NVIDIA and AI For Science

CUCAI 2026 | University of Toronto | 2026 Mar 7 | Jonathan Dursi

NVIDIA AI & HPC Platform

Chips Purpose-Built for AI Supercomputing
GPU | CPU | DPU | NIC | NVLink Switch | IB Switch | Enet Switch

AI for Science(s)

cuDSS
CAE

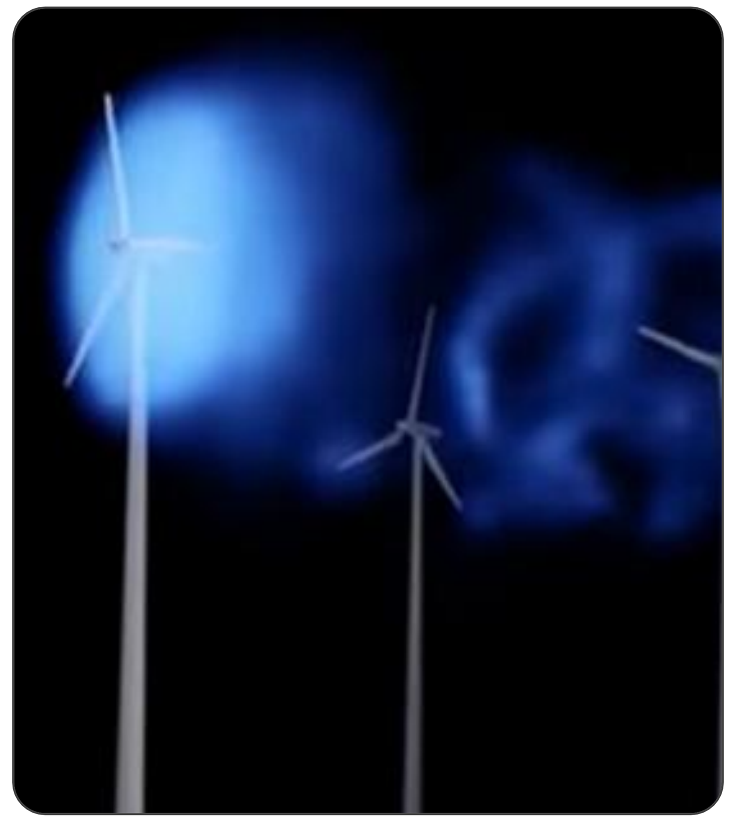

PhysicsNeMo
AI Physics

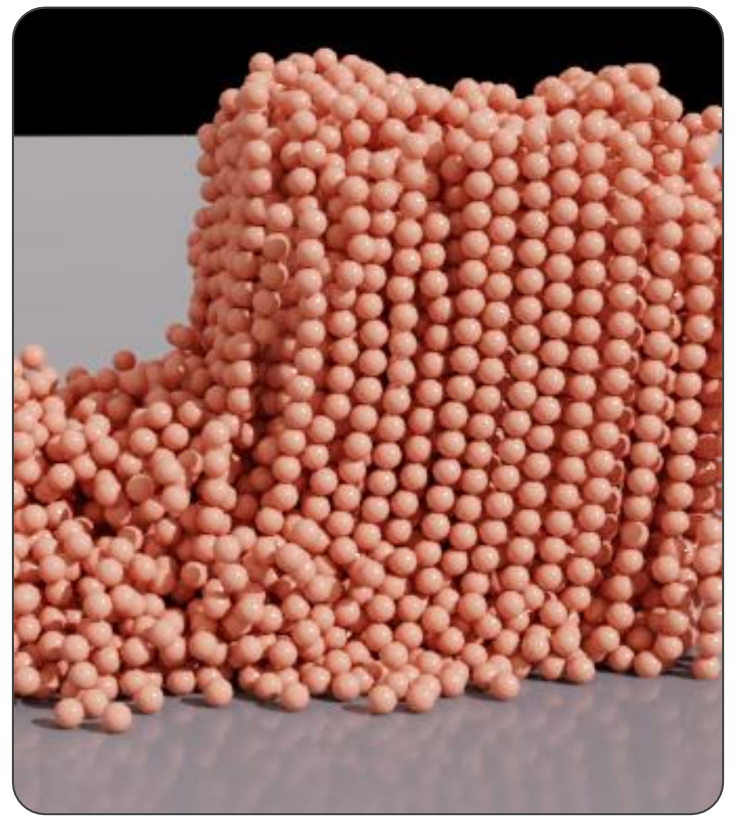

Warp
Physical Simulation

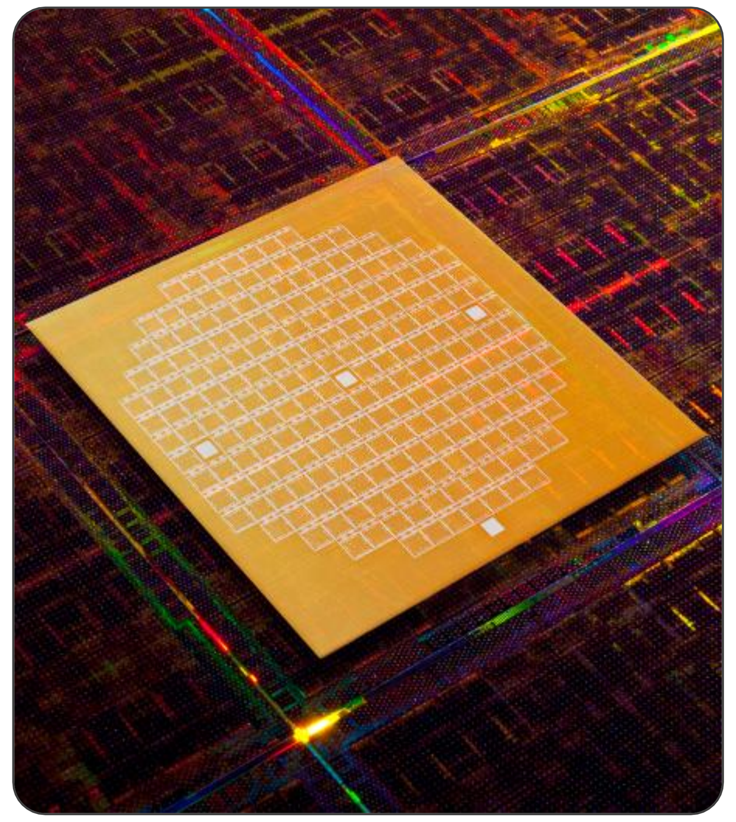

cuLitho
Computational Lithography

ALCHEMI
AI Materials Science

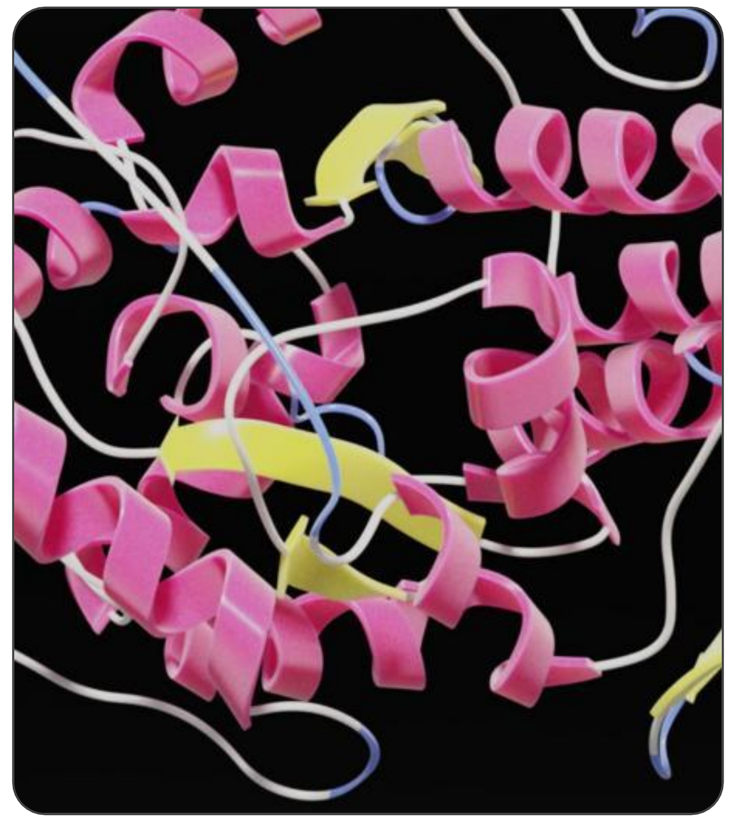

cuEquivariance
Drug & Materials Discovery

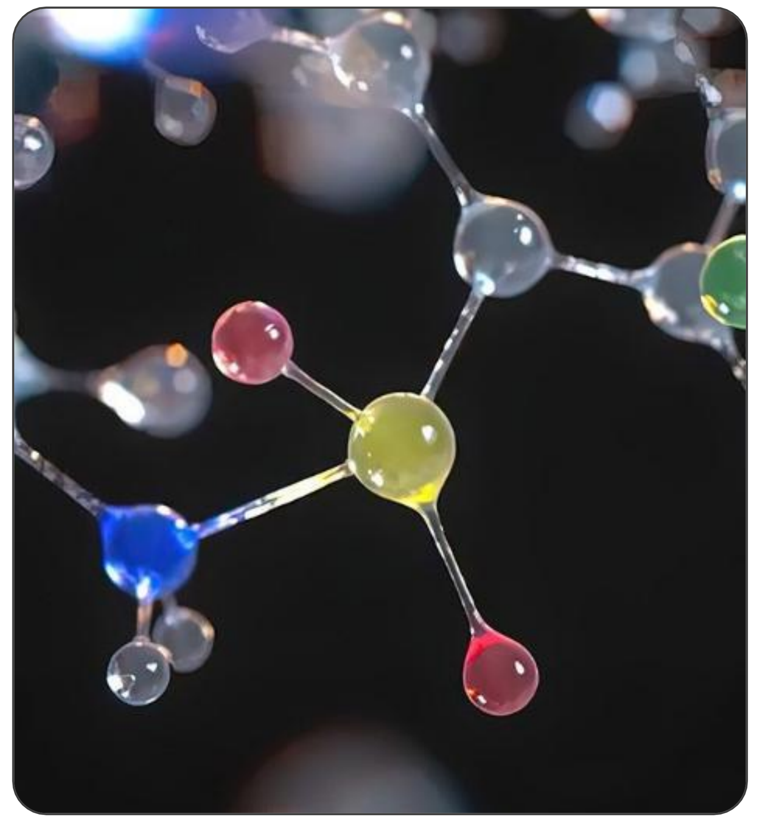

Parabricks
Gene Sequencing

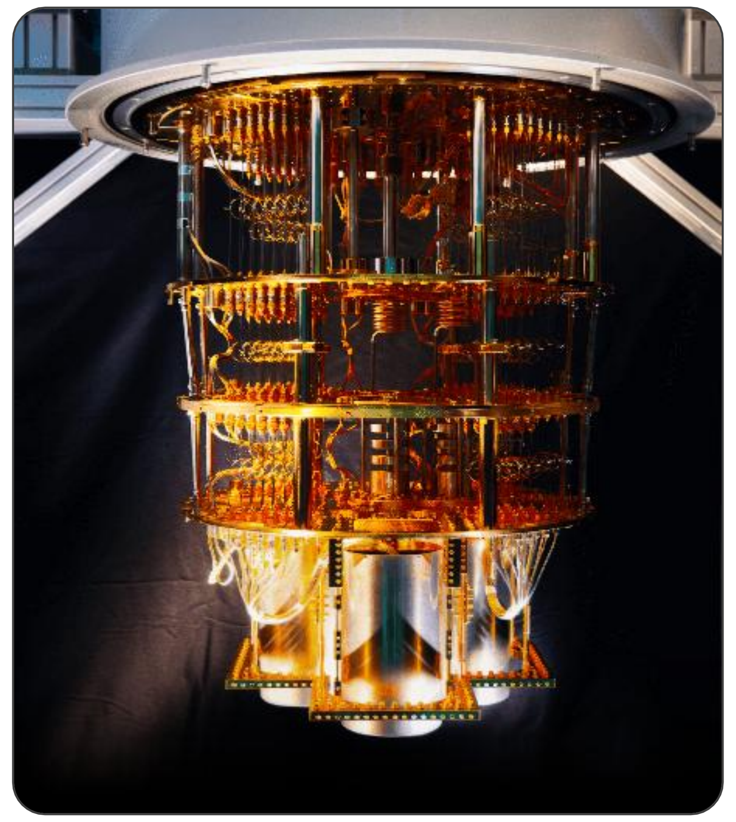

CUDA-Q
Quantum Computing

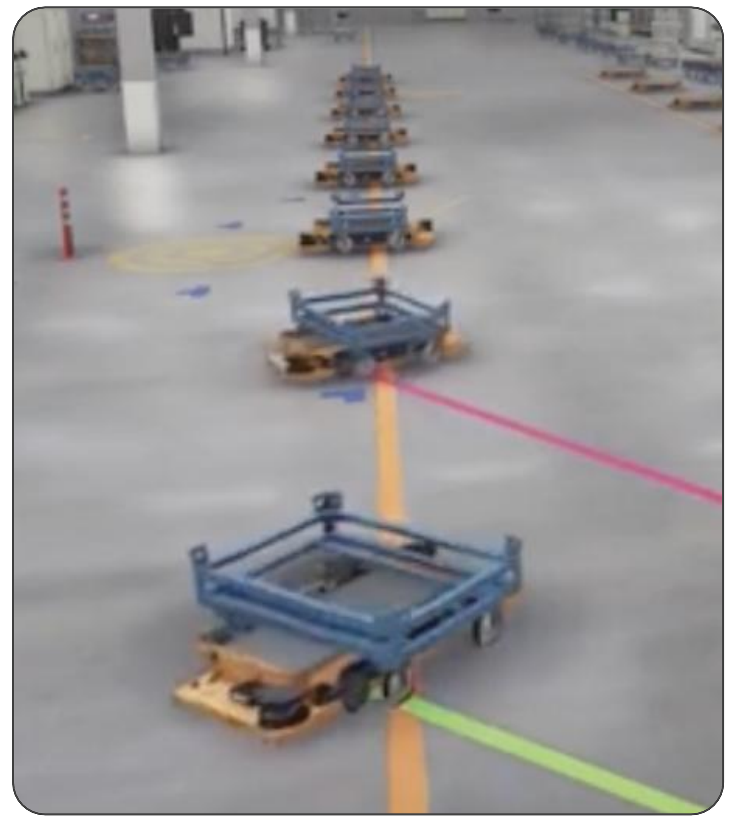

cuOpt
Decision Optimization

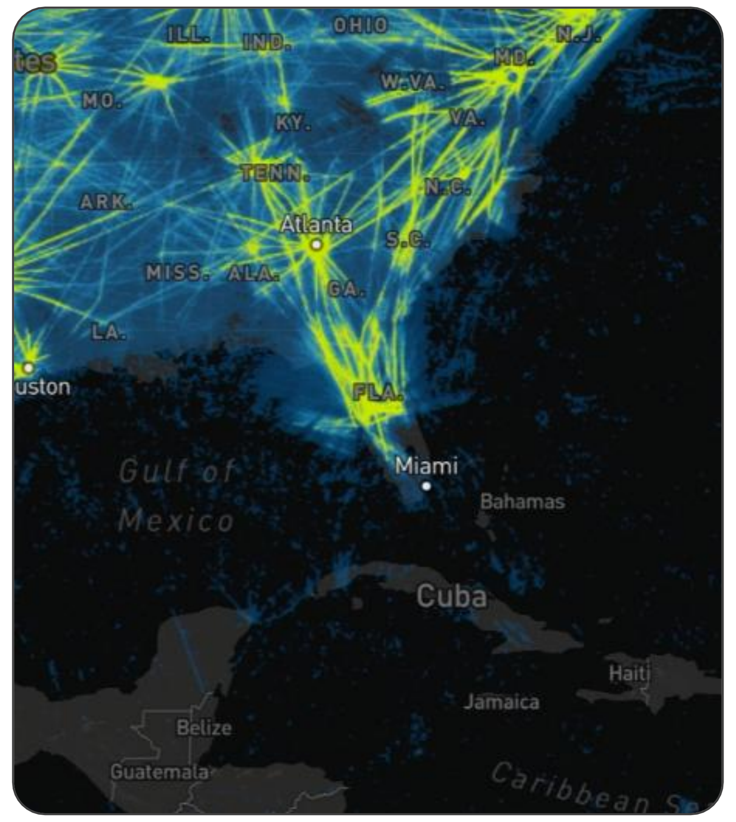

cuDF
Data Processing

cuPyNumeric
Numerical Computing

Holoscan
Edge HPC

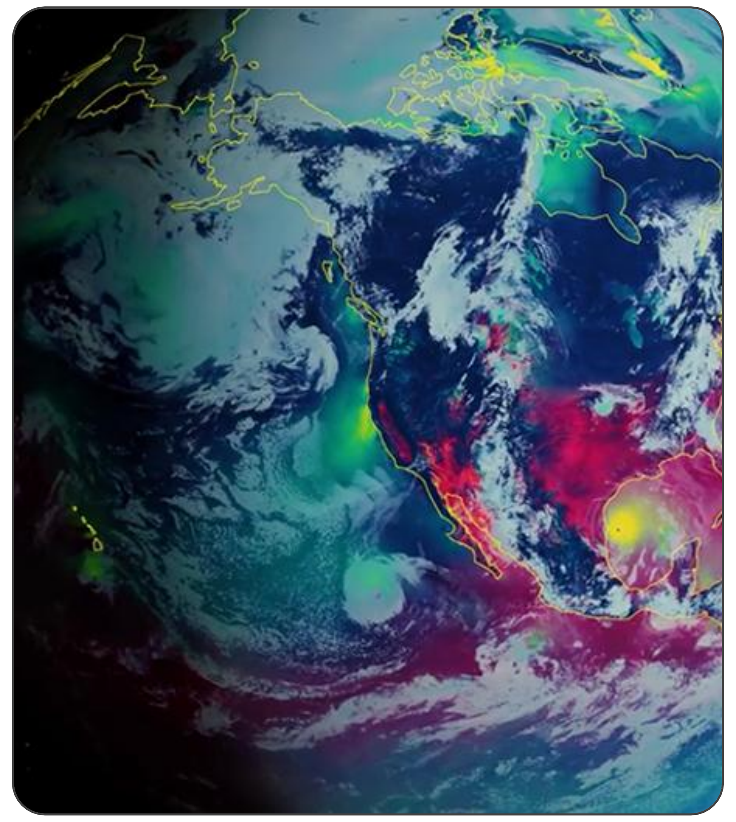

Earth-2
Weather Analytics

“

One of the most amazing things that anyone has ever said to me came from a quantum chemist -‘because of NVIDIA’s work, I can do my life’s work *in my lifetime*’

NVIDIA founder and CEO, Jensen Huang

”

AI Accelerates Scientific Discovery

With Traditional Science, Discovery Takes Years

AERODYNAMICS

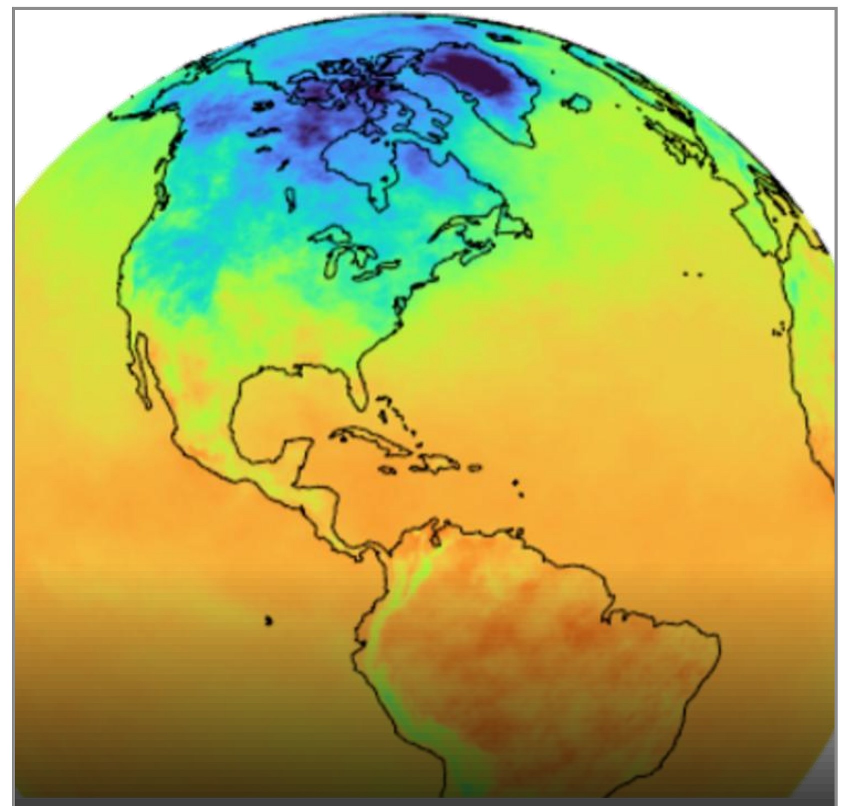

CLIMATE AND WEATHER

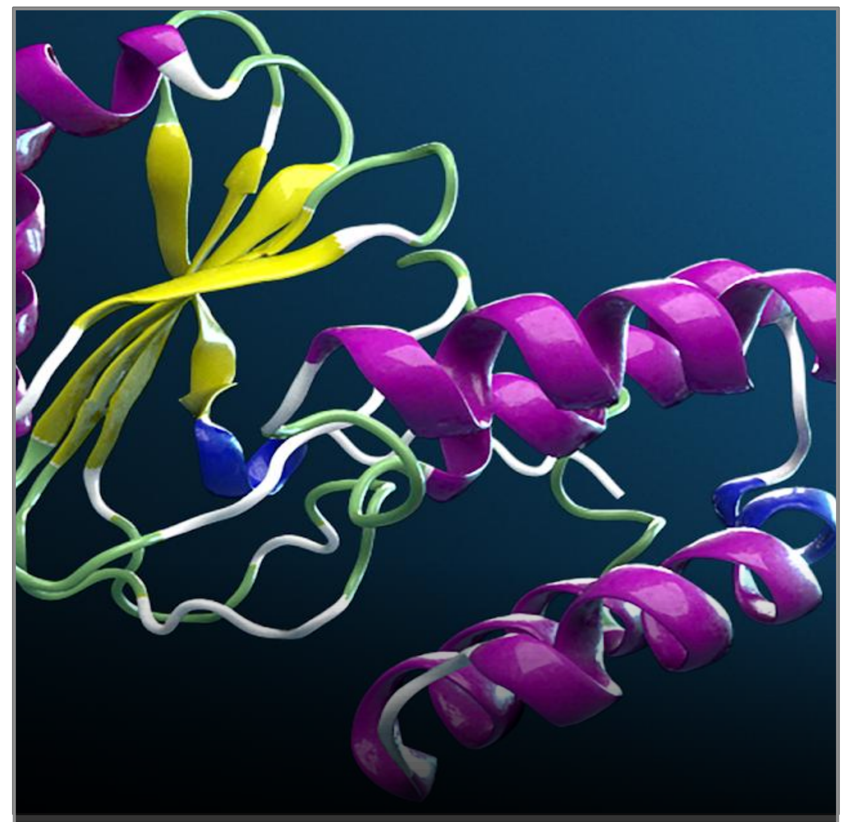

DRUG DISCOVERY

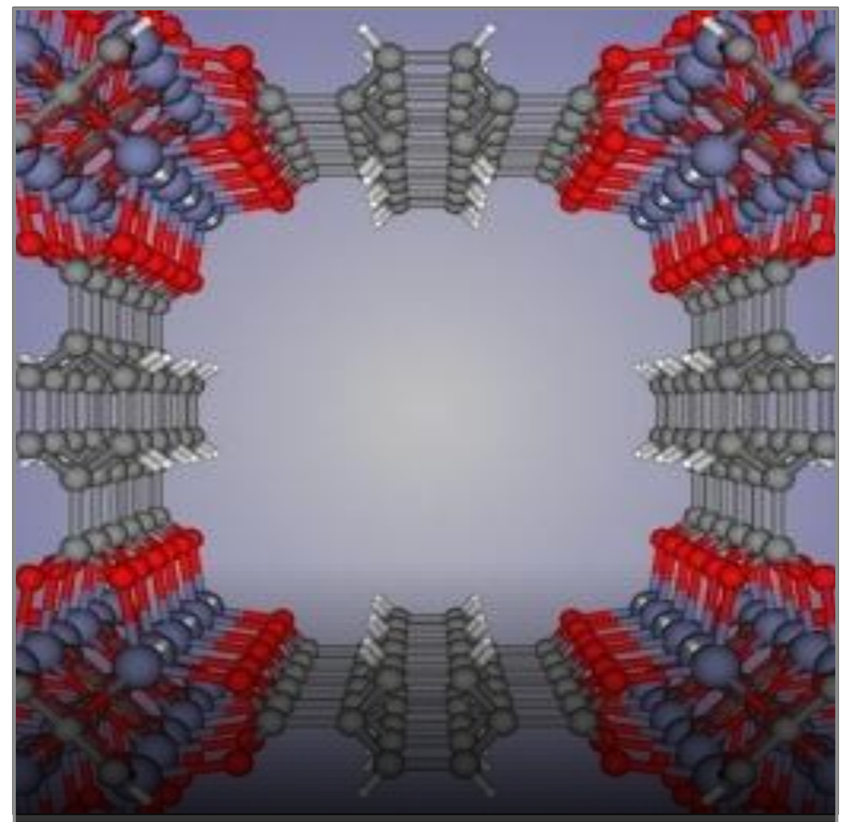

MATERIALS DISCOVERY

Across Domains, AI Shrinks Time-to-Science

Chemistry
Demis Hassabis, John Jumper for AlphaFold
David Baker for Rosetta (protein design)

Physics
Geoff Hinton and John Hopfield
for artificial neural networks

Groundbreaking Achievements with NVIDIA

Photo by Hans Reiners on Unsplash

Challenges - Saturating performance in traditional HPC

Simulations are getting larger and more complex

Traditional solution methods are:

- Computationally Expensive
- Plagued by Domain Discretization Techniques
- Not suitable for Data-assimilation or Inverse problems

