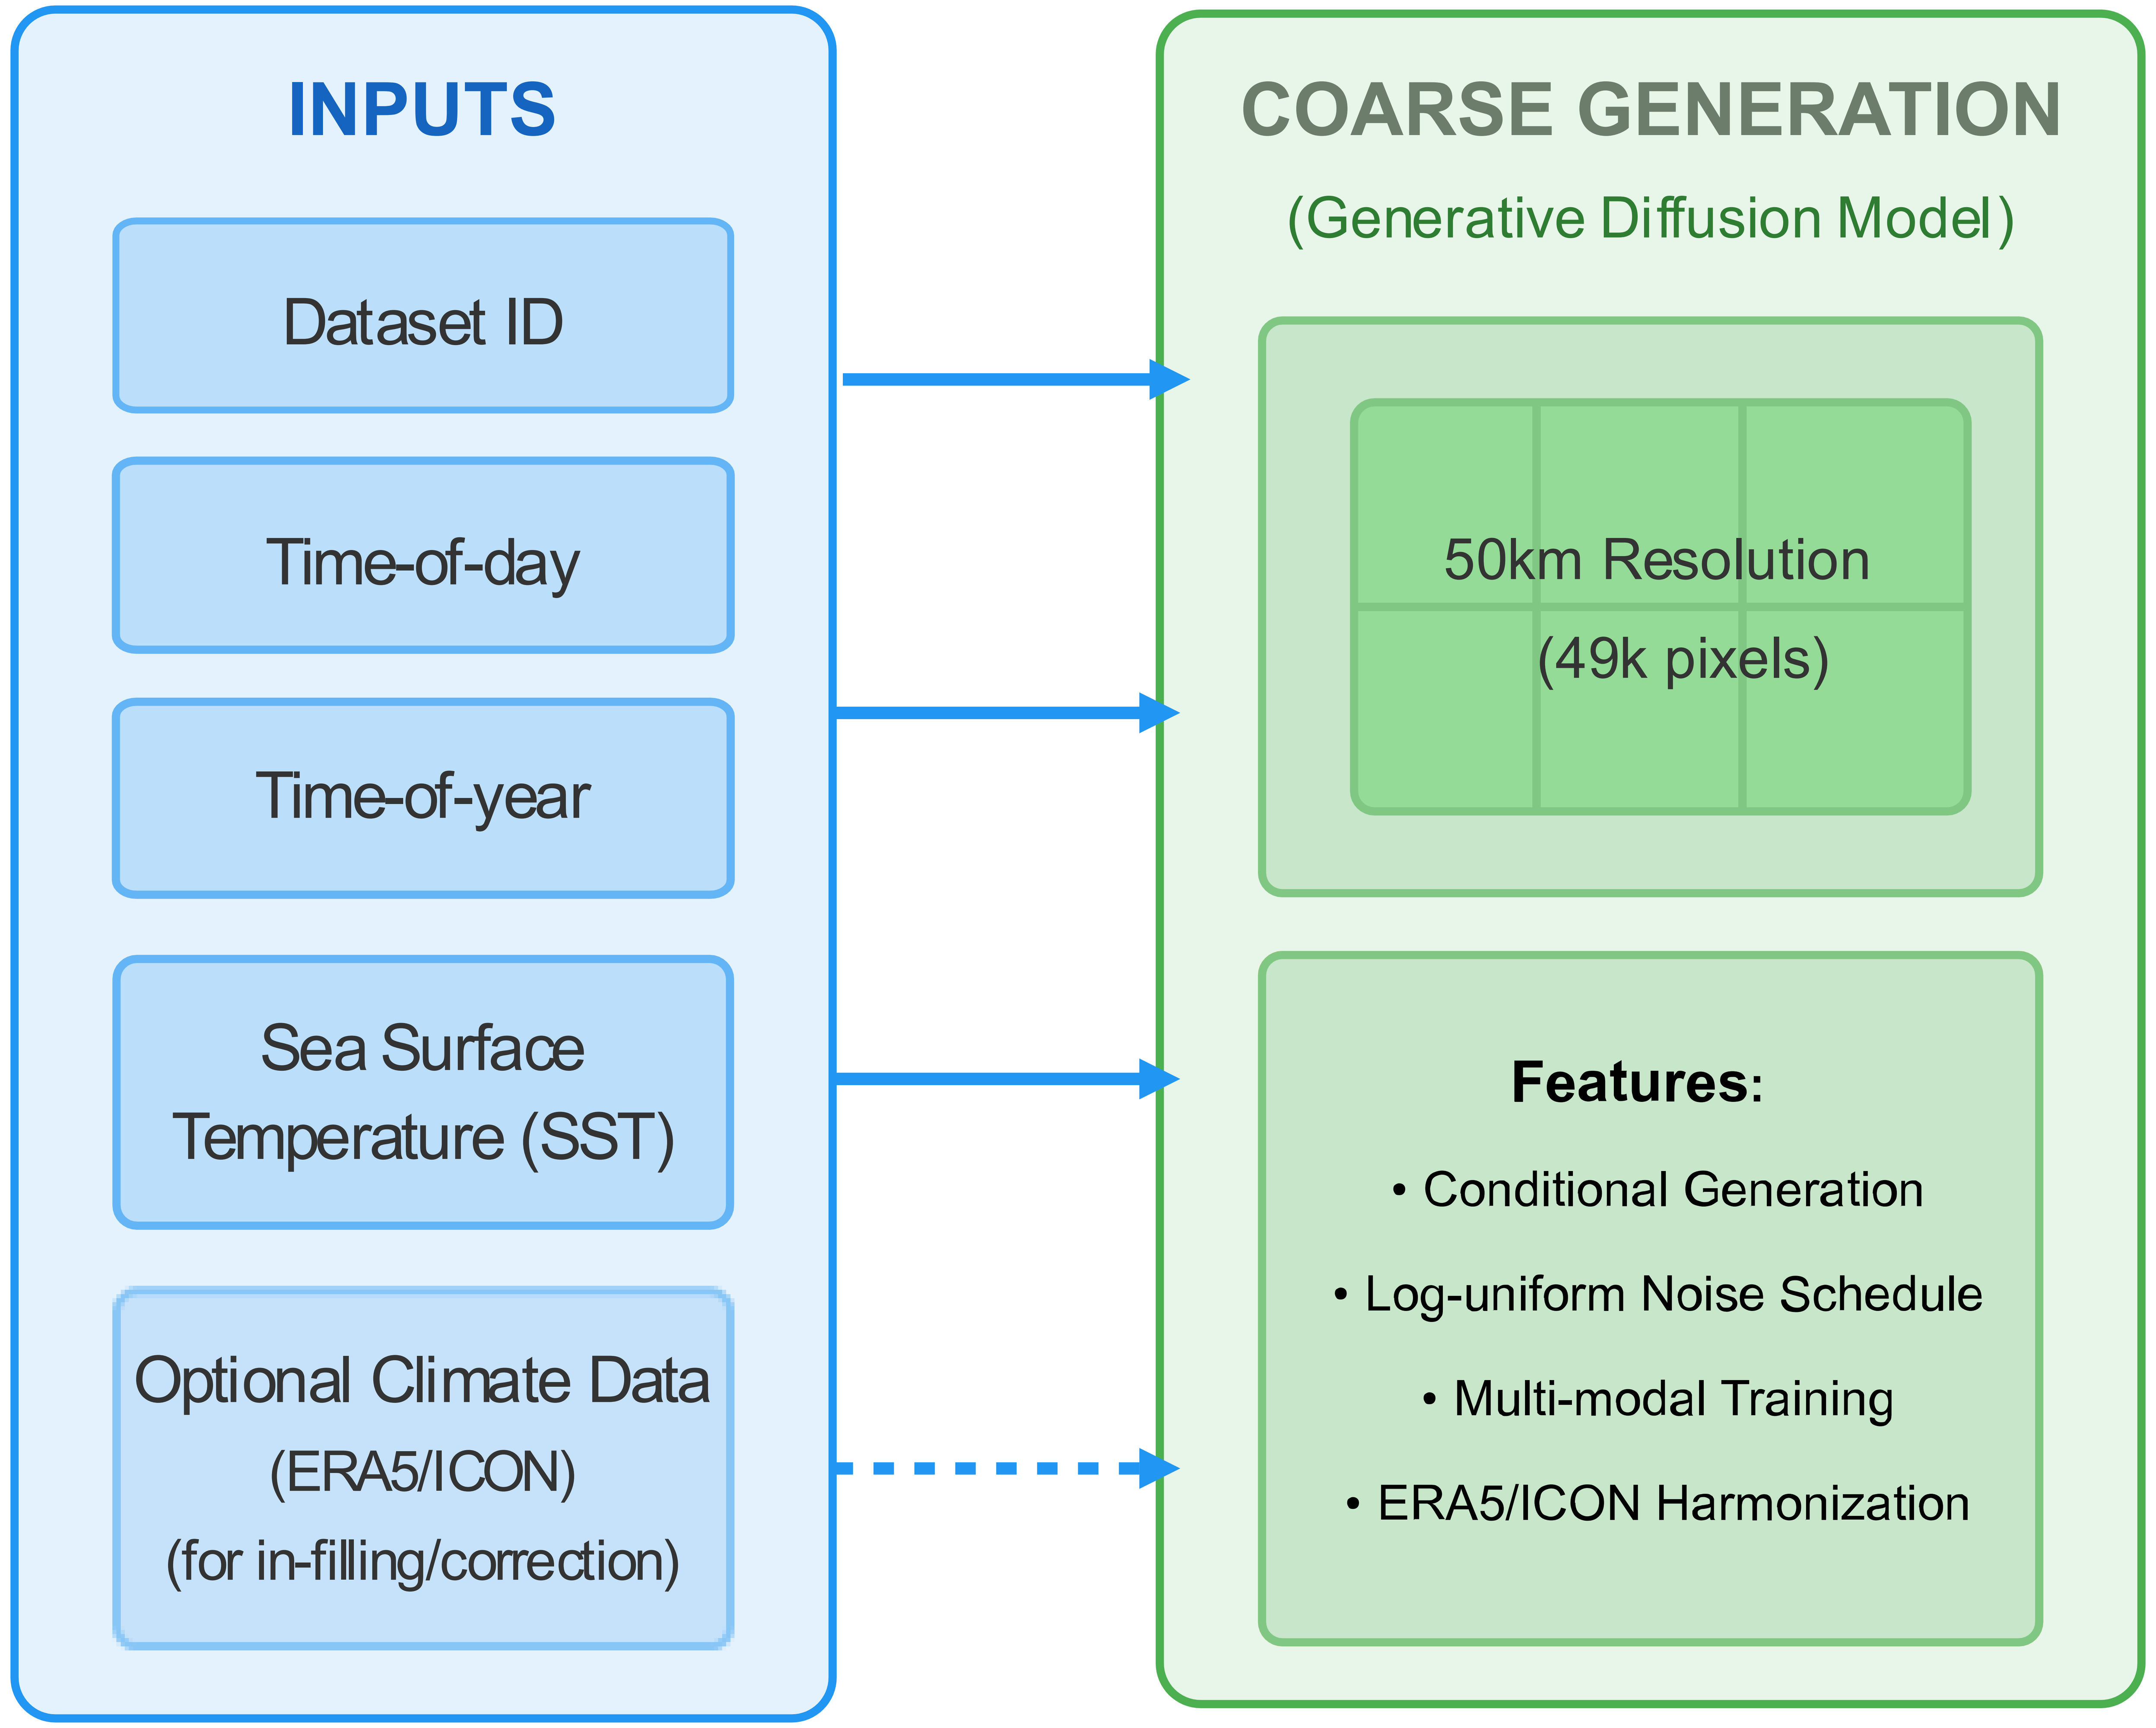

Validating Variability: Tropical Cyclone Frequency HeatMap

Correct Climatological Action Centers

Validating Variability: Detectability of the Diurnal Rainfall Cycle

Signal-to-Noise Ratio of the Estimated Diurnal Rainfall Amplitude

Correct Structure over Tropical Continents & Detectable Stratocumulus Drizzle Cycle

