Google DeepMind

Applying AI to complex simulation problems in the physical sciences

Alvaro Sanchez Gonzalez

18-Jan-2024

Simulation is fundamental to science and technology

Largest supercomputers in the world (Nov 2019)

#1. "*Summit*" @ Oak Ridge: "A Sneak Peek at 19 Science **Simulations** for the Summit Supercomputer in 2019"

1. *Evolution of the universe*
2. *Whole-cell simulation*
3. *Inside a nuclear reactor*
4. *Post-Moore's Law graphene circuits*
5. *Formation of matter*
6. *Cell's molecular machine*
7. *Unpacking the nucleus*
8. *Mars landing*
9. *Deep learning for microscopy*
10. *Elements from star explosions*
11. *Cancer data*
12. *Earthquake resilience for cities*
13. *Nature of elusive neutrinos*
14. *Extreme weather with deep learning*
15. *Flexible, lightweight solar cells*
16. *Virtual fusion reactor*
17. *Unpredictable material properties*
18. *Genetic clues in the opioid crisis*
19. *Turbulent environments*

Simulation is fundamental to science and technology

Largest supercomputers in the world (Nov 2019)

#1. "*Summit*" @ Oak Ridge: "A Sneak Peek at 19 Science **Simulations** for the Summit Supercomputer in 2019"

#2 "*Sierra*" @ Lawrence Livermore: "[nuclear] **simulation** in lieu of underground testing"

#3 "*Sunway TaihuLight*" @ NSC, Wuxi: "**simulated** the Universe with 10 trillion digital particles"

#4 "*Tianhe-2A*" @ NSC, Guangzhou: "main application … is for computational fluid dynamics (CFD) … aircraft **simulations**"

#5 "*Frontera*" @ TACC: "high-resolution climate **simulations**, molecular dynamics models with millions of atoms"

#6 "*Piz Daint*" @ CSCS: "**simulate** processes for projects in geophysics, materials science, chemistry, … climate modeling"

#7 "*Trinity*" @ Los Alamos: "A trillion-particle **simulation**? No sweat for the Trinity supercomputer at Los Alamos"

(#8 "*ABCI*" @ AIST, Japan: not simulation, but deep learning)

#9 "*SuperMUC-NG*" @ Leibniz Supercomputing Centre: "Researchers Visualize the Largest Turbulence **Simulation** Ever"

#10 "*Lassen*" @ Lawrence Livermore: "The system is designated for unclassified **simulation** and analysis"

Why *learn* simulation?

Engineered simulators:

1. Substantial effort to build
2. Substantial resources to run
3. Only as accurate as the designer
4. Not always suitable for solving inverse problems

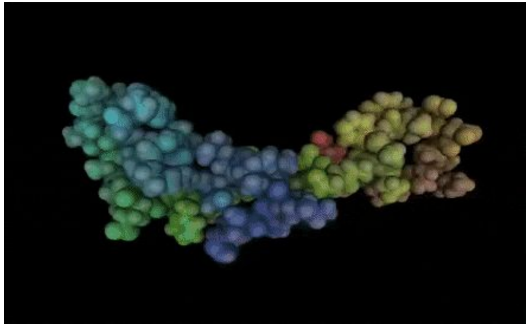

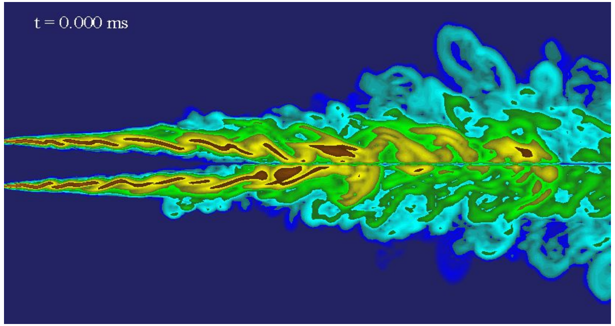

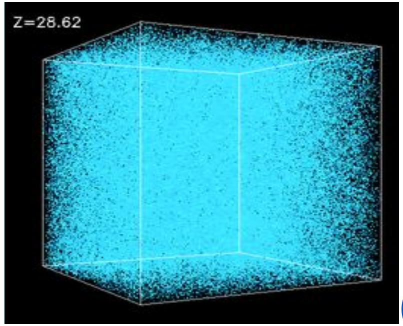

Why *learn* simulation?

Engineered simulators:

1. Substantial effort to build
2. Substantial resources to run
3. Only as accurate as the designer
4. Not always suitable for solving inverse problems

Learned simulators:

1. Shared architectures
2. Accuracy-efficiency trade off
3. As accurate as the available data
4. Gradient-based planning
5. Interpretable models!*

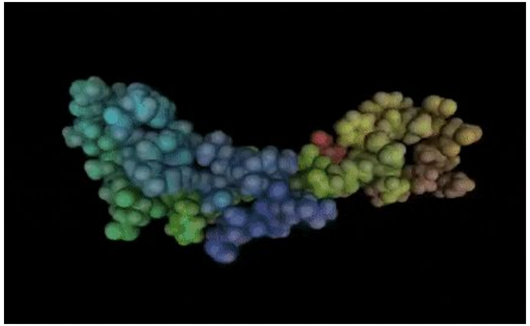

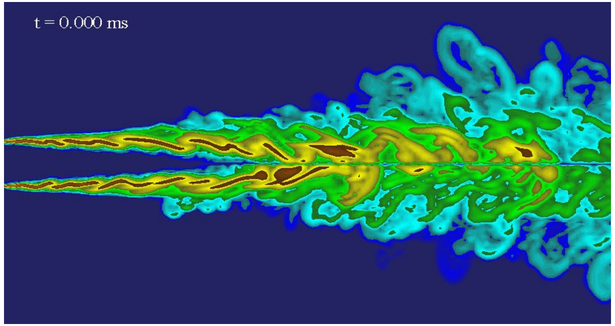

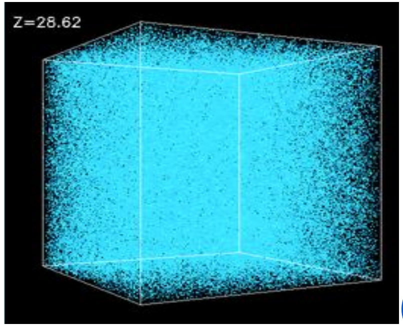

*"Discovering Symbolic Models from Deep Learning with Inductive Biases"
Cranmer et al., NeurIPS 2020

Content

Learning AI simulation for:

1. Graphics and engineering
2. Scientific turbulence
3. Global weather

Big credit to all co-authors!

Learning simulation for graphics and engineering

"Learning to Simulate Complex Physics with Graph Networks" (ICML 2020) link
Alvaro Sanchez-Gonzalez, Jonathan Godwin*, Tobias Pfaff*, et al.*
Video page: sites.google.com/view/learning-to-simulate

"Learning Mesh-Based Simulation with Graph Networks" (ICLR 2021) link
Tobias Pfaff, Meire Fortunato*, Alvaro Sanchez-Gonzalez*, Peter Battaglia*
Video page: sites.google.com/view/meshgraphnets

"Inverse Design for Fluid-Structure Interactions using Graph Network Simulators" (NeurIPS 2022) link
Kelsey Allen, Tatiana Lopez Guevara*, Kimberly Stachenfeld*, et al.*
Video page: sites.google.com/corp/view/optimizing-designs

- Focus on **general principles behind the design**
  - Applicable to other domains

"Learning to Simulate Complex Physics with Graph Networks" (ICML 2020)
Video page: sites.google.com/view/learning-to-simulate

Water simulation (SPH)

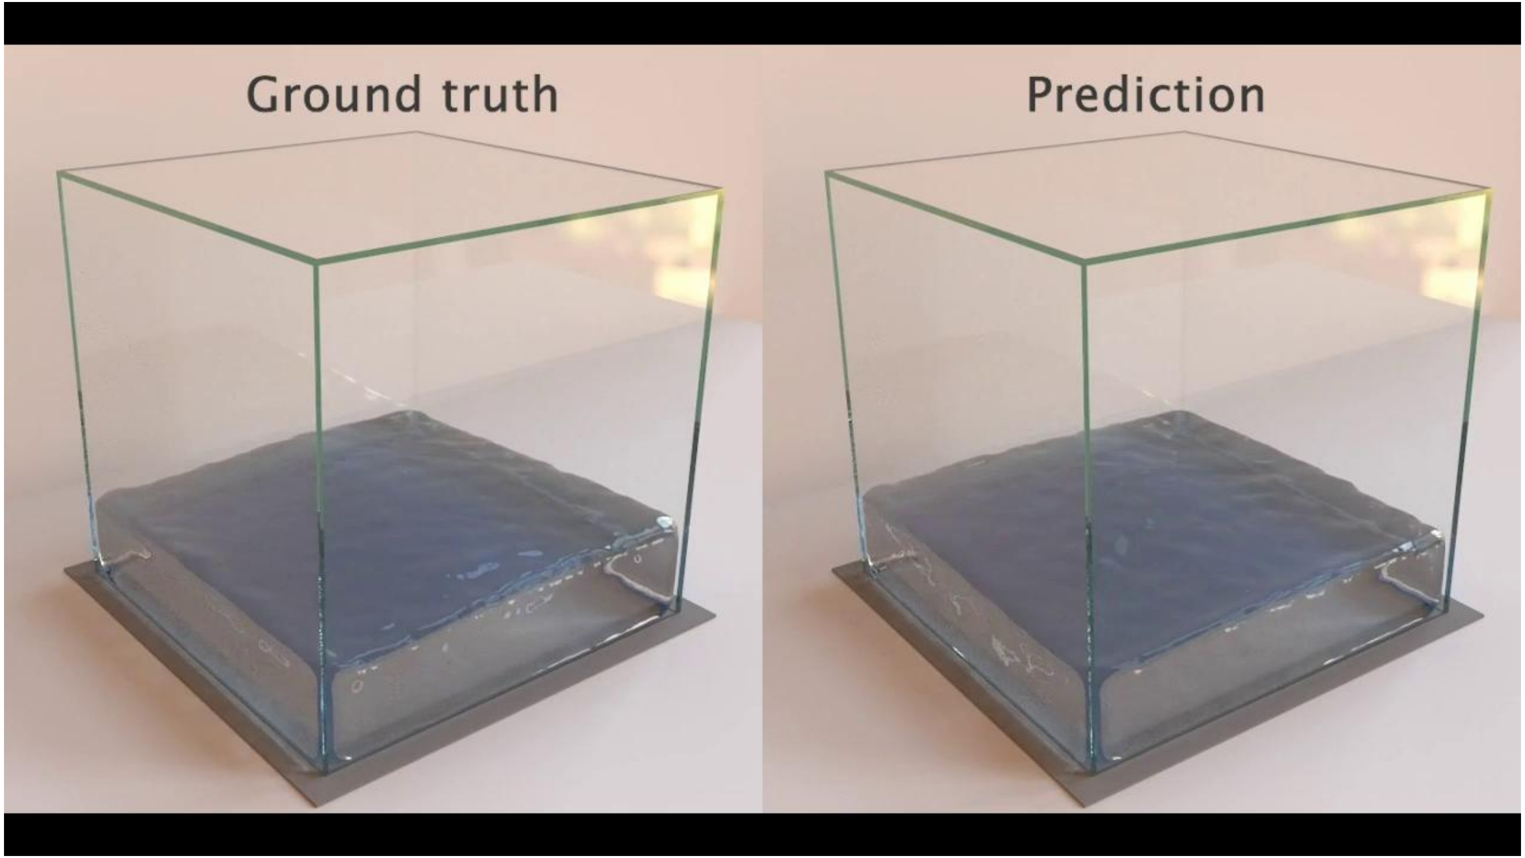

“Learning to Simulate Complex Physics with Graph Networks” (ICML 2020)
Video page: sites.google.com/view/learning-to-simulate

Water simulation (SPH)

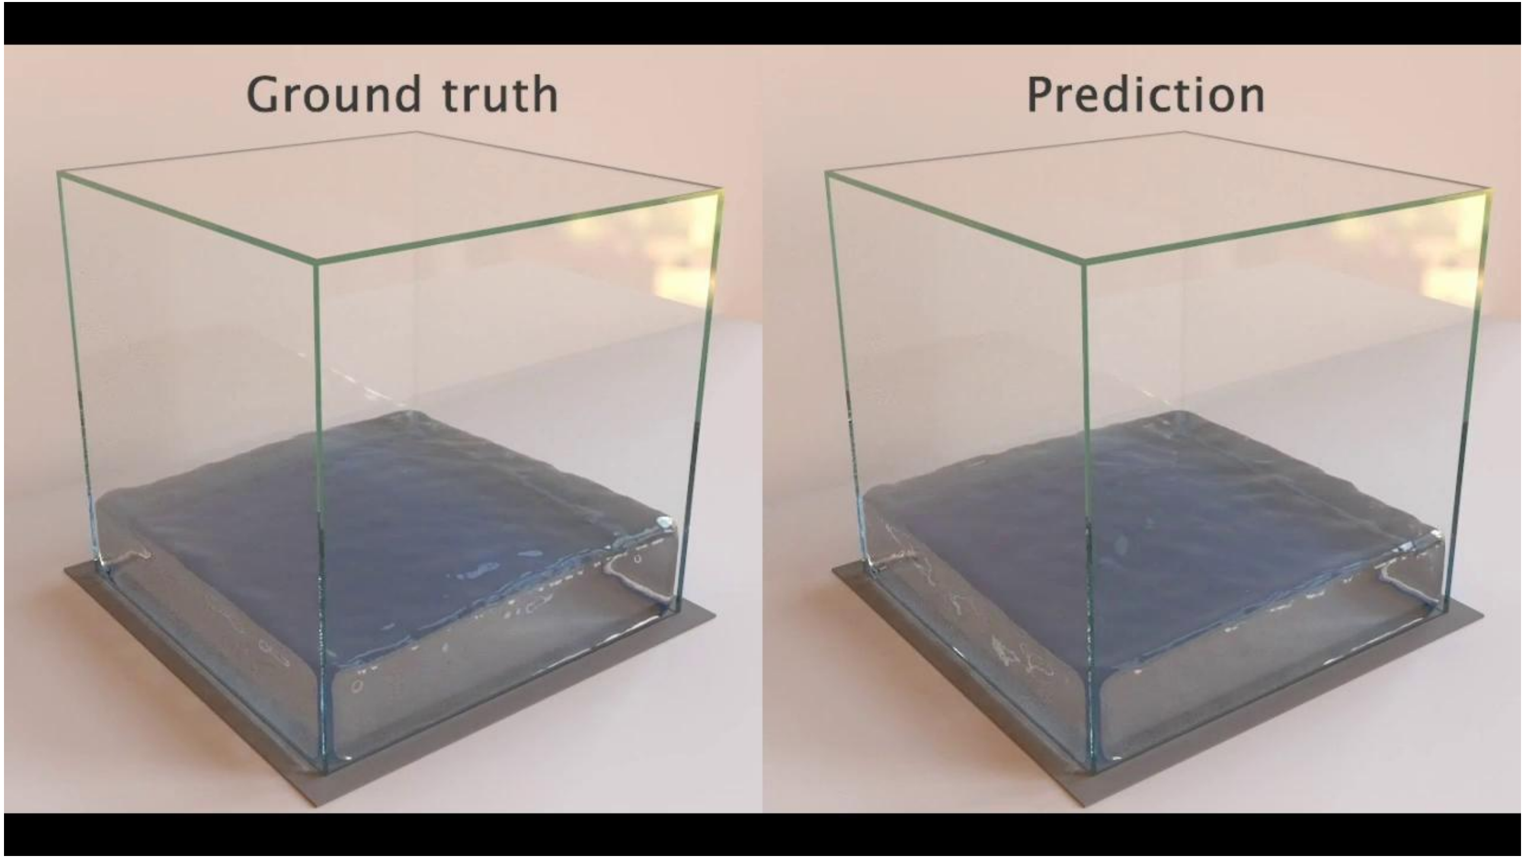

Multiple materials

"Learning to Simulate Complex Physics with Graph Networks" (ICML 2020)
Video page: sites.google.com/view/learning-to-simulate

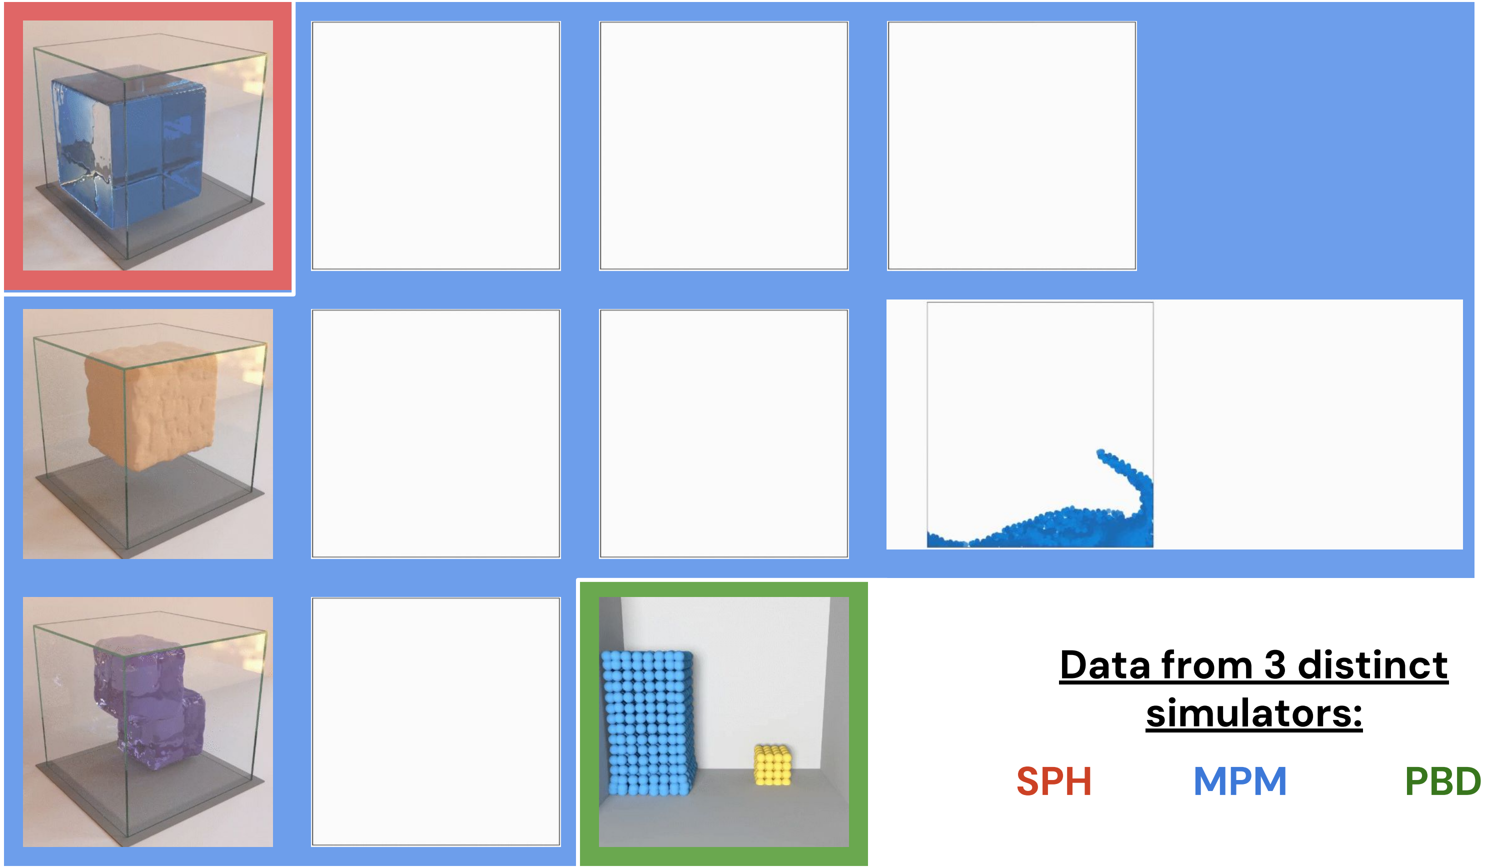

Data from 3 distinct simulators:

SPH MPM PBD

“Learning Mesh-Based Simulation with Graph Networks” (arXiv, under review)
Video page: sites.google.com/view/meshgraphnets

Cloth simulation (ArcSim)

- Triangular dynamic mesh
- Lagrangian representation

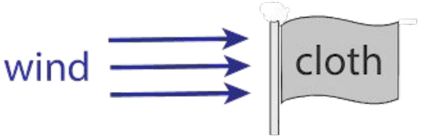

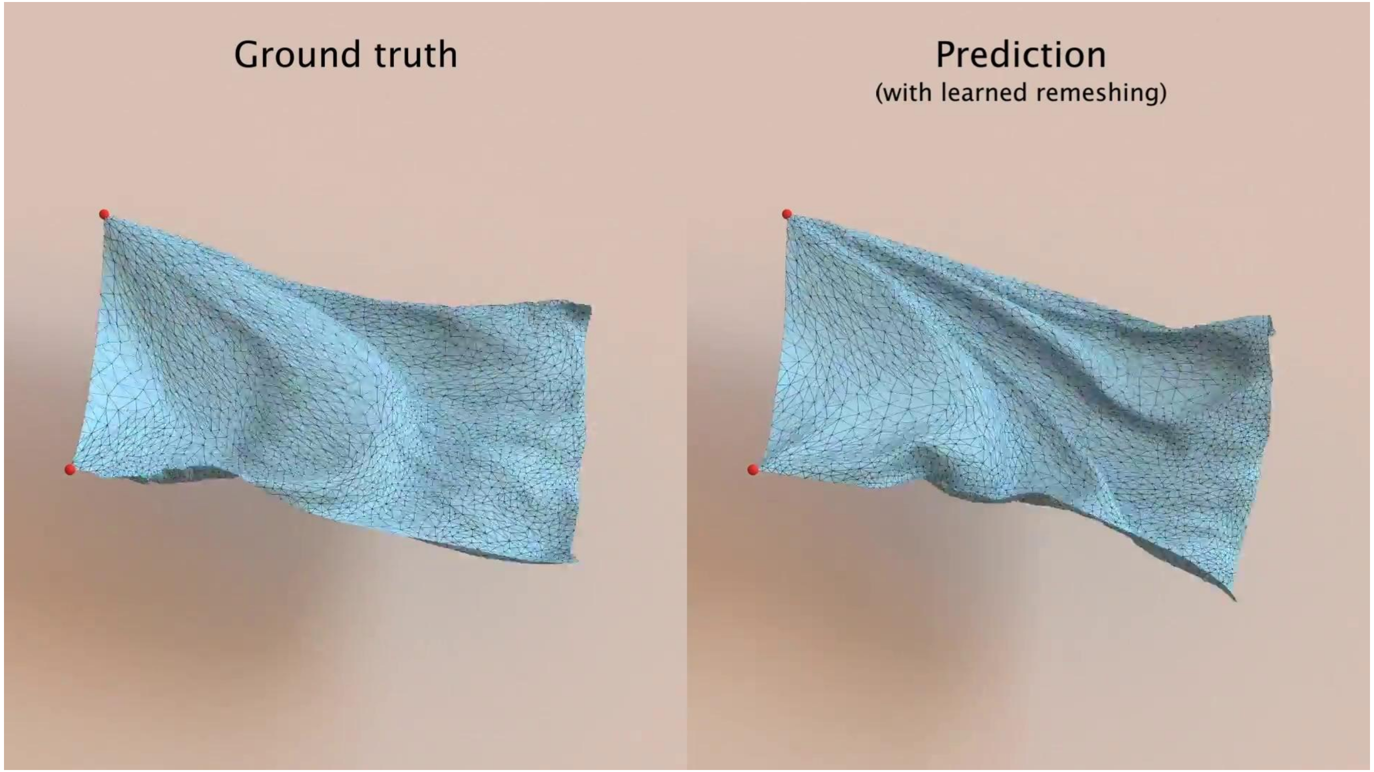

“Learning Mesh-Based Simulation with Graph Networks” (arXiv, under review)
Video page: sites.google.com/view/meshgraphnets

Structural dynamics (COMSOL)

- Tetrahedral mesh
- Lagrangian representation
- Quasi-static simulation

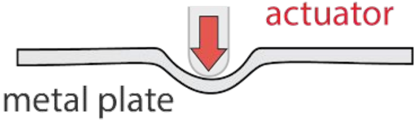

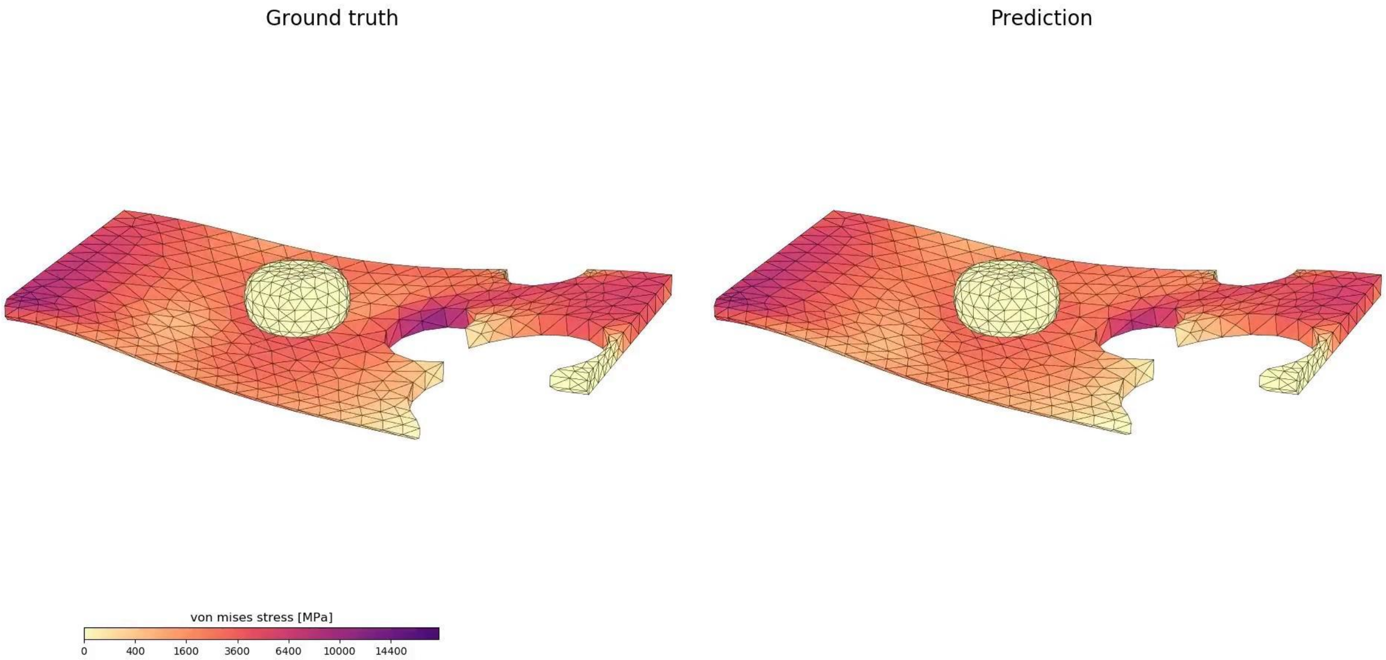

“Learning Mesh-Based Simulation with Graph Networks” (arXiv, under review)
Video page: sites.google.com/view/meshgraphnets

Incompressible fluids (COMSOL)

- Navier-Stokes
- Eulerian representation

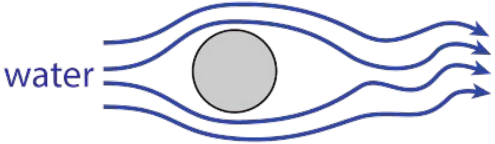

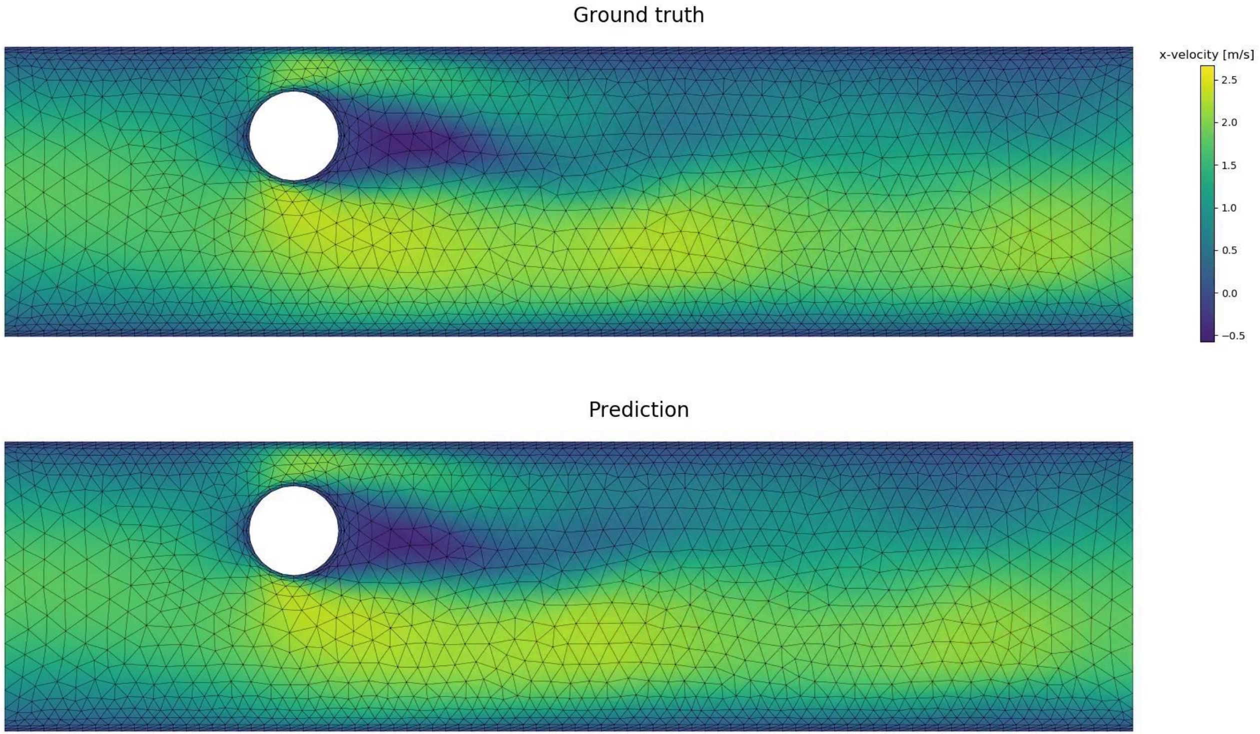

Aerodynamics (SU2)

- Navier-Stokes
- Eulerian representation

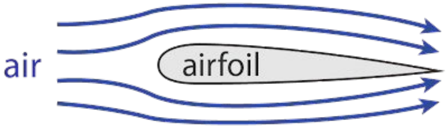

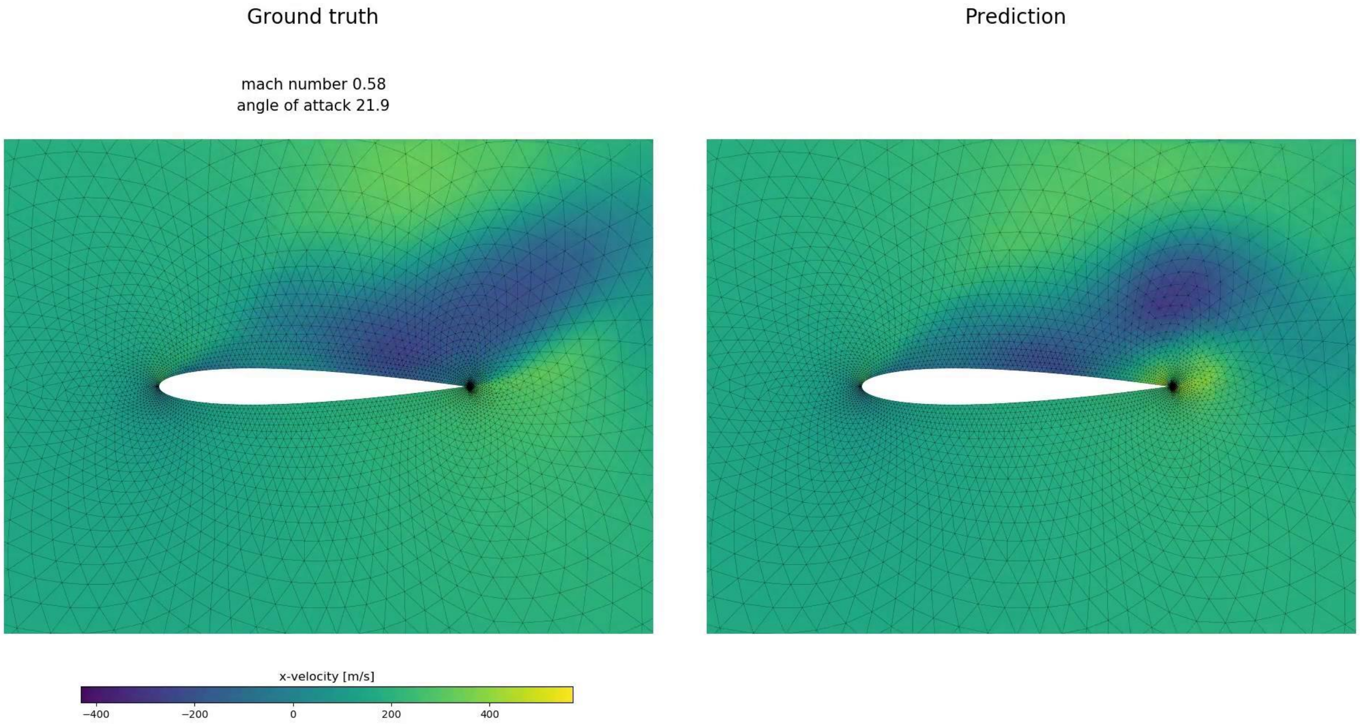

Why Graph Network based simulators?

- **Adaptability**
  - Same model → Vastly different materials and domains
- **Data efficiency**
  - < 1000 training trajectories
- **Performance**
  - Latest model: ~10 to 100 times faster than ground truth simulator
- **Generalization**

“Learning to Simulate Complex Physics with Graph Networks” (ICML 2020)
Video page: sites.google.com/view/learning-to-simulate

Graph Network-Based Simulators

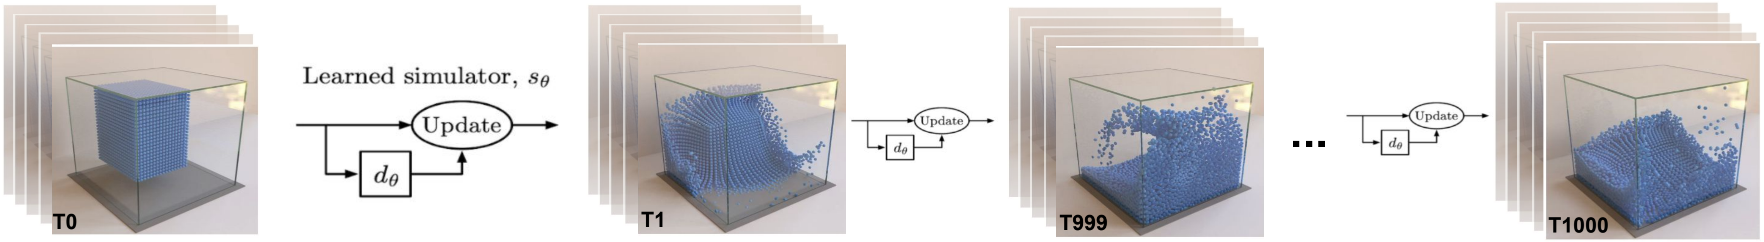

Input

Prediction 1

Prediction 2

Prediction n

Design principle: **neural networks trained on small datasets are dumb, let’s give them some extra info!**

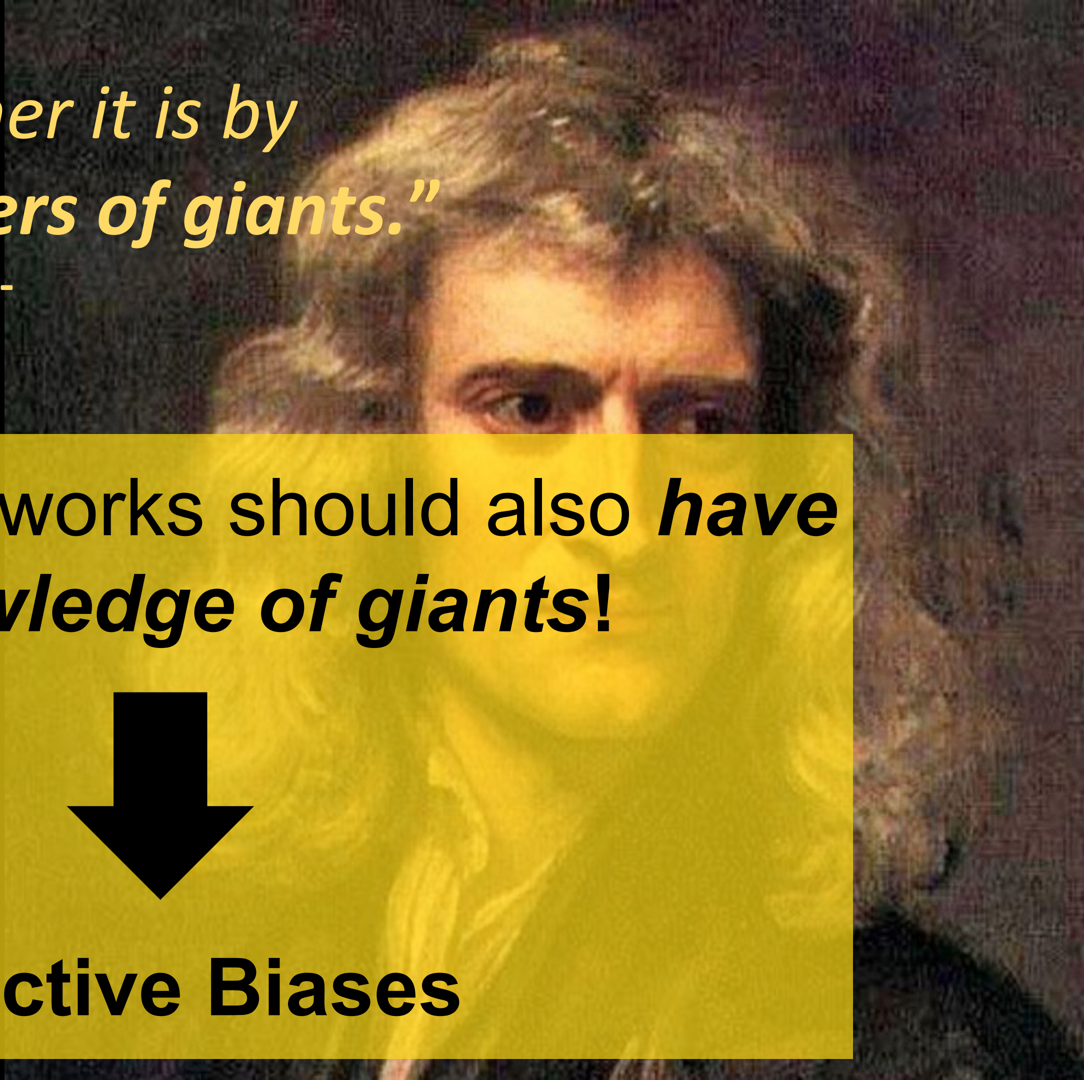

“If I have seen further it is by ***standing on the shoulders of giants.”***

-Sir Isaac Newton-

Our Neural Networks should also ***have the knowledge of giants*****!**

Inductive Biases

Physics-inspired inductive biases

"An **inductive bias** allows a learning algorithm to prioritize one solution (or interpretation) over another."

Mitchell, T. M.. *The need for biases in learning generalizations. (1980)*

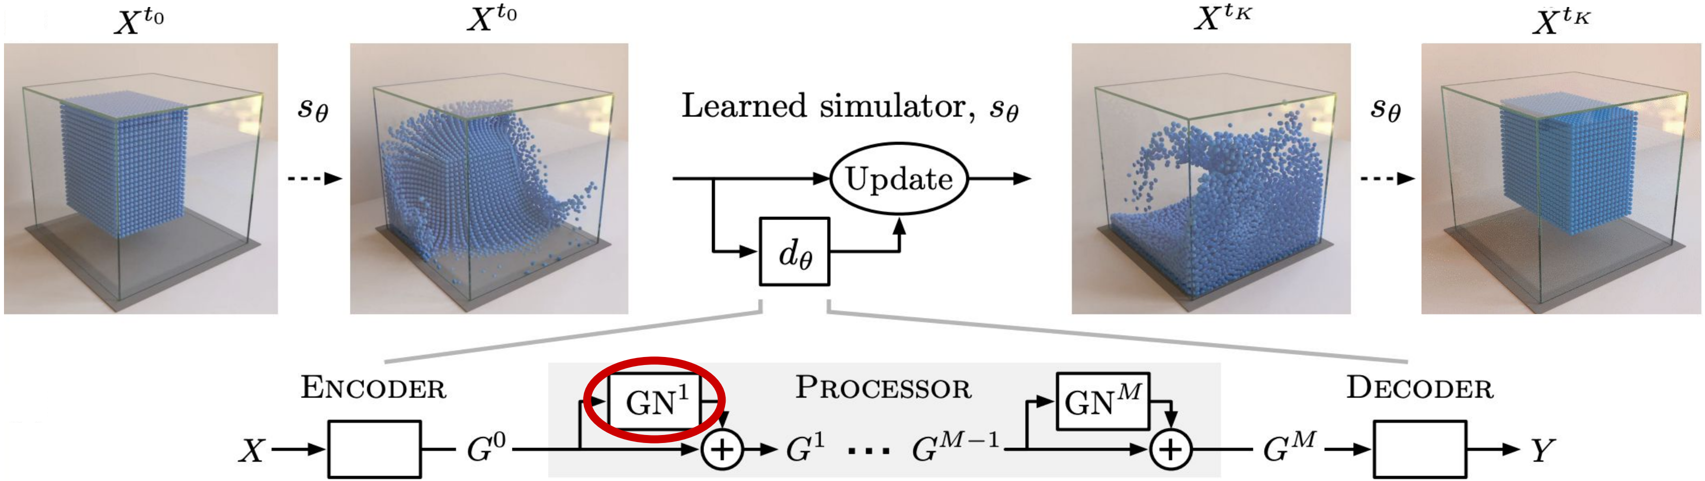

Local interactions

Spatial equivariance

Space homogeneity

Pairwise interactions

Superposition principle

Differential equations

Permutation equivariance

Interaction Networks (Battaglia et al., 2016, NeurIPS)

Graph Networks (Battaglia et al., 2016, arXiV)

Graph Networks

A neural network that can predict properties of a graph

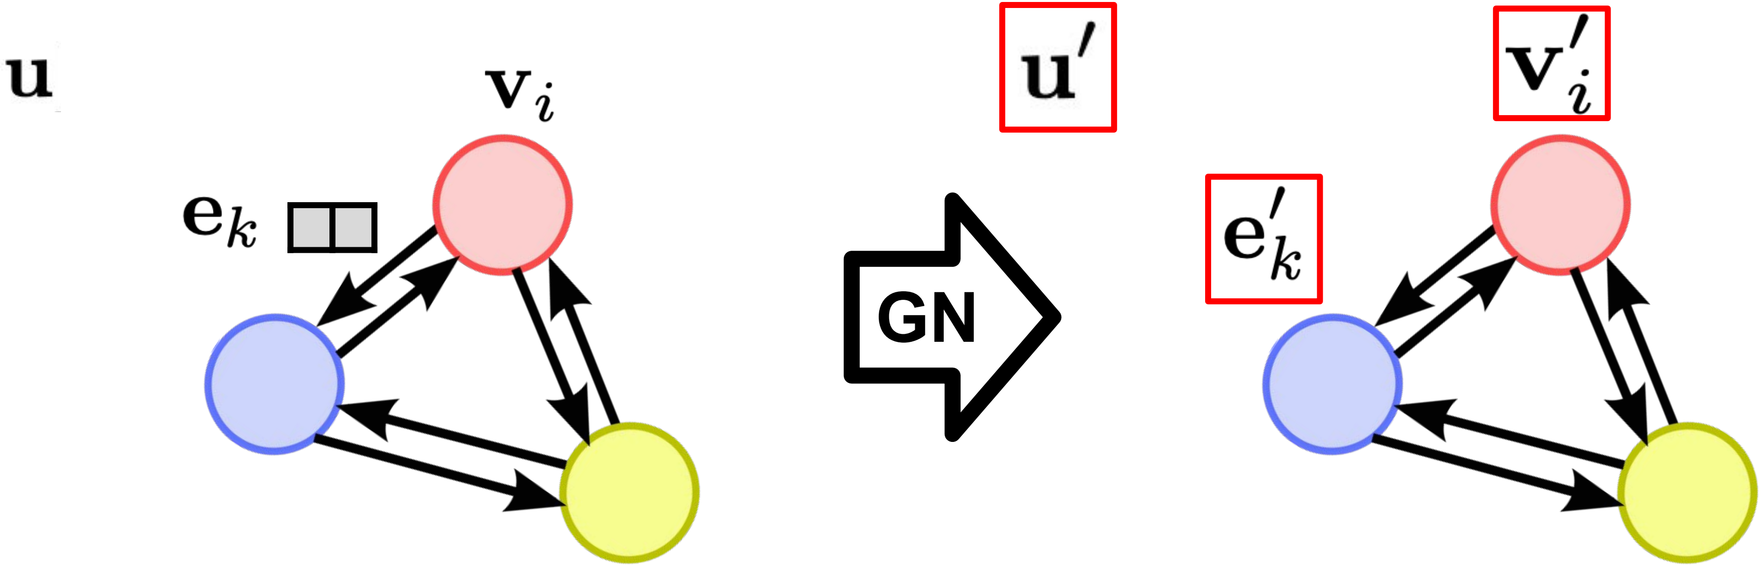

Encoder

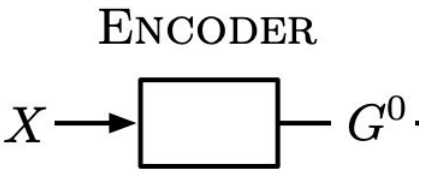

- Transform the inputs into a graph
  - Add connectivity with a certain radius *R*

Local (spatial) interactions

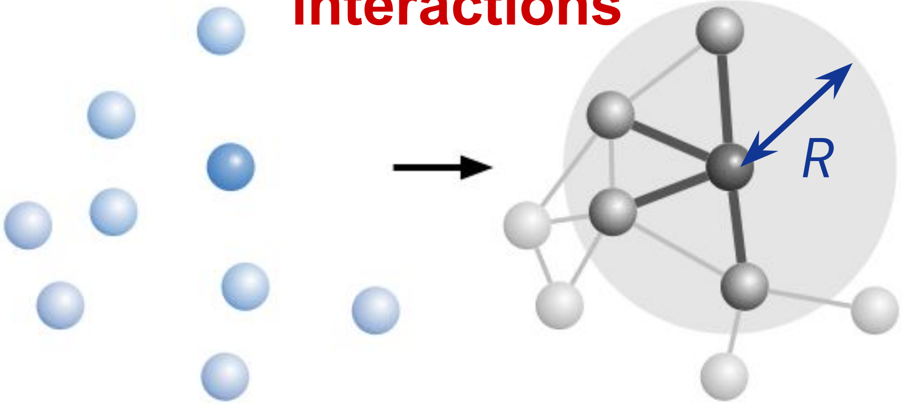

Processor: Deep Graph Network

- Increase range of communication

Local interactions

After 1 MP steps | After 2 MP steps | After 3 MP steps | After *n* MP steps

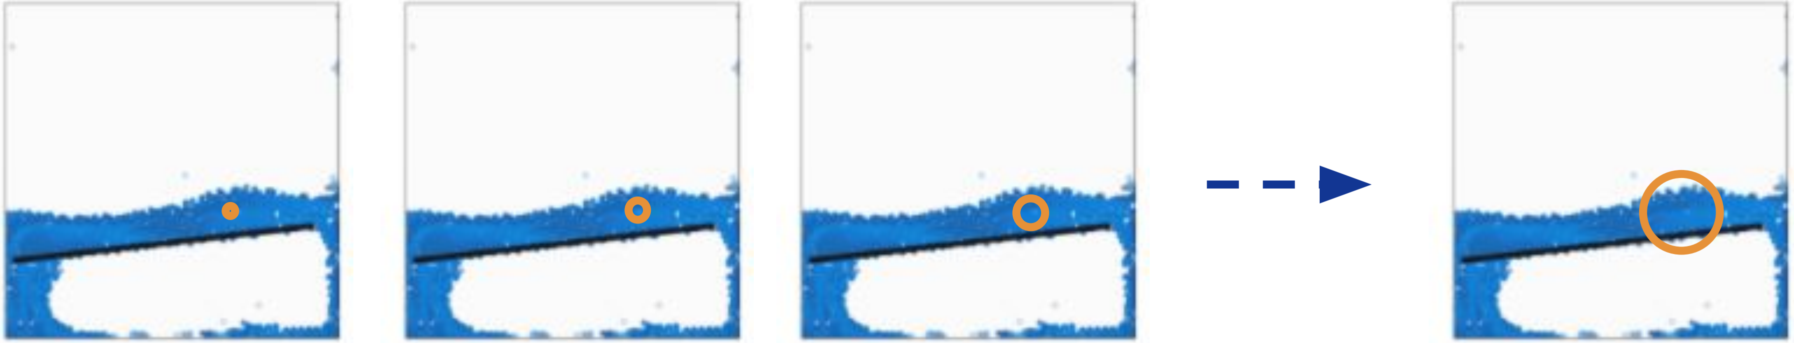

Decoder and update

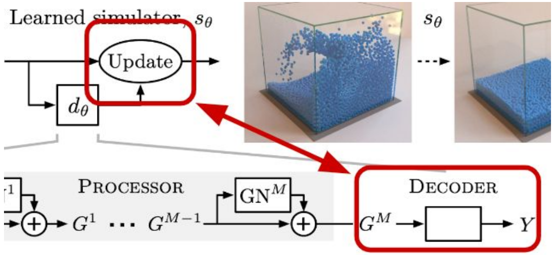

Differential equations prior: add a cheap integrator!

Lagrangian Newtonian

Particle-based fluids

Cloth simulation

$$\frac{d\,\dot{\mathbf{x}}}{dt} = f(\mathbf{x}, \dot{\mathbf{x}})$$

Euler integrator:

$x^{t+1} = x^t + (x^t - x^{t-1}) + \mathbf{NN}(x^t, v^t)$

Quasistatic

Structural mechanics

$$\frac{d\,\mathbf{x}}{dt} = f(\mathbf{x})$$

Eulerian Navier Stokes

Aerodynamics

$$\frac{d\,\dot{\mathbf{v}}}{dt} = f(\mathbf{v}, \rho)$$

$$\frac{d\,\dot{\rho}}{dt} = f(\mathbf{v}, \rho)$$

Incompressible fluid

$$\frac{d\,\dot{\mathbf{v}}}{dt} = f(\mathbf{v})$$

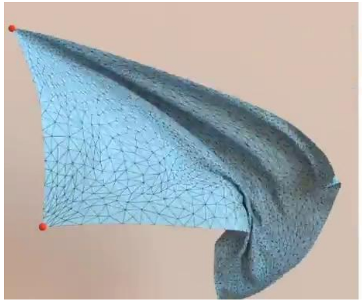

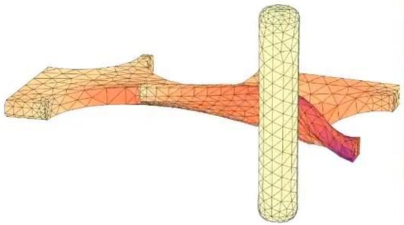

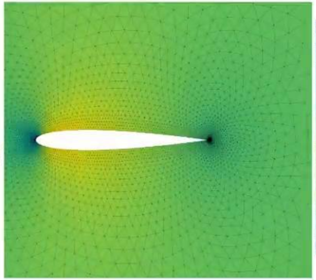

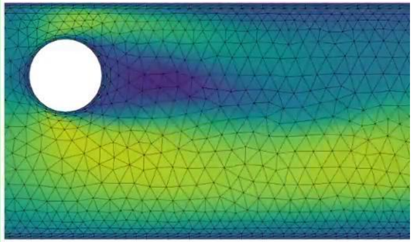

“Learning to Simulate Complex Physics with Graph Networks” (ICML 2020)
Video page: sites.google.com/view/learning-to-simulate

Generalization to many more particles

Training

1 x 1 domain
2k particles
600 steps

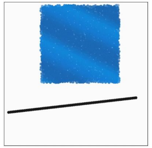

“Learning to Simulate Complex Physics with Graph Networks” (ICML 2020)

Video page: sites.google.com/view/learning-to-simulate

Generalization to many more particles

Training

1 x 1 domain
2k particles
600 steps

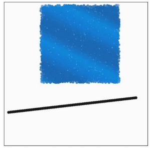

Generalization

2 x 2 domain
28k particles
2500 steps

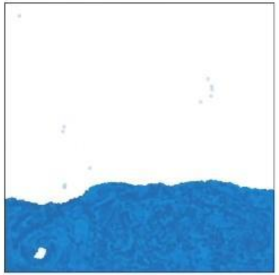

Generalization

8 x 4 domain
85k particles
5000 steps

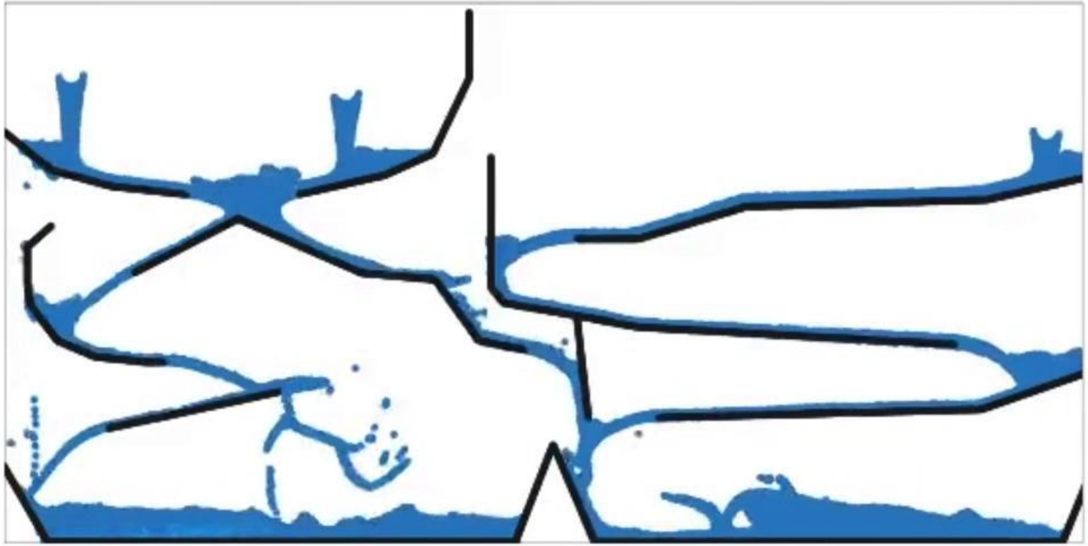

Generalization to different meshes

Training

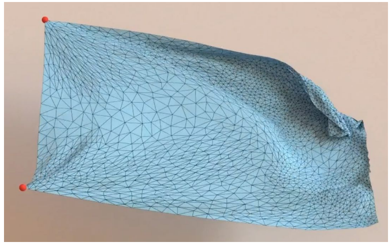

Generalization