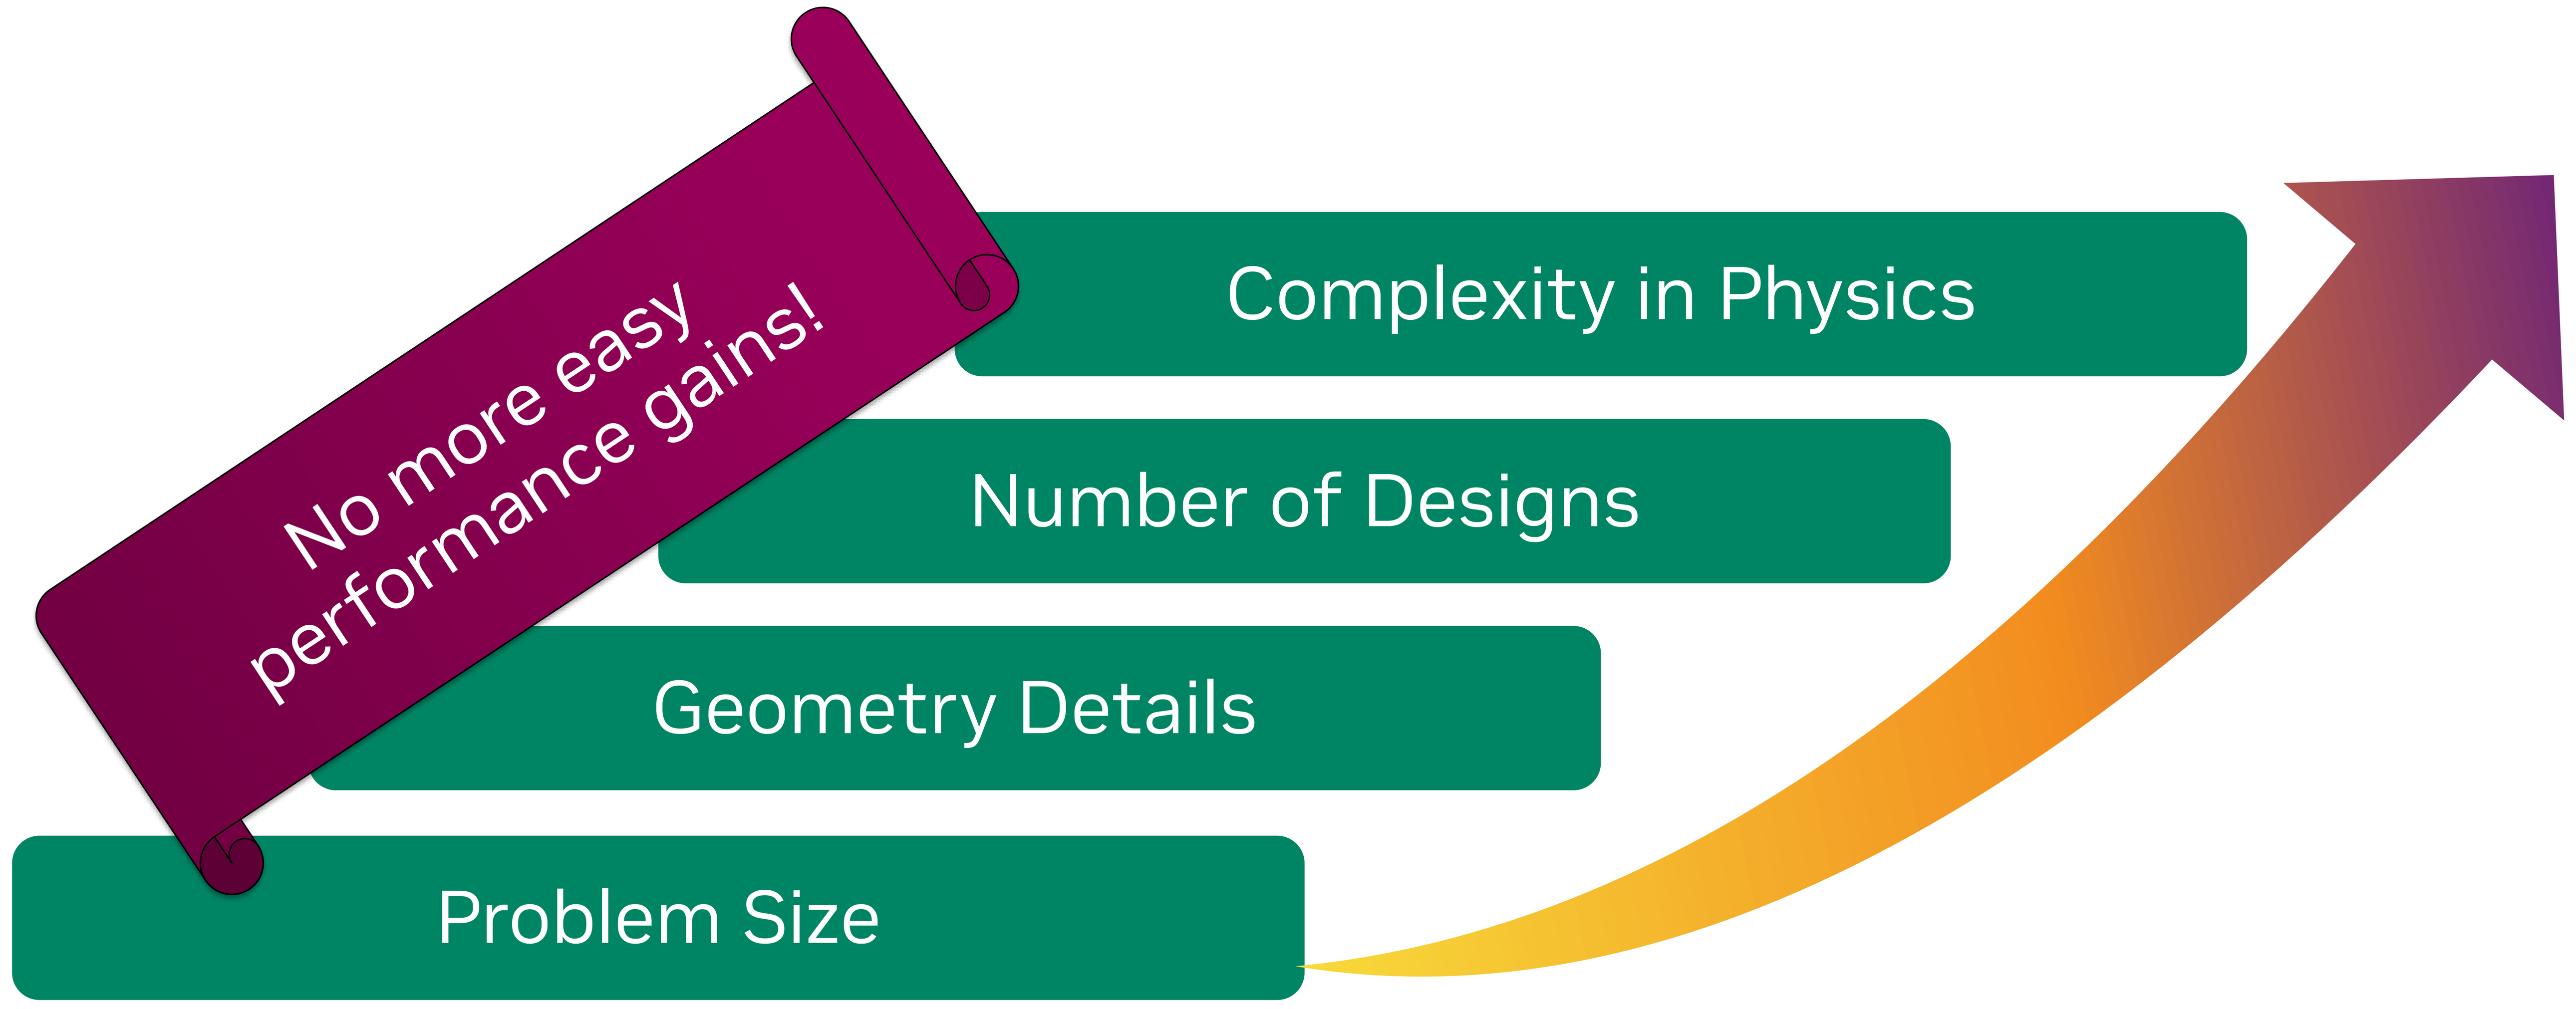

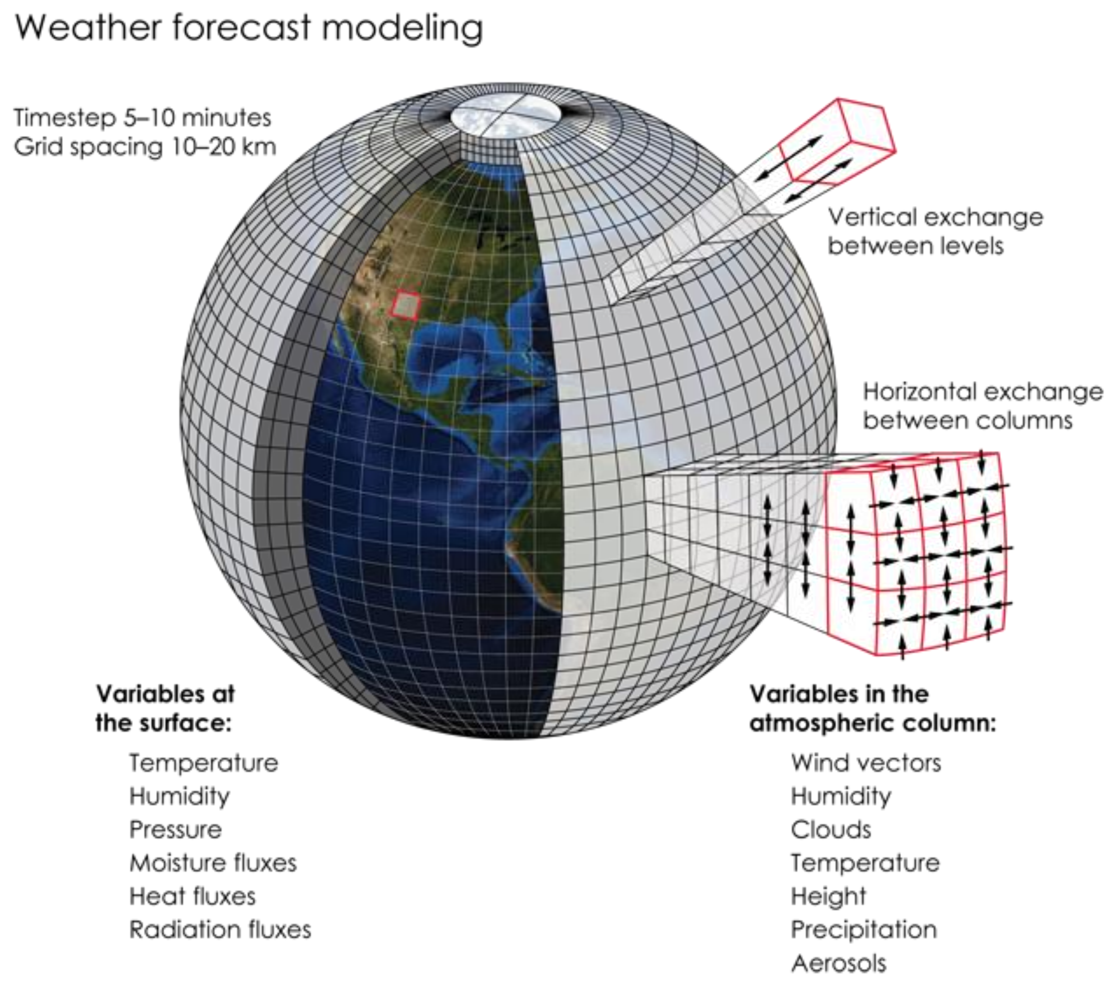

Atmospheric flows

Finite Volume Methods, Sub-grid scale modeling

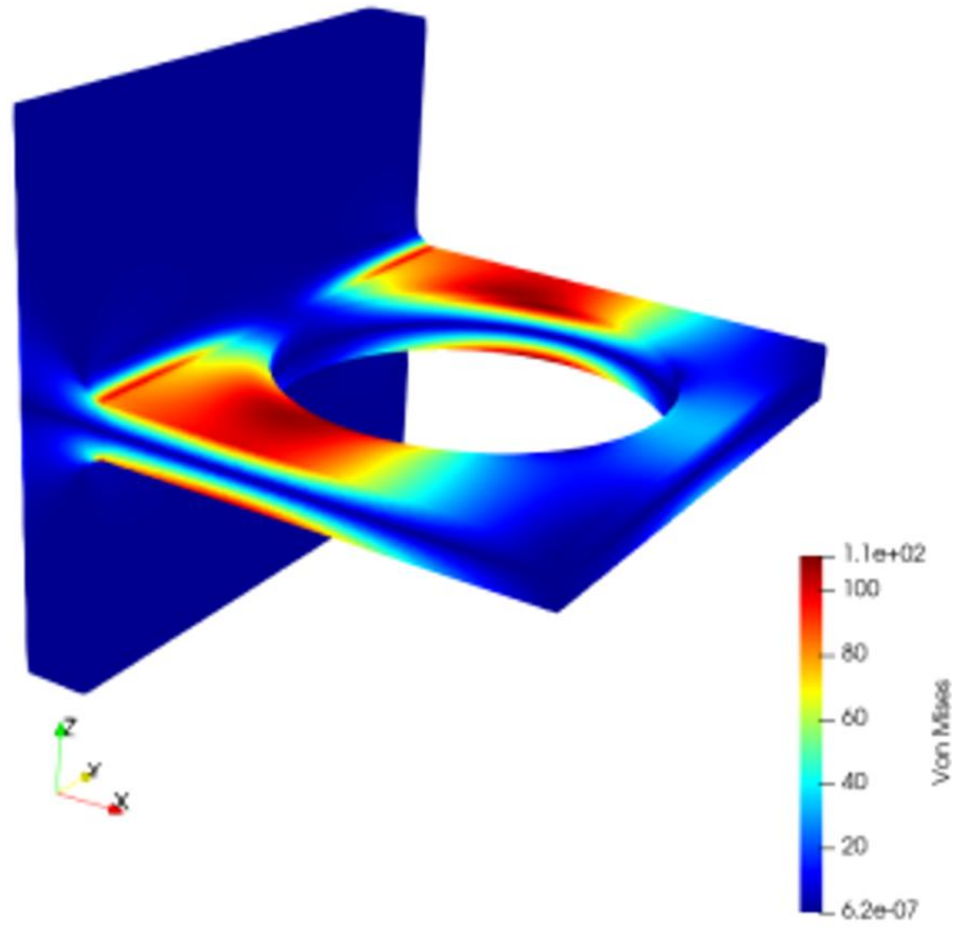

Structural Mechanics

Finite Element Methods

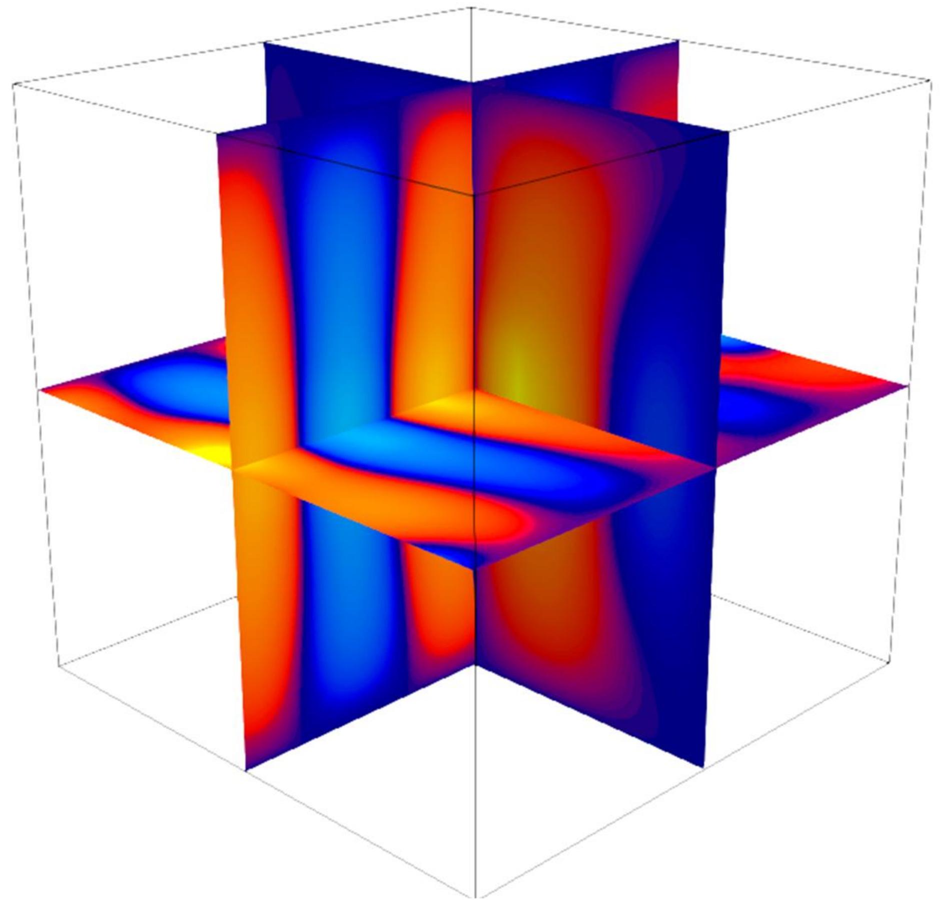

Electromagnetics

Finite Element and Frequency domain methods

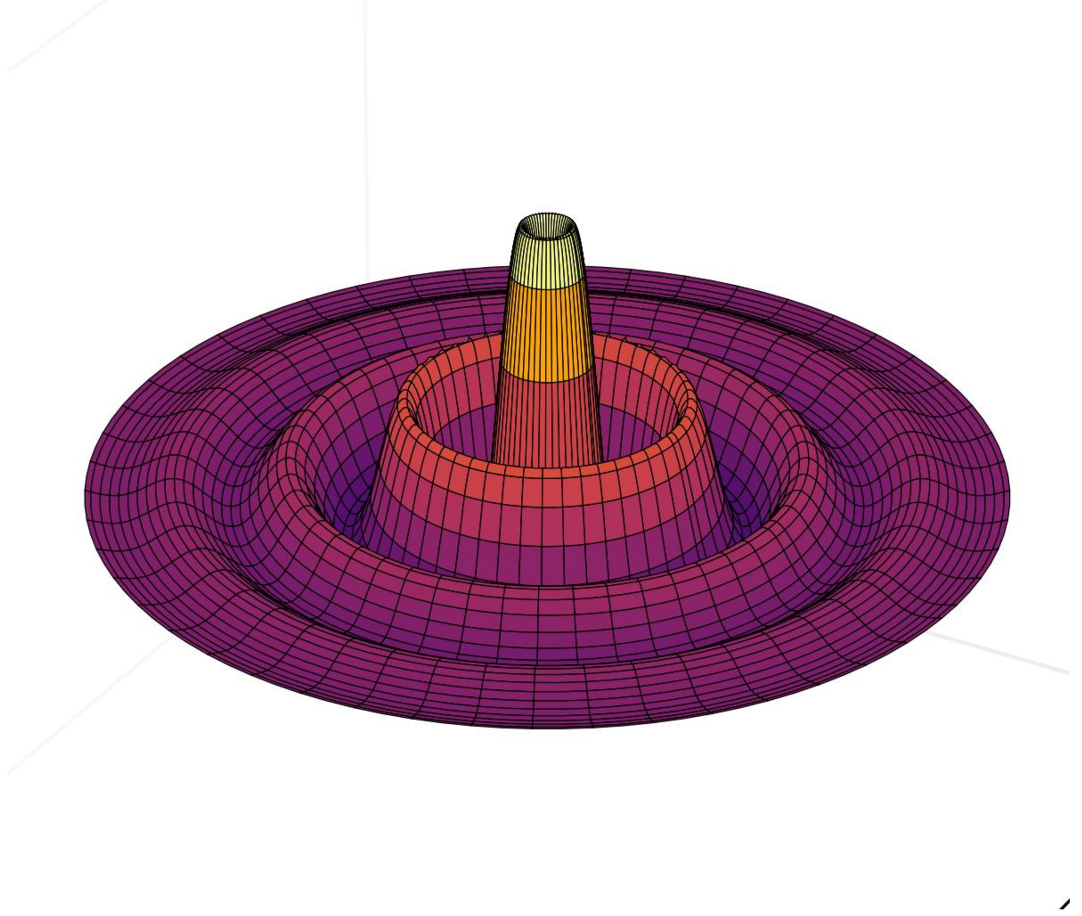

Vibrations / Acoustics

Finite Difference Methods

Universal Function Approximation

Use a NN as an approximation function

- When we're solving a PDE, we're looking for a function
- But multi-layer perceptrons can approximate arbitrary functions, given "enough" (surprisingly few) nodes, under various constraints given the details of the function, activations, *etc*.
- If this is new to you, looking into Universal Approximation Theorem(s) is a good starting point for the mathematics minded, or for the computational physics crowd (like me), I like a post on Jane Street's blog using ReLU activations to model arbitrary piecewise-linear functions
- Other kinds of networks can do this too! e.g. Kolmo.gorov-Arnold Networks This is part of the broad research efforts around PINN-type methods

Raissi, Perdikaris, & Karniadakis (2019)

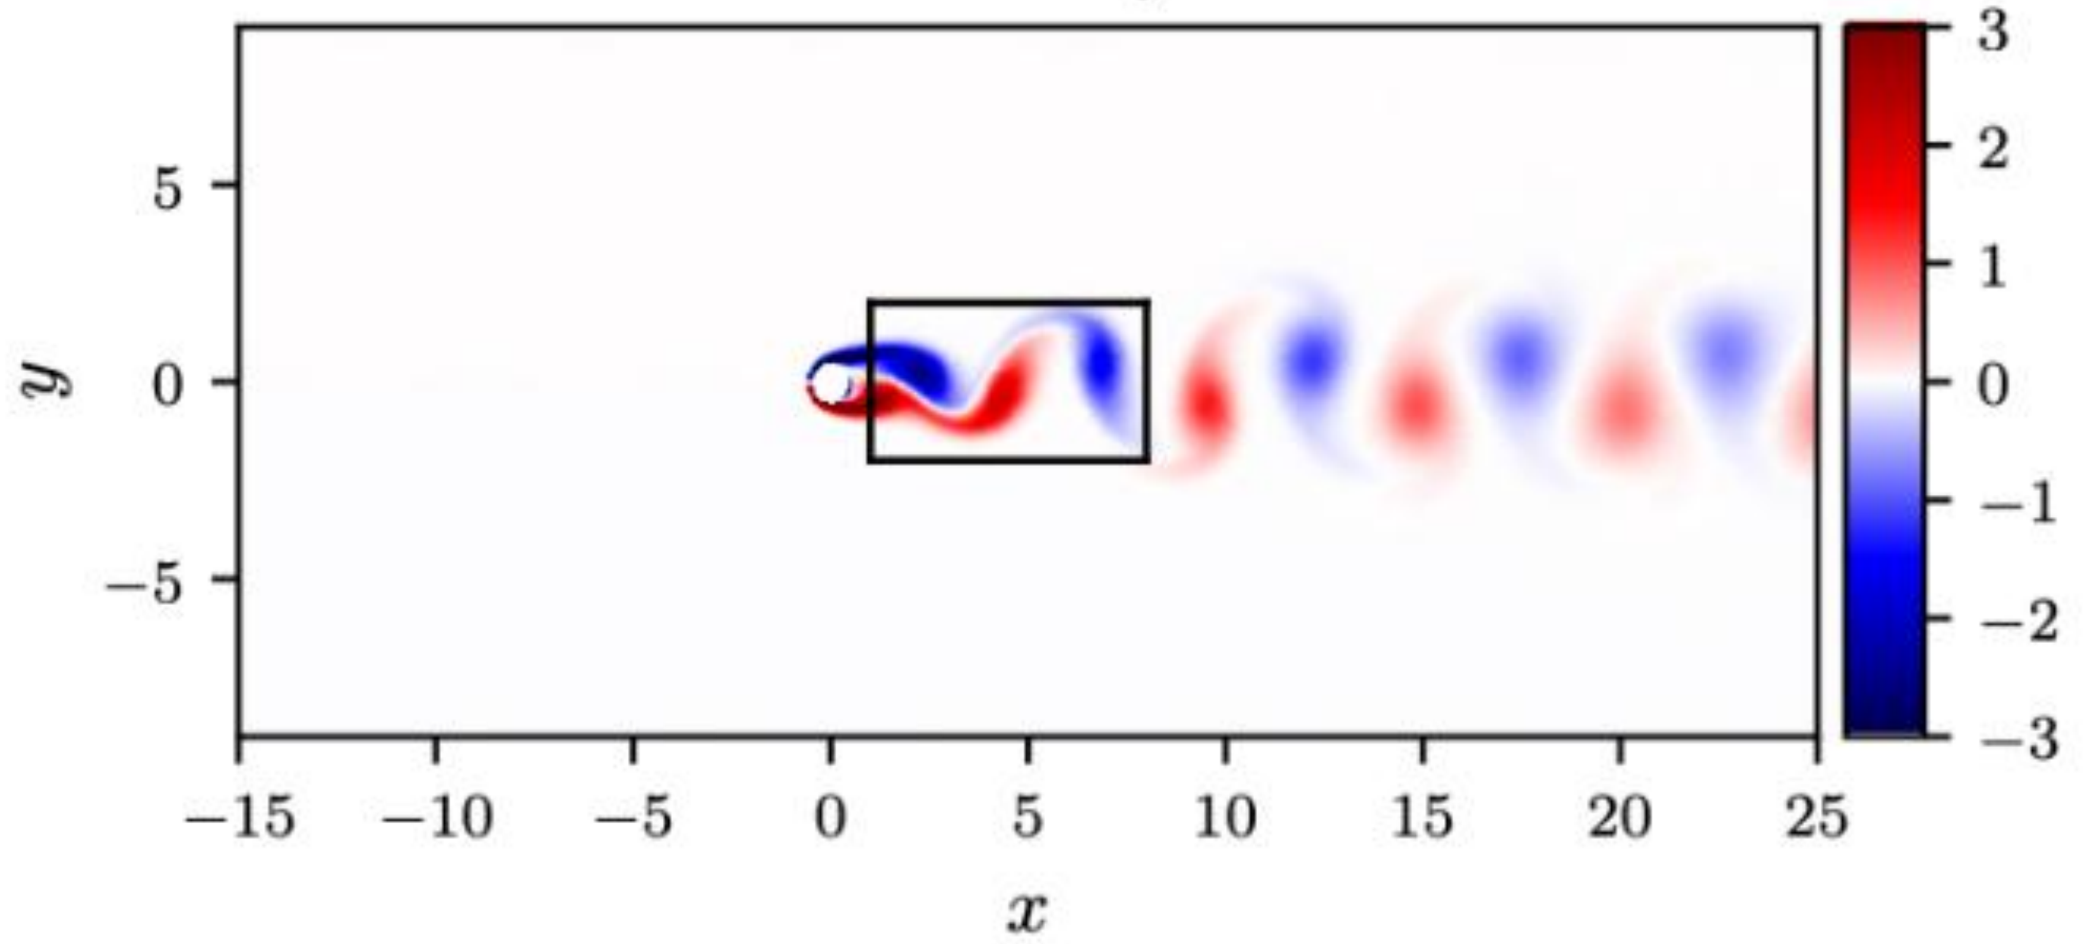

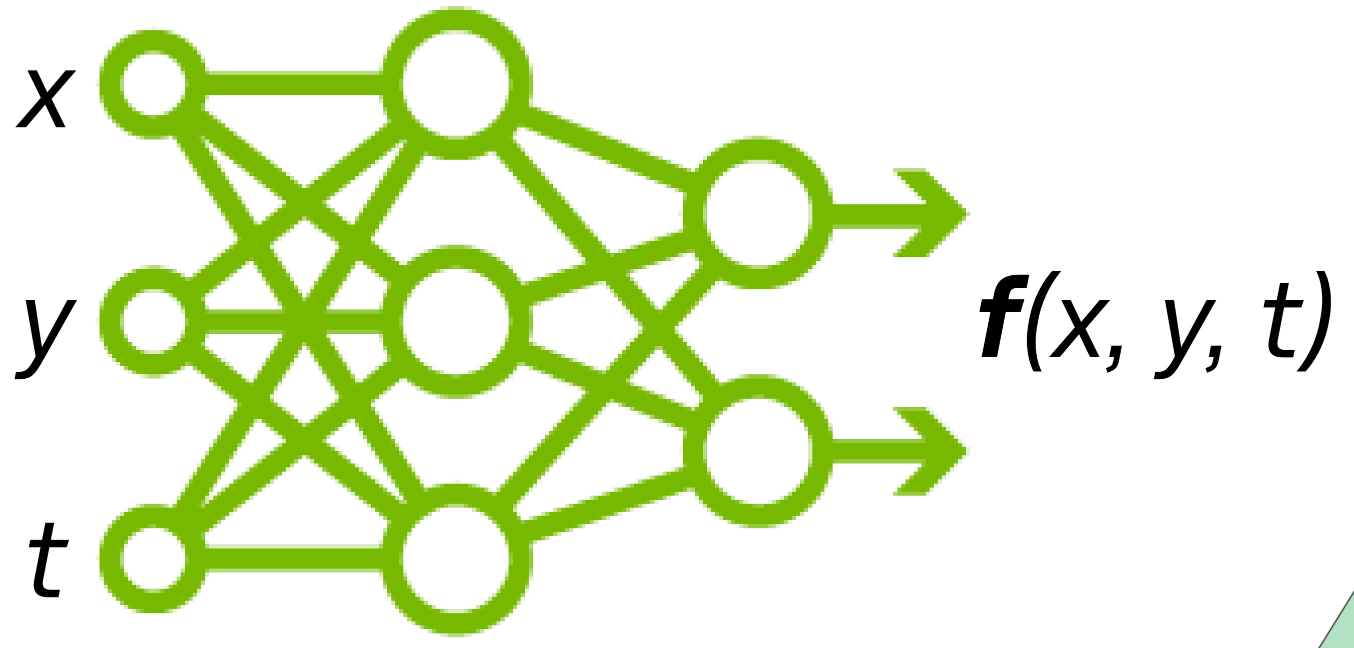

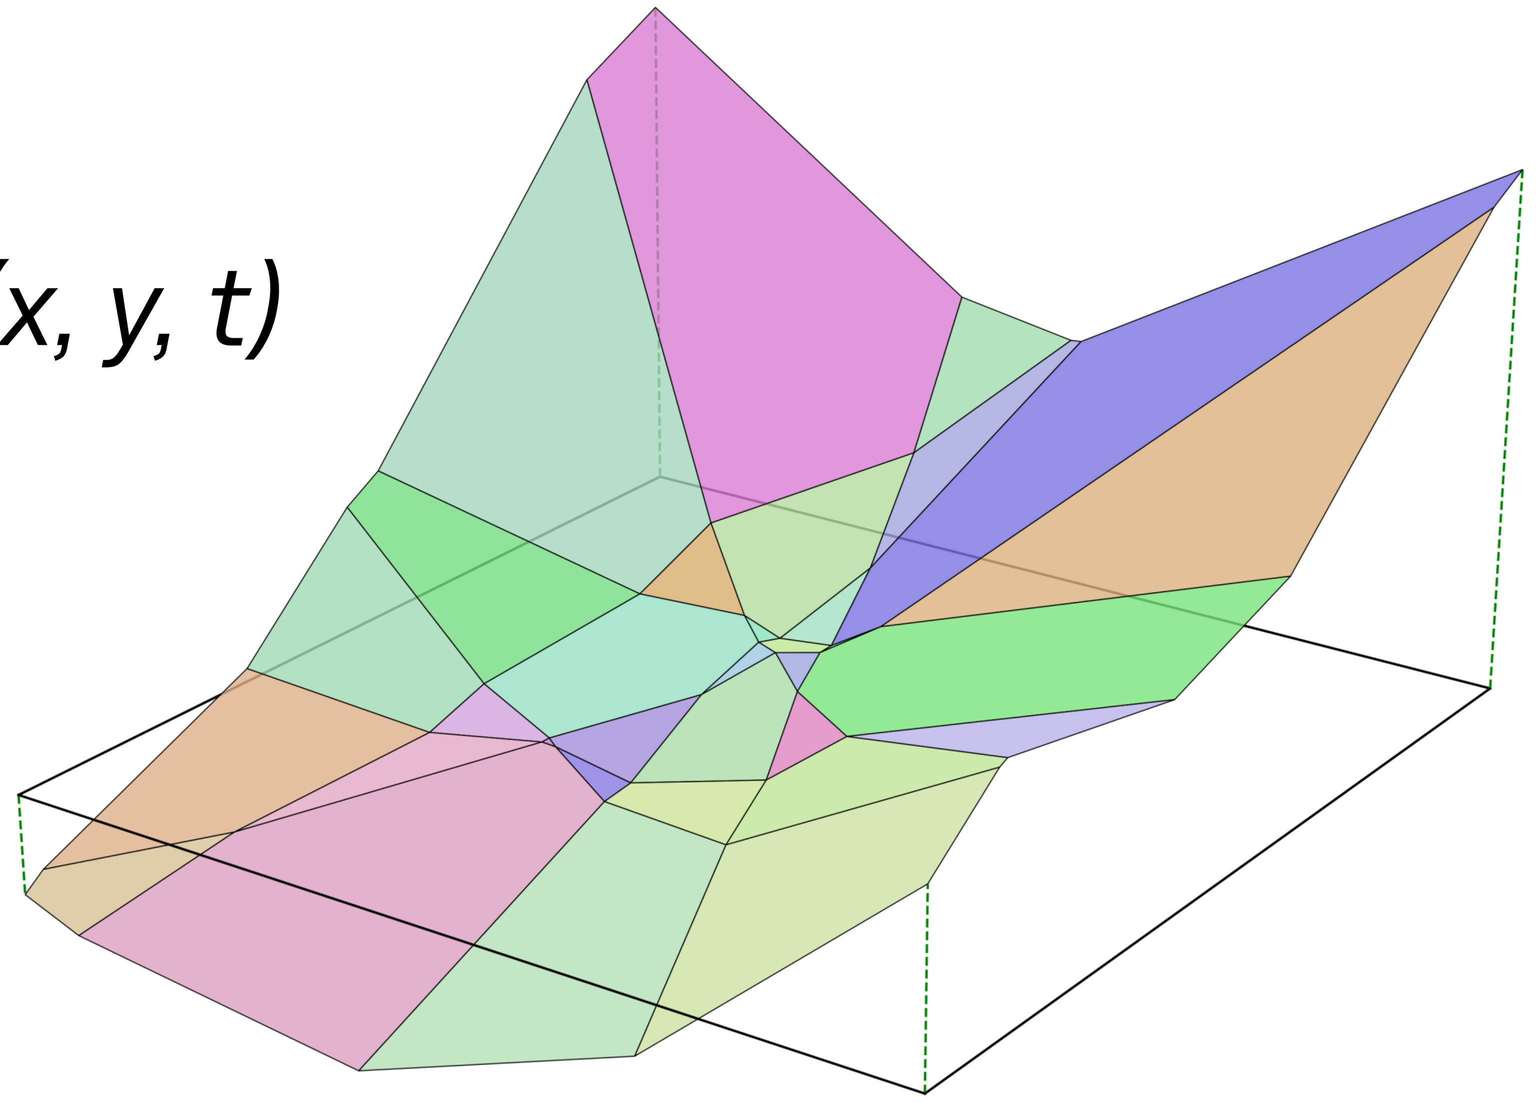

https://blog.janestreet.com/visualizing-piecewise-linear-neural-networks/

A Machine Learning Accelerator for Science

Naively using Deep Learning for scientific problems has its own challenges:

- Lack of interpretability
- Poor generalization – accurately models the process within vicinity of training data, but not away from this data.
- Lack of training data

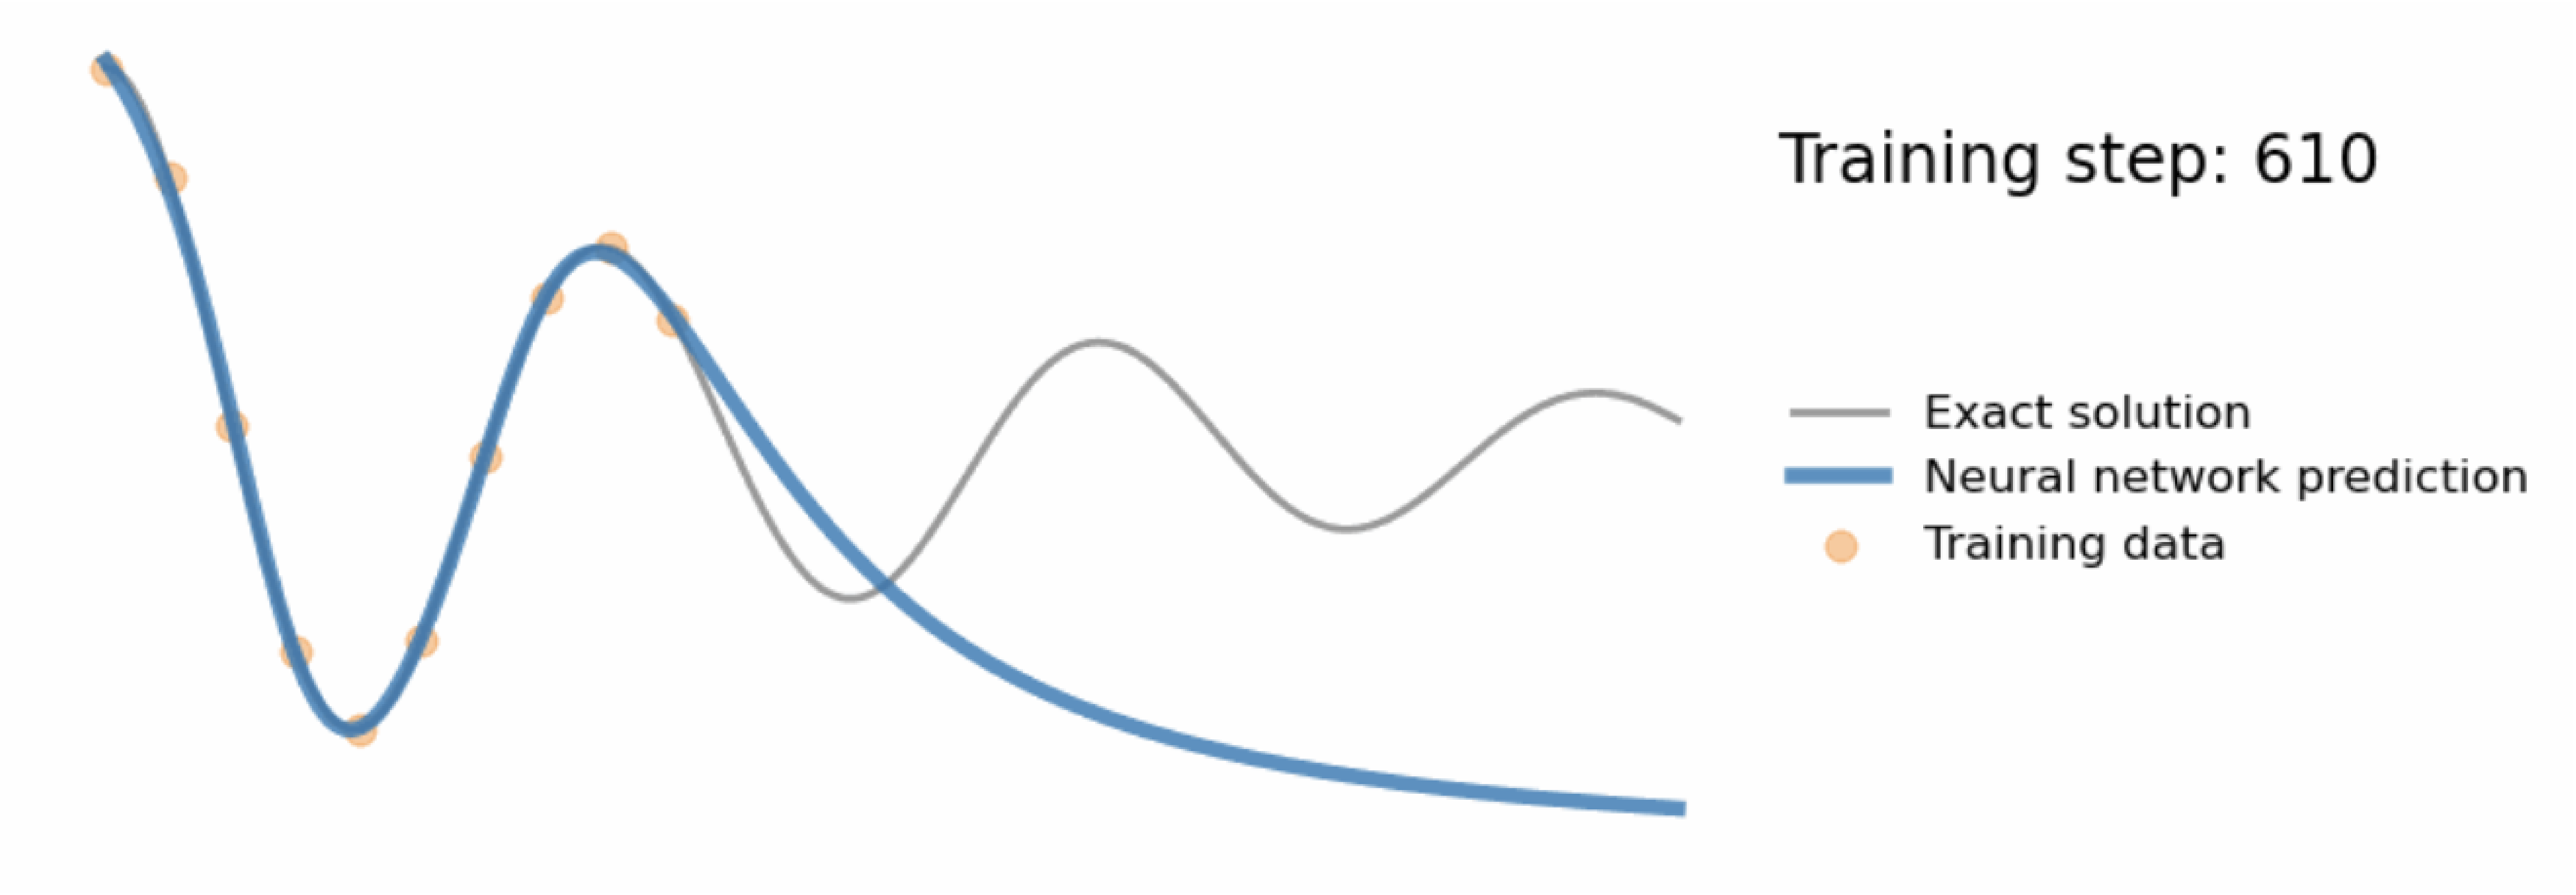

Source: https://benmoseley.blog/my-research/so-what-is-a-physics-informed-neural-network/
ETH Zürich 2024 AI in the Sciences and Engineering Master's course

Scientific Machine Learning

- What if we can somehow add the underlying physics to our ML workflows?
- This way we can take advantage of both the power of Neural Networks + our prior scientific knowledge.
- Advantages:
  - More accurate predictions even if we have sparse or no data, as we can leverage the physical laws.
  - Better generalization as the workflow is not confined to the training data only.
  - Useful in scenarios where data collection is challenging or expensive.

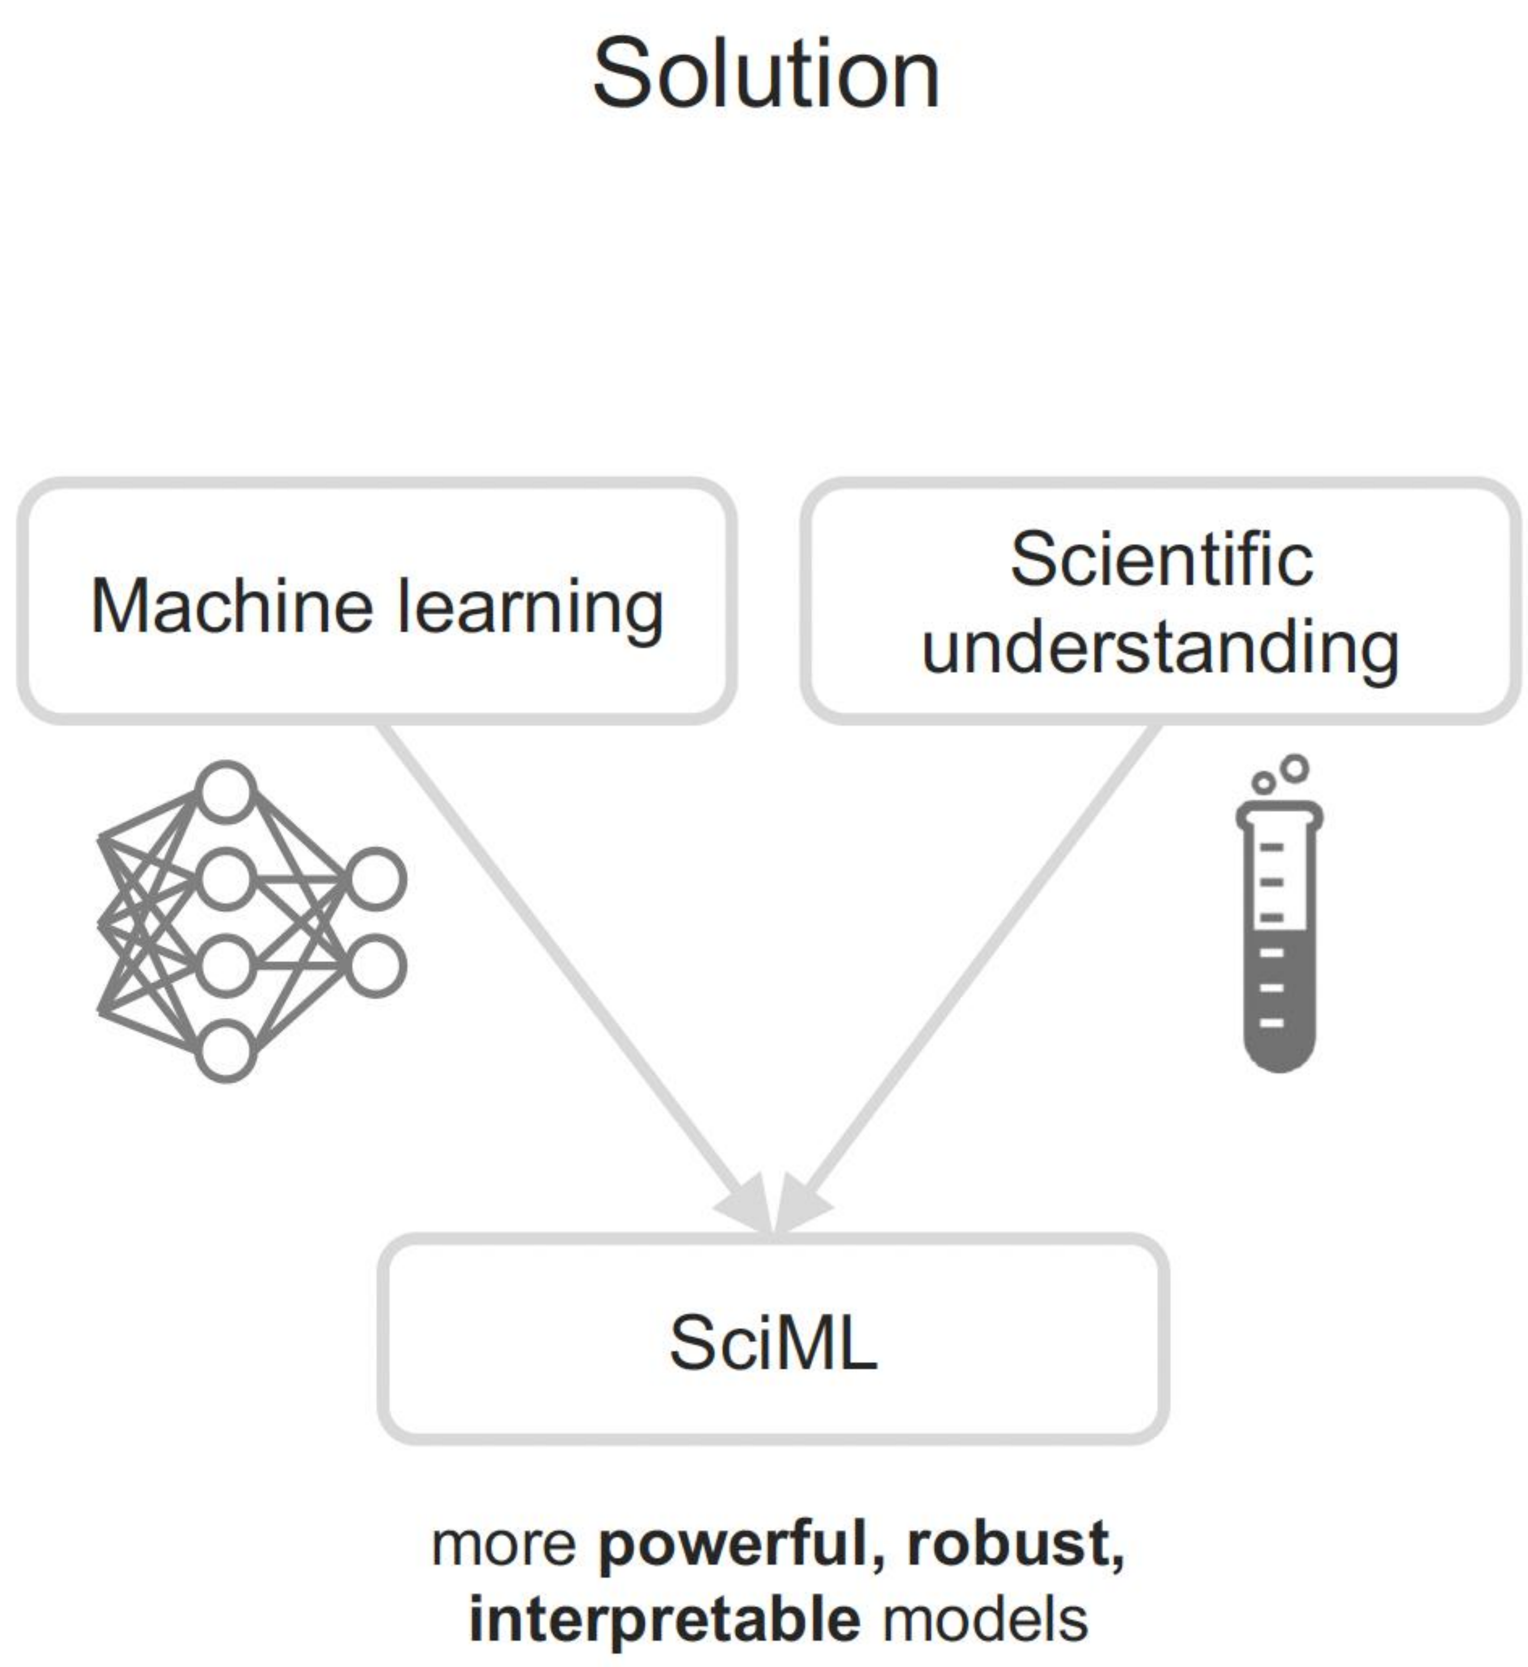

Source: https://benmoseley.blog/my-research/so-what-is-a-physics-informed-neural-network/
ETH Zürich 2024 AI in the Sciences and Engineering Master's course
https://www.engineered-mind.com/engineering/physics-informed-neural-networks-pinns/

How Would Learning An Operator Even Work?

Handling mappings between infinite-dimensional objects

- It turns out there are universal approximation to nonlinear operator theorems, too (*e.g.* Chen & Chen (1995) IEEE Tran. Neur. Net.)
- How do you handle an infinite dimensional input?
- Very traditional numerical approach – we don't
  - Project into some reduced-dimensionality space in some way
  - Work in that reduced space
  - Extrapolate back into the full space
- In contrast to PINNs, which needed no real training data, operator learning requires lots of training samples of input and output functions
  - Experimental data
  - Surrogate data such as simulations
- Training is very expensive, but can then inference very quickly with different ICs (and maybe BCs)

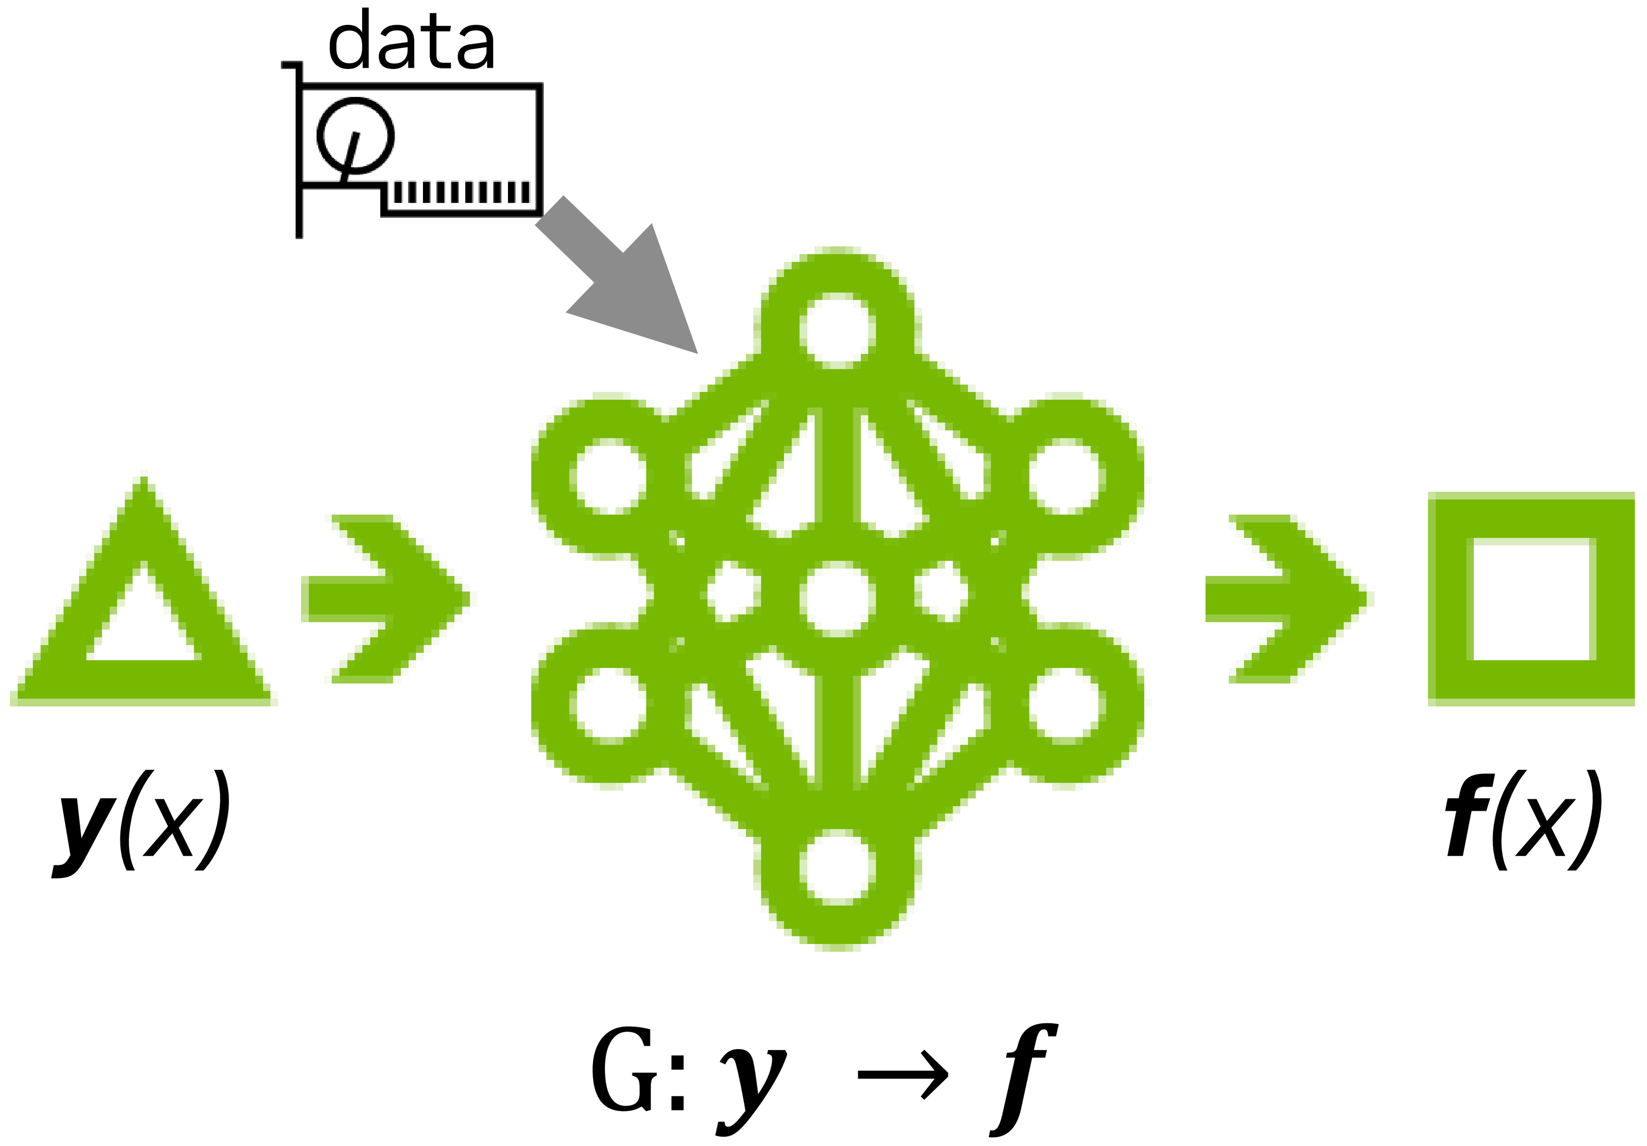

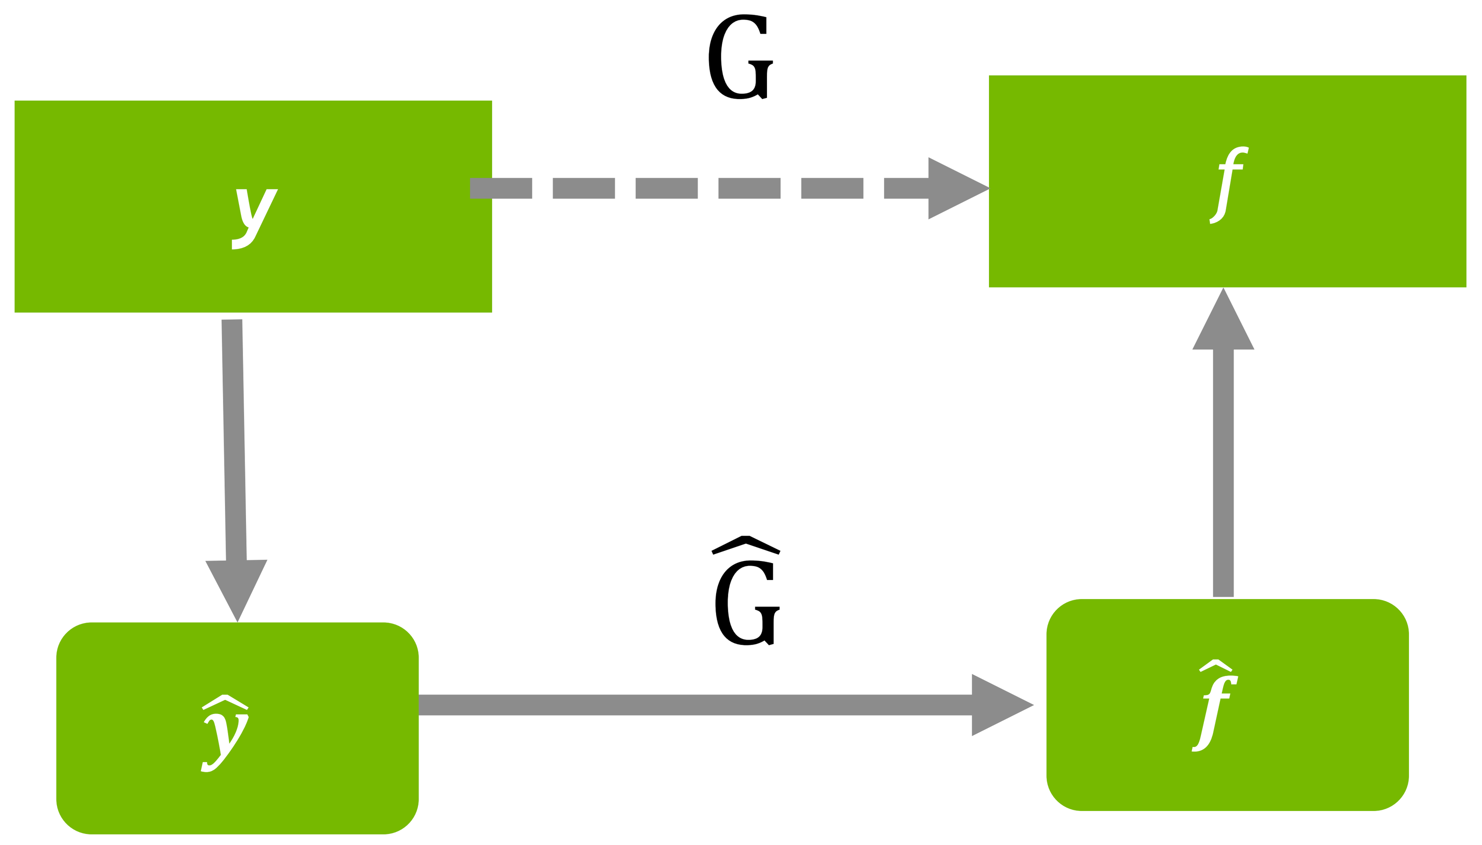

What is PhysicsNeMo?

Open-Source Framework for developing Physics AI algorithms

- Enterprise grade framework/toolkit for building physics AI models
  - Built-in physics AI training pipelines for CAE
    - Model architectures tuned for Physics AI
    - GPU accelerated data pipelines for engineering data formats and data structures
    - Utilities for injecting Physics – PDEs, BCs, Geometry constraints
  - 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

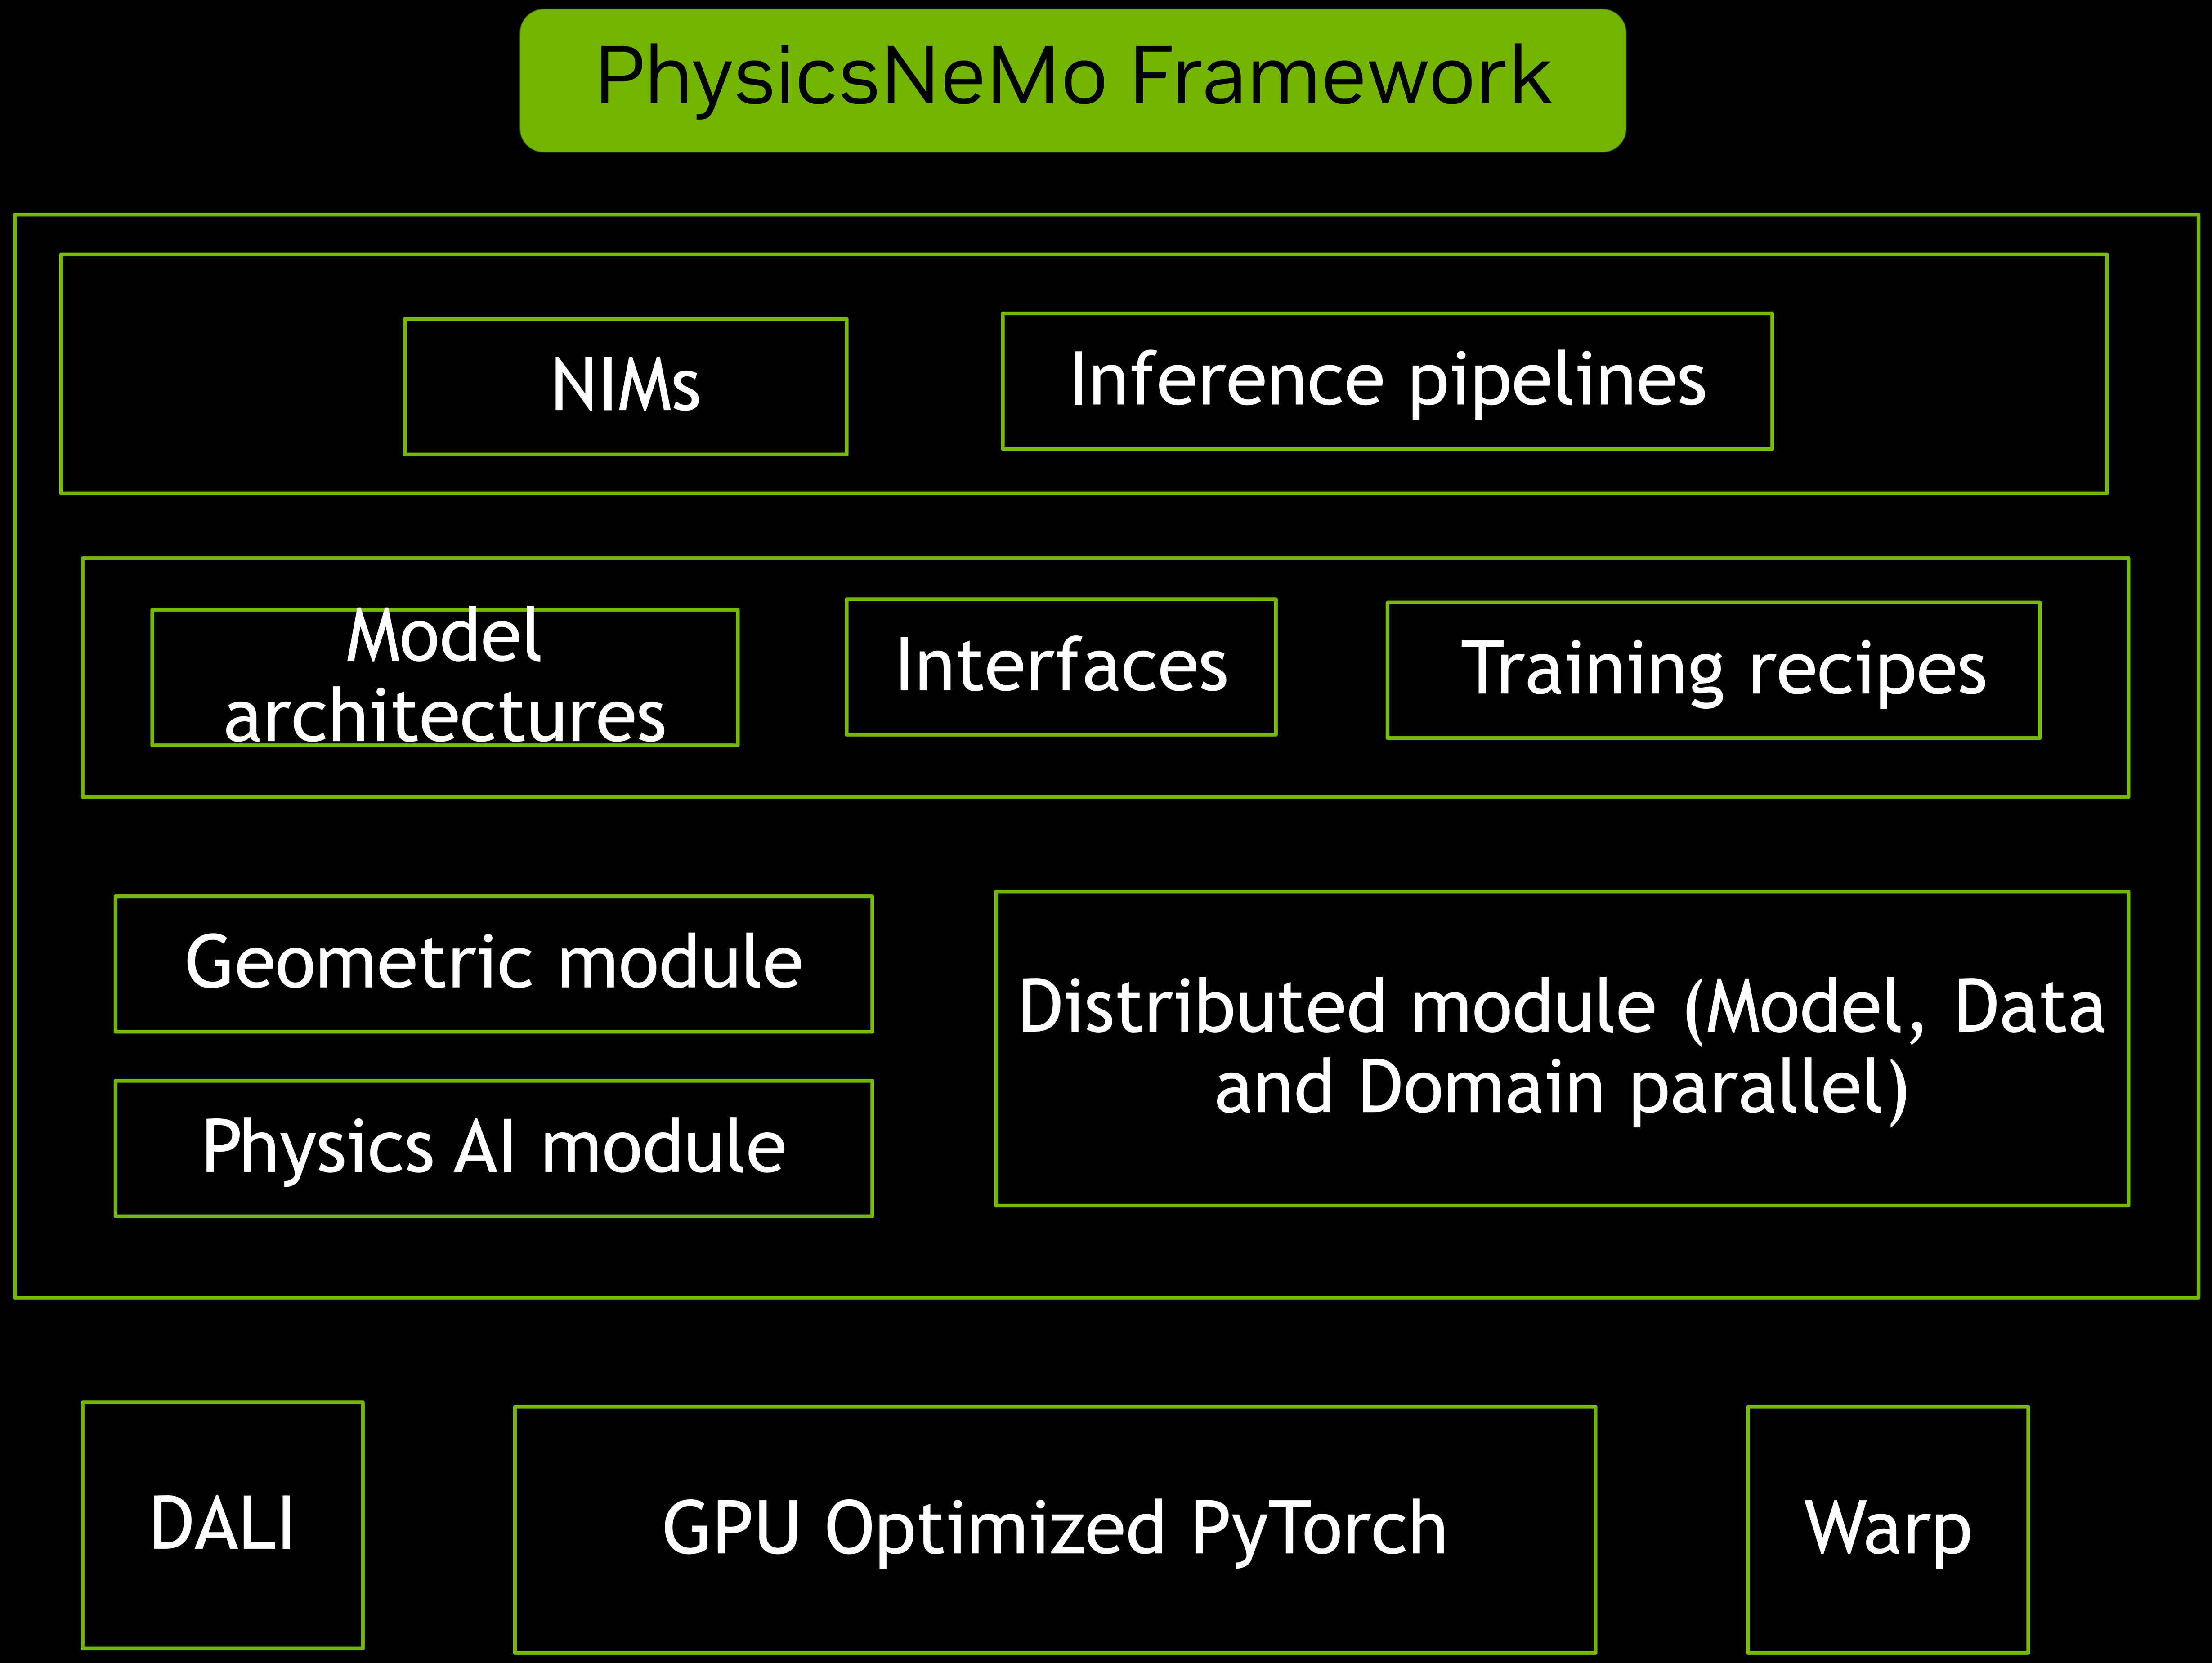

https://github.com/NVIDIA/PhysicsNeMo

What can Physics AI enable?

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

Built On Top Of PyTorch

Features that can aid data and/or physics driven problems

Data & Physics oriented utils

- Performance Enhancements
CUDA graphs, kernel fusion, JIT compilation, data parallelism, model parallel, etc.

- Pre-built Network Architectures
Diffusion Models, Neural Hash Encoding, Neural Operators, Graph Networks, DCT-RNN, several variants of MLPs etc.

- Hydra Configuration
Hyper-parameter tuning and customization

- Data Pipeline
For very large data-driven problems using Zarr, NVComp, GDS

- Data and Inference Tools
Pre-defined datasets for common data formats (VTK, HDF5, …). Model export functions to TensorRT and Triton

- Integration with Other Products
Omniverse, PySDF, NVFuser, Triton, Tensor RT, DALI, Warp, etc.

Physics oriented utils

- Geometry Module
Integrated, parameterized geometry module with point cloud/SDF

- Symbolic PDE Loss Construction
Automated PDE loss construction using SymPy API

- Automated Optimized Gradient Calculations
Automatic gradient calculations for physics-informed learning with optimizations such as FuncTorch, AMP16, mesh free derivatives etc.

- Convergence and Stabilization Methods
Mass balance control planes, loss balancing schemes, AdaHessian support, learning rate annealing, etc.

- Exact Boundary Enforcement
Exact enforcement of continuity or boundary conditions

- Variational Learning
Solving PDE systems using variational formulations

PhysicsNeMo Resources

PhysicsNeMo Product page

Get an Overview – webinar

Download PhysicsNeMo Framework

PhysicsNeMo Documentation

Getting started guide

Reference samples 1

Reference samples 2

Reference blogs

Self paced DLI Course

CFD

Use case	Model
Vortex Shedding	MeshGraphNet
Ahmed Body Drag prediction	MeshGraphNet
Navier-Stokes Flow	RNN
Gray-Scott System	RNN
Darcy Flow	FNO
Darcy Flow using Nested-FNOs	Nested-FNO
Darcy Flow (Data + Physics Driven) using DeepONet approach	FNO (branch) and MLP (trunk)
Darcy Flow (Data + Physics Driven) using PINO approach (Numerical gradients)	FNO
Stokes Flow (Physics Informed Fine-Tuning)	MeshGraphNet and MLP

Weather

Use case	Model	AMP	CUDA Graphs	Multi-GF
Medium-range global weather forecast using FCN-SFNO	FCN-SFNO	YES	NO	YES
Medium-range global weather forecast using GraphCast	GraphCast	YES	NO	YES
Medium-range global weather forecast using FCN-AFNO	FCN-AFNO	YES	YES	YES
Medium-range and S2S global weather forecast using DLWP	DLWP	YES	YES	YES

Healthcare

Use case	Model	Transient
Cardiovascular Simulations	MeshGraphNet	YES
Brain Anomaly Detection	FNO	YES

Molecular Dymanics

Use case	Model	Transient
Force Prediciton for Lennard Jones system	MeshGraphNet	NO

Generative

Use case	Model	Multi-GPU	Multi-
Generative Correction Diffusion Model for Km-scale Atmospheric Downscaling	CorrDiff	YES	YES

Additional examples

In addition to the examples in this repo, more Physics-ML usecases and examples can be referenced from the Modulu

Introductory

Use case	Model
Lid Driven Cavity Flow	Fully Connected MLP PINN
Anti-derivative	Data and Physics informed DeepO
Darcy Flow	FNO, AFNO, PINO
Spring-mass system ODE	Fully Connected MLP PINN
Surface PDE	Fully Connected MLP PINN

Turbulence

Use case	Model
Taylor-Green	Fully Connected MLP PINN
Turbulent channel	Fourier Feature MLP PINN
Turbulent super-resolution	Super Resolution Network, Pix2P

Electromagnetics

Use case	Model	Level	
Waveguide	Fourier Feature MLP PINN	Intermediate	St

Solid Mechanics

Use case	Model	
Plane displacement	Fully Connected MLP PINN, VPINN	In

Design Optimization

Use case	Mo
2D Chip	Fully Connected MLP PINN
3D Three fin Heatsink	Fully Connected MLP PINN
FPGA Heatsink	Multiple Models (including Four etc.)
Limerock Industrial Heatsink	Fourier Feature MLP PINN

Geophyscis

Use case	Model	Leve
Reservoir simulation	FNO, PINO	Advance
Seismic wave	Fully Connected MLP PINN	Intermec
Wave equation	Fully Connected MLP PINN	Intermec

Healthcare

Reference sample catalogue

https://github.com/NVIDIA/PhysicsNeMo/tree/main/examples

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

CFD

Use case	Model	Transient
Vortex Shedding	MeshGraphNet	YES
Ahmed Body Drag prediction	MeshGraphNet	NO
Navier-Stokes Flow	RNN	YES
Gray-Scott System	RNN	YES
Darcy Flow	FNO	NO
Darcy Flow using Nested-FNOs	Nested-FNO	NO
Darcy Flow (Data + Physics Driven) using DeepONet approach	FNO (branch) and MLP (trunk)	NO
Darcy Flow (Data + Physics Driven) using PINO approach (Numerical gradients)	FNO	NO
Stokes Flow (Physics Informed Fine-Tuning)	MeshGraphNet and MLP	NO

Weather

Use case	Model	AMP	CUDA Graphs	Multi-GPU	Multi-Node
Medium-range global weather forecast using FCN-SFNO	FCN-SFNO	YES	NO	YES	YES
Medium-range global weather forecast using GraphCast	GraphCast	YES	NO	YES	YES
Medium-range global weather forecast using FCN-AFNO	FCN-AFNO	YES	YES	YES	YES
Medium-range and S2S global weather forecast using DLWP	DLWP	YES	YES	YES	YES

Healthcare

Use case	Model	Transient
Cardiovascular Simulations	MeshGraphNet	YES
Brain Anomaly Detection	FNO	YES

Molecular Dymanics

Use case	Model	Transient
Force Prediciton for Lennard Jones system	MeshGraphNet	NO

Generative

Use case	Model	Multi-GPU	Multi-Node
Generative Correction Diffusion Model for Km-scale Atmospheric Downscaling	CorrDiff	YES	YES

Additional examples

In addition to the examples in this repo, more Physics-ML usecases and examples can be referenced from the Modulus-Sym examples.

Introductory

Use case	Model	Level	Attributes
Lid Driven Cavity Flow	Fully Connected MLP PINN	Introductory	Steady state, Multi-GPU
Anti-derivative	Data and Physics informed DeepONet	Introductory	Steady state, Multi-GPU
Darcy Flow	FNO, AFNO, PINO	Introductory	Steady state, Multi-GPU
Spring-mass system ODE	Fully Connected MLP PINN	Introductory	Steady state, Multi-GPU
Surface PDE	Fully Connected MLP PINN	Introductory	Steady state, Multi-GPU

Turbulence

Use case	Model	Level	Attributes
Taylor-Green	Fully Connected MLP PINN	Intermediate	Steady state, Multi-GPU
Turbulent channel	Fourier Feature MLP PINN	Intermediate	Steady state, Multi-GPU
Turbulent super-resolution	Super Resolution Network, Pix2Pix	Intermediate	Steady state, Multi-GPU

Electromagnetics

Use case	Model	Level	Attributes
Waveguide	Fourier Feature MLP PINN	Intermediate	Steady state, Multi-GPU

Solid Mechanics

Use case	Model	Level	Attributes
Plane displacement	Fully Connected MLP PINN, VPINN	Intermediate	Steady state, Multi-GPU

Design Optimization

Use case	Model	Level	Attributes
2D Chip	Fully Connected MLP PINN	Advanced	Steady state, Multi-GPU
3D Three fin Heatsink	Fully Connected MLP PINN	Advanced	Steady state, Multi-Node
FPGA Heatsink	Multiple Models (including Fourier Feature MLP PINN, SIRENS, etc.)	Advanced	Steady state, Multi-Node
Limerock Industrial Heatsink	Fourier Feature MLP PINN	Advanced	Steady state, Multi-Node

Geophyscis

Use case	Model	Level	Attributes
Reservoir simulation	FNO, PINO	Advanced	Steady state, Multi-Node
Seismic wave	Fully Connected MLP PINN	Intermediate	Steady state, Multi-Node
Wave equation	Fully Connected MLP PINN	Intermediate	Steady state, Multi-Node

Healthcare

Use case	Model	Level	Attributes
Aneurysm modeling using STL geometry	Fully Connected MLP PINN	Intermediate	Steady state, Multi-Node

Atmospheric Science

High-Resolution Climate Prediction is a Computational Challenge

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

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

NVIDIA Earth-2

Platform for Accelerated Weather and Climate Modelling

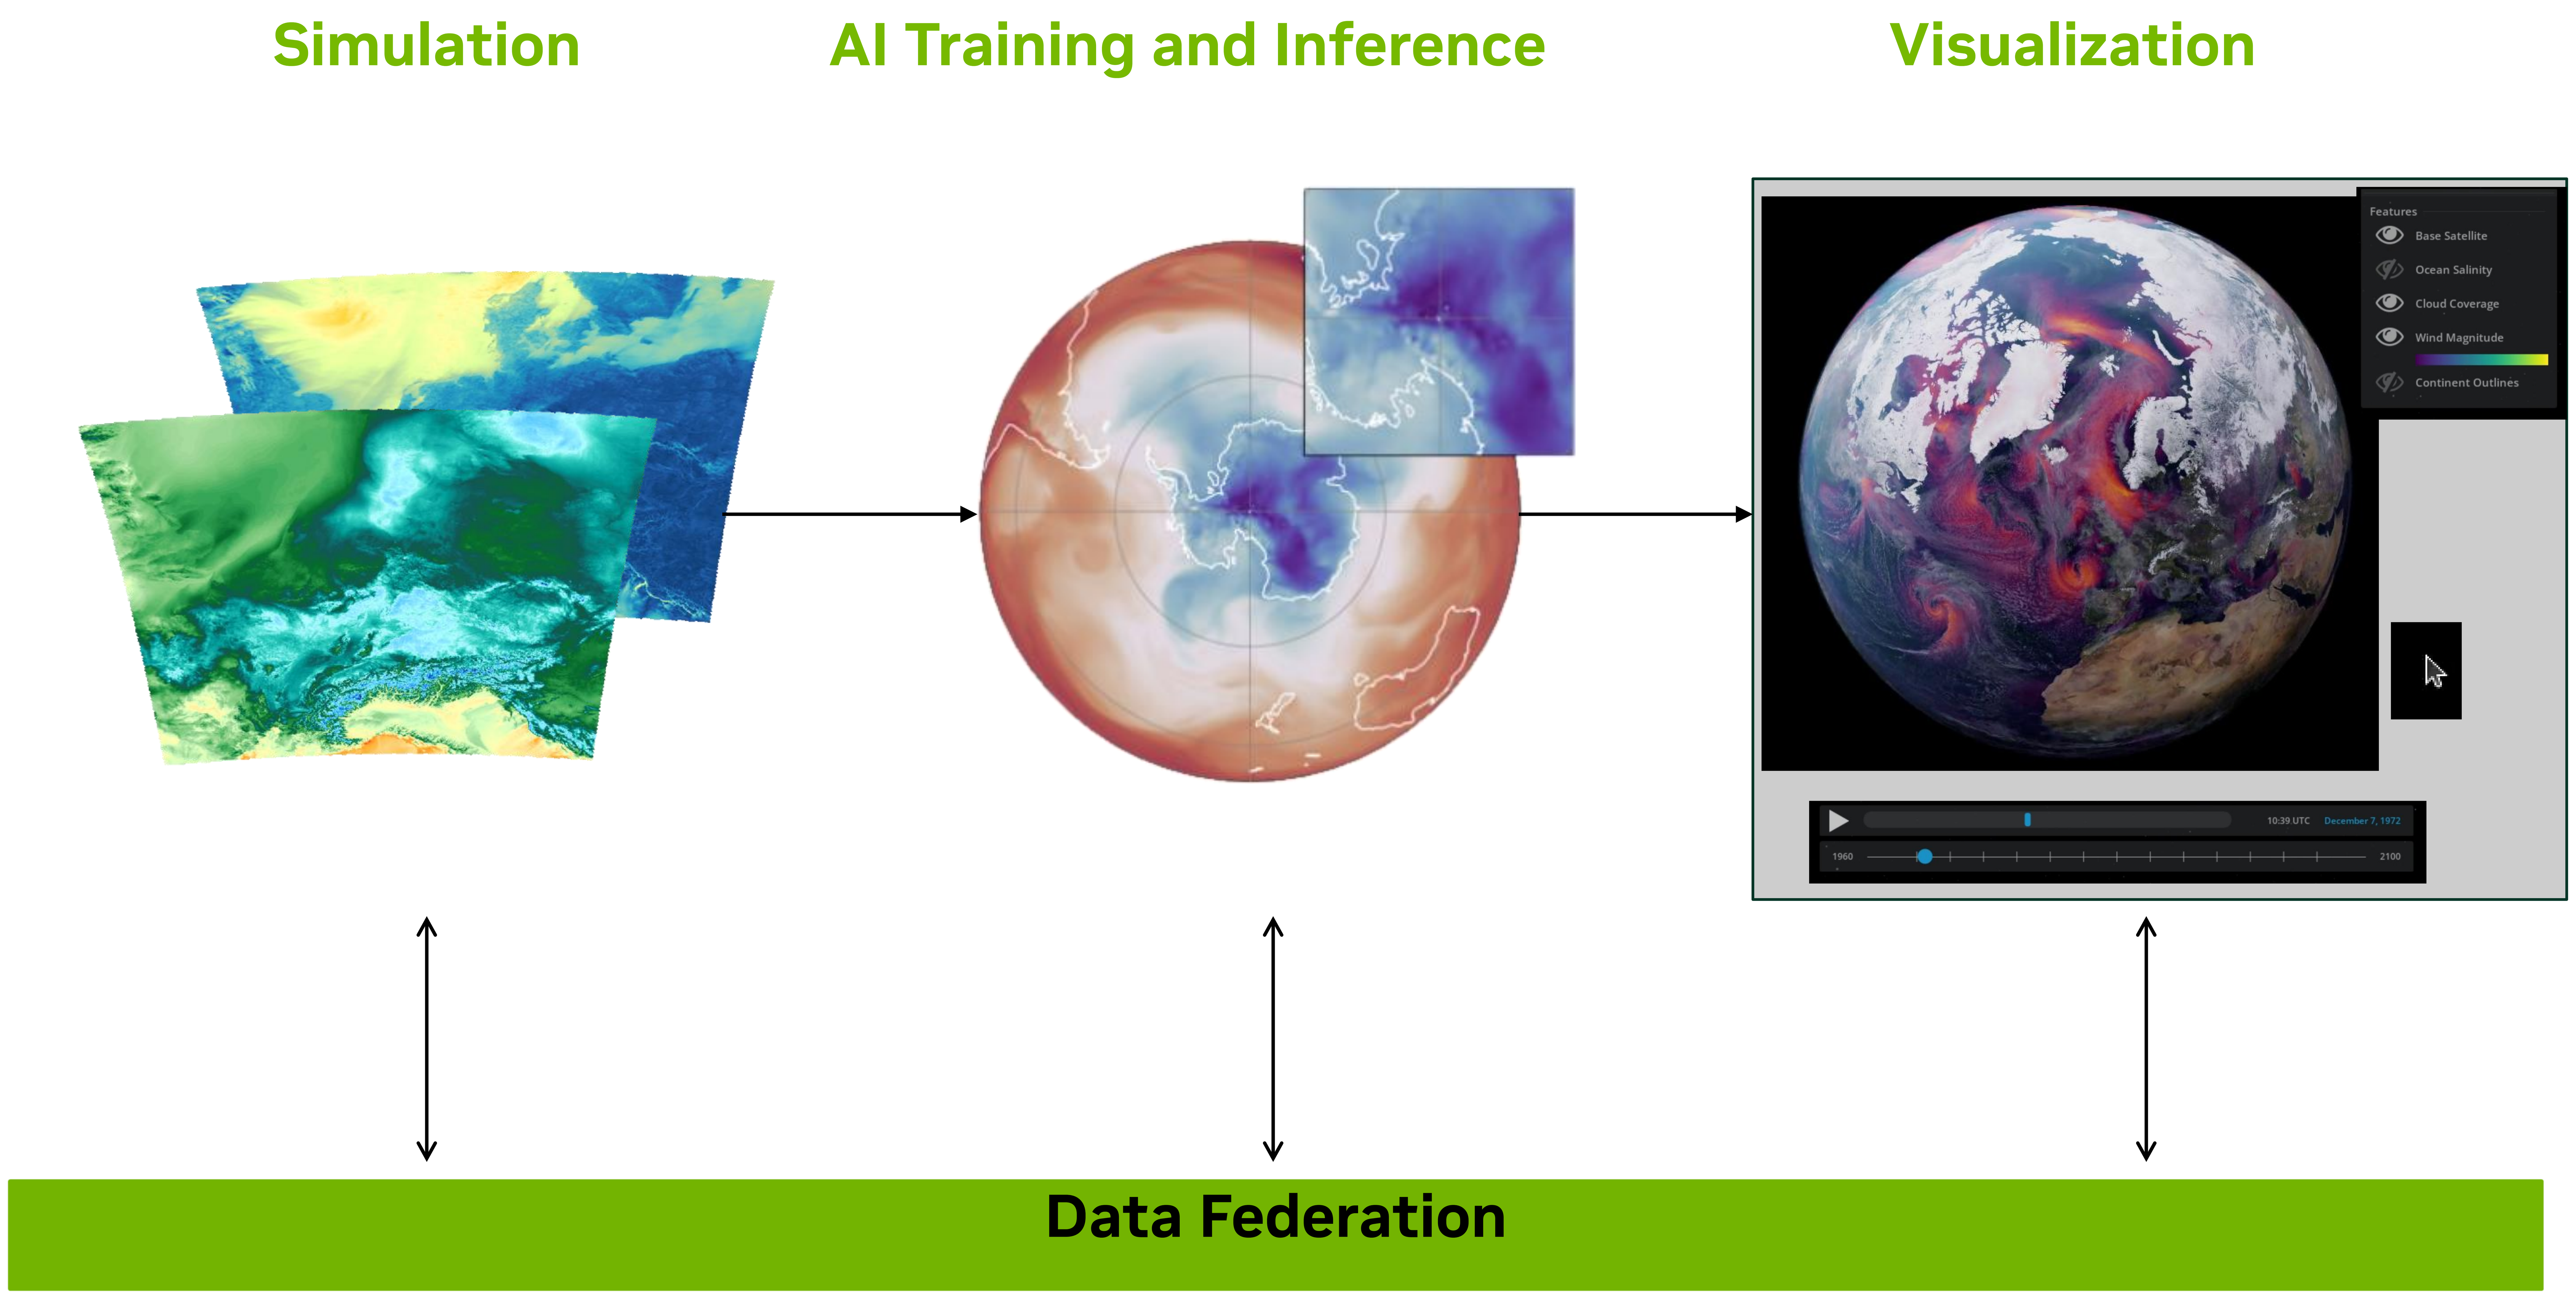

Enabling Accelerated Weather Intelligence for Accurate, Hyperlocal Forecasts

NVIDIA Earth-2 accelerates Tomorrow.io's computation from hours to minutes

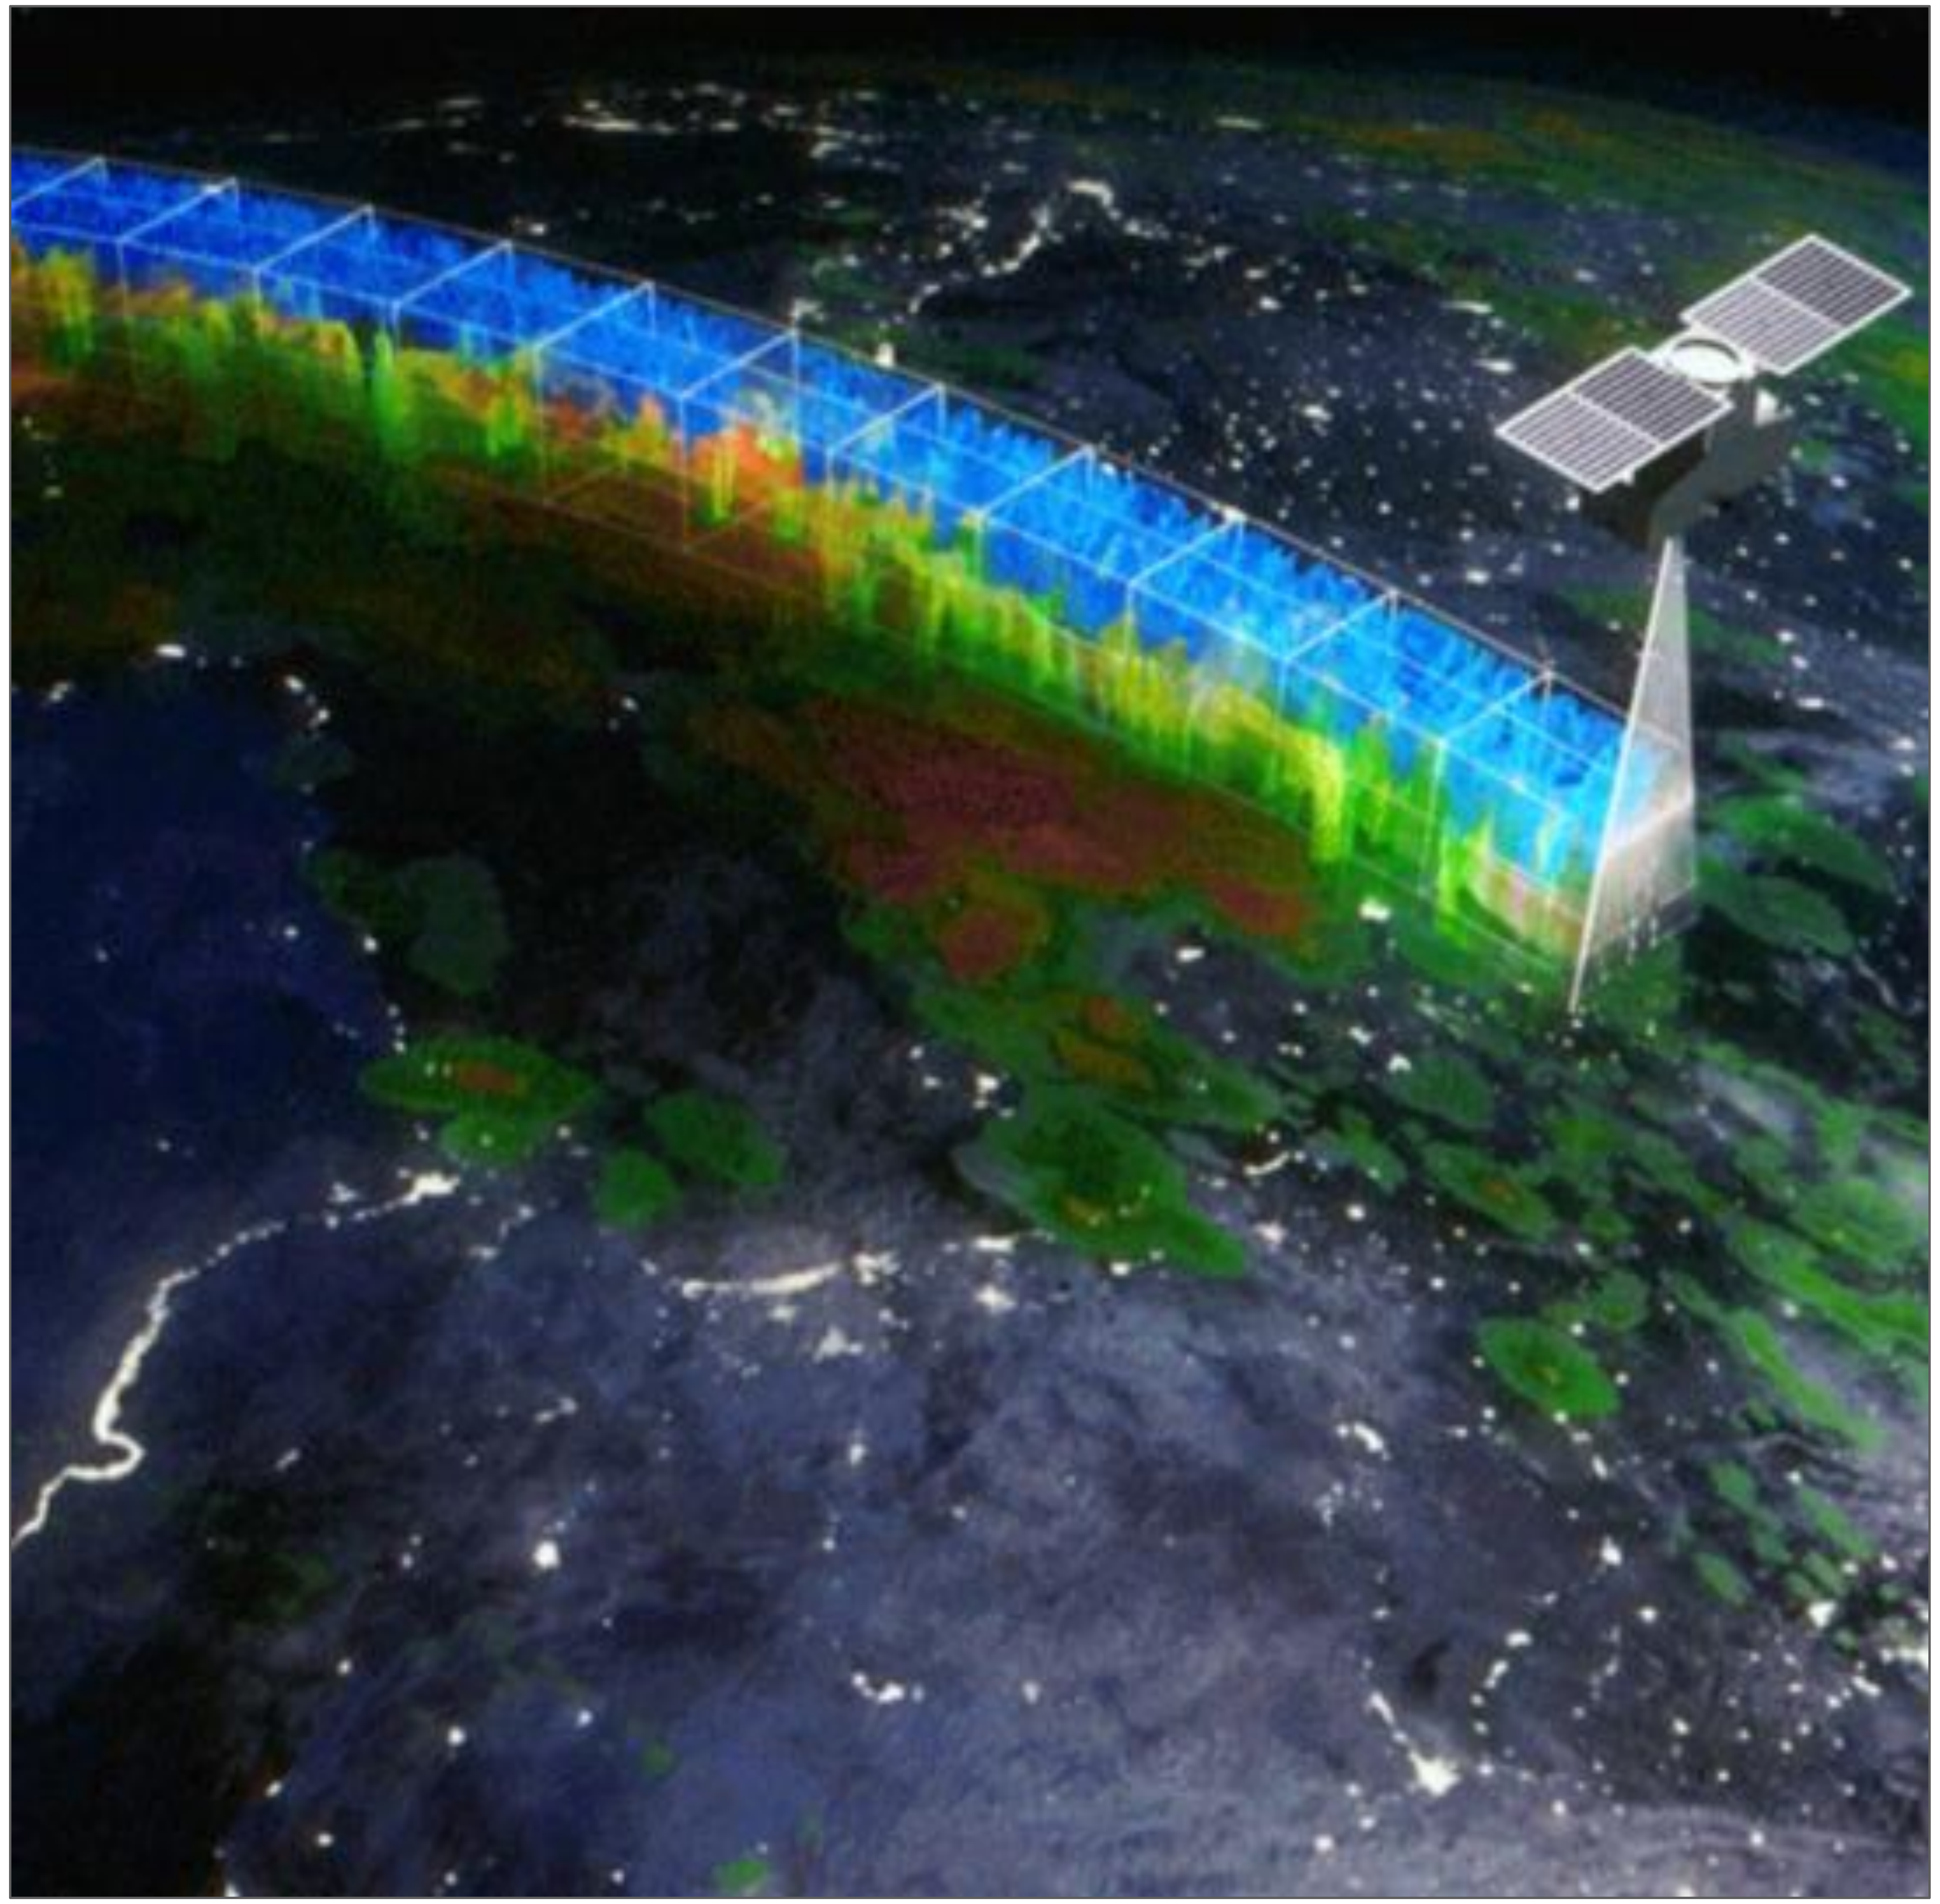

Satellite Imagery

Map-based Weather View in Tomorrow.io Platform

Collaborated with Taiwan CWA to Predict Super Typhoons

1,000x faster and 3,000x more energy efficient than numerical modeling, 13x finer resolution

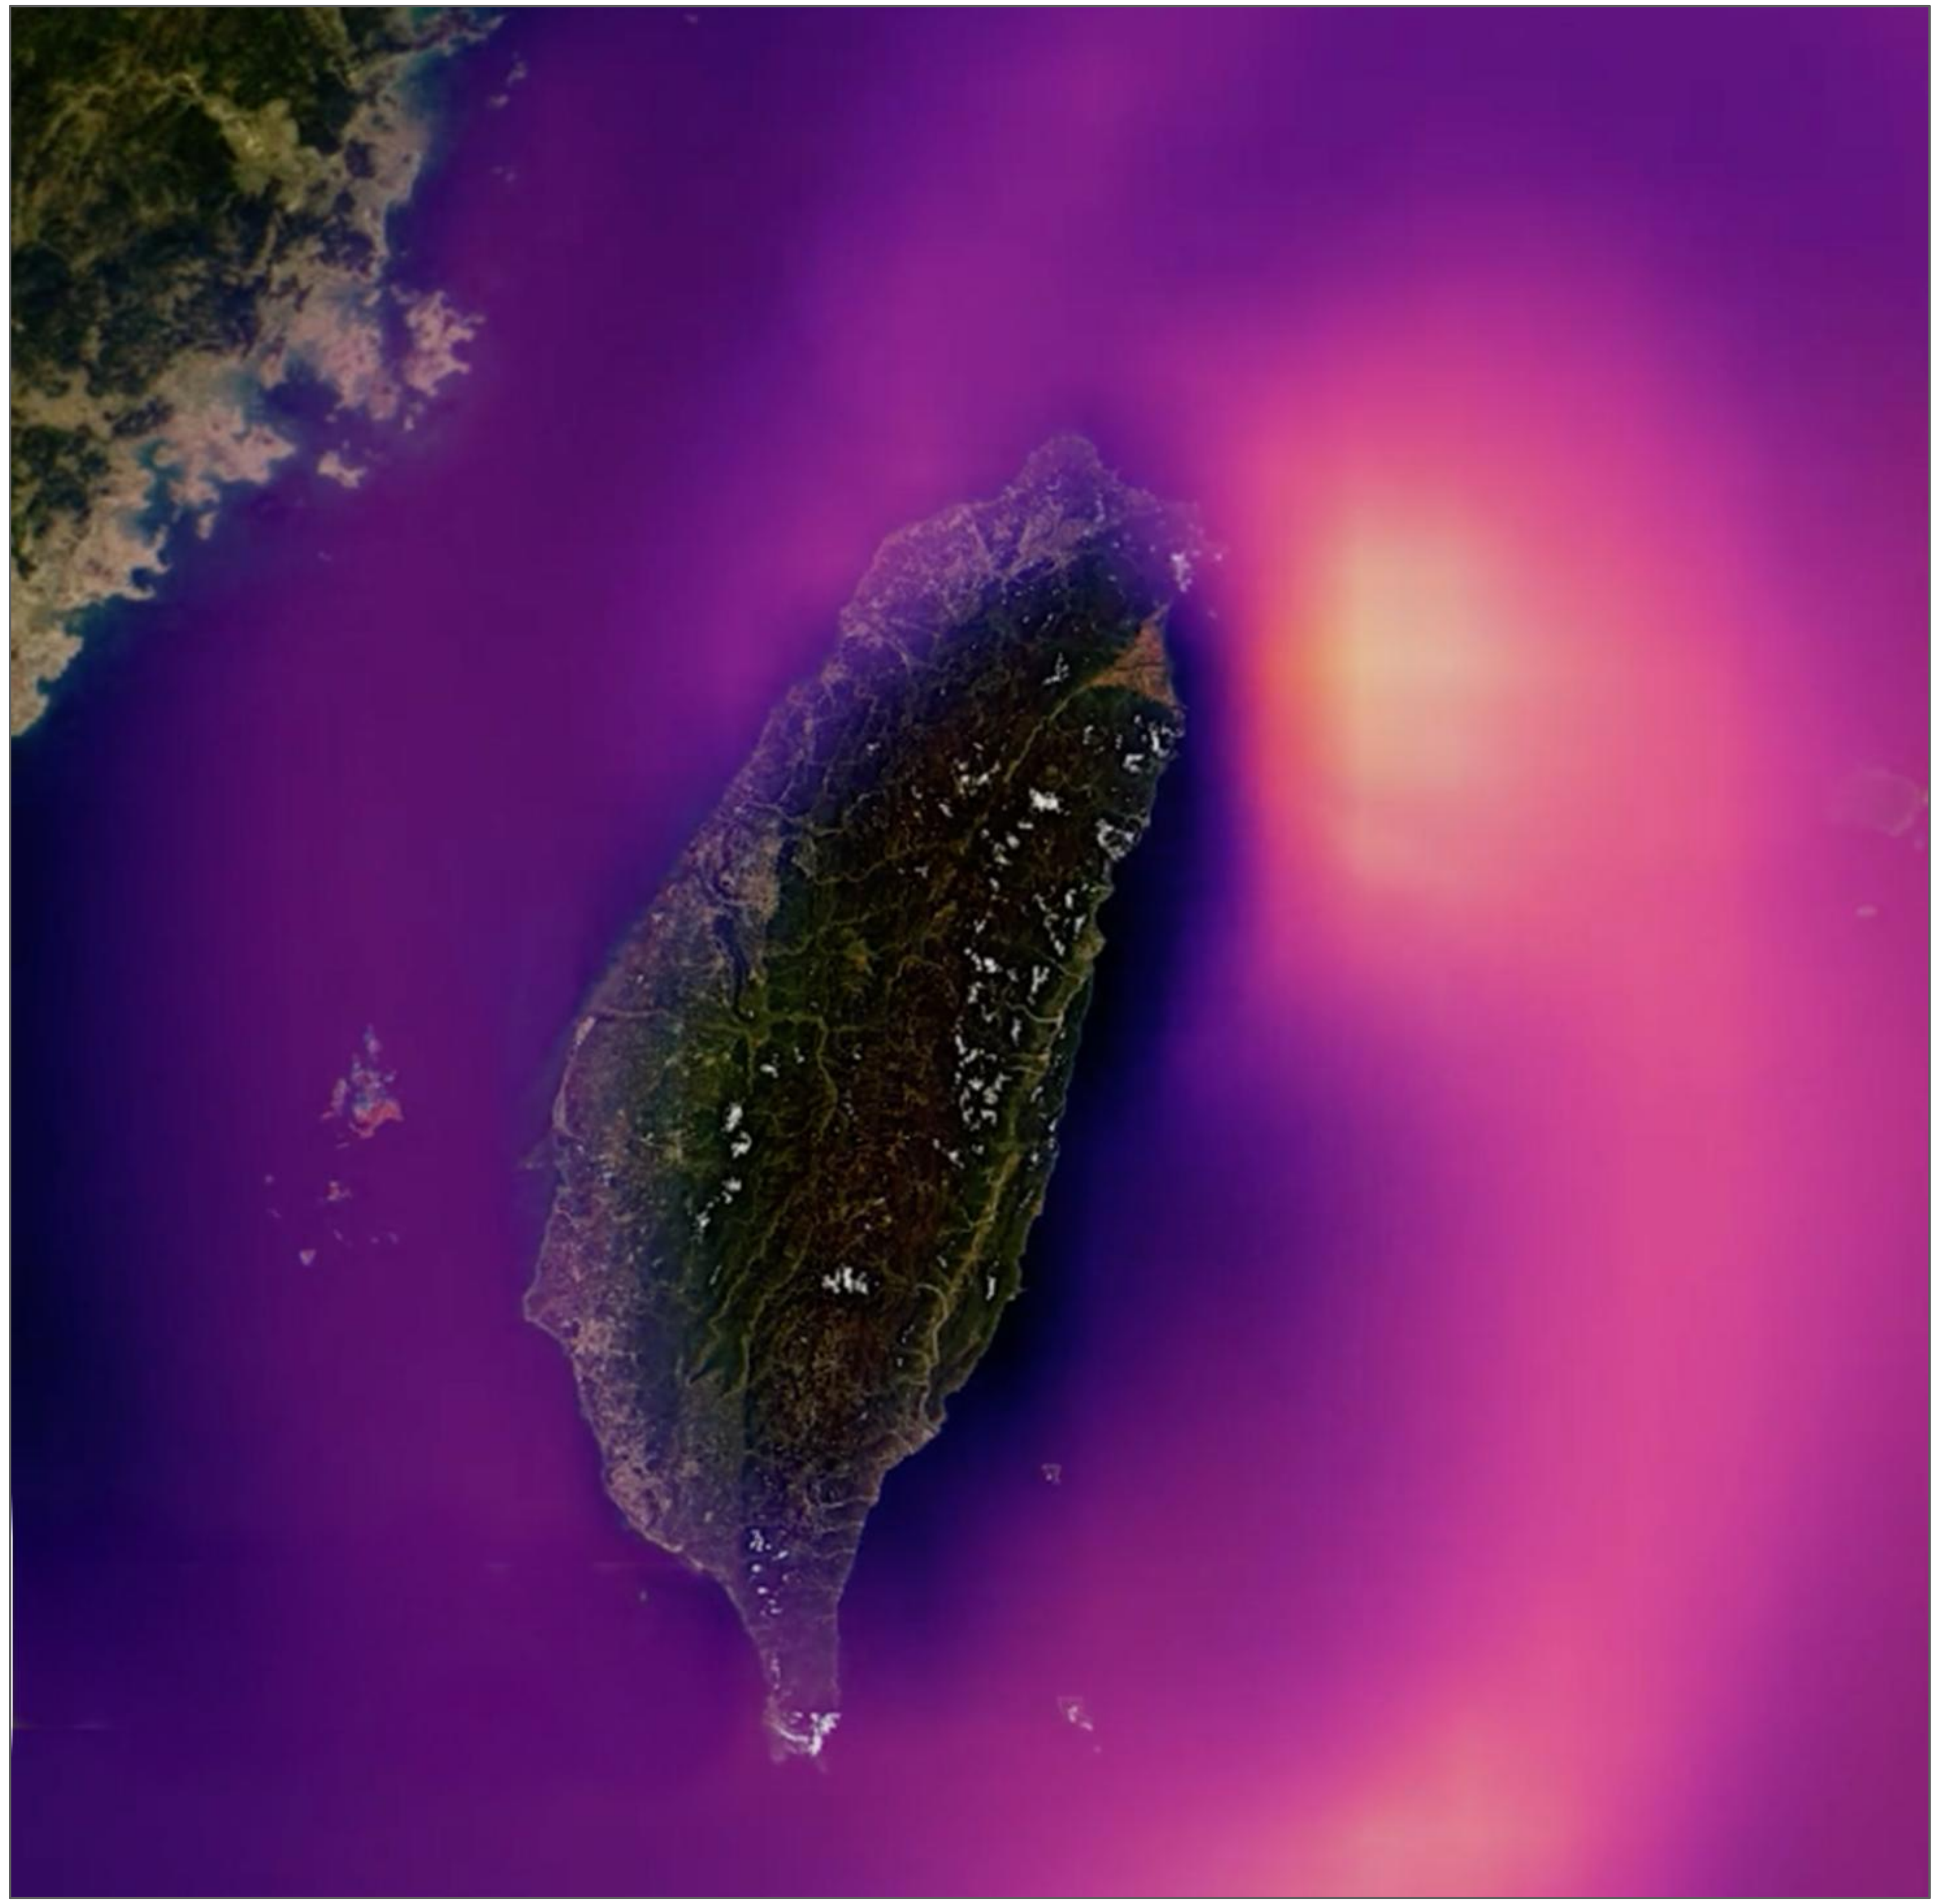

Global AI Forecast at 25km Resolution

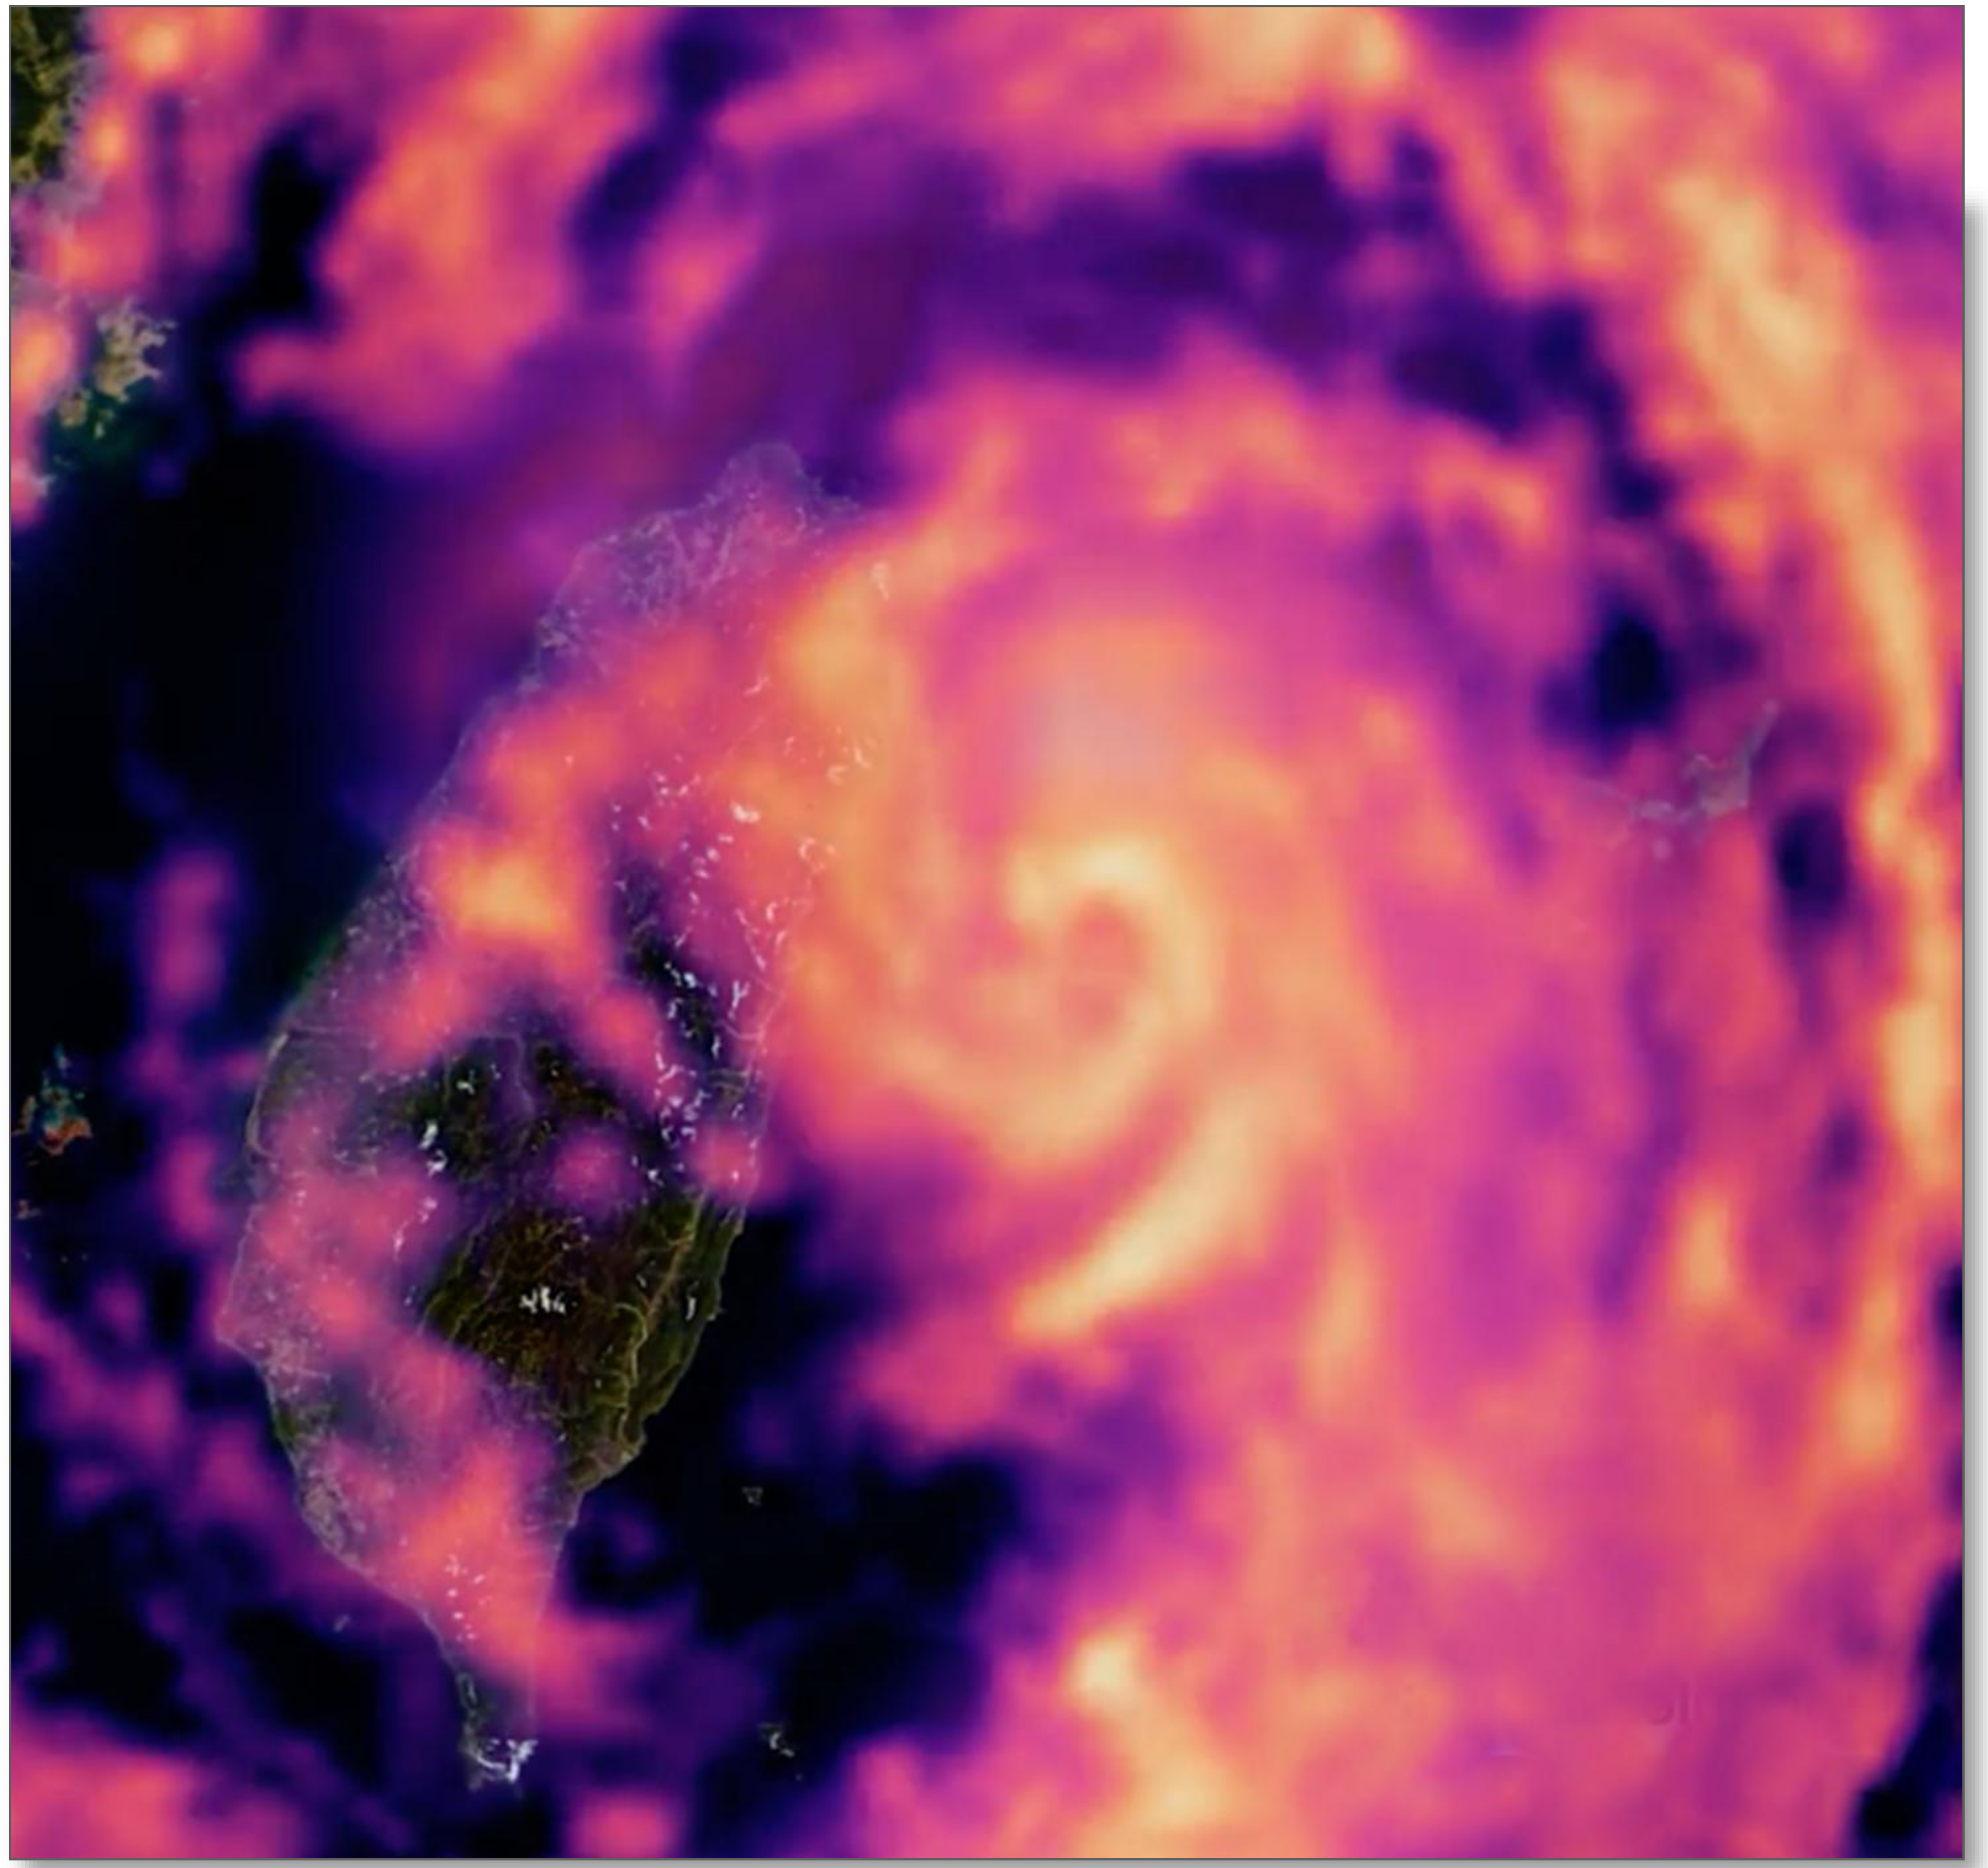

CorrDiff at 2km Resolution

NVIDIA

Earth-2 “Climate In A Bottle” Foundation AI Model

World’s First Generative AI Climate Model at Km-Scale Resolution

INPUTS

Time-of-Day

Day-of-year

Sea-surface
Temperature

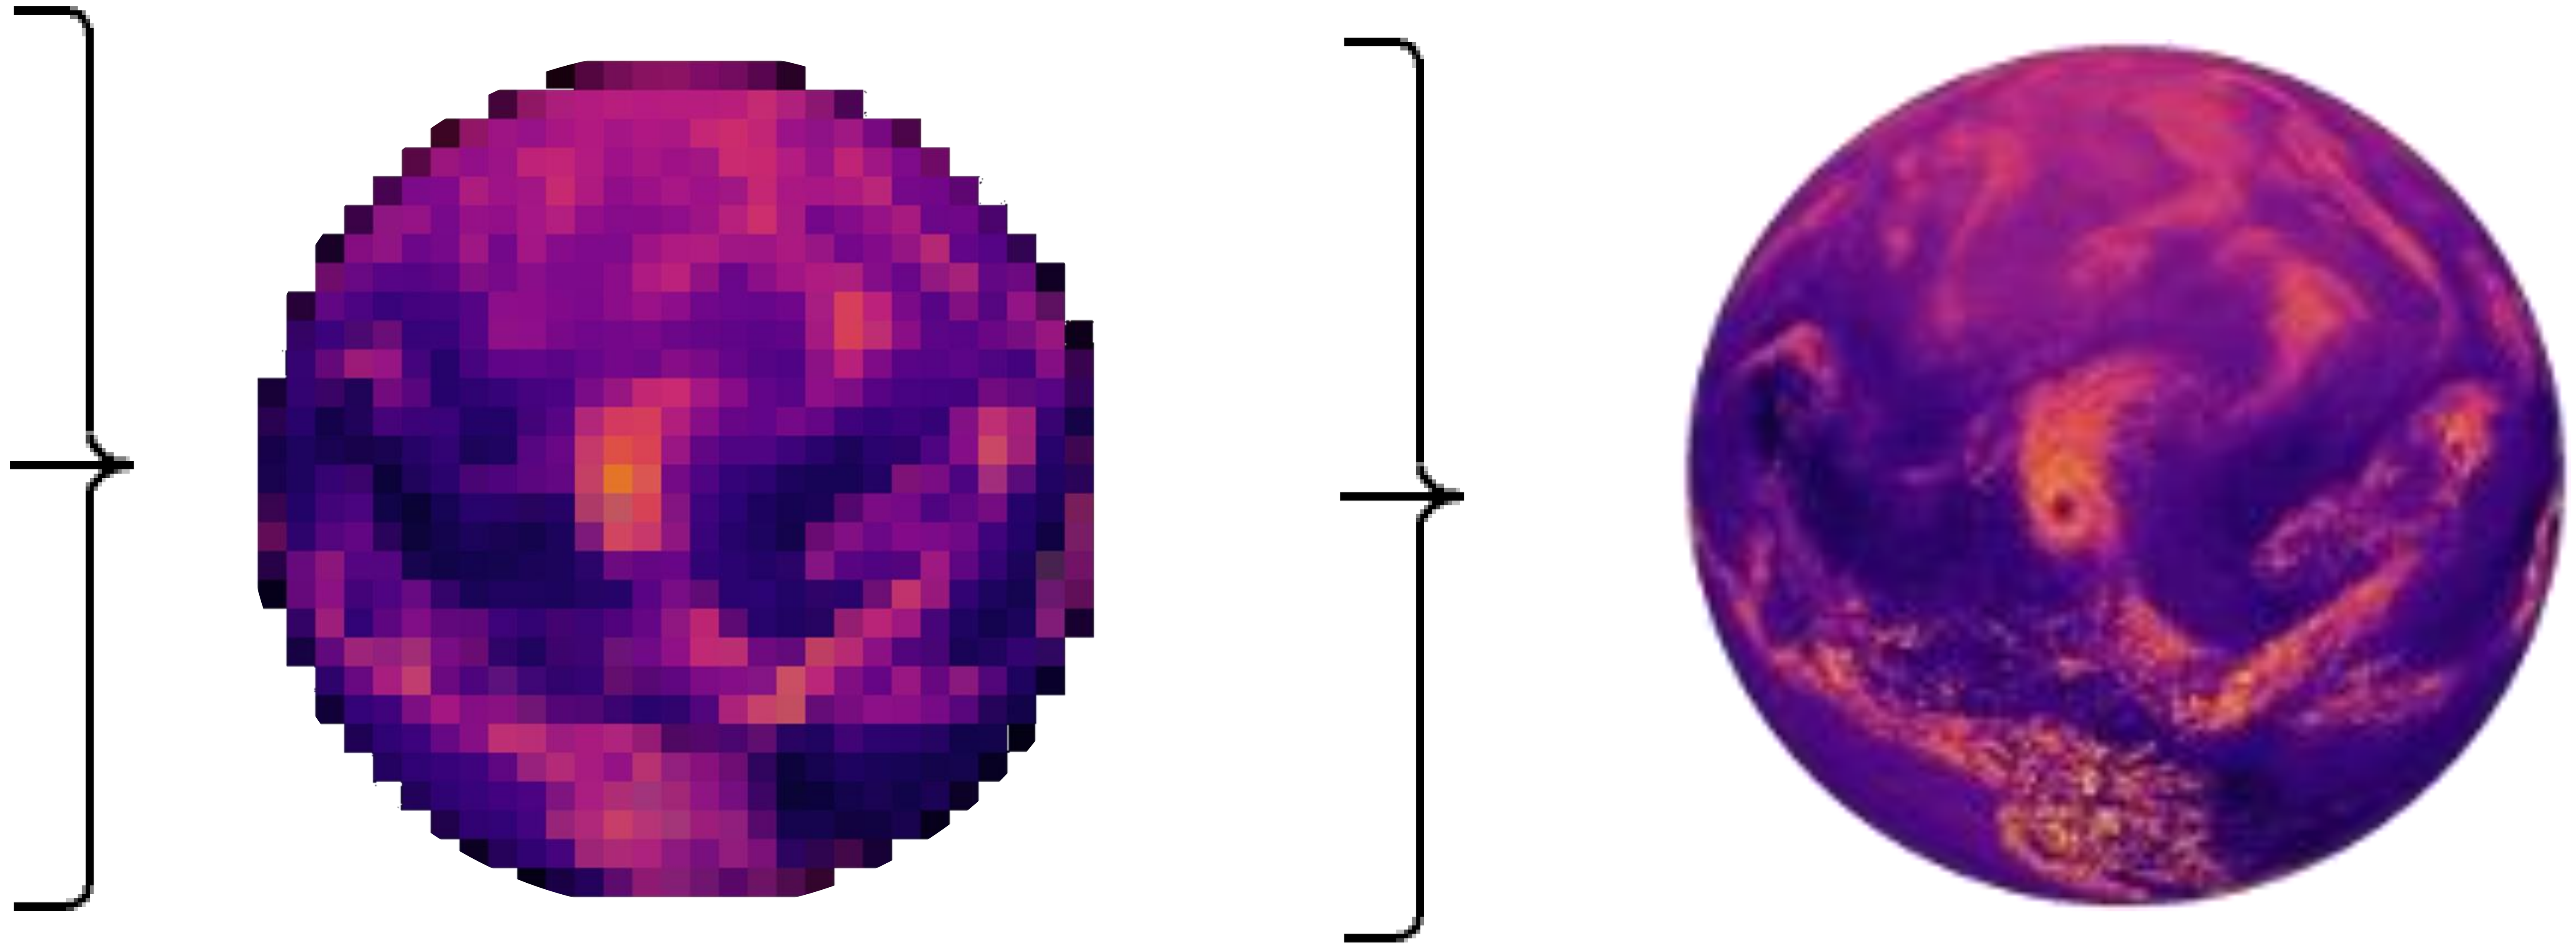

100 km

5 km

Super-resolution

Correcting Climate
Models

What-if Exploration

Data Compression

NVIDIA.

Earth-2 Powering Next-Gen Extreme Weather Simulation

AI Weather Generation

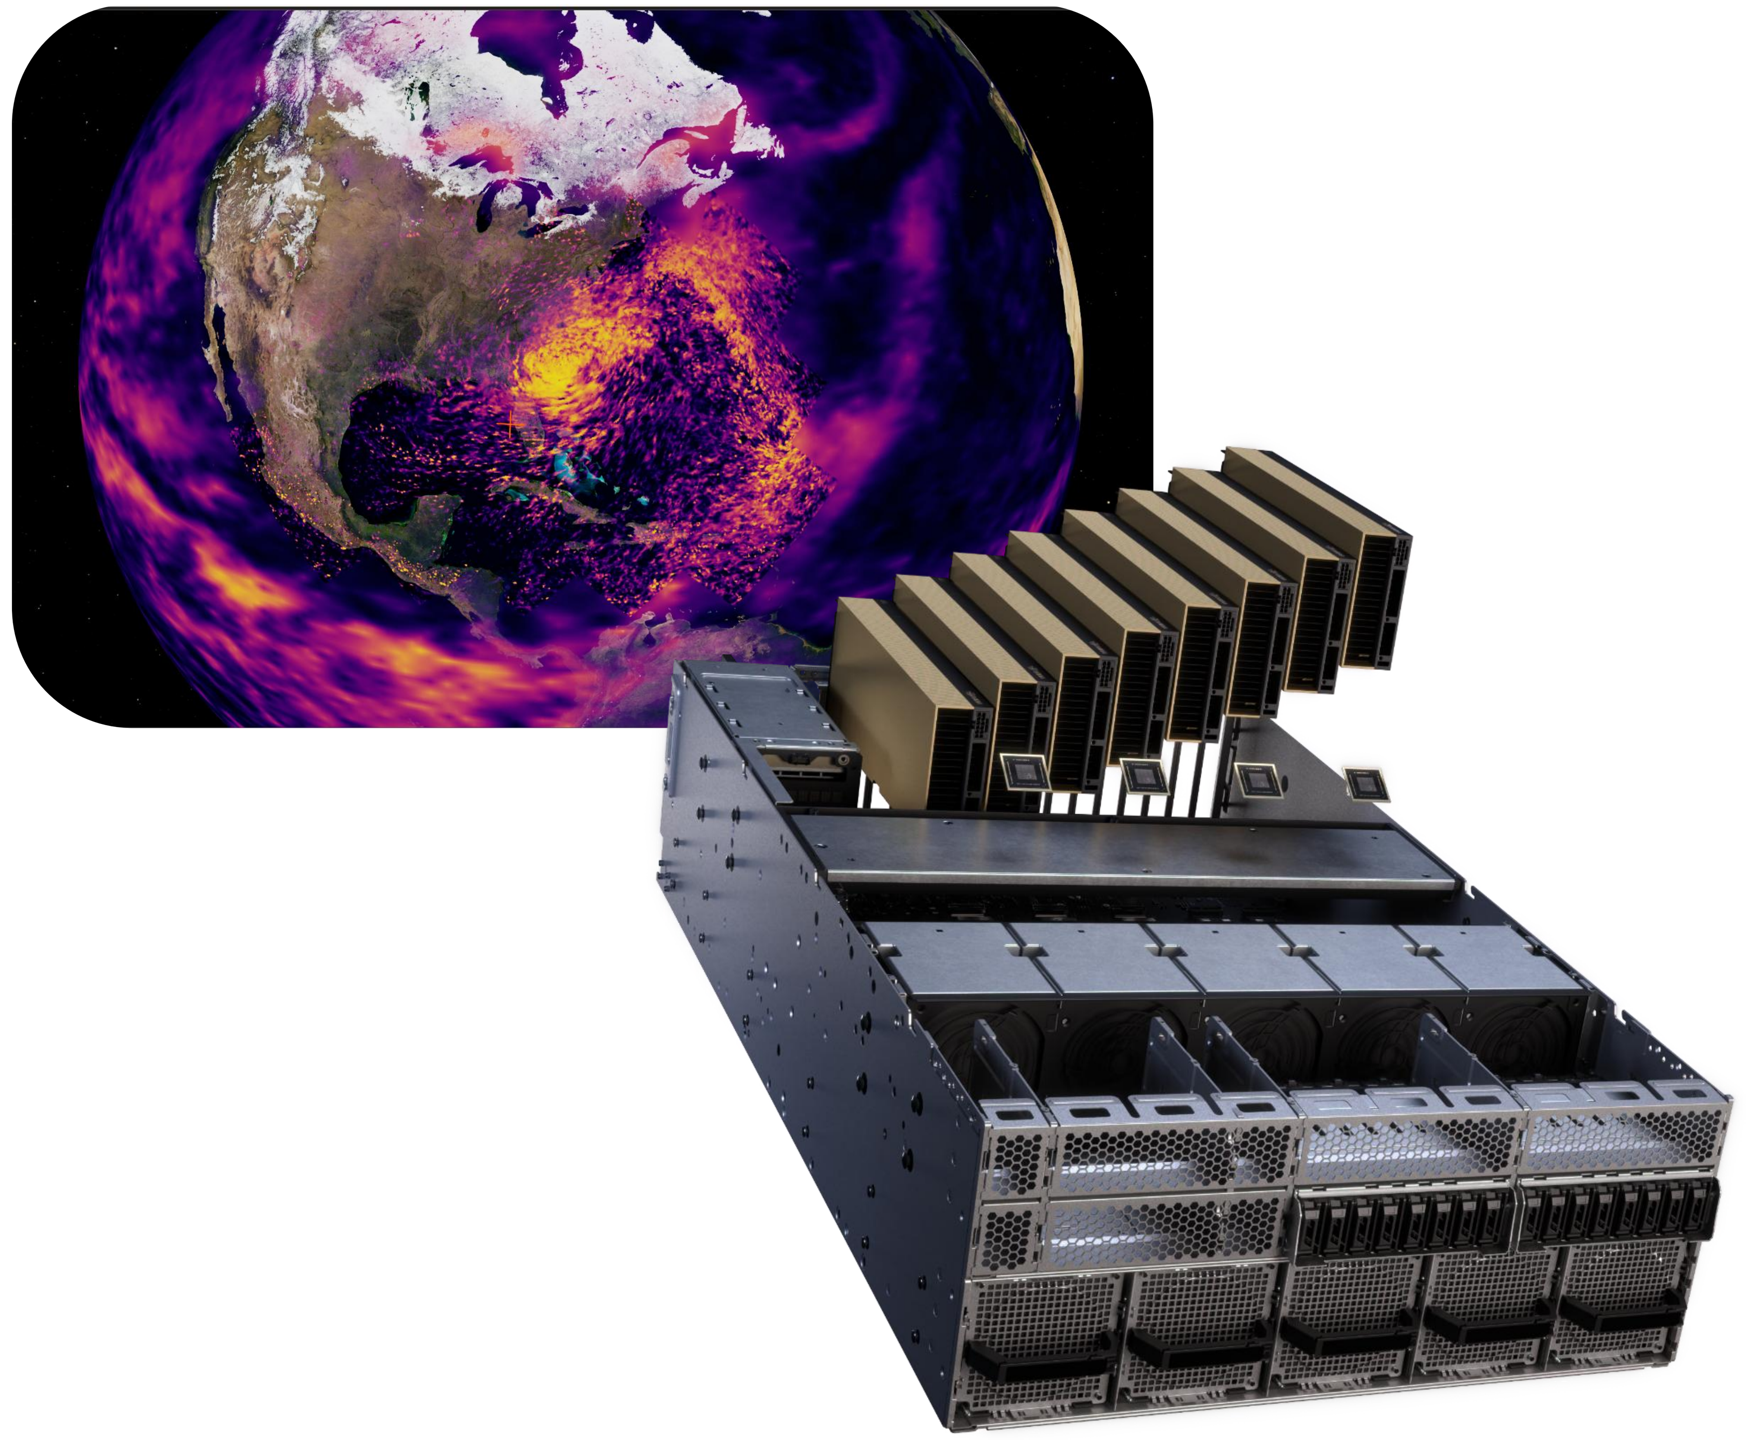

RTX PRO SERVER

Workflow for Conditional Weather Generation

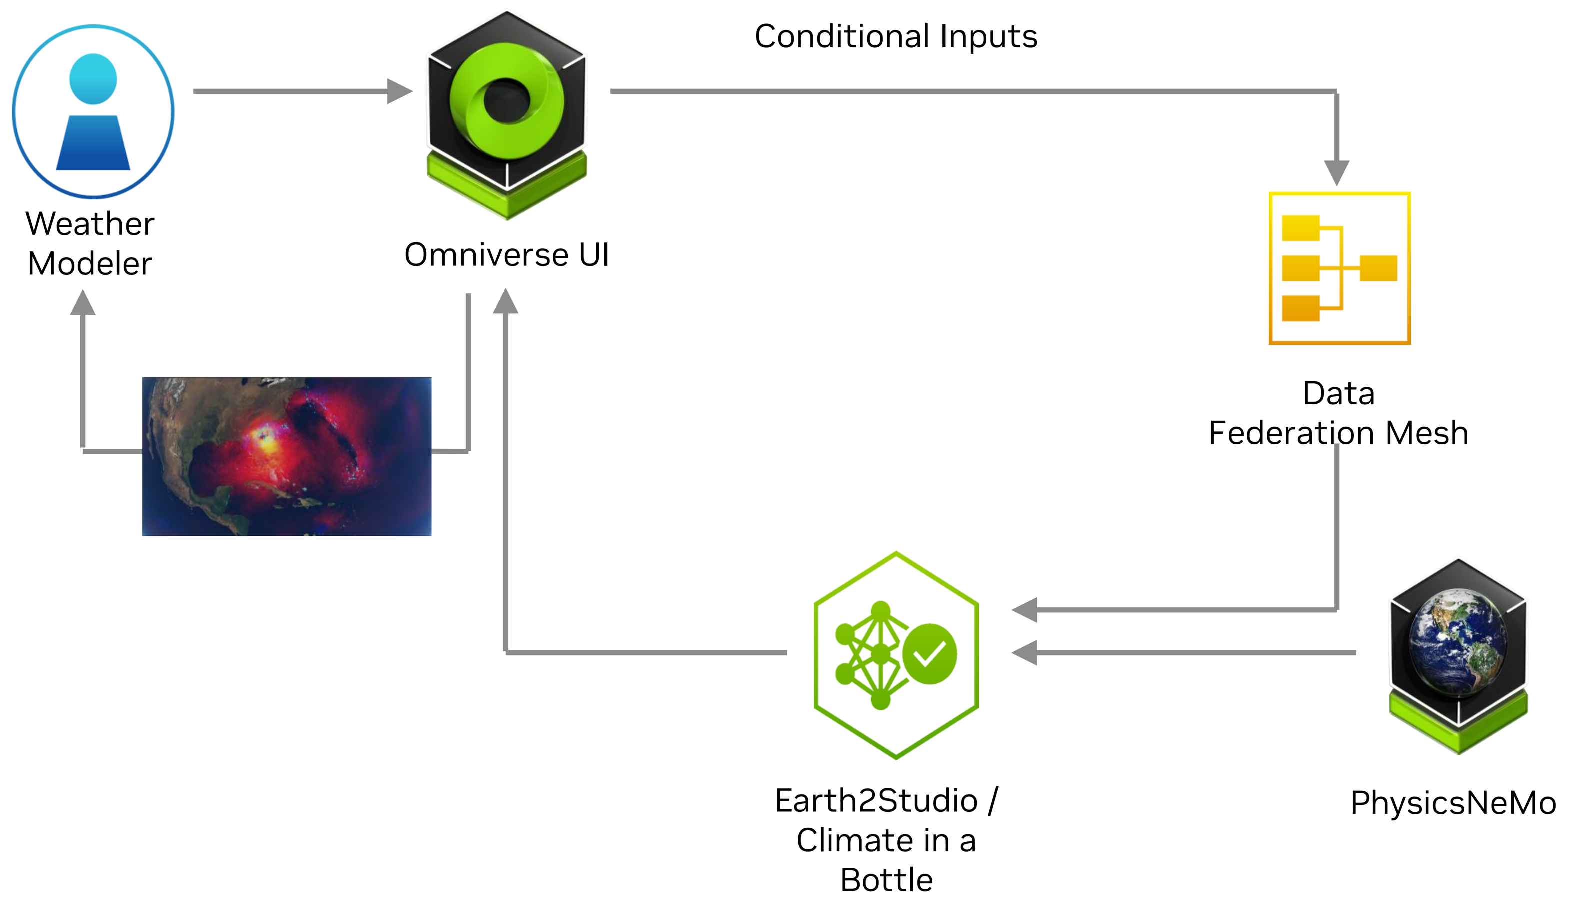

Score-Based Data Assimilation

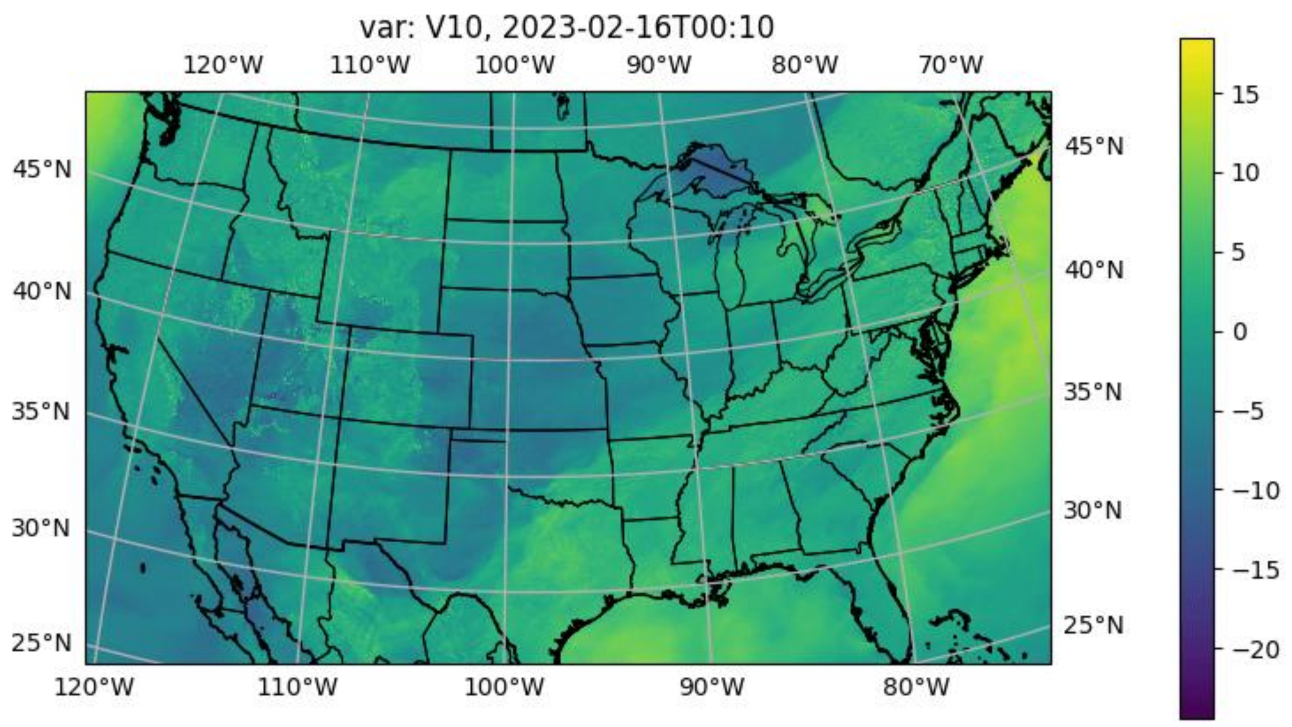

Earth-2 | PhysicsNeMo | Omniverse

Earth2Studio

Library for rapid exploration and validation of pretrained models

- Python library built on PhysicsNeMo for **AI weather model inference**
- Building blocks and recipes for specific tasks
    - medium range, seasonal forecasts with diagnostics
    - Ensemble forecasting
    - Generative downscaling
    - Validation case studies and examples
- Built-in data pipelines e.g., CDS, ARCO, GFS, HRRR, and local/custom
- Select from **NIMs, pretrained models or bring your own**
- Train your custom models in PhysicsNeMo
- Open-source library and enterprise support available through NVIDIA AI Enterprise (NVAIE)

OSS: https://github.com/NVIDIA/earth2studio

Examples: https://nvidia.github.io/earth2studio/examples/

Medium-range forecasting

- FourCastNet (AFNO/SFNO) Publication
- PhysicsNeMo-GraphCast (GNN) Publication
- AIFS Publication
- Aurora Publication
- FengWu (Transformer) Publication
- FuXi (Transformer) Publication
- PanguWeather (Transformer) Publication
- StormCast (Diffusion) Publication

Subeasonal-to-seasonal forecasting

- DLWP (CNN) Publication
- DLWP HEALPix (CNN) Publication

Generative downscaling

- CorrDiff (Diffusion) Publication
- Climate in a Bottle Publication

AI-on-top

- Precipitation (AFNO)
- Tropical cyclones (CNN) Publication
- Atmospheric rivers (CNN) Publication
- Temporal interpolation (AFNO) Publication

Edge Supercomputing

Next Generation Instruments are Producing Increasing Amounts of Data

Complexity of Experiments is Booming and Human Insight is Now the Bottleneck

Radio Astronomy

ngVLA – 244 Dishes

100 Petabytes per Year

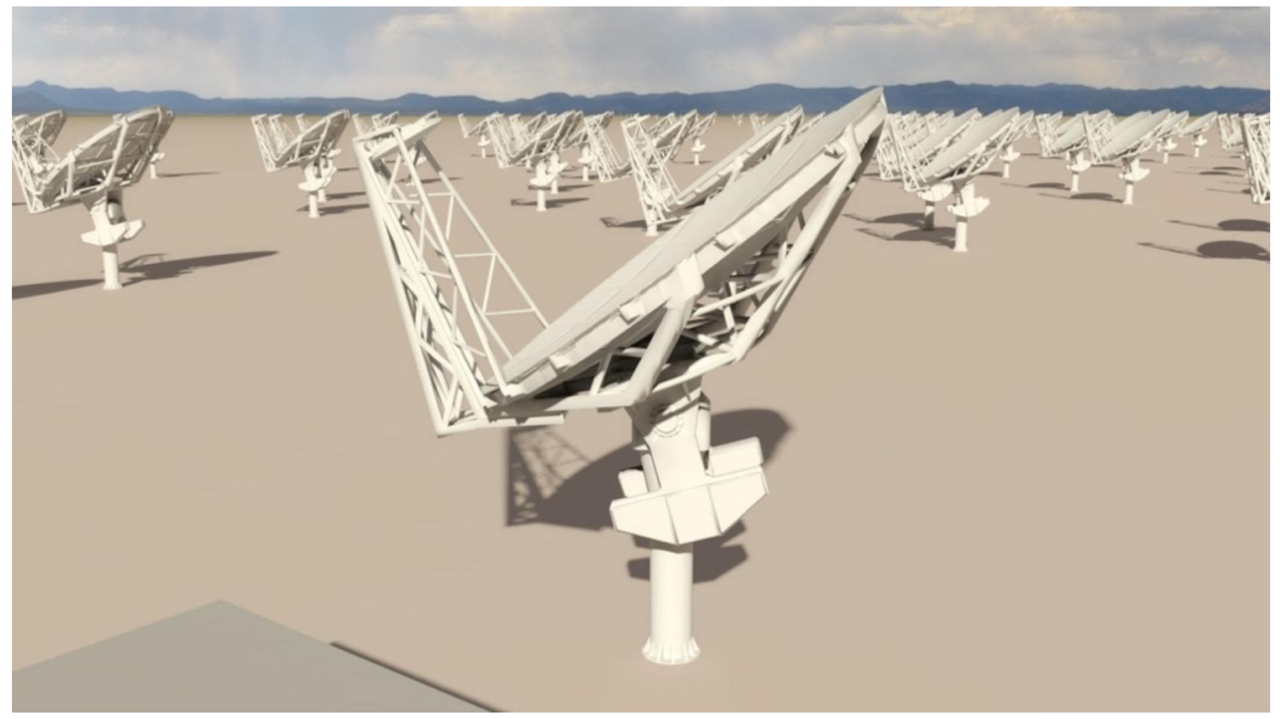

SKA – 200 Dishes

1 Exabyte per Year

Particle Physics

High Luminosity LHC

100x Data than Higgs Boson Discovery

Advanced Laser Systems

100x Increase in Repetition Rate

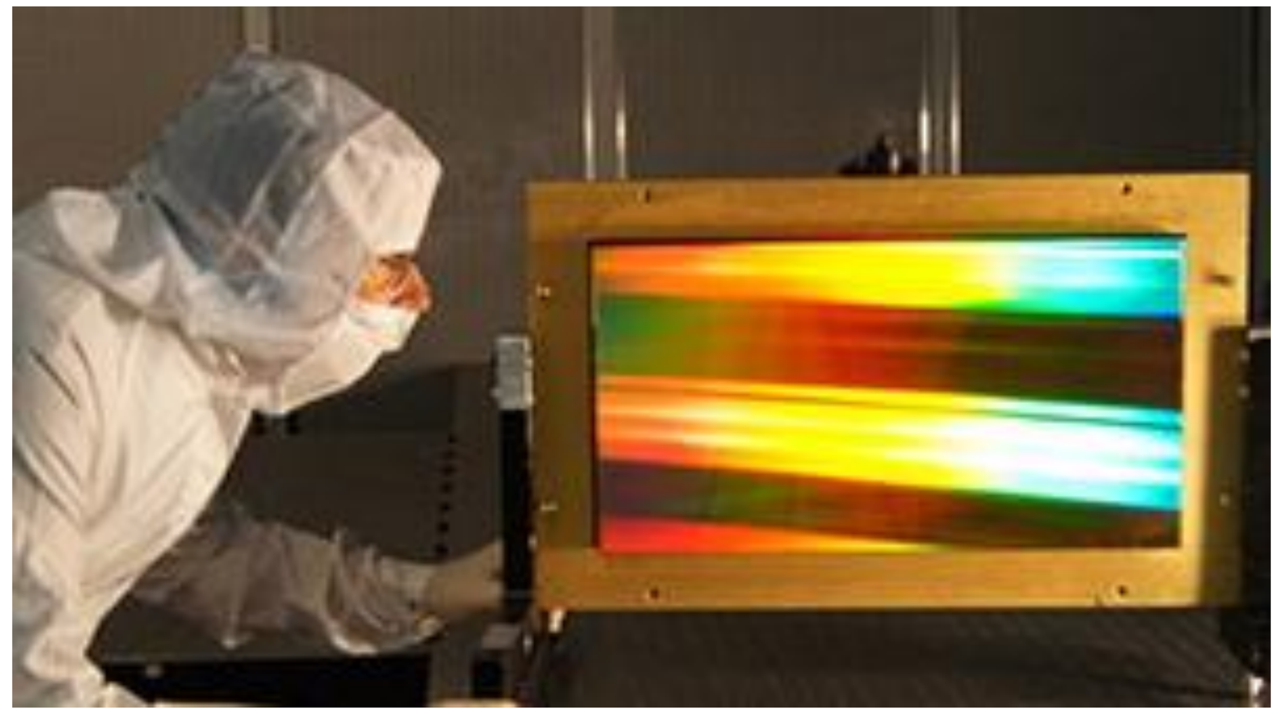

Light Sources

APS-U – >60 beamlines

100-200 Petabytes per Year

Free Electron Laser LCLS-II

1 MHz Rep-Rate Upgrade

NVIDIA Holoscan

SDK for Building AI-Enabled Sensor Processing Applications

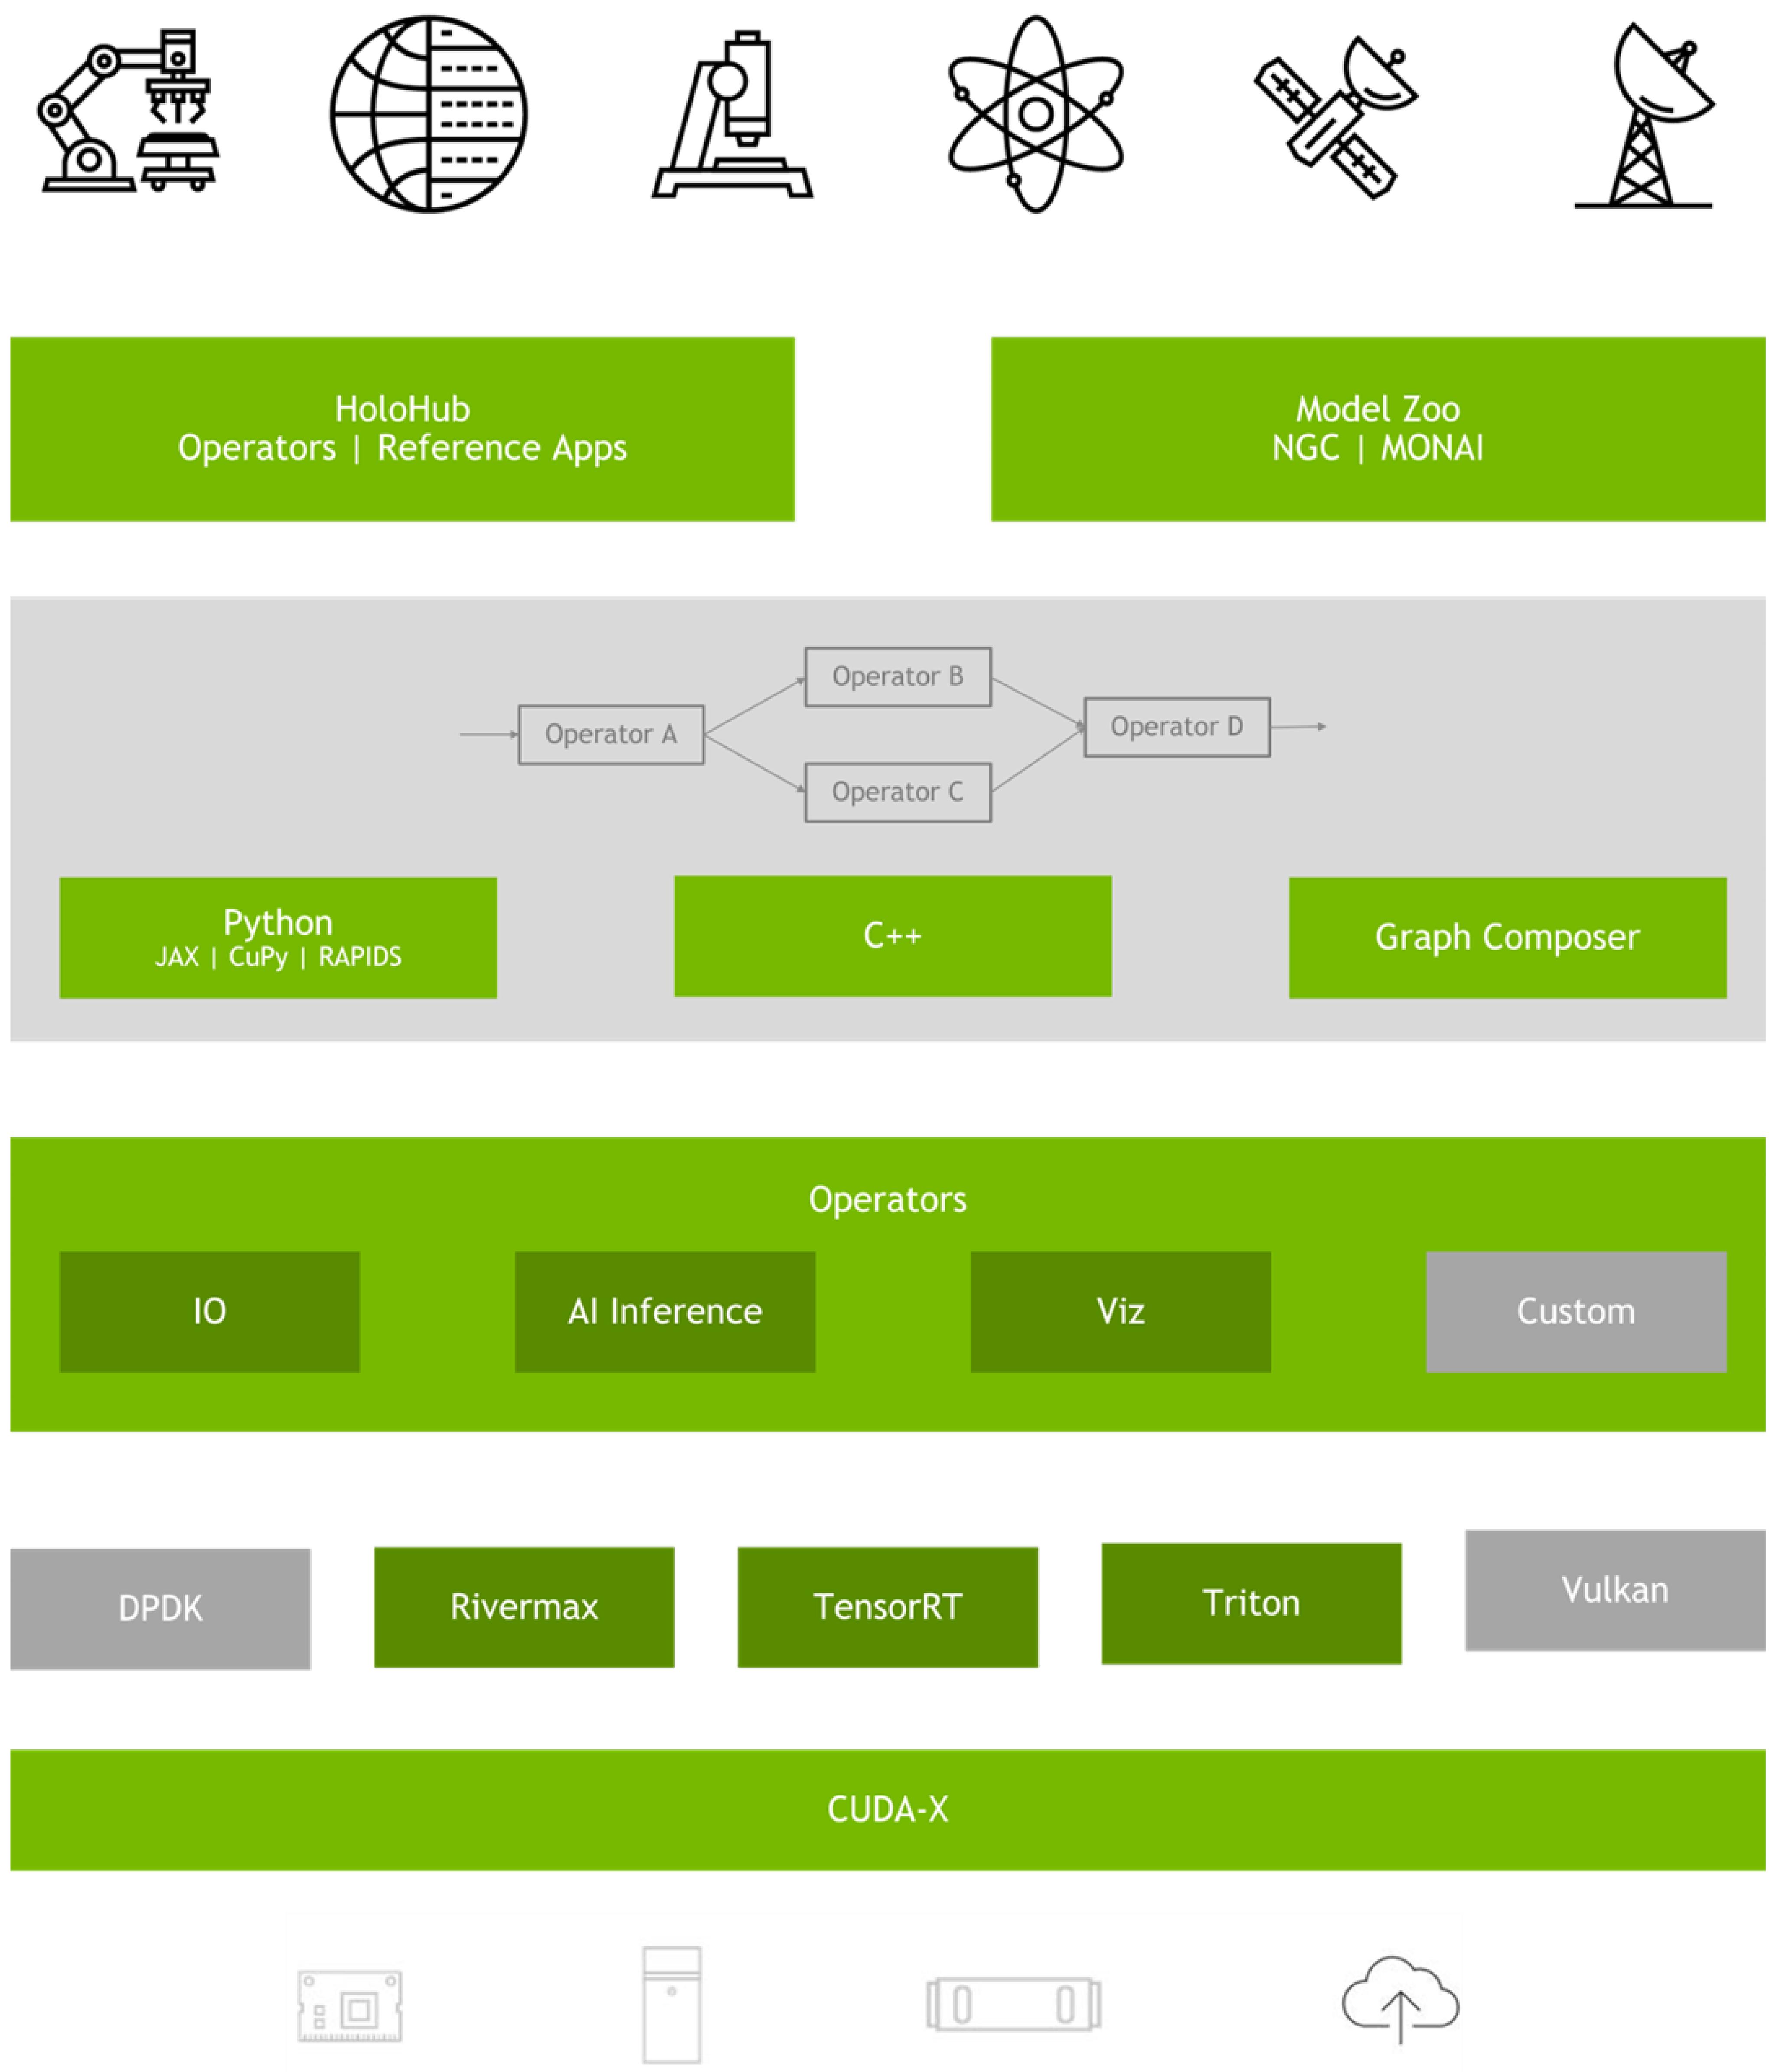

Features

- C++ and Python APIs for domain agnostic sensor data processing workflows
- Scalable from IGX (ARM + GPU) to DGX (x86 + A100)
- Sample applications to jump-start ML/AI-enabled and Accelerated Computing sensor pipelines with Holohub
- AI Inference with pluggable backends such as TensorRT
- Apache 2 Licensed and Available on GitHub

Benefits

- Simplifies sensor I/O to GPU
- Simplifies the deployment of an AI model in a streaming pipeline
- Provides customizable, reusable, and flexible components to build and deploy GPU-accelerated algorithms
- Scale workloads with Holoscan Cloud Native

Real Time Nanoscale Imaging with NVIDIA Holoscan

Brookhaven National Laboratory NSLS-II, DECTRIS

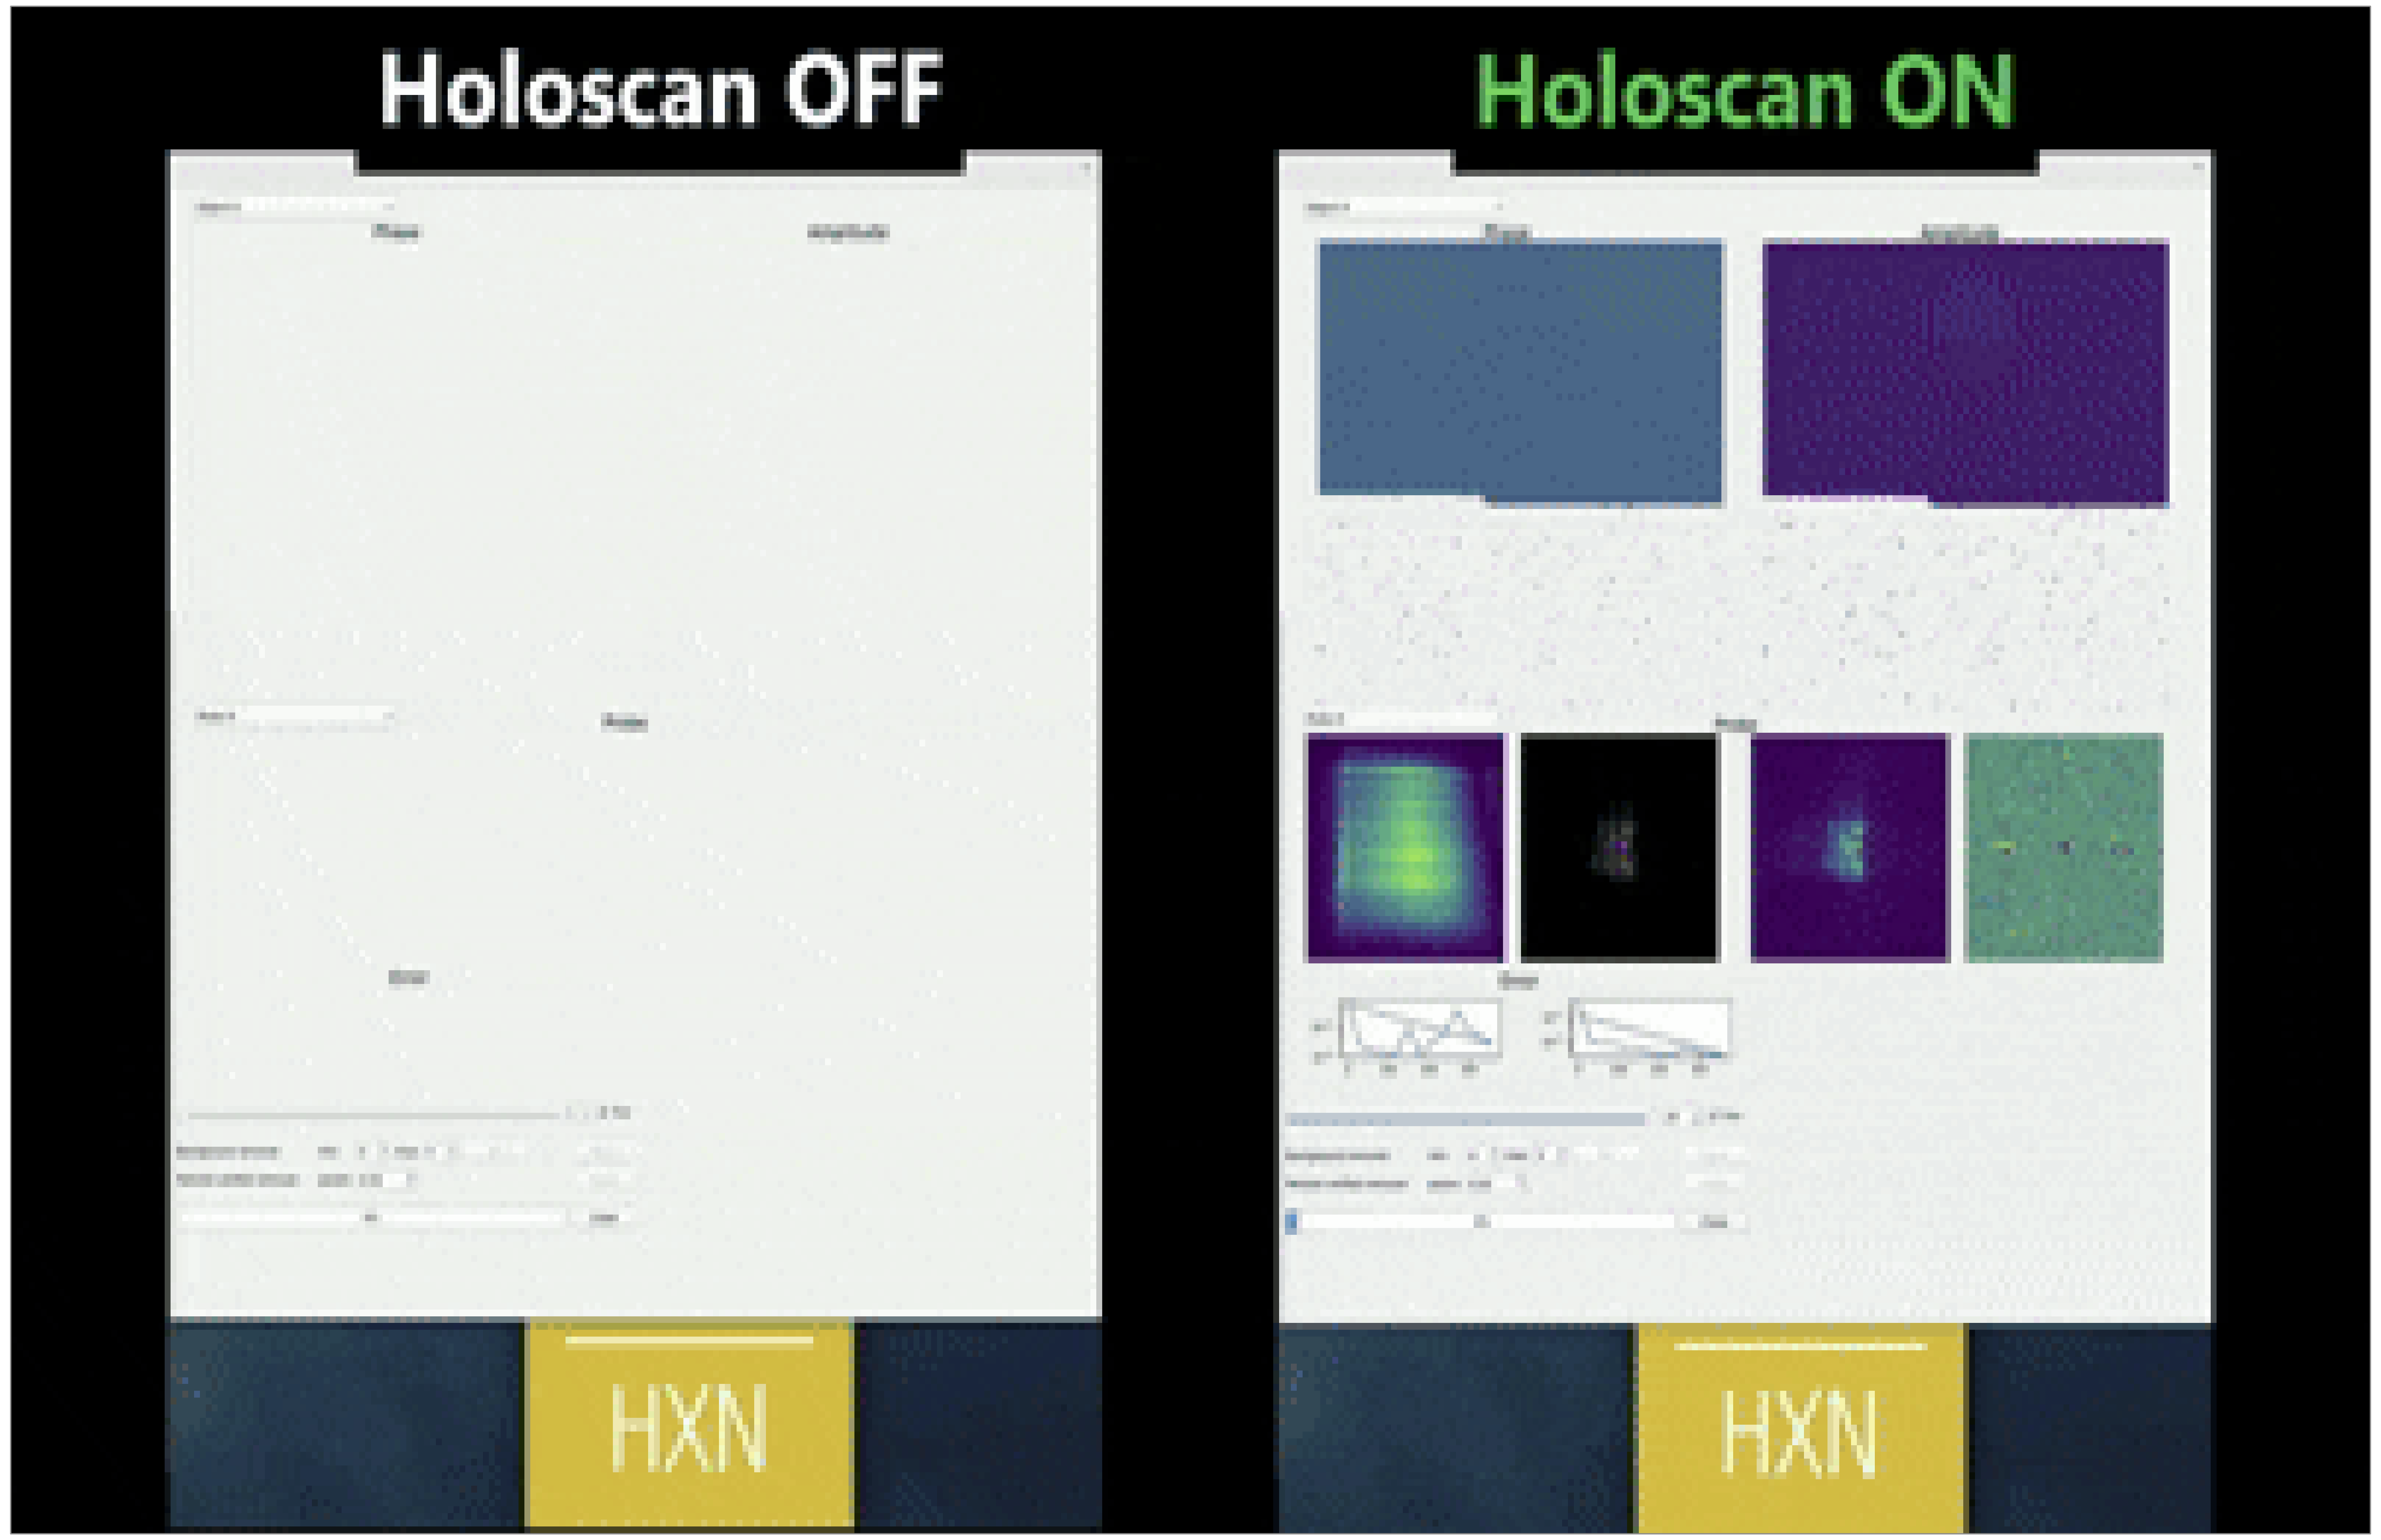

Demonstration of Offline versus Online Ptychographic Imaging

Self-assembly from nanoparticles, potentially useful as optical and mechanical metamaterials

For visualization purposes, both offline and online videos have been sped up by 2x

- **Current beamline scan takes about 8 seconds, with reconstruction taking 1-2 minutes**
  - Requires full dataset from x-ray detector and disk storage prior to reconstruction
- **Impact of Edge HPC and Holoscan for real time imaging**
  - Real time reconstruction provides instantaneous user feedback on experiment results
  - Ability to handle higher data rates yields higher resolution experiments and improved science products
  - Laying foundation for AI-driven agents and autonomous experimentation, including AI based scan patterns, leading to the **discovery of new materials at unprecedented rates**

Qian Chen group, Department of Materials Science and Engineering, University of Illinois, Urbana-Champaign

HPC and AI at the Edge

AI, Combined with, Streaming Data and Unprecedented Rates Welcomes New Scientific Discoveries

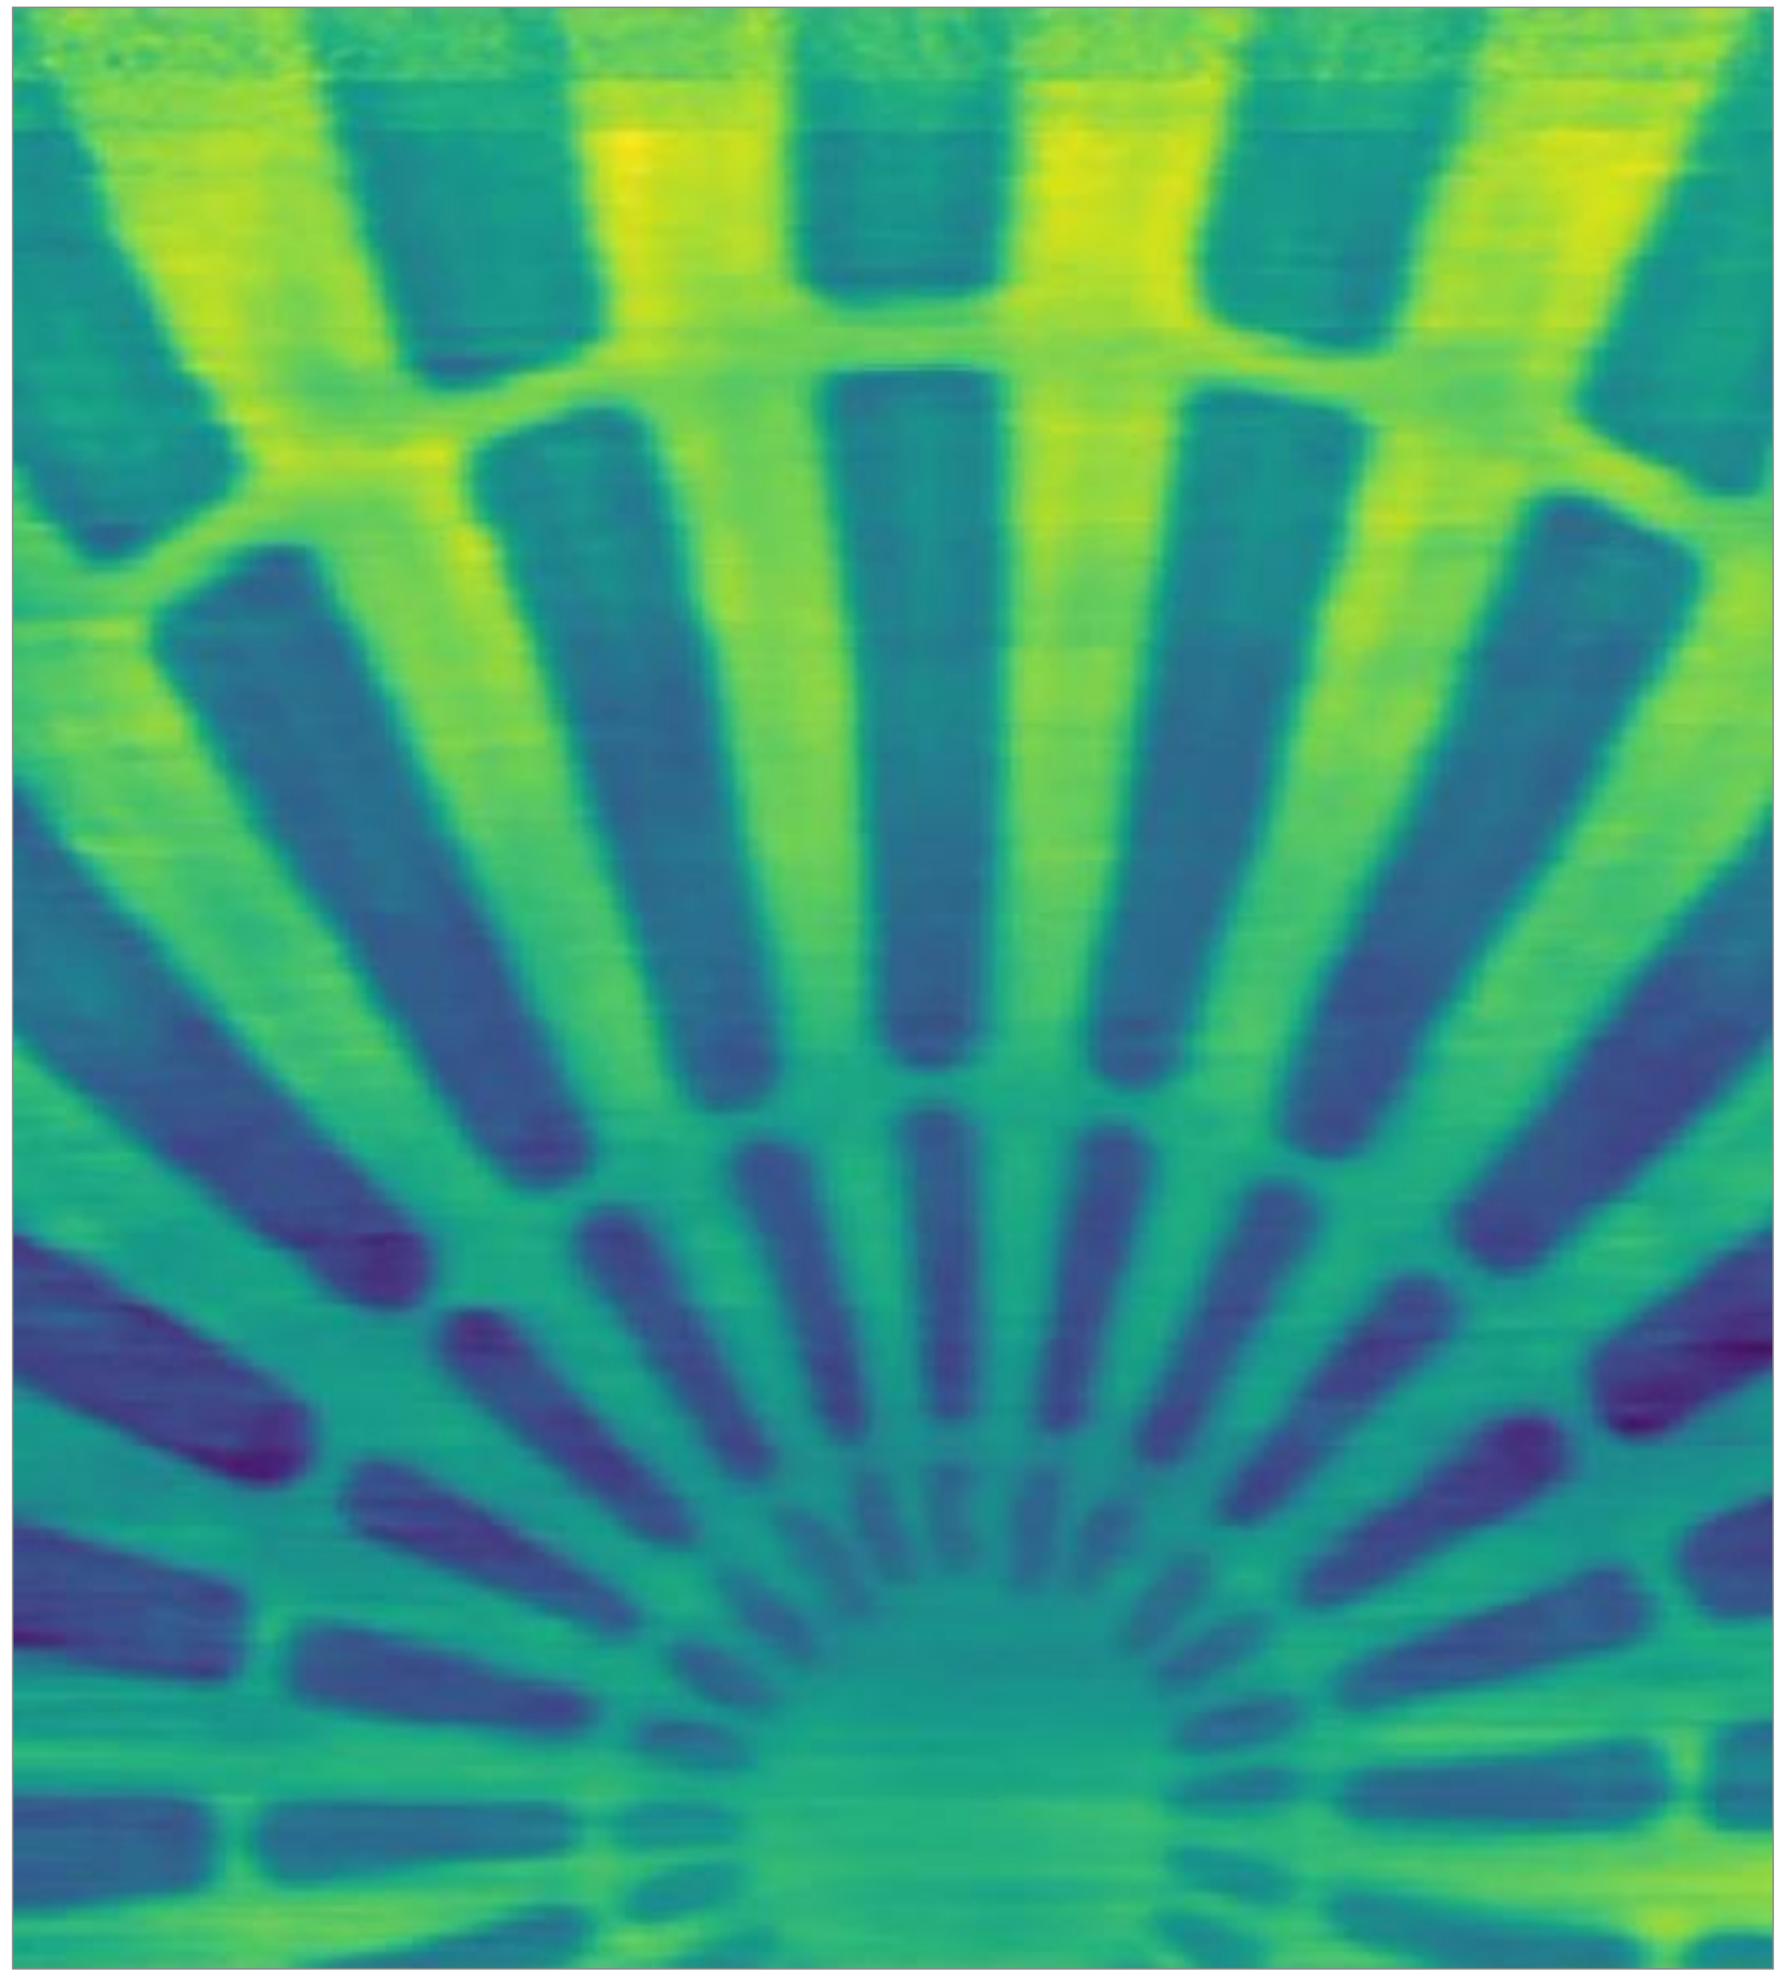

Nanoscale Imaging of Materials
Real Time Ptychography with NSLS-II and DECTRIS

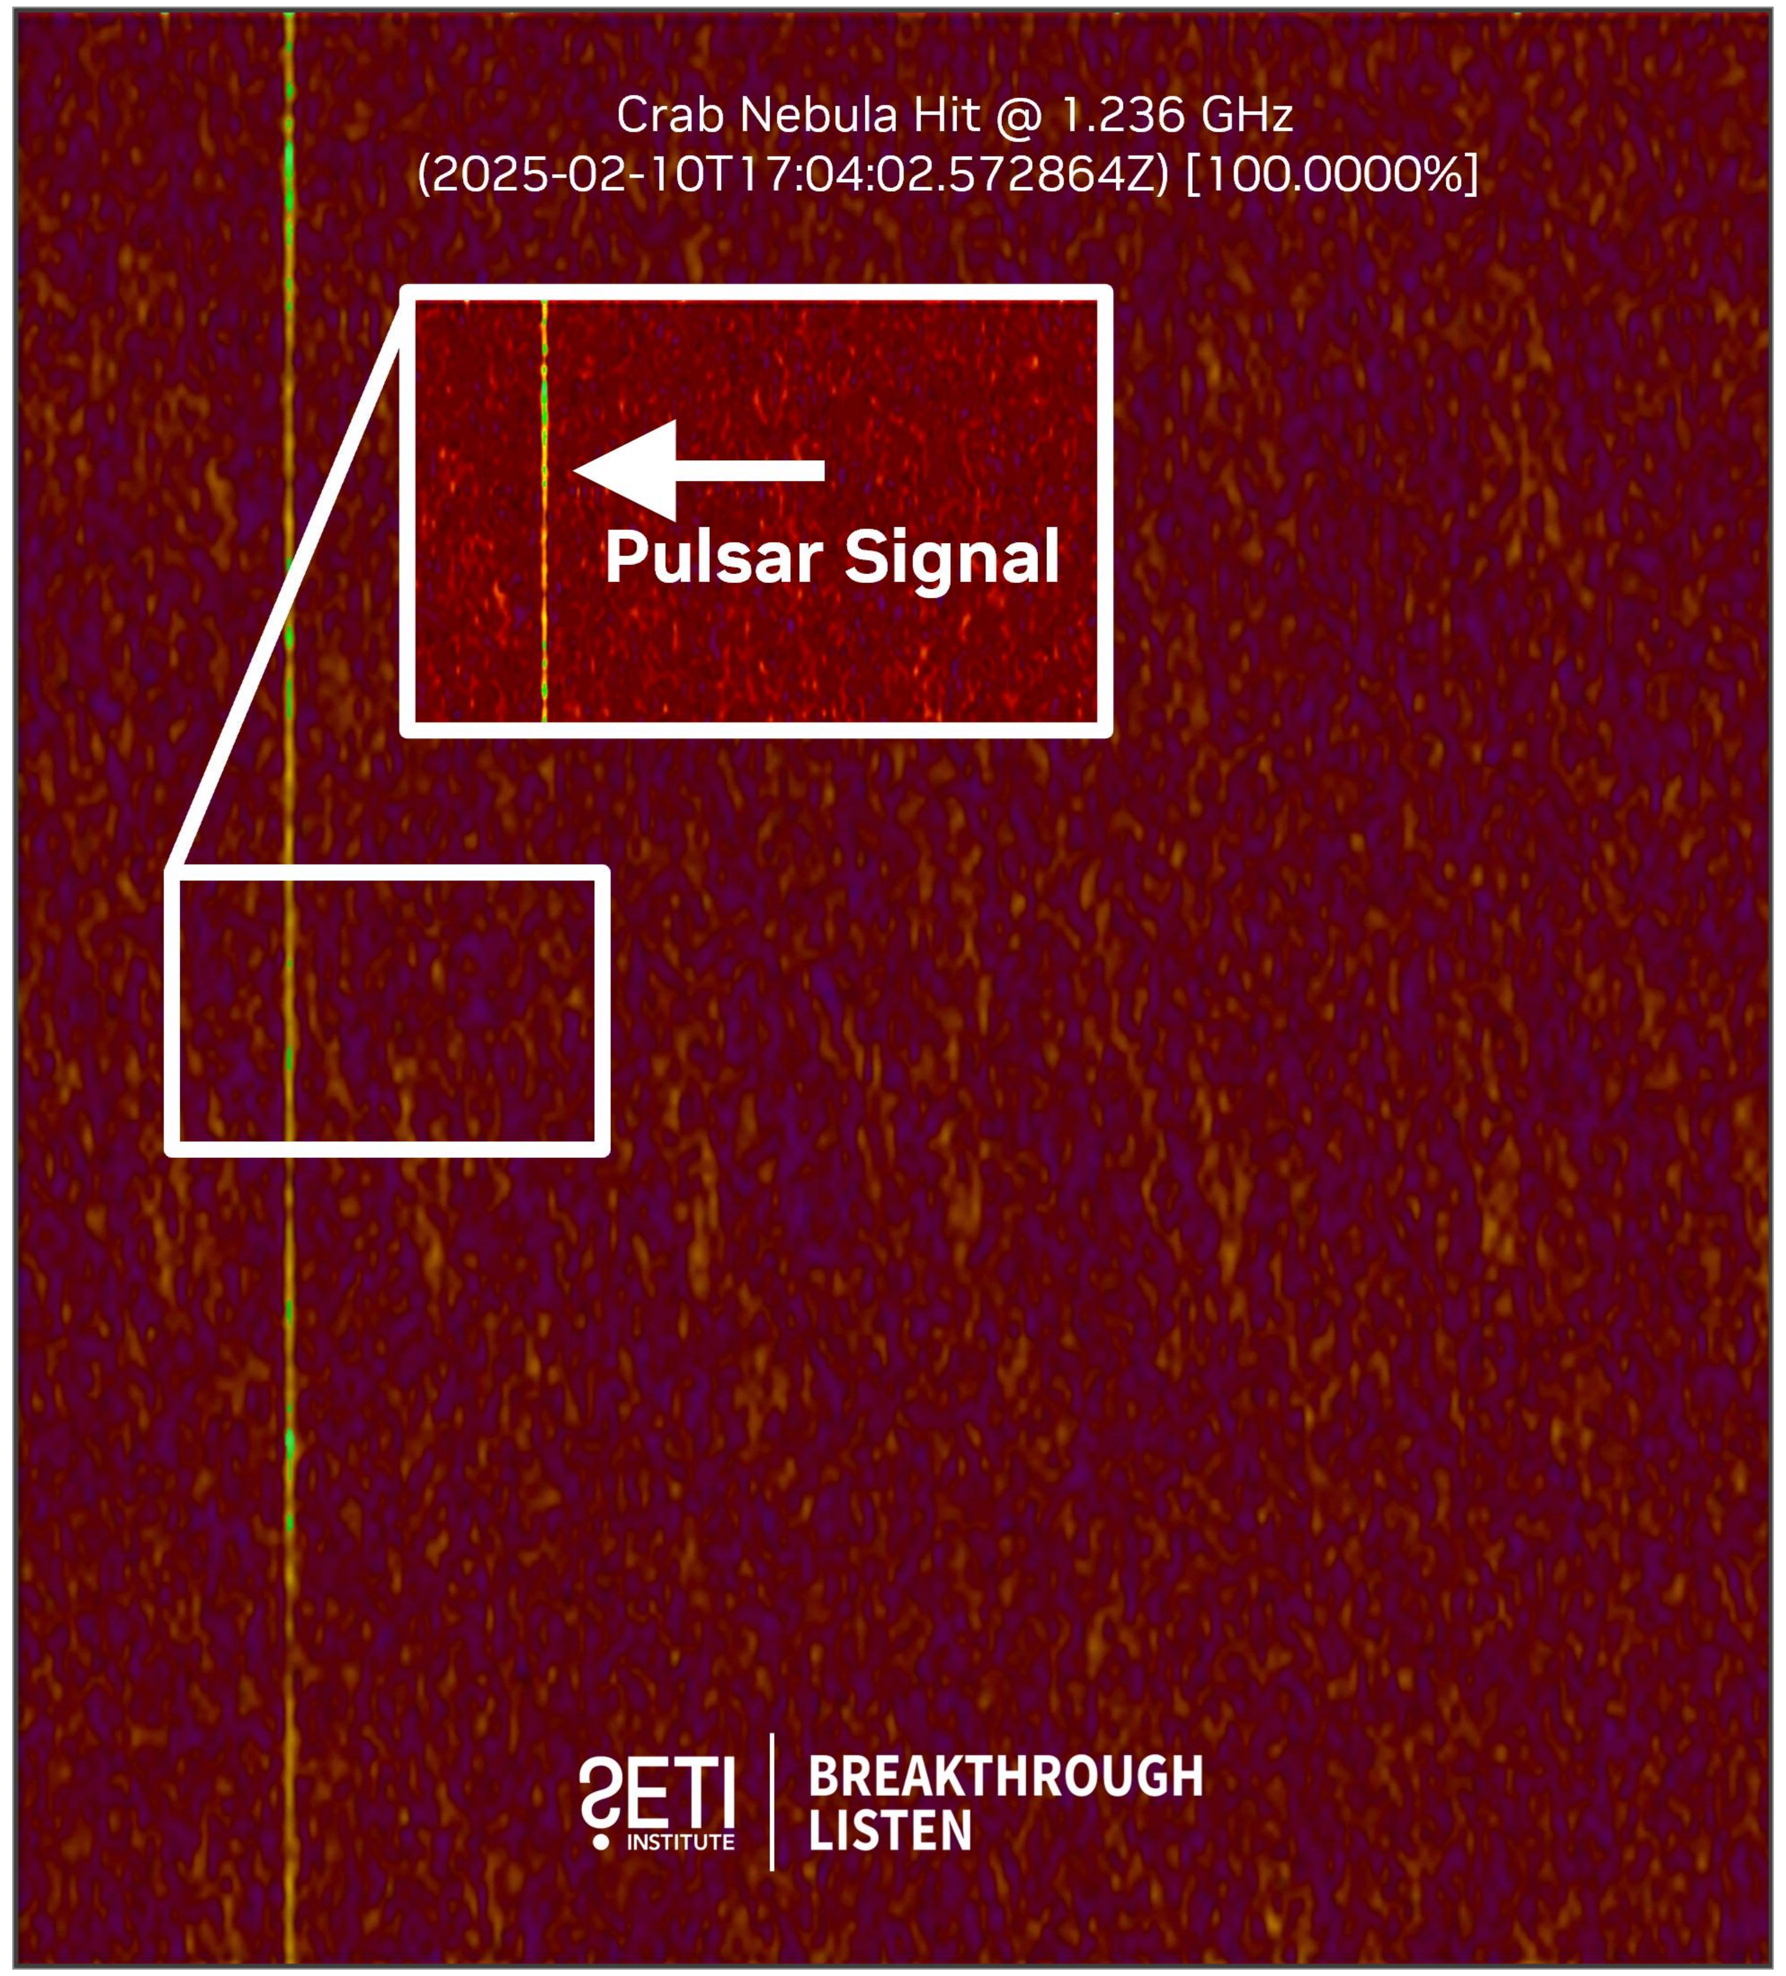

AI Search for Pulsars and Technosignatures
AI as a Detector with the SETI Institute and the Allen Telescope Array

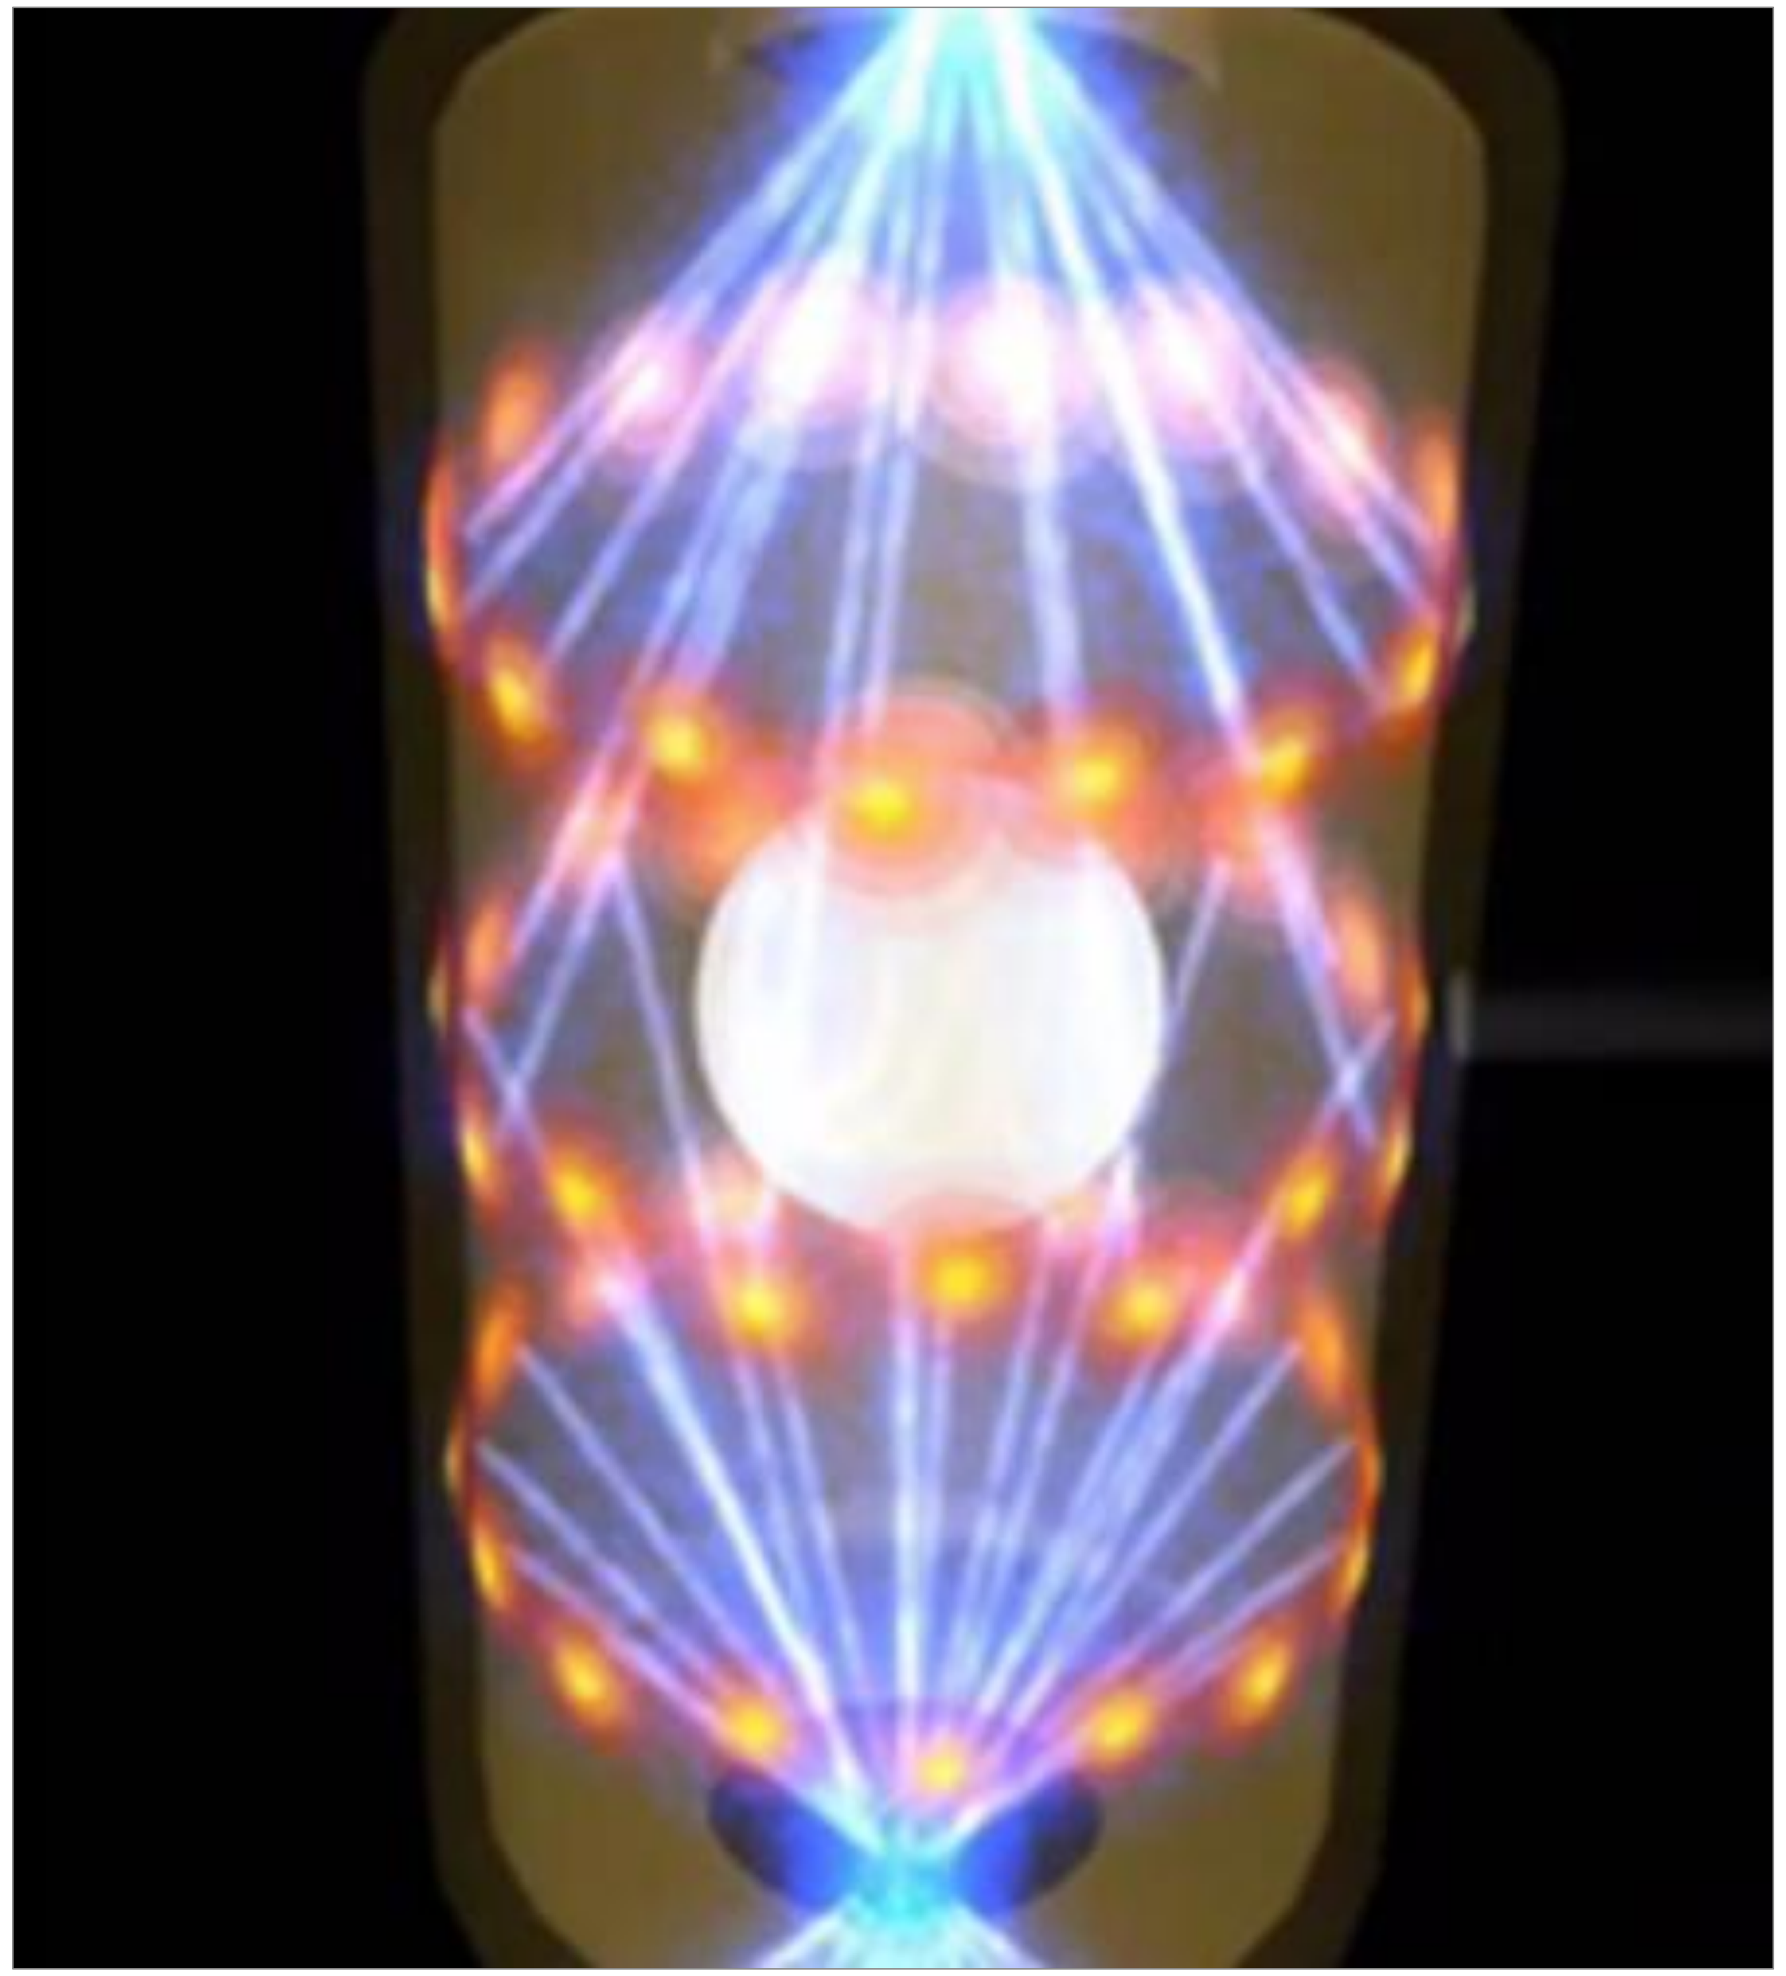

Self Driving Experiments
AI Guided Fusion Research with Lawrence Livermore National Laboratory

NVIDIA Holoscan Enables Real-Time Signal Detection with AI

AI Enabled Real Time Signal Detection

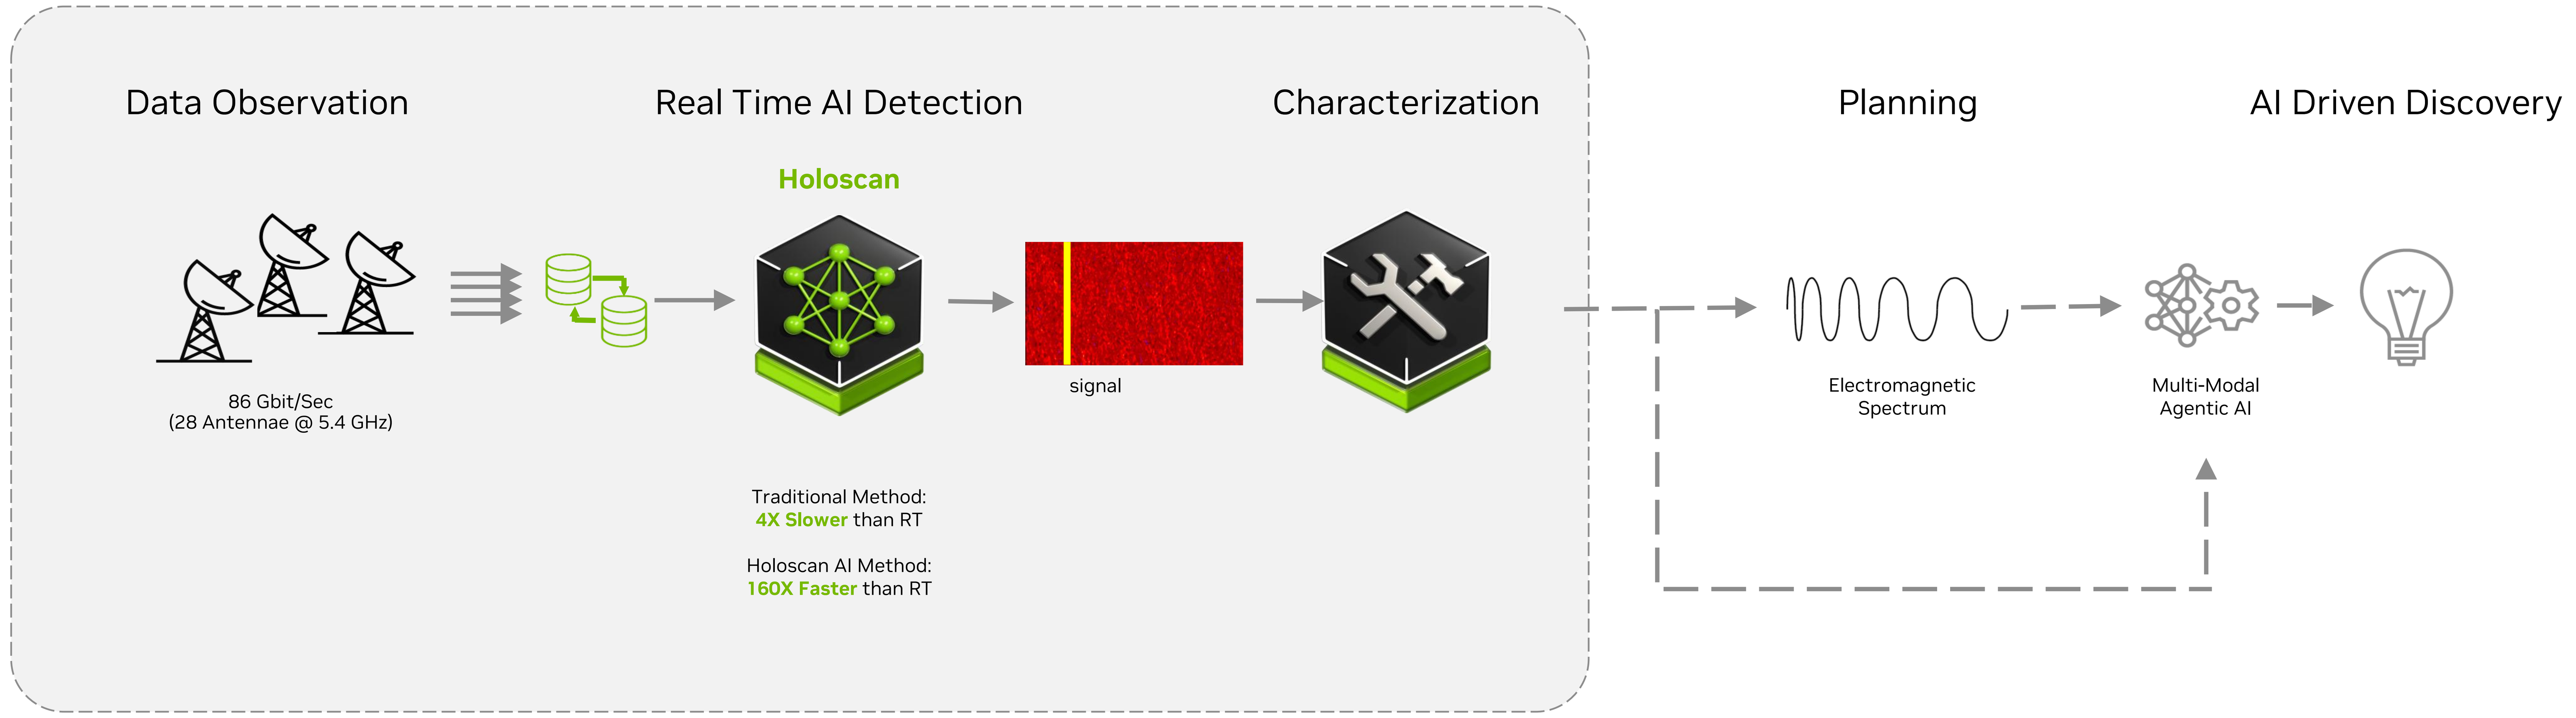

Holoscan | CUDA-X | TensorRT | DOCA

Computational Engineering

For illustration purposes only

Quantum

Quantum-Accelerated Supercomputing

Supercomputers are the foundation of Quantum R&D

Simulation

- Quantum computers are small and error-prone -> simulation is an essential tool
- **Today:** Powerful simulators enable algorithm and application R&D – new approaches (e.g. tensor networks)
- **Future:** Digital twins of quantum computers for design and architecture optimization

HPC Quantum Integration

- Useful quantum computing will be hybrid
- **Today:** Enable domain scientists to start developing for QPUs, enable quantum researchers to use accelerated computing
- **Future:** quantum computers will integrate tightly with supercomputers as accelerators and be co-programmed

AI for Quantum

- Error correction, calibration, control, compilation are challenging computationally, real-time compute often needed
- Accelerated computing and AI can solve these problems
- **Today:** Enable AI research for all of the above
- **Future:** Hybrid Quantum+AI supercomputer with low-latency link

NVIDIA Quantum

Powering the Global Quantum Computing Community

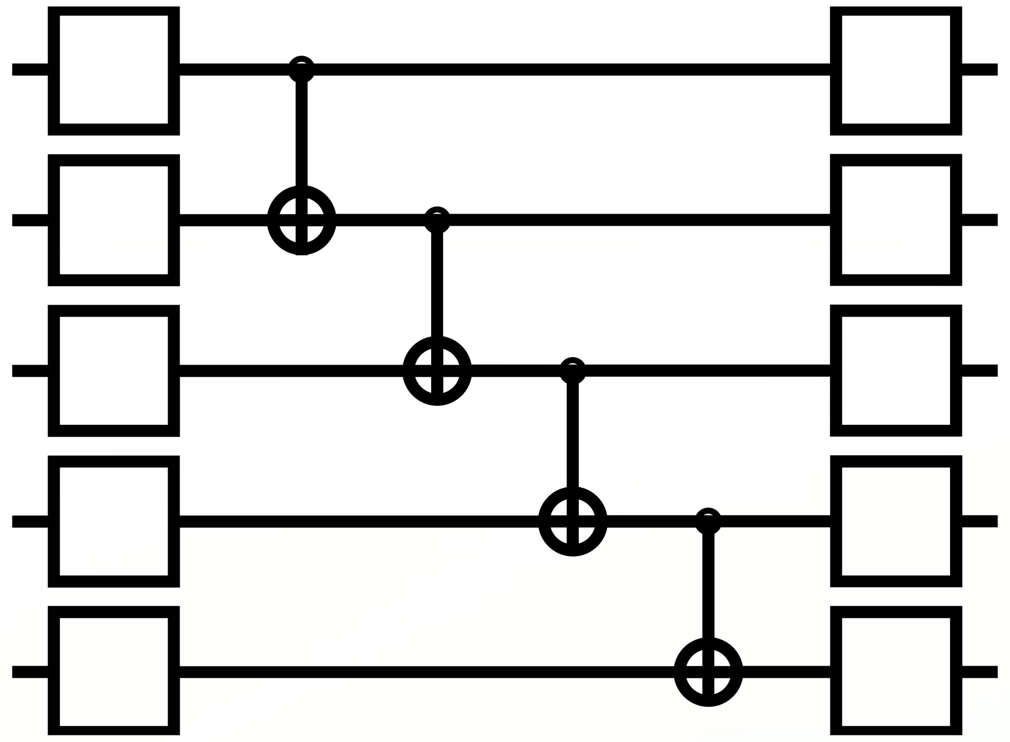

Simulation

Algorithm Design, Resource Estimation, QPU Design

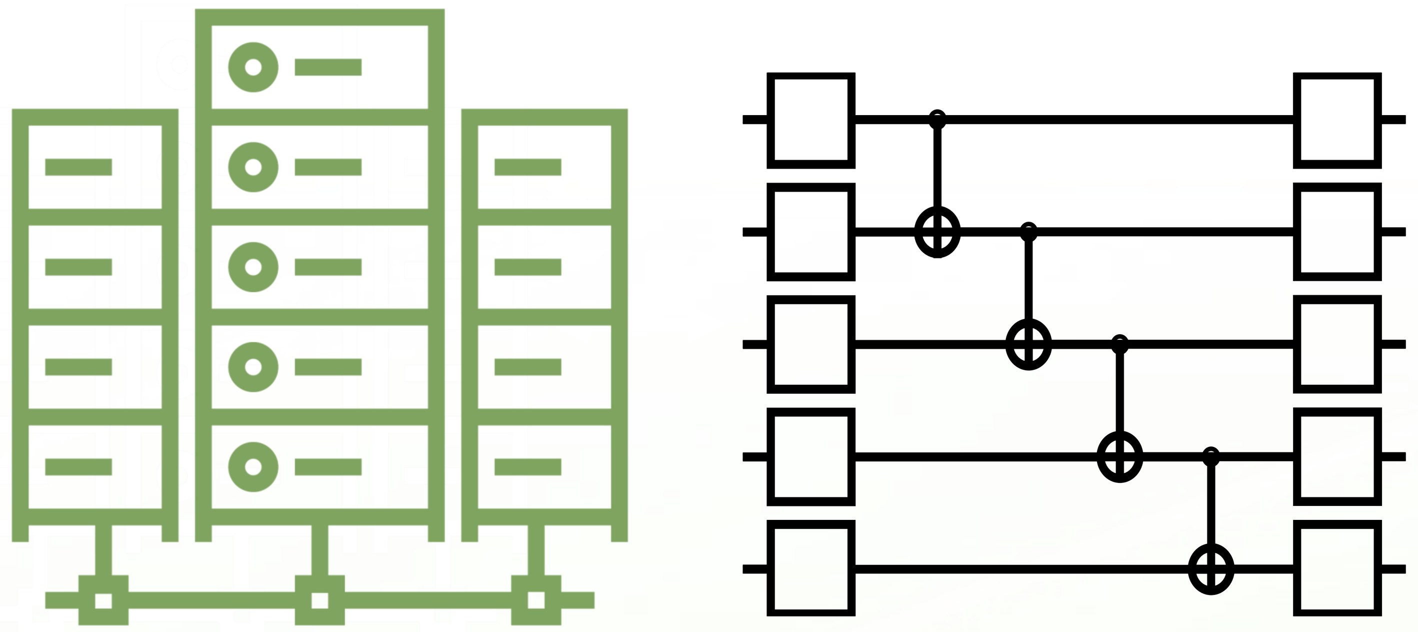

HPC Quantum Integration

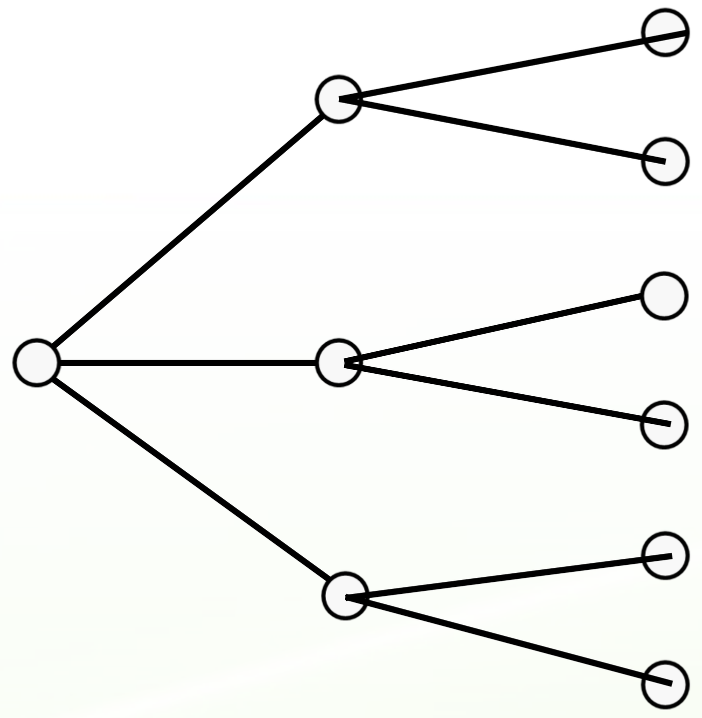

AI for Quantum

QEC, Calibration, Algorithms

CUDA-Q

Libraries | Programming Model | Tools | Infrastructure

cuQuantum

Quantum Simulation

DGX-Q

Quantum Integrated Computing

GPU Supercomputing

The GPT Quantum Eigensolver

The University of Toronto, St. Jude Children’s Research Hospital, and NVIDIA

- Developed a novel Generative Pre-Trained Transformer-based (GPT) method for computing the ground-state energy of molecules of interest
- The first demonstration of a GPT-generated quantum workload
- A powerful example of leveraging AI to accelerate quantum computing
- Executed using CUDA-Q on A100 GPUs on Perlmutter
- Opens the door to a wide variety of novel Generative Quantum Algorithms (GQAs) for drug discovery, materials science, and environmental challenges

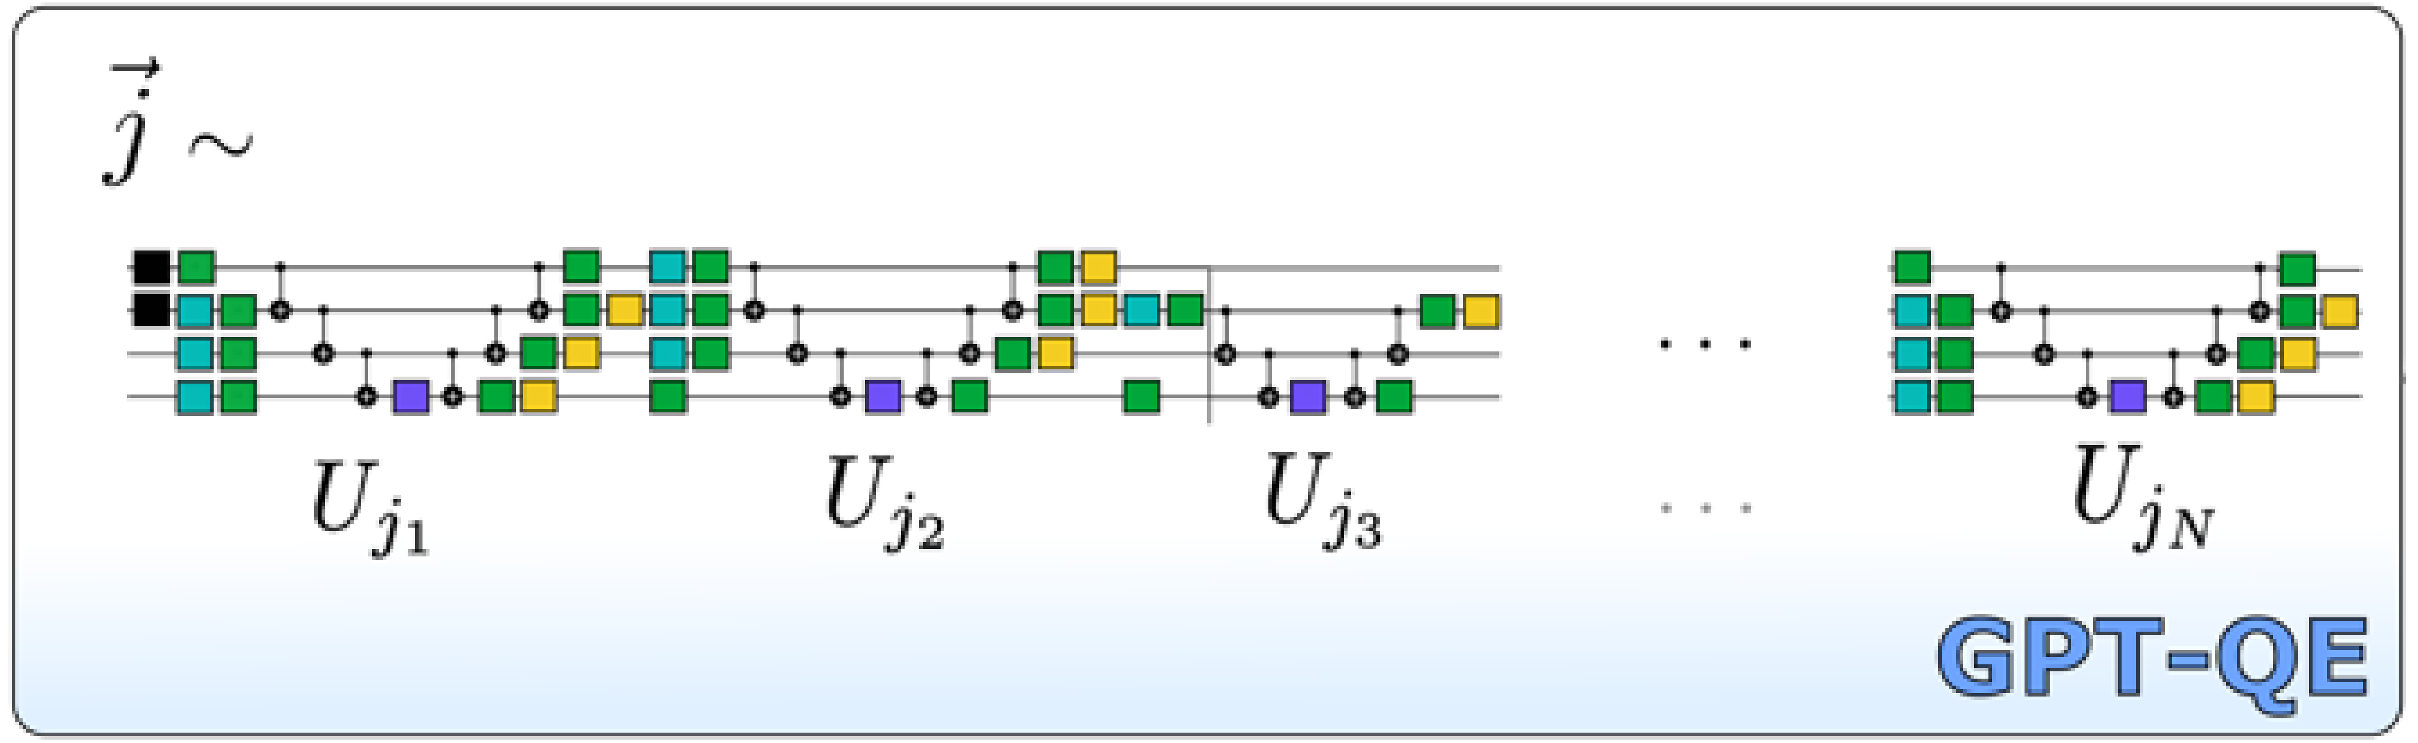

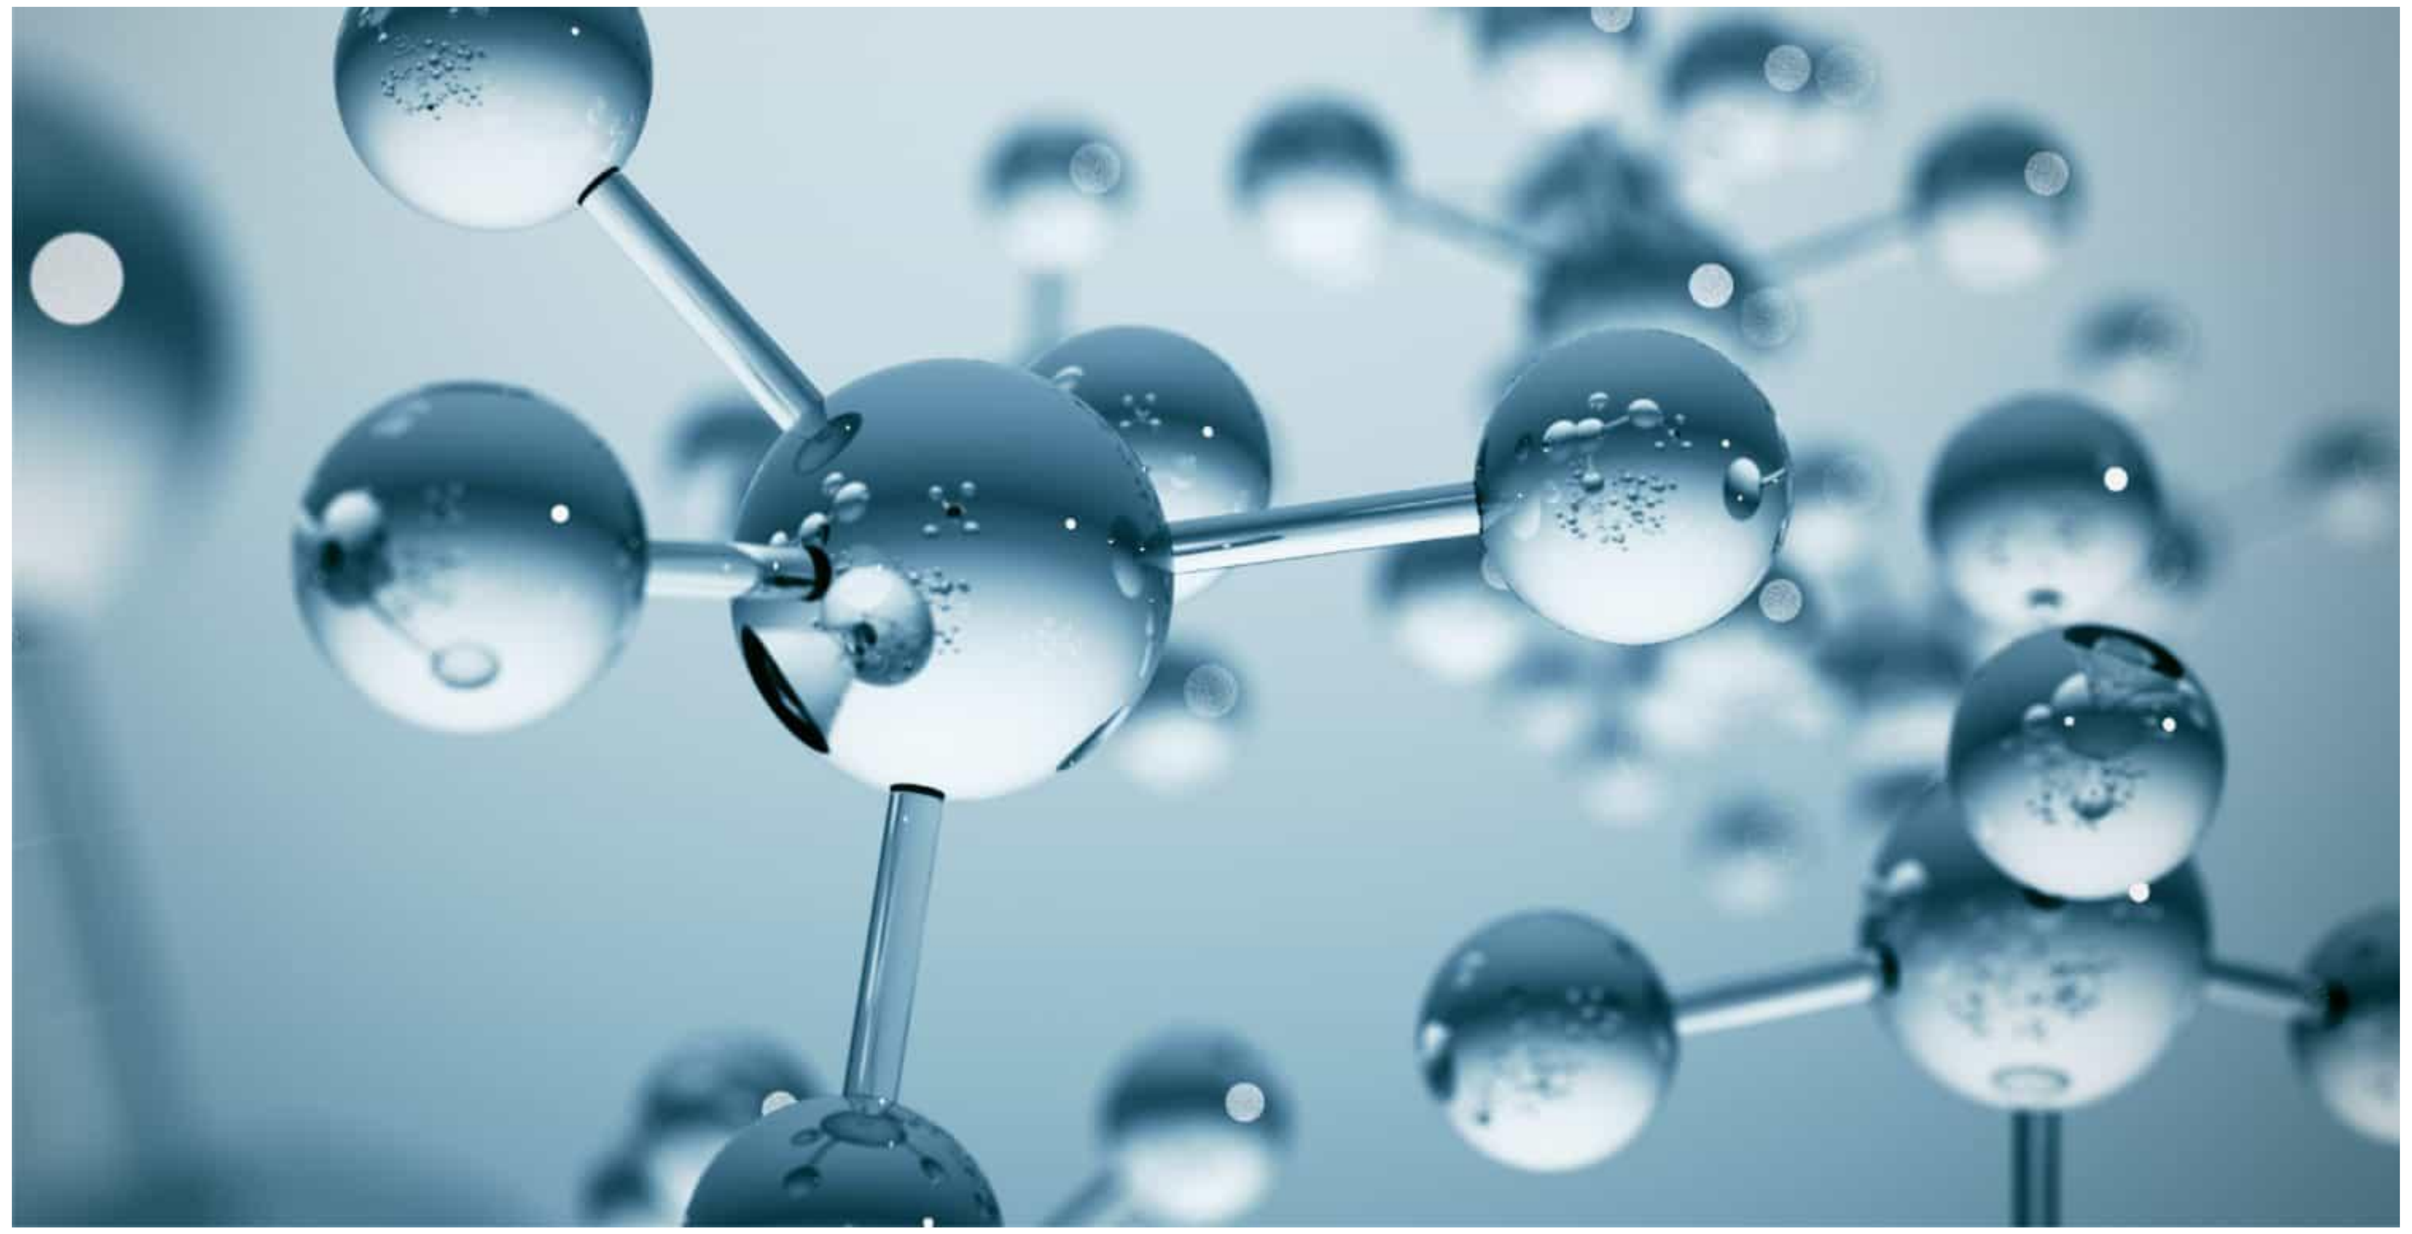

https://arxiv.org/pdf/2401.09253.pdf

UNIVERSITY OF TORONTO

St. Jude Children’s® Research Hospital
Finding cures. Saving children.

Announcing CUDA-Q Academic

Quantum Curriculum and Workforce Development Partnerships

- Coursework designed in partnership with academic institutions
- Access workshops from anywhere with GPU-acceleration in the cloud
- Active learning in distributed, quantum-accelerated computing and CUDA-Q
- Interactive Jupyter notebooks feature lectures, explanations, exercises, and assessments

Arizona State University | Carnegie Mellon University | Dartmouth Engineering | Quantum Center | Fordham University | Mälardalen University

Northwestern | Pittsburgh Supercomputing Center | Universitat Politècnica de València | Princeton University | Purdue University | Robert Morris

TUM Chair for Design Automation | Technion Israel Institute of Technology | UC Davis | University of Cambridge Physical Computation Laboratory | The University of Chicago | University of Florida

University of Illinois Urbana-Champaign | University of Pittsburgh | Universitat Politècnica de Catalunya BarcelonaTech

Drug Discovery

NVIDIA BioNeMo Platform

NVIDIA BioNeMo Blueprints Are Reference Workflows for Drug Discovery

Available on build.nvidia.com

Multimodal PDF Data Extraction for Enterprise RAG

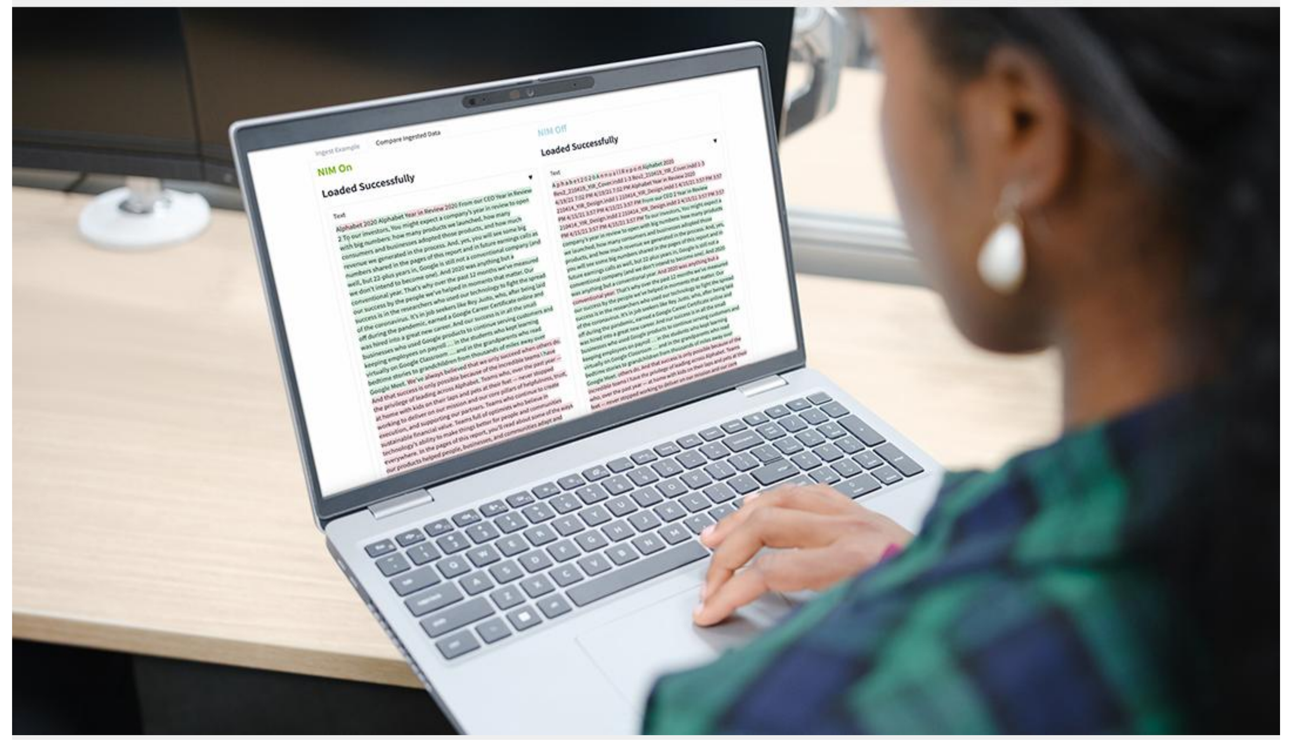

Generative Protein Binder Design for Drug Discovery

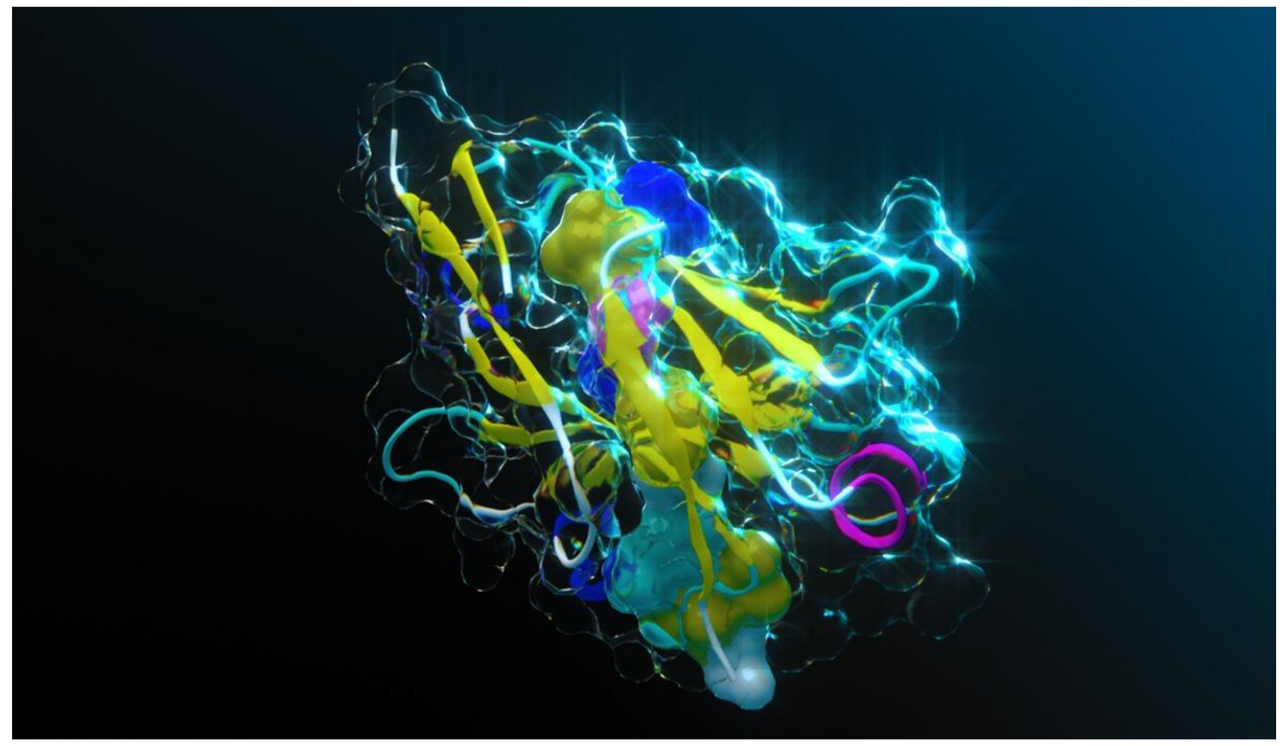

Generative Virtual Screening for Drug Discovery

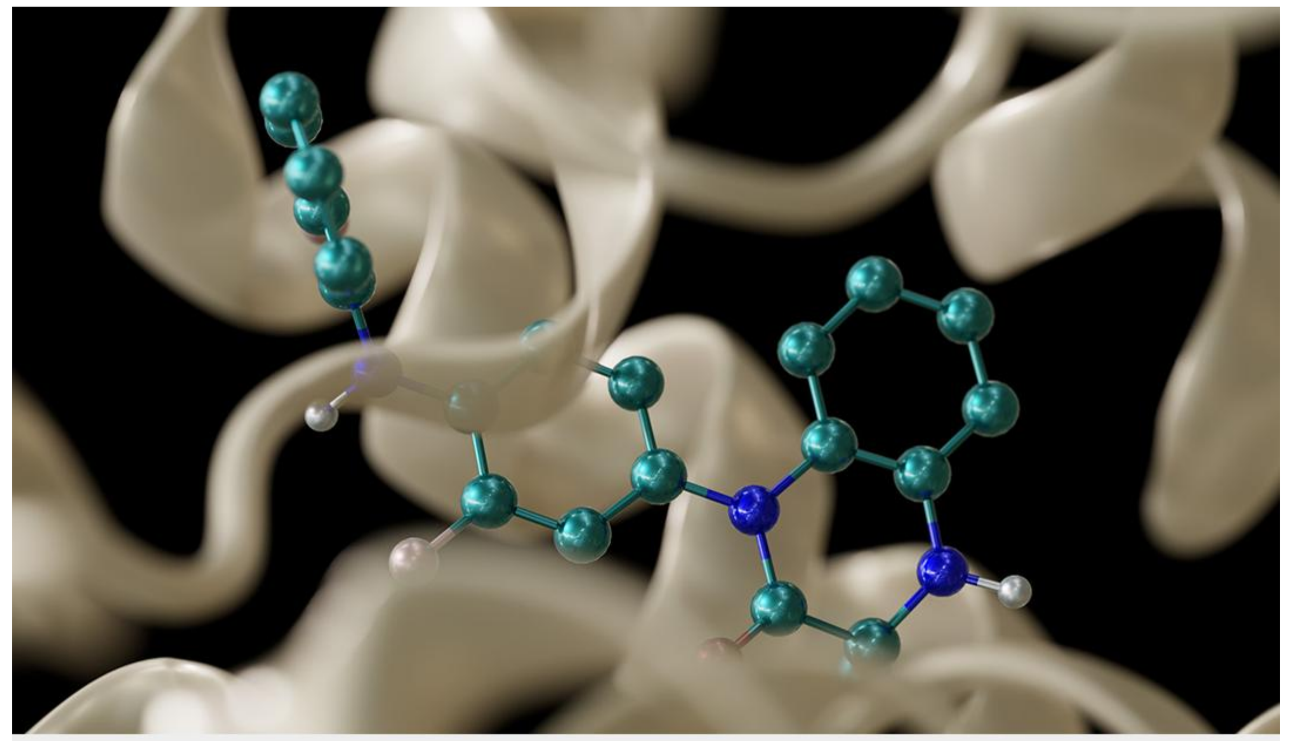

...

monthly release

NVIDIA NIM Agent Blueprint

Example Application	Sample Data	Reference Code	Architecture	Customization Tools	Orchestration Tools
Interactive experience that can be easily replicated	Public data for workflow testing	Reference code for constructing workflows	Reference architecture including API definitions, NIM, and more	Customize and evaluate models	Deploy and manage workflow microservices

Generative Protein Binder Design

Models: MSA, AlphaFold2, RFdiffusion, ProteinMPNN, Alphafold2-Multimer

Use Cases

- Design novel proteins to bind to a target of interest and/or perform a desired function.

Value Prop:

- Efficiently explore vast protein design space for structures with precise functional features
- SOTA speed and scalability over other approaches
- Identify high-affinity binders from a smaller set of designs

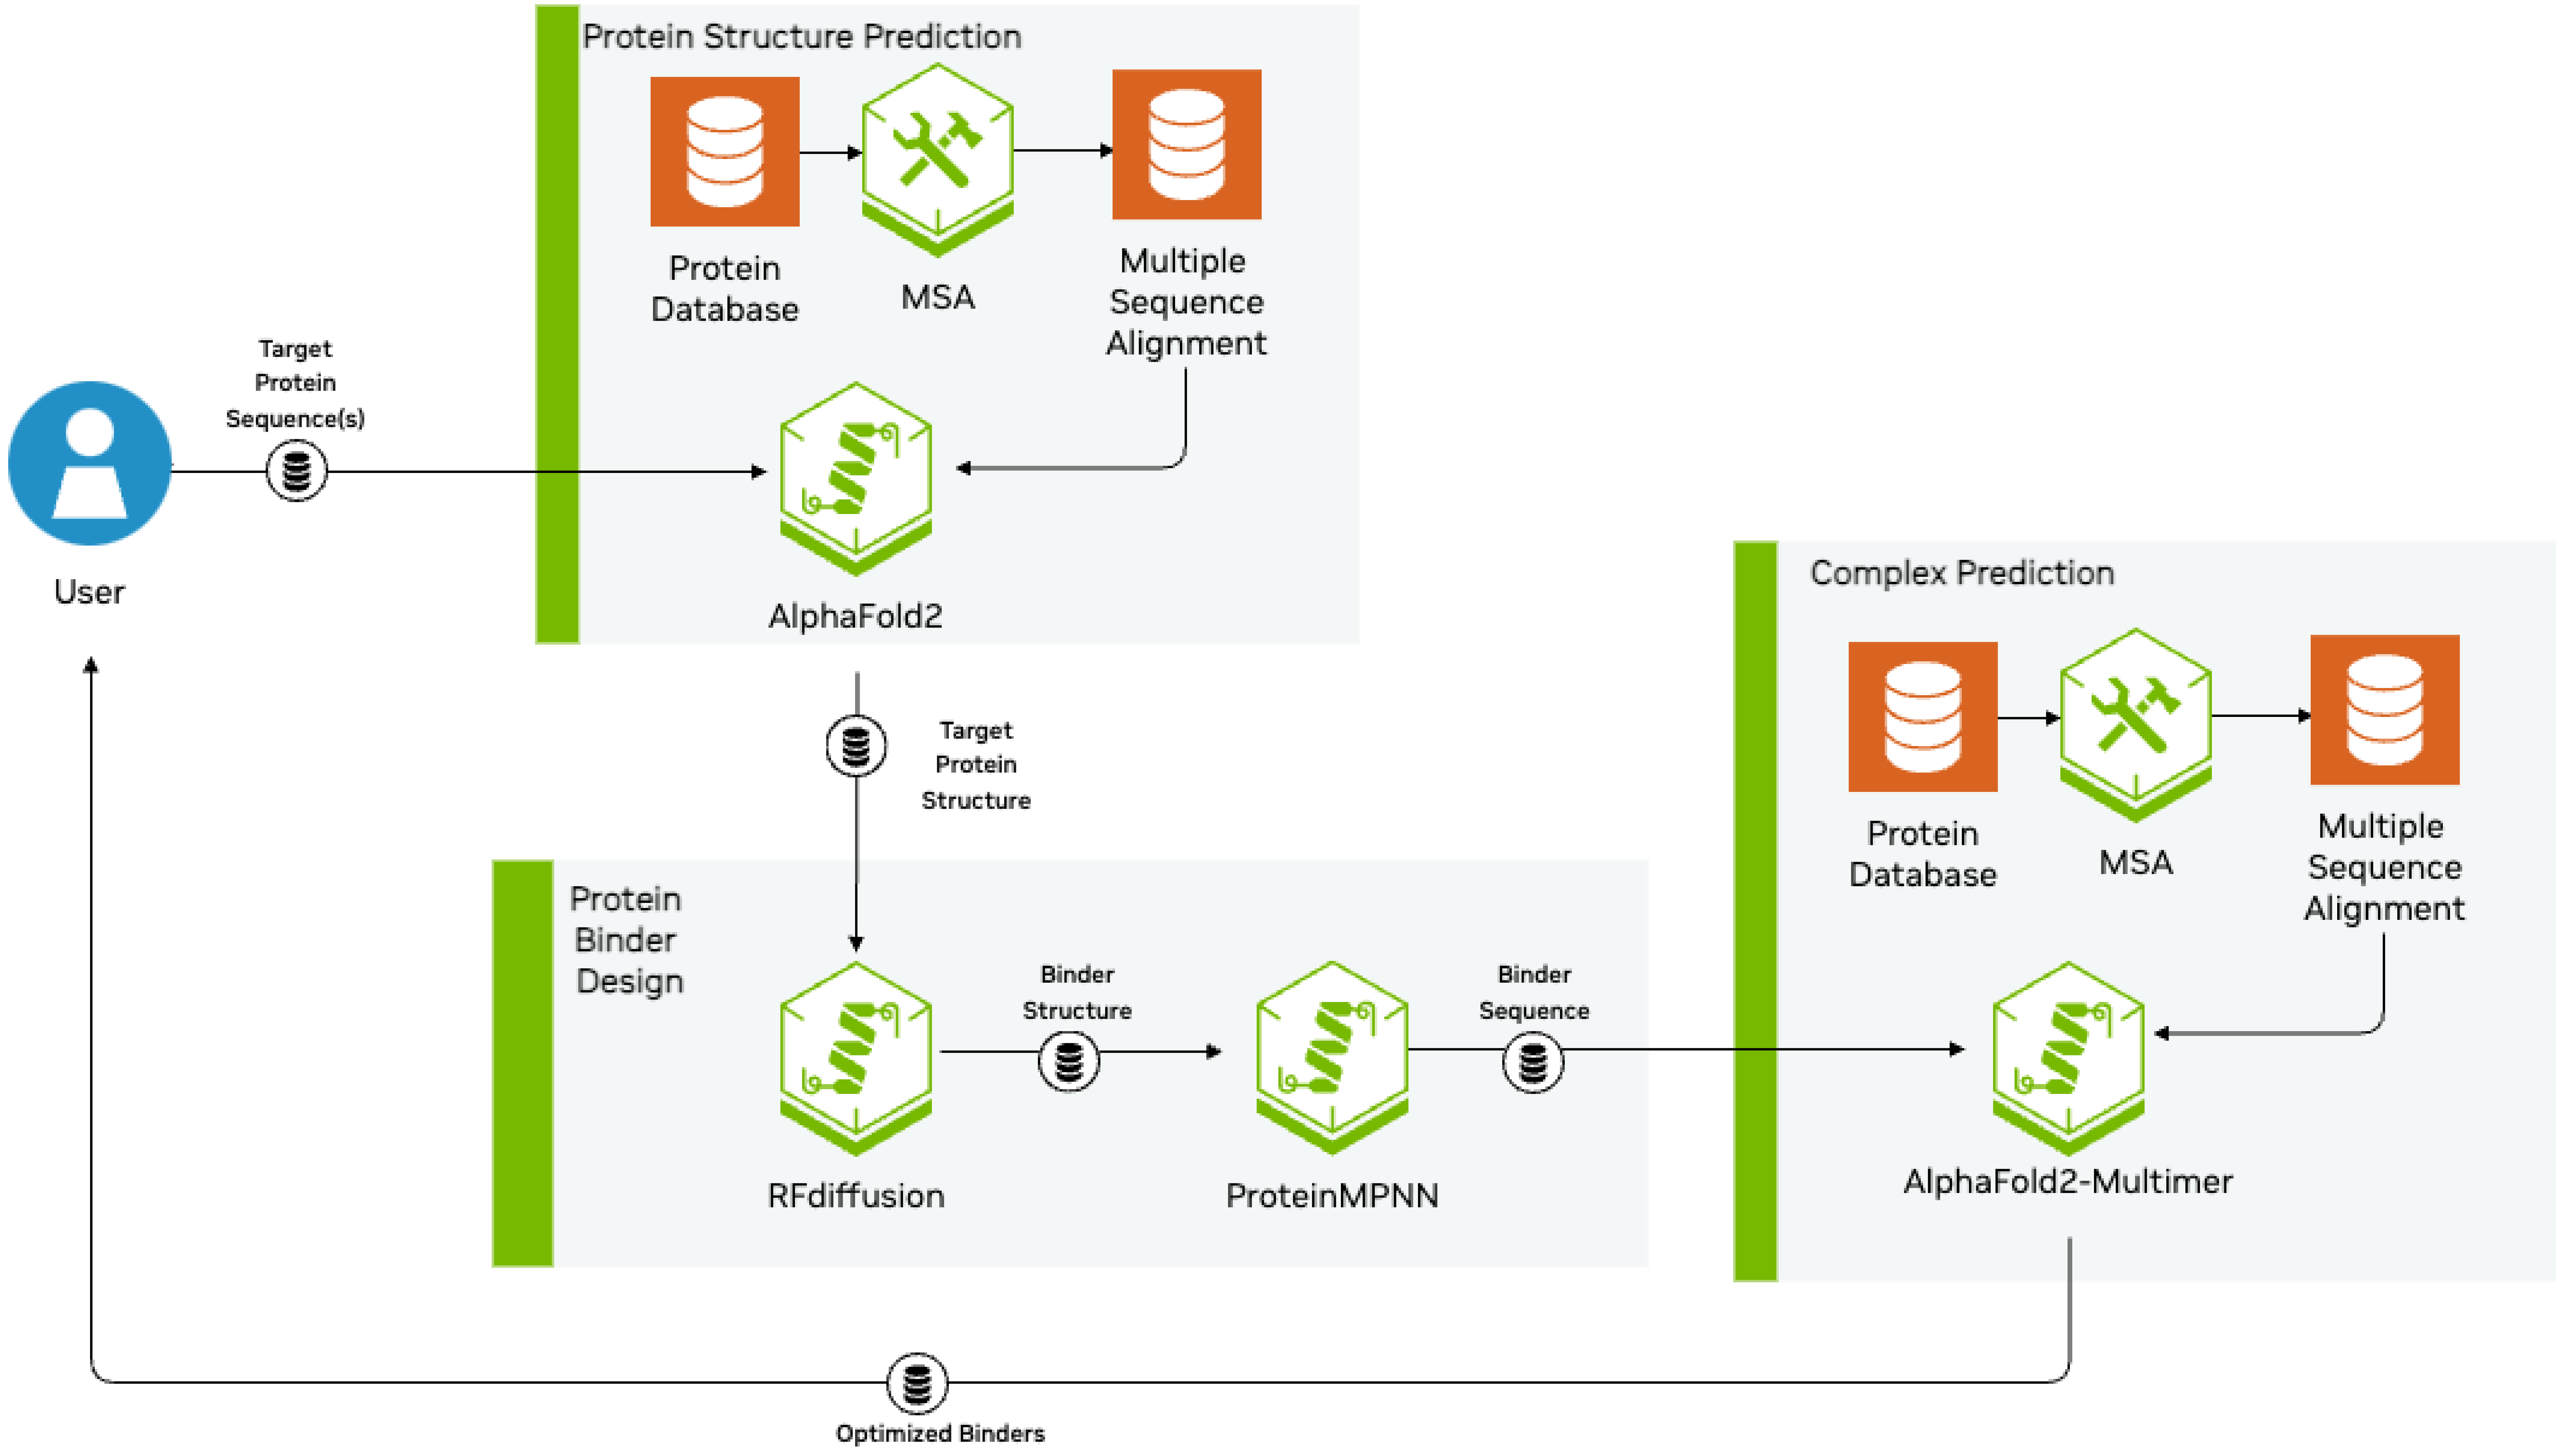

Working @ NVIDIA

Your Future Starts Here

NVIDIA is hiring! Our work in AI and the metaverse is transforming the world's largest industries and profoundly impacting society. Help shape what comes next.

- For internships and new college graduate roles **www.nvidia.com/university**
- General hiring areas **https://nvda.ws/nvur-jd**
- In the meantime, follow us on **LinkedIn**, **Instagram** and **NVIDIA Blog** to stay connected!

Employee Benefits—Setting Us Apart (Canada)

Interns

- Health Benefits (six months or longer)
- Emergency PTO
- Housing Stipend (in-office)
- Ergonomic & Internet Stipend (remote)
- Holiday Pay & Free Days
- Equity (ESPP)
- Guidance Resources EAP
- Discounts & Services
- Intern Events

New College Graduate

- Health Benefits
- PTO & Emergency PTO
- Holiday Pay & Free Days
- Equity (RSUs & ESPP)
- Ergonomic & Internet Stipend (remote)
- Parental Leave
- Guidance Resources EAP
- Tuition Reimbursement
- RRSP Matching
- Discounts & Services

Find Your Perfect Fit

Search for your area of interest at **nvda.ws/nvur-jd** and apply at **www.nvidia.com/university**

Hardware

- ASIC Design
- Verification
- Physical Design/VLSI
- Mixed Signal Design
- Digital Circuit Design

Architecture

- Computer Architecture
- Deep Learning Computer Architecture

Systems Software

- System Software
- Graphics Systems Software
- Compiler
- Firmware & Embedded Software
- Software Security

Software

- Development Tools
- Cloud
- Tools Infrastructure
- Machine Learning Operations

Autonomous Vehicles and Robotics

- Autonomous Vehicles
- Robotics

Deep Learning

- Deep Learning Applications & Algorithms
- Deep Learning Frameworks & Libraries

Research

- Research (PhD)
- Applied Research (BS, MS, PhD)

Business Operations

- Business Operations (MBA)

NVIDIA