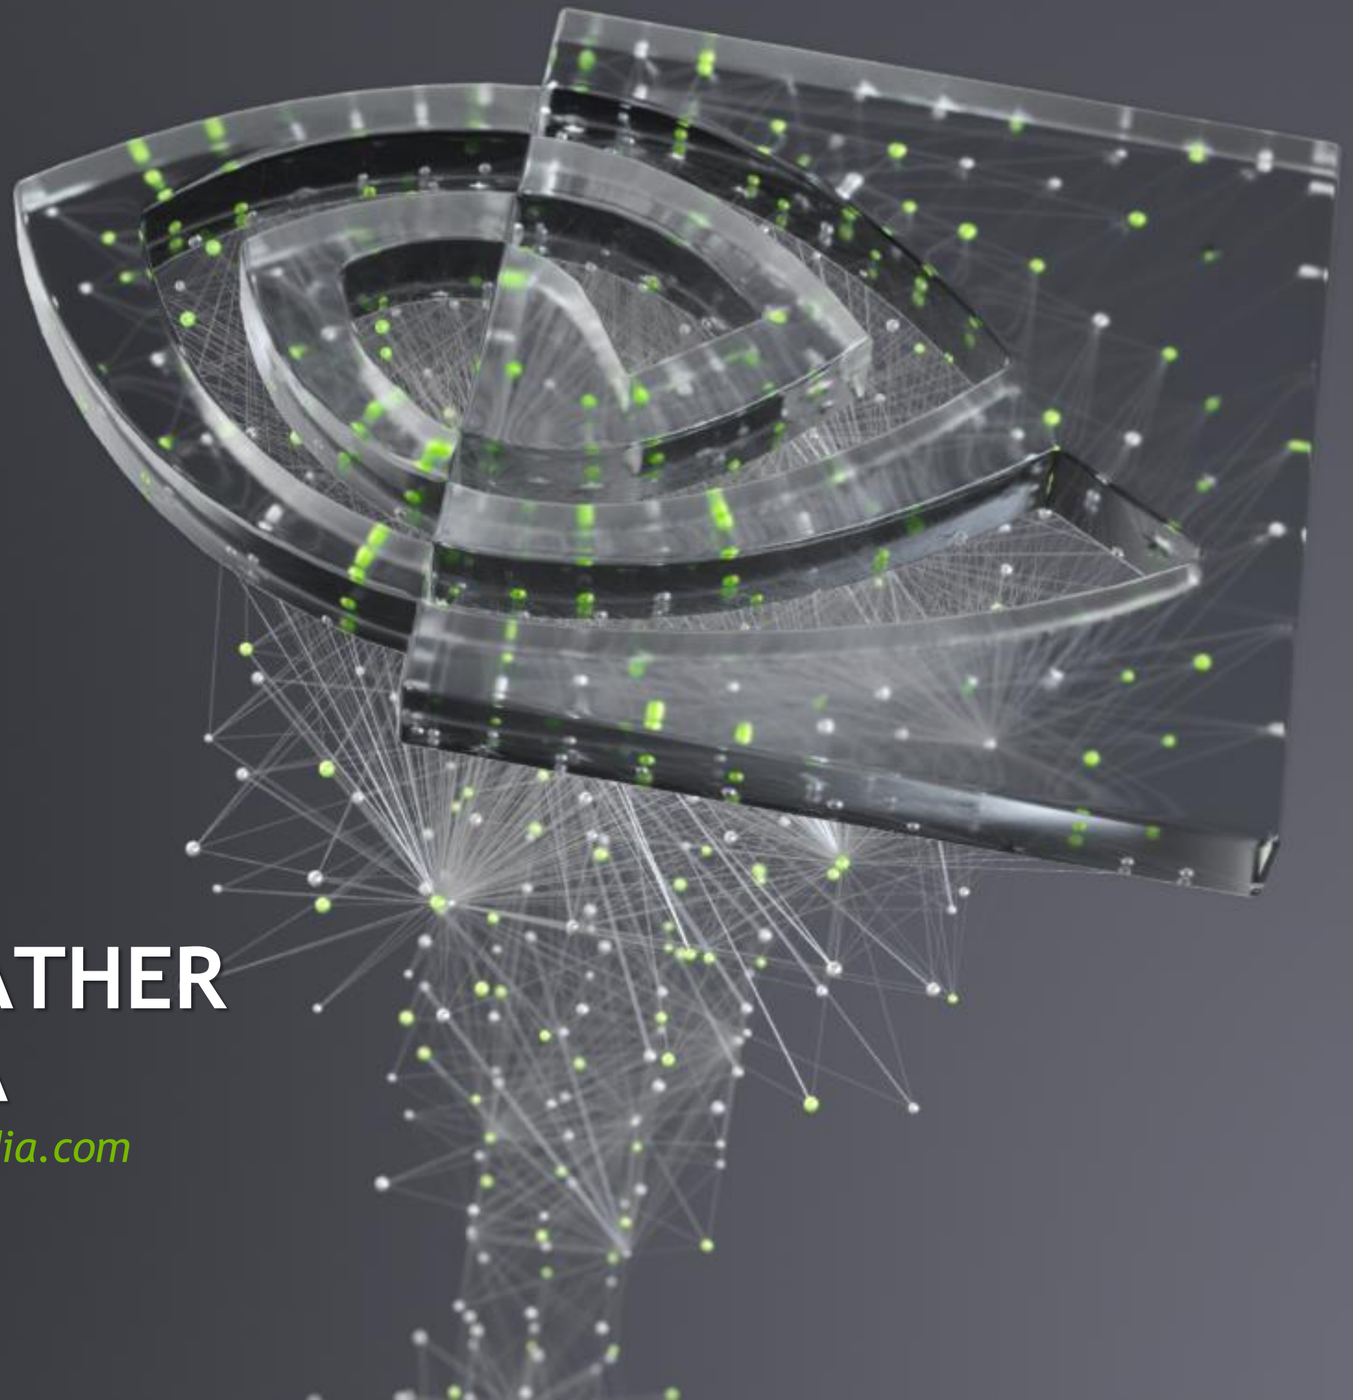

DRIVING NUMERICAL WEATHER PREDICTION WITH NVIDIA

Jeff Adie (on behalf of Peter Messmer) – *jadie@nvidia.com*

AGENDA

Modelling and Simulation
AI
Visualisation
A Digital Twin Earth

IFS, 1.4km global run on Summit
Cloud and streamlines generated in ParaView
Data courtesy of Nils Wedi, ECMWF

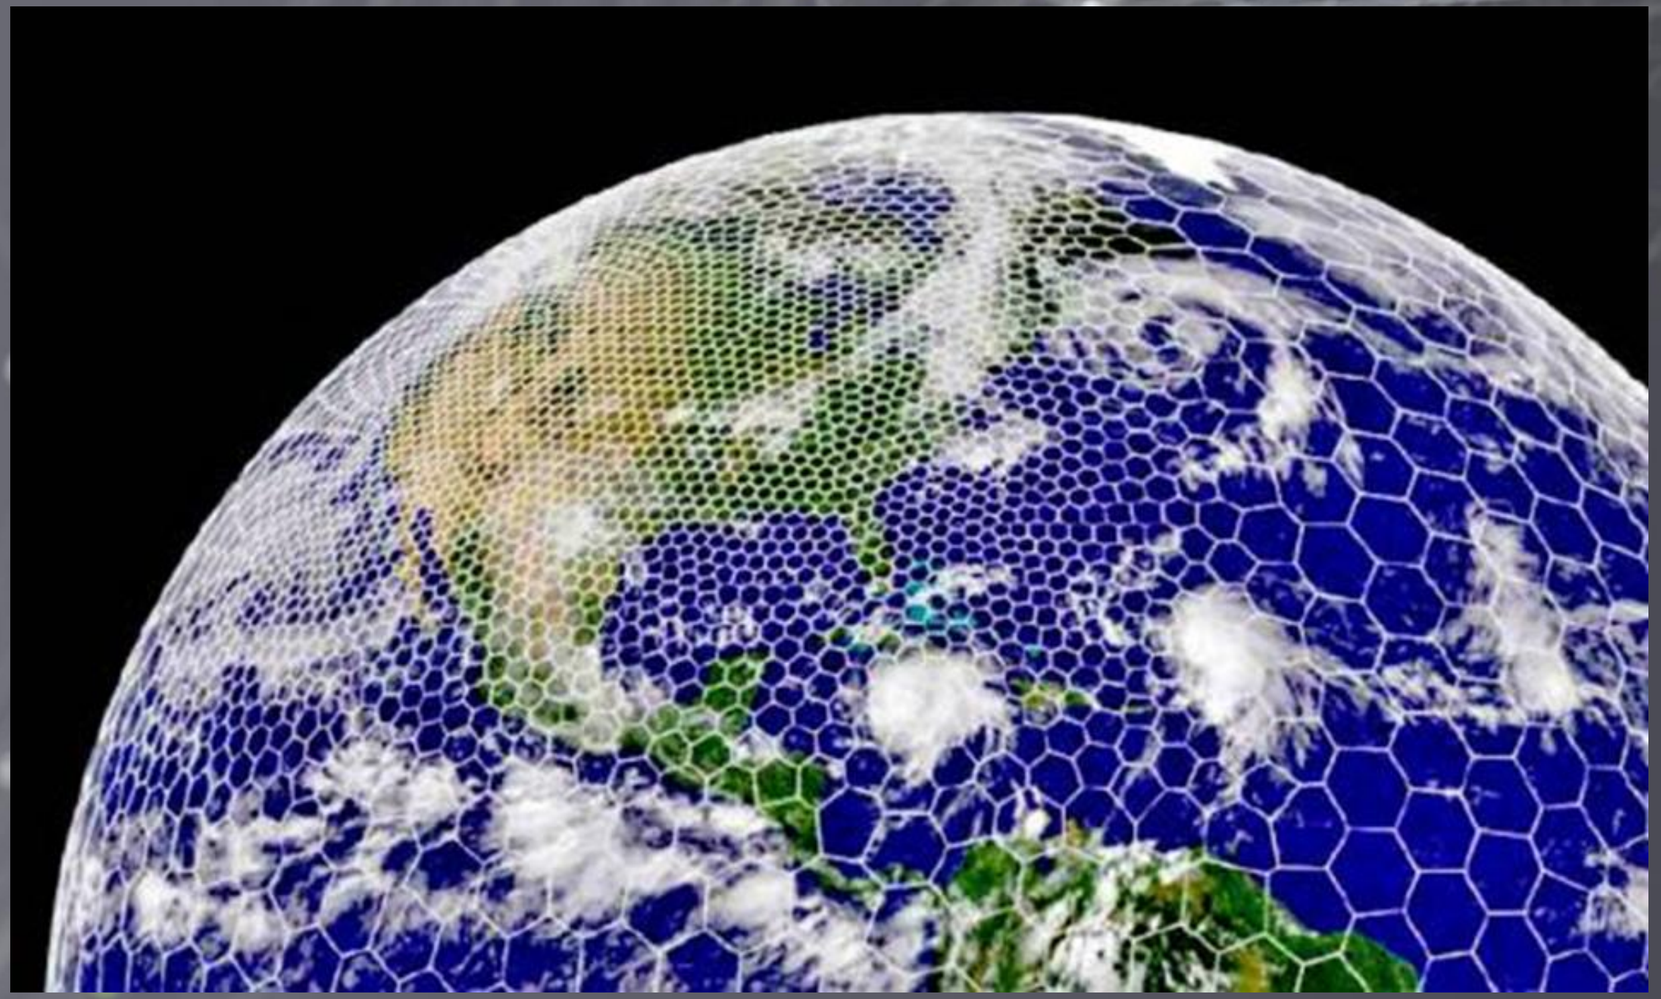

MODELLING AND SIMULATION

EARTH SYSTEM MODELS PROVIDE THE DATA

Physical models are as important as ever, and more accurate simulations are needed.

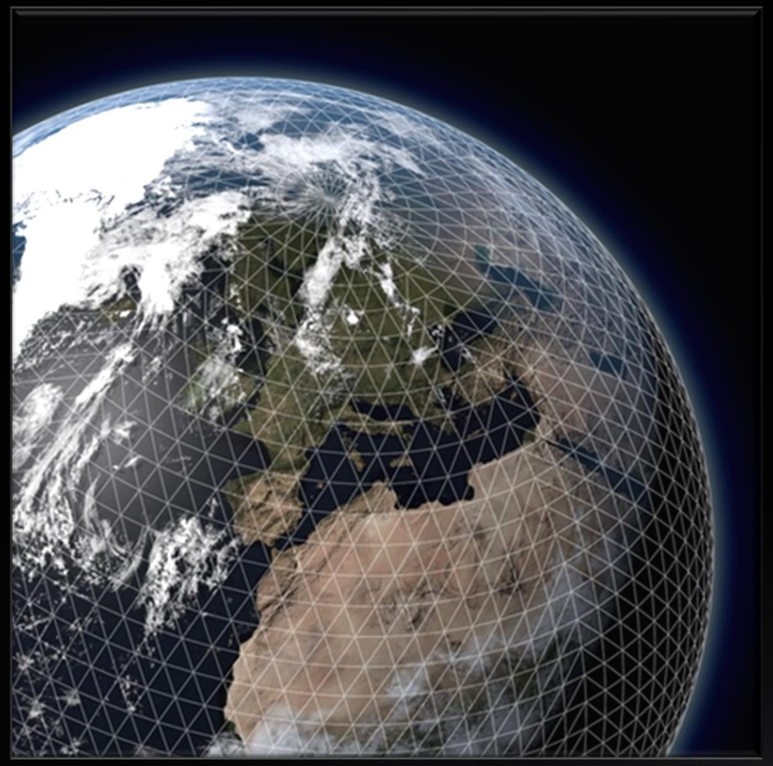

Cloud resolving weather and climate models

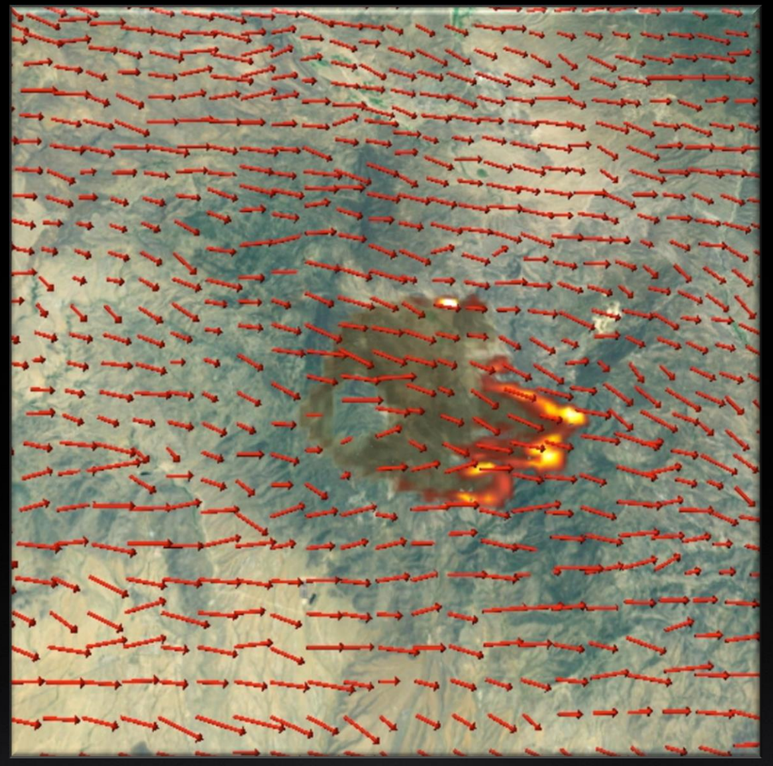

Physical models: floods, fires, glaciers, tsunamis, storms

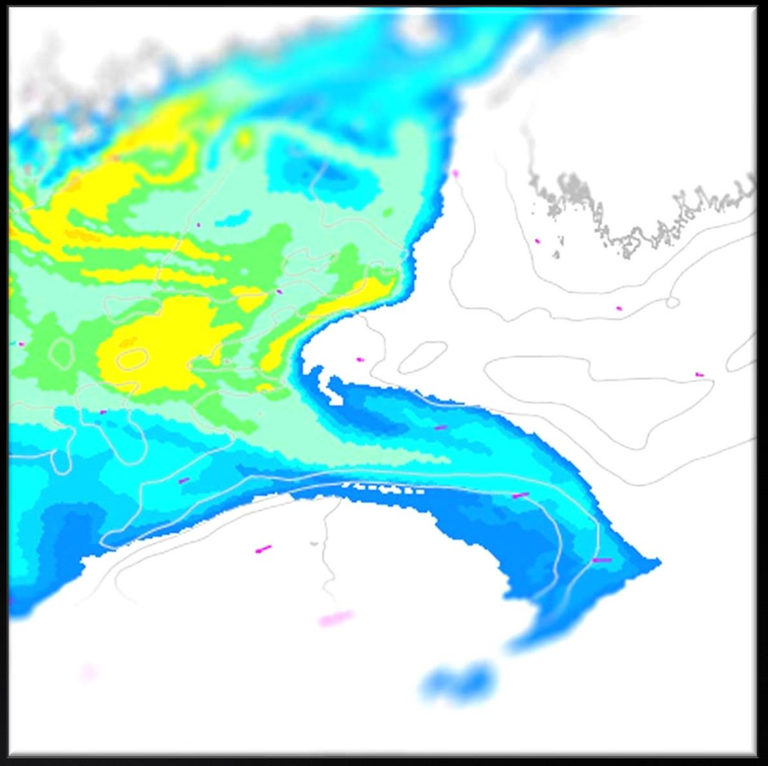

Dynamical models of power grids, transportation networks, biology

http://phys.org/news/2015-12-ncar-wildland-colorado.html
https://oceanservice.noaa.gov/hazards/hab/east-coast.html