Trends - Climate Change of Heat Waves: Spatial Pattern

SST Conditioning Enough to Capture Secular Trend - though it is underestimated

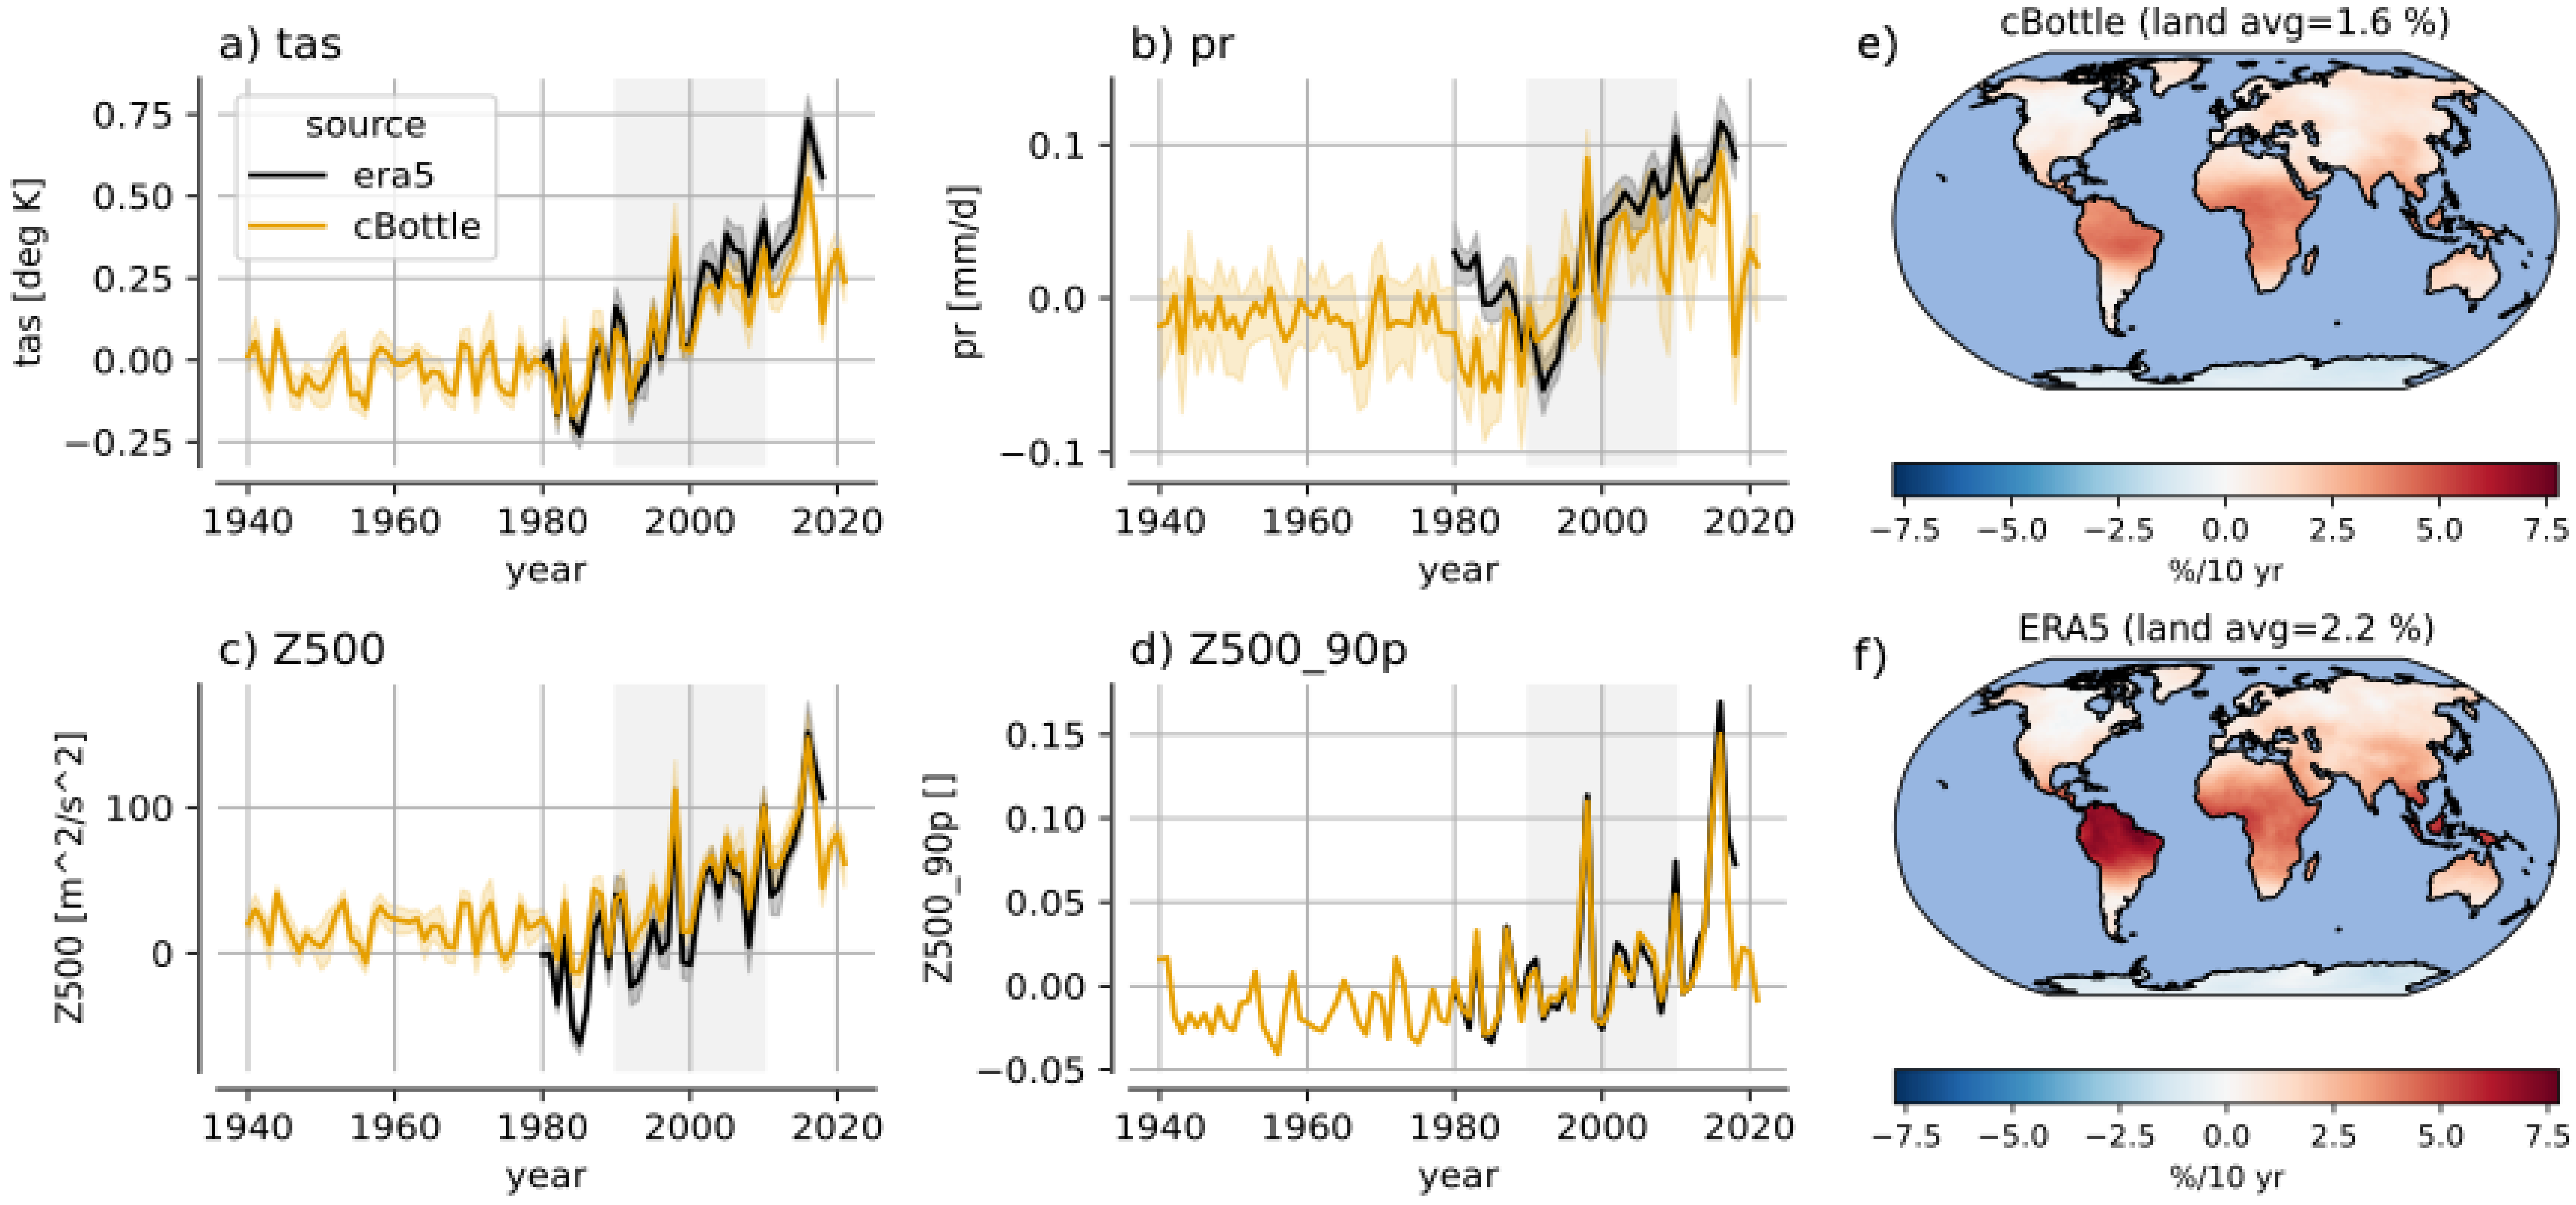

Component 2: Super-Resolution from 50k to 12.5M Pixels

HPX64 → HPX1024

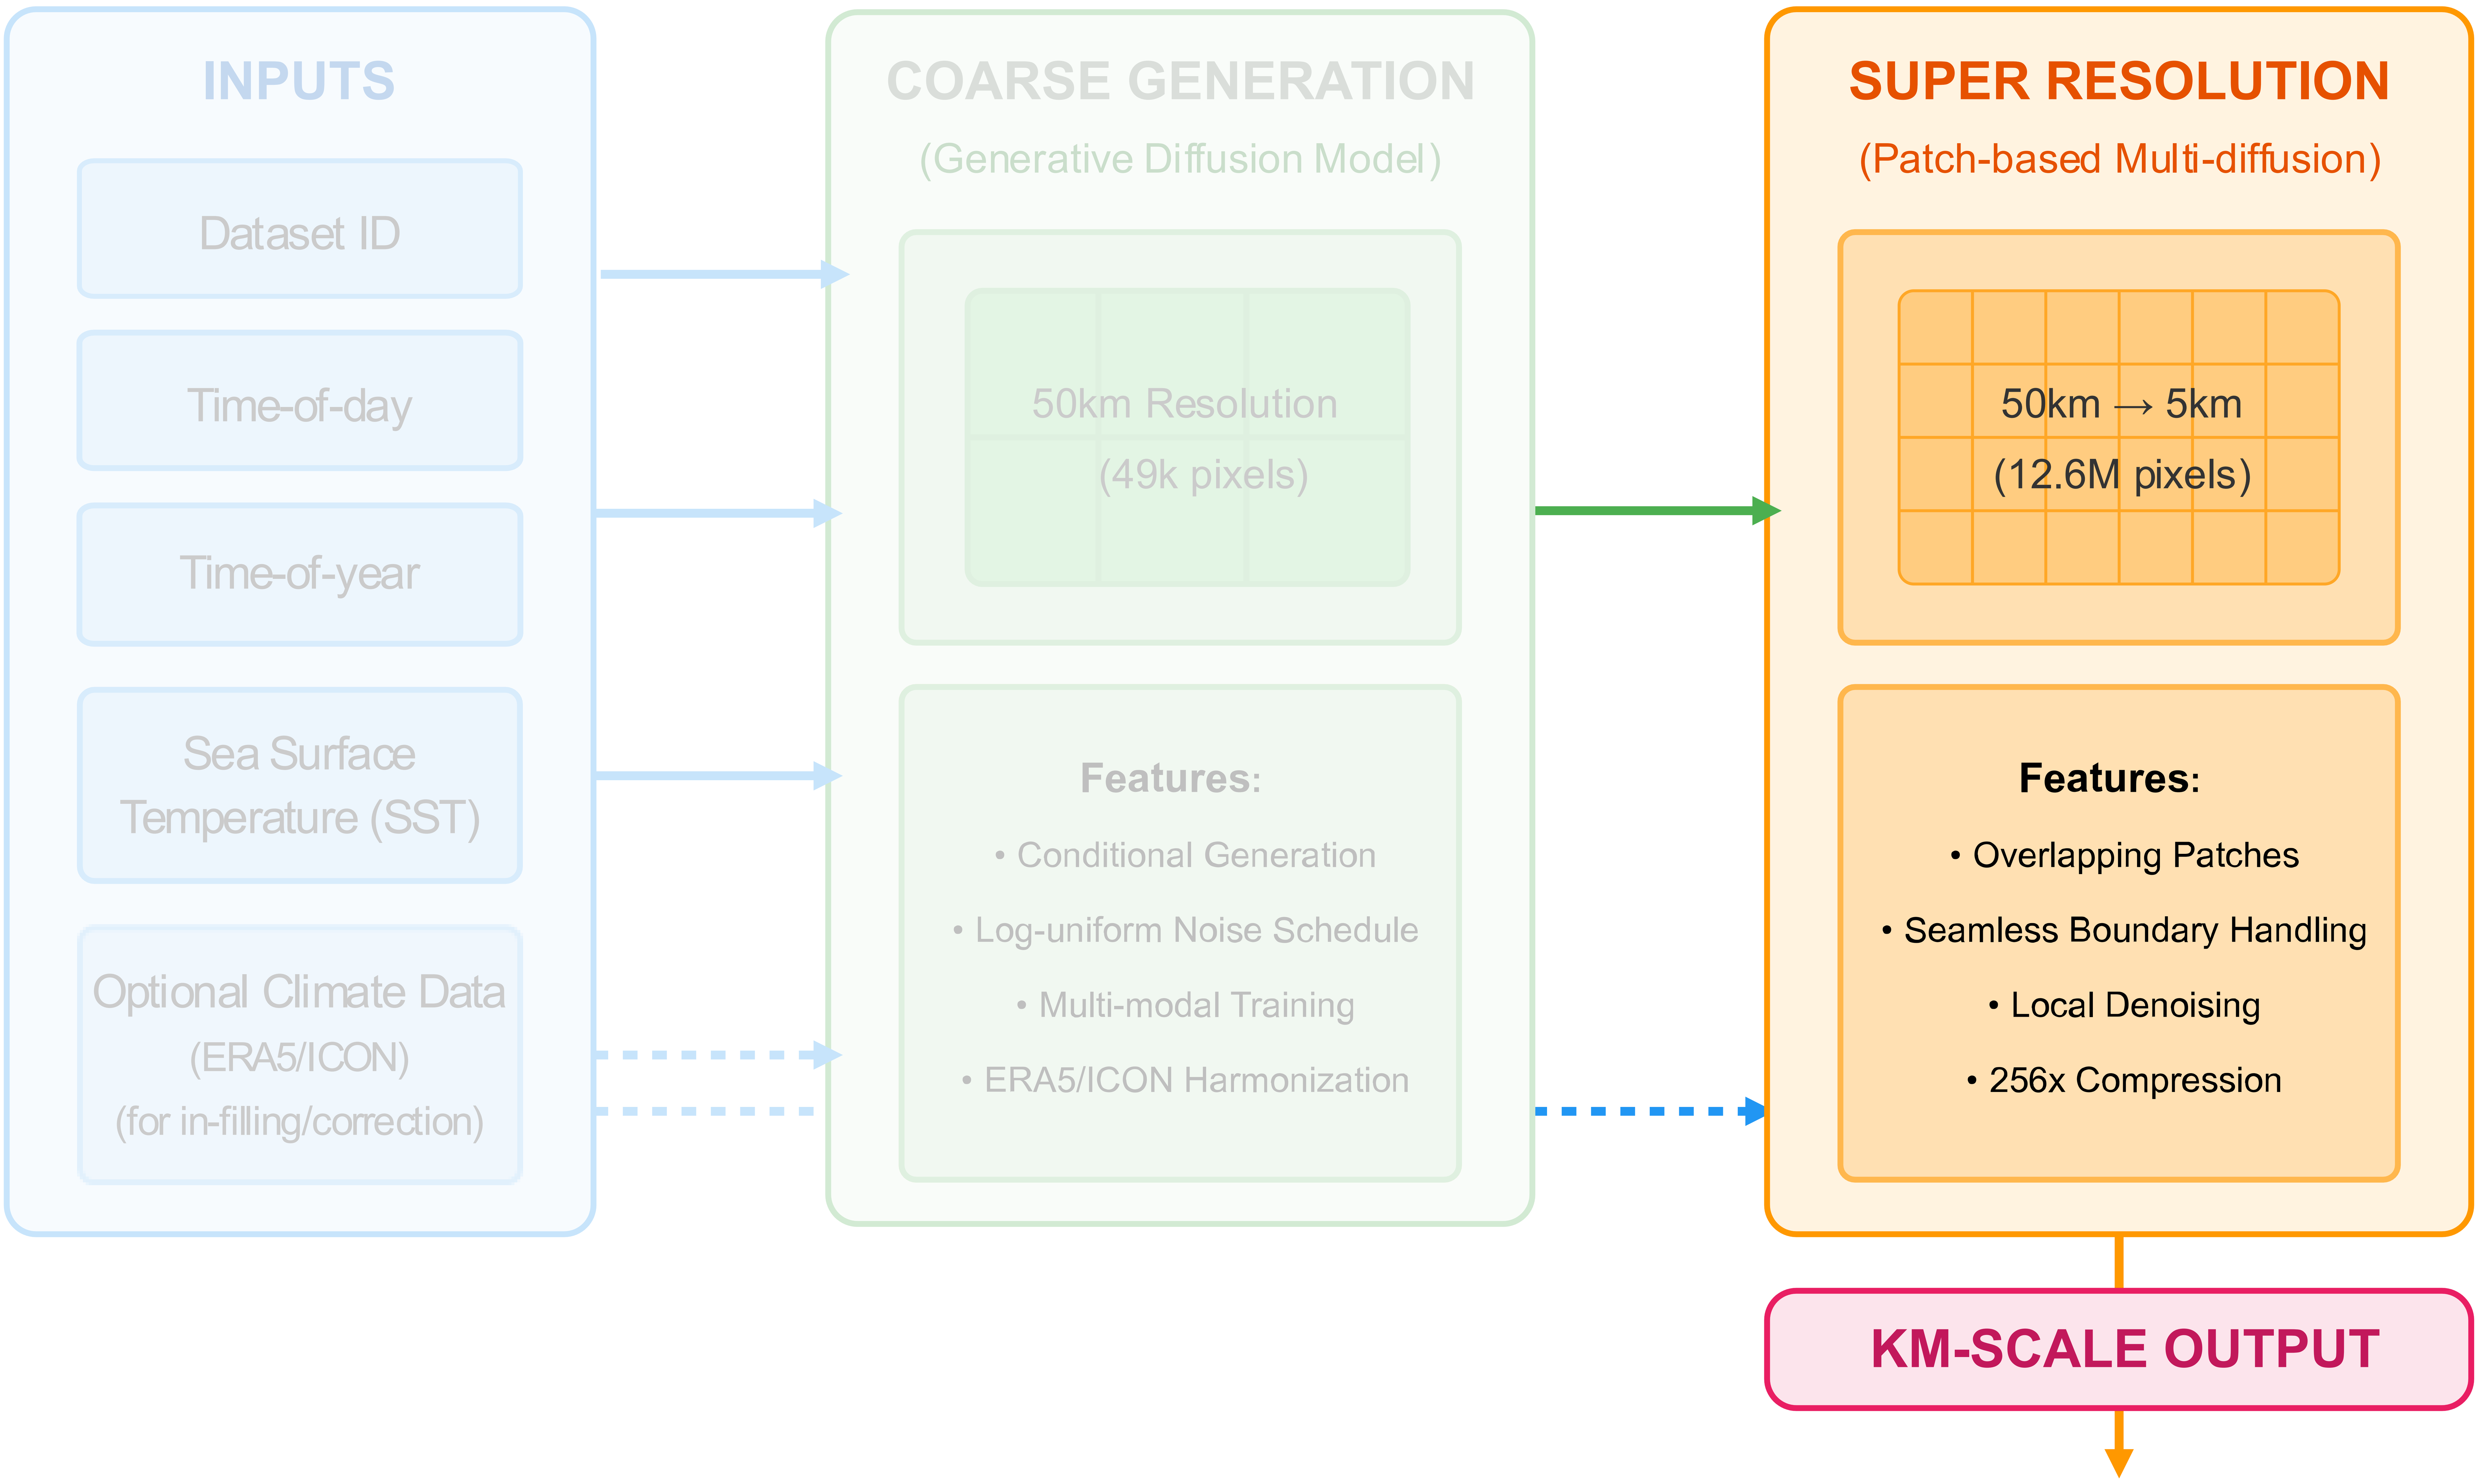

Challenge: 12.5M Pixels is a Lot

Directly Applying a Diffusion Model Would Require Over 2,000 GB of GPU Memory

OLR – from ICON Cycle 3 (HPX1024)

Local Patch-Based Approach

2D multi-diffusion on HEALPix patches for coherent generation

(a)

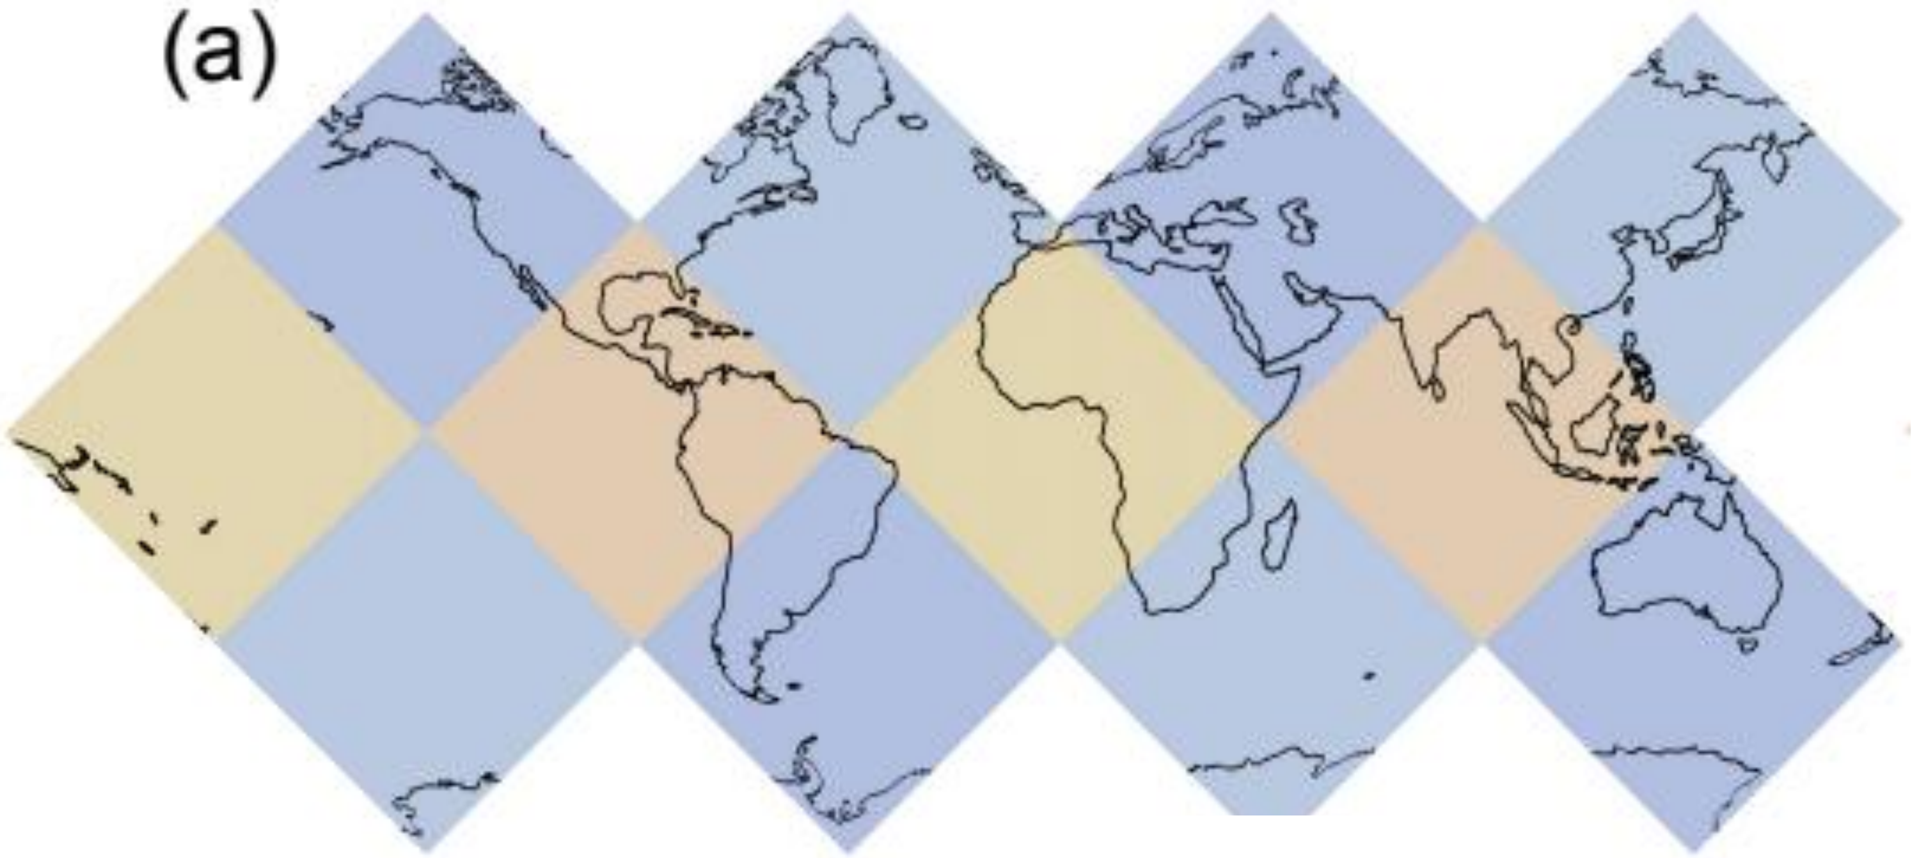

(f)

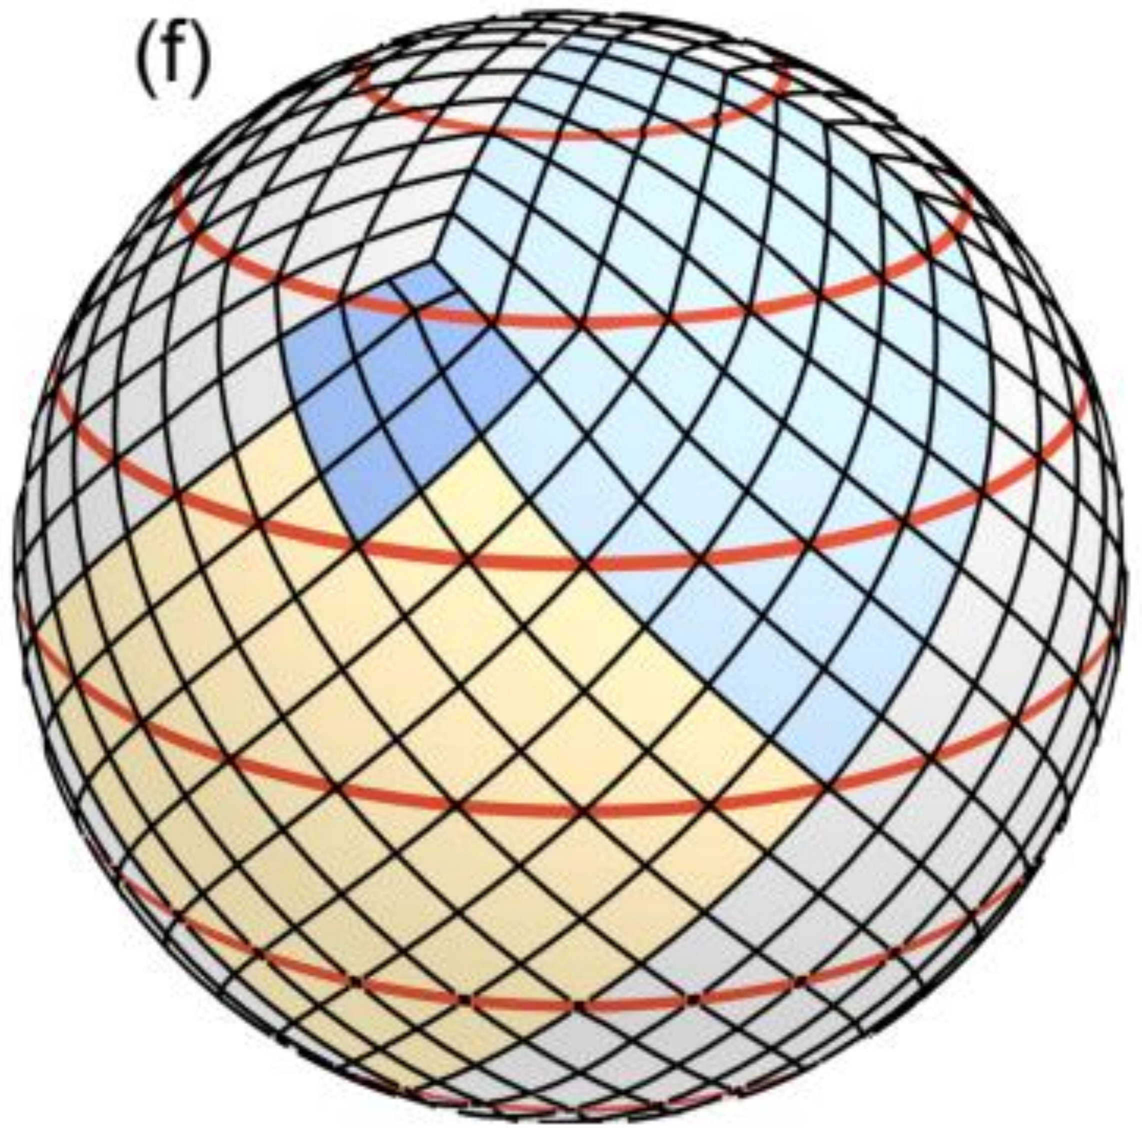

(g)

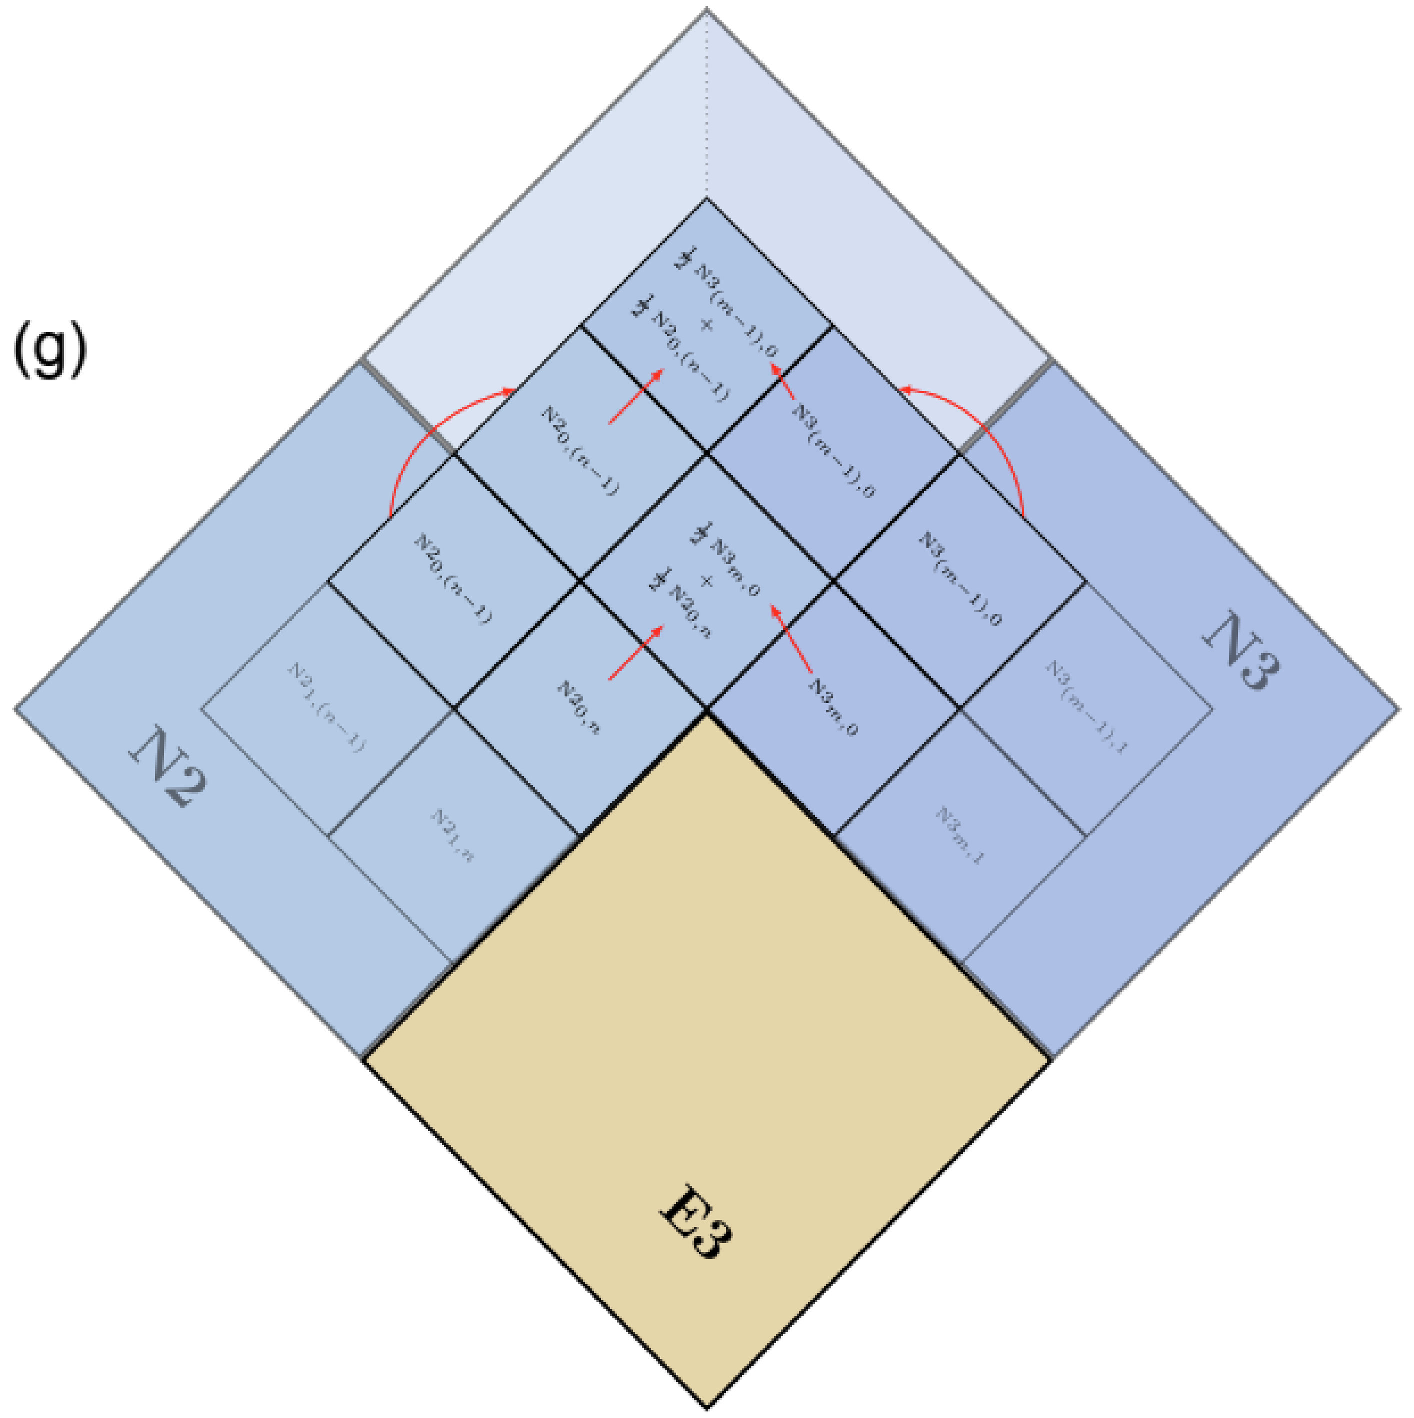

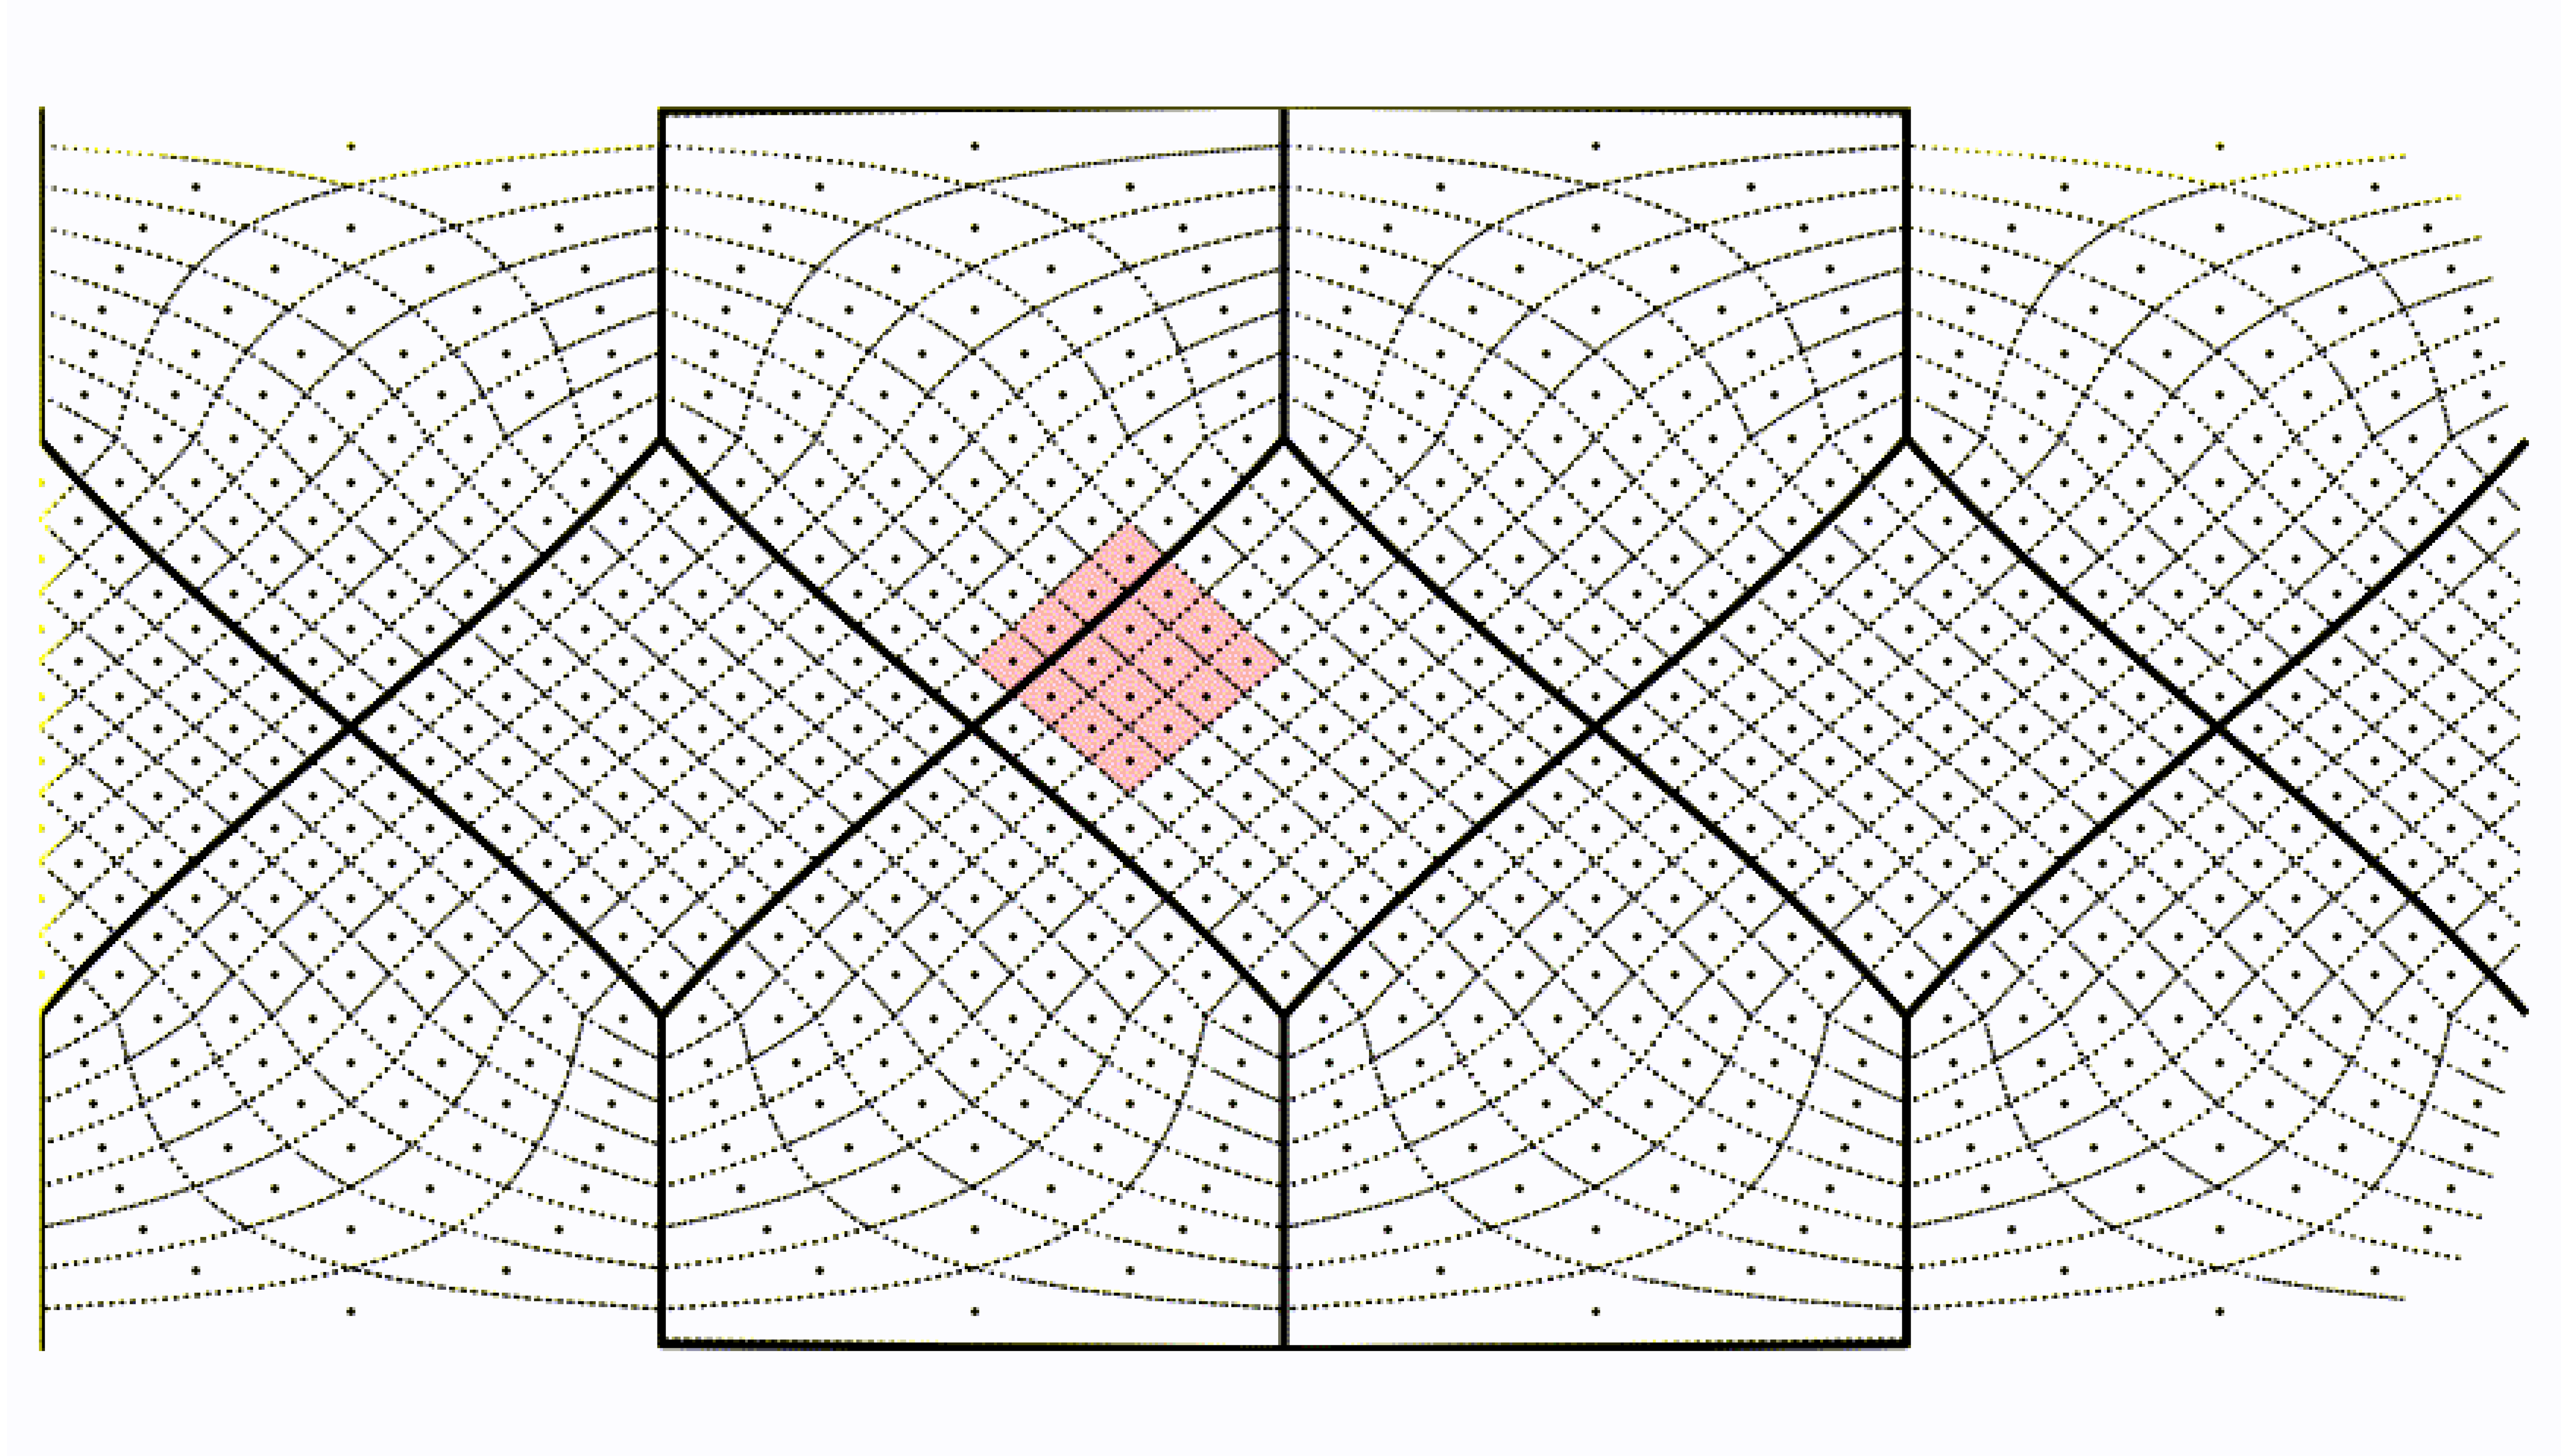

Putting it Together

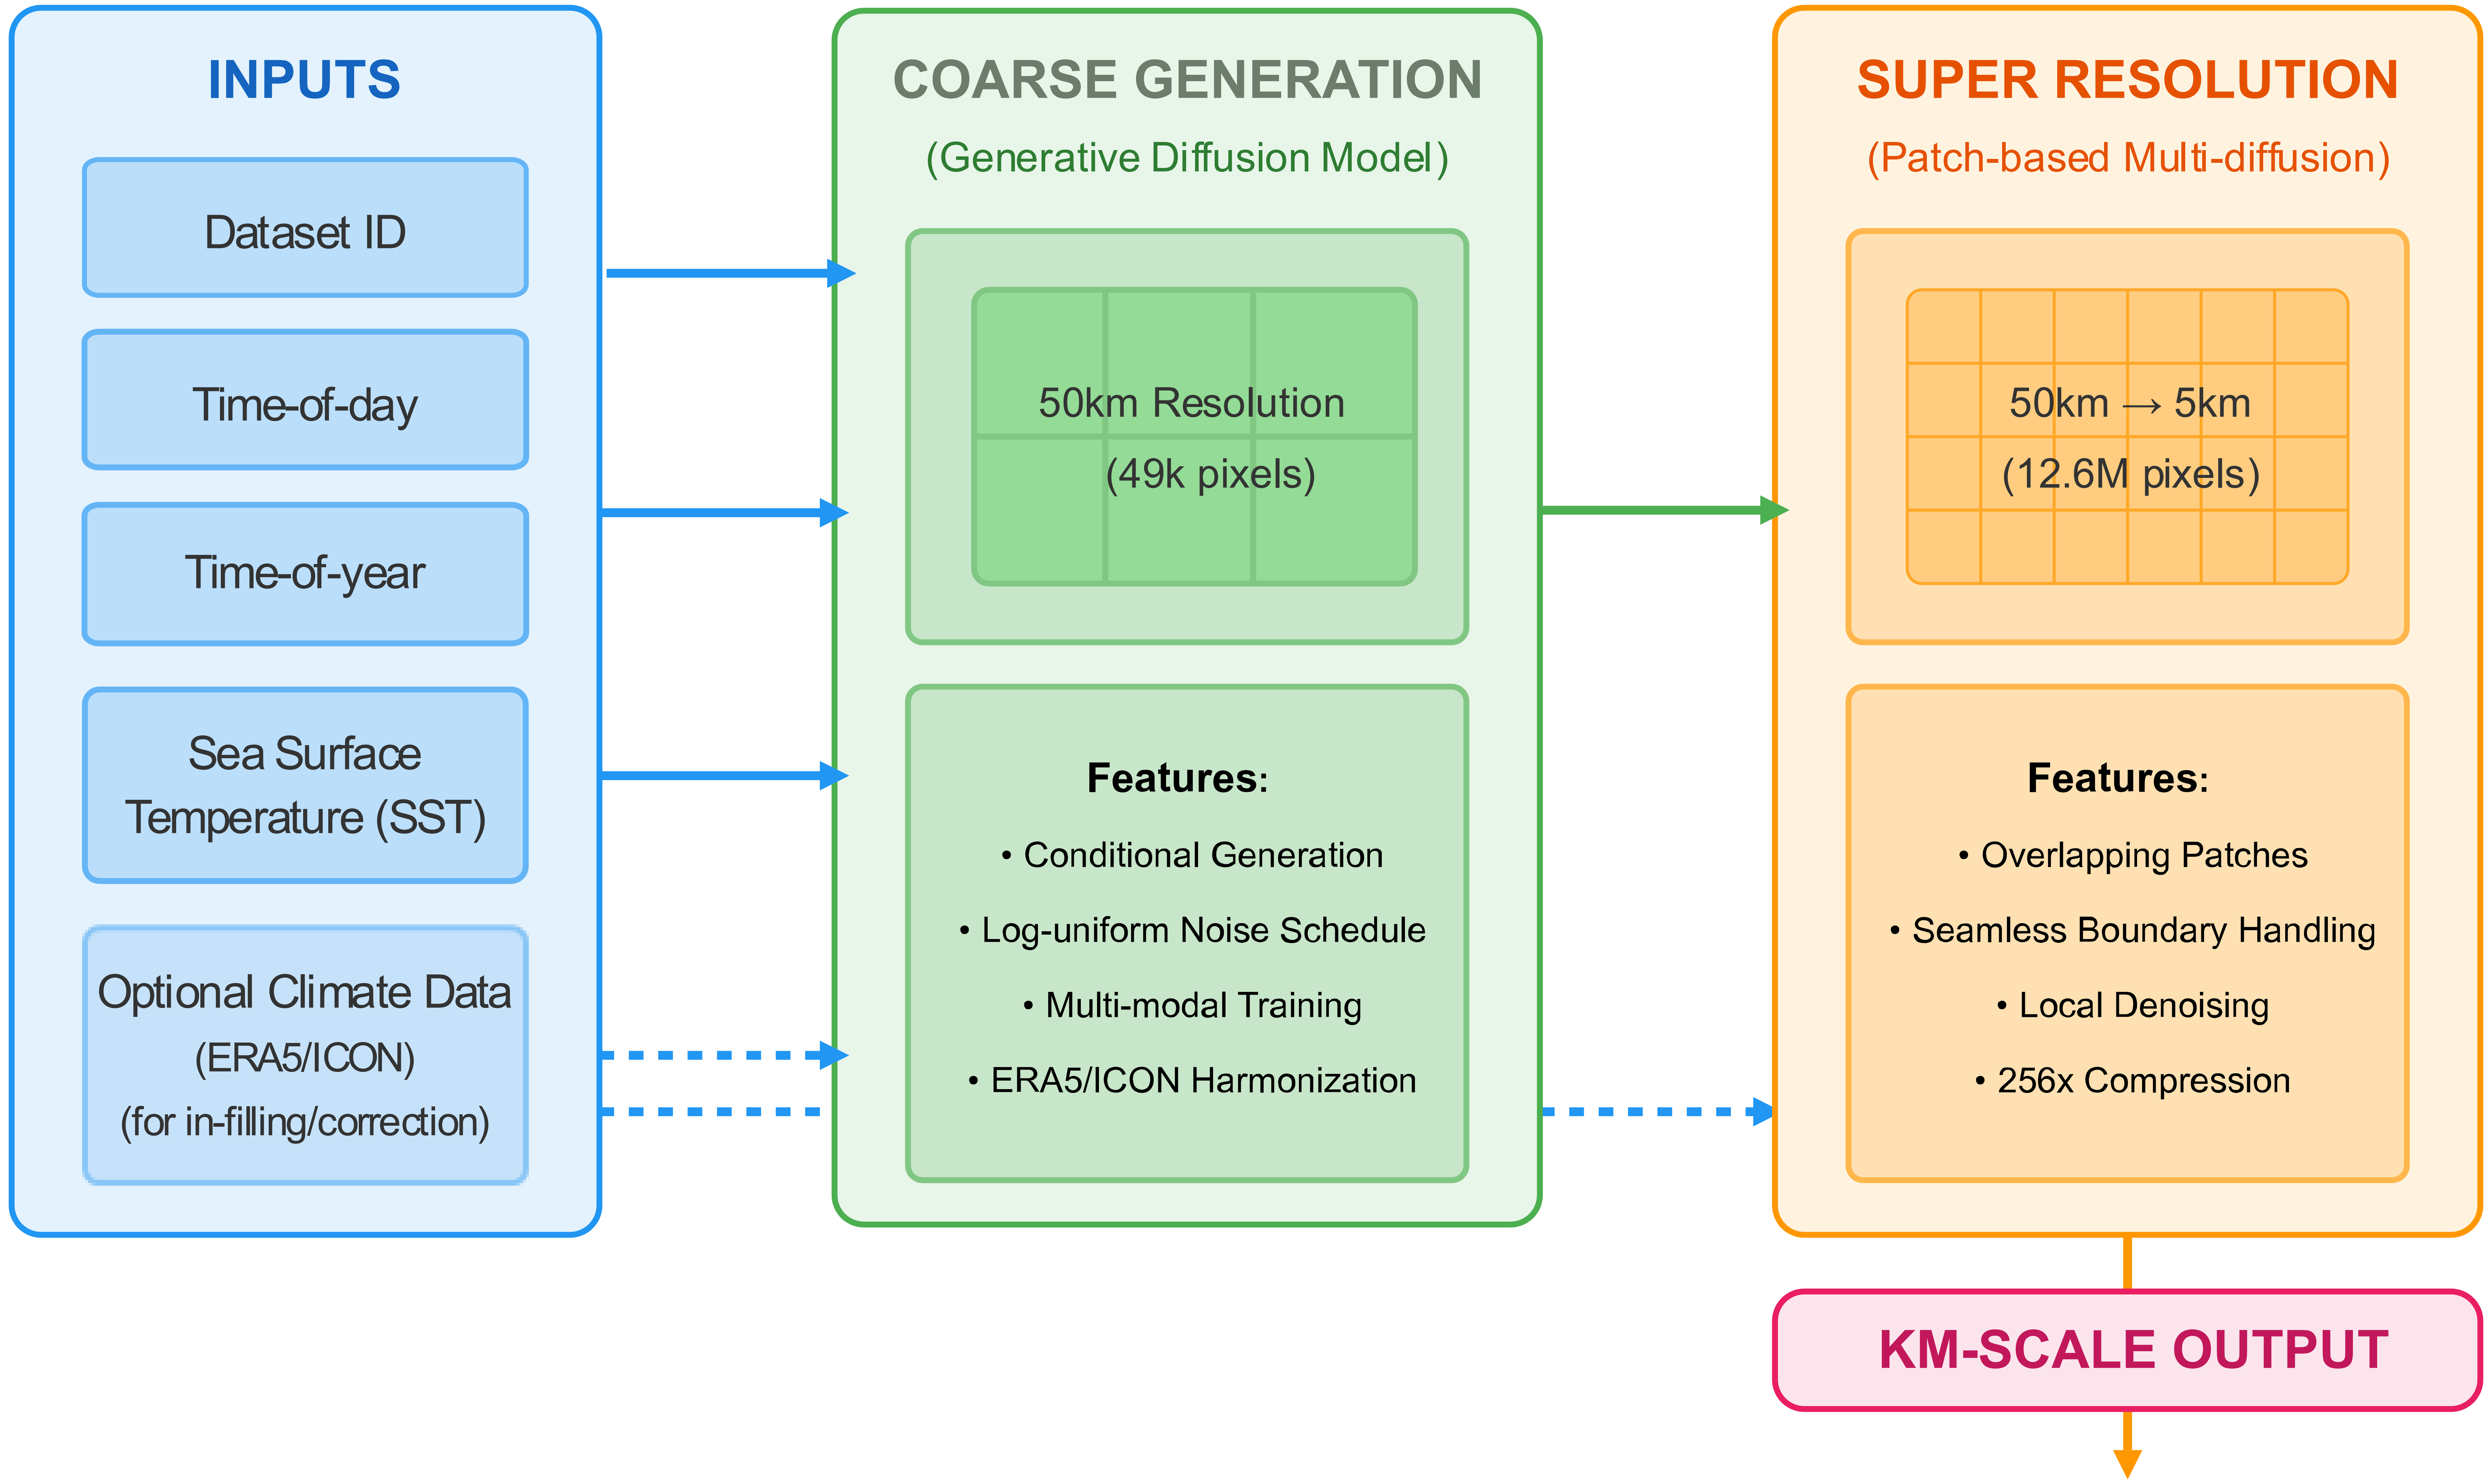

Multi-Modal Inference Tricks

In the theme of Foundation Modeling

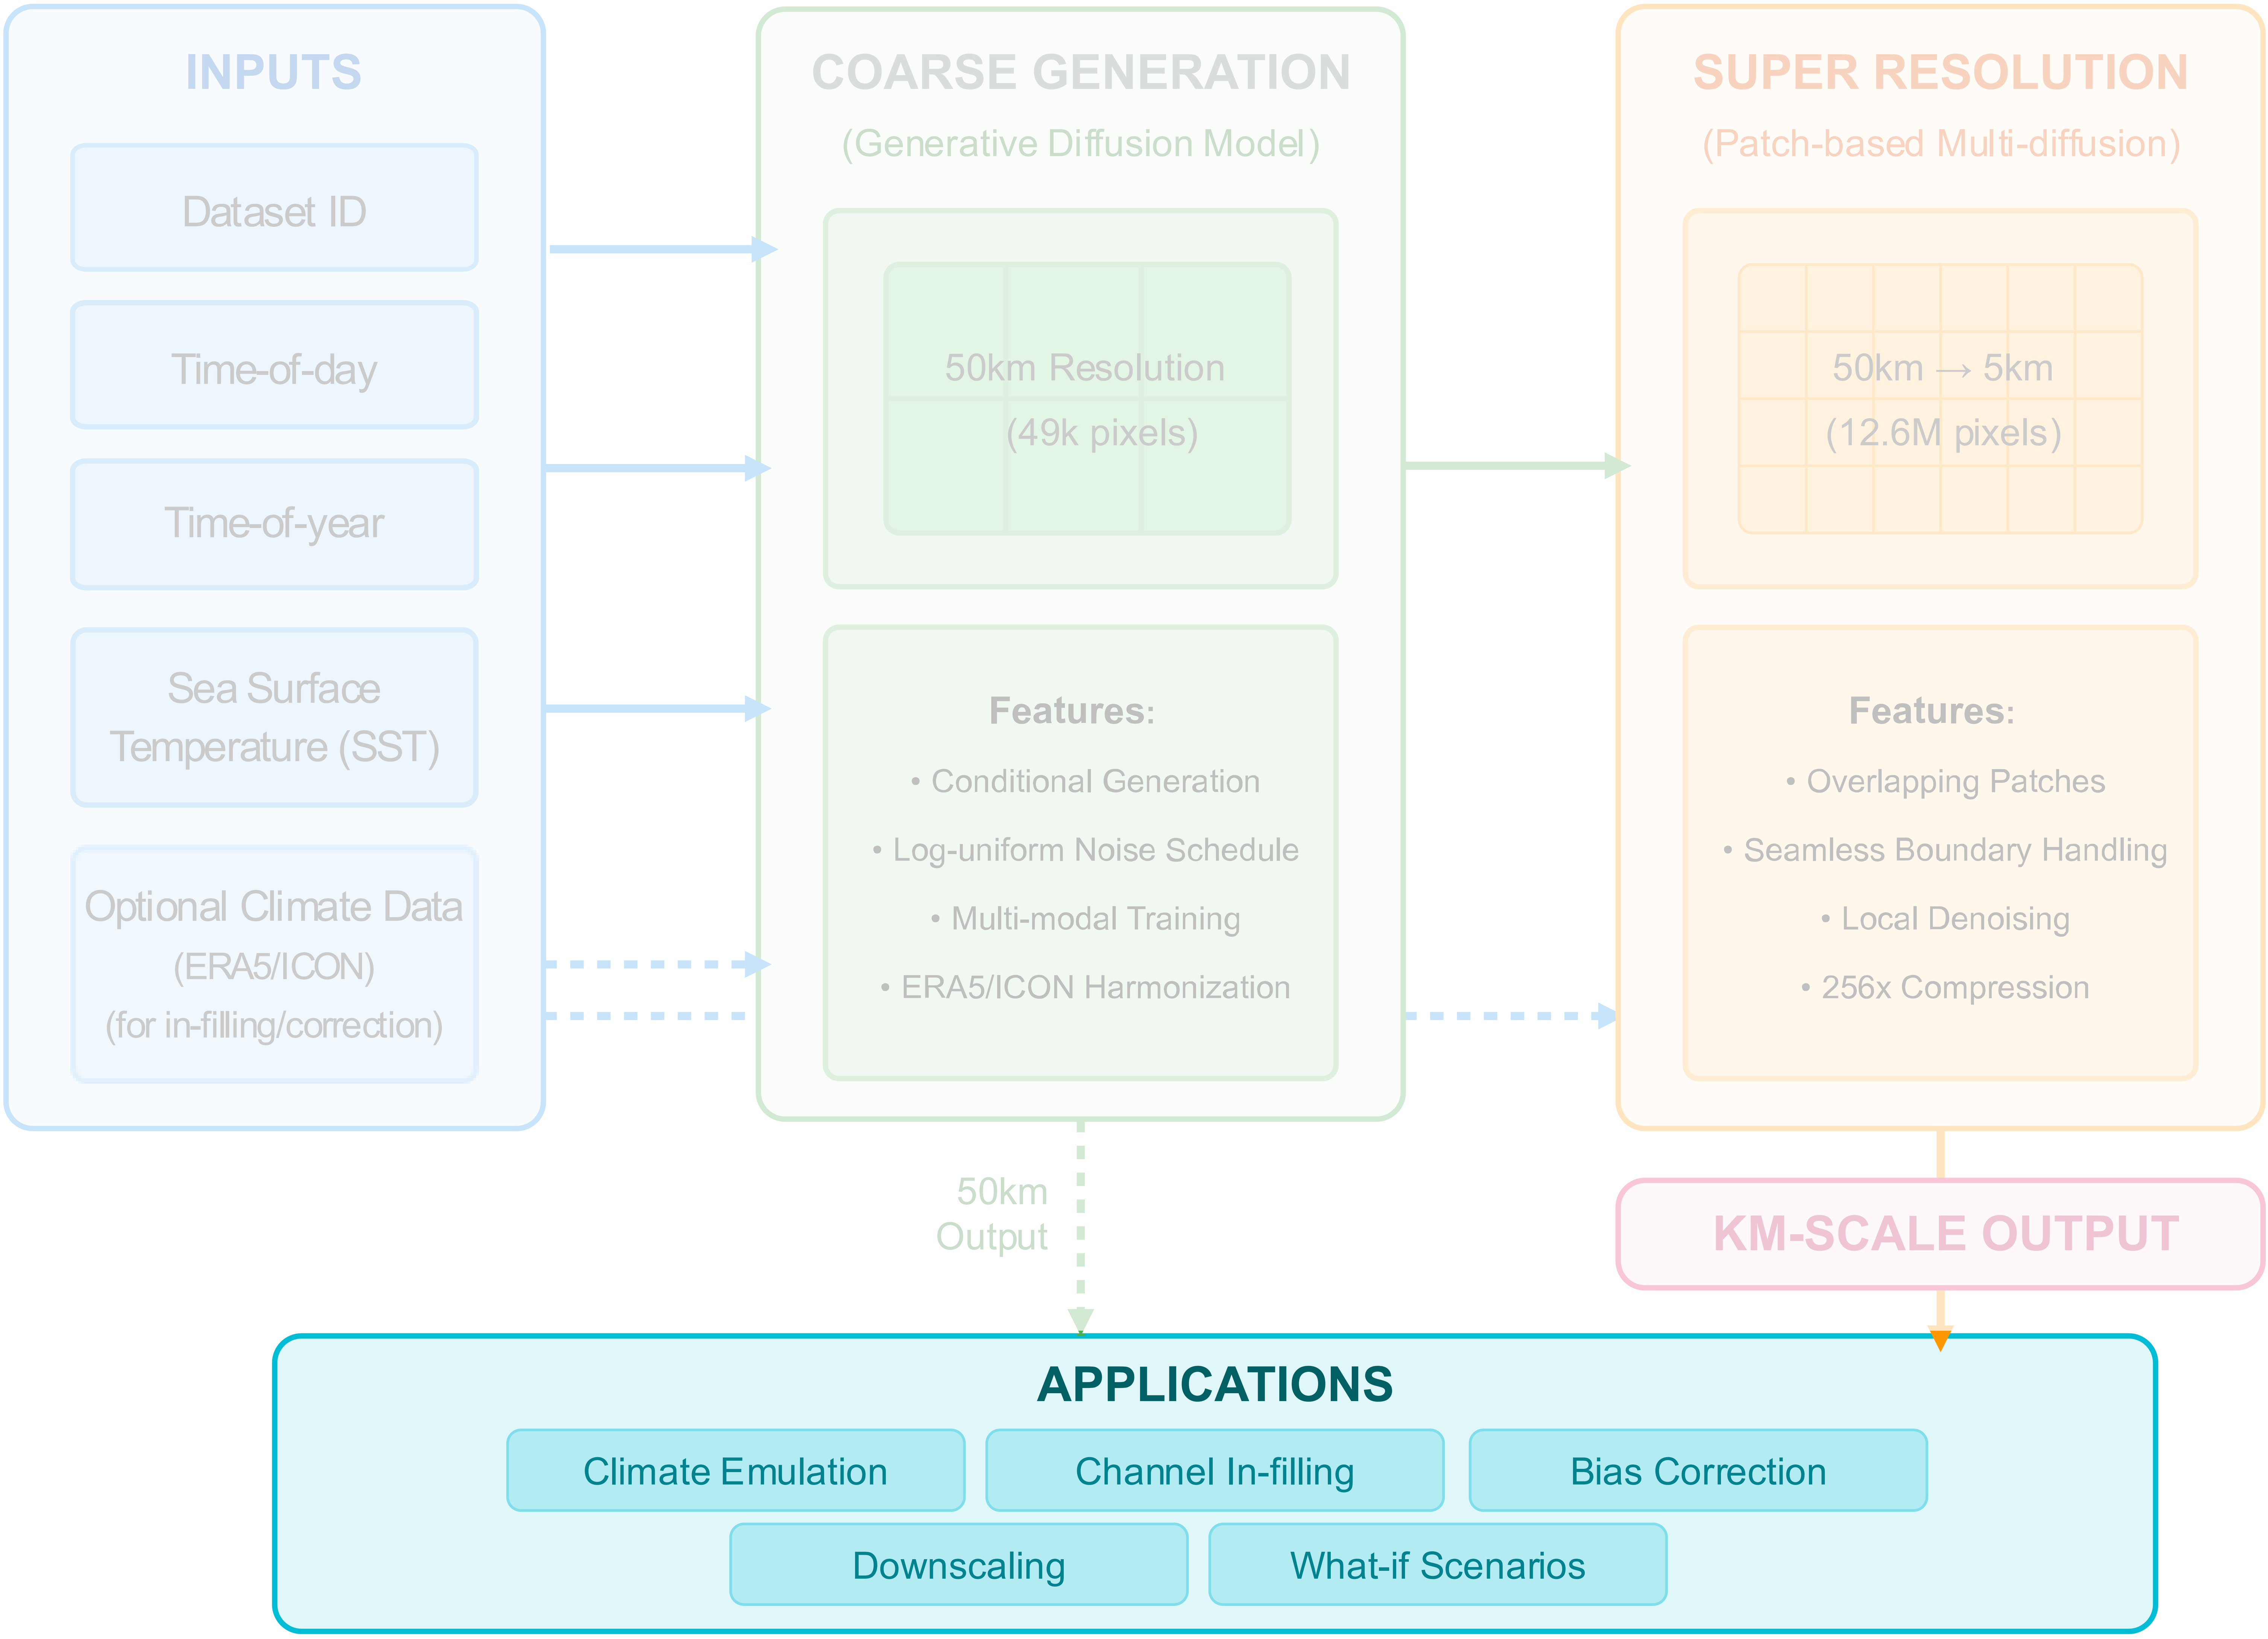

NVIDIA

Task 1. Downscaling: Adding ICON-learnt Cloud Textures to ERA5

Upwelling Solar Radiation

1. ERA5 HPX64 Input

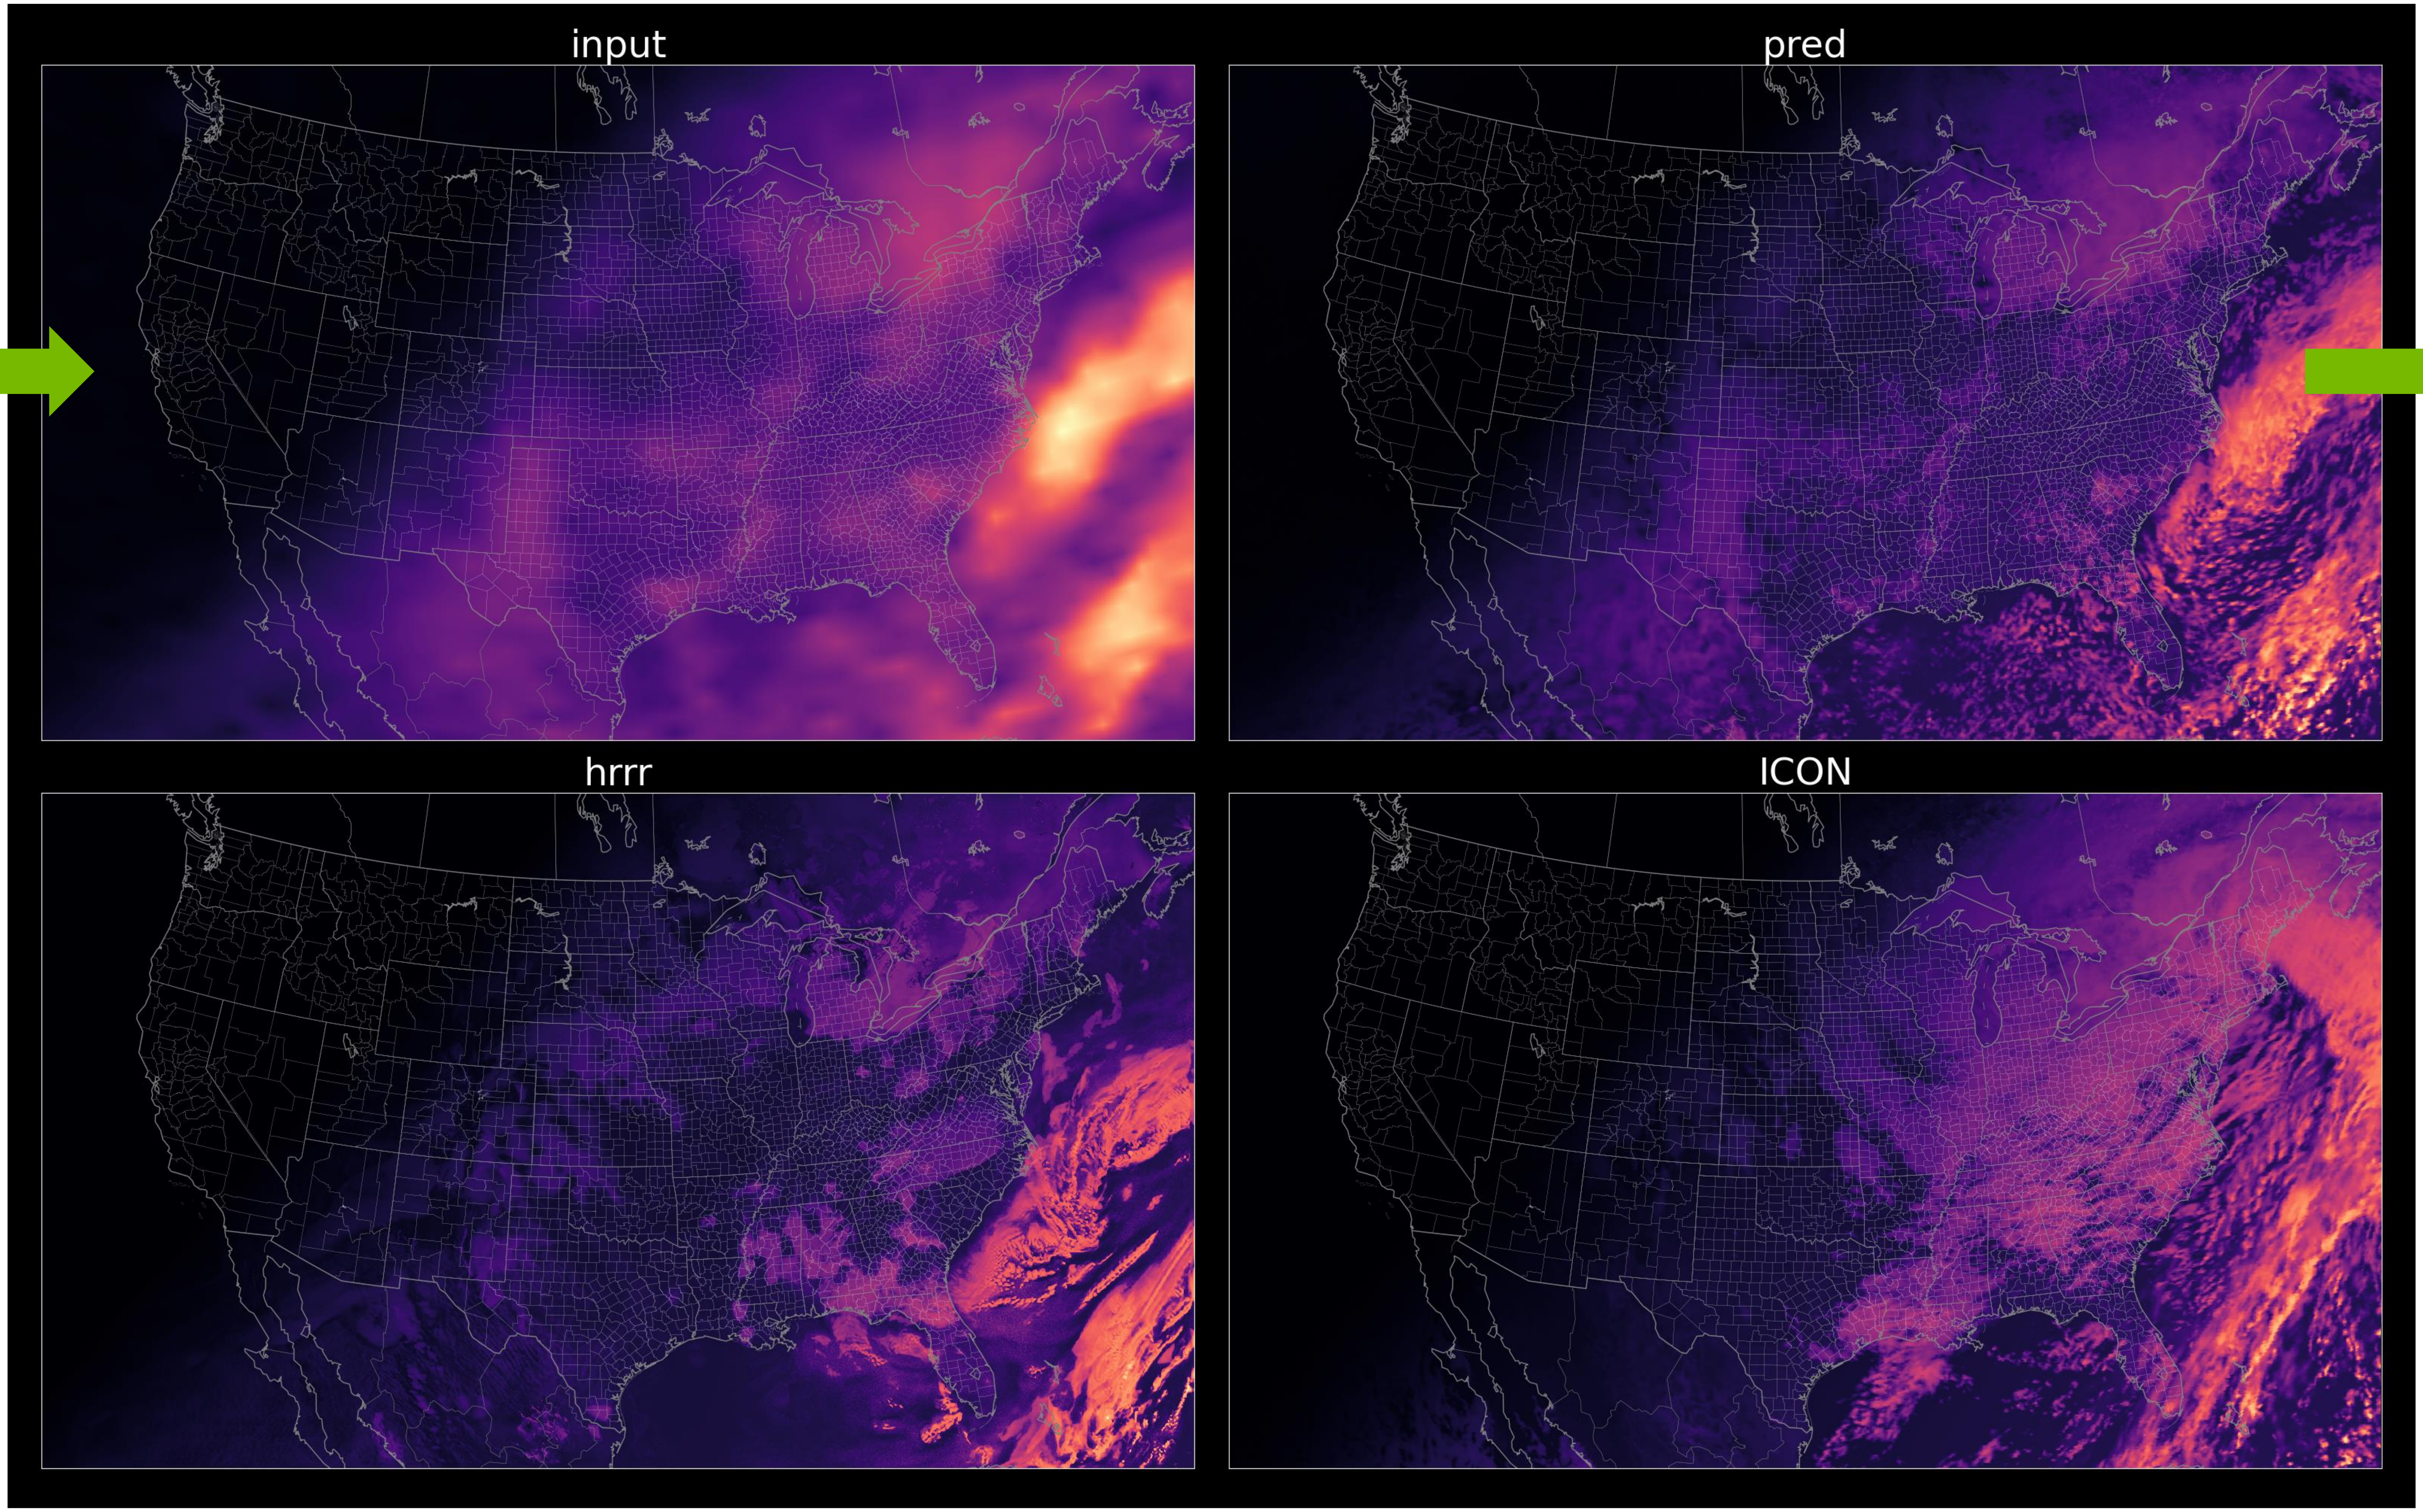

2. cBottle in-paints km-scale detail

Task 2. In-Filling ERA5's Corrupted Shortwave Radiation Channel

Using Data Learnt from ICON Modality

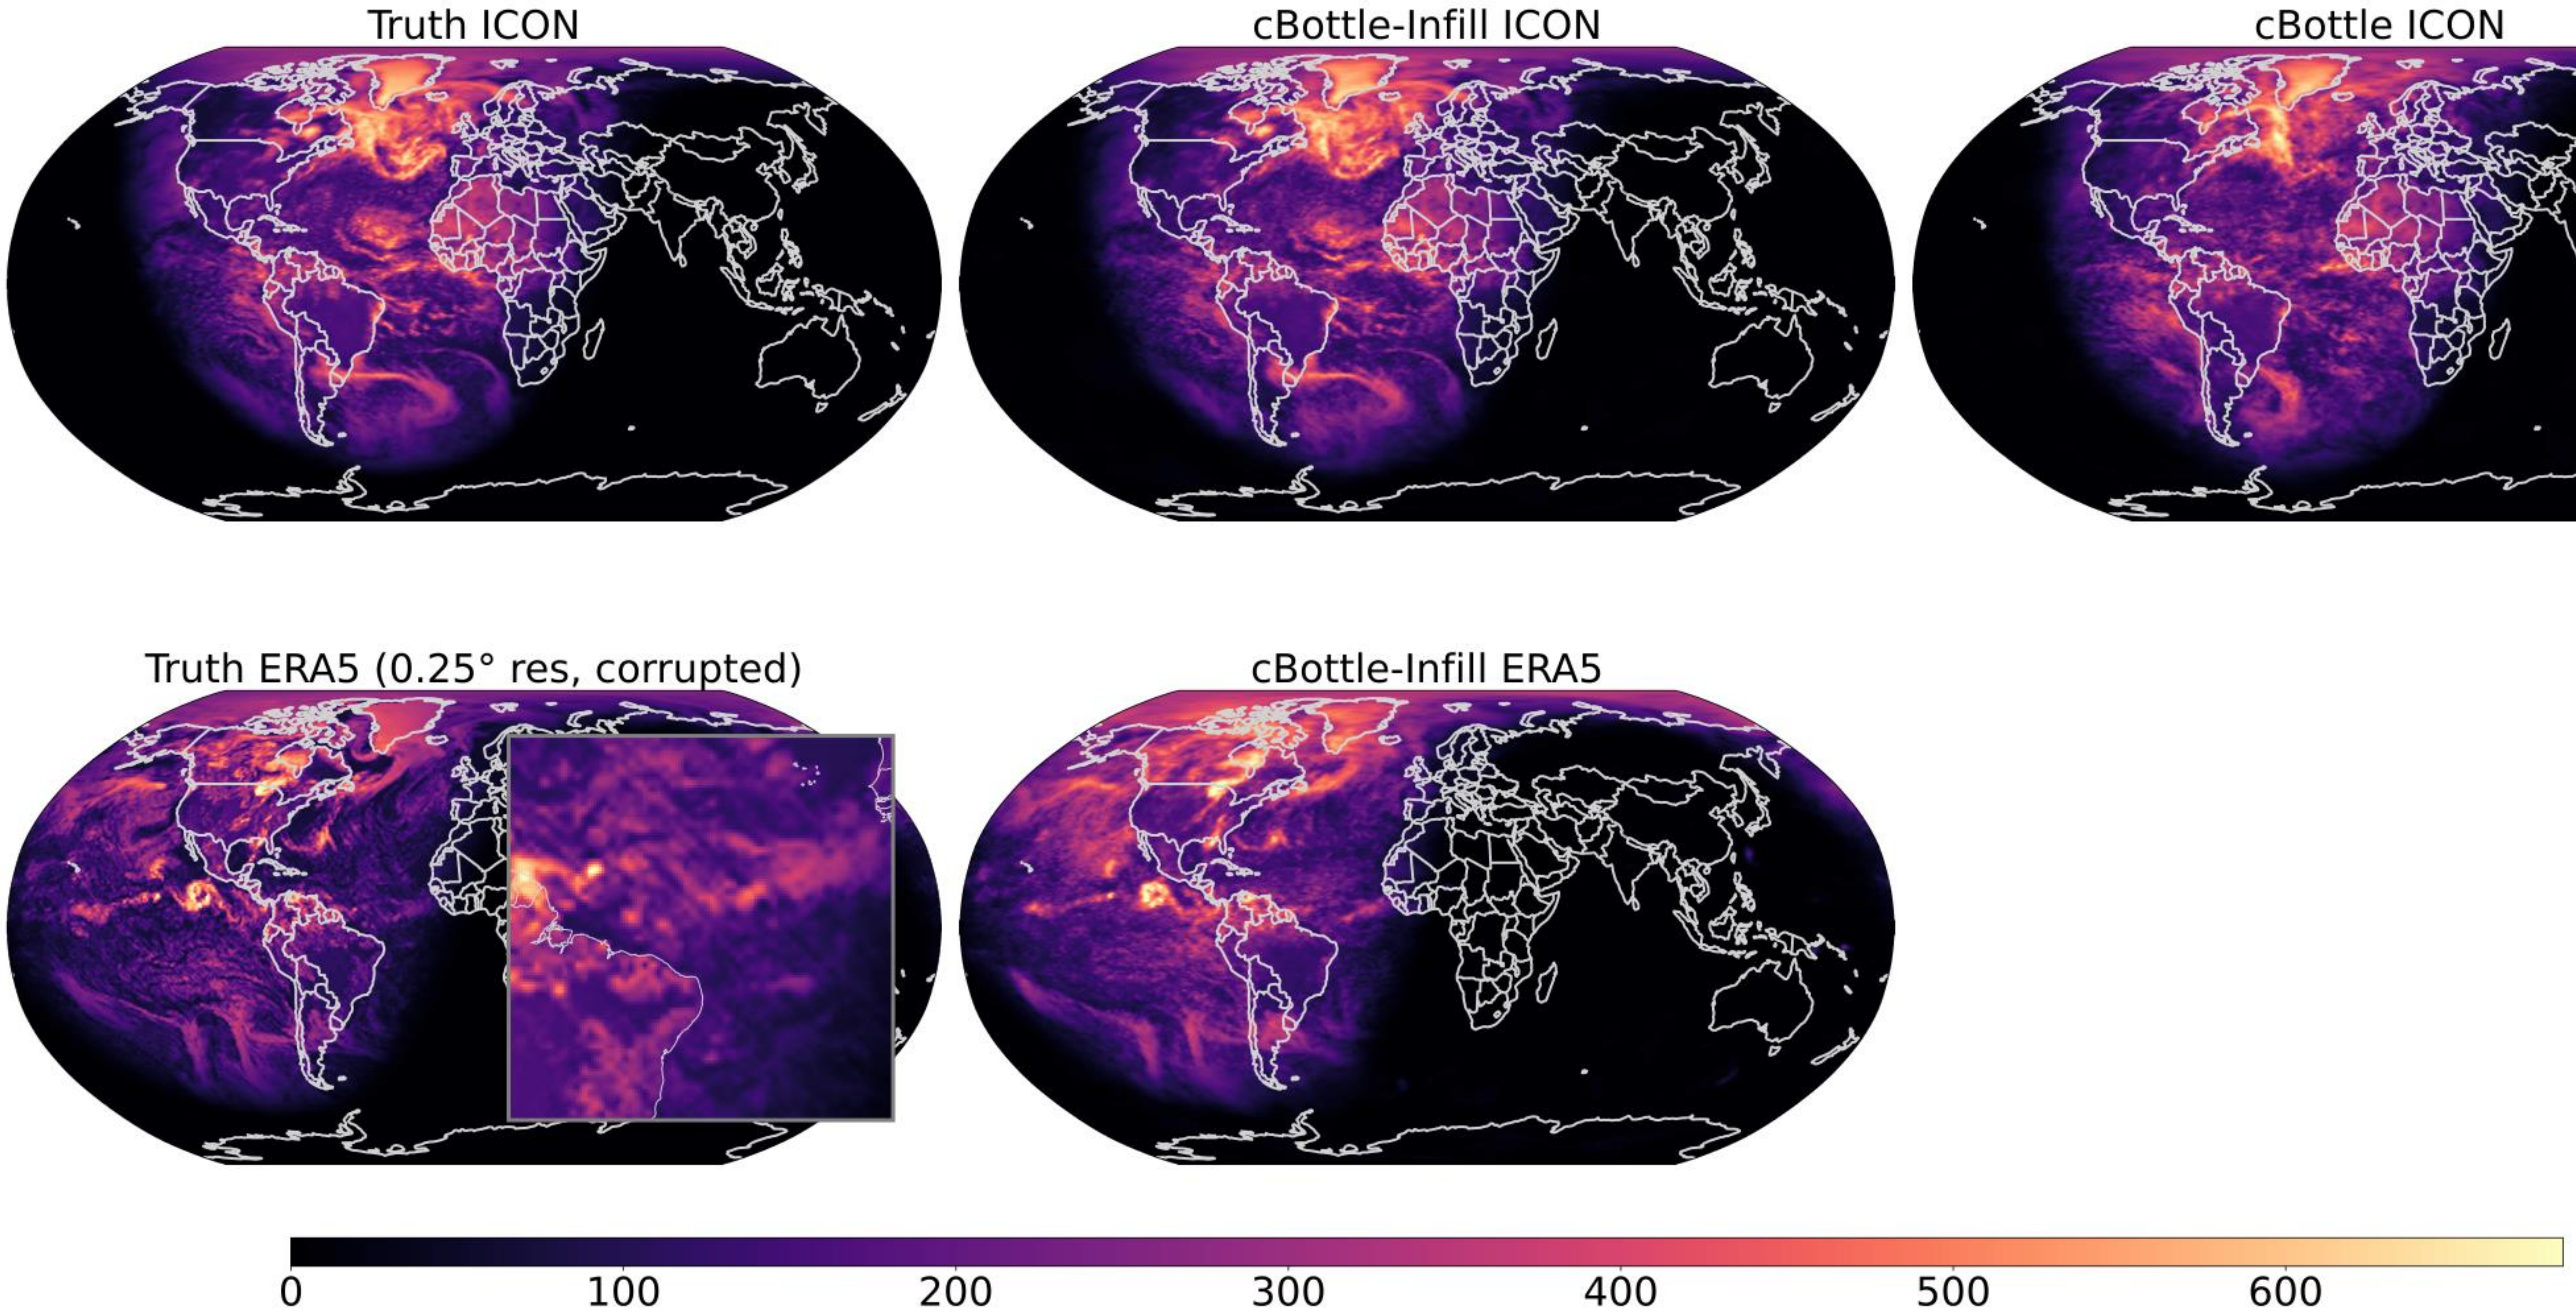

Task 3. De-Biasing ICON’s Liquid Cloud Water Statistics

While Preserving Its Weather Patterns. Using Data Learnt from the ERA5 Modality

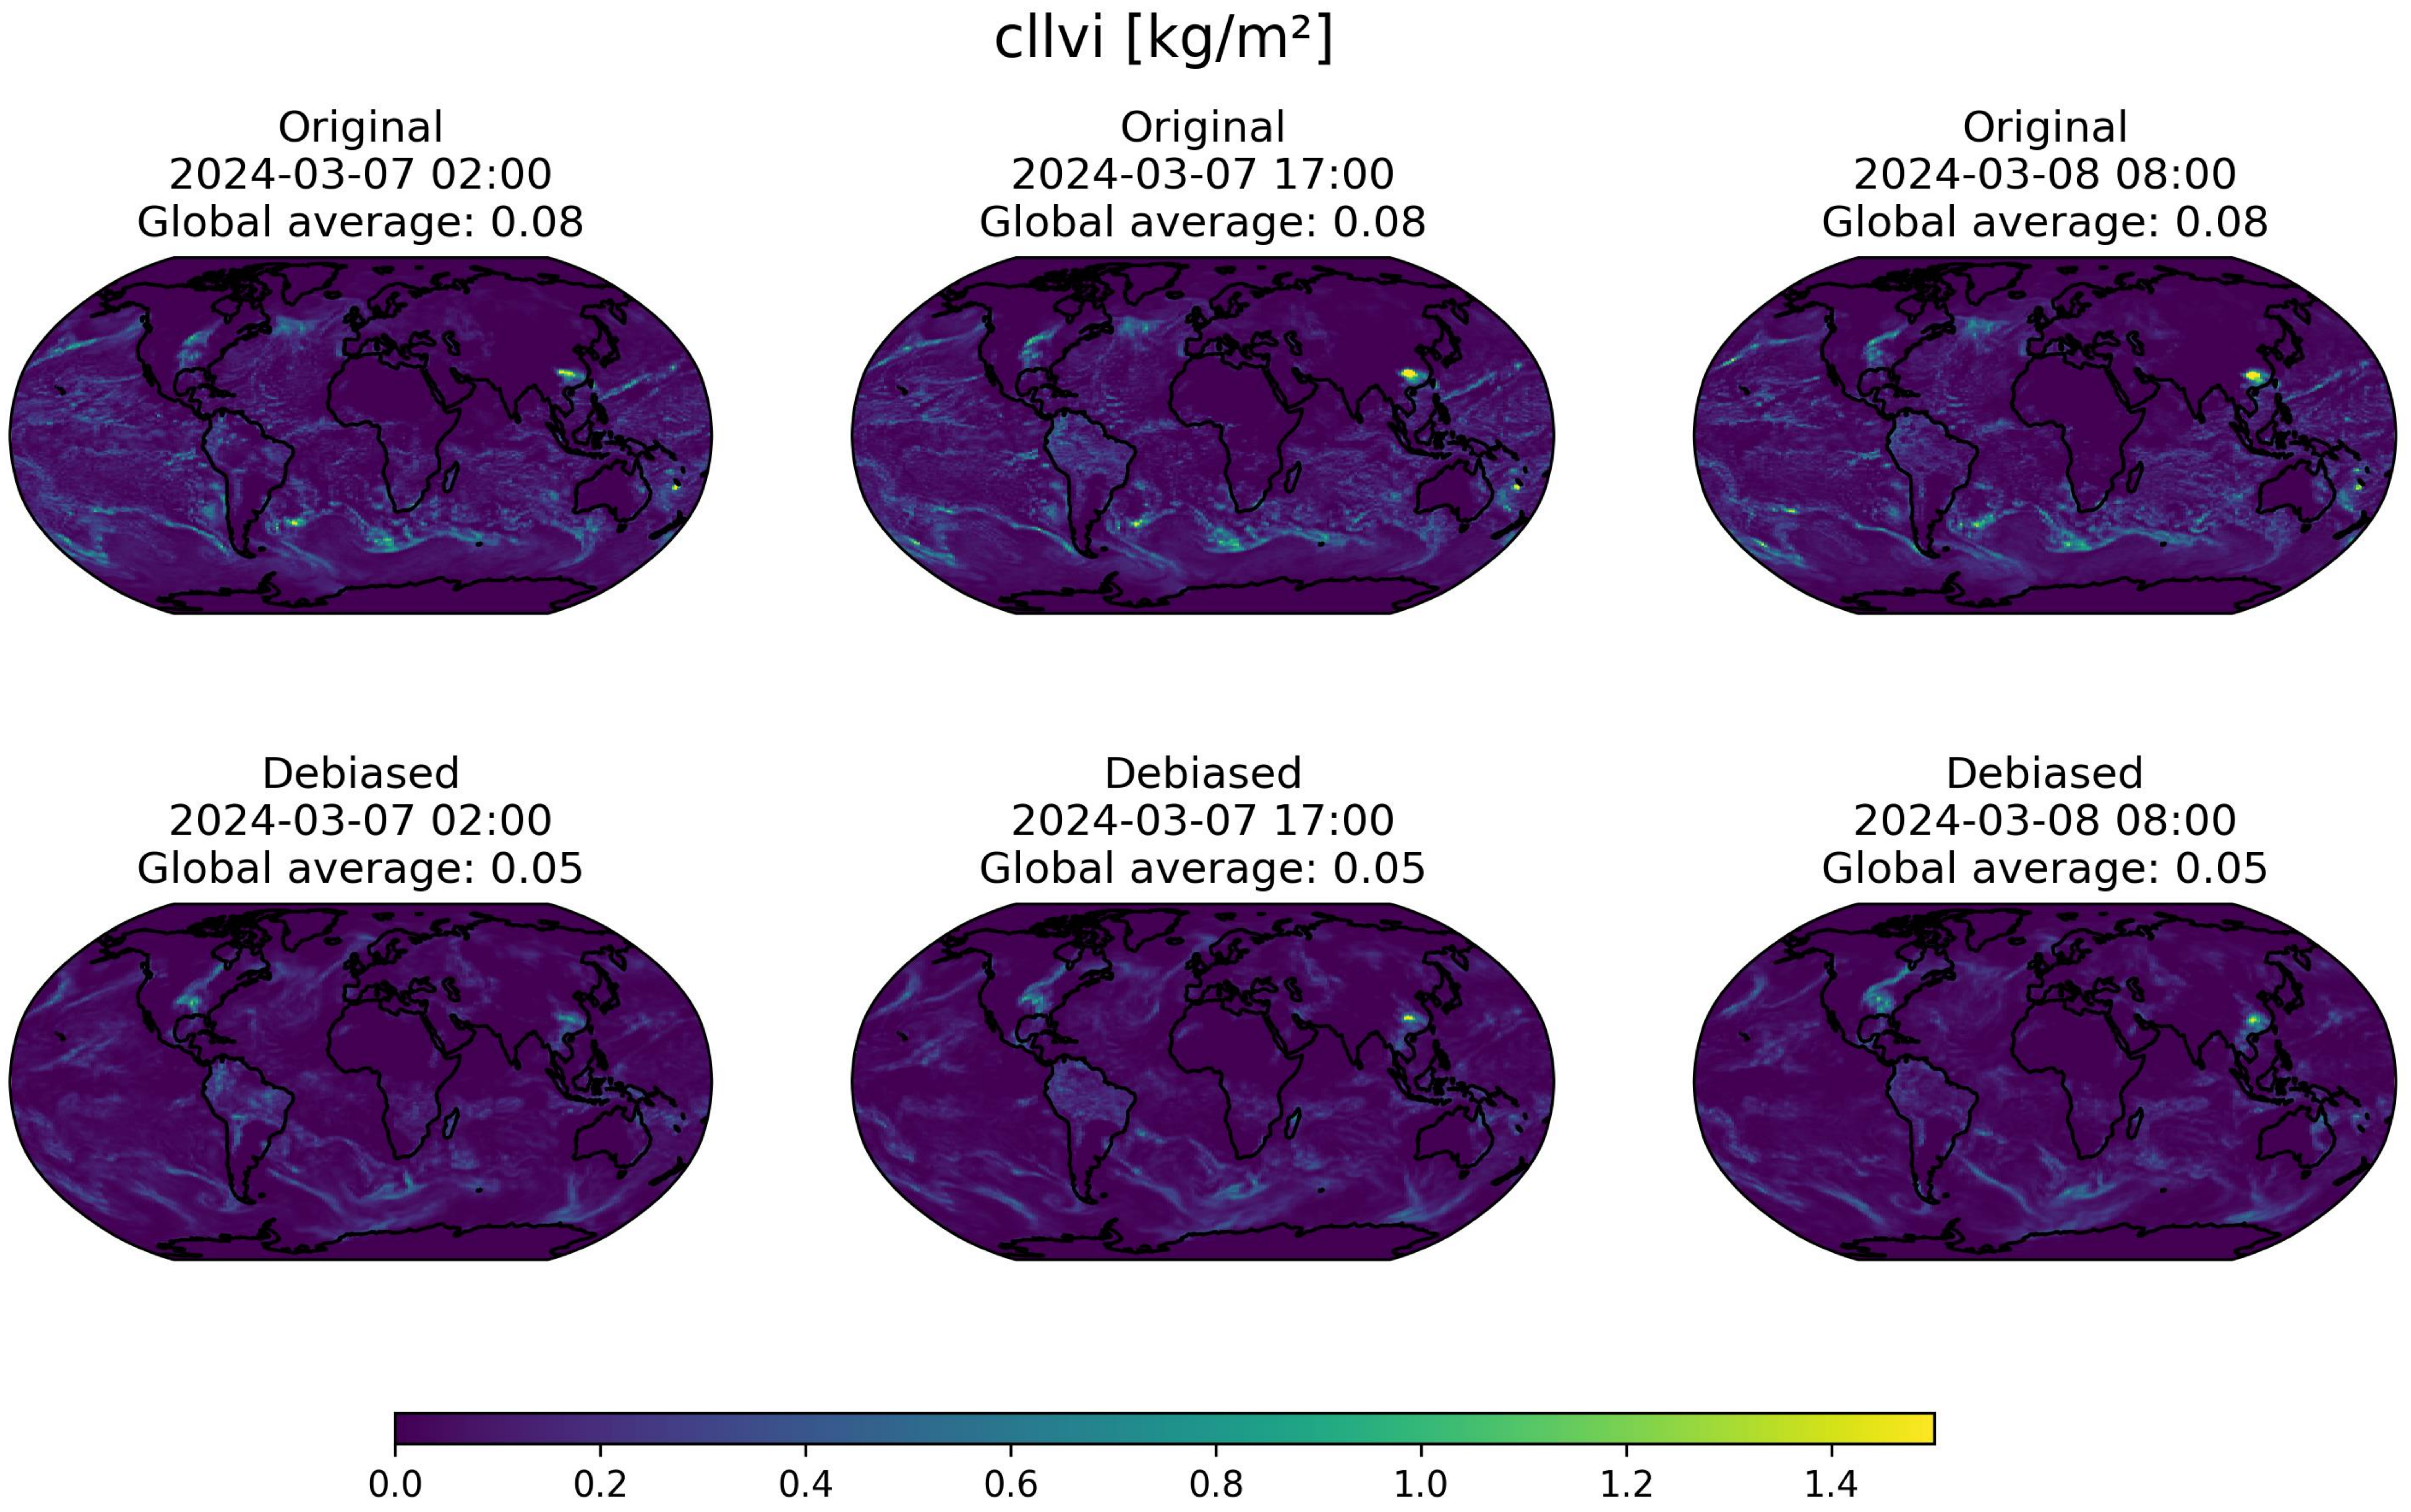

Climate in a Bottle

Impact:
Democratizes high-resolution climate-data access

- Enables interactive climate exploration
- Fixes corrupted/missing climate data
- Creates "what-if" scenarios instantly
- Transforms climate science from needing massive data archives to generating realistic climate data anywhere, anytime

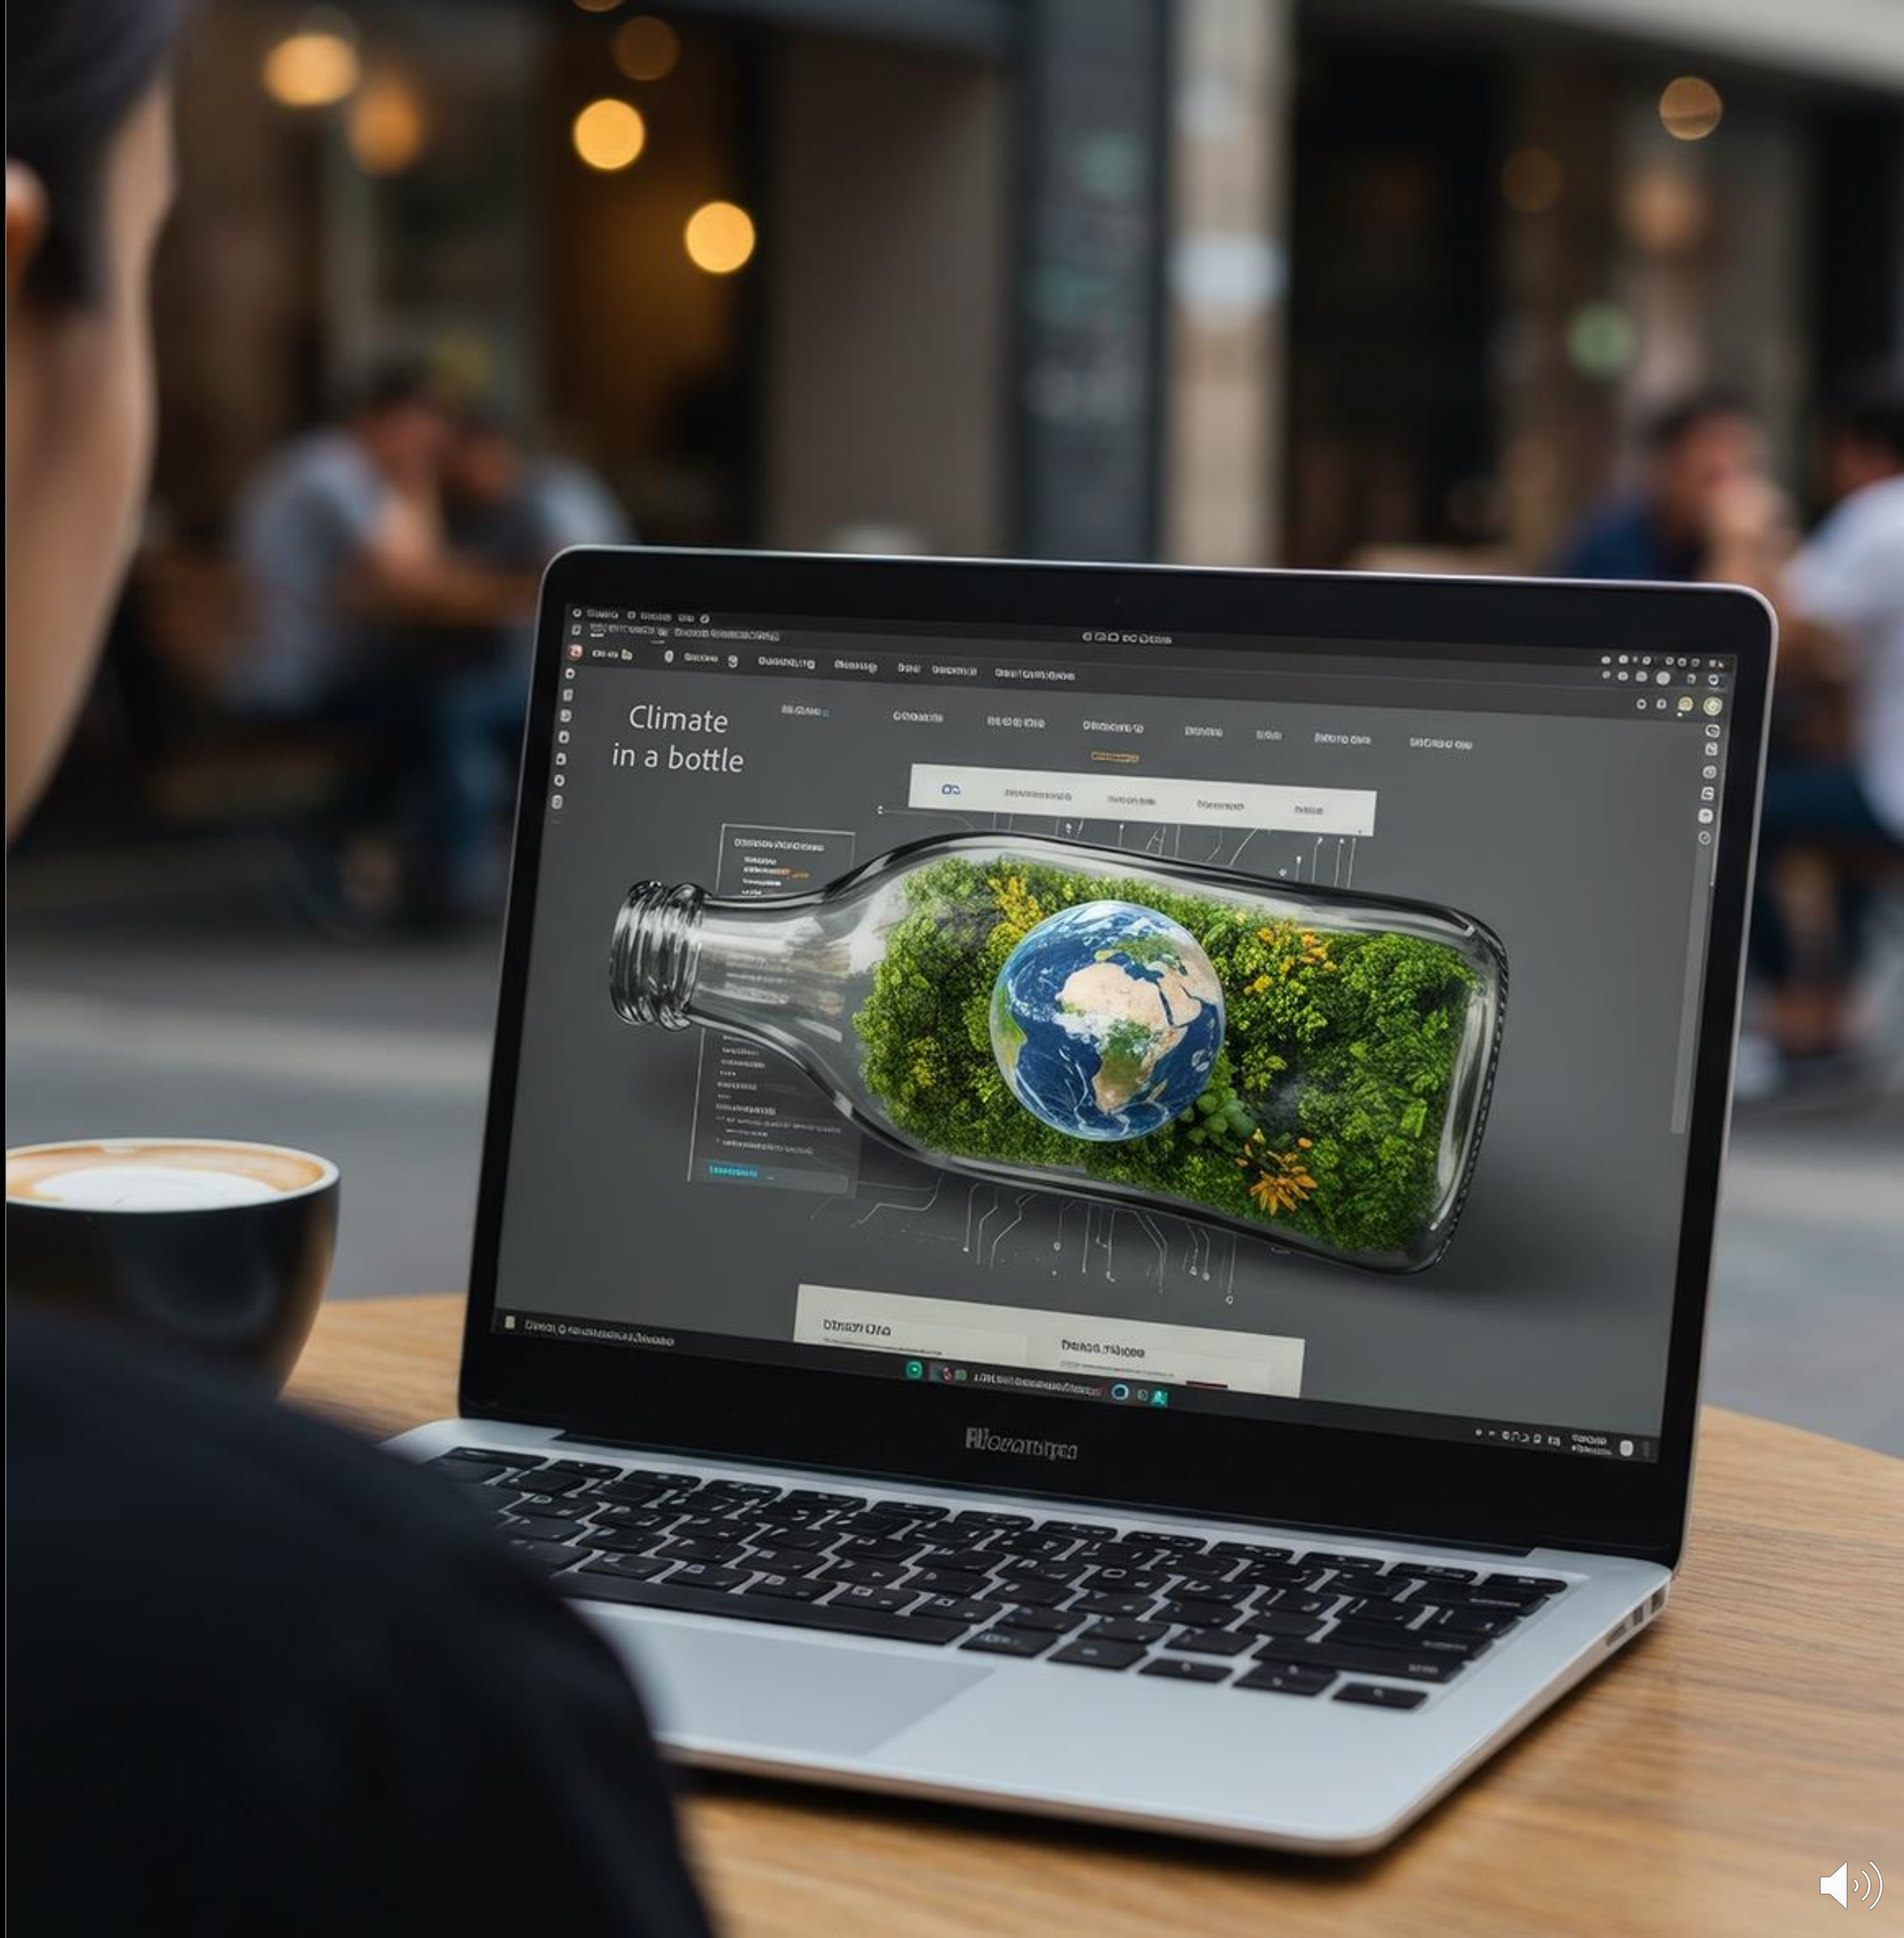

Create Window

Globe View

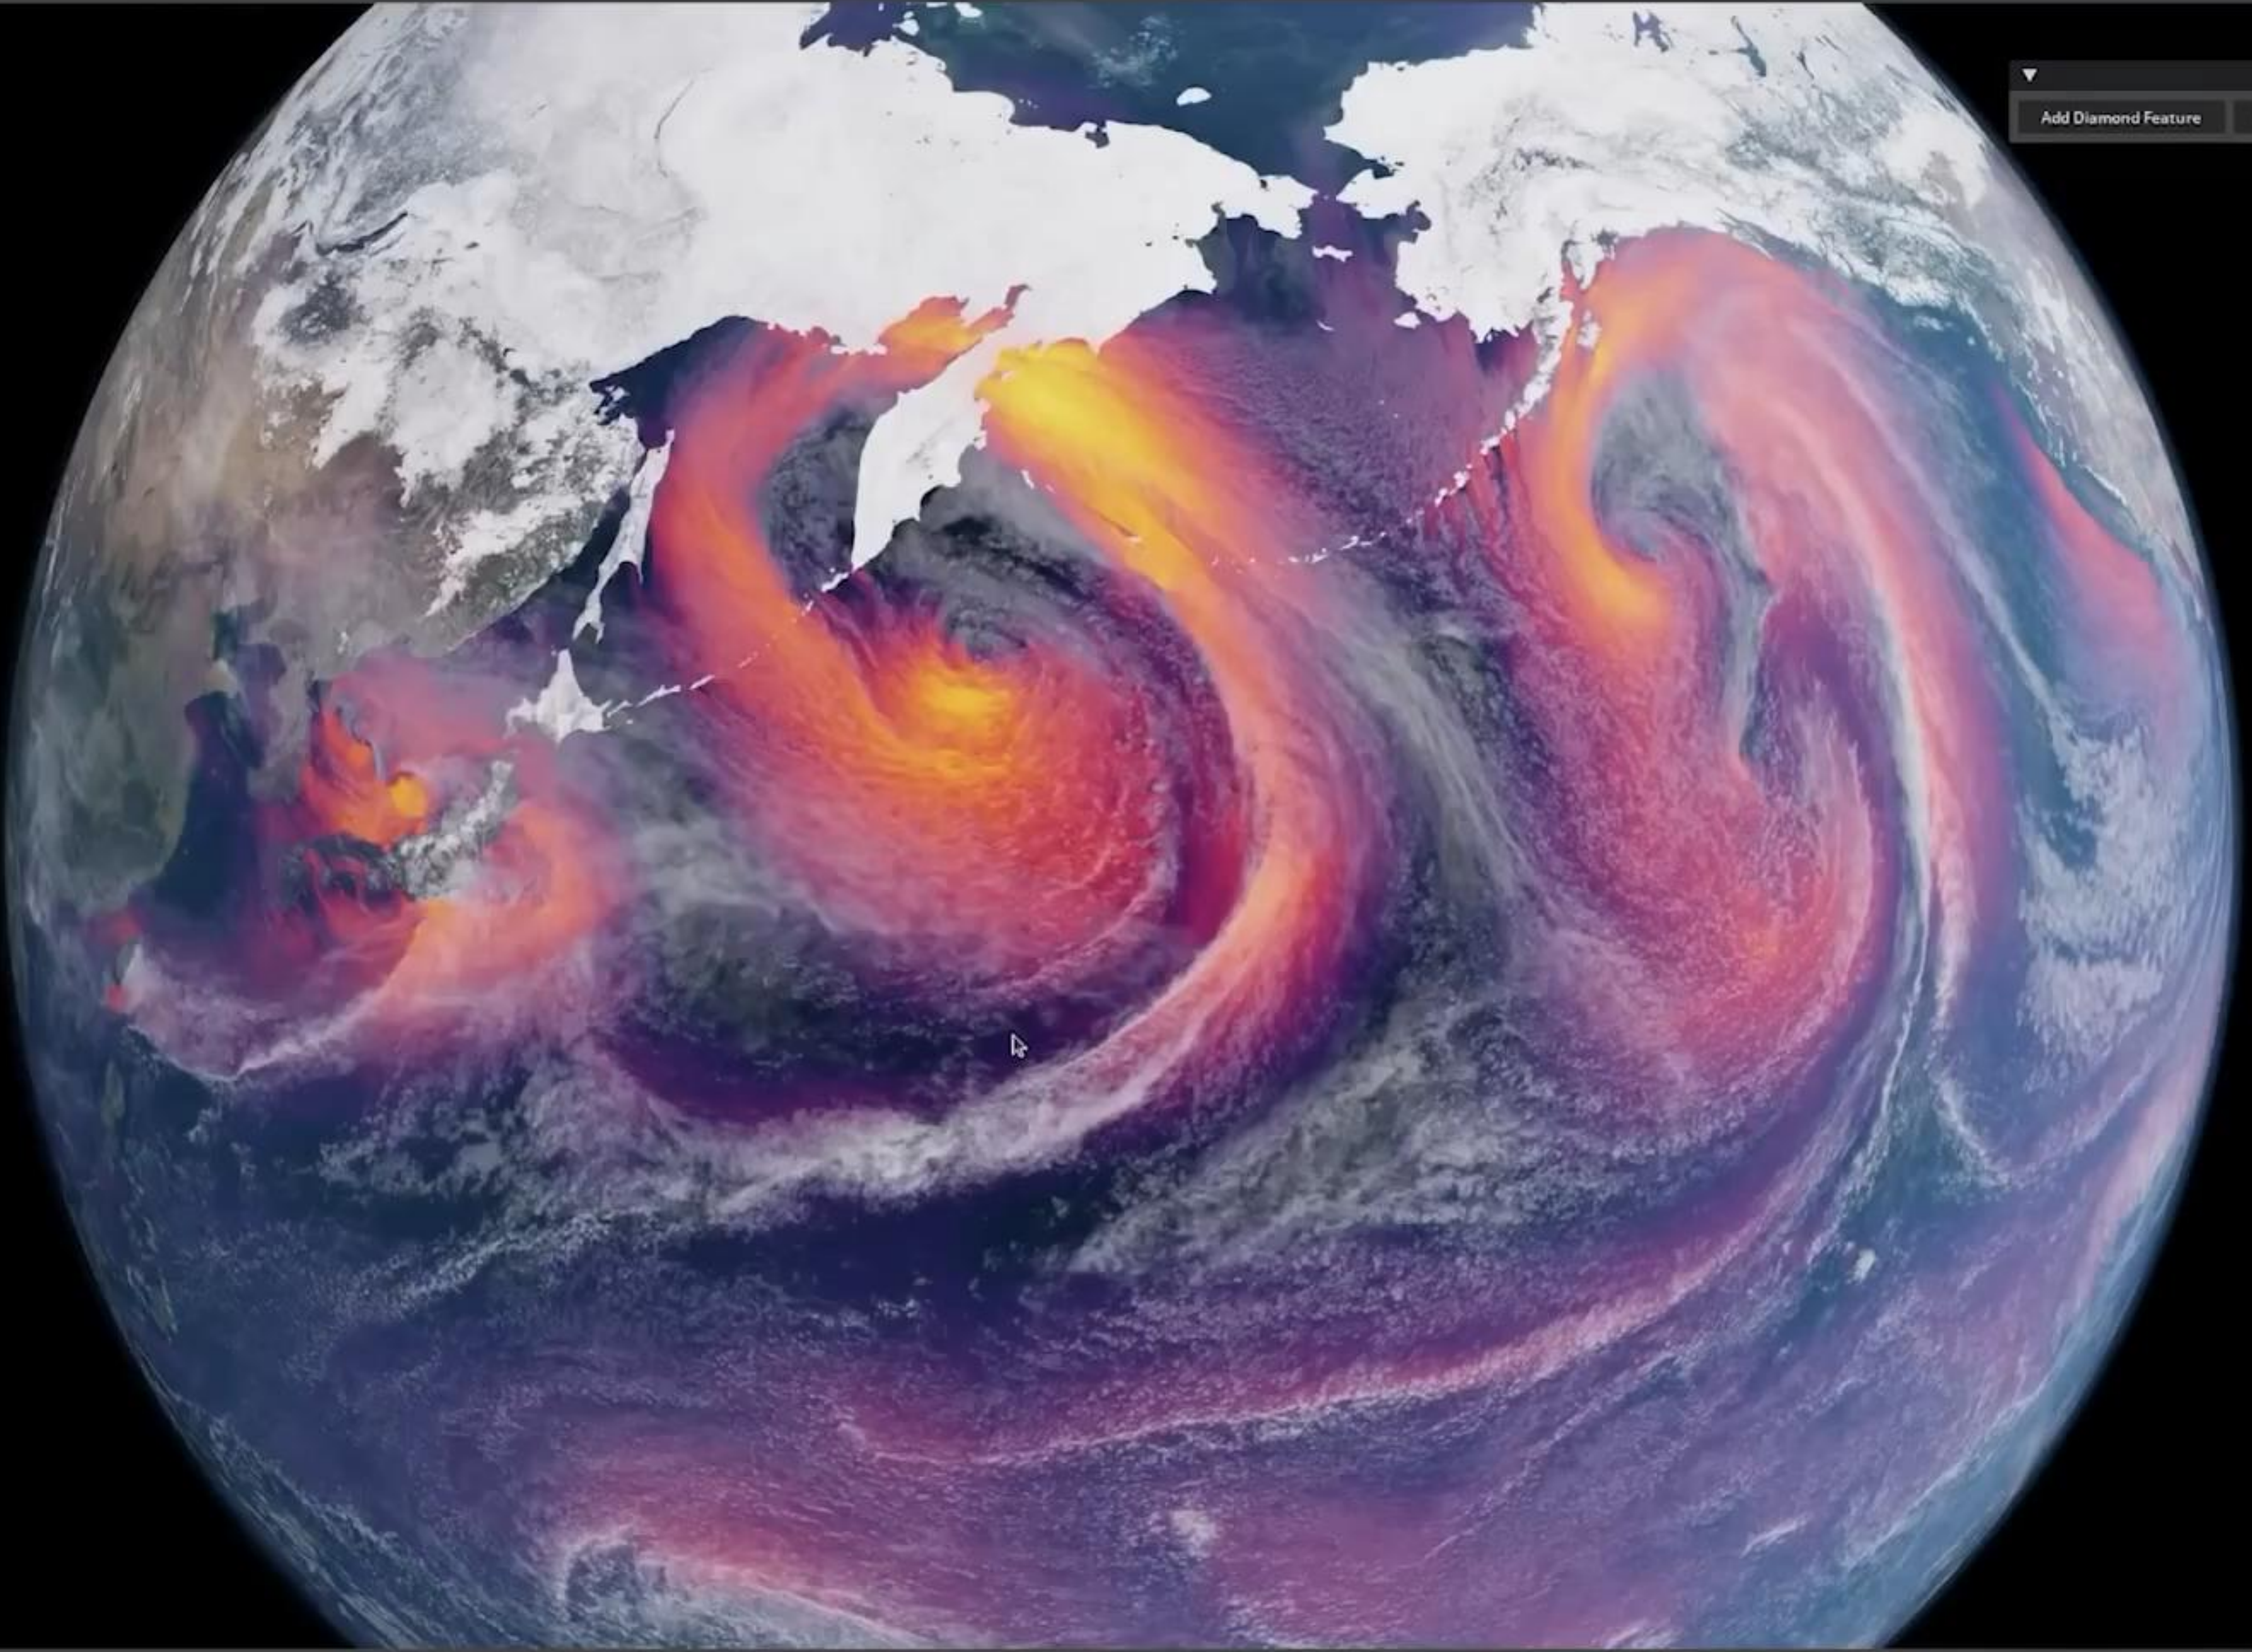

Image Feature Tests

Add Diamond Feature | Add LatLong Feature | Clear Features


File Edit Create Window Tools Profiler Layout Help
Layer
Viewport
RTX - Real-Time
Camera
Stage Lights
EBC Source Selection
EBC Data Source
R2B8
R2B9
R2B10
R2B11
Visualization
Emission Intensity:
3000
Toggle Wind
Toggle Clouds
Toggle Base
Caching
Precache
Stage
Render Settings
Search
Name (Old to New)
World (defaultPrim)
Looks
Plane
null_light
Camera
base
scripts
Property
Content
NVIDIA Assets
Asset Stores (beta)
Samples
Environments
Materials
Console
Import
Search
Videos
VirtualBox VMs
workspace
datasets
material
mdl-sdk
nvidia
OV
dynamic_texture_tes
test_001
Name
Date
Size
03/07/2023 10:16AM
12.46 KB
flat_earth_v001.usd
Date Modified03/07/2023 10:13AM
Created by
Modified by
File size
128.82 KB
Checkpoints
238
238
24
FPS
NVIDIA

AI-Accelerated Physical Modelling for Weather, Climate, and Engineering at NVIDIA

Ira Shokar, PhD | Applied Scientist | NVIDIA | ishokar@nvidia.com