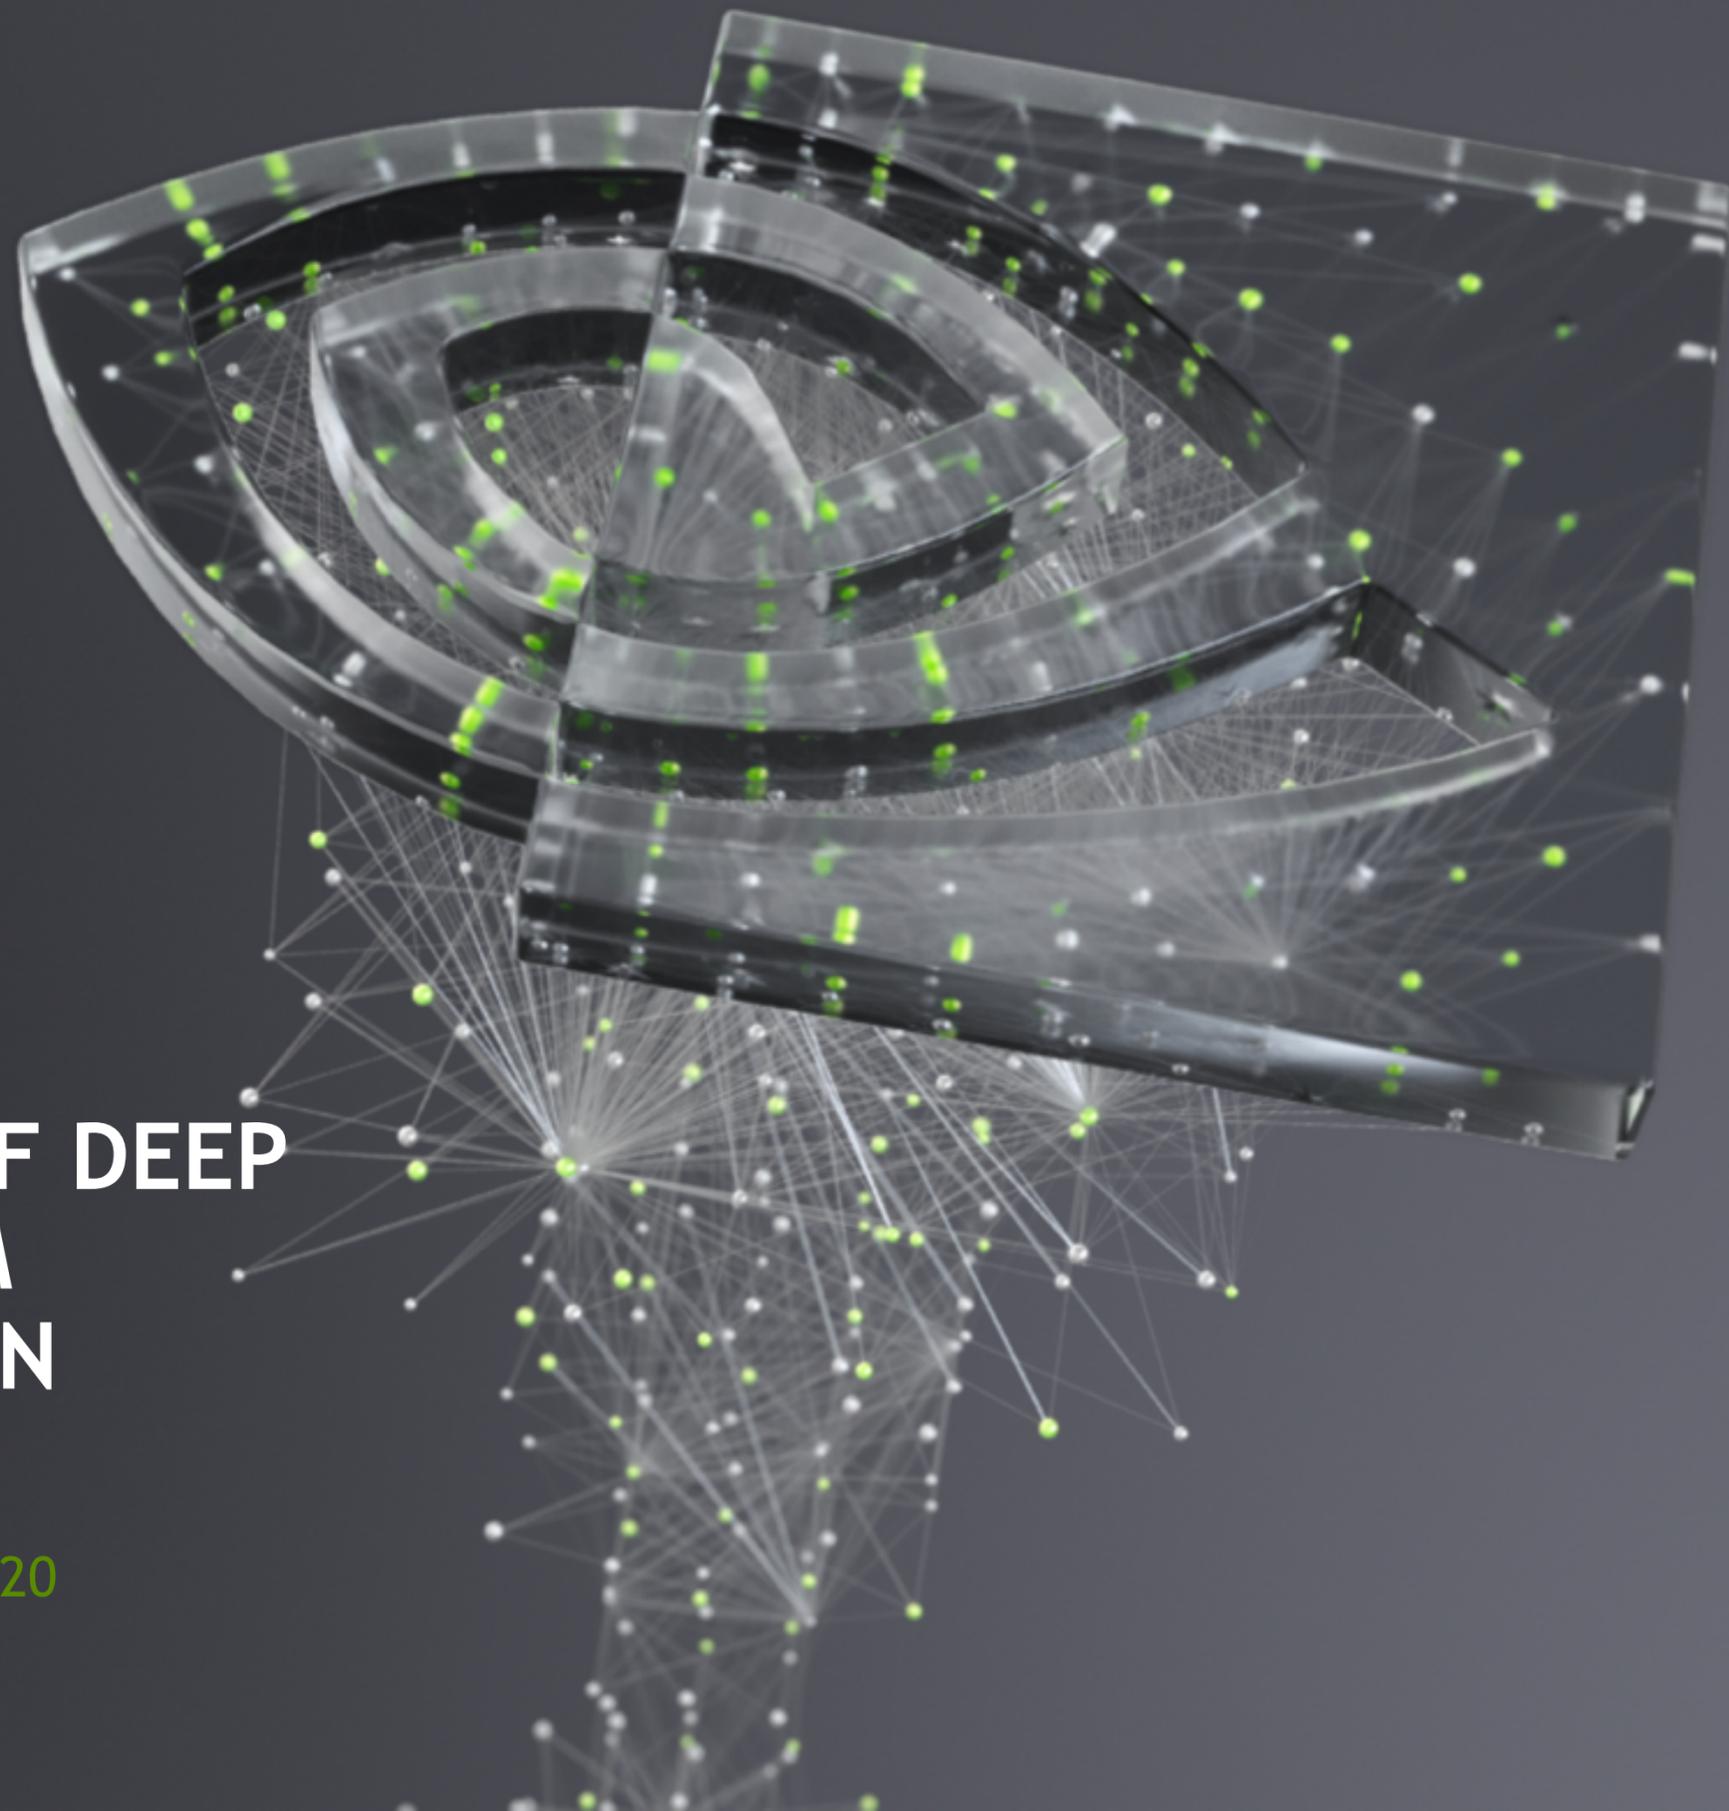




# EXPLORING THE FRONTIERS OF DEEP LEARNING FOR EARTH SYSTEM OBSERVATION AND PREDICTION

Dr. David M. Hall, Senior Data Scientist, NVIDIA

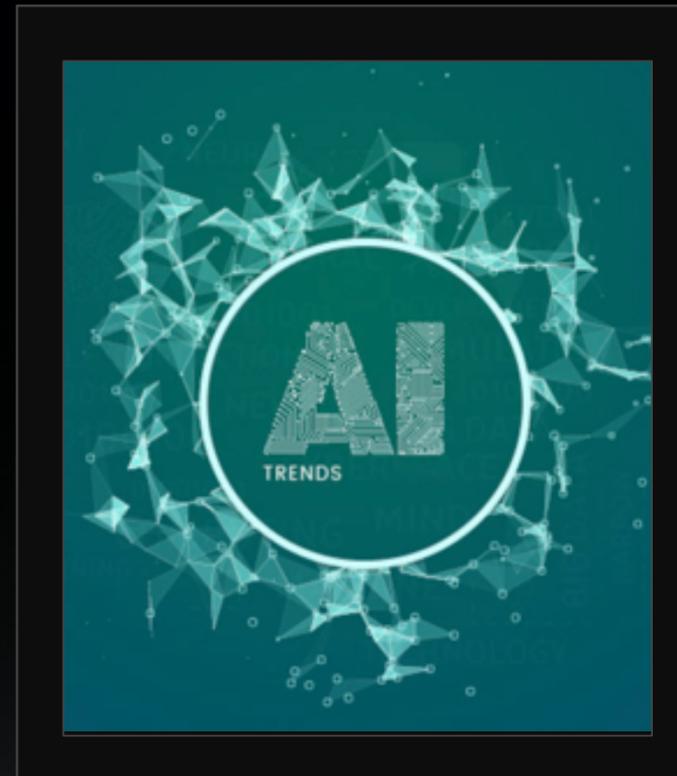
ECMWF-ESA Workshop on Machine Learning, Oct, 2020



# THE FRONTIERS OF DEEP LEARNING



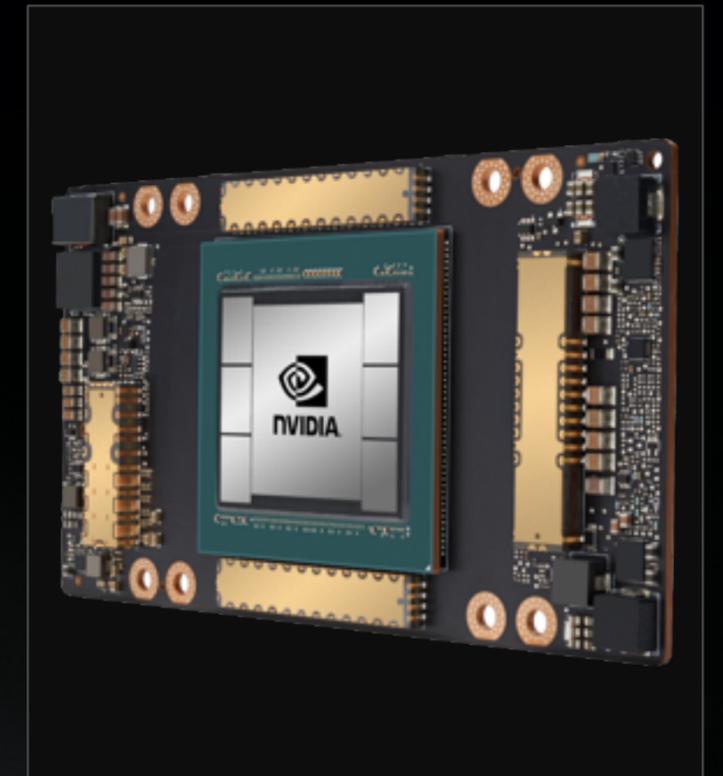
Applications



AI Trends



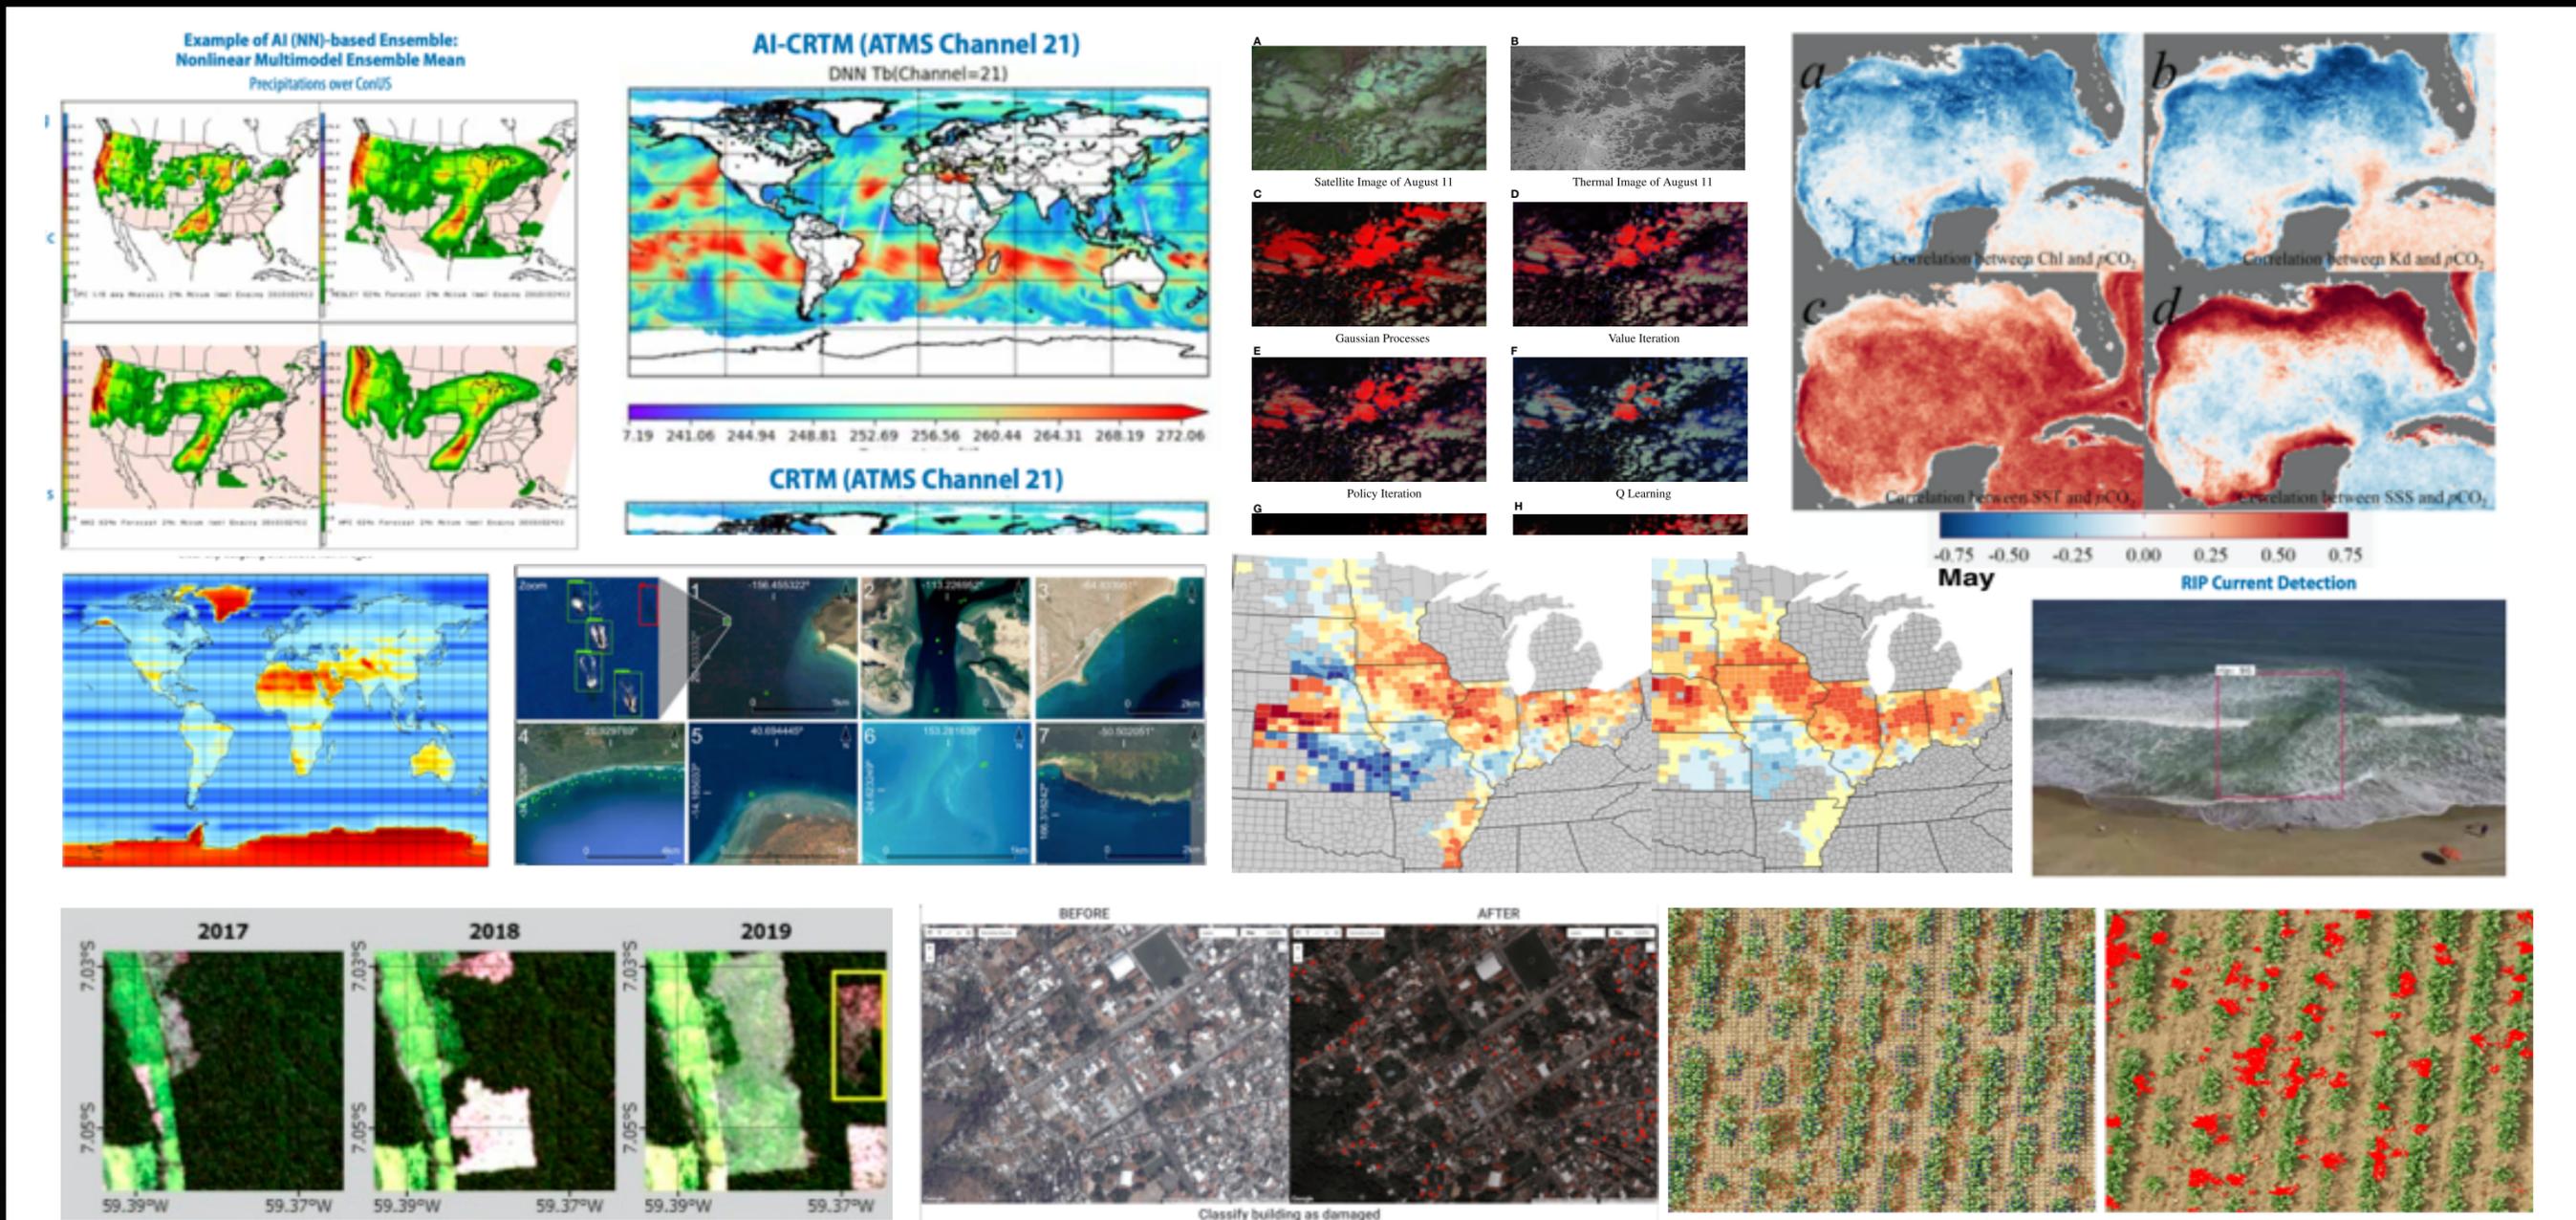
Scientific Challenges



Hardware and Tools

# RAPID ADOPTION

Deep learning is being rapidly adopted by the Earth System Science community



# NSF AI INSTITUTES

Seven new Artificial Intelligence institutes, including one for Weather and Climate

## NSF AI Institute for Research on Trustworthy AI in Weather, Climate and Coastal Oceanography



[NSF announcement](#)

# NOAA CENTER FOR AI

Official strategy and dedicated center focusing on AI

## NOAA Artificial Intelligence Strategy

Analytics for Next-Generation Earth Science

### ***NOAA Artificial Intelligence Strategy***

**Goal 1: Establish an efficient organizational structure and processes to advance AI across NOAA.**

**Goal 2: Advance AI research and innovation in support of NOAA's mission.**

**Goal 3: Accelerate the transition of AI research to applications.**

**Goal 4: Strengthen and expand AI partnerships.**

**Goal 5: Promote AI proficiency in the workforce.**

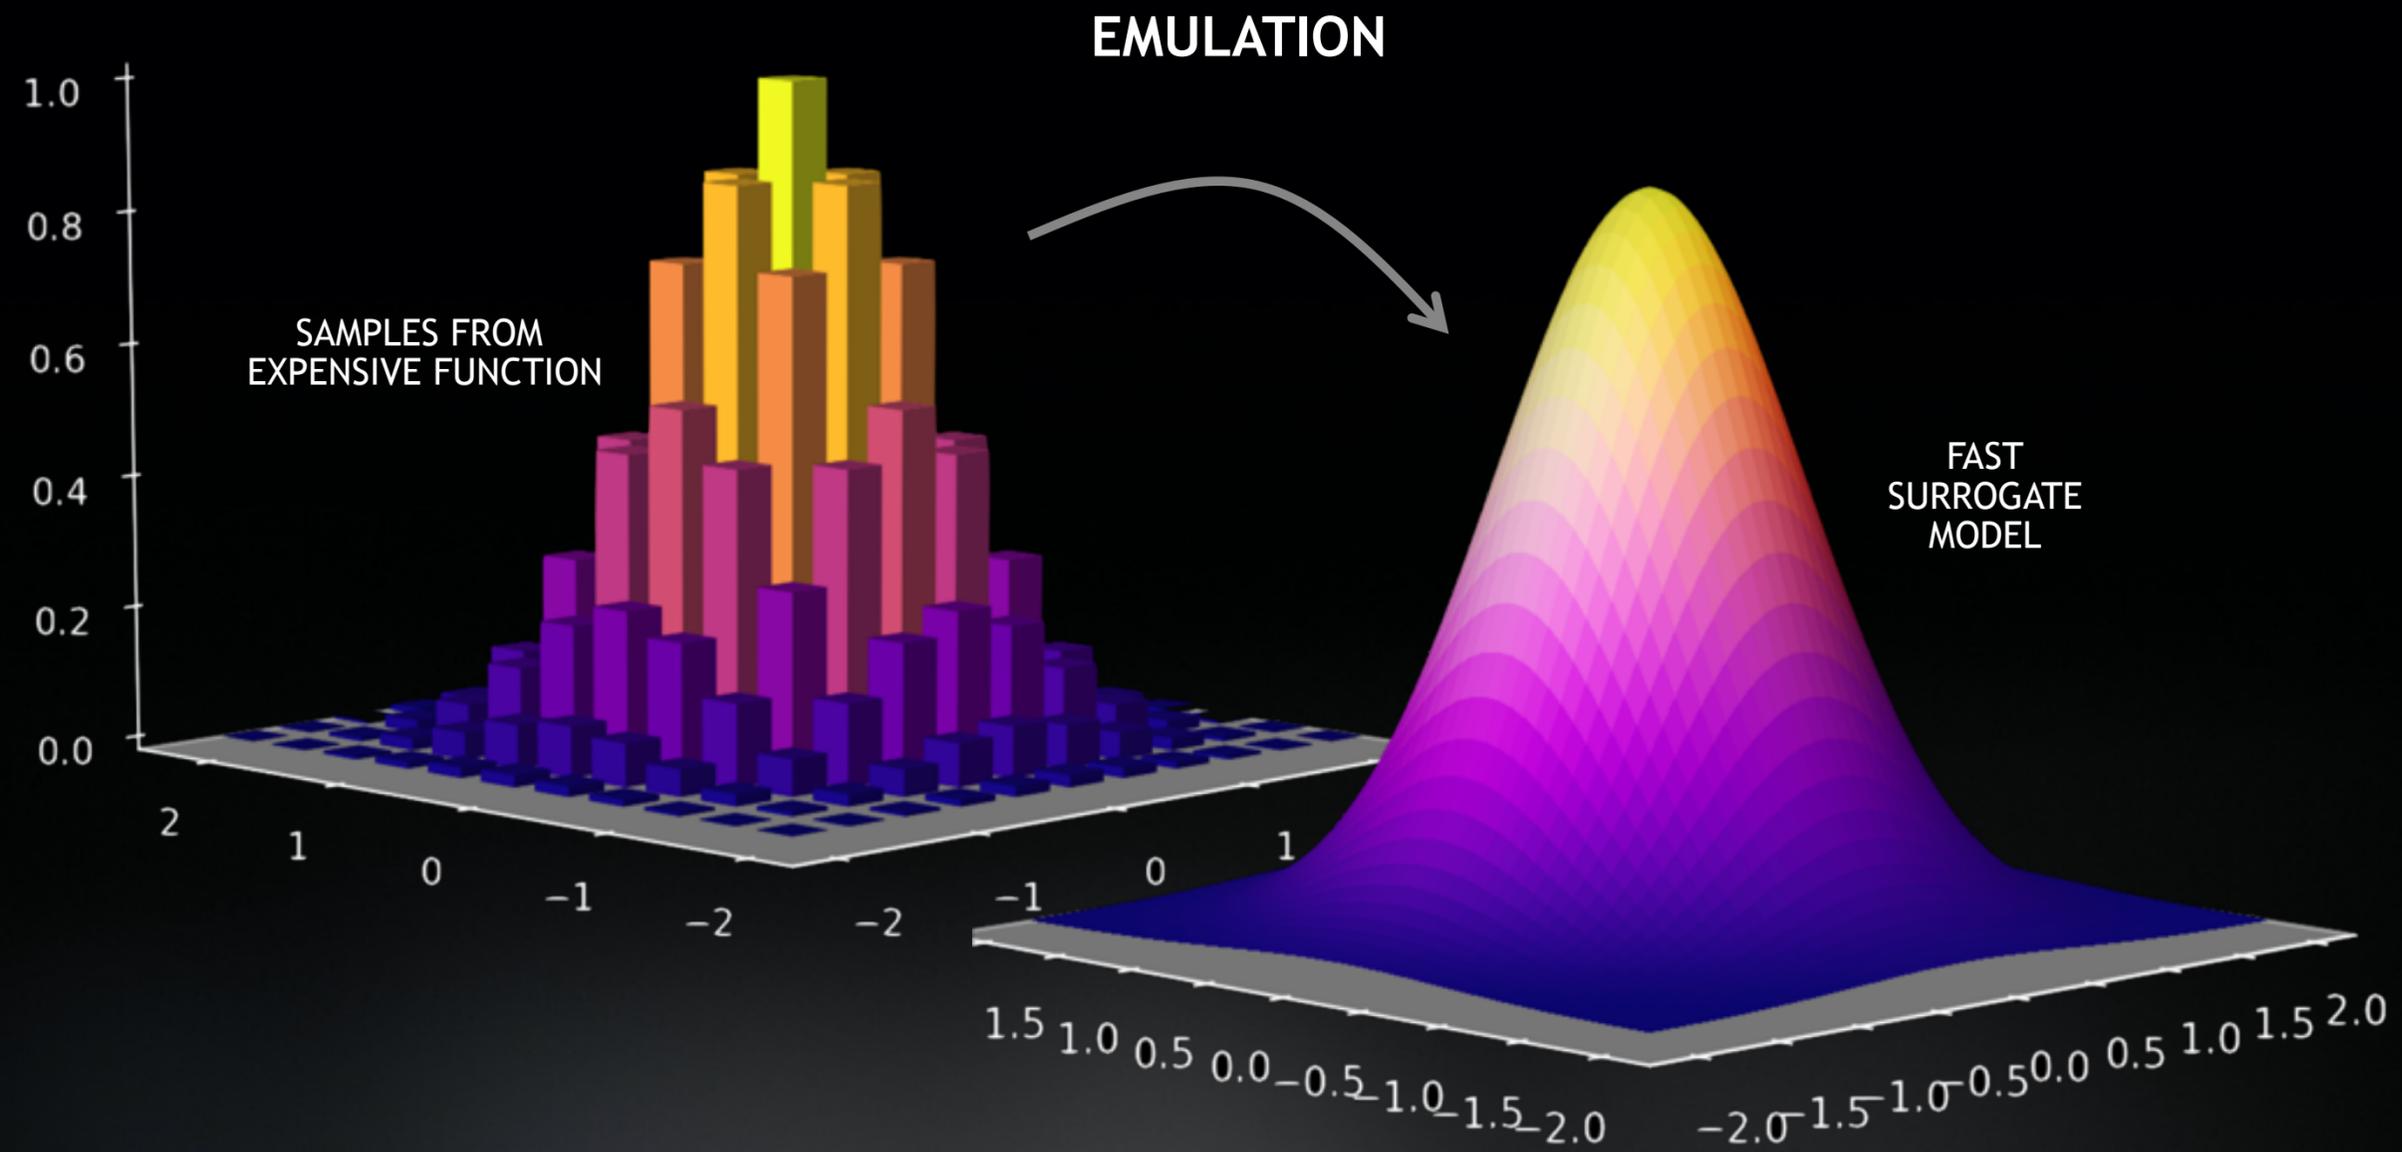
<https://nrc.noaa.gov/LinkClick.aspx?fileticket=0I2p2-Gu3rA%3D&tabid=91&portalid=0>