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

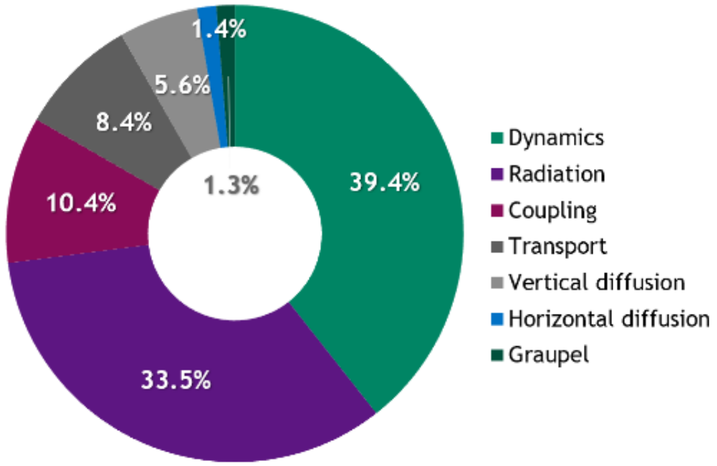

Time distribution per step on 2xEPYC 7742

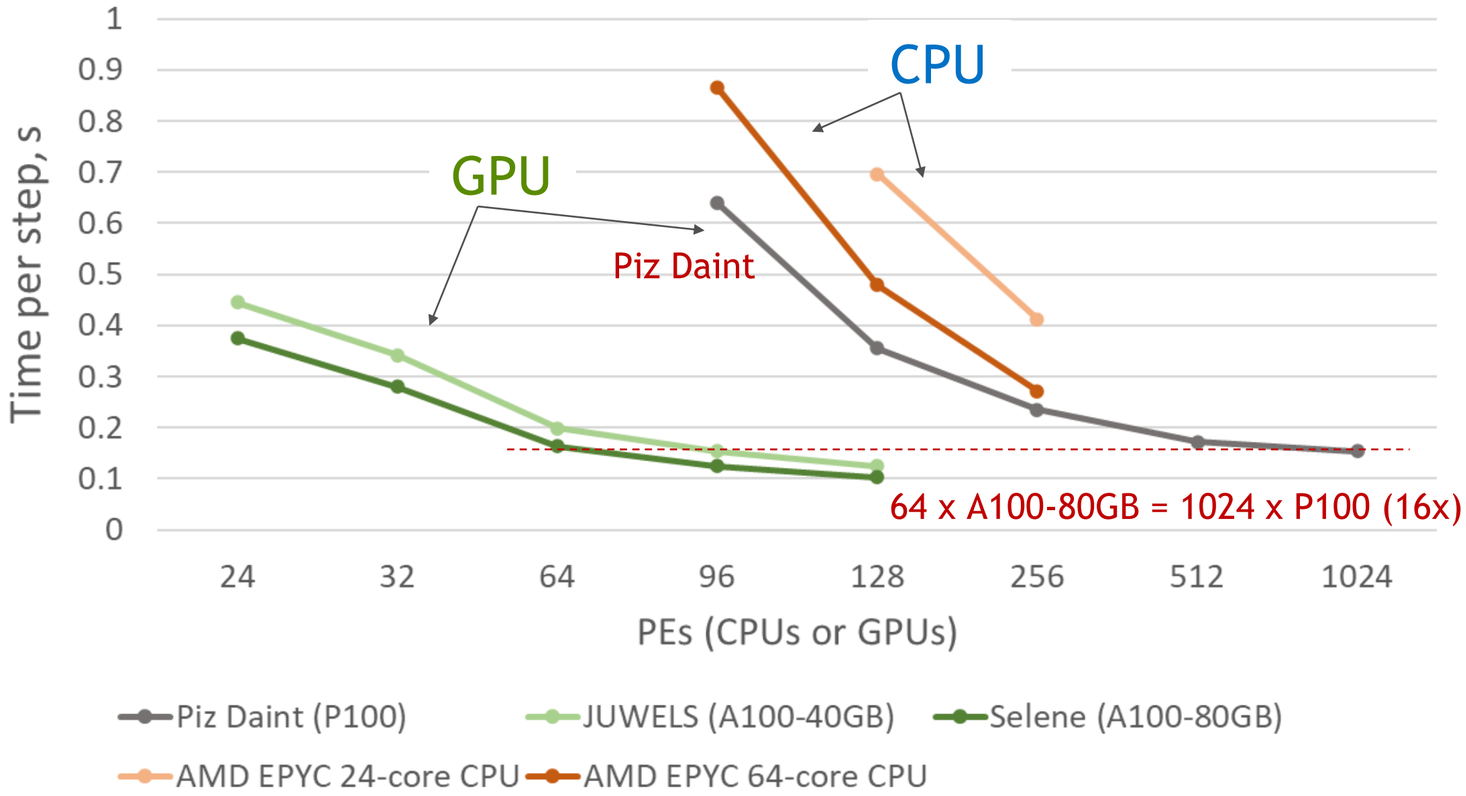

Source: NVIDIA, March 2021 Dr. Dmitry Alexeev

MPAS GPU Scalability on ORNL Summit

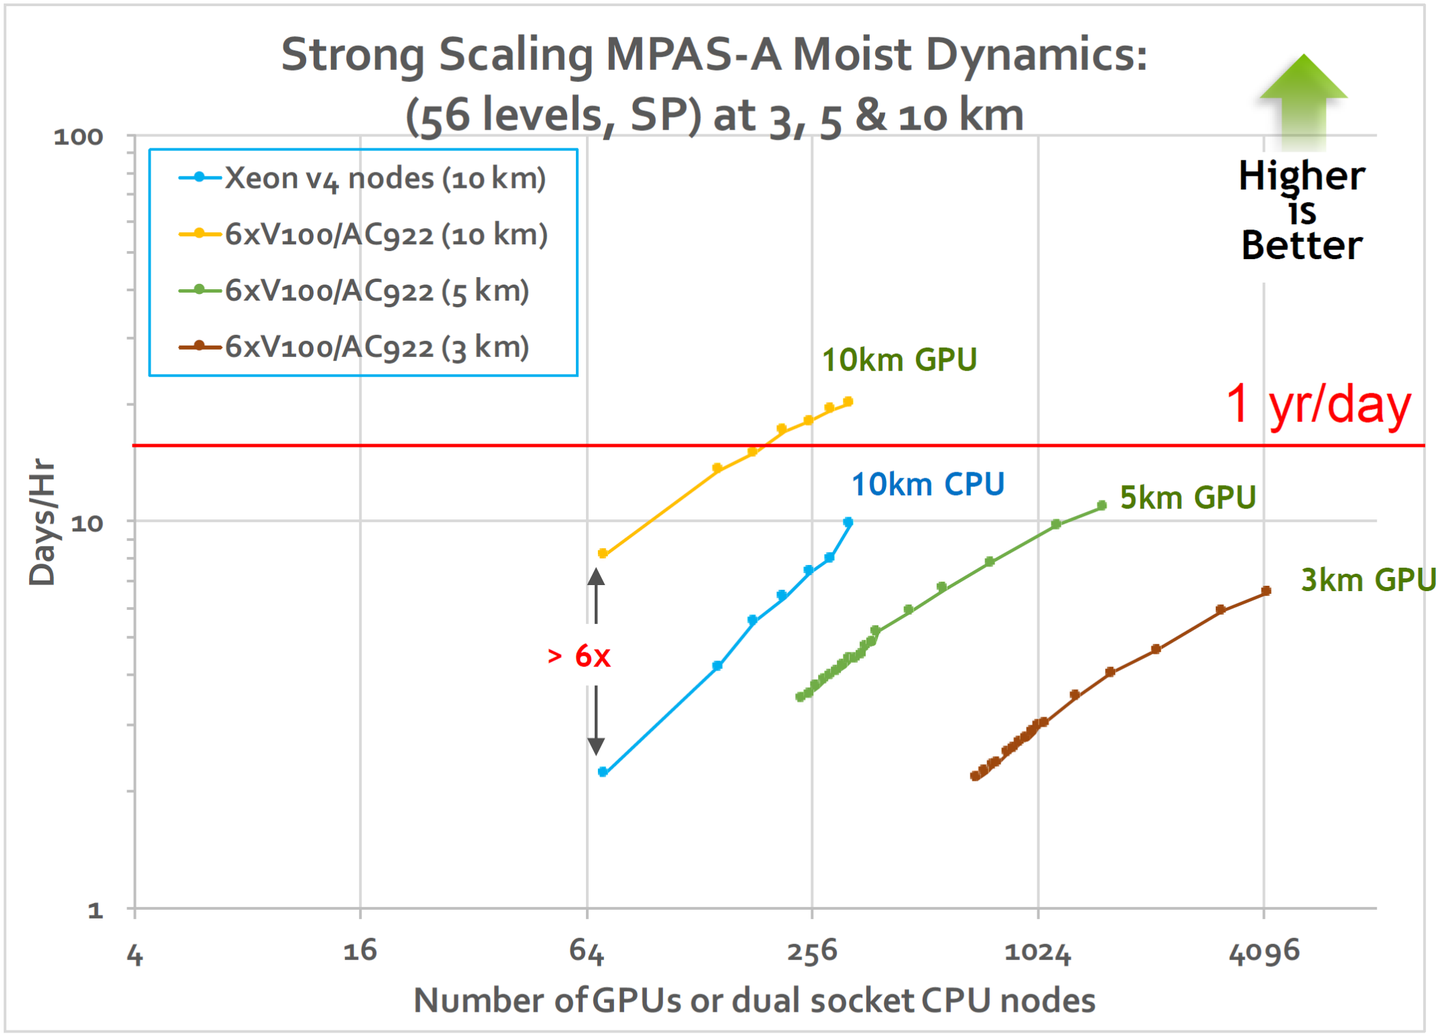

AMS 2020

12 – 16 Jan 2020, Boston, USA

An Implementation of MPAS-Atmosphere Running on GPUs

Dr. Raghu Kumar, et al.

- ORNL Summit GPU system (V100)
- NCAR Cheyenne CPU system (BDW)

An Implementation of MPAS-Atmosphere Running on GPUs [PDF (1MB)]

ECMWF Scaling on ORNL Summit ~2000 Nodes

ECMWF

Batched matrix multiplication: speedup in ESCAPE1 dwarf

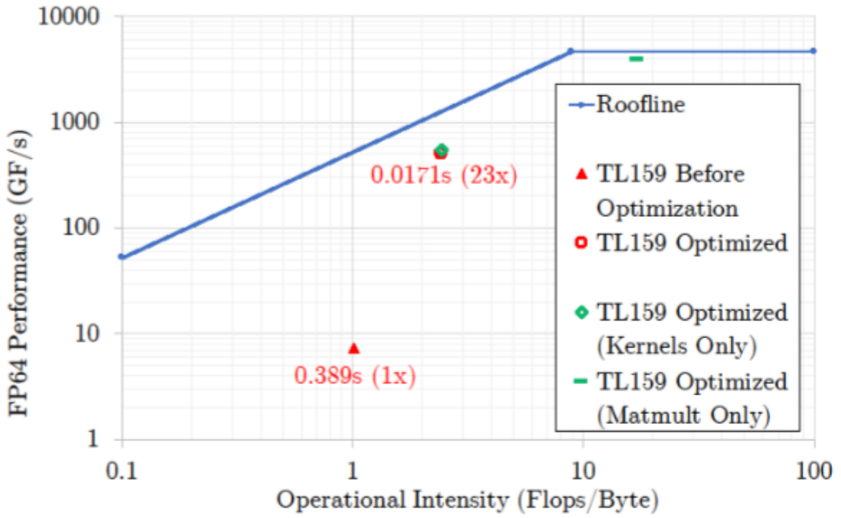

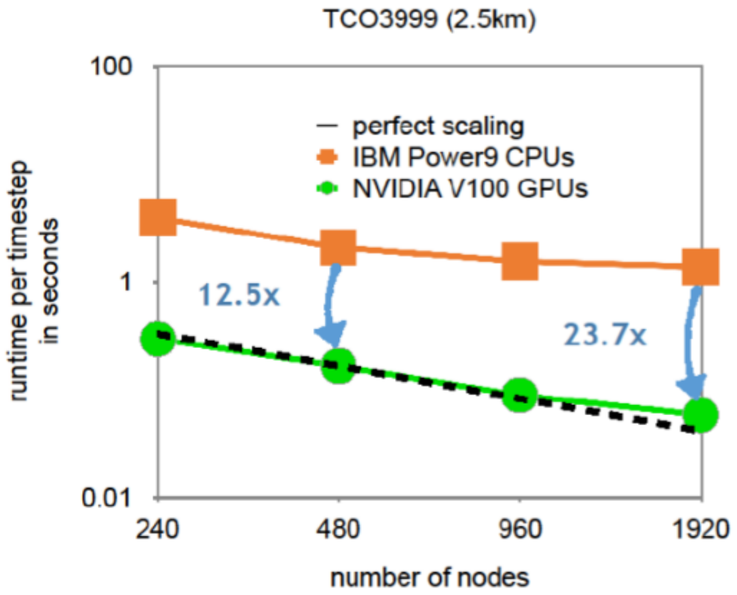

- added zero operations increase operational intensity
- overall huge speedup
- should also have strong positive effect on strong scaling

ECMWF EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS

20

Spectral Harmonics Dwarf = 23.7x

Batched Legendre Transform (GEMM)

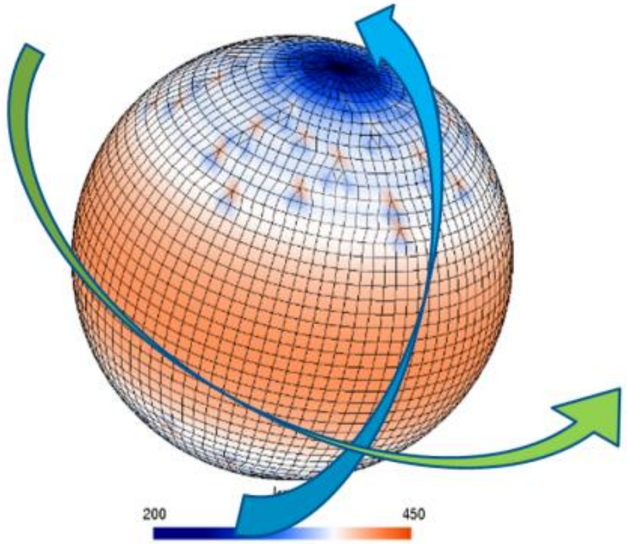

https://www.ecmwf.int/en/about/media-centre/science-blog/2019/oak-ridge-leadership-computing-facility-user-workshop-2019

NVIDIA

WRF Scaling on ORNL Summit 4096 Nodes

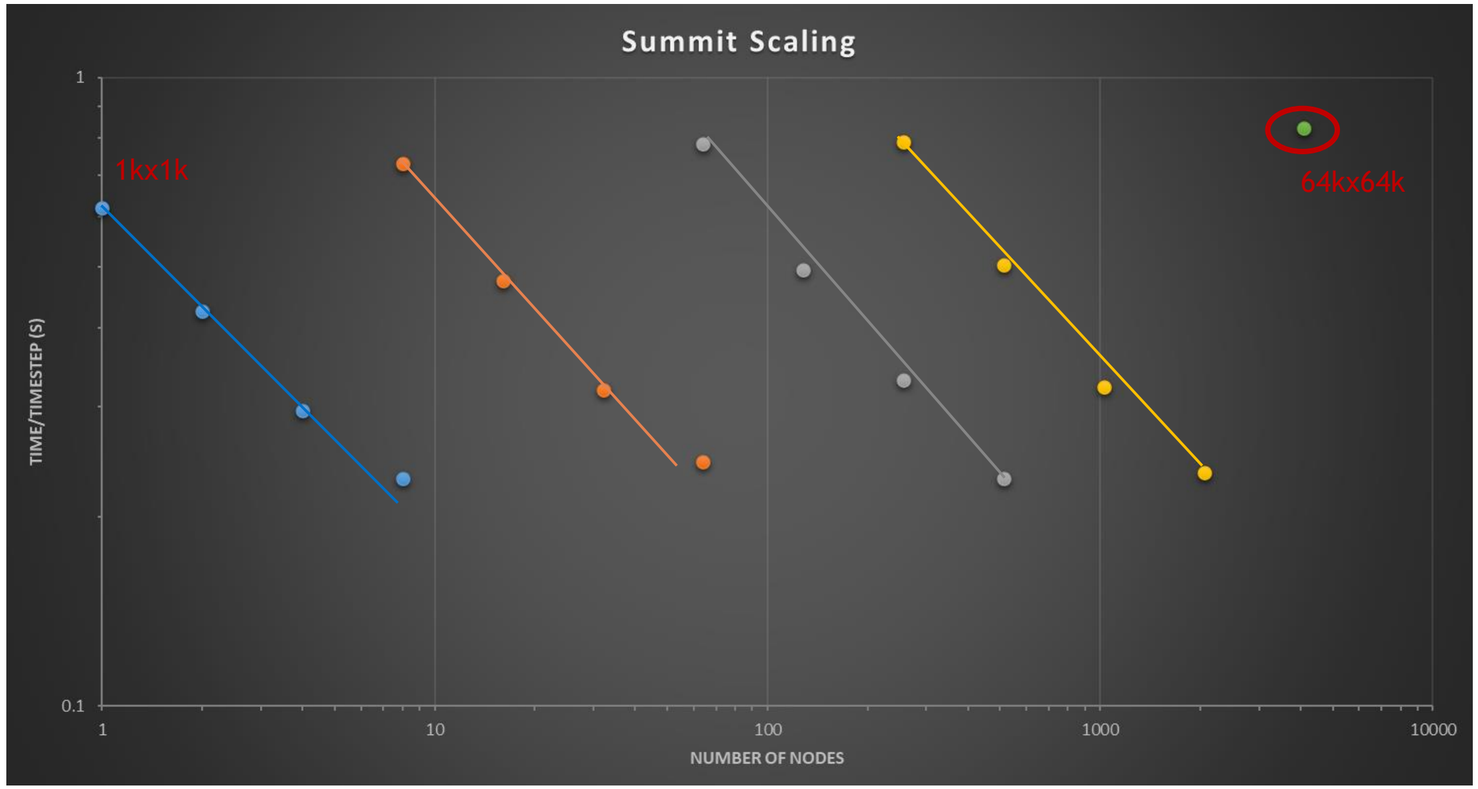

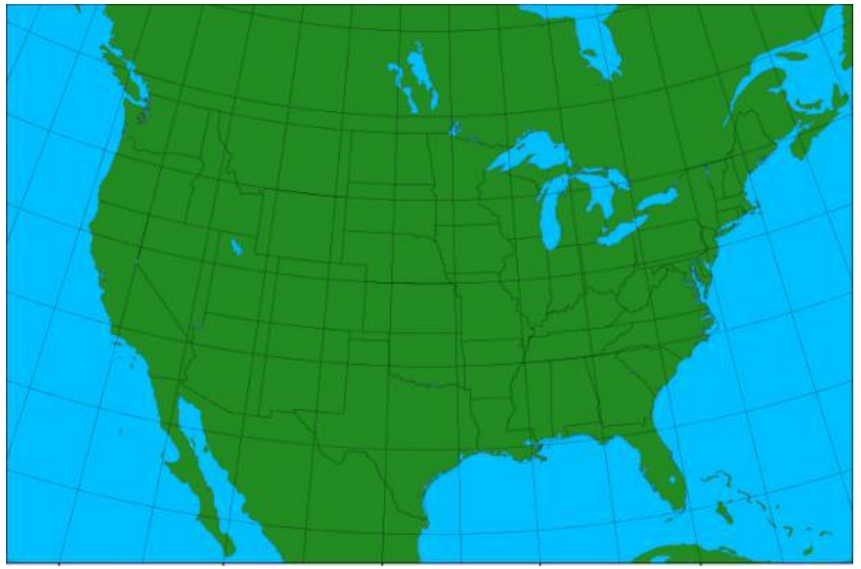

Data limited!

Mesh size (128 levels)	Gridpoints (B)	Input (GB)	24hr, 1hr history (GB)
1024x1024	0.13	9.4	242
2048x2048	0.54	37.6	966
4096x4096	2.15	150.3	3,865
8192x8192	8.59	601.3	15,462
16384x16384	34.36	2,405	61,848
32678x32768	137.4	9,621	247,390
65536x65536	549.8	38,483	989,560

100m LES CONUS, 128 levels, 550 Billion Cells

WRF THE WEATHER RESEARCH & FORECASTING MODEL

NVIDIA/TQI 2021: S. Posey, G. Sever, J. Adie

NVIDIA.

NVIDIA'S GRACE CPU

An Arm GPU for Giant-Scale AI & HPC

NVIDIA GRACE CPU

Purpose-Built to Train the World's Largest Models

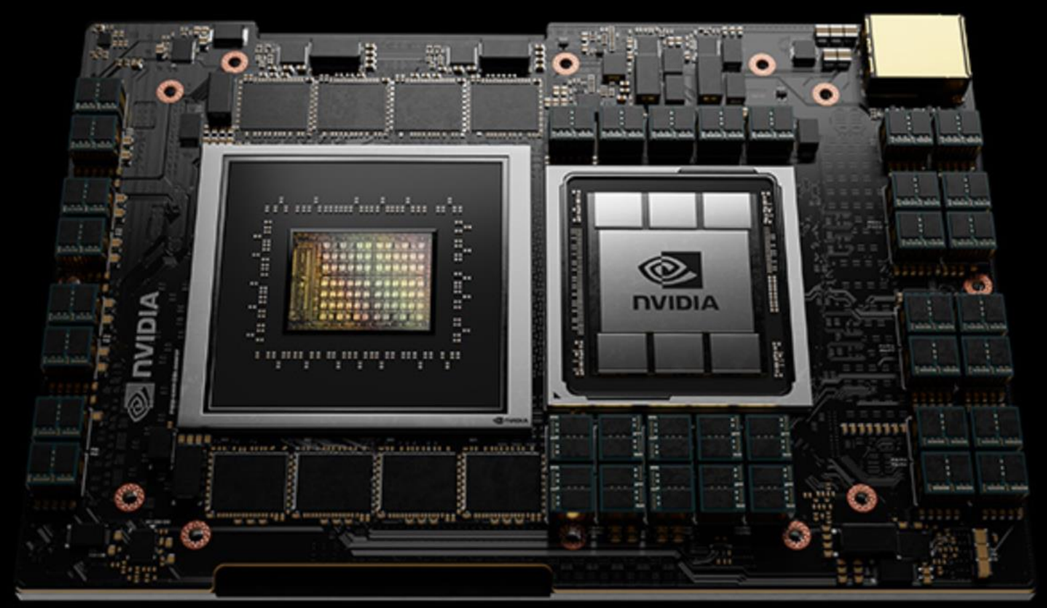

CSCS Next-Gen HPC System Based on Arm + GPU

CSCS MIGRATION FROM PIZ DAINT TO ALPS

20 Exaflops of AI

Accelerated w/ NVIDIA Grace CPU and NVIDIA A-NEXT GPU

HPC and AI For Scientific and Commercial Apps

Advance Weather, Climate, and Material Science

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Hewlett Packard Enterprise

NVIDIA

HYPERSCALE CLIMATE MODELS ?

A cloud-resolving climate model will require 10 million times the performance

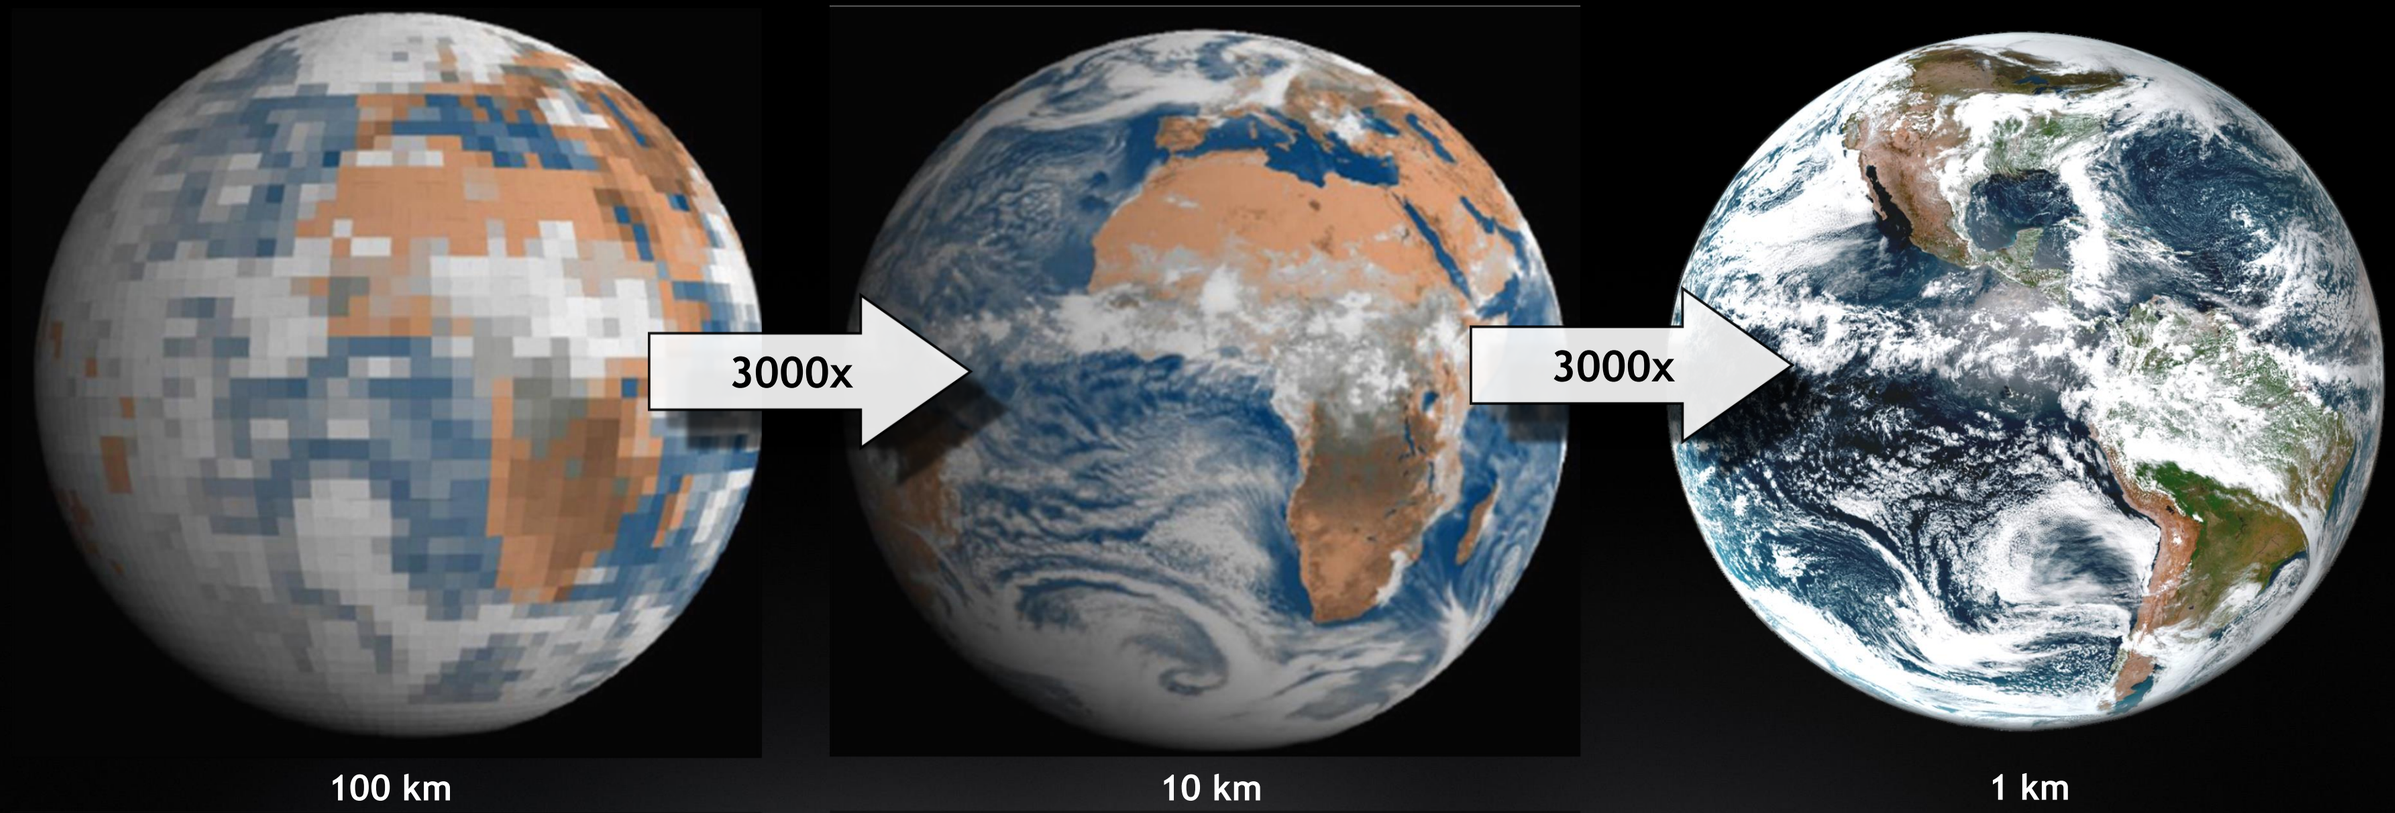

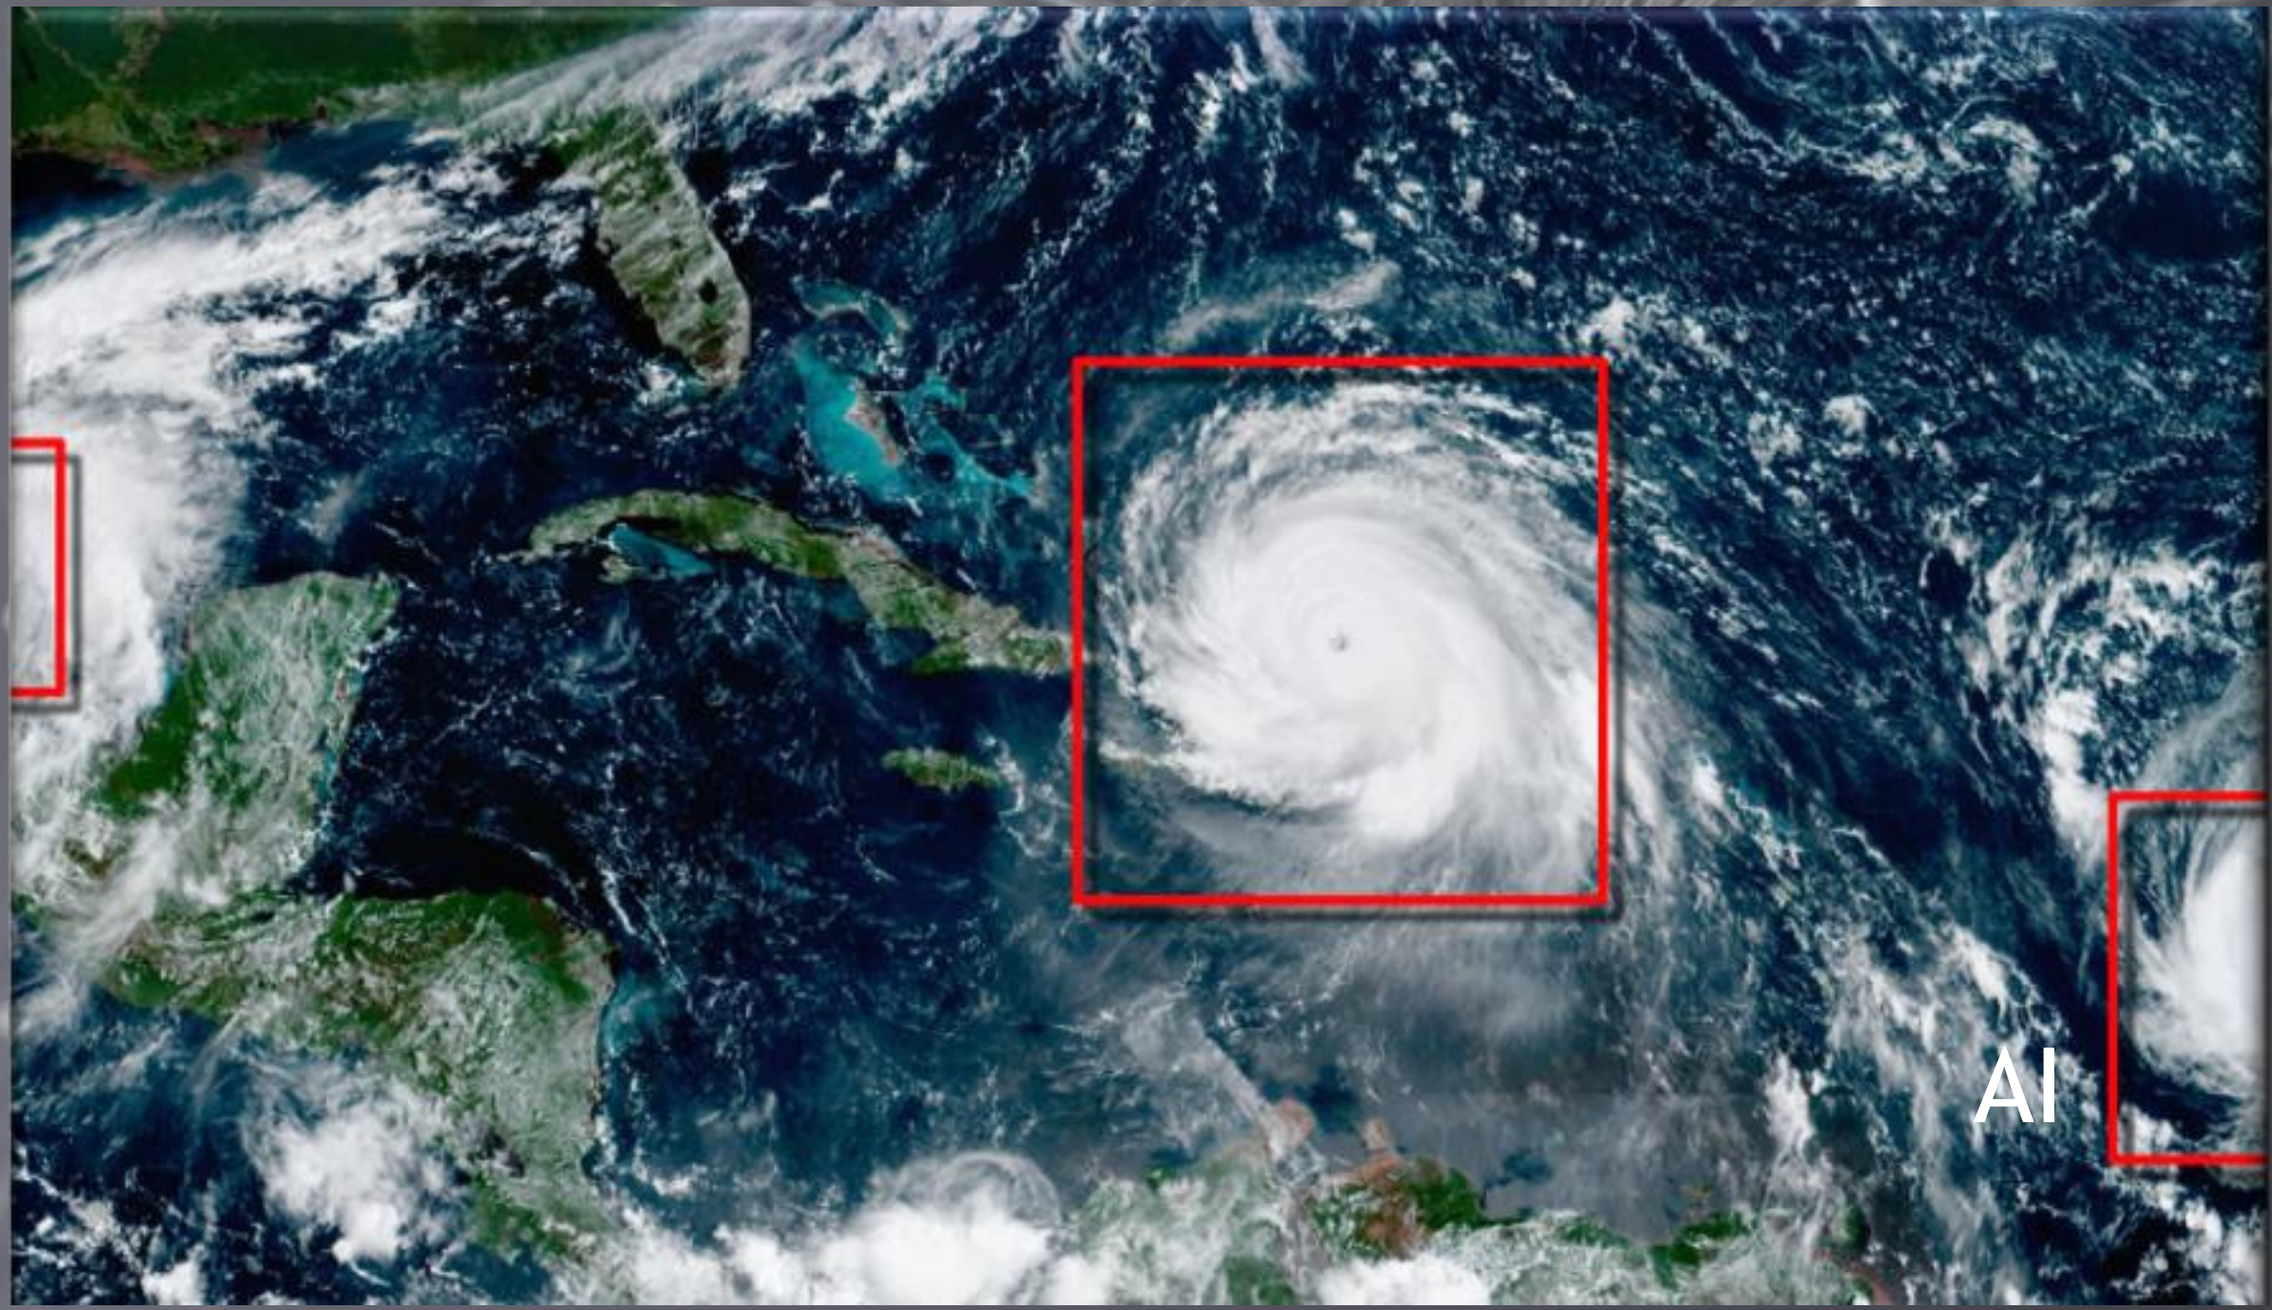
AI

DATA-DRIVEN MODELS FOR INTERACTIVITY

Data driven models provide the speed and interactivity needed to explore what-if scenarios

Omniverse: Interactive Experience

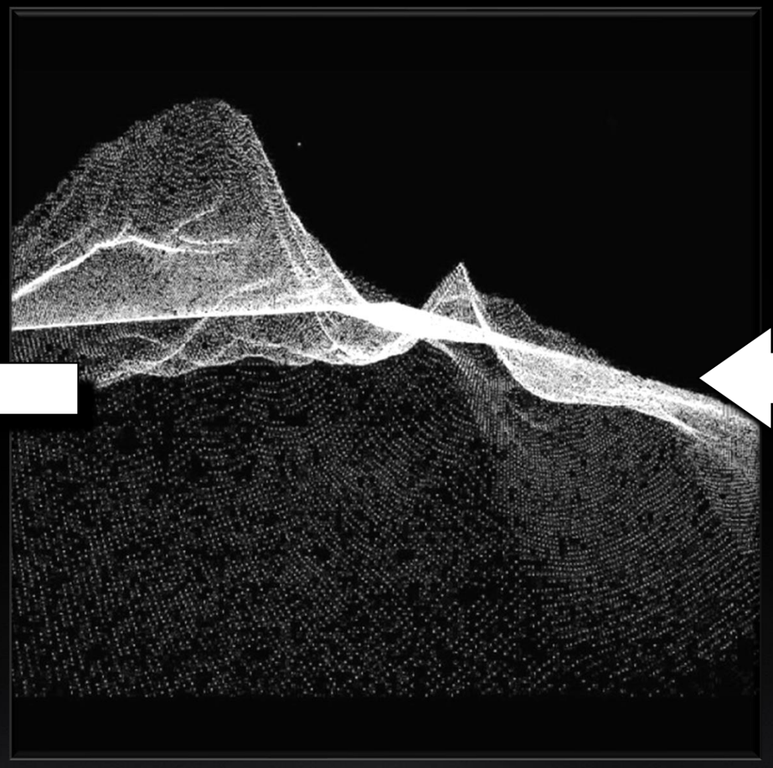

Data-Driven Models using AI, SimNet, and FNOs

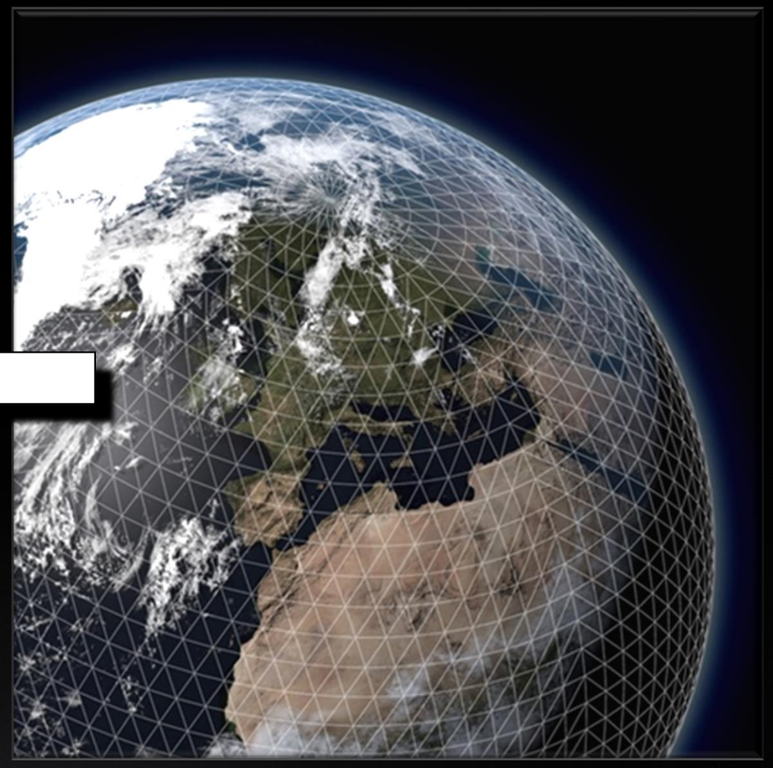

Earth Engines: Physics Based Predictions

ATOS: FULL-MODEL EMULATION

Million:1 speedups of weather and climate simulations

Purely data-driven medium-range weather forecasting achieves comparable skill to physical models at similar resolution

Stephan Rasp
Department of Informatics
Technical University of Munich
Munich, Germany
stephan.rasp@tum.de

Nils Thuerey
Department of Informatics
Technical University of Munich
Munich, Germany
nils.thuerey@tum.de

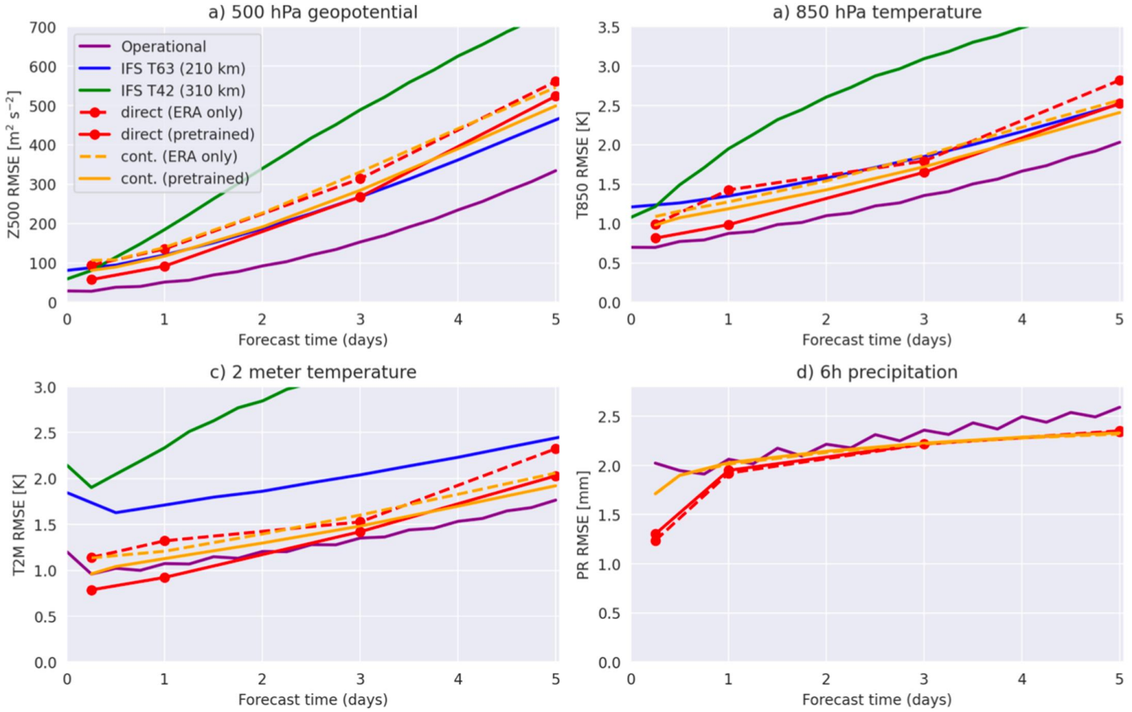

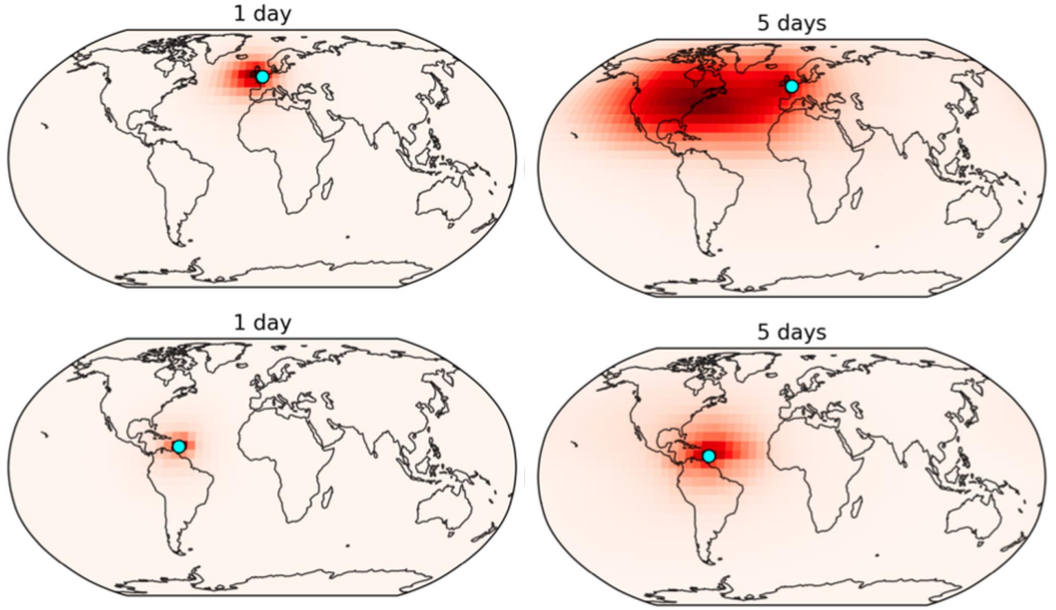

ATOS: WHOLE-MODEL EMULATION

Large ensembles improve prediction of extremes

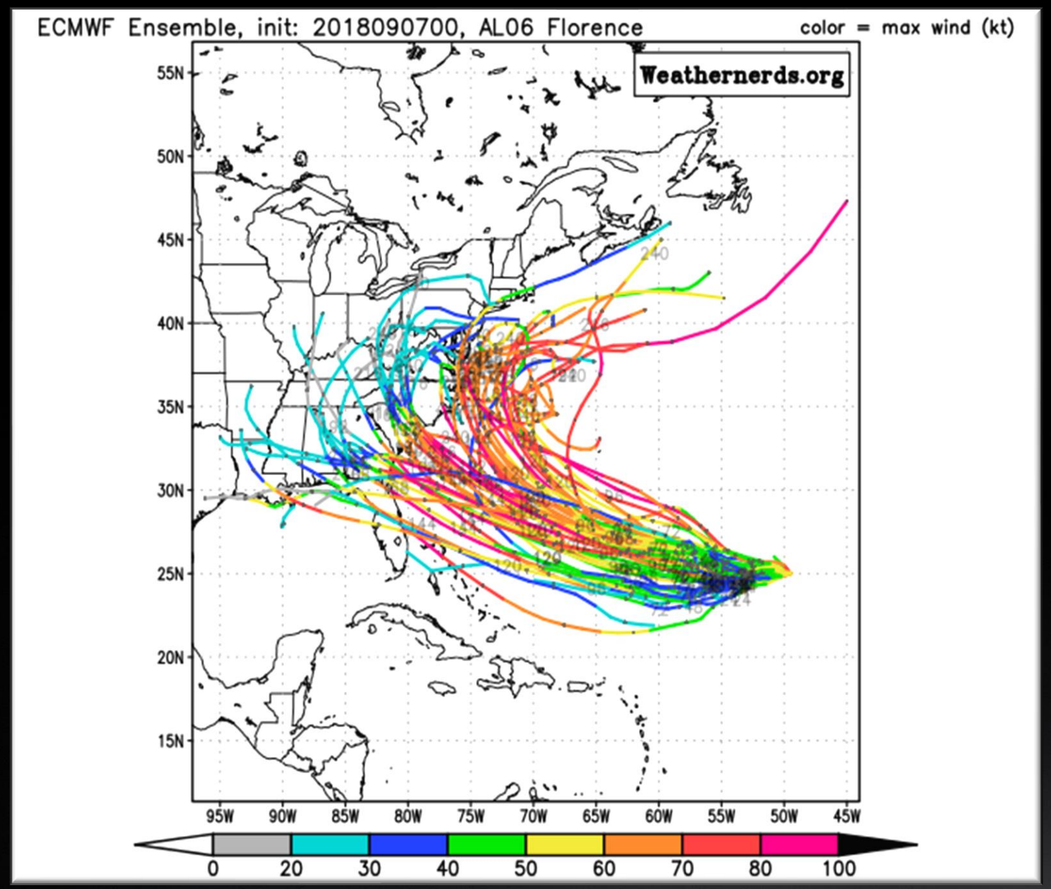

Ensemble Prediction: Hurricane Florence

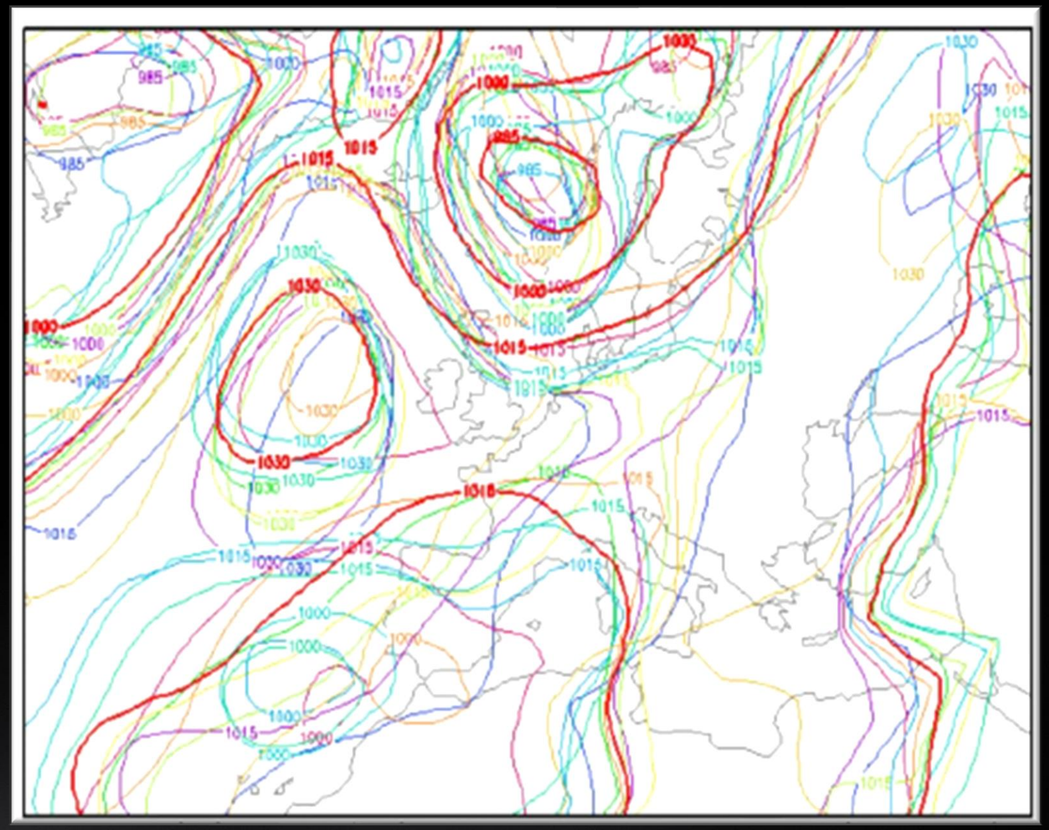

Ensemble prediction, Surface Pressure

ECMWF AND NOAA

Pre-operational Automated Tropical Cyclone Detection

Prediction example (High-Res data, 0.25°)

- Training: ~27 years in range [1990-2018]
- Validation: 3 years [1999, 2009, 2019]
- Cosine labels + MSE loss
- Run time: ~4 days on 14GPU's

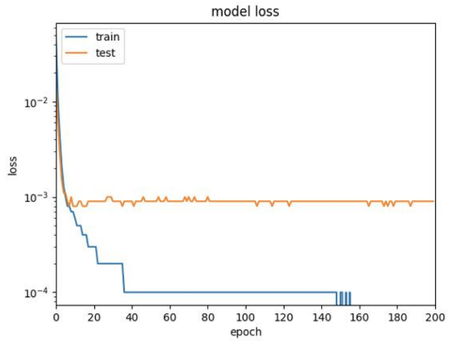

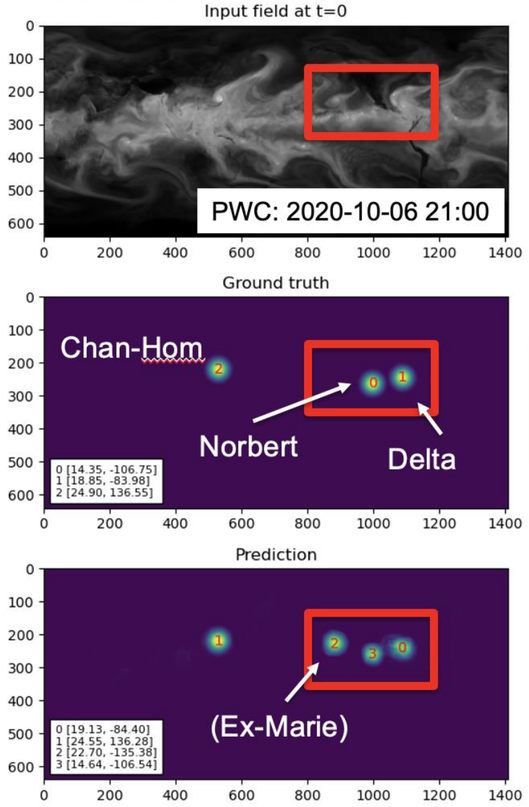

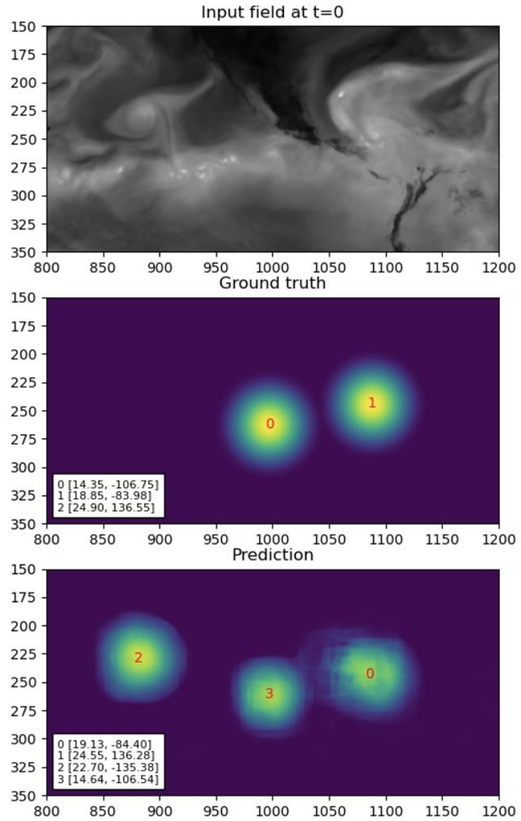

ECMWF EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS 7

DKRZ AND U-BERLIN

Nature Geoscience: Using NVIDIA's In-Painting to reconstruct missing climate data

Artificial intelligence reconstructs missing climate information

Christopher Kadow[1,2] ✉, David Matthew Hall[3] and Uwe Ulbrich[2]

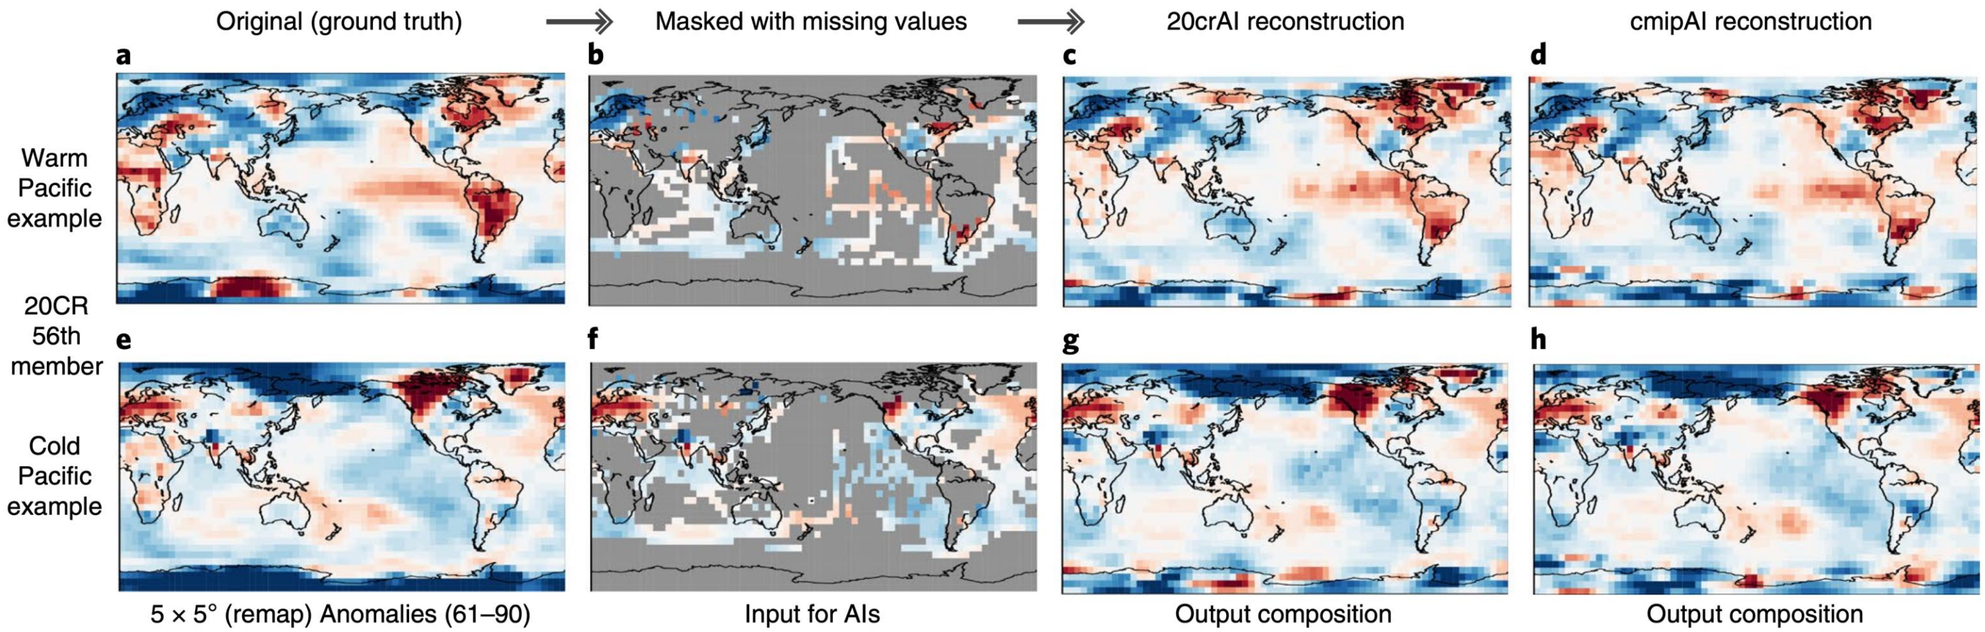

ECMWF AND ATOS

Improving Data Assimilation through Temperature Bias Correction

Training Results: time series

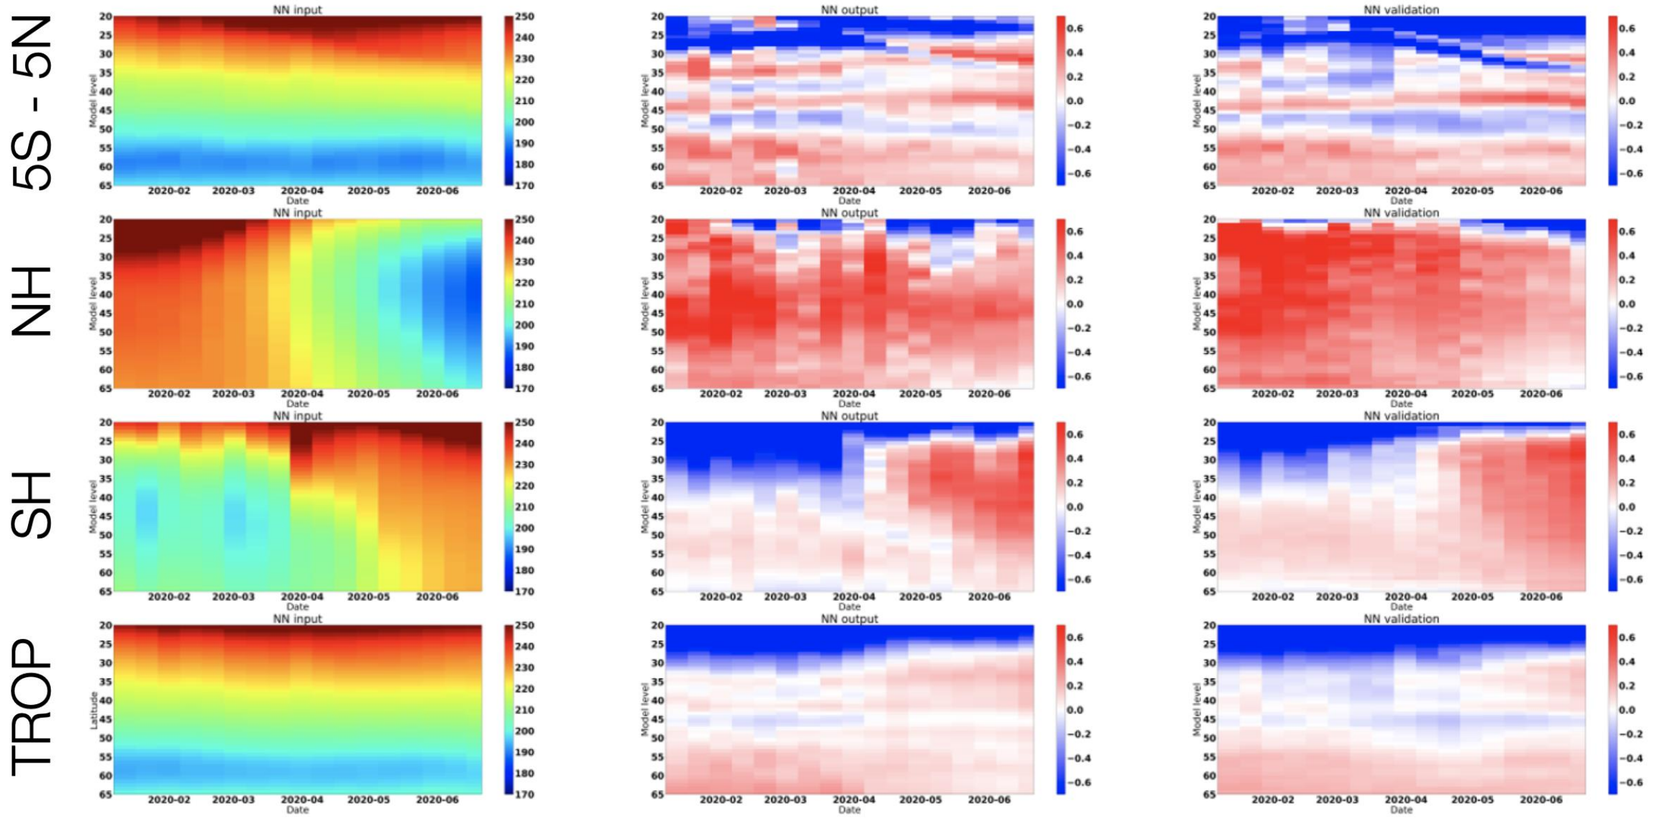

PHYSICS INFORMED NEURAL NETS

SimNet and FNOs can be used to improve the efficiency of data-driven models

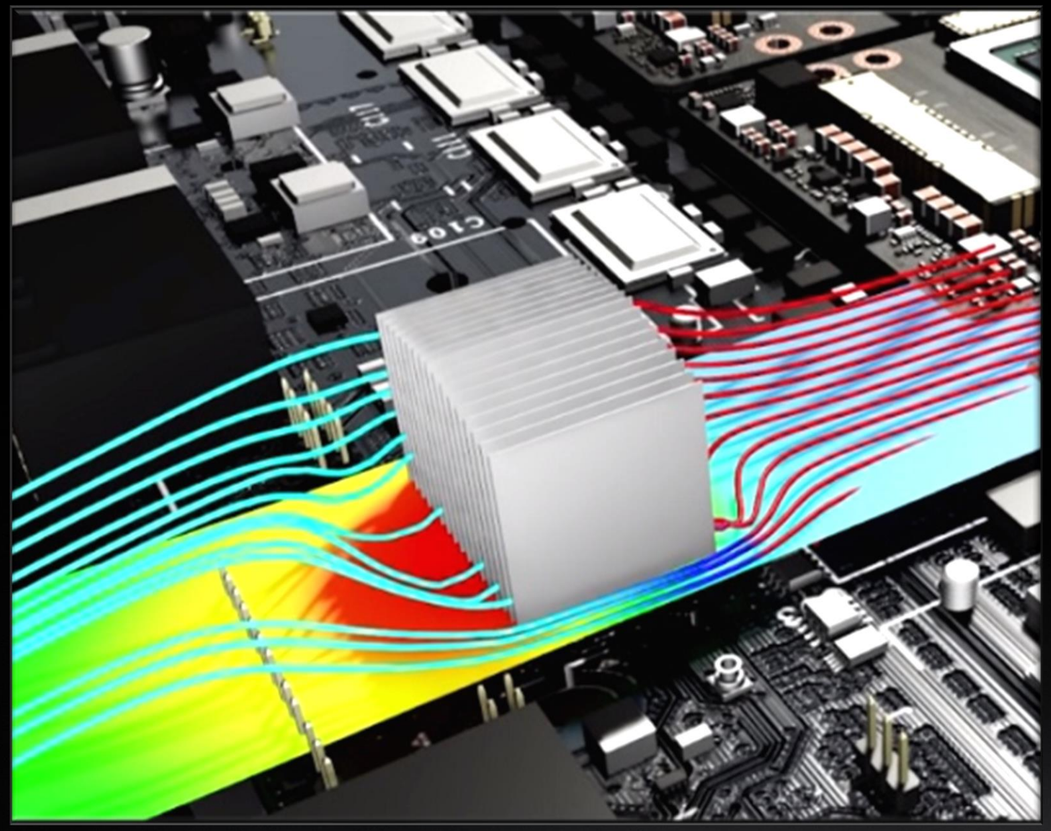

SIIMNET: Physics Informed Neural Nets

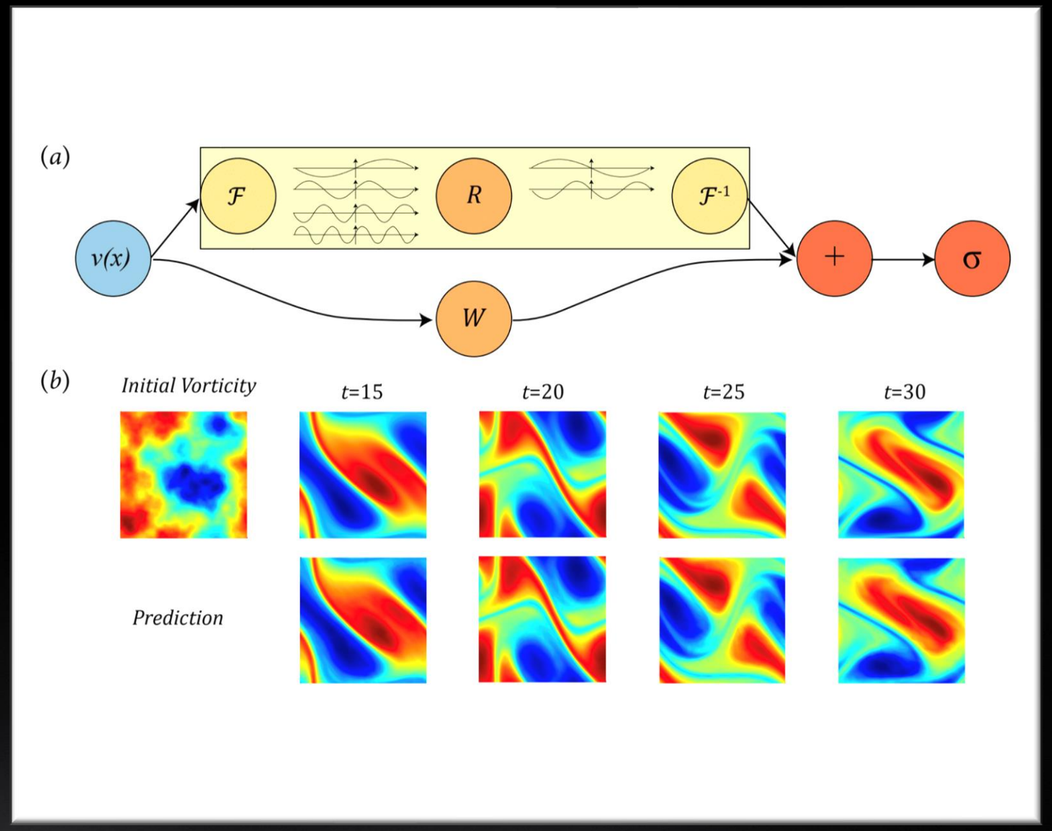

Fourier Neural Operators

VIZUALISATION

OMNIVERSE FOR AN INTUITIVE INTERFACE

Interactive, Intuitive, Real-time, Flexible, Collaborative

Omniverse Marbles Demo

Omniverse Weather Demo

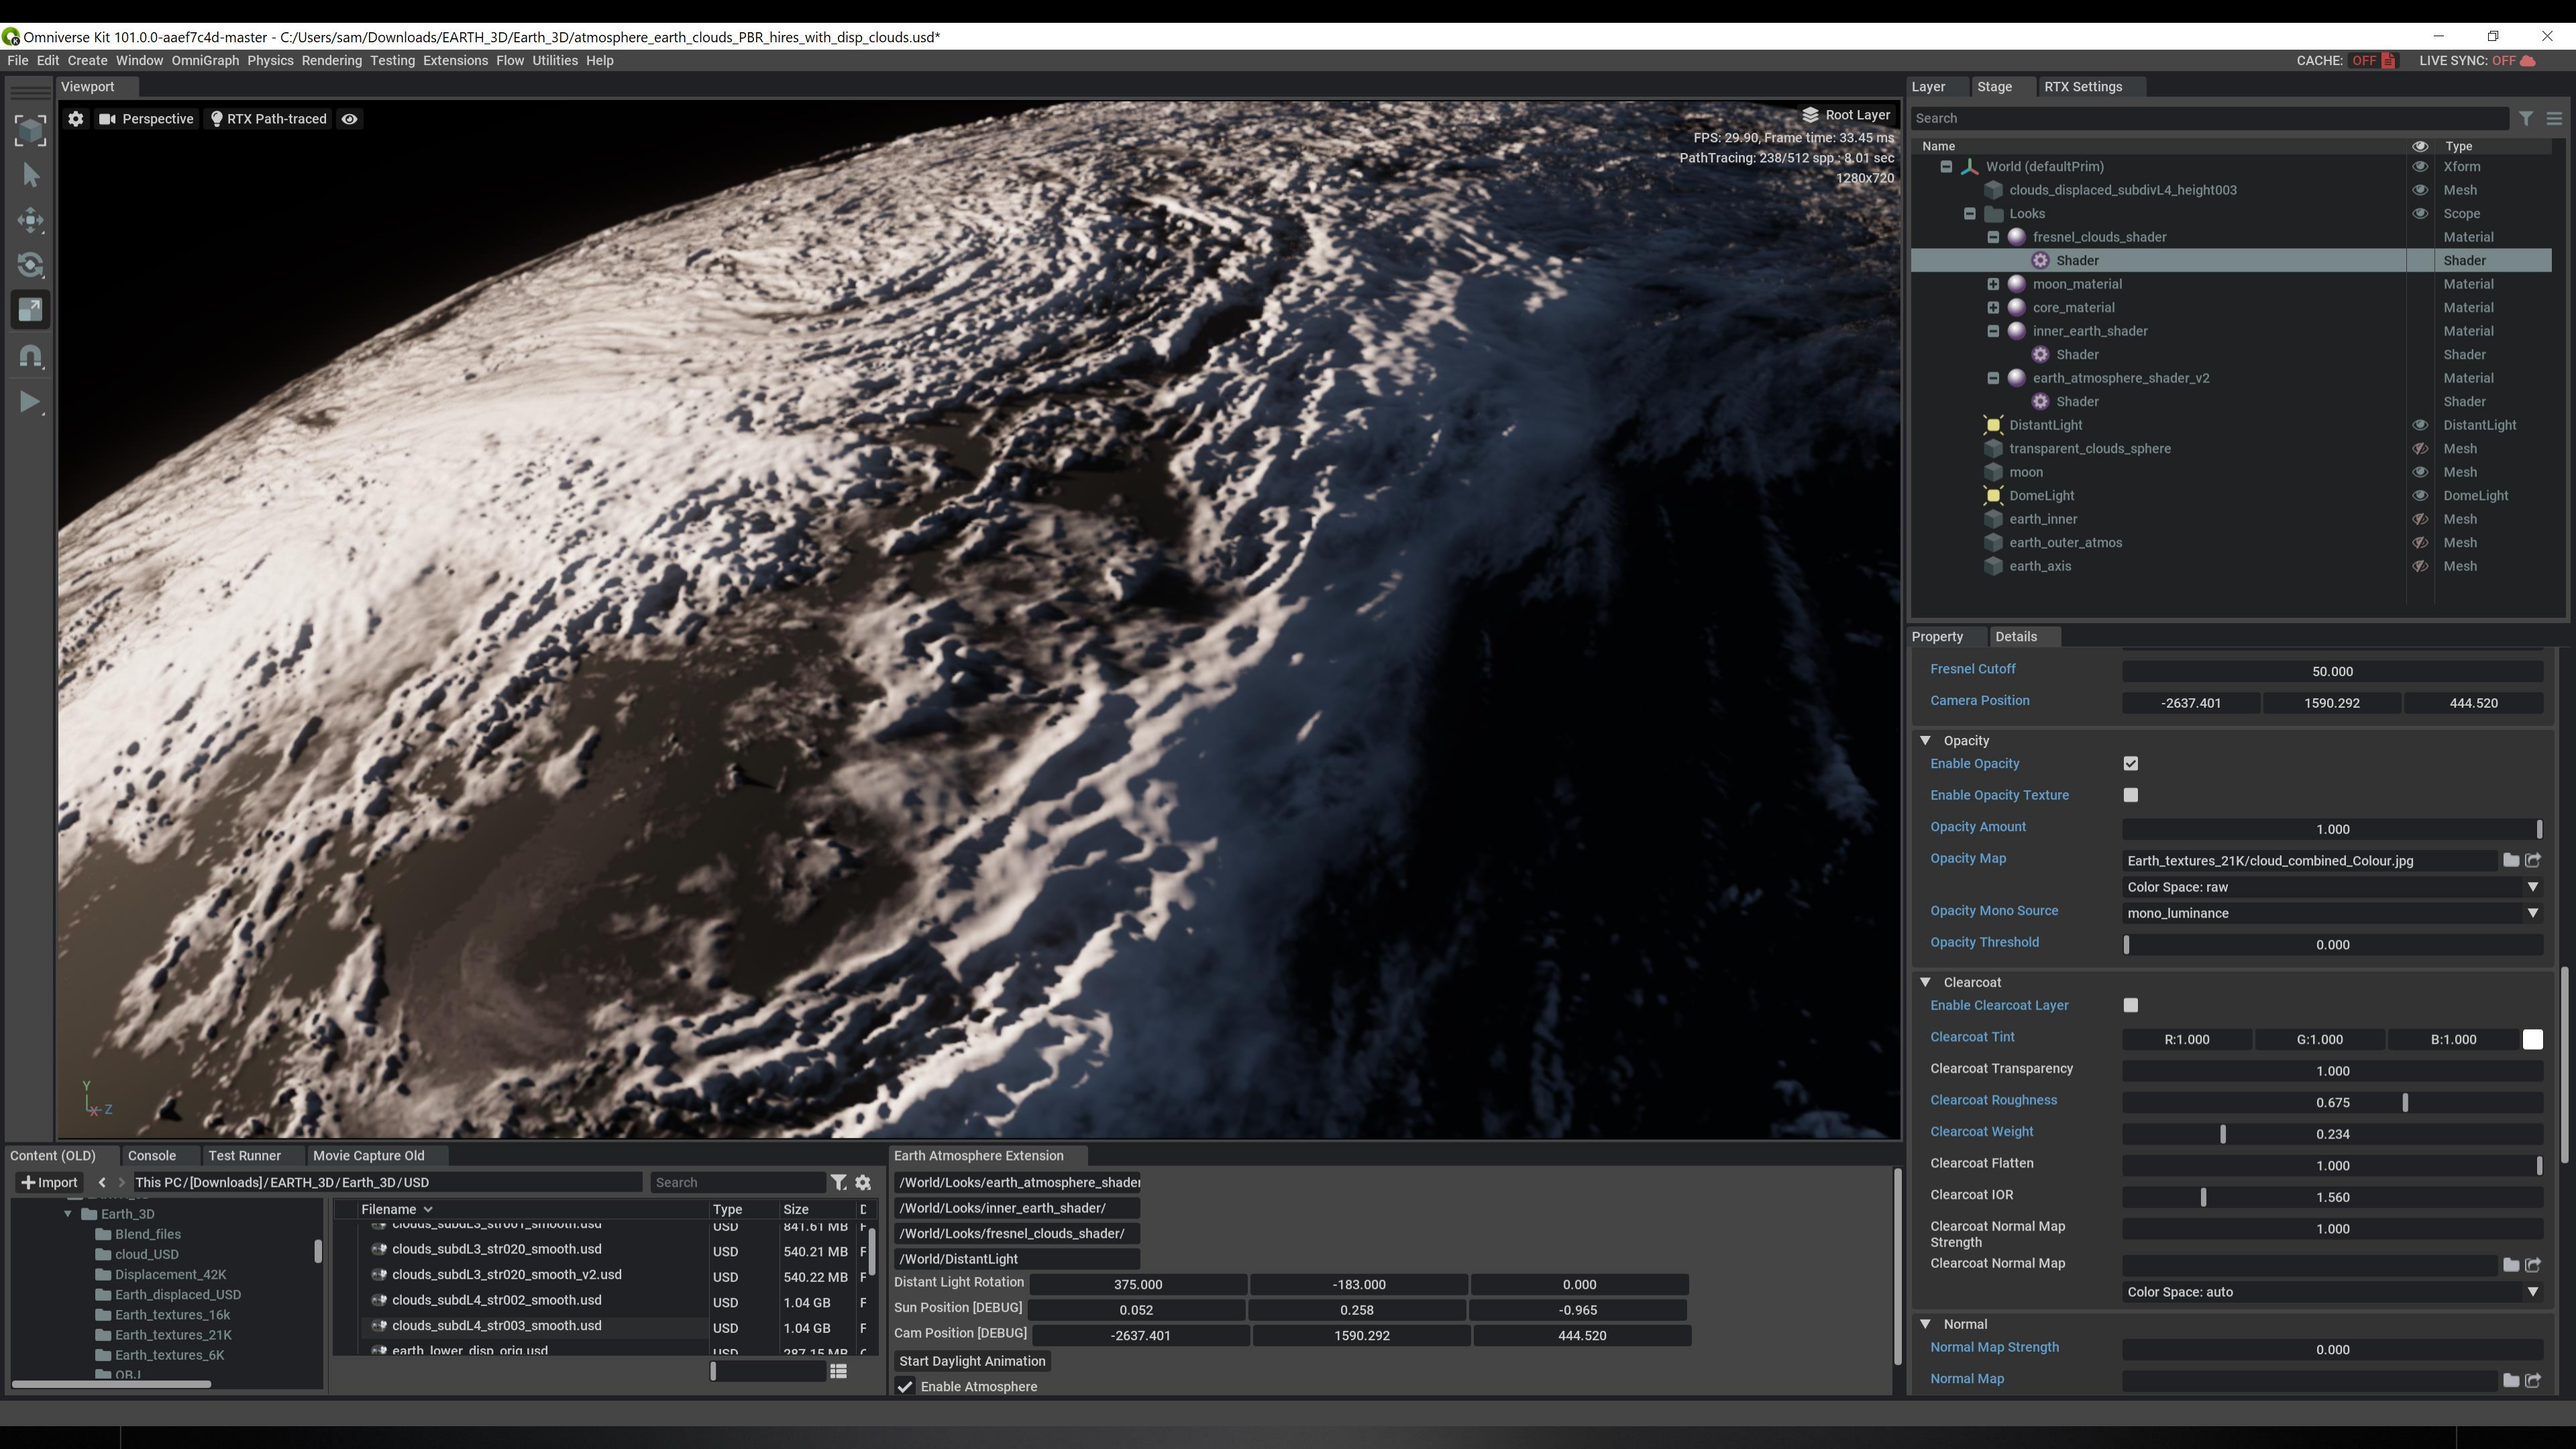
Omniverse Kit 101.0.0-aaef7c4d-master - C:/Users/sam/Downloads/EARTH_3D/Earth_3D/atmosphere_earth_clouds_PBR_hires_with_disp_clouds.usd*
File Edit Create Window OmniGraph Physics Rendering Testing Extensions Flow Utilities Help
CACHE: OFF
LIVE SYNC: OFF
Viewport
Perspective
RTX Path-traced
Root Layer
FPS: 29.90, Frame time: 33.45 ms
PathTracing: 238/512 spp - 8.01 sec
1280x720
Layer
Stage
RTX Settings
Search
Name
Type
World (defaultPrim)
Xform
clouds_displaced_subdivL4_height003
Mesh
Looks
Scope
fresnel_clouds_shader
Material
Shader
Shader
moon_material
Material
core_material
Material
inner_earth_shader
Material
Shader
Shader
earth_atmosphere_shader_v2
Material
Shader
Shader
DistantLight
DistantLight
transparent_clouds_sphere
Mesh
moon
Mesh
DomeLight
DomeLight
earth_inner
Mesh
earth_outer_atmos
Mesh
earth_axis
Mesh
Property
Details
Fresnel Cutoff
50.000
Camera Position
-2637.401
1590.292
444.520
Opacity
Enable Opacity
Enable Opacity Texture
Opacity Amount
1.000
Opacity Map
Earth_textures_21K/cloud_combined_Colour.jpg
Color Space: raw
Opacity Mono Source
mono_luminance
Opacity Threshold
0.000
Clearcoat
Enable Clearcoat Layer
Clearcoat Tint
R:1.000
G:1.000
B:1.000
Clearcoat Transparency
1.000
Clearcoat Roughness
0.675
Clearcoat Weight
0.234
Clearcoat Flatten
1.000
Clearcoat IOR
1.560
Clearcoat Normal Map Strength
1.000
Clearcoat Normal Map
Color Space: auto
Normal
Normal Map Strength
0.000
Normal Map
Content (OLD)
Console
Test Runner
Movie Capture Old
Import
This PC/[Downloads]/EARTH_3D/Earth_3D/USD
Search
Earth_3D
Blend_files
cloud_USD
Displacement_42K
Earth_displaced_USD
Earth_textures_16k
Earth_textures_21K
Earth_textures_6K
OBJ
Filename
Type
Size
clouds_subdL3_str020_smooth.usd
USD
540.21 MB
clouds_subdL3_str020_smooth_v2.usd
USD
540.22 MB
clouds_subdL4_str002_smooth.usd
USD
1.04 GB
clouds_subdL4_str003_smooth.usd
USD
1.04 GB
Earth Atmosphere Extension
/World/Looks/earth_atmosphere_shader
/World/Looks/inner_earth_shader/
/World/Looks/fresnel_clouds_shader/
/World/DistantLight
Distant Light Rotation
375.000
-183.000
0.000
Sun Position [DEBUG]
0.052
0.258
-0.965
Cam Position [DEBUG]
-2637.401
1590.292
444.520
Start Daylight Animation
Enable Atmosphere

Atmospheric Simulation Data in Omniverse

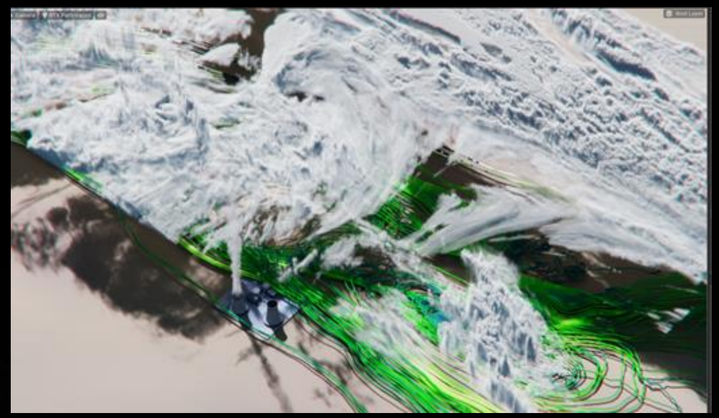
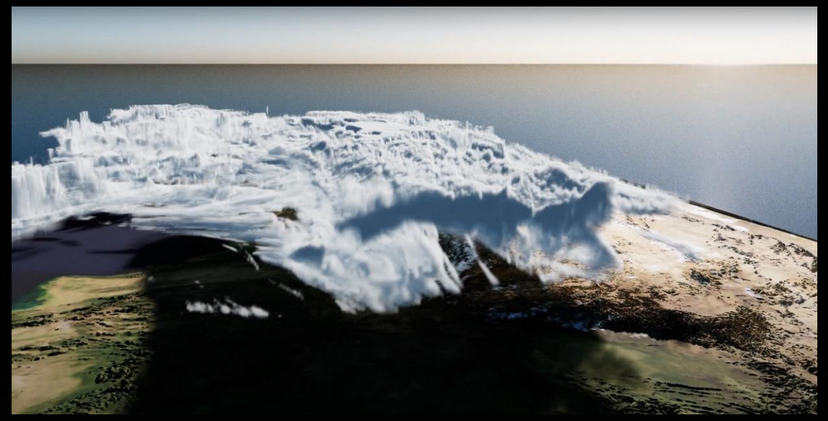
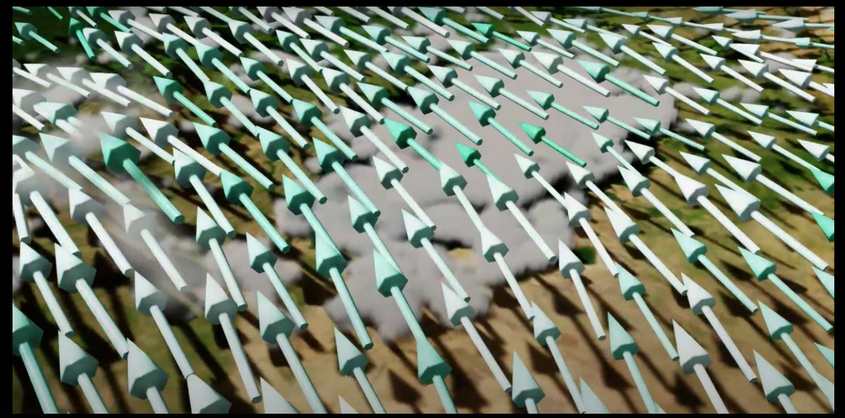
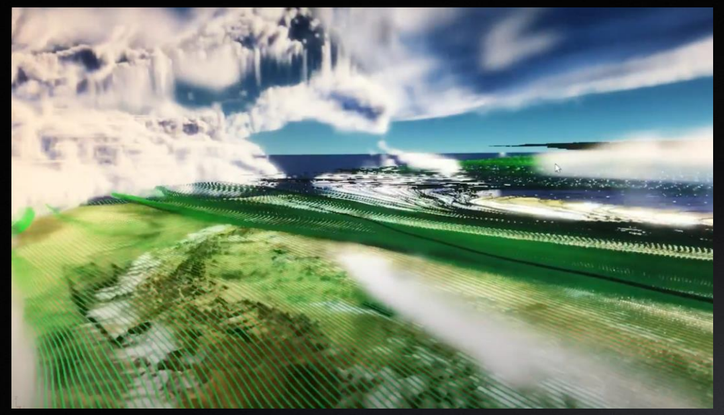

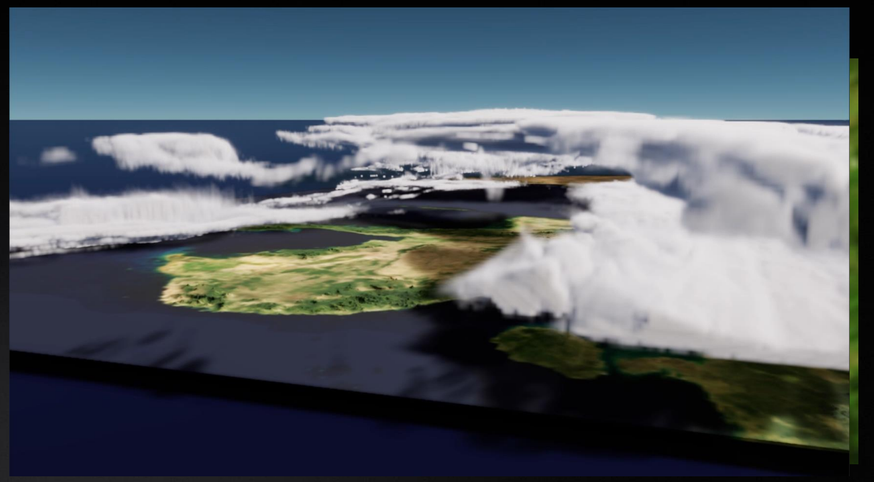

IFS (INCITE), ICON

WHAT IF?

Playing through Scenarios with Real Time Physics

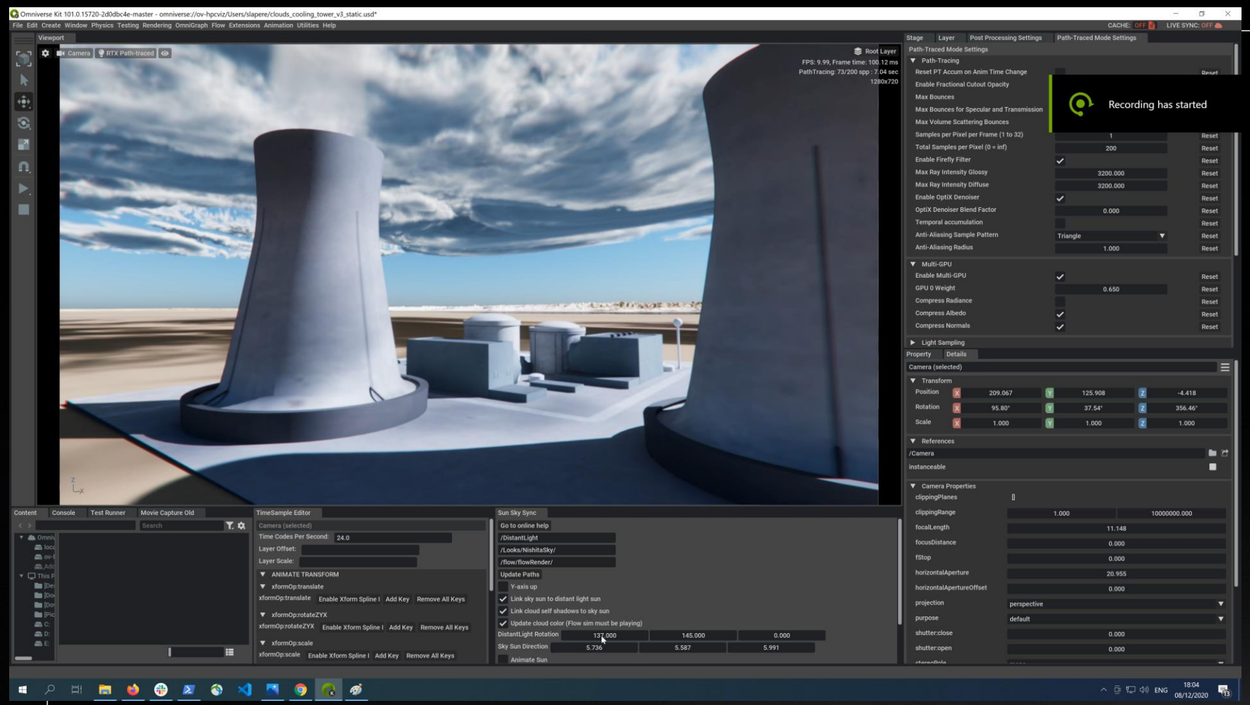

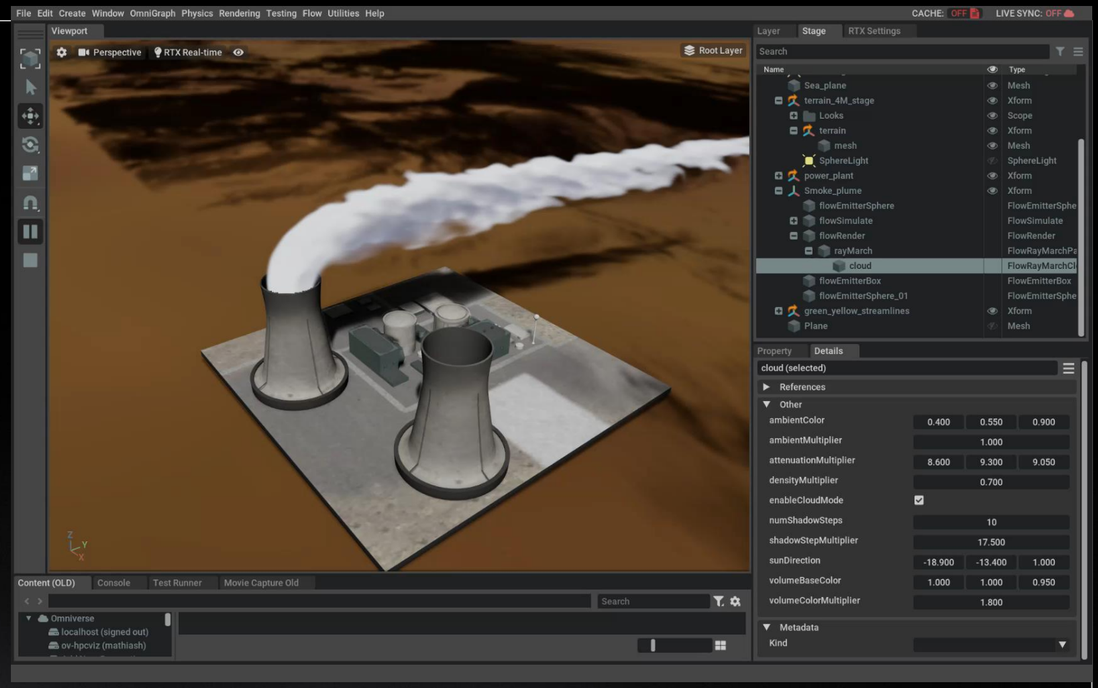

A DIGITAL TWIN EARTH

MAJOR CLIMATE-MODELING INITIATIVES

Destination Earth, VULCAN, CLIMA, LEAP

Destination Earth

AI2 Climate Model

Ai2
Allen Institute for AI

CLIMA

LEAP

Columbia University

DESTINATION-EARTH

Project DestinE envisions what Earth-system modeling could be

Intuitive User Interface

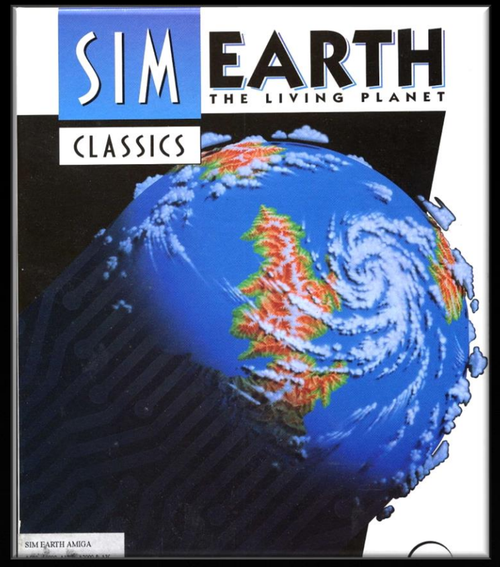

What-if Q & A

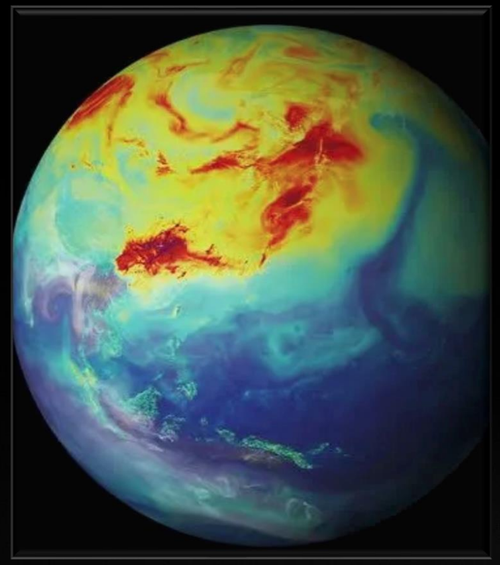

Storm-resolving Models

Unified Observations

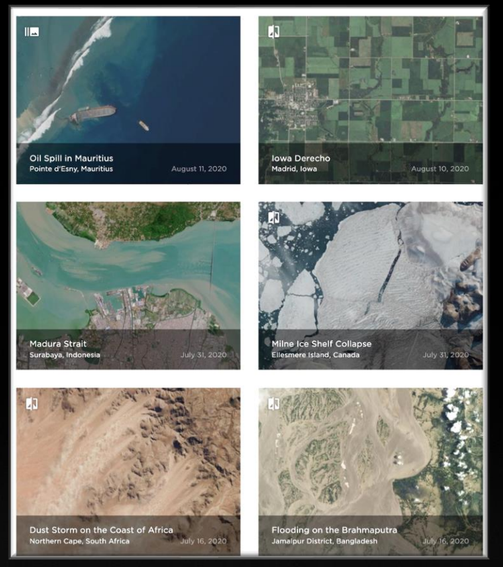

Exascale Compute

https://visgallery.ucar.edu/global-weather-simulated-by-a-cloud-resolving-weather-prediction-model/
https://www.planet.com/gallery/

NVIDIA.

DESTINATION-EARTH

NVIDIA has the technologies needed to make this vision a reality

Intuitive User Interface	What-if Q & A	Storm-resolving Models	Unified Observations	Exascale Compute
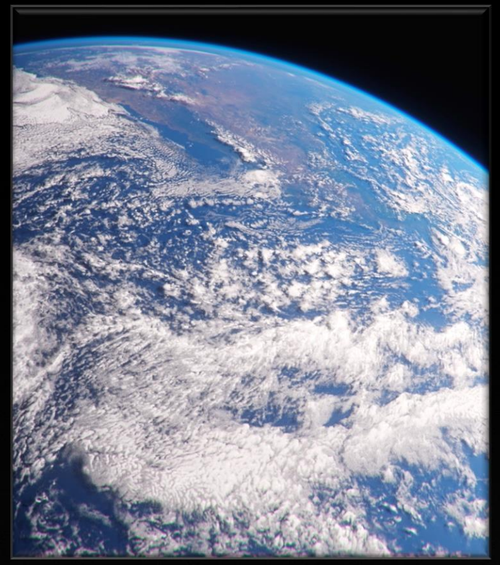	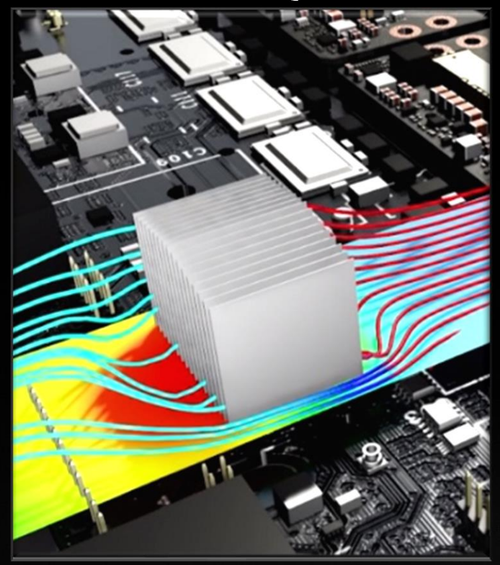			
OMNIVERSE	AI / SIMNET / FNO	GPU Accel + Hardware	OV NUCLEUS SERVER	Superpod + Grace CPU

https://visgallery.ucar.edu/global-weather-simulated-by-a-cloud-resolving-weather-prediction-model/
https://www.planet.com/gallery/

SUMMARY

- HPC Modelling and Simulation is key for determining future extreme events
- Data driven models provides quick answers for what-if scenarios
- Visualisation provides insight into our data, both models and observations
- Combining all of these together into a digital twin can provide a platform for improved climate science
- NVIDIA has the tools and capabilities to tackle all of these aspects.

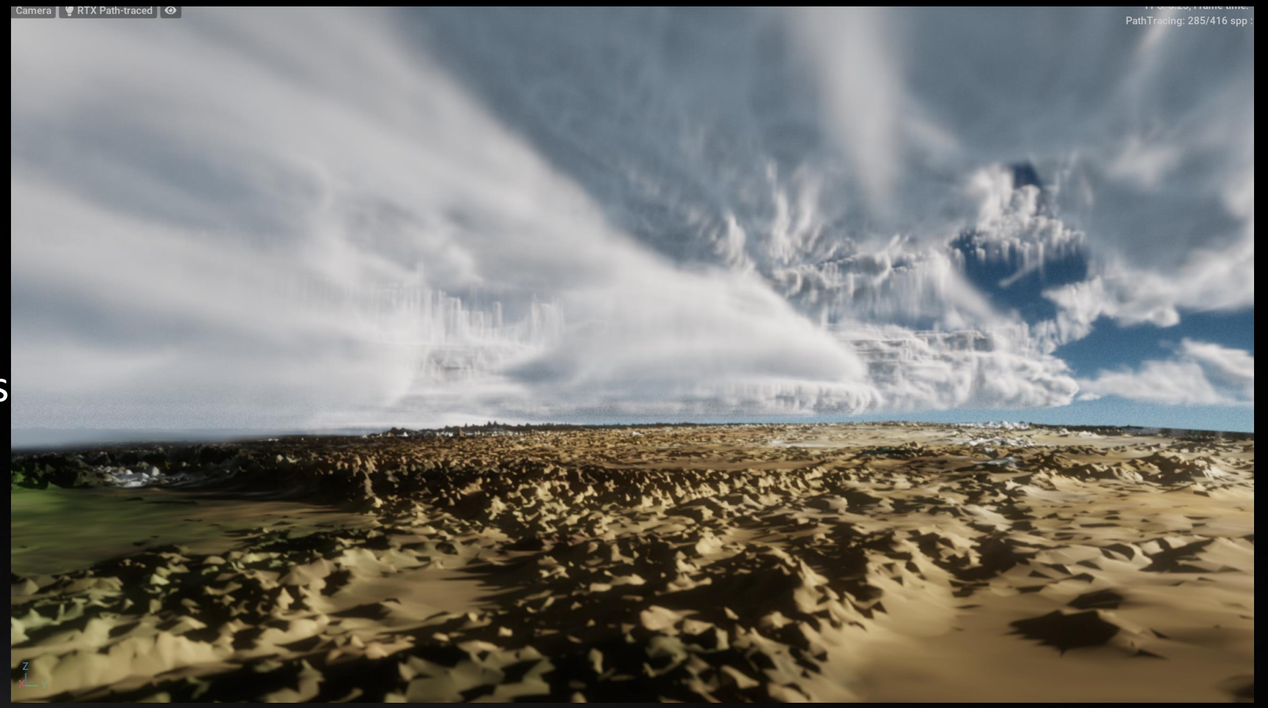

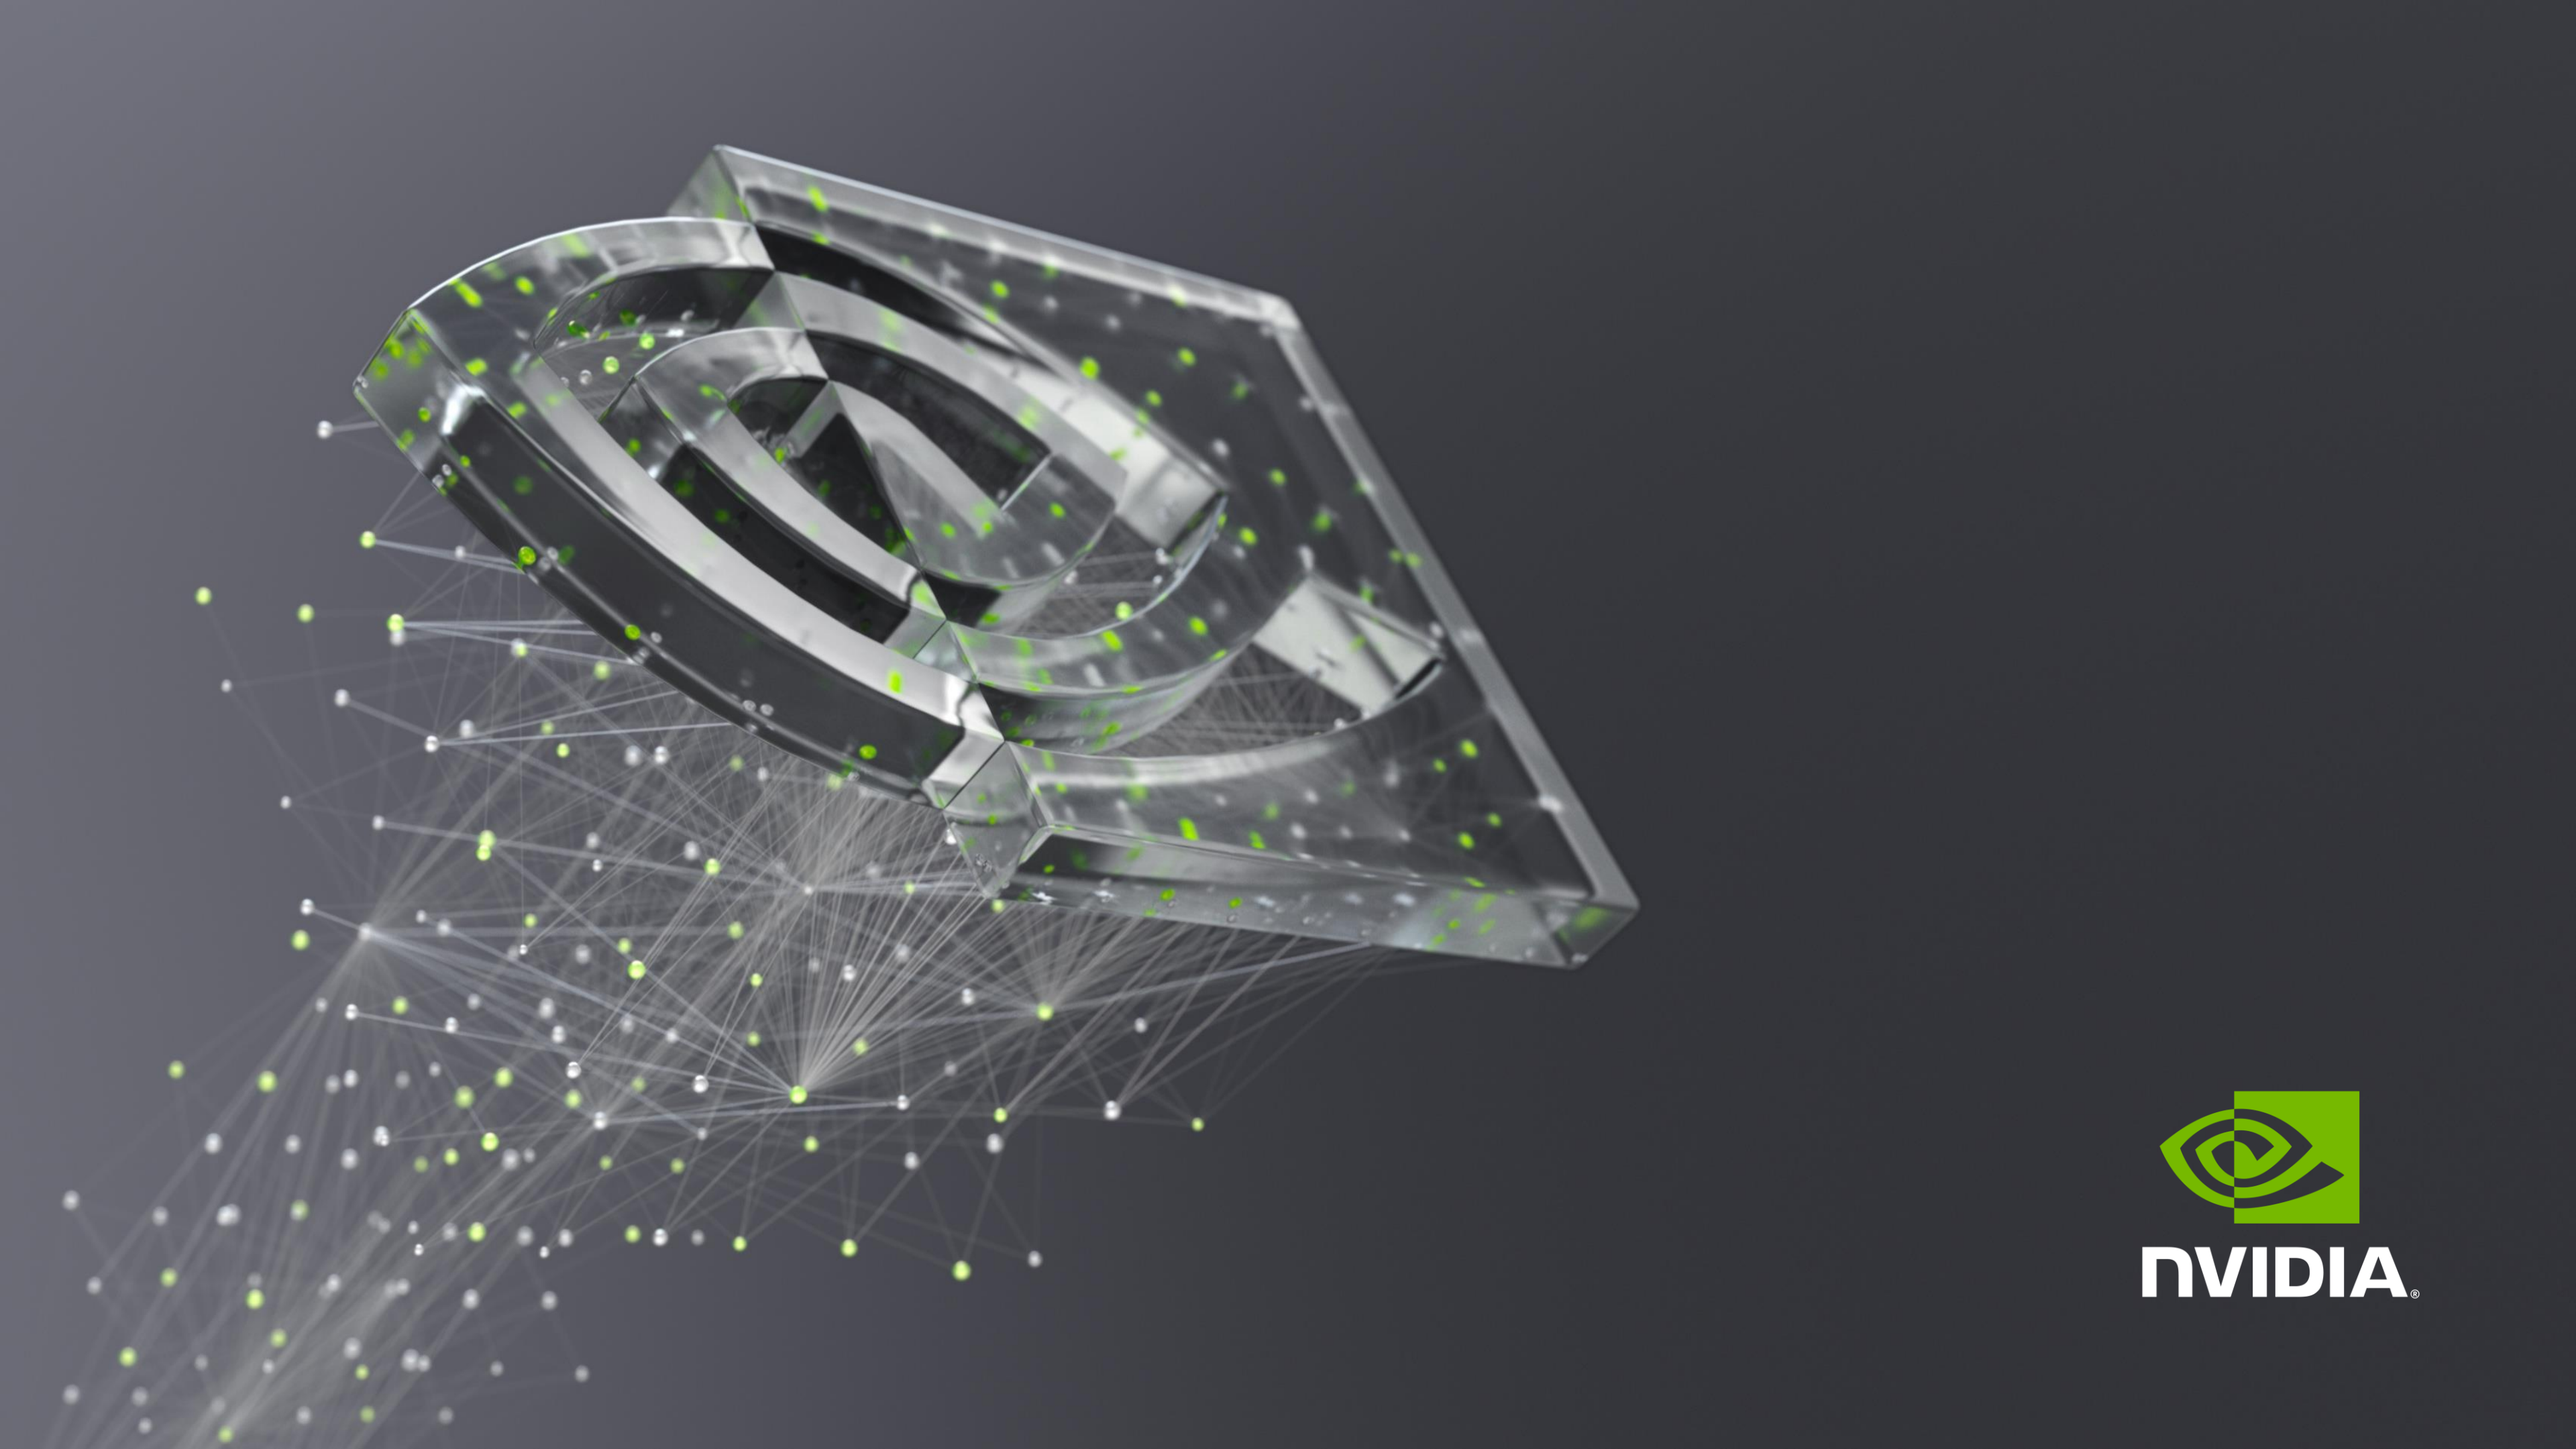
NVIDIA