# APPLICATIONS IN EARTH SYSTEM SCIENCE

# ACCELERATED PHYSICS

Using surrogate models to speed up existing code



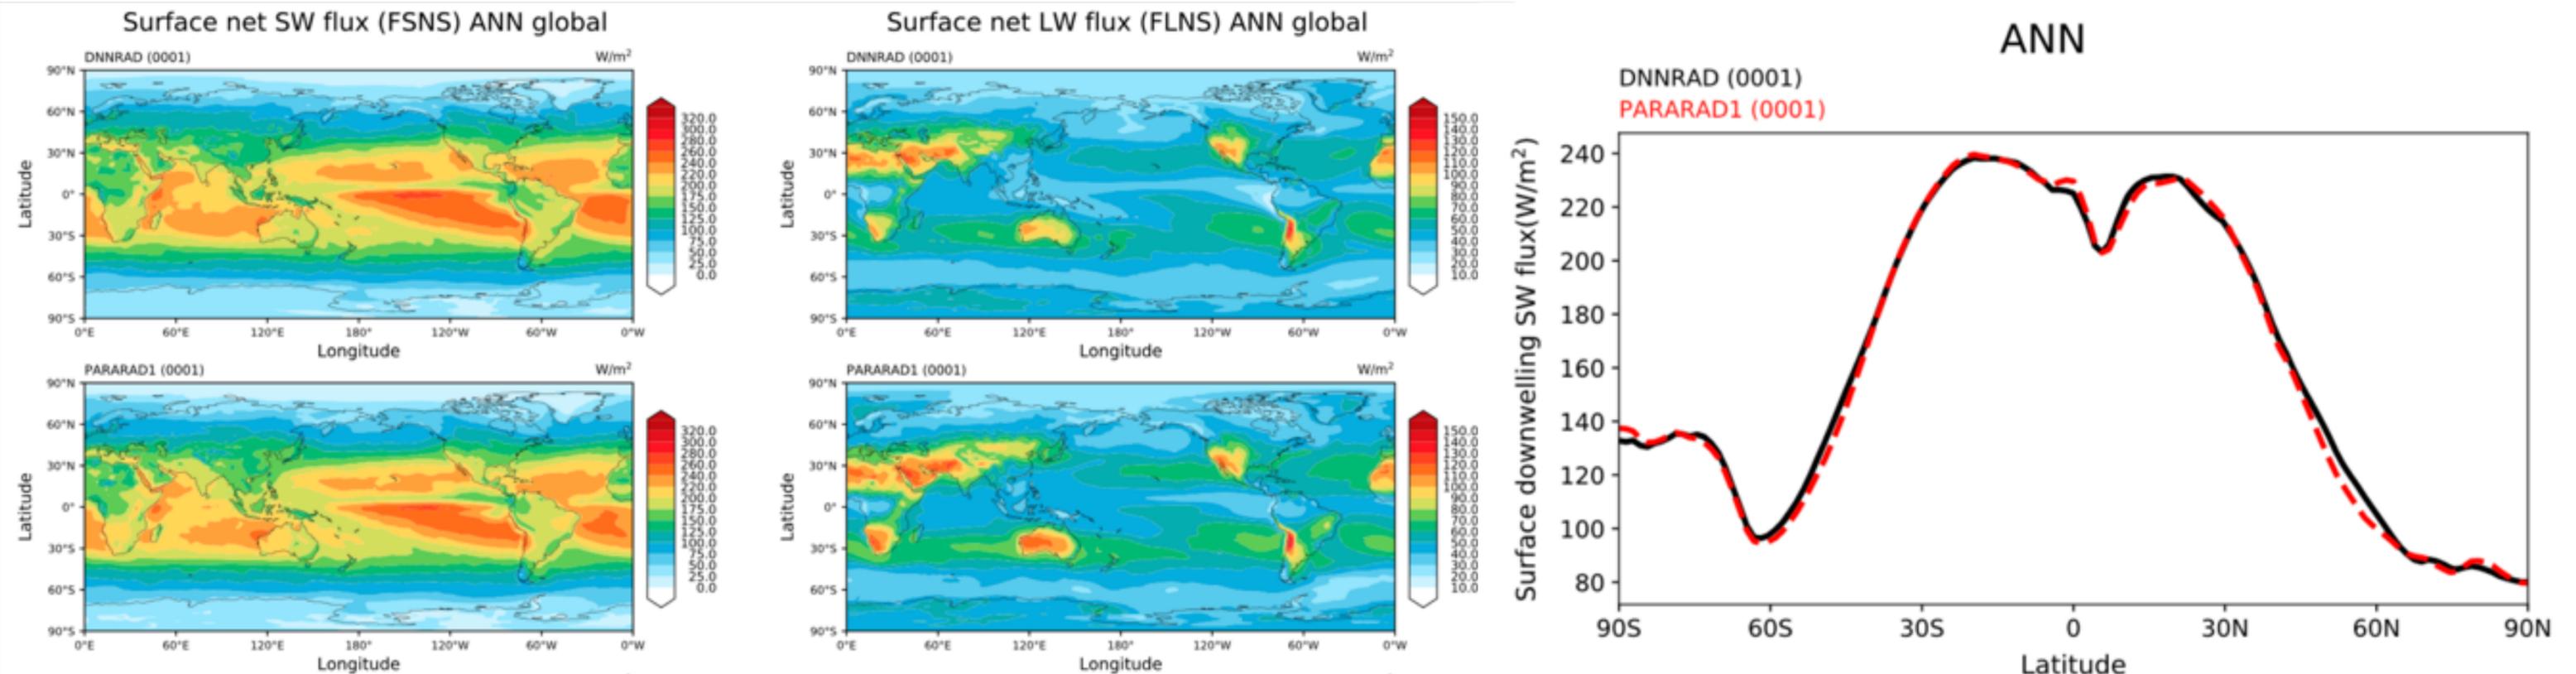
# ACCELERATED PHYSICS PARAMETERIZATIONS

## Emulation of E3SM Super-parameterized SW and LW Radiation

Using Deep Neural Networks as Cost-Effective Surrogate Models for Super-Parameterized E3SM Radiative Transfer

Anikesh Pal<sup>1</sup>, Salil Mahajan<sup>2</sup>, and Matthew R. Norman<sup>1</sup>

<sup>1</sup>National Center for Computational Sciences, Oak Ridge National Laboratory, Oak Ridge, TN, USA, <sup>2</sup>Computational Earth Sciences Oak Ridge National Laboratory, Oak Ridge, TN, USA



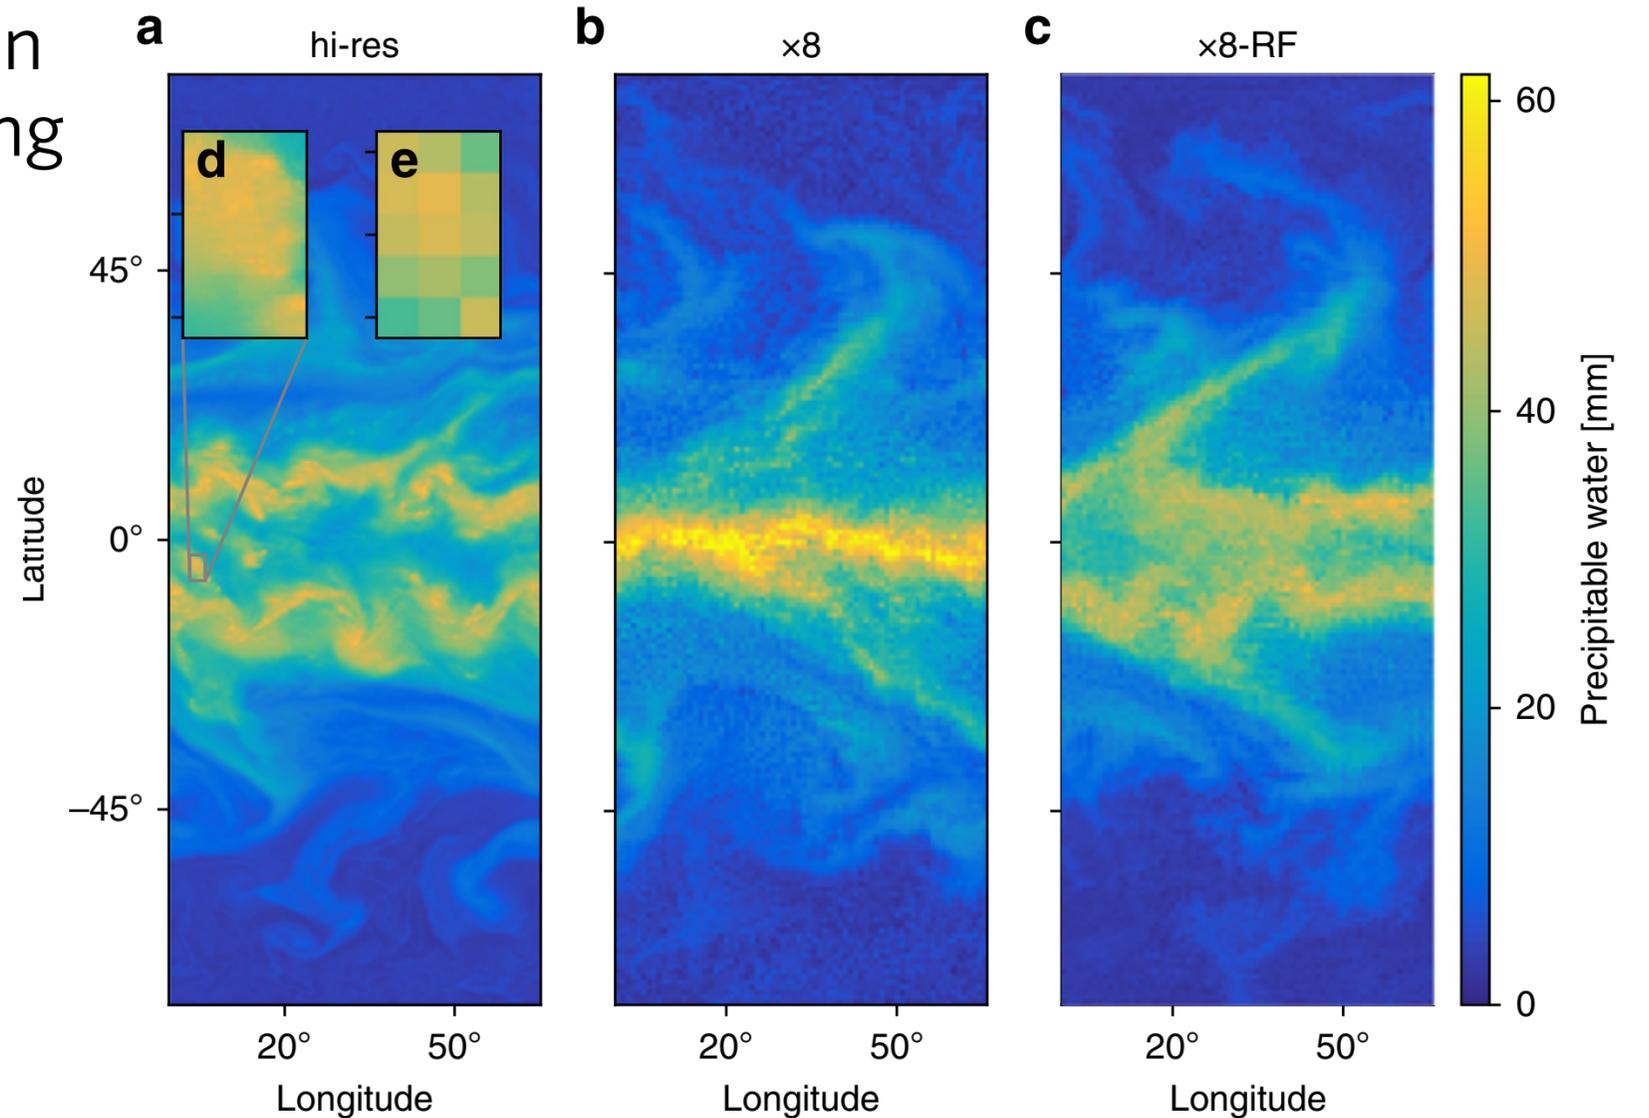
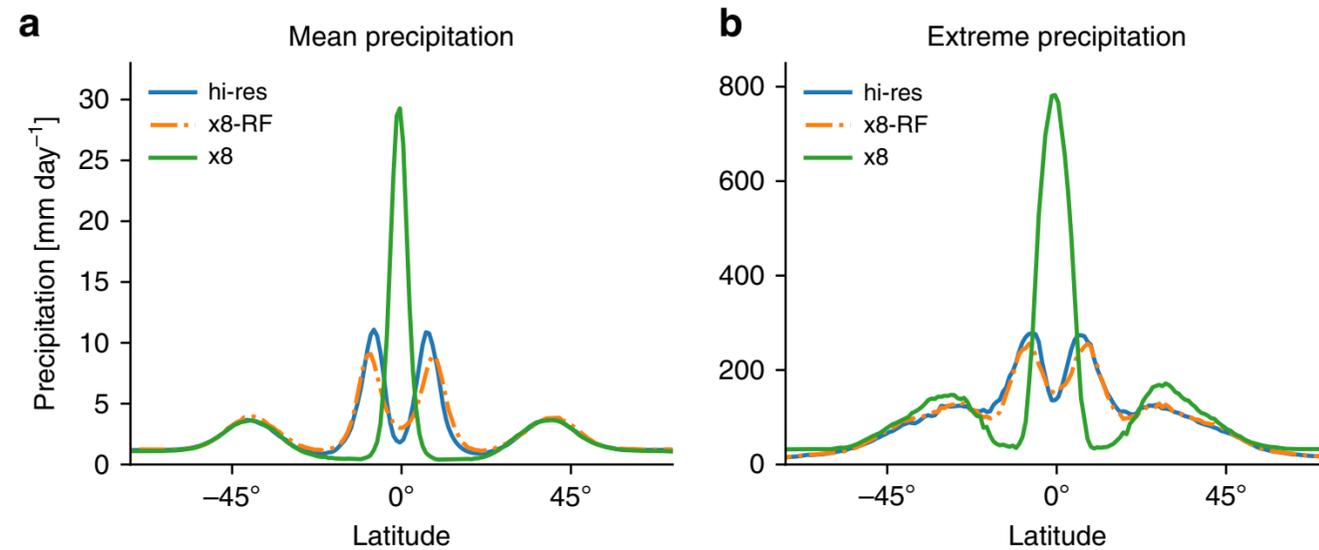
8-10x Speedup in SW and LW Radiative Transfer Calculations

# IMPROVED PHYSICS PARAMETERIZATIONS

Building more accurate physics parameterizations

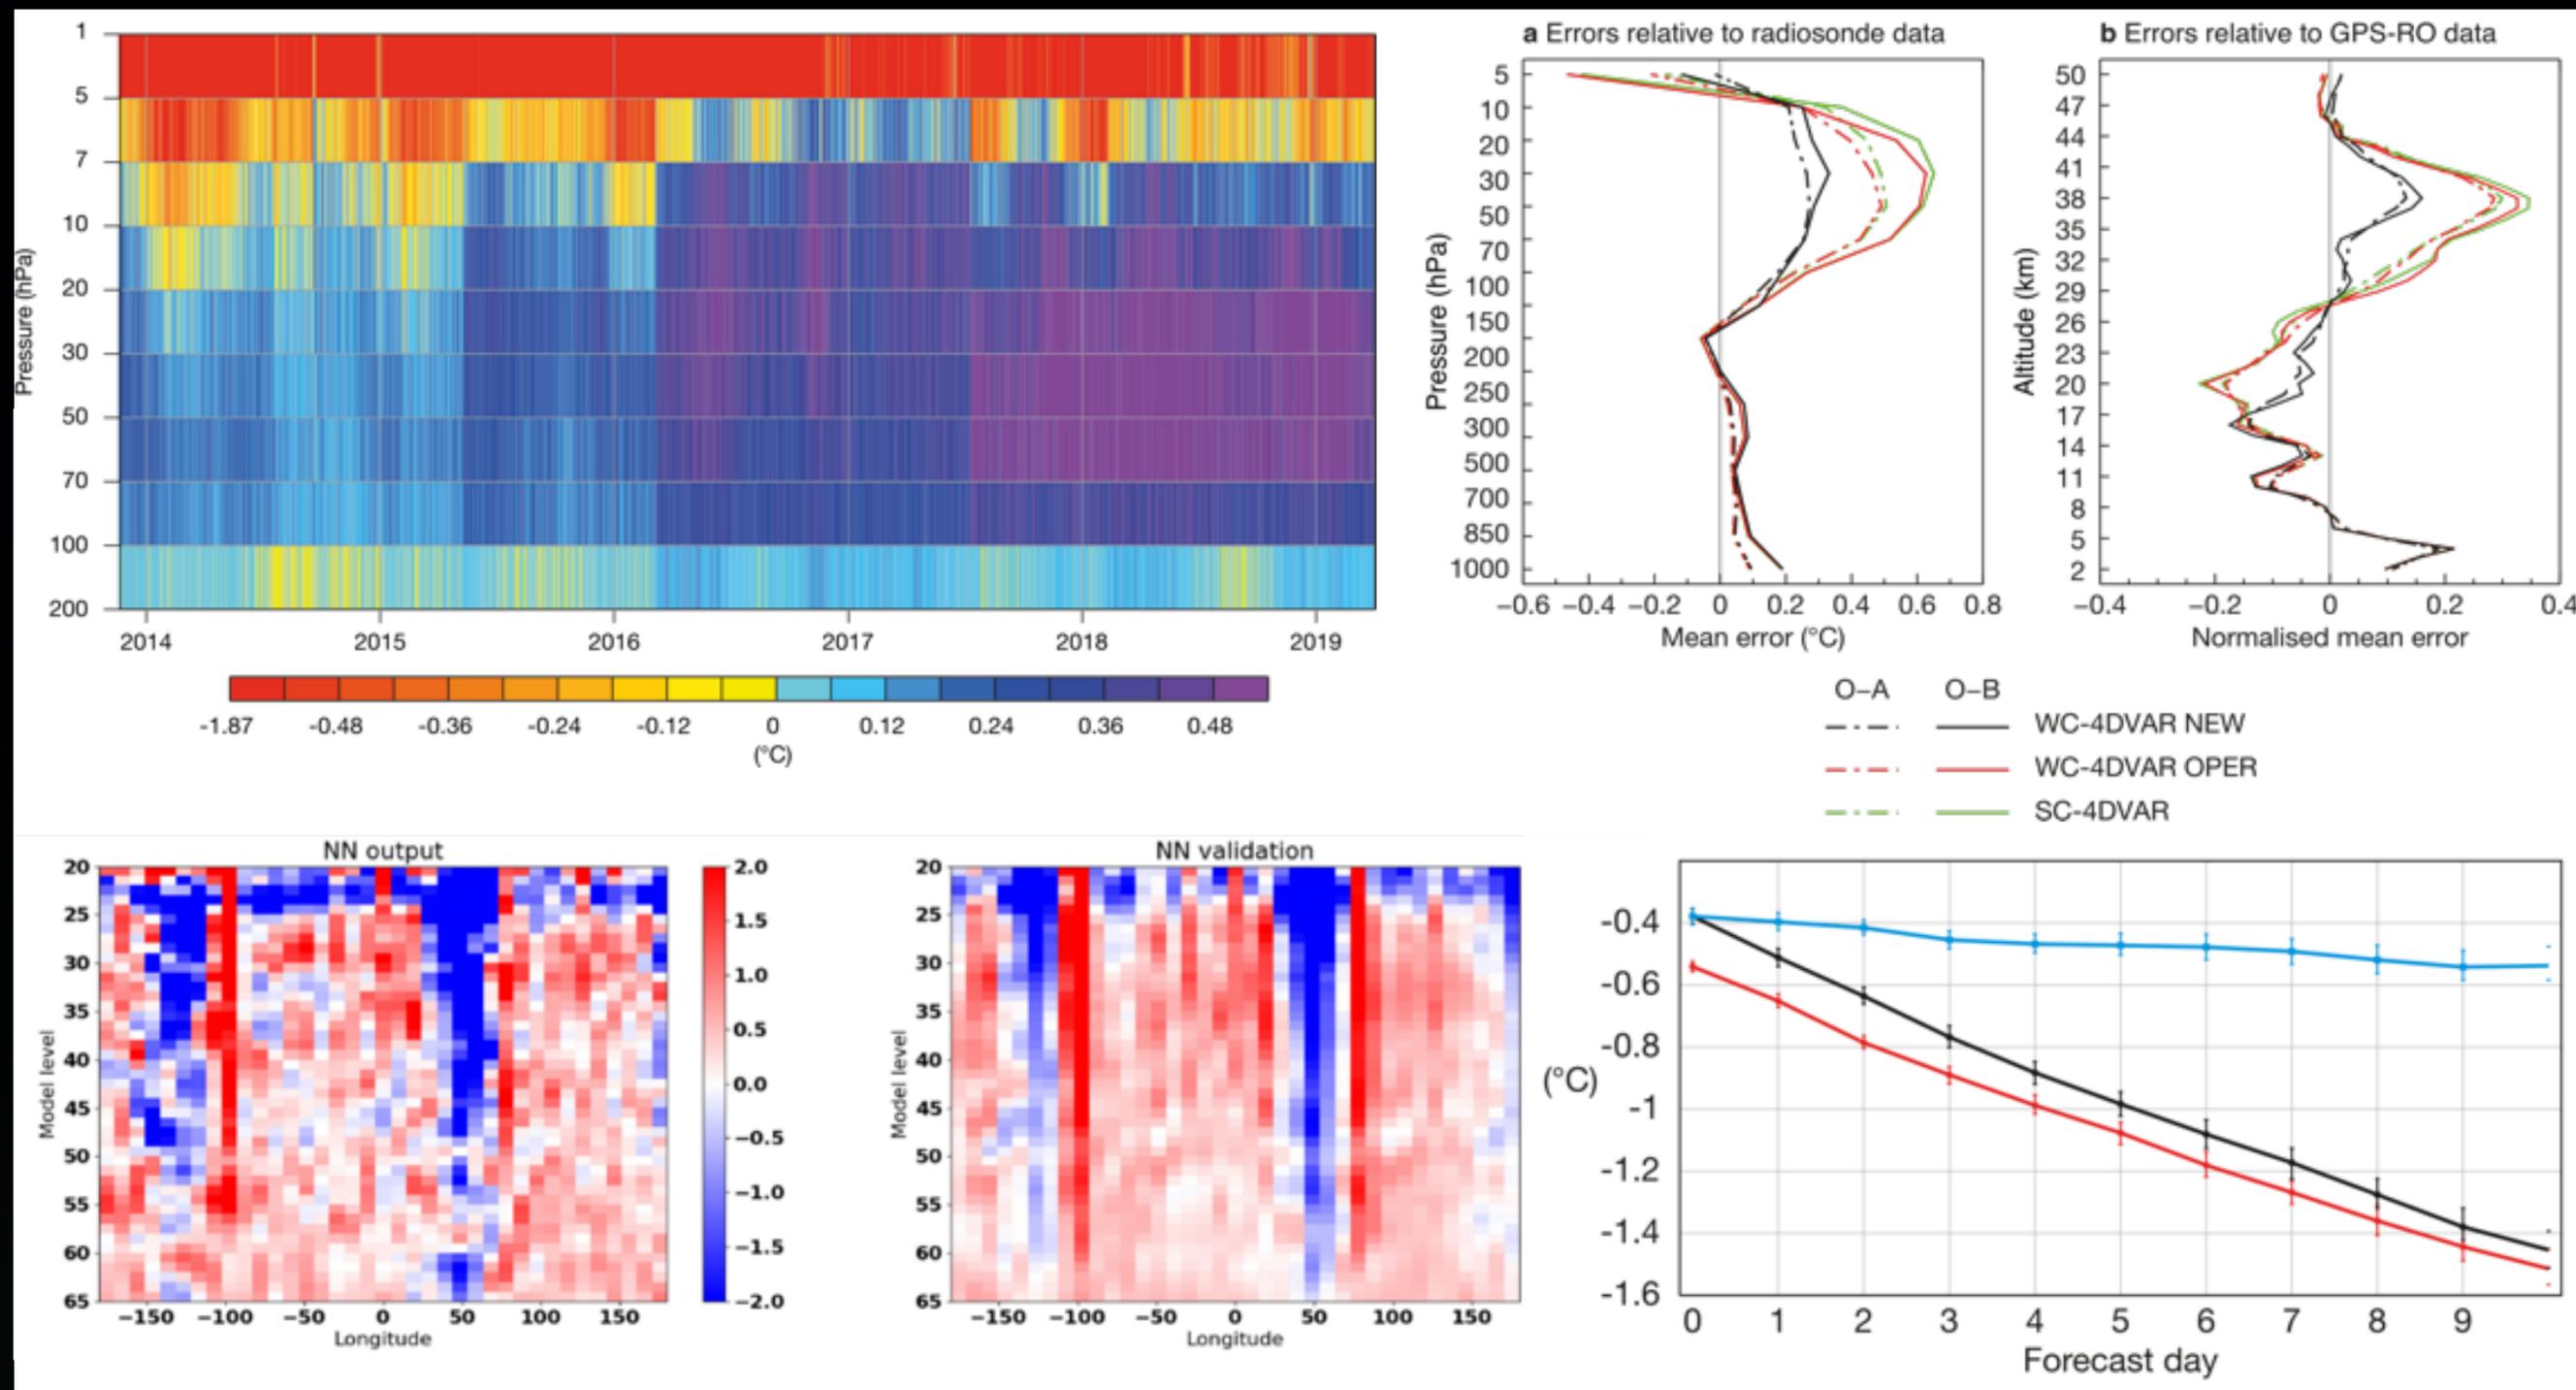
Stable machine-learning parameterization of subgrid processes for climate modeling at a range of resolutions

Janni Yuval<sup>1</sup> & Paul A. O’Gorman<sup>1</sup>



# BIAS CORRECTION

Work with ECMWF to remove IFS Model bias

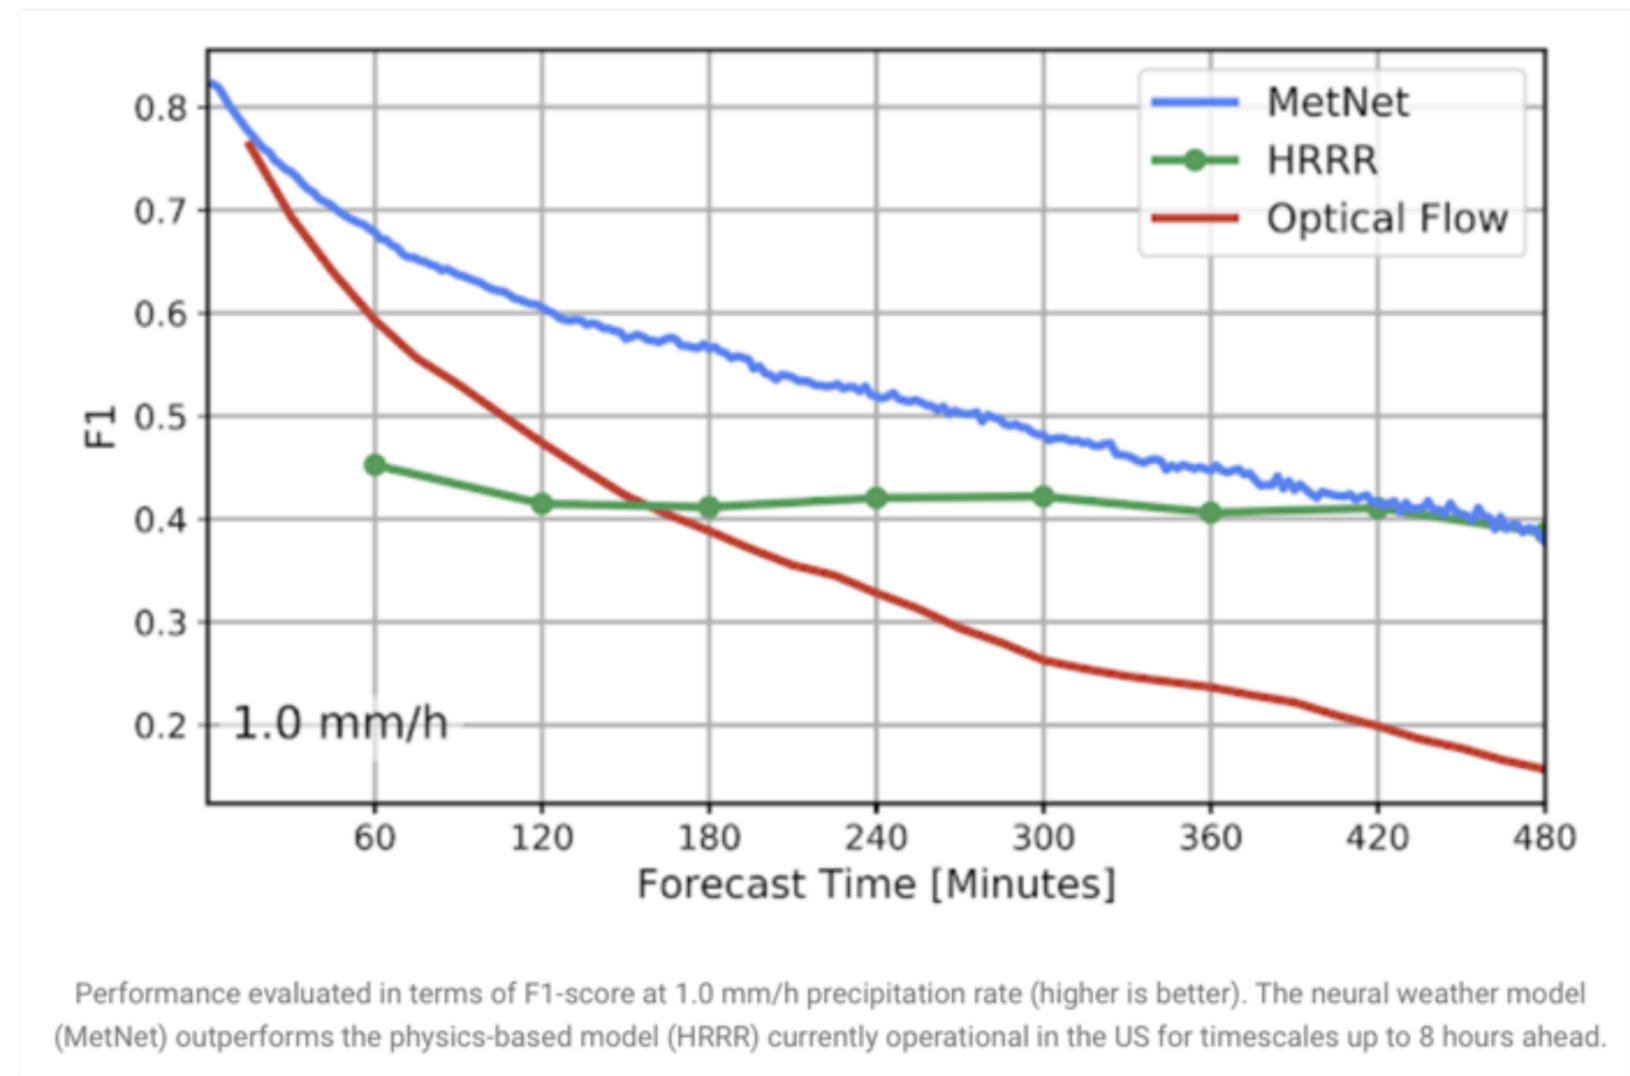
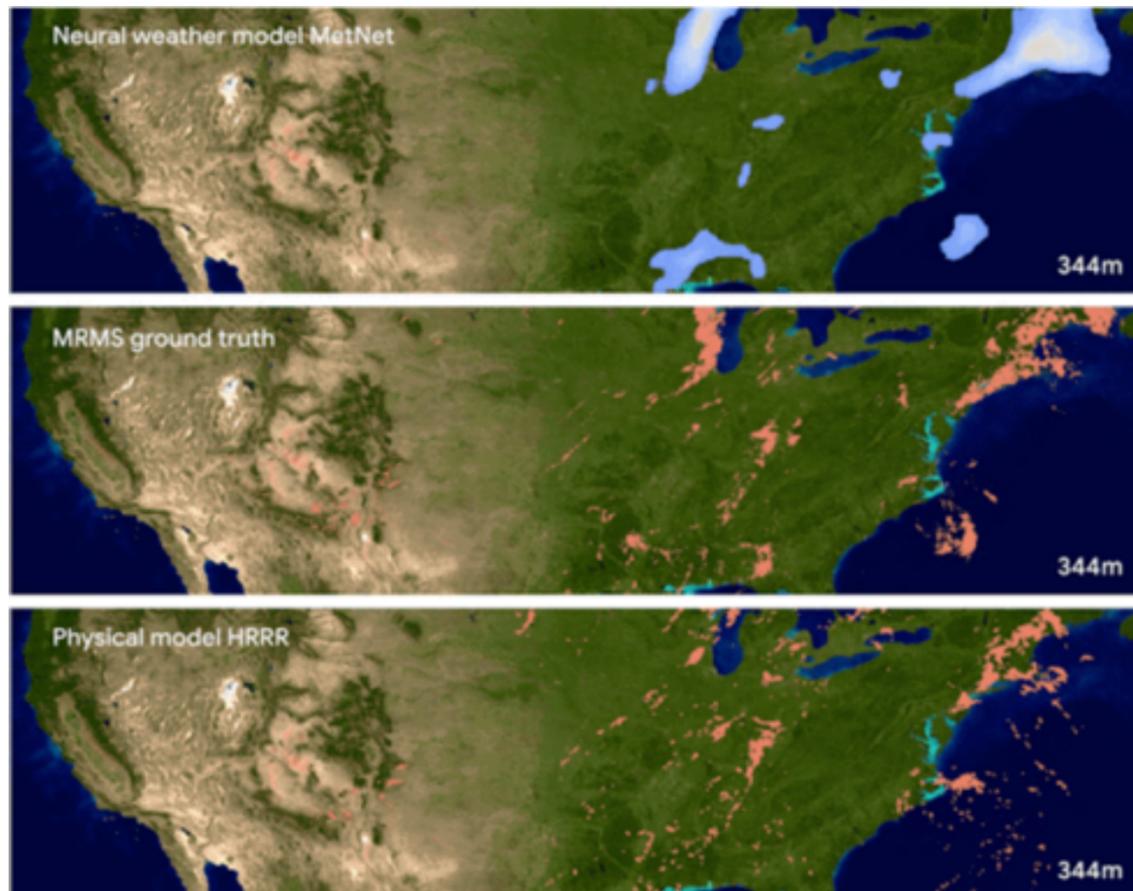


# NOWCASTING

MetNet: Can beat physical models up to 8 hours

## MetNet: A Neural Weather Model for Precipitation Forecasting

GOOGLE RESEARCH



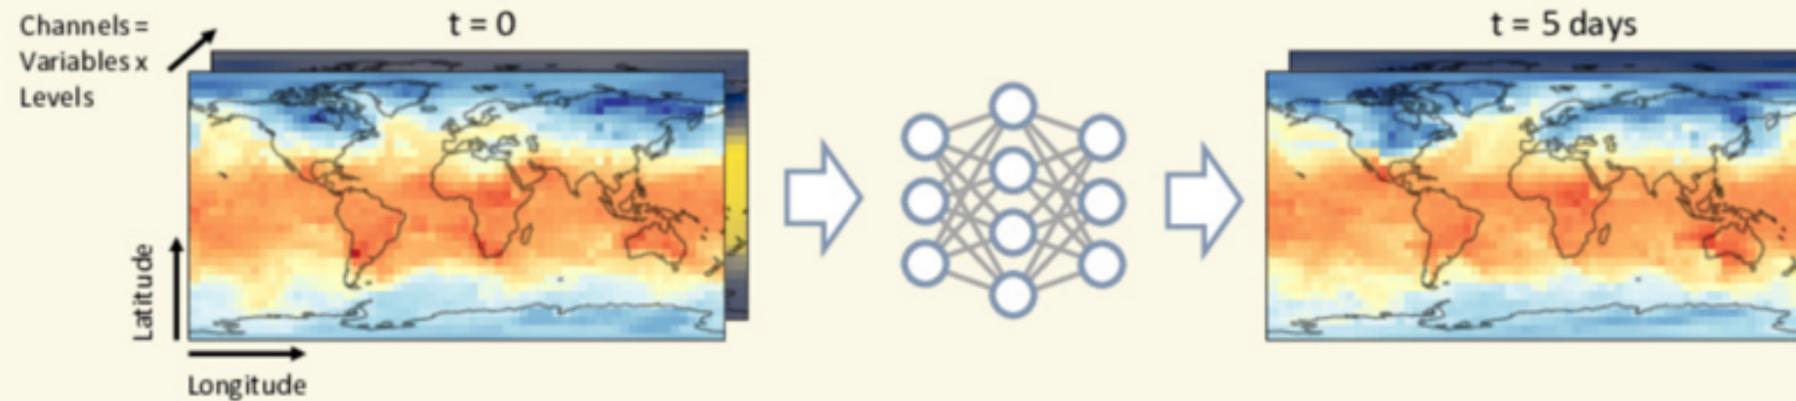
# MEDIUM-RANGE WEATHER FORECASTING

## Weather Bench: A Standardized Benchmark for Data Driven Forecasts

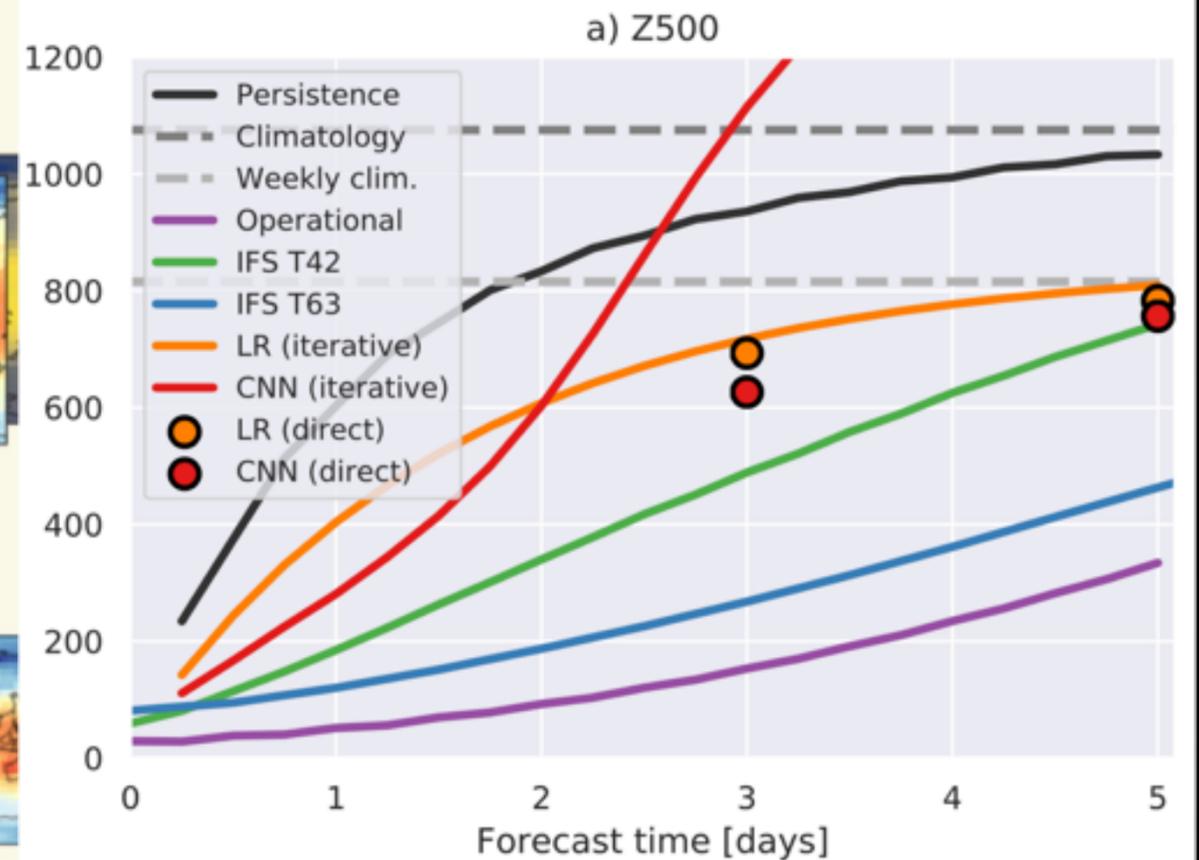
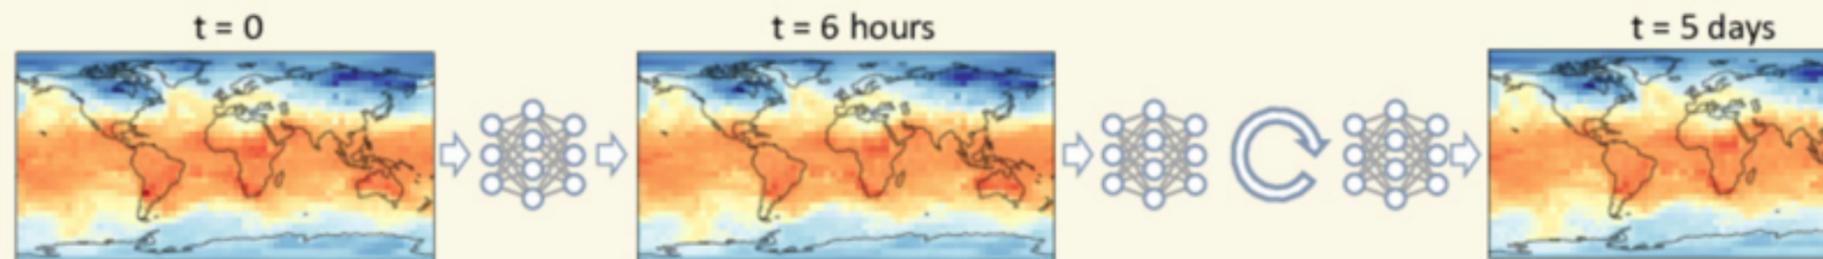
WeatherBench: A benchmark dataset for data-driven weather forecasting

Stephan Rasp, Peter D. Dueben, Sebastian Scher, Jonathan A. Weyn, Soukayna Mouatadid, Nils Thuerey

a) Direct prediction



b) Iterative prediction

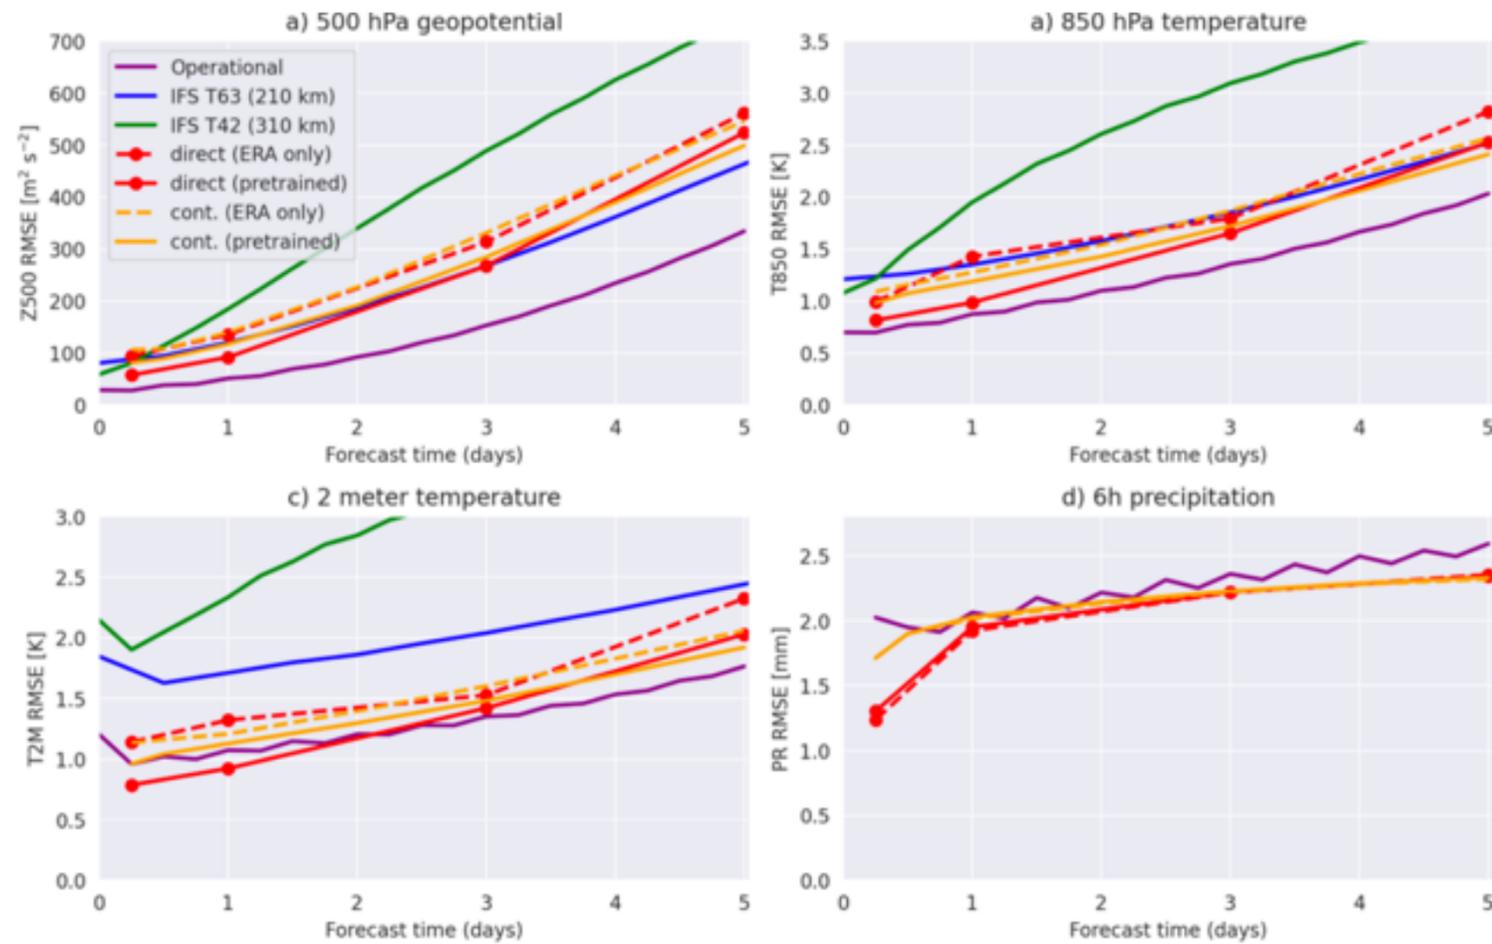


<https://arxiv.org/abs/2002.00469>

# MEDIUM-RANGE WEATHER FORECASTING

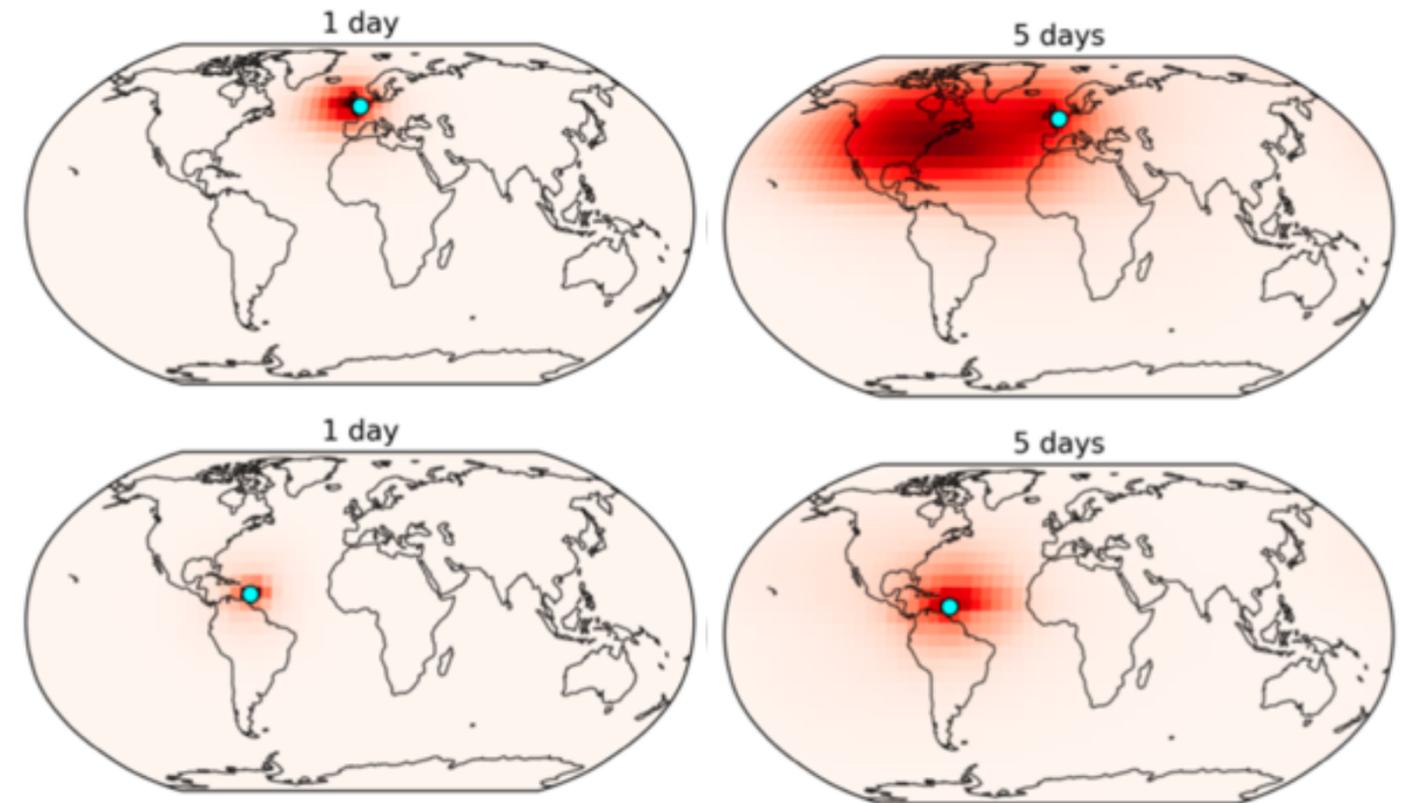
Data Driven models match accuracy of physical models on weather-bench

## 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

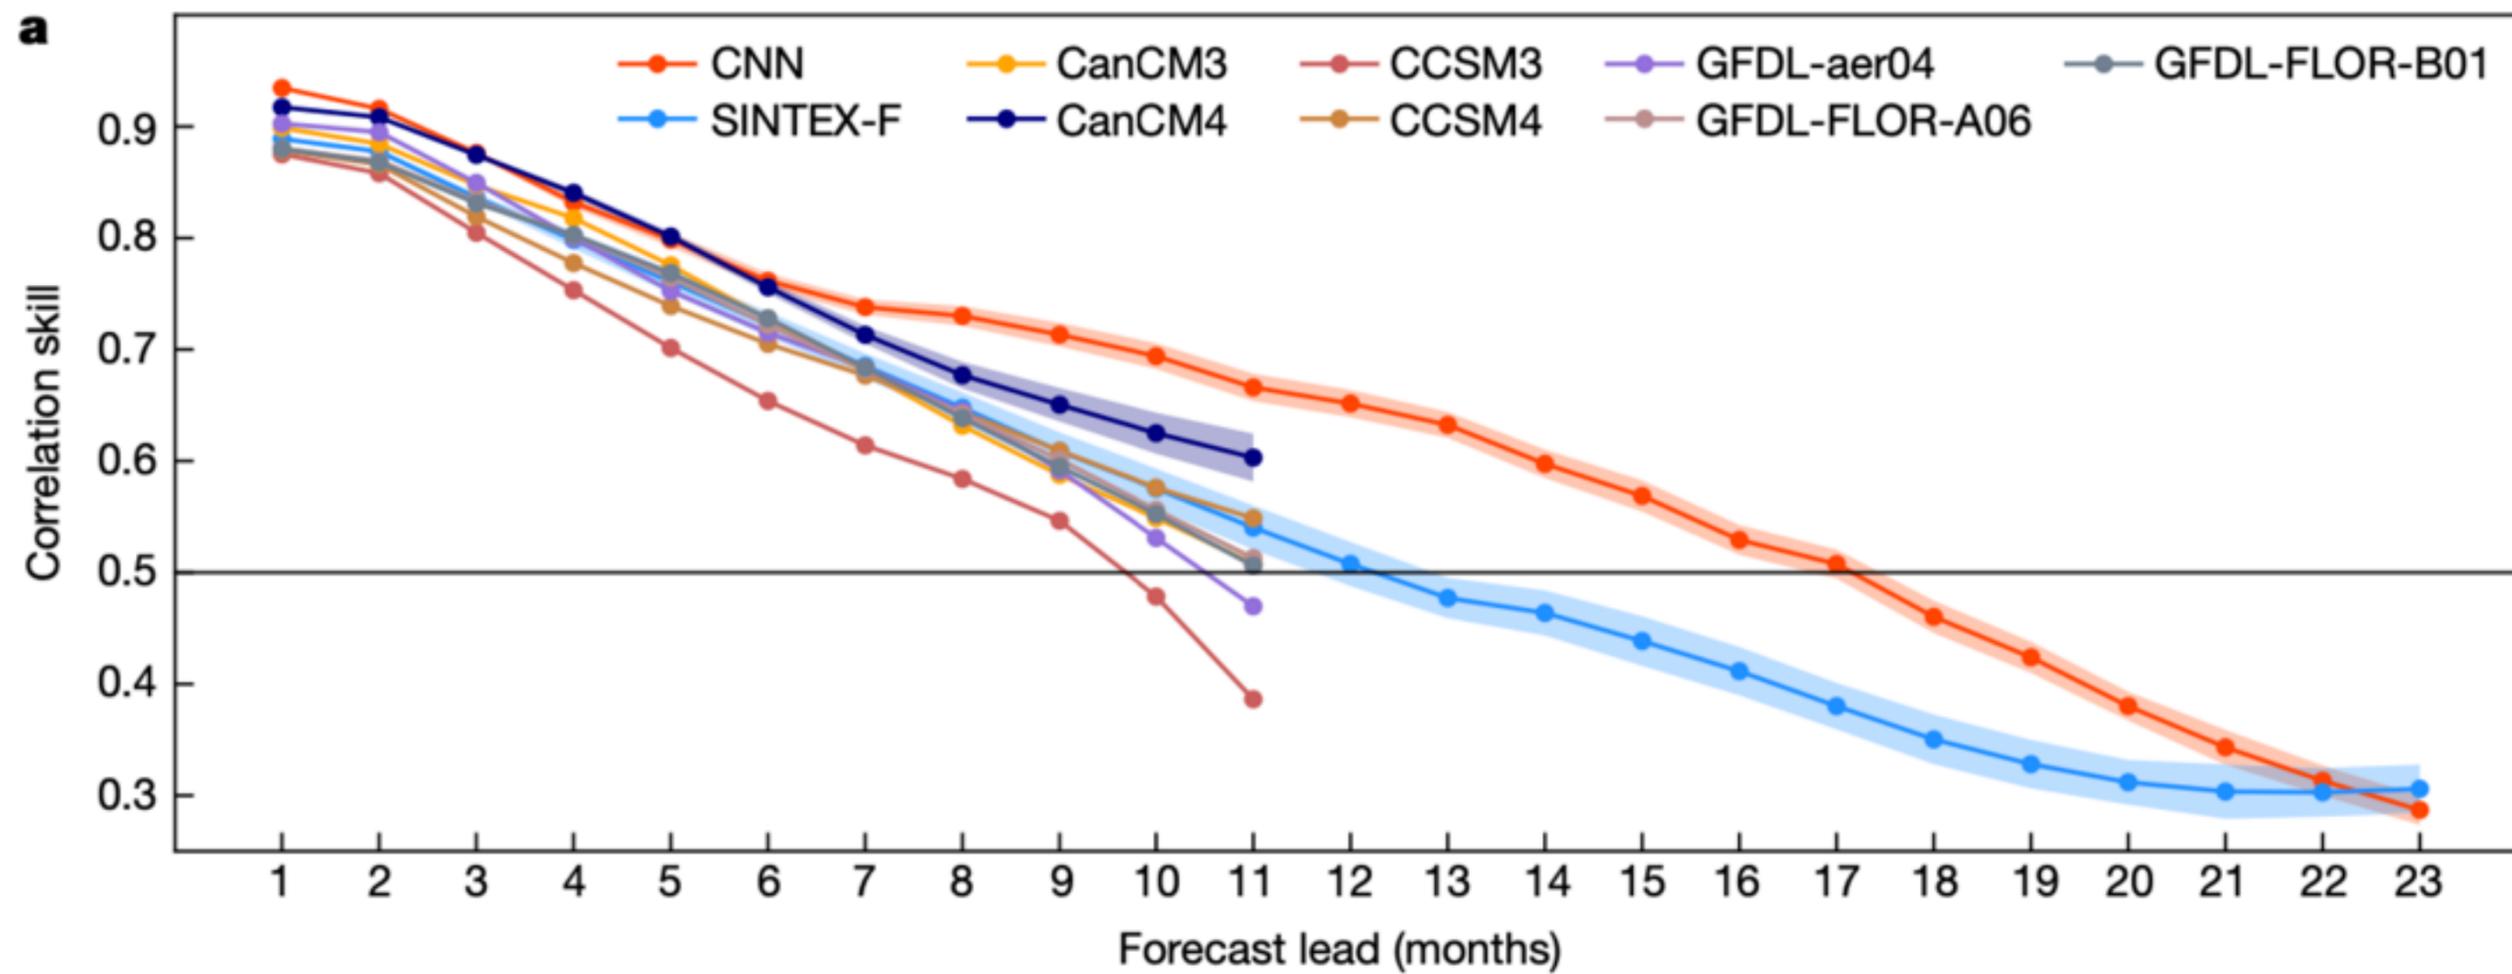


# EL NIÑO PREDICTION

Deep learning model able to predict El Niño up to 18 months in advance

## Deep learning for multi-year ENSO forecasts

Yoo-Geun Ham<sup>1\*</sup>, Jeong-Hwan Kim<sup>1</sup> & Jing-Jia Luo<sup>2,3</sup>



<https://www.nature.com/articles/s41586-019-1559-7.pdf>

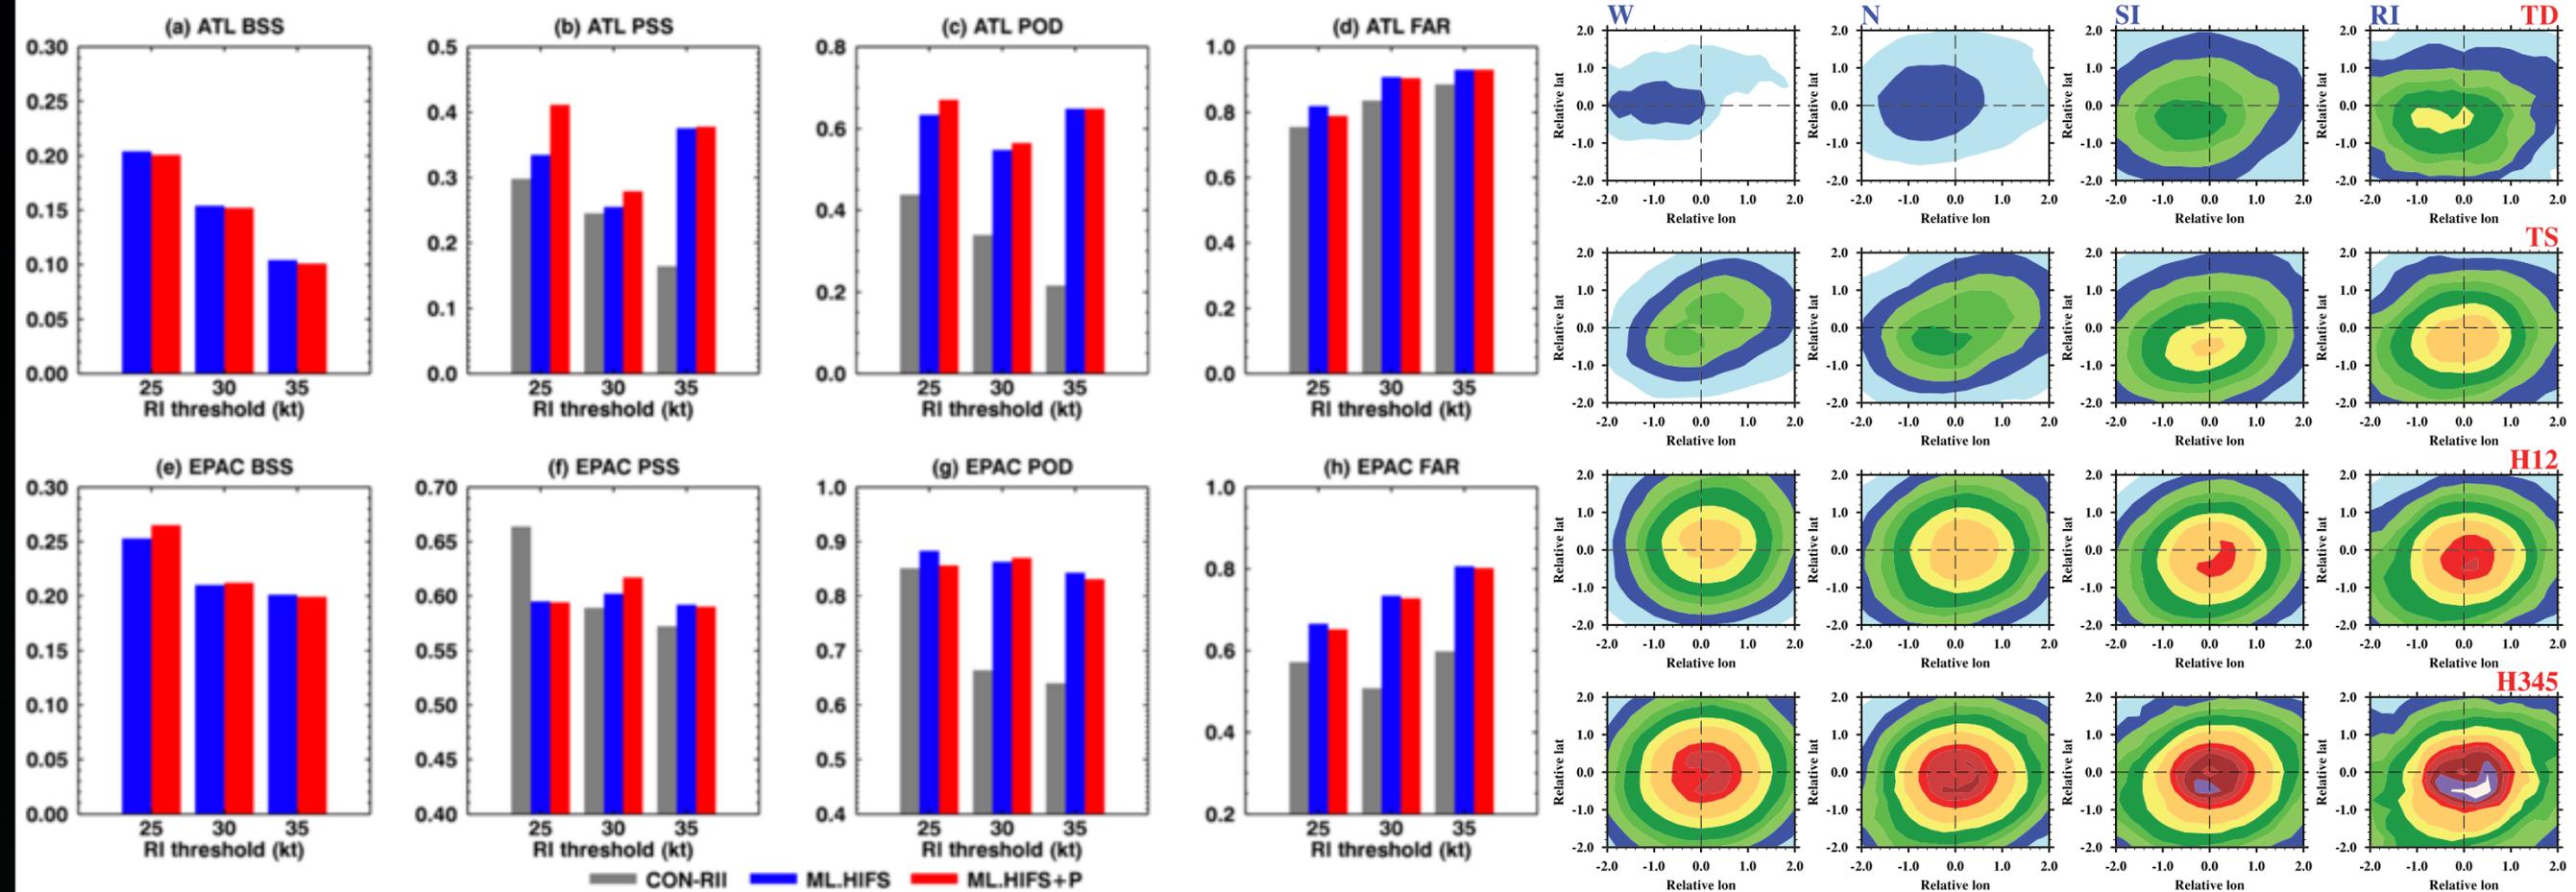
# RAPID INTENSIFICATION

Machine Learning improves detection by 40-200 % over operational consensus

## Applying Satellite Observations of Tropical Cyclone Internal Structures to Rapid Intensification Forecast With Machine Learning

Hui Su<sup>1</sup>, Longtao Wu<sup>1</sup>, Jonathan H. Jiang<sup>1</sup>, Raksha Pai<sup>2</sup>, Alex Liu<sup>3</sup>, Albert J. Zhai<sup>4</sup>, Peyman Tavallali<sup>1</sup>, and Mark DeMaria<sup>5</sup>

<sup>1</sup>Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA, <sup>2</sup>IBM Global Business Services, Armonk, NY, USA, <sup>3</sup>RMDS Lab, Pasadena, CA, USA, <sup>4</sup>Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA, USA, <sup>5</sup>National Hurricane Center, Miami, FL, USA



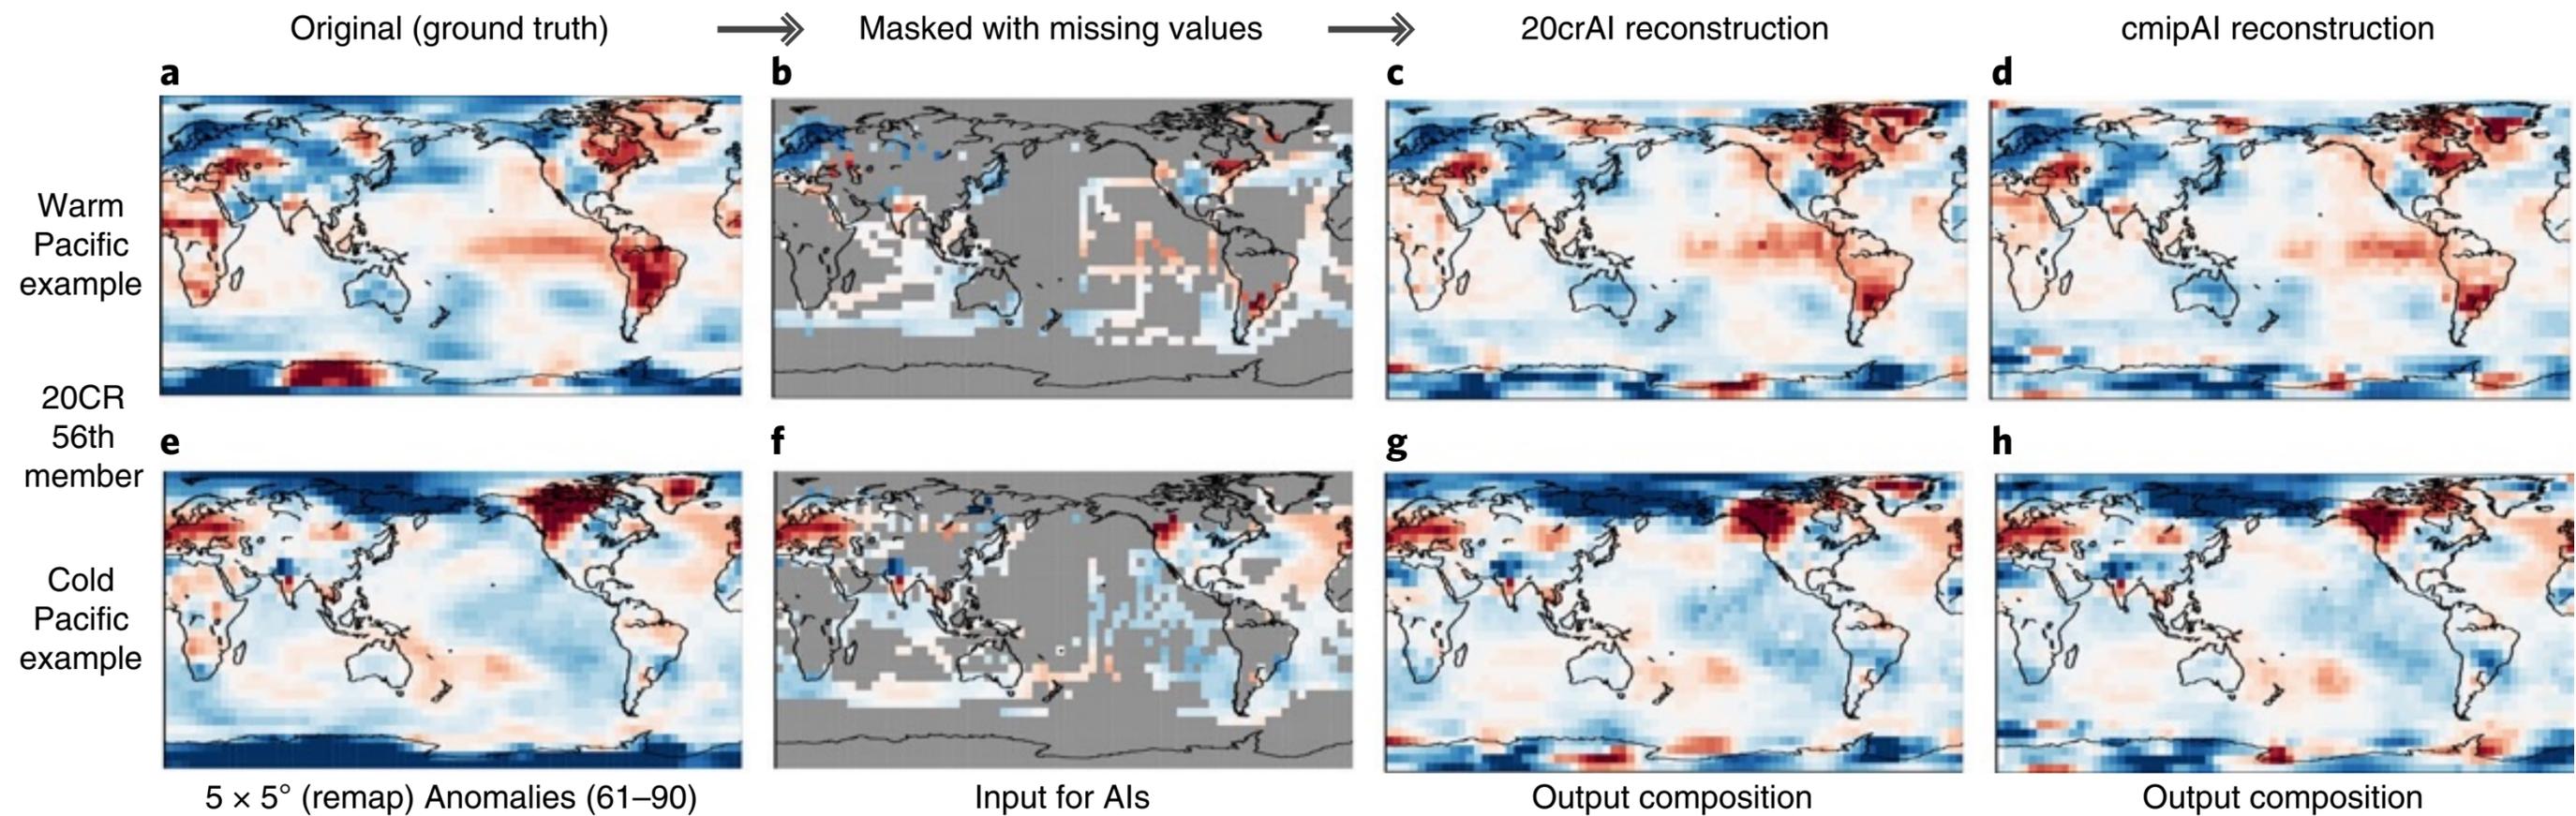
<https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2020GL089102>

# FILLING-IN MISSING CLIMATE OBSERVATIONS

DL transfer learning + inpainting beats Kriging and PCA

## Artificial intelligence reconstructs missing climate information

Christopher Kadow<sup>1,2</sup>, David Matthew Hall<sup>3</sup> and Uwe Ulbrich<sup>2</sup>





# SCIENTIFIC CHALLENGES

# SCIENTIFIC CHALLENGES

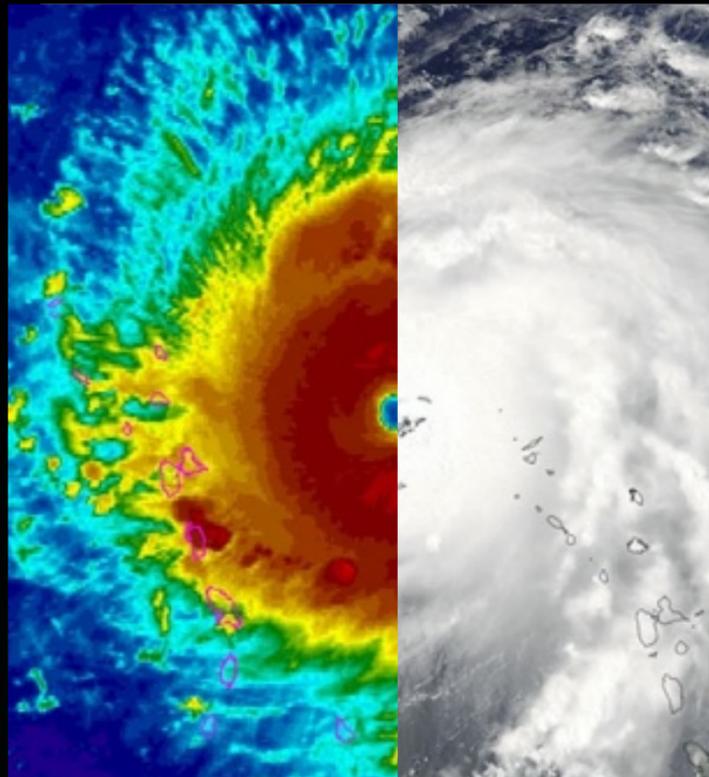
Problems that arise when applying AI to science

- Data Labelling
- Limited Data
- Enforcing Physical Constraints
- Uncertainty Quantification
- Trustworthiness and Interpretability
- HPC – AI Coupling
- Loss of Dynamic Range
- Data Movement
- Numerical Stability
- Generalization



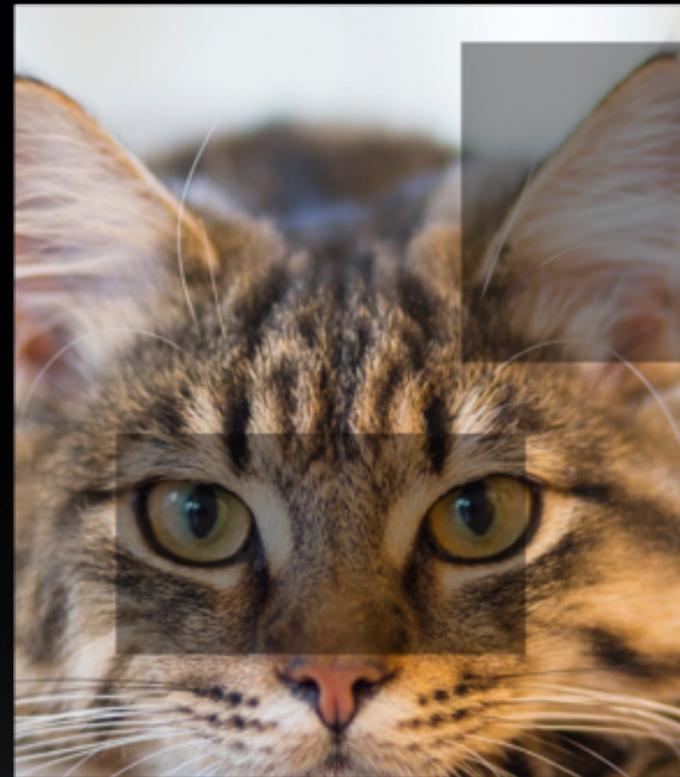
# DATA LABELLING

How can we get enough labelled data?



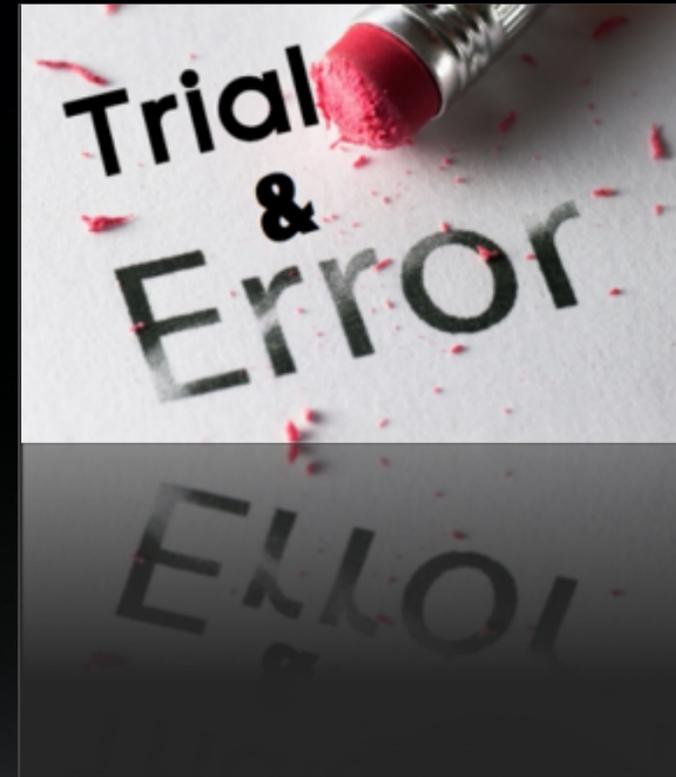
## Data Fusion

Using one data source as the label for another



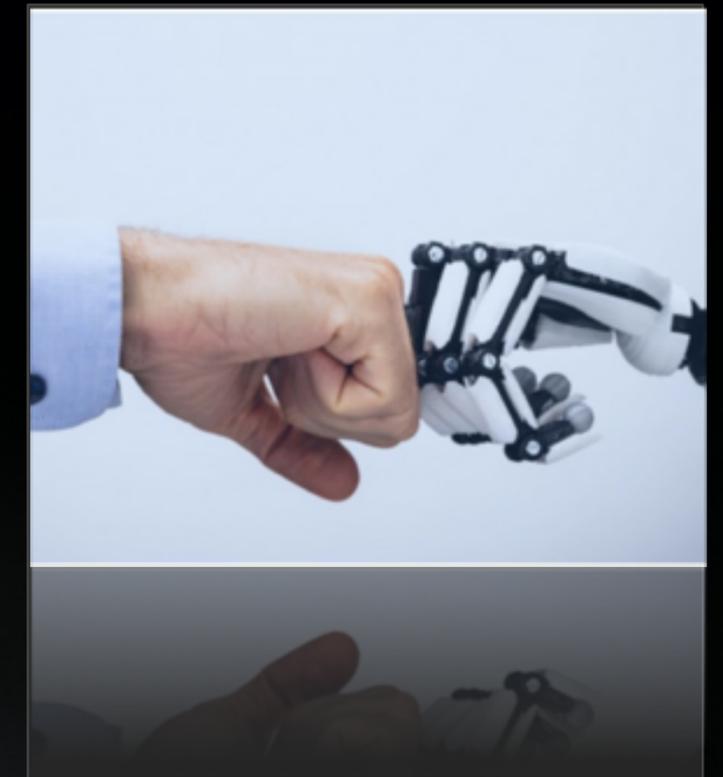
## Self-Supervised Learning

Predicting input B from input A



## Reinforcement Learning

Obtaining labels directly from the environment or simulation



## Active Learning

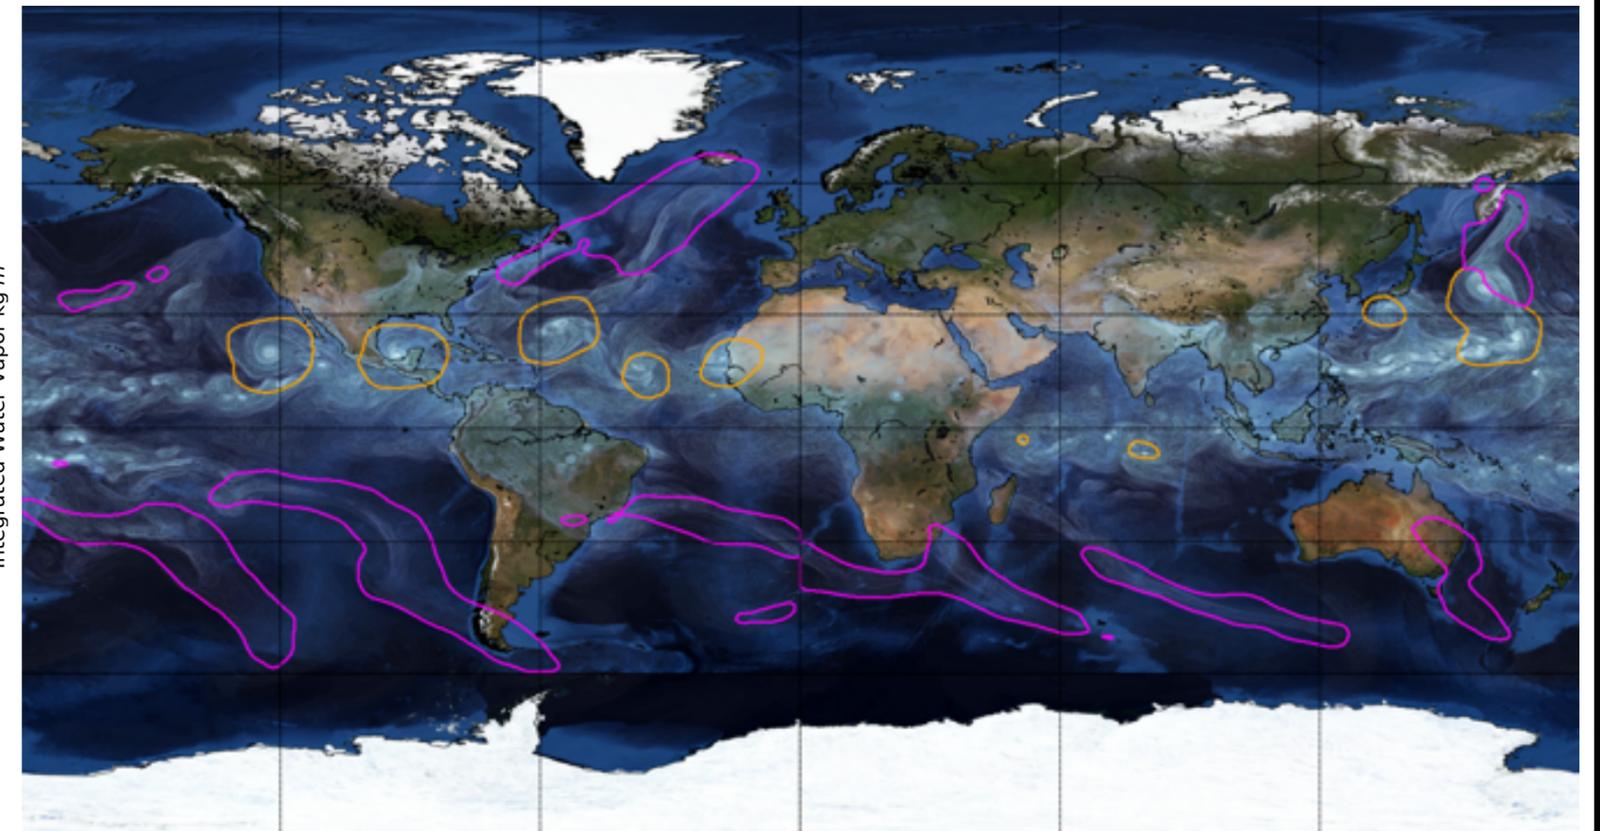
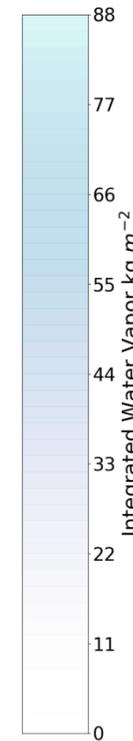
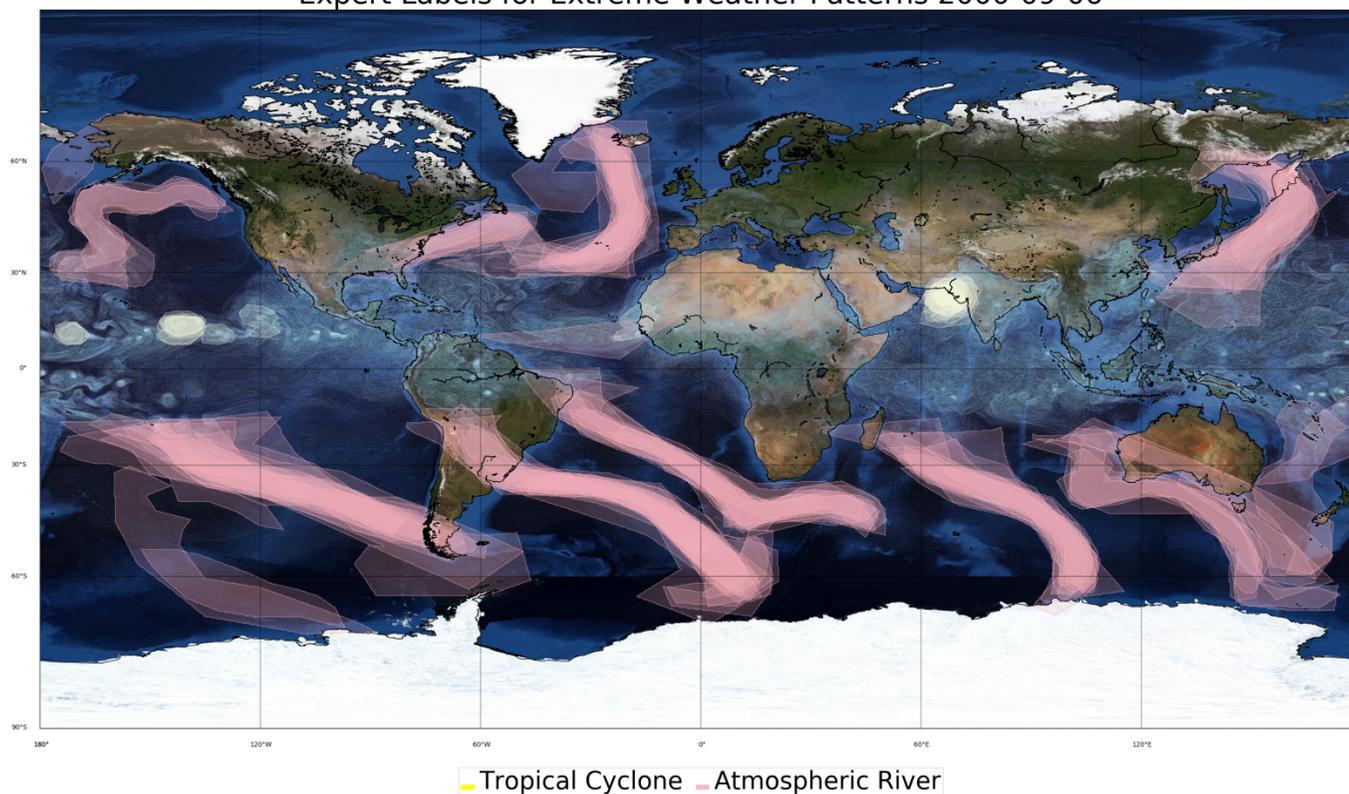
Using human machine iteration to make labelling easier

# DATA LABELLING

## Online Extreme-Weather Labelling Tool

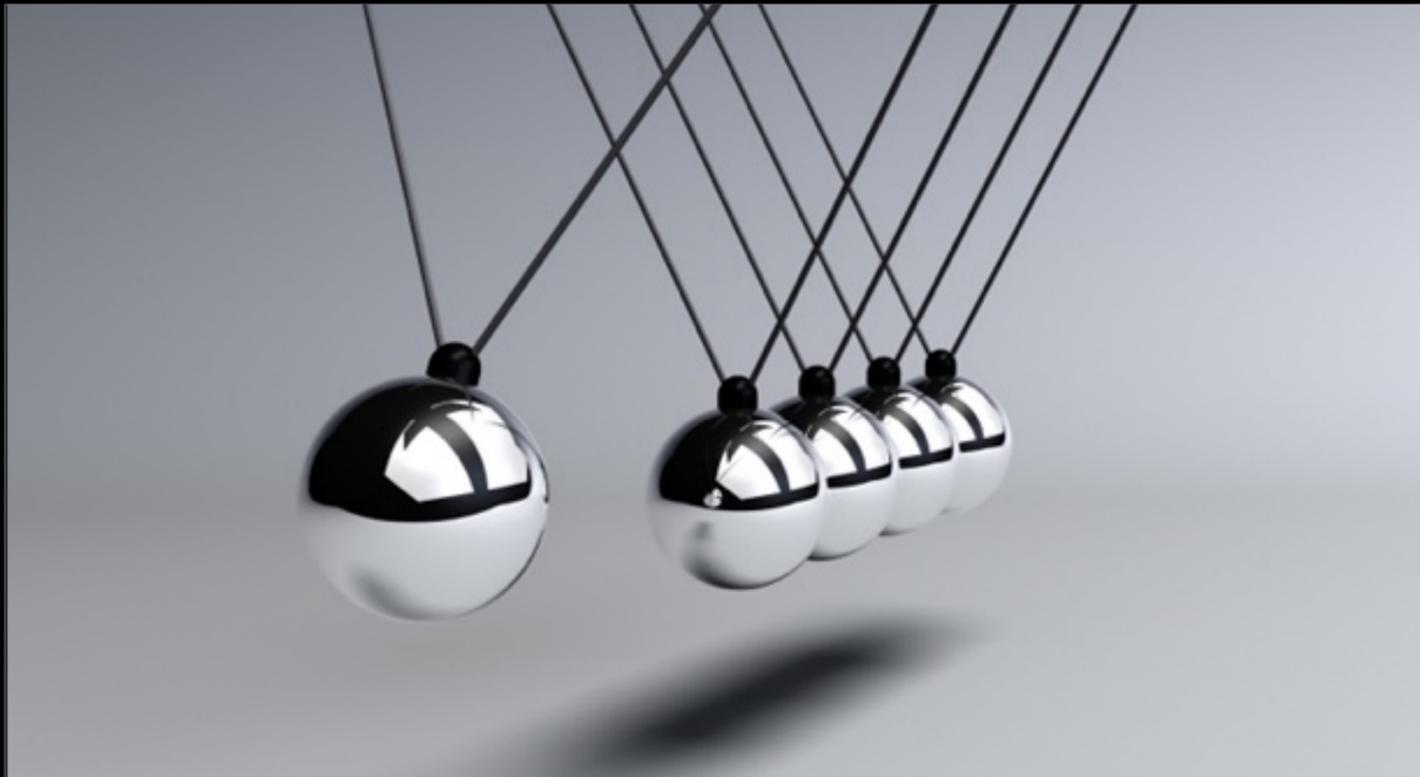
**ClimateNet: an expert-labelled open dataset and Deep Learning architecture for enabling high-precision analyses of extreme weather**

Expert Labels for Extreme Weather Patterns 2000-09-06

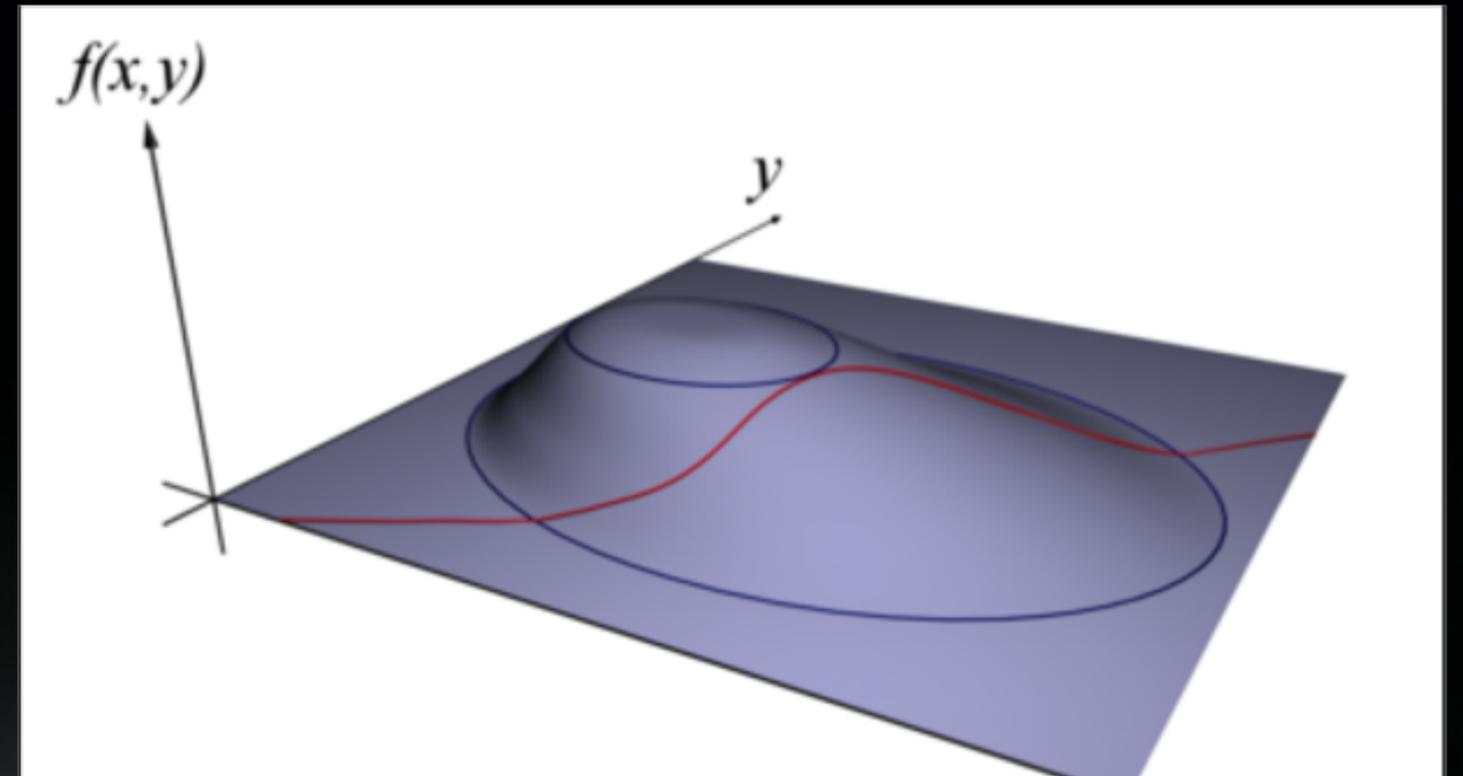


# ENFORCING PHYSICAL CONSTRAINTS

Is the solution physically correct?



Conservation of Mass, Momentum, Energy, Incompressibility,  
Turbulent Energy Spectra, Translational Invariance



Loss Penalization, Hard Constraints, Projective Methods,  
Differentiable Programs

# ENFORCING PHYSICAL CONSTRAINTS

## Physics Informed Neural Nets

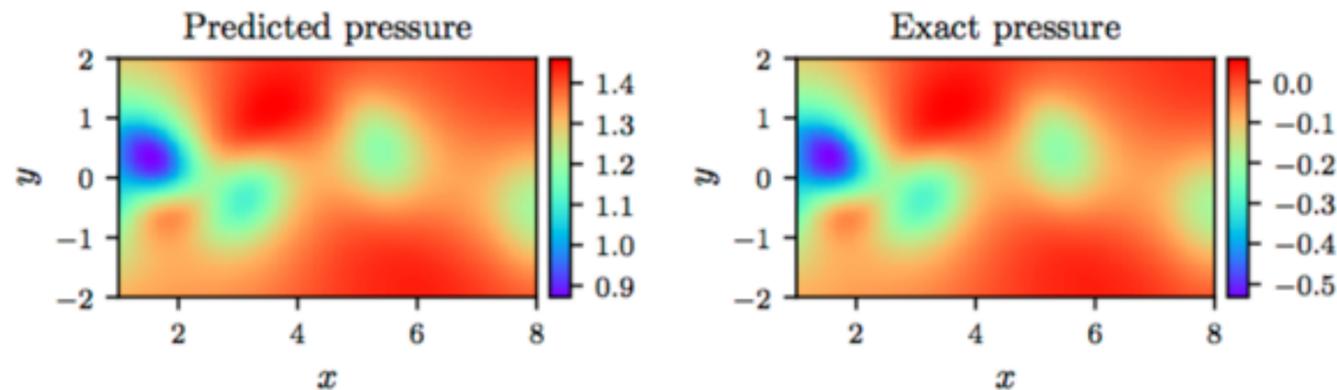
Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations

M. Raissi<sup>a</sup>, P. Perdikaris<sup>b,\*</sup>, G.E. Karniadakis<sup>a</sup>

<sup>a</sup> Division of Applied Mathematics, Brown University, Providence, RI, 02912, USA

<sup>b</sup> Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA, 19104, USA

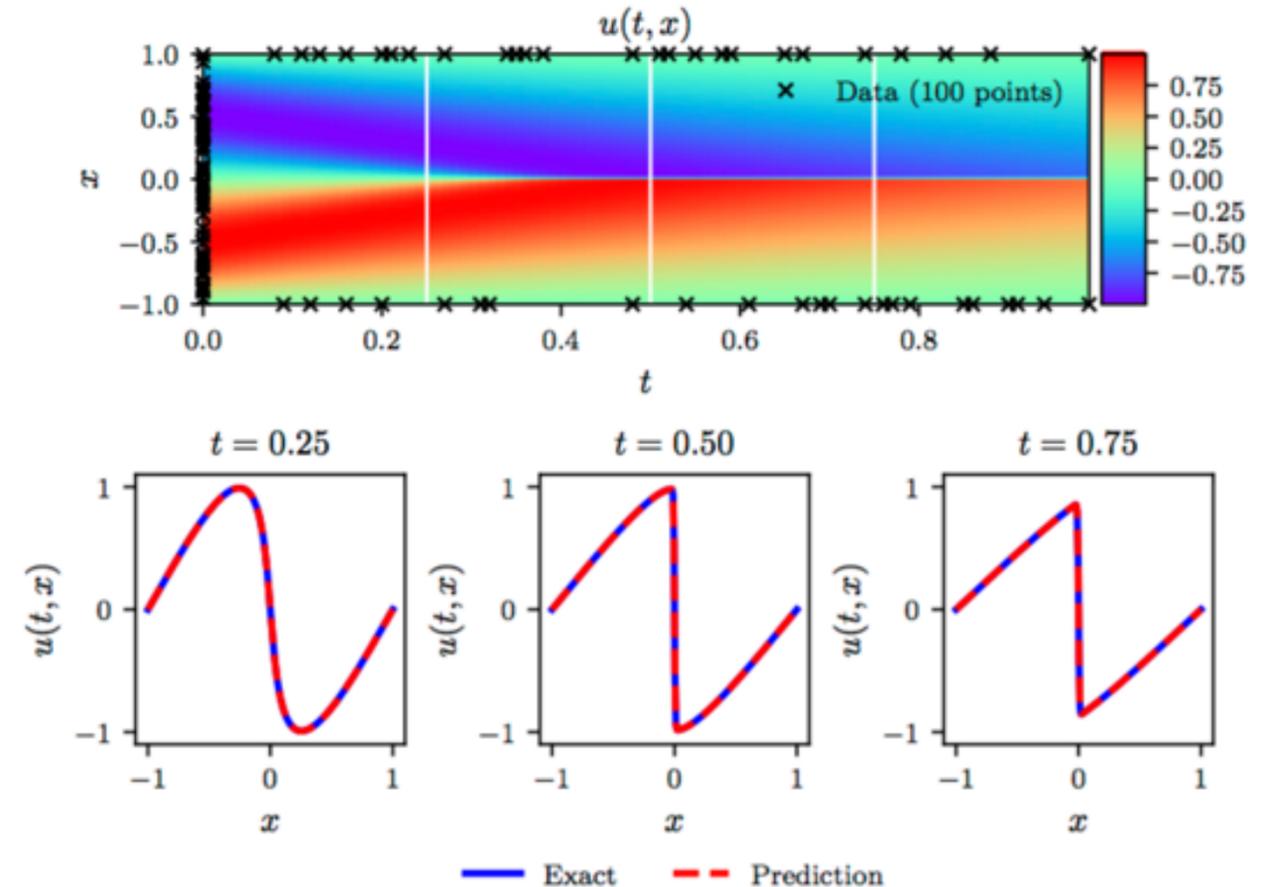
### Data-driven Discovery of Nonlinear Partial Differential Equations



Correct PDE	$u_t + (uu_x + vv_y) = -p_x + 0.01(u_{xx} + u_{yy})$ $v_t + (uv_x + vv_y) = -p_y + 0.01(v_{xx} + v_{yy})$
Identified PDE (clean data)	$u_t + 0.999(uu_x + vv_y) = -p_x + 0.01047(u_{xx} + u_{yy})$ $v_t + 0.999(uv_x + vv_y) = -p_y + 0.01047(v_{xx} + v_{yy})$
Identified PDE (1% noise)	$u_t + 0.998(uu_x + vv_y) = -p_x + 0.01057(u_{xx} + u_{yy})$ $v_t + 0.998(uv_x + vv_y) = -p_y + 0.01057(v_{xx} + v_{yy})$

### Data-driven Solutions of Nonlinear Partial Differential Equations

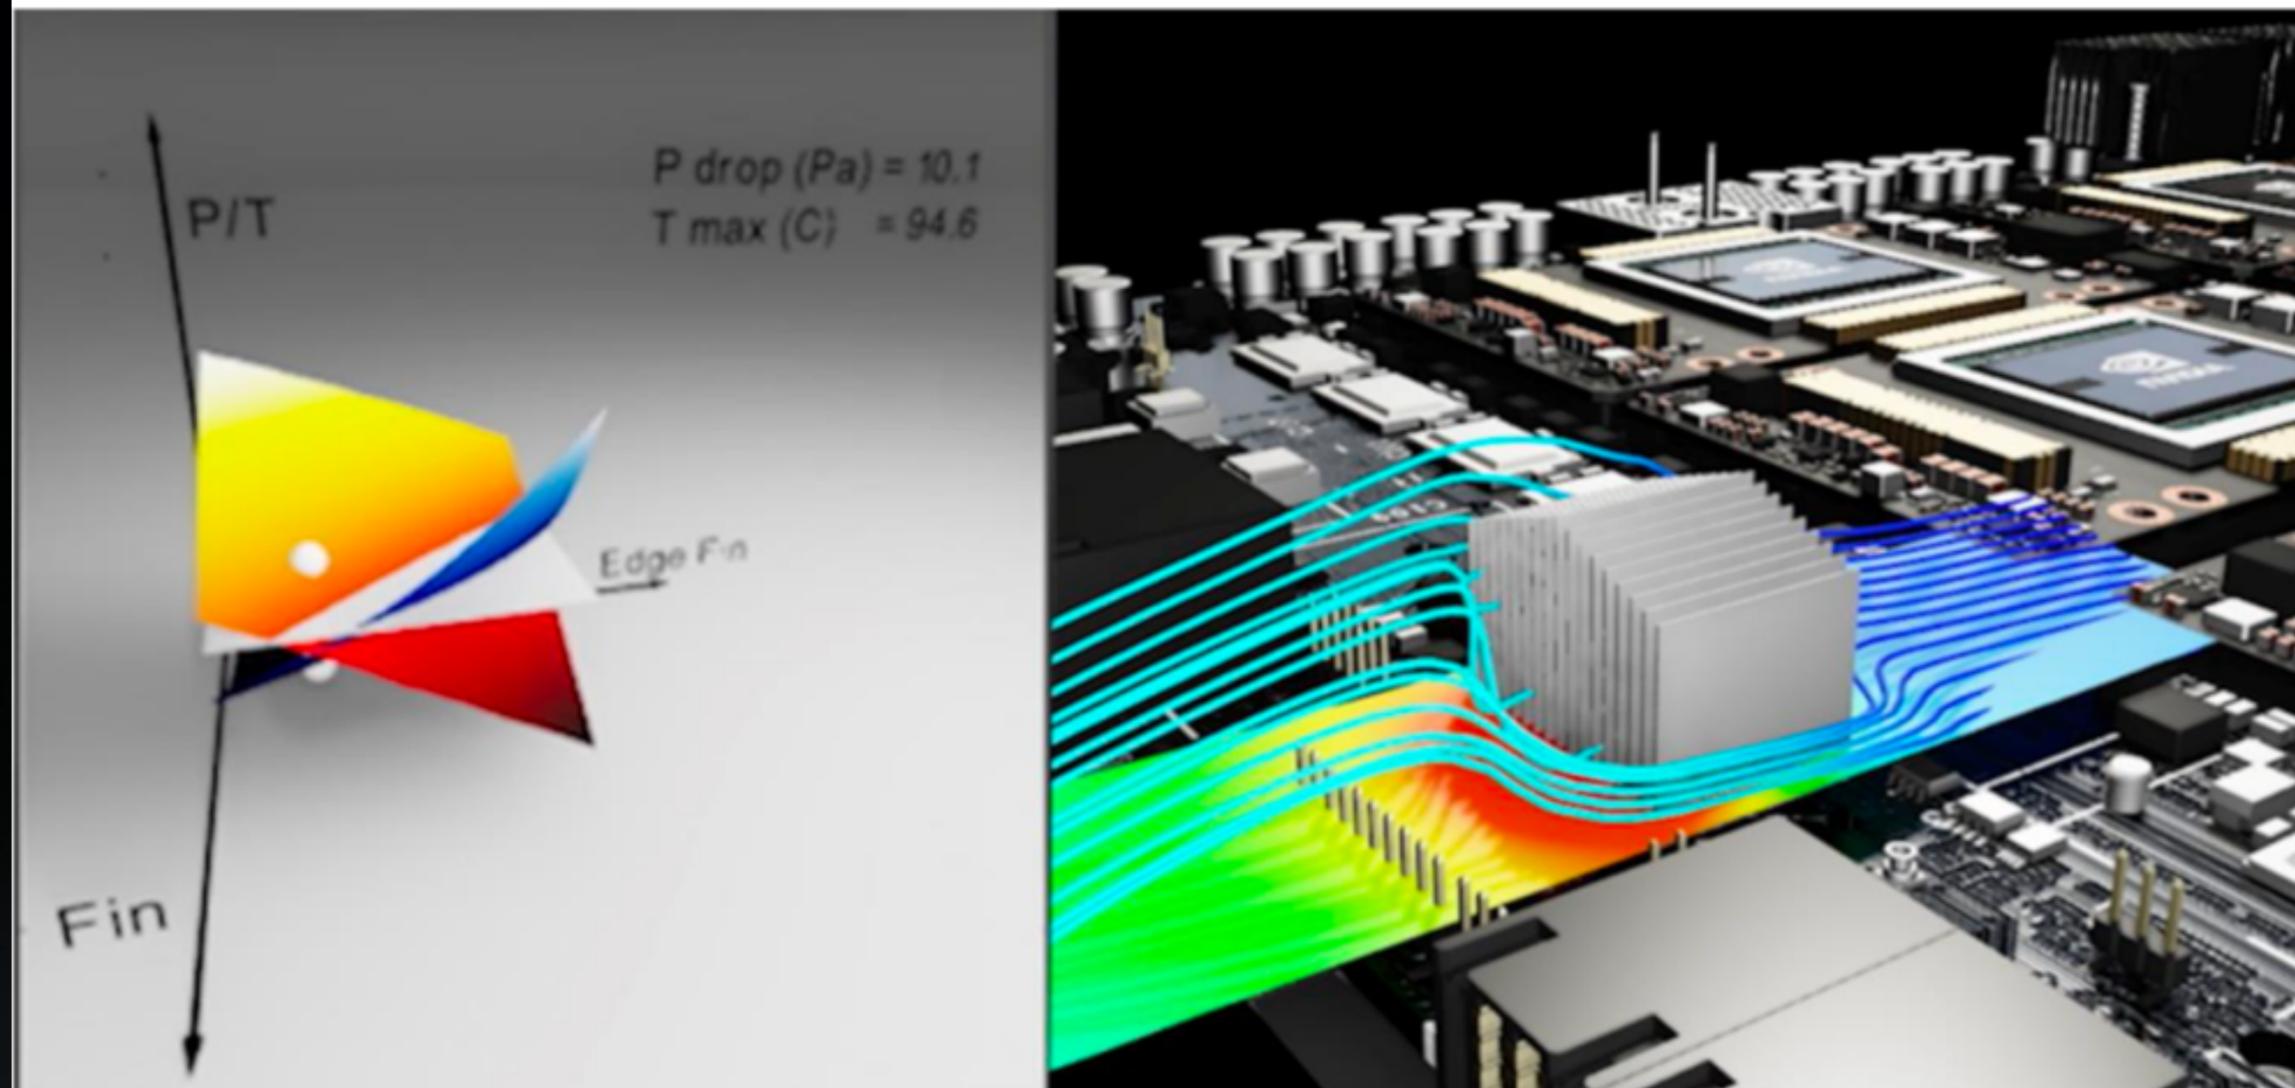
The following figure summarizes our results for the data-driven solution of the Burgers' equation.



# ENFORCING PHYSICAL CONSTRAINTS

NVIDIA SimNet

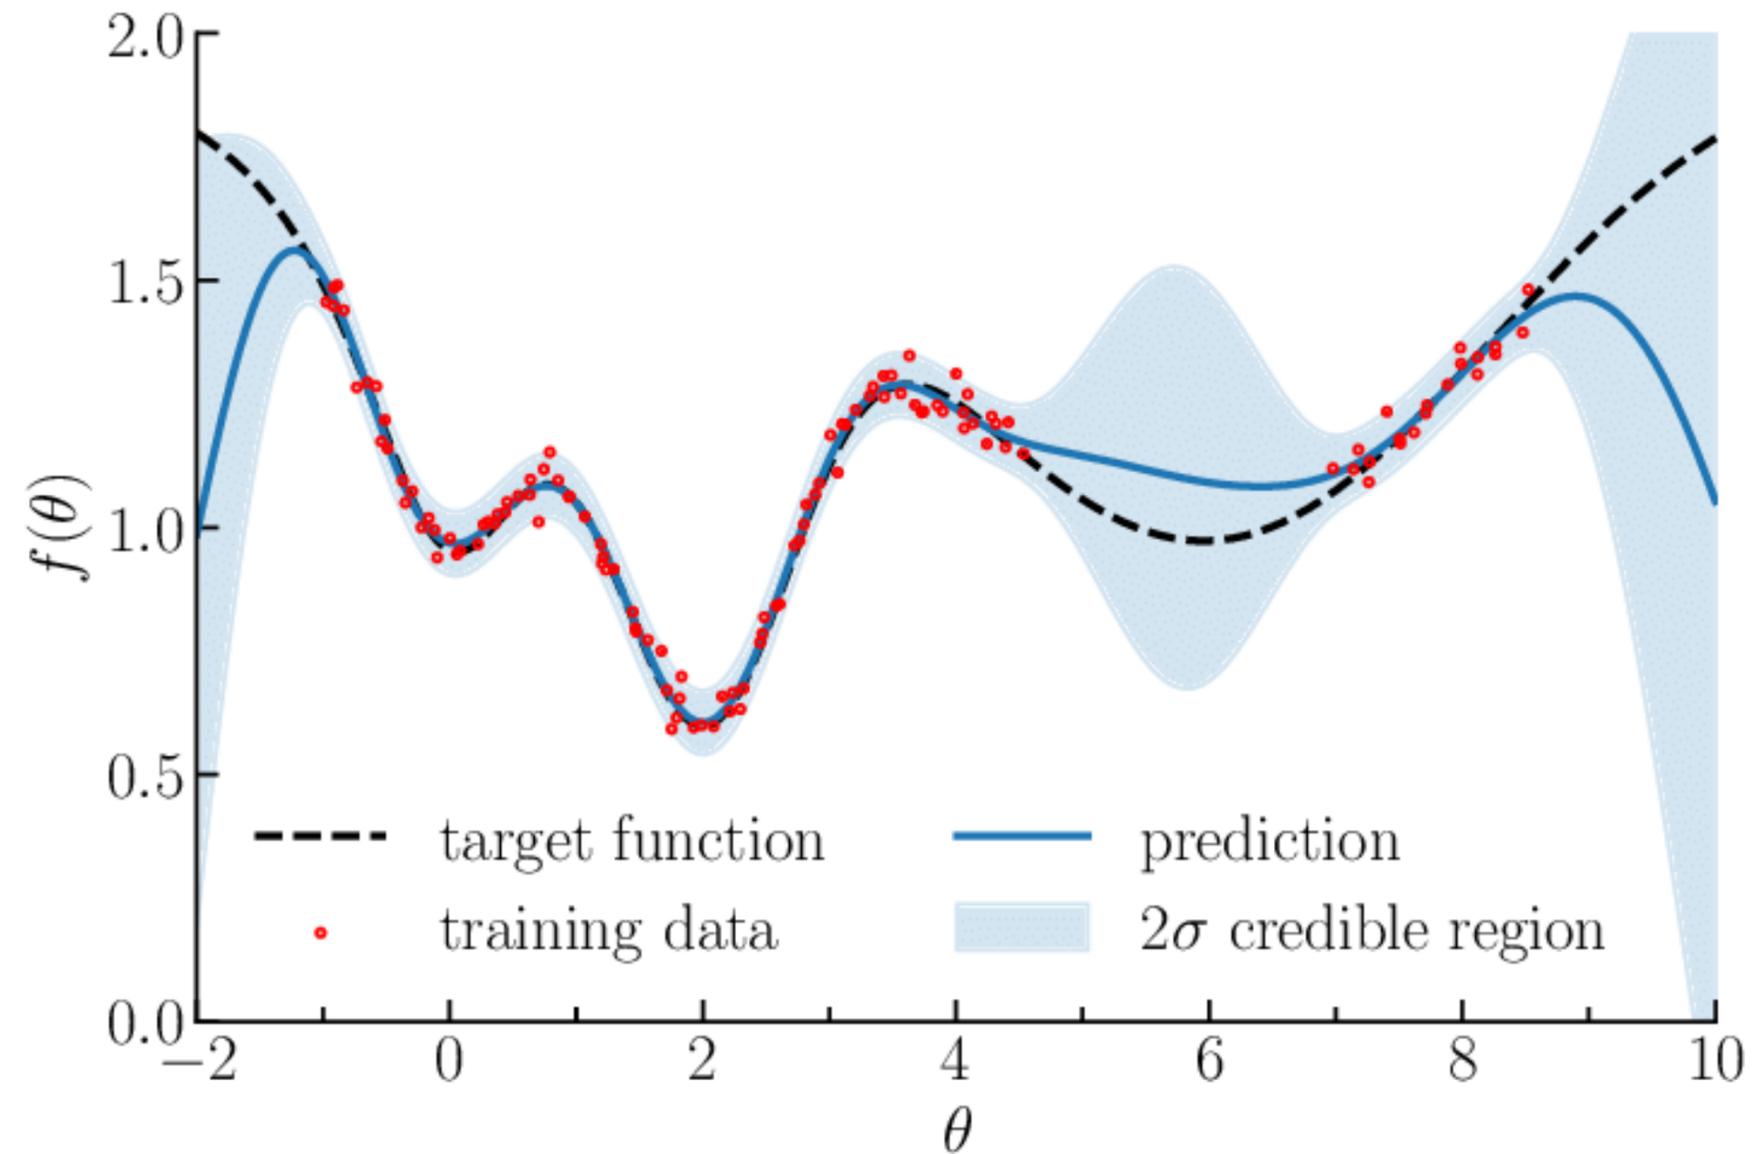
A Neural Network Based Partial Differential Equation Solver



<https://www.youtube.com/watch?v=Oq2Mpi5pF1w>

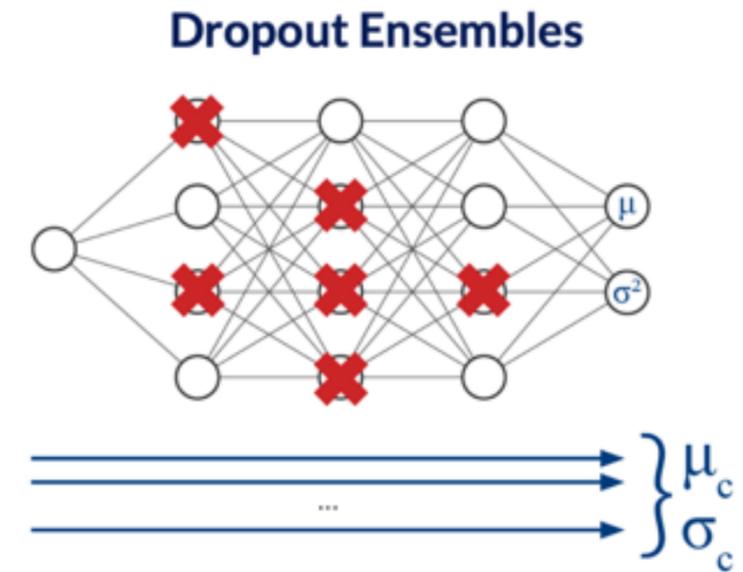
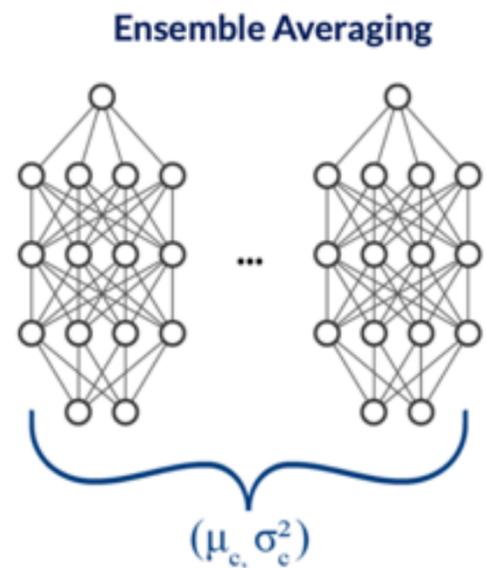
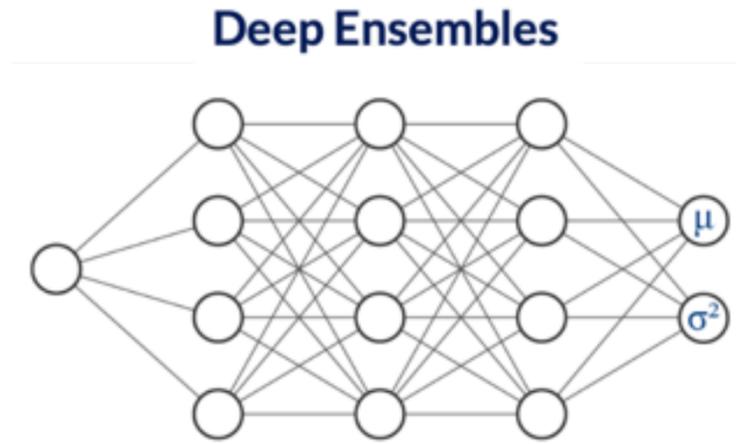
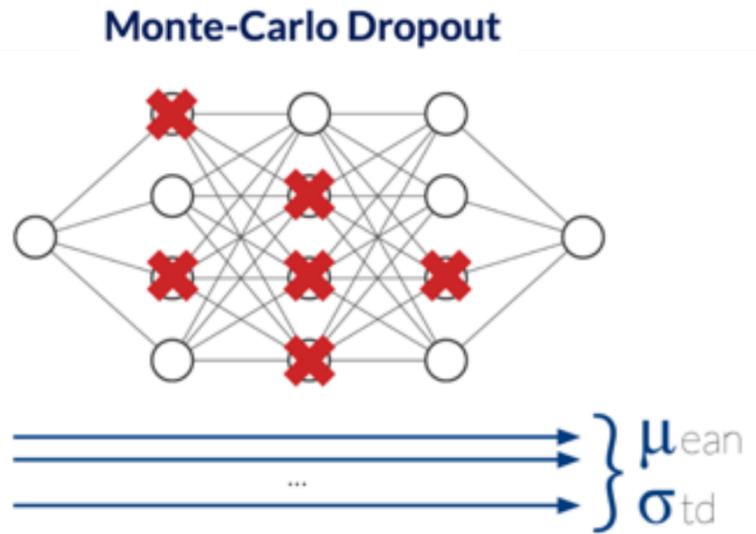
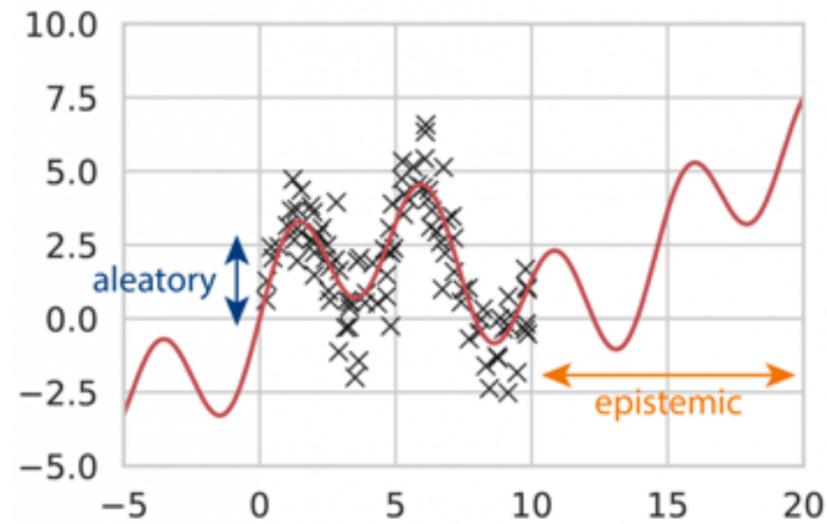
# UNCERTAINTY QUANTIFICATION

How certain is the prediction?



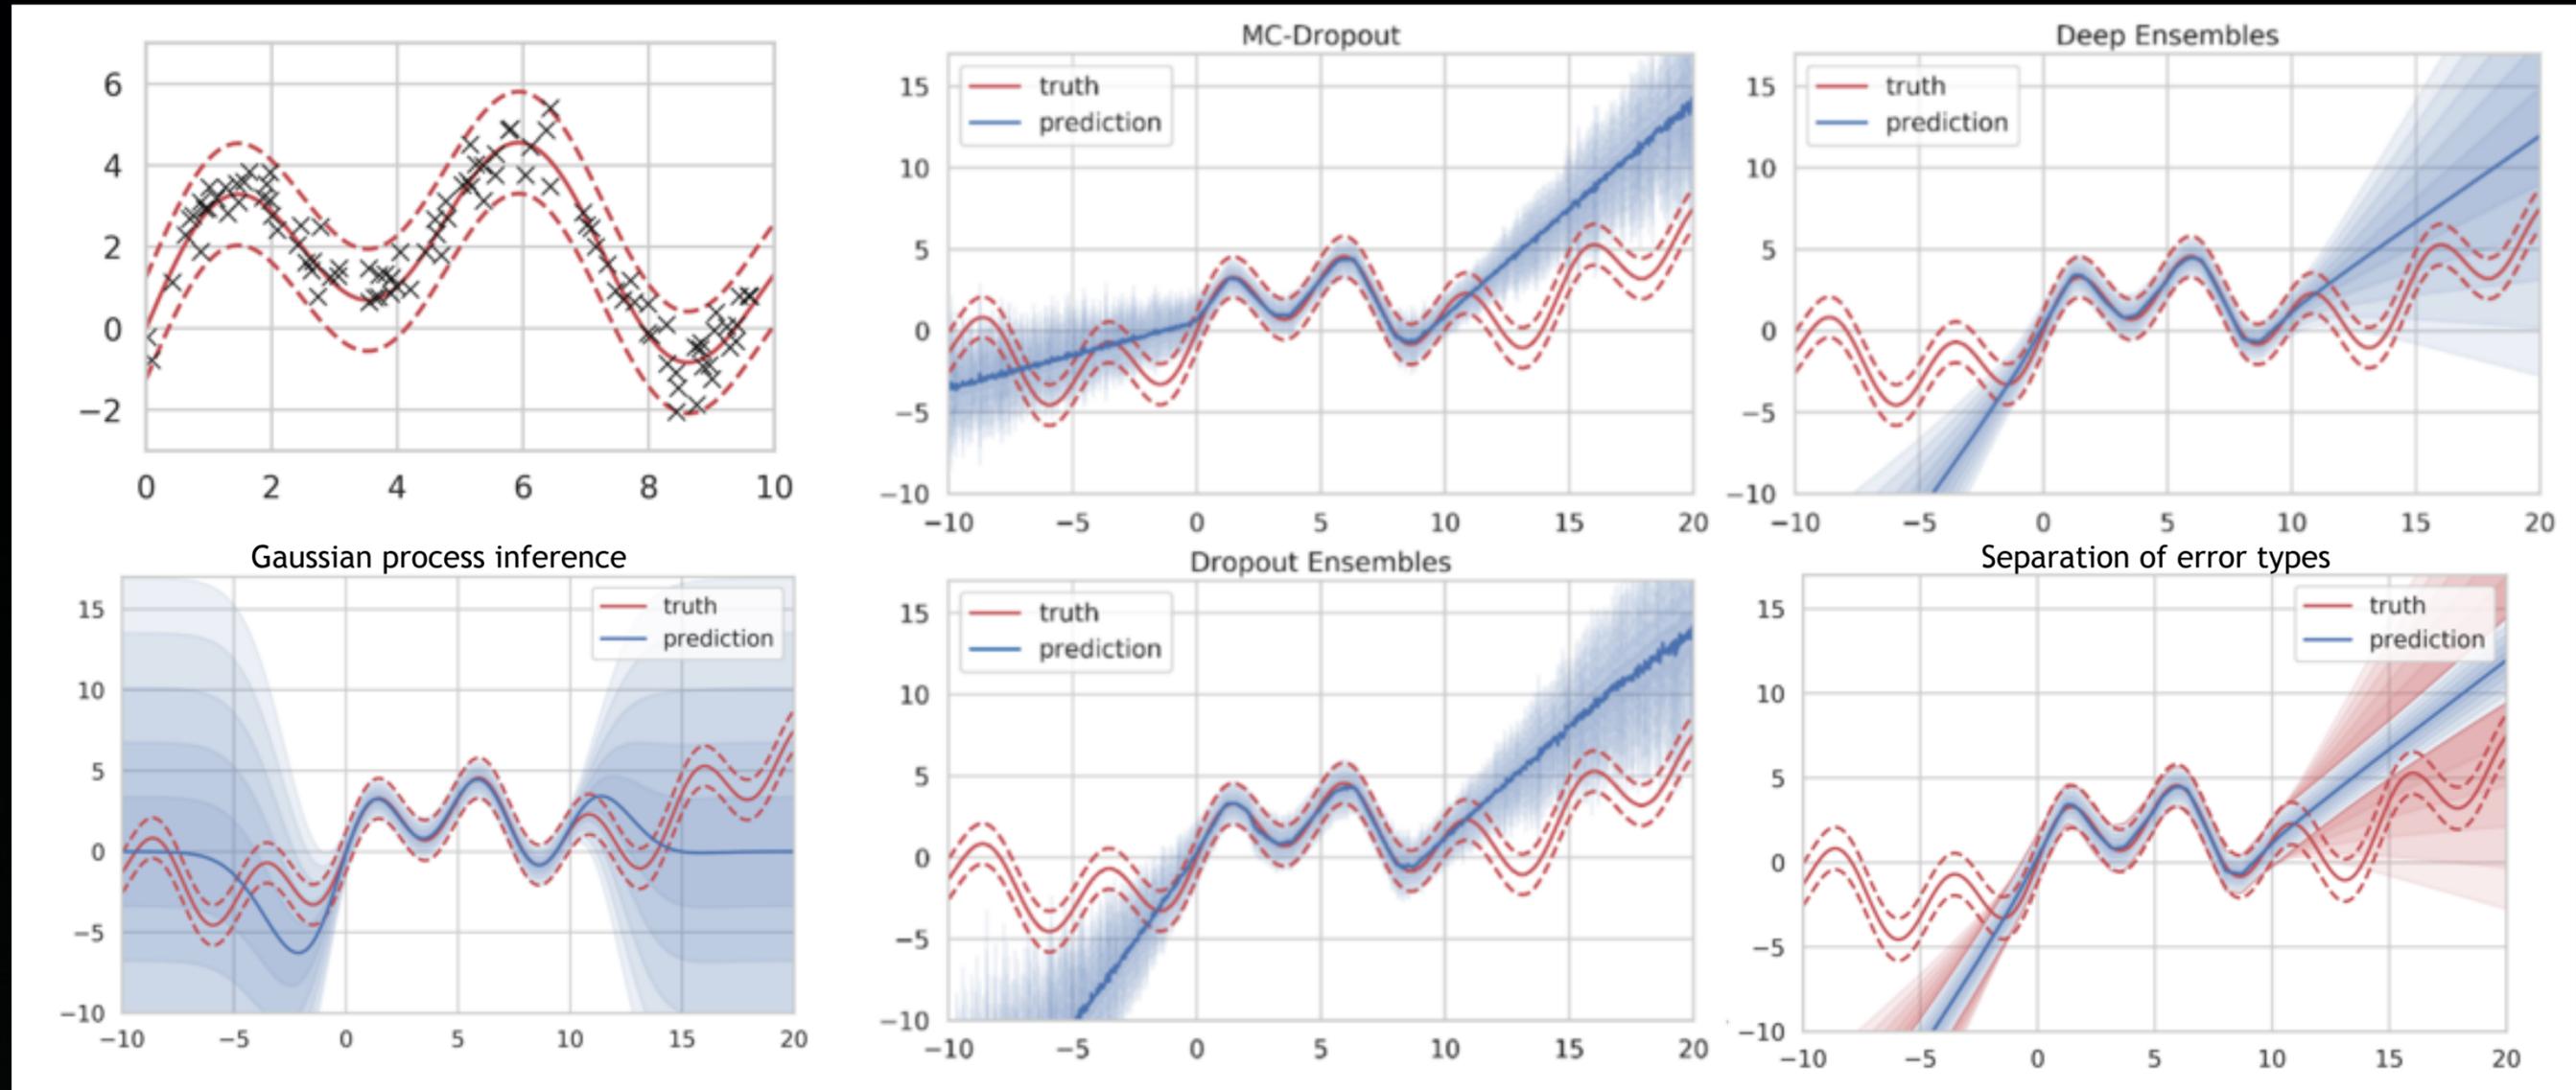
# UNCERTAINTY QUANTIFICATION

## Methods for quantifying uncertainty



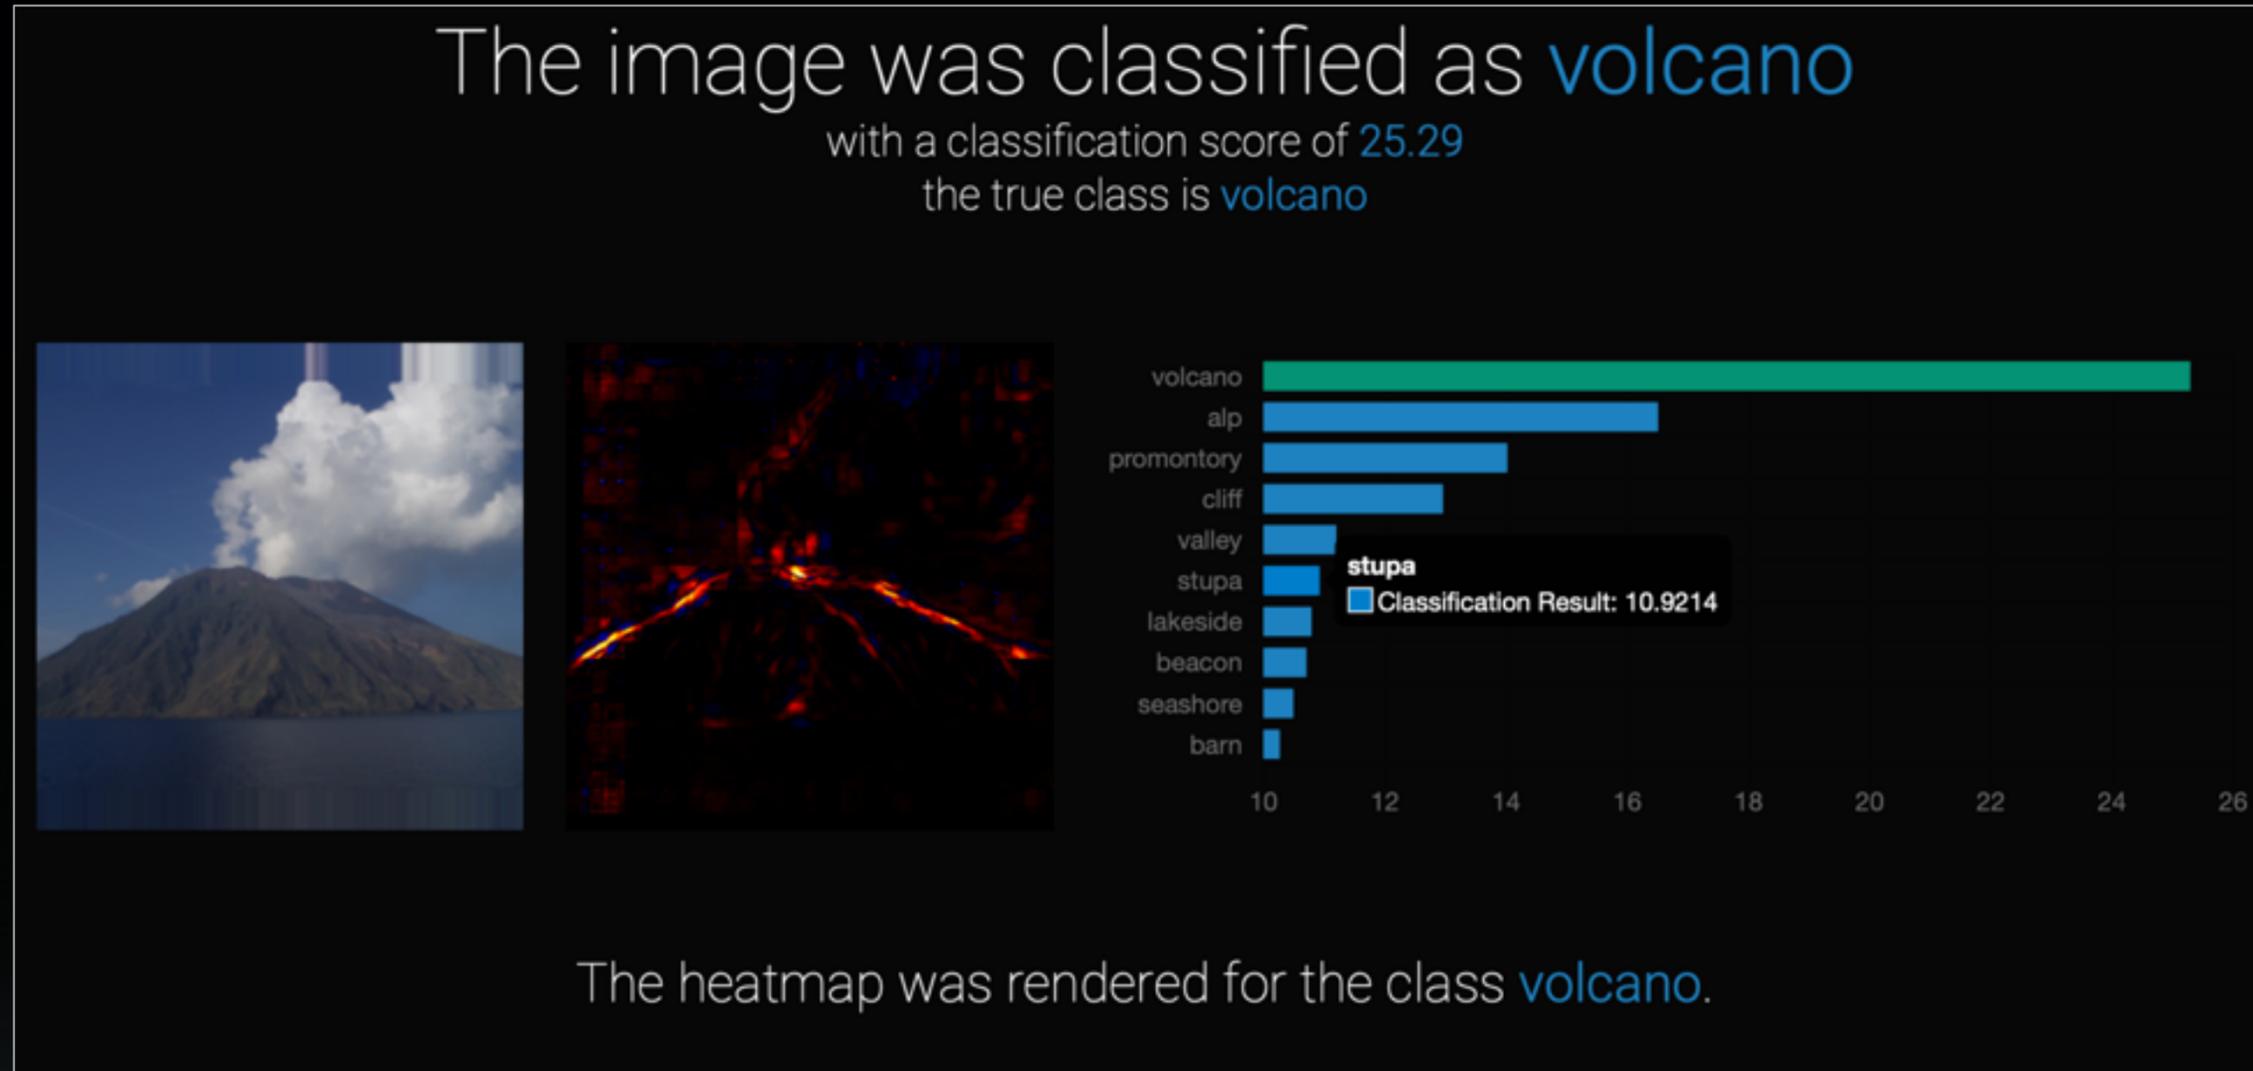
# UNCERTAINTY QUANTIFICATION

## Methods for quantifying uncertainty



# INTERPRETABILITY

What criteria were used in this prediction?



Layer-wise Relevance Propagation (LRP)

# INTERPRETABILITY

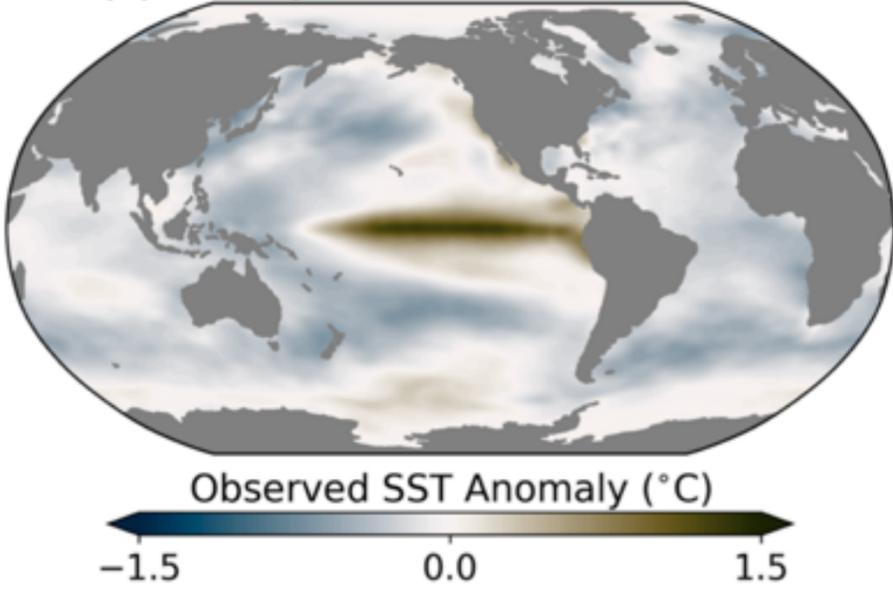
## Backwards optimization and Layerwise Relevance Propagation applied to ENSO

### Physically Interpretable Neural Networks for the Geosciences: Applications to Earth System Variability

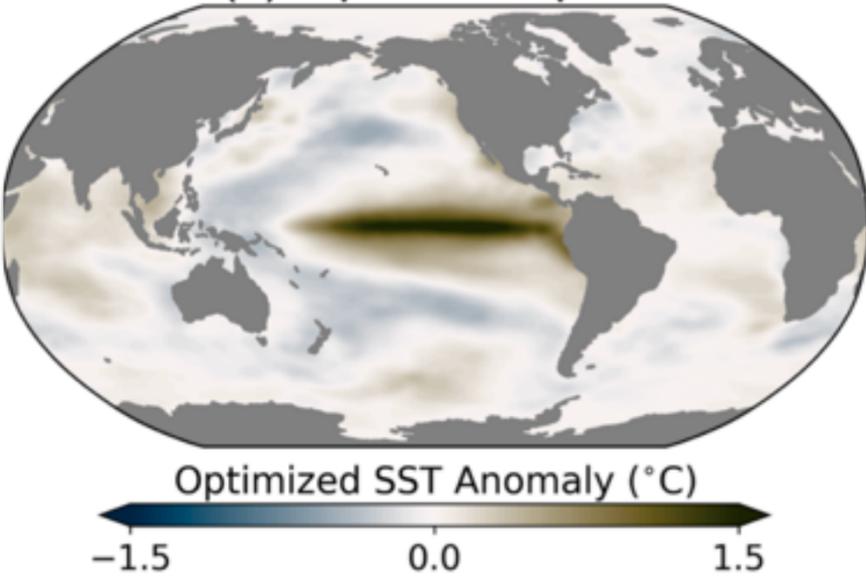
Benjamin A. Toms<sup>1</sup> , Elizabeth A. Barnes<sup>1</sup> , and Imme Ebert-Uphoff<sup>2,3</sup> 

<sup>1</sup>Department of Atmospheric Science, Colorado State University, Fort Collins, CO, USA, <sup>2</sup>Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO, USA, <sup>3</sup>Cooperative Institute for Research in the Atmosphere, Colorado State University, Fort Collins, CO, USA

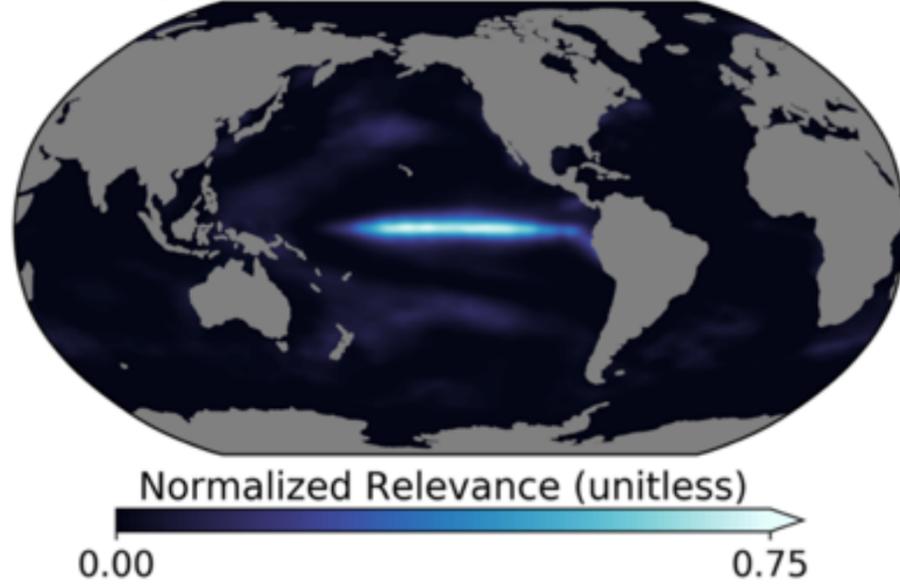
(c) Composite Observations



(a) Optimal Input



(b) Layerwise Relevance



# INTERPRETABILITY

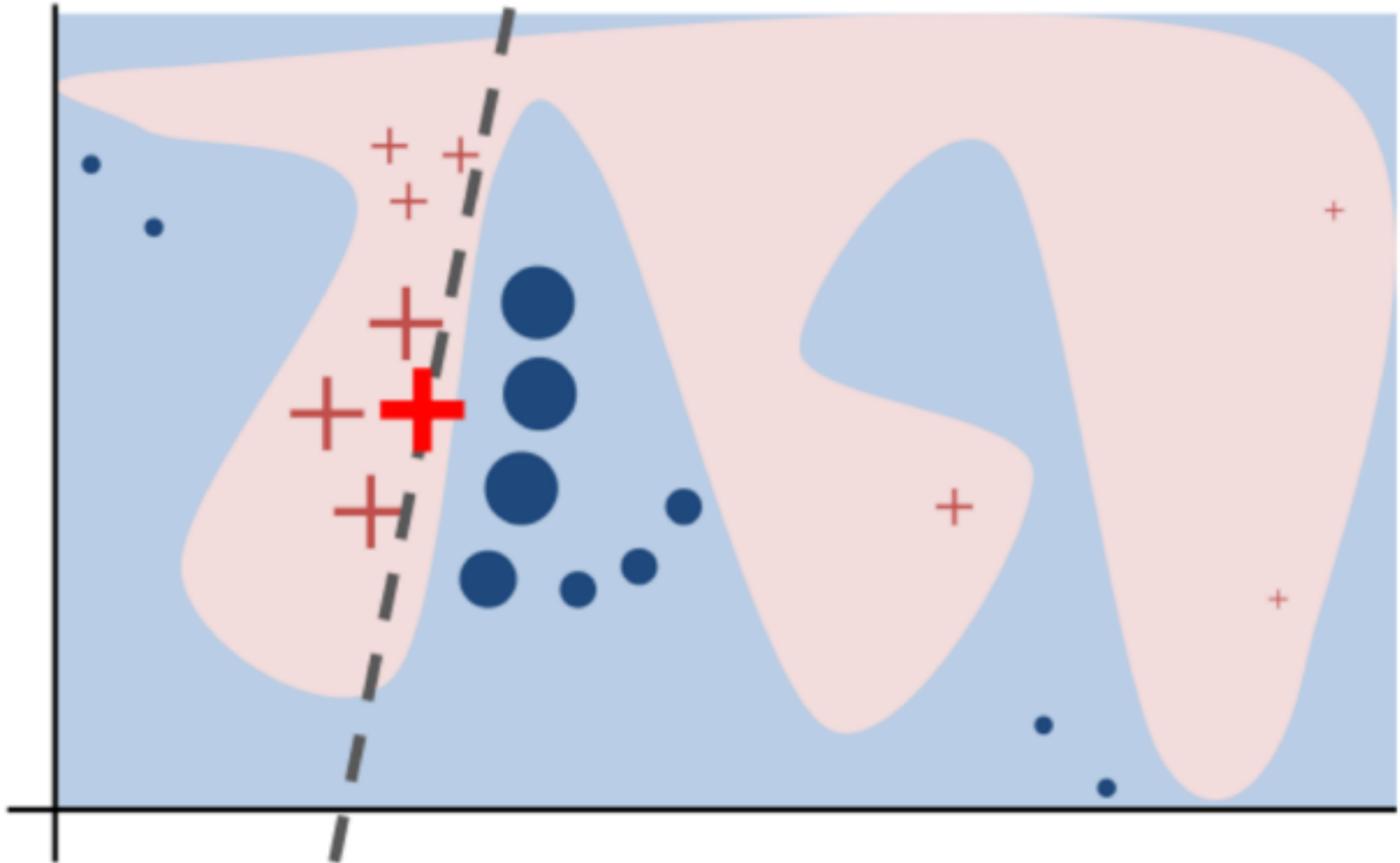
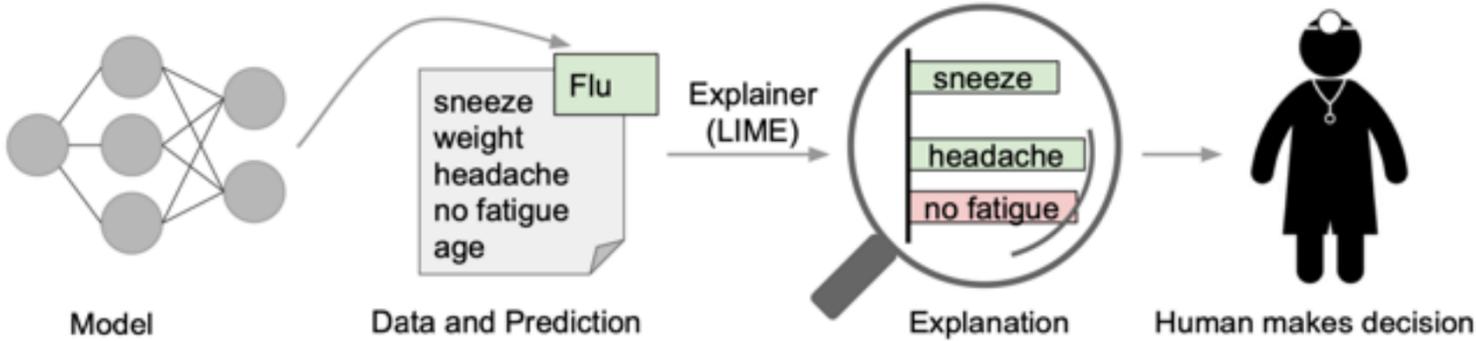
## LIME: Local interpretable model-agnostic explanations

**“Why Should I Trust You?”  
Explaining the Predictions of Any Classifier**

Marco Tulio Ribeiro  
University of Washington  
Seattle, WA 98105, USA  
marcotcr@cs.uw.edu

Sameer Singh  
University of Washington  
Seattle, WA 98105, USA  
sameer@cs.uw.edu

Carlos Guestrin  
University of Washington  
Seattle, WA 98105, USA  
guestrin@cs.uw.edu



# FORTRAN / AI COUPLING

How can I glue my AI and HPC code together?



+



# FORTRAN / AI COUPLING

Many solutions. None ideal.

Use Julia. Call Fortran

  
  
DiffEqFlux.jl  
GPUArrays

Use C++ Instead

  
8 Language Bindings  
Deep integration into Python and support for Scala, Julia, Clojure, Java, C++, R and Perl.  
GluonCV   
GluonNLP   
GluonTS 

Ok, but limited

  
  
Fortran  
Implemented layers

- Dense
- Dropout
- Batch Normalization

Missing: Native API

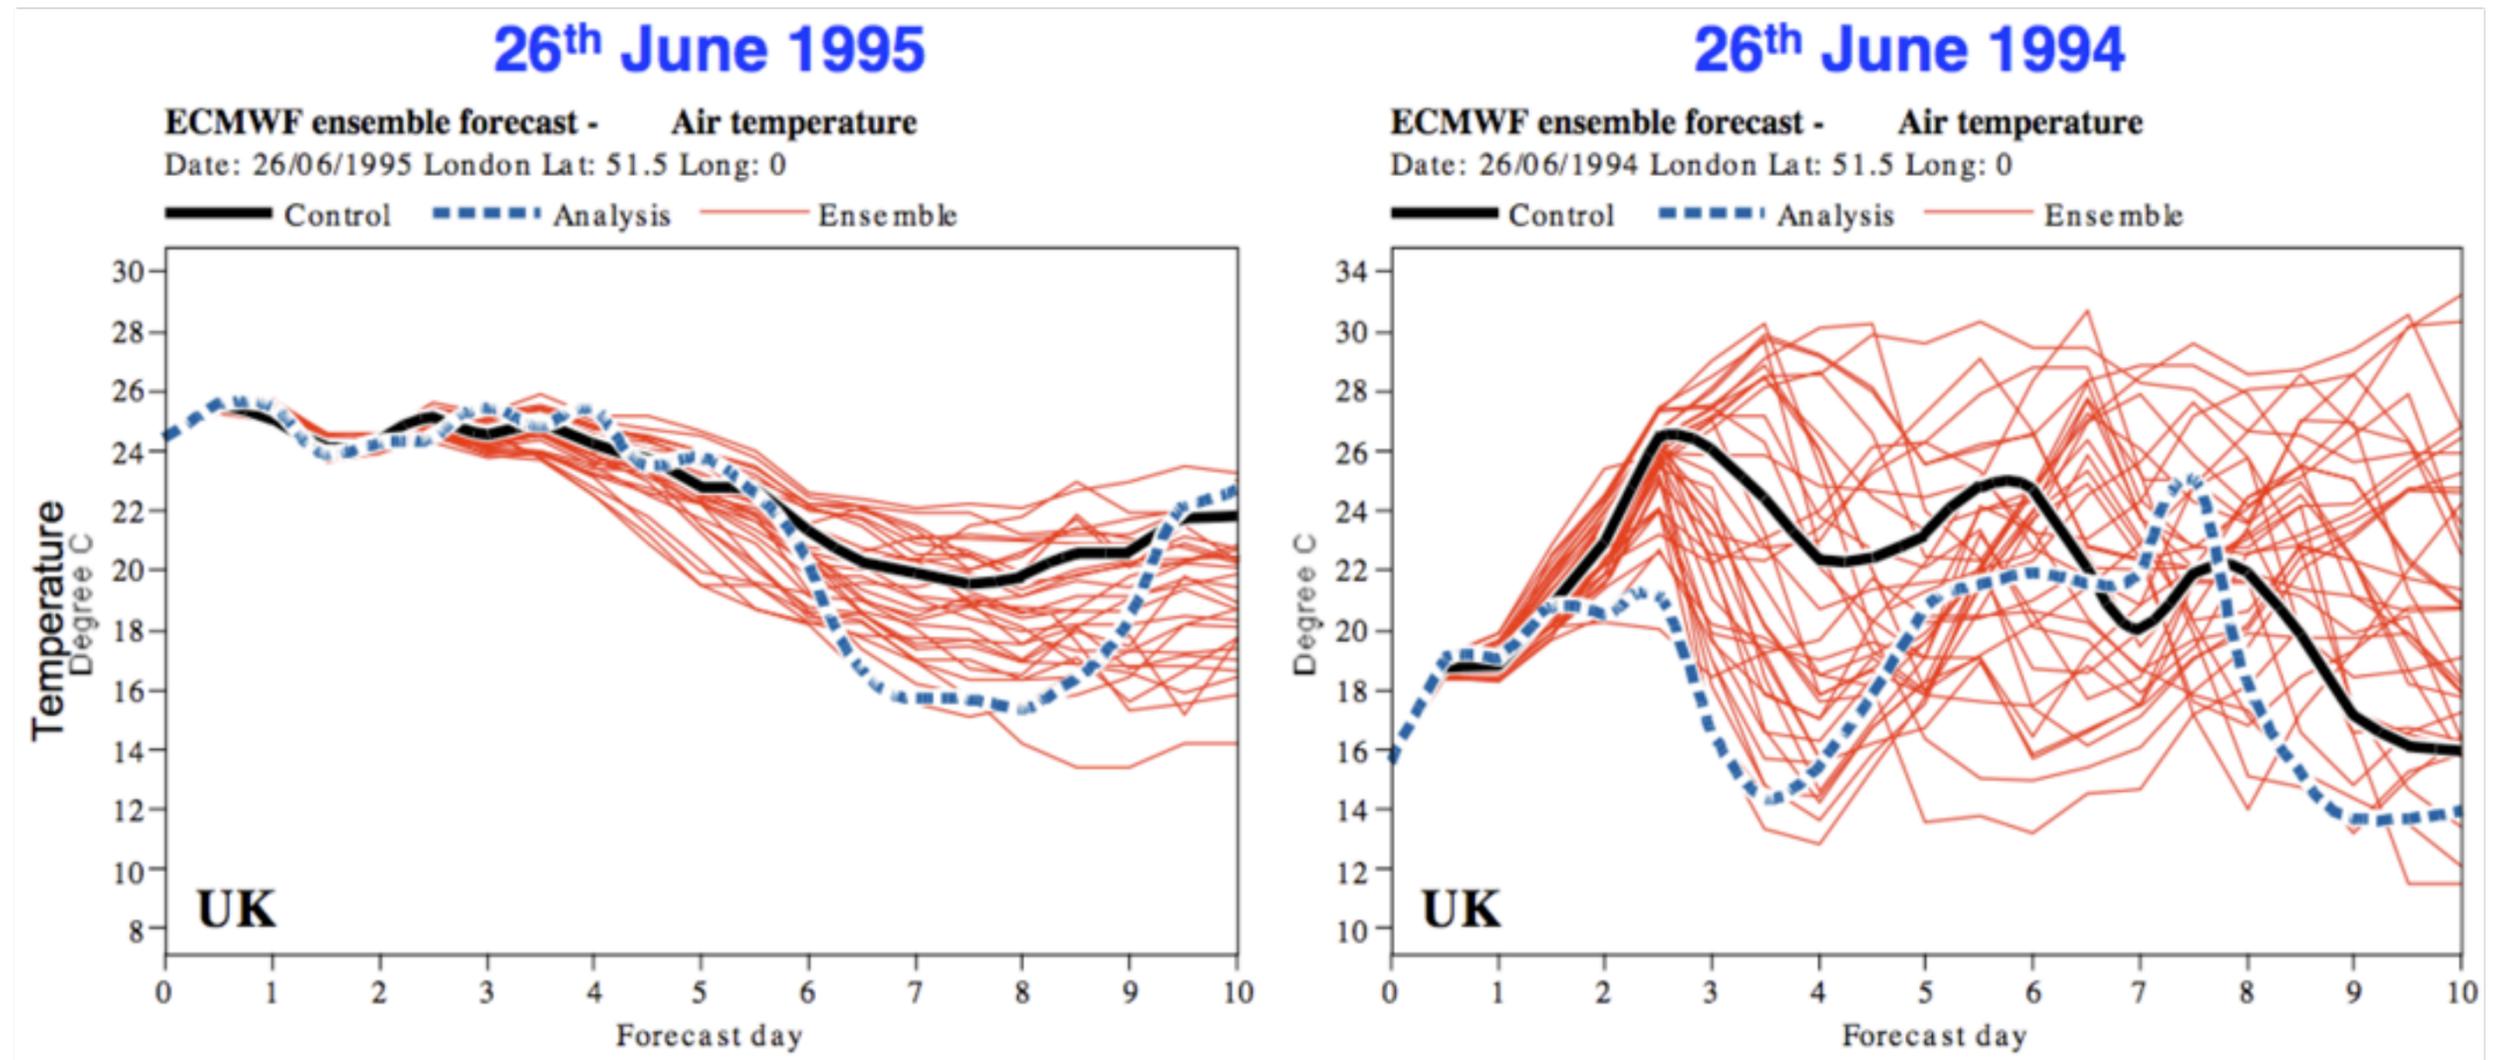


Missing: Fortran Bindings

  
  
TensorFlow

# OVERCOMING LOSS OF DYNAMIC RANGE

Ensemble mean always has less variability than the ensemble members



# OVERCOMING LOSS OF DYNAMIC RANGE

## Stochastic Parameterizations

### Machine Learning for Stochastic Parameterization: Generative Adversarial Networks in the Lorenz '96 Model

David John Gagneil<sup>1</sup> , Hannah M. Christensen<sup>1,2</sup> , Aneesh C. Subramanian<sup>3</sup> ,  
and Adam H. Monahan<sup>4</sup> 

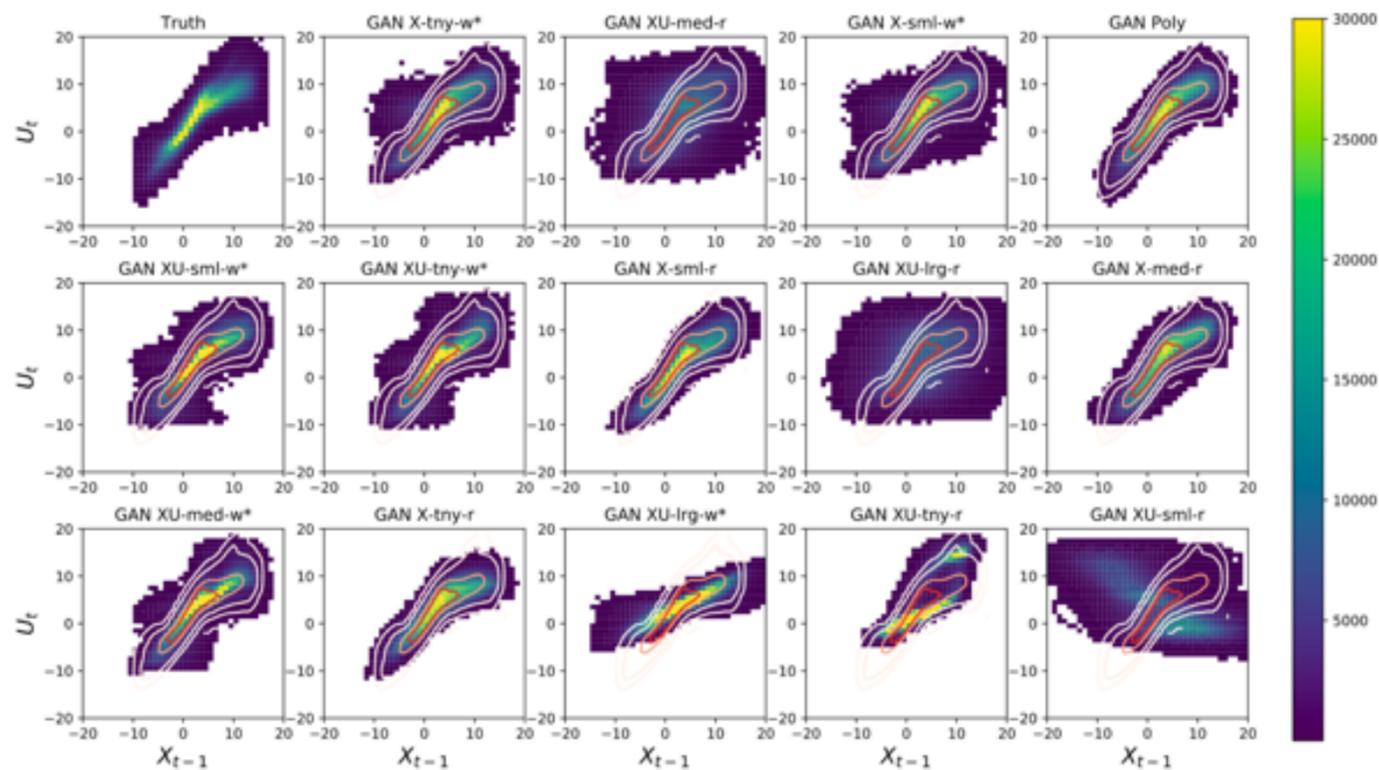
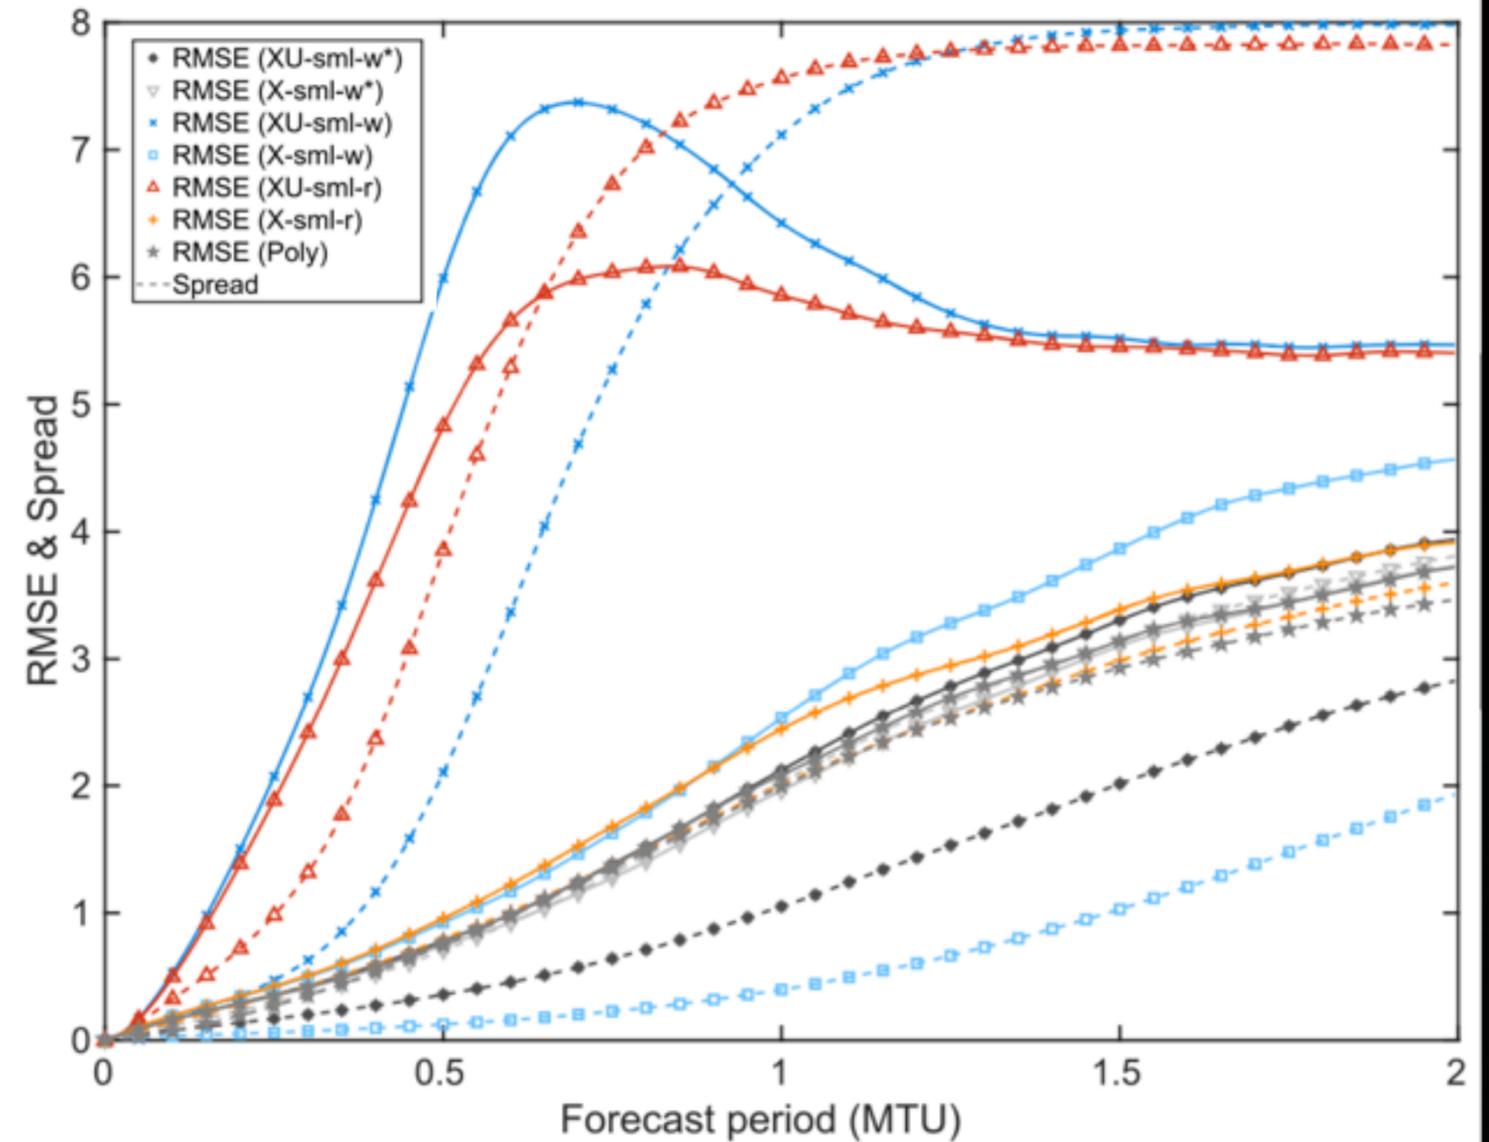


Figure 3. Joint distributions (2-D histograms) of  $X_{t-1}$  and  $U_t$  for each GAN configuration. The truth joint distribution is overlaid in red contours on each forecast model distribution. The distributions are ordered from left to right descending in terms of their relative total marginal Hellinger distances.





# TRENDS AND BREAKTHROUGHS

# SELF-SUPERVISION

Babies learn about the world without large labelled datasets. AI can too.

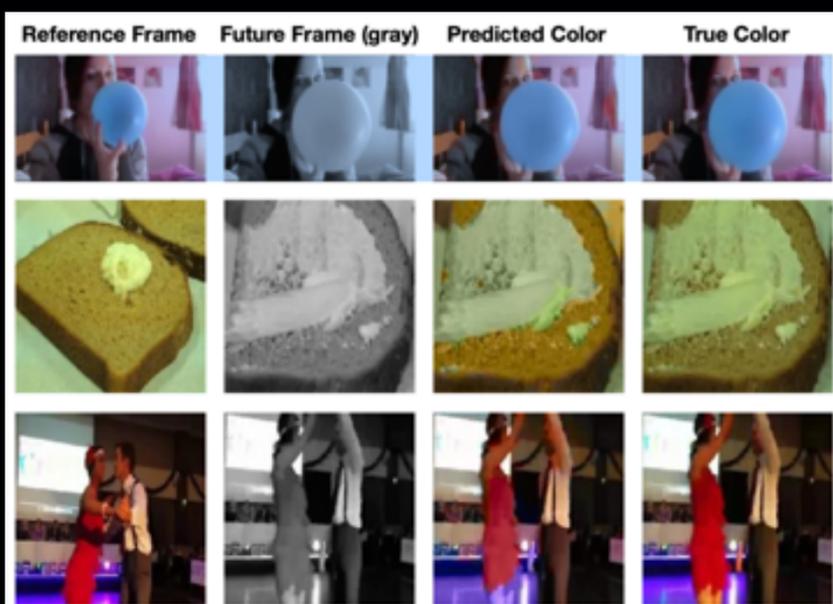




# SELF-SUPERVISION

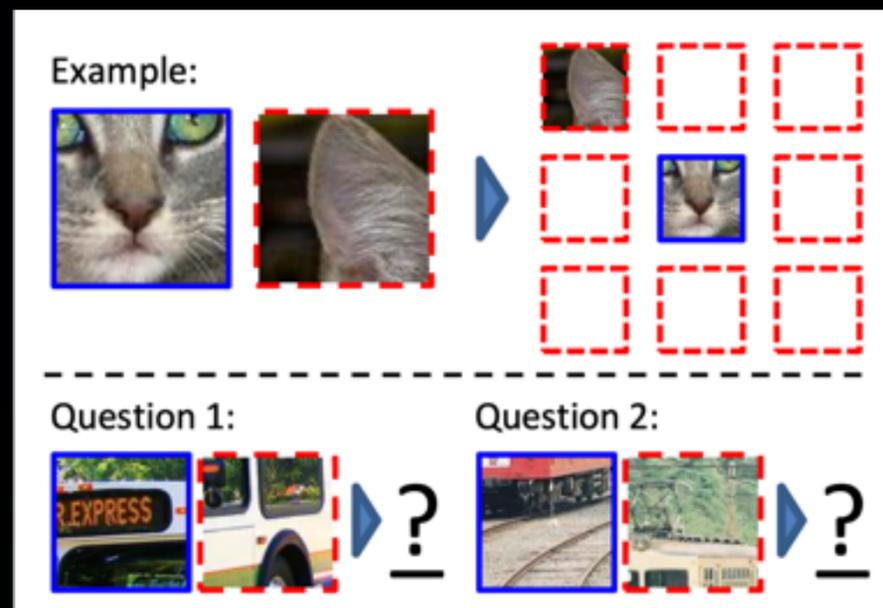
Pretext tasks build up an internal representation

PREDICT COLOR

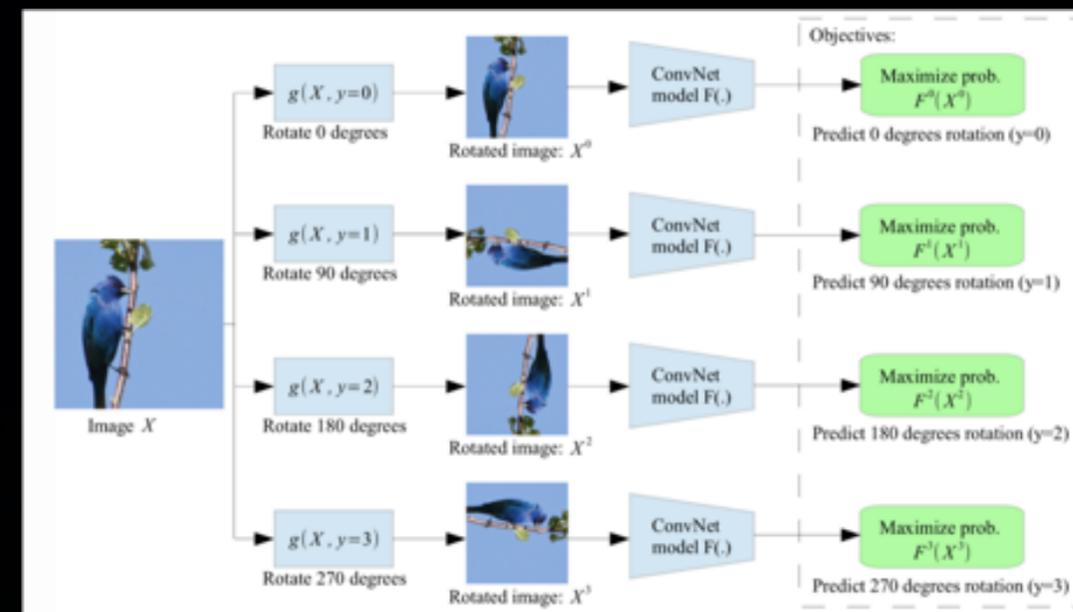


<https://arxiv.org/pdf/1806.09594.pdf>

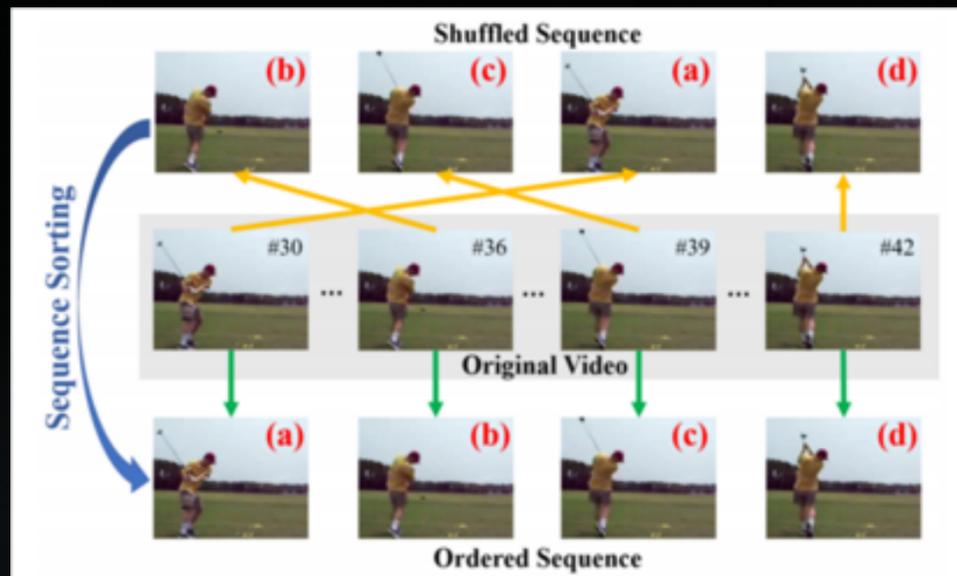
PREDICT SPATIAL RELATIONSHIPS



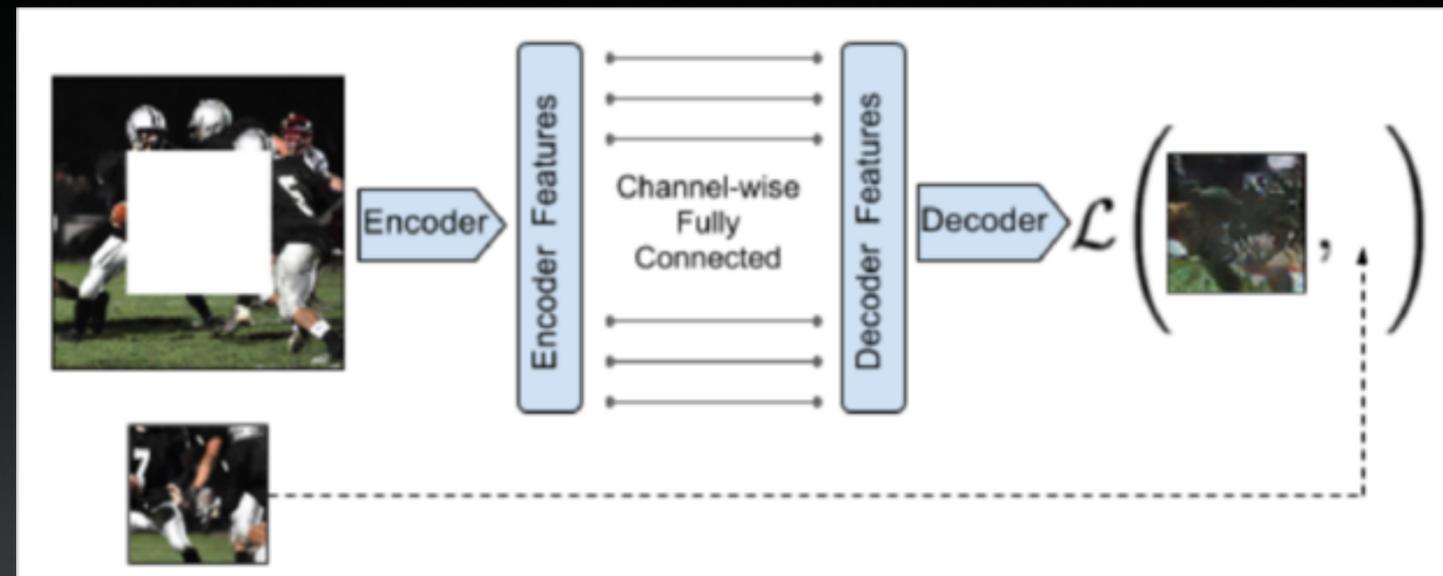
PREDICT ORIENTATION



PREDICT TEMPORAL ORDER

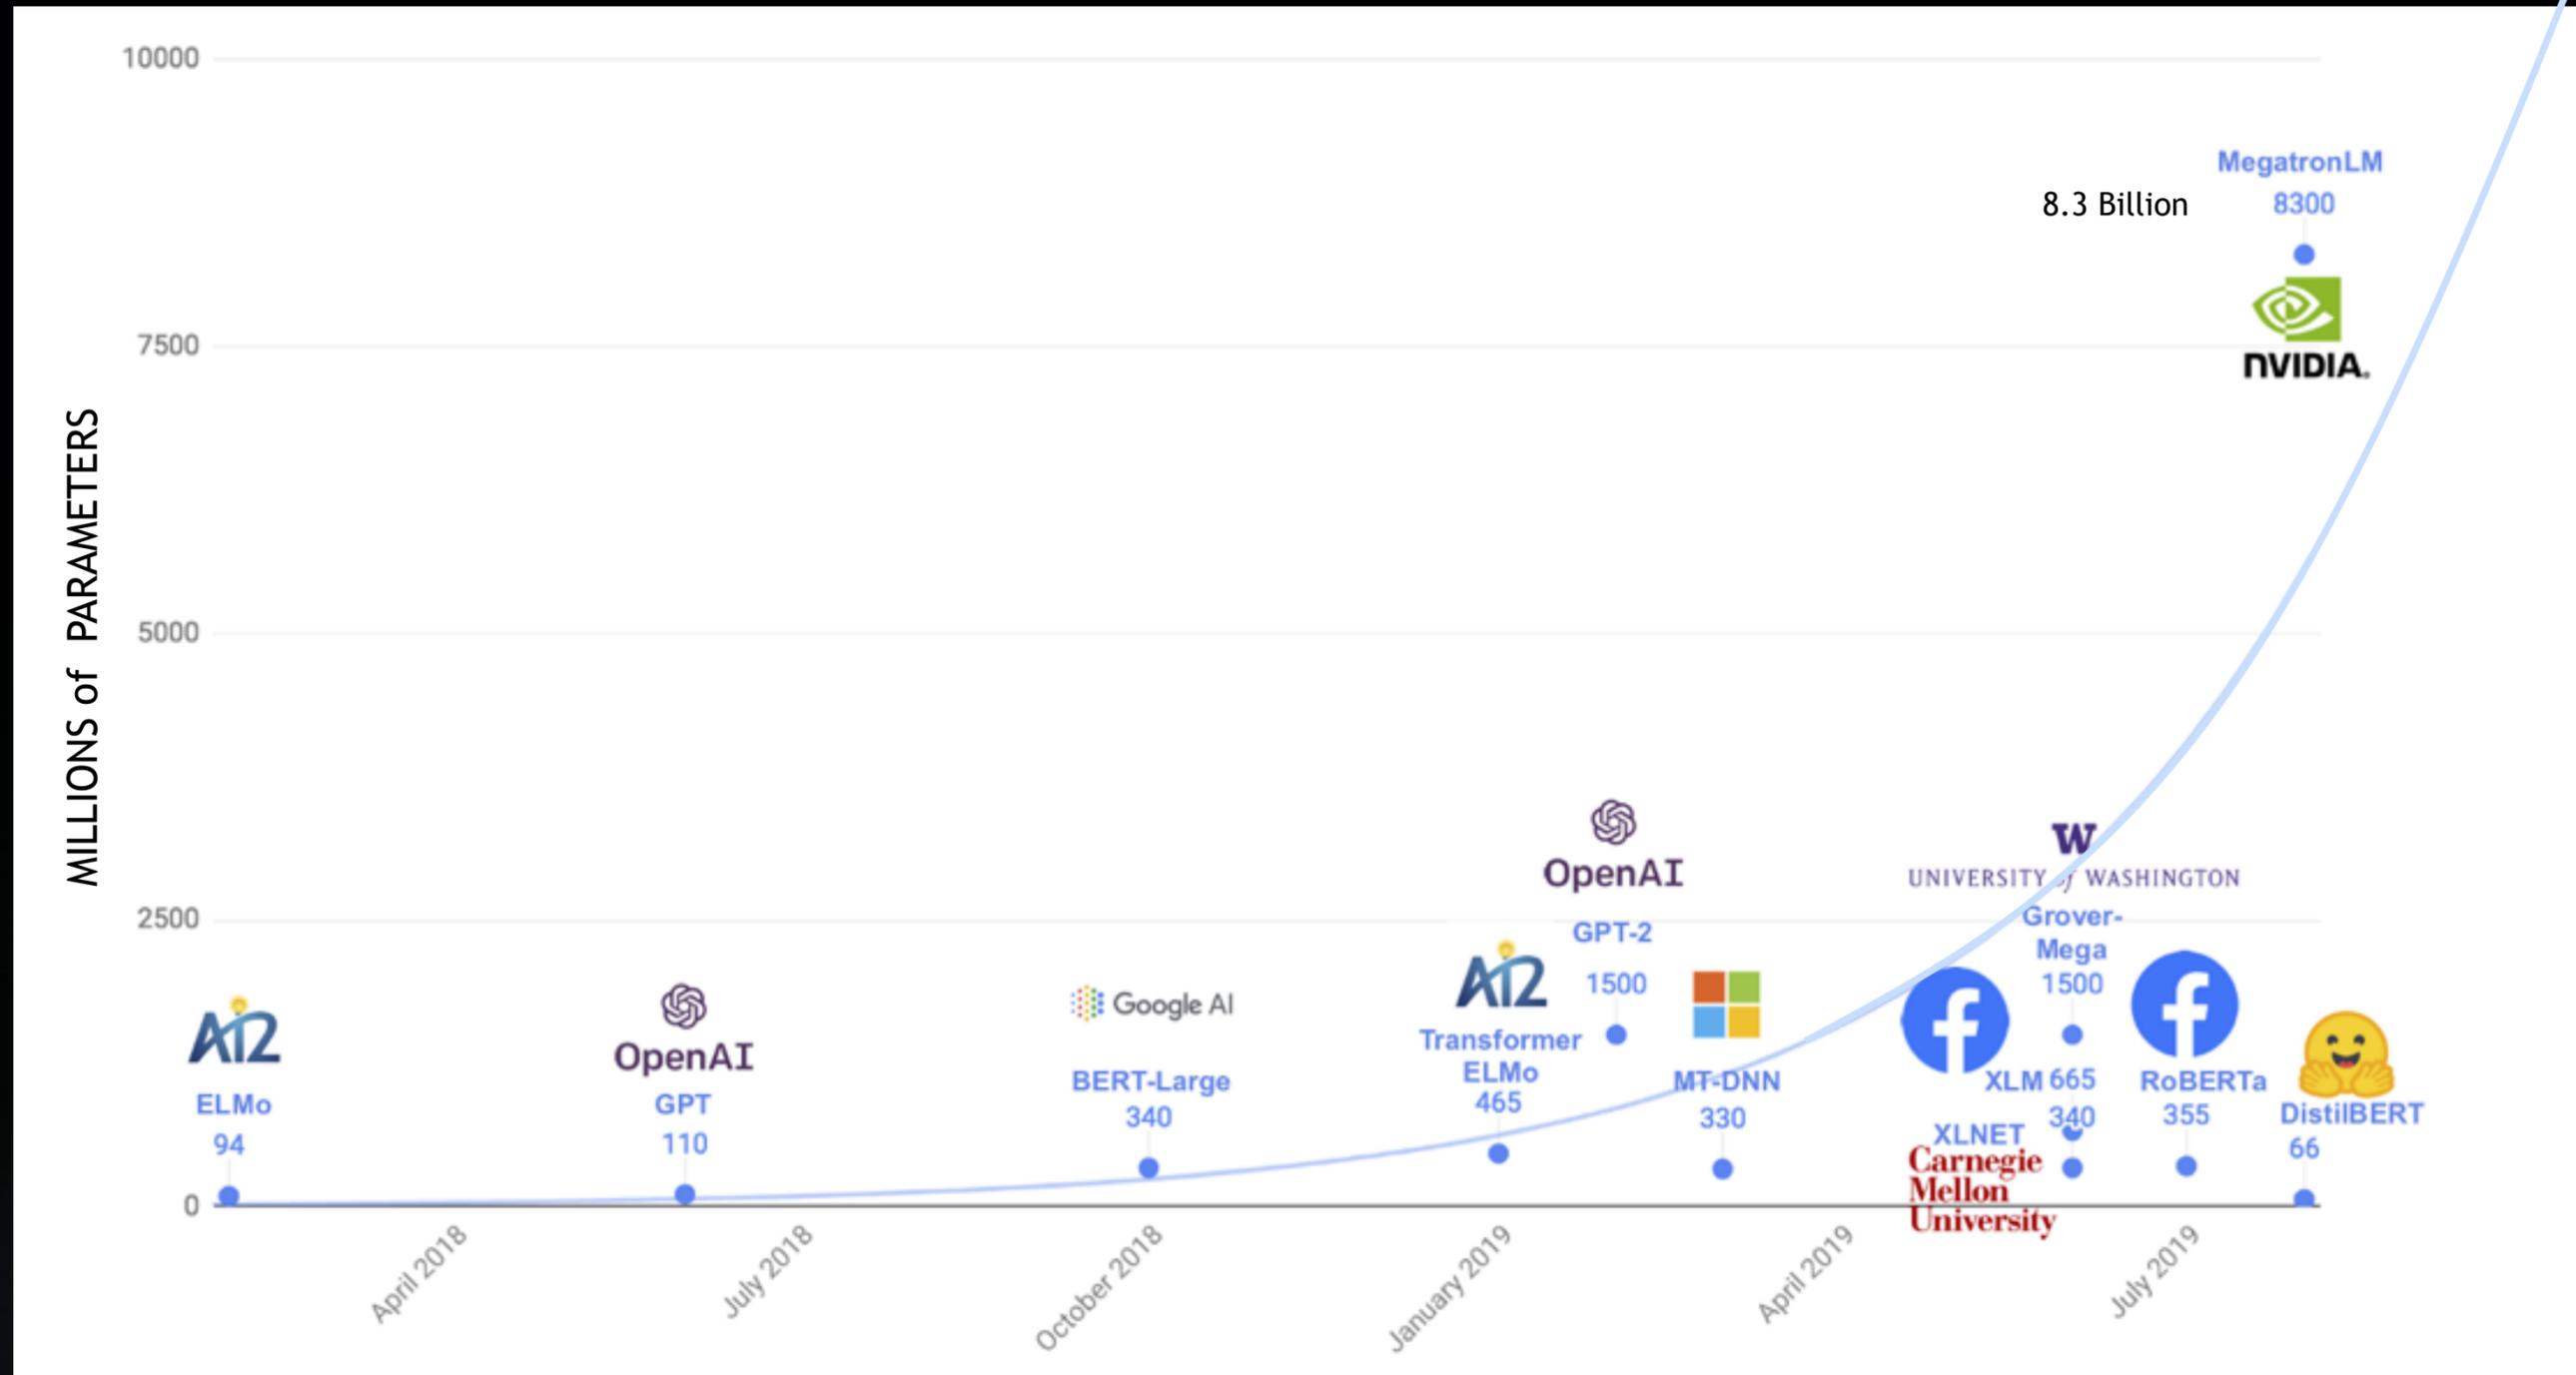


PREDICT MISSING PIECES



# THE TRANSFORMER

Has enabled a series of enormous language models



# GPT-3

(July 2020) Current king of the languages models. Based upon the transformer

## Language Models are Few-Shot Learners

The three settings we explore for in-context learning

### Zero-shot

The model predicts the answer given only a natural language description of the task. No gradient updates are performed.

```
1 Translate English to French: ← task description
2 cheese => ..... ← prompt
```

### One-shot

In addition to the task description, the model sees a single example of the task. No gradient updates are performed.

```
1 Translate English to French: ← task description
2 sea otter => loutre de mer ← example
3 cheese => ..... ← prompt
```

### Few-shot

In addition to the task description, the model sees a few examples of the task. No gradient updates are performed.

```
1 Translate English to French: ← task description
2 sea otter => loutre de mer ← examples
3 peppermint => menthe poivrée ←
4 plush girafe => girafe peluche ←
5 cheese => ..... ← prompt
```

What is the average number of influencers each user is subscribed to?

```
1 SELECT
2   avg(count)
3 FROM
4   (
5     SELECT
6       user_id,
7       count(*)
8     FROM
9       subscribers
10    GROUP BY
11    user_id
```

Equation description

integral from a to b of f(t) with respect to t = F of b minus F of a

Translate

$$\int_a^b f(t) dt = \int_a^b \frac{F(b) - F(a)}{t} dt$$

How many carbon atoms are there in Benzene?

There are 6 carbon atoms in Benzene.

[More info on this / Show more results](#)

large text that says "WELCOME TO MY NEWSLETTER" and a blue button that sa Generate

```
<h1 style={{fontSize: 50, color: 'white'}}>WELCOME TO MY NEWSLETTER</h1><button
style={{color: 'white', backgroundColor: 'blue'}}>Subscribe</button>
```

WELCOME TO MY NEWSLETTER

Subscribe

Build Keras Models

Build a model to classify images into 5 groups. The dataset has 25000 images, with an input shape of 500x500.

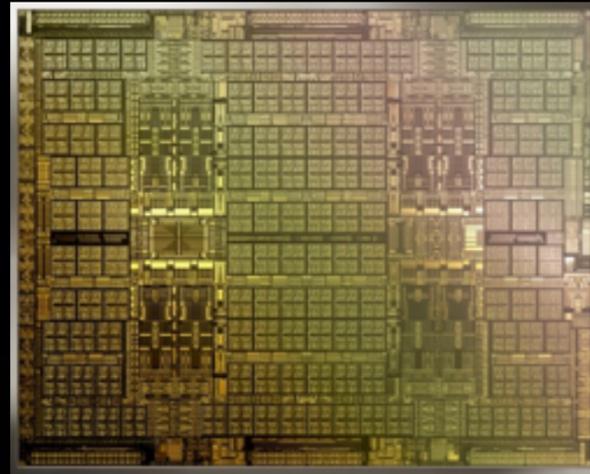
Generate Model

```
from keras.models import Sequential
from keras.layers import Conv2D, MaxPooling2D,
Dropout, Flatten, Dense, Activation,
BatchNormalization
model = Sequential()
model.add(Conv2D(32, (5, 5)))
```

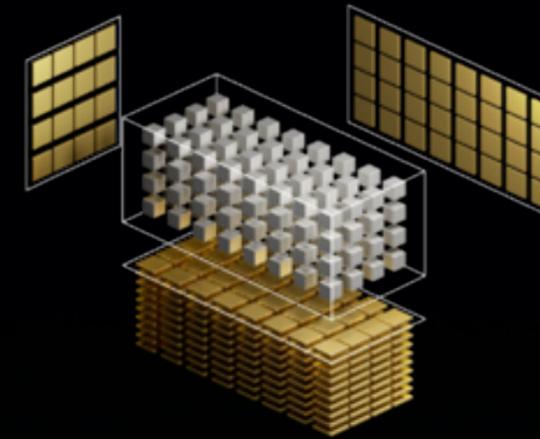


# NEW HARDWARE AND SOFTWARE TOOLS

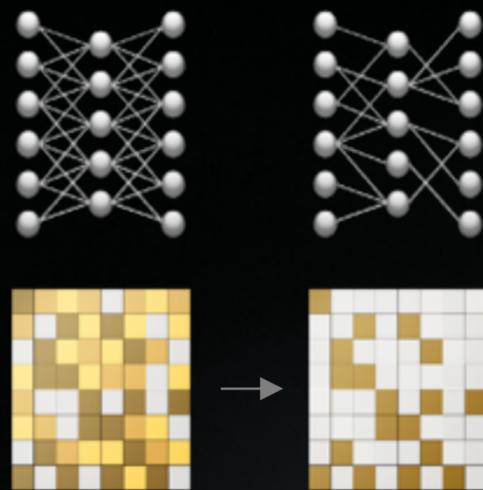
# 5 MIRACLES OF A100



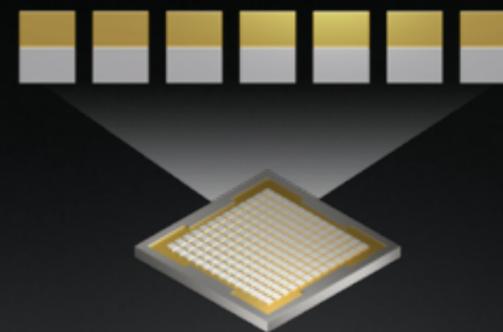
**NVIDIA Ampere Architecture**  
World's Largest 7nm chip  
54B XTORS, HBM2



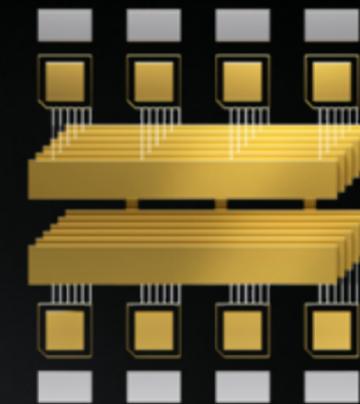
**3rd Gen Tensor Cores**  
Faster, Flexible, Easier to use  
20x AI Perf with TF32  
2.5x HPC Perf



**New Sparsity Acceleration**  
Harness Sparsity in AI Models  
2x AI Performance



**New Multi-Instance GPU**  
Optimal utilization with right sized GPU  
7x Simultaneous Instances per GPU



**3rd Gen NVLINK and NVSWITCH**  
Efficient Scaling to Enable Super GPU  
2X More Bandwidth

# OMNIVERSE

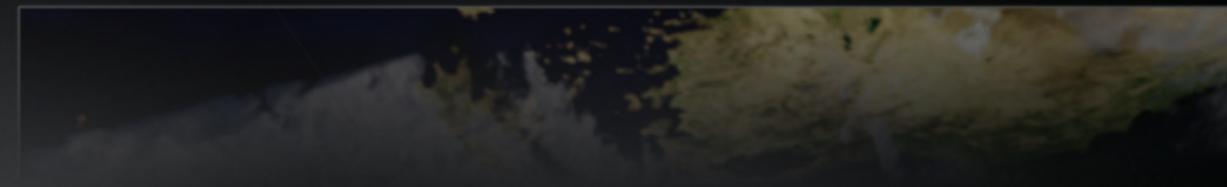
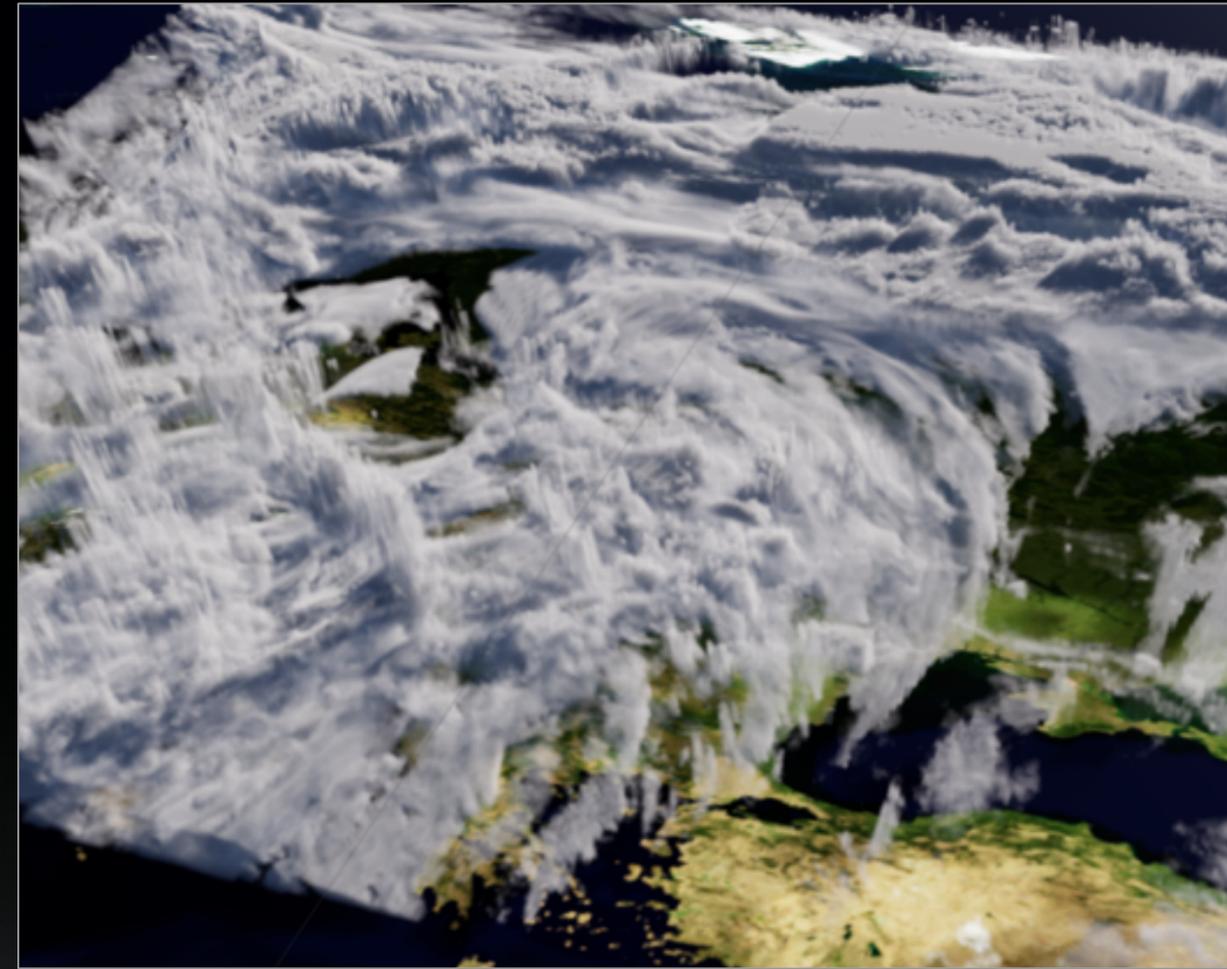
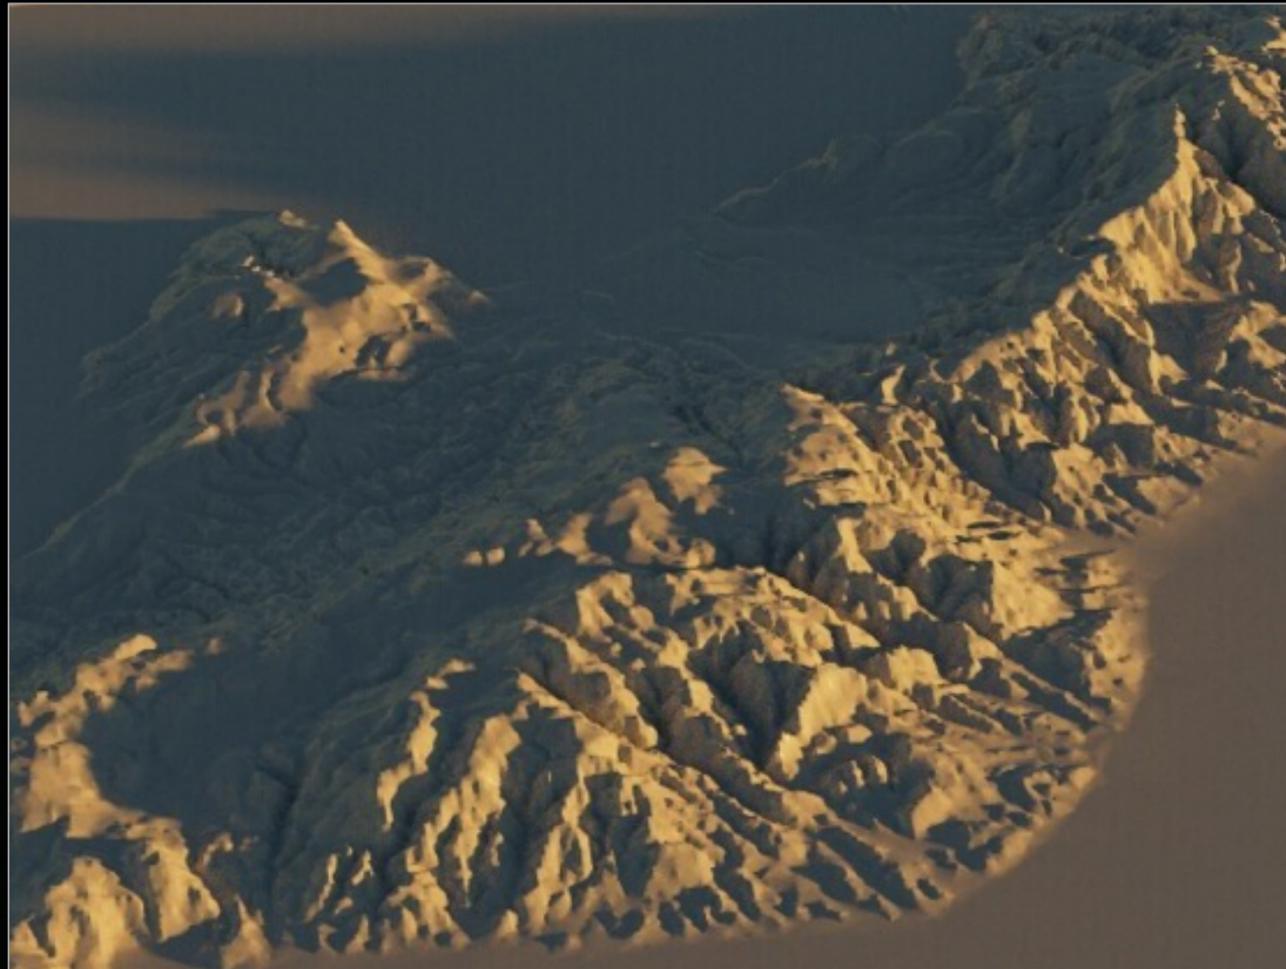
Interactive Raytracing for Data Visualization And Graphical User Interfaces



<https://developer.nvidia.com/nvidia-omniverse-platform>

# OMNIVERSE

Interactive Raytracing for Data Visualization And Graphical User Interfaces



# PYTORCH LIGHTNING

API to standardize and accelerate your PyTorch models

## Common Use Cases

- 16-bit training
- Computing cluster (SLURM)
- Child Modules
- Debugging
- Experiment Logging
- Experiment Reporting
- Early stopping
- Fast Training
- Model Hooks
- Hyperparameters
- Learning Rate Finder
- Multi-GPU training
- Multiple Datasets
- Saving and loading weights
- Optimization
- Performance and Bottleneck Profiler
- Single GPU Training
- Sequential Data
- Training Tricks
- Transfer Learning
- TPU support
- Test set
- Inference in Production

 PyTorch



# PyTorch Lightning

Scale your models. Write less boilerplate.

The lightweight PyTorch wrapper for ML researchers

# PYTORCH-LIGHTNING BOLTS

Lightning implementations of popular models, optimized for GPUs



## PyTorch Lightning Bolts

Models and model components

```
from pl_bolts.models import VAE, LogisticRegression, ImageGPT
from pl_bolts.models.resnets import resnet50
```

SOTA models

```
from pl_bolts.models.self_supervised import SimCLR
from pl_bolts.models.self_supervised import SwAV
```

Callbacks

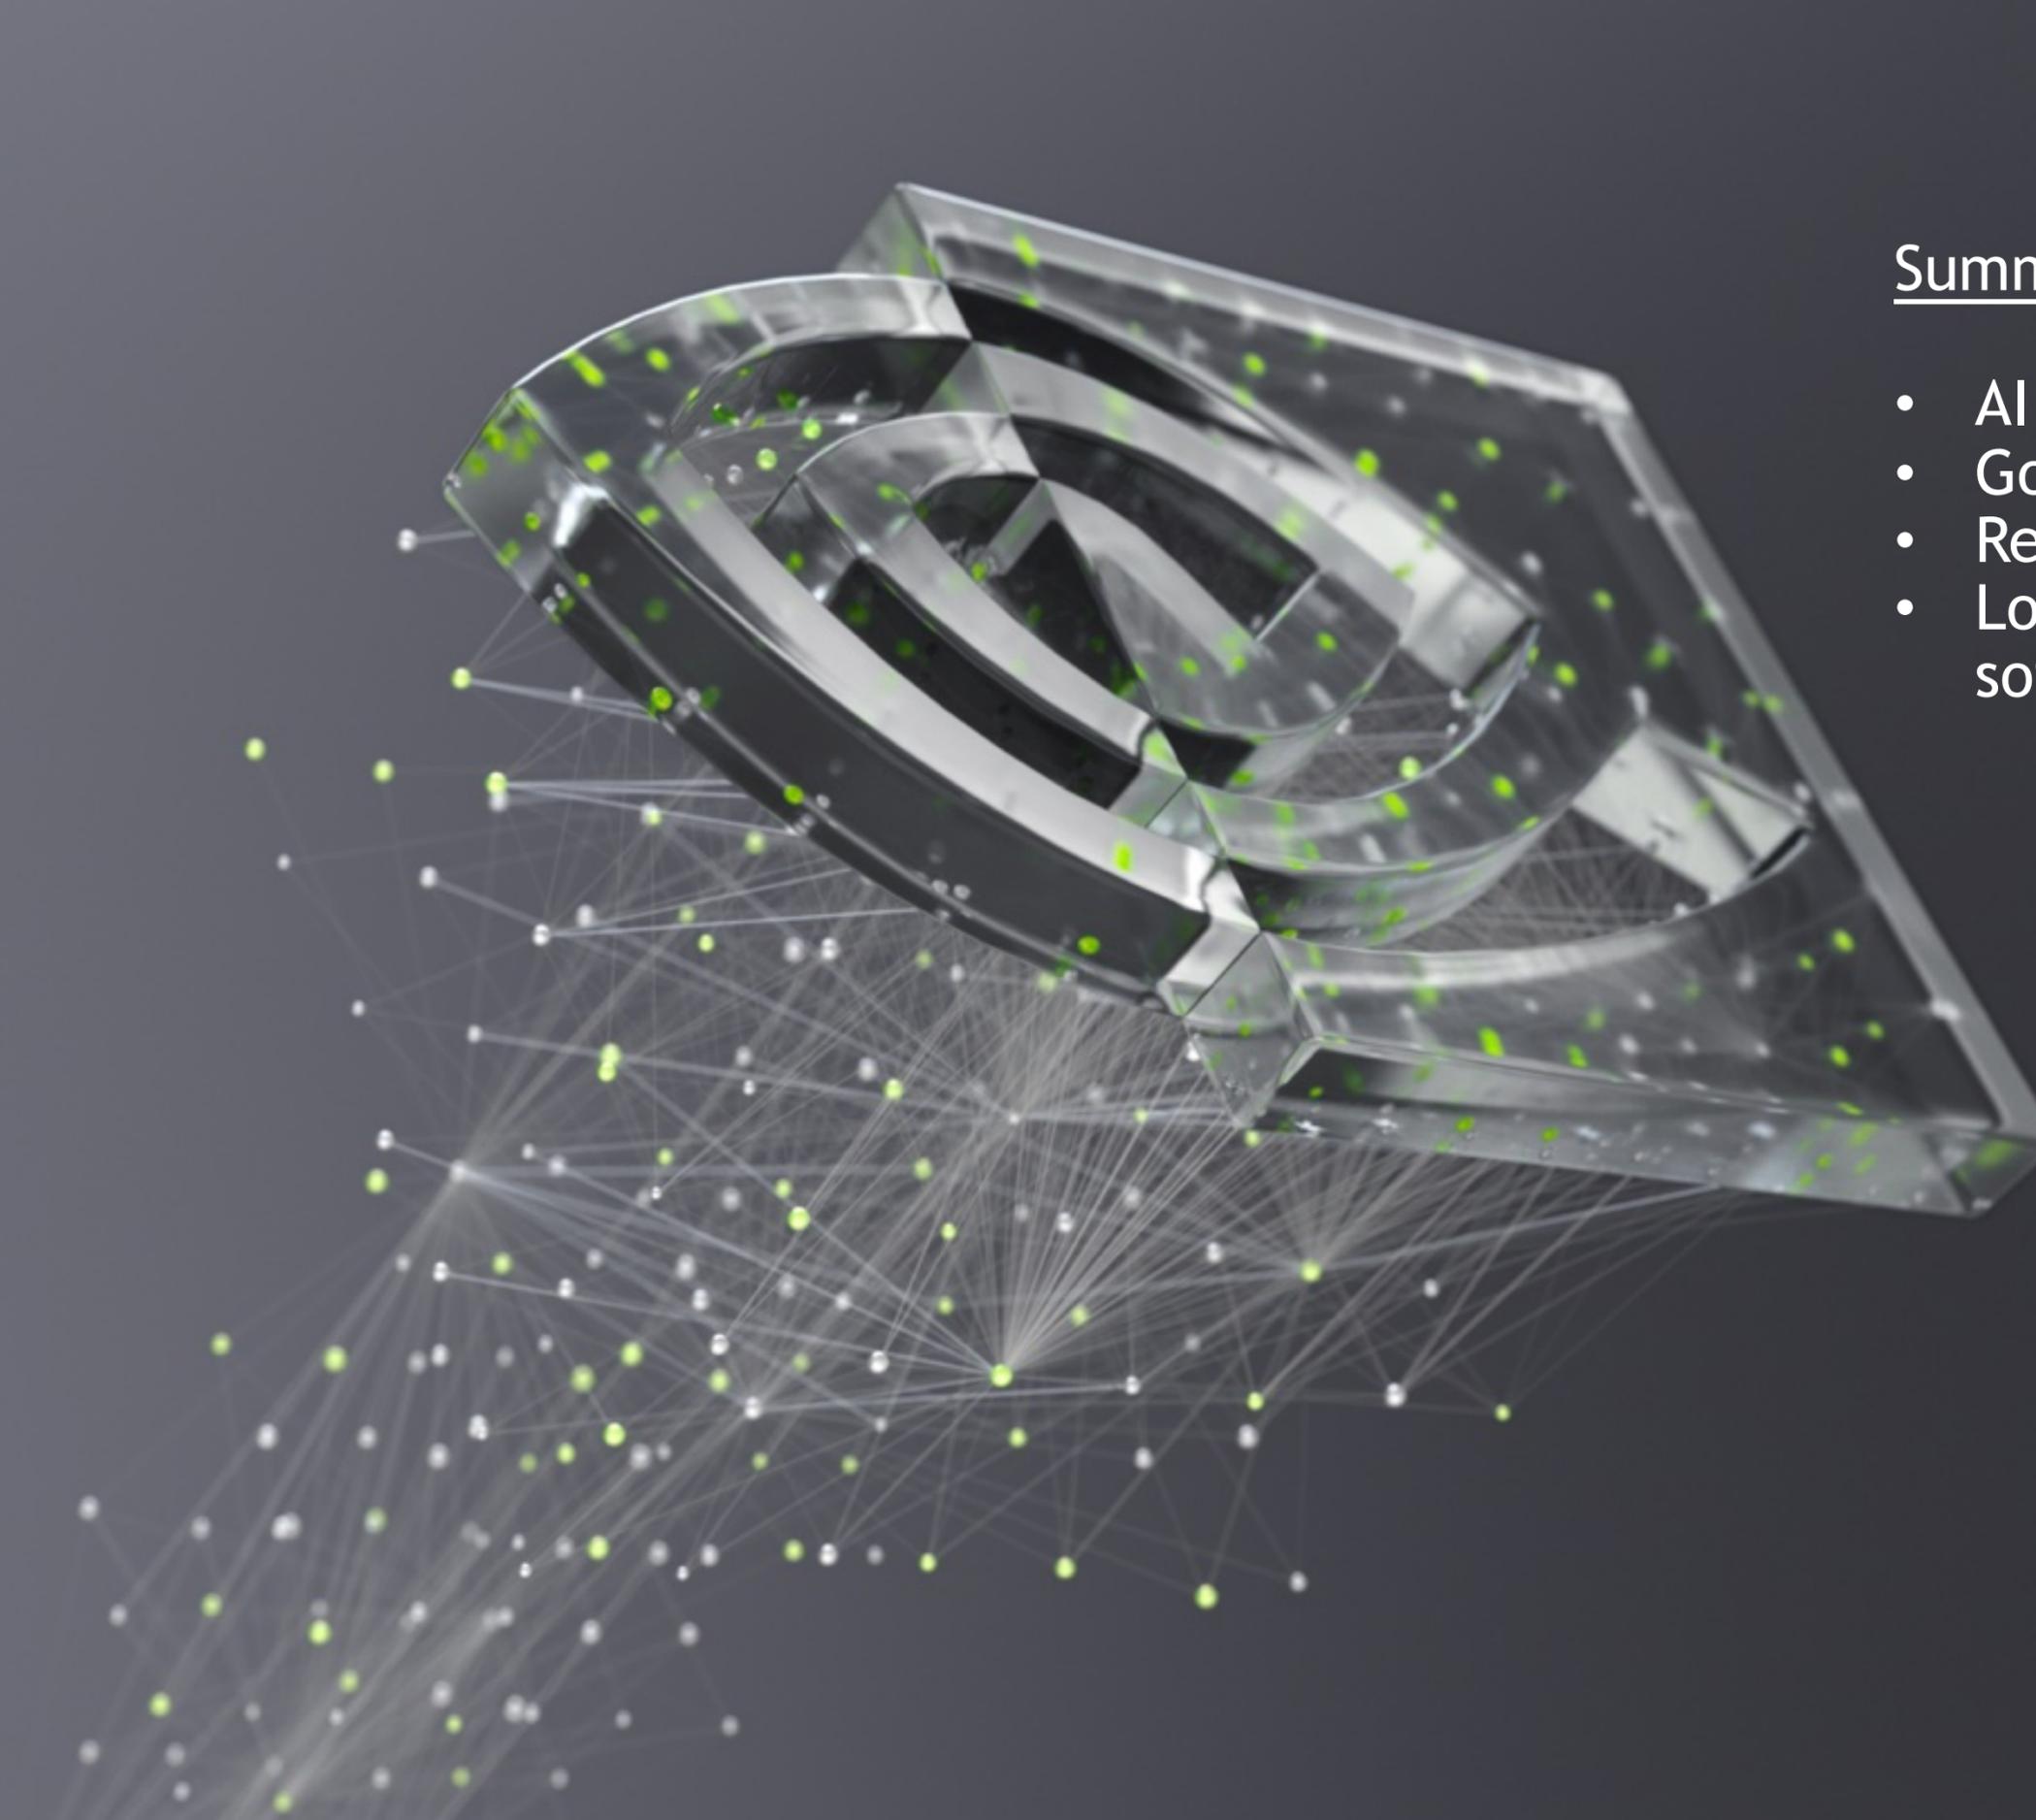
```
from pl_bolts.callbacks.self_supervised import SSLOnlineEvaluator
```

DataModules

```
from pl_bolts.datamodules import CIFAR10DataModule, FashionMNISTDataModule
```

Mix and match your data and lightning code

```
model = ImageGPT(datamodule=FashionMNISTDataModule(PATH))
```



## Summary

- AI for science is advancing rapidly!
- Good progress on scientific challenges
- Reviewed major trends in AI
- Looked at powerful new hardware and software tools you can use



dhall@**nvidia**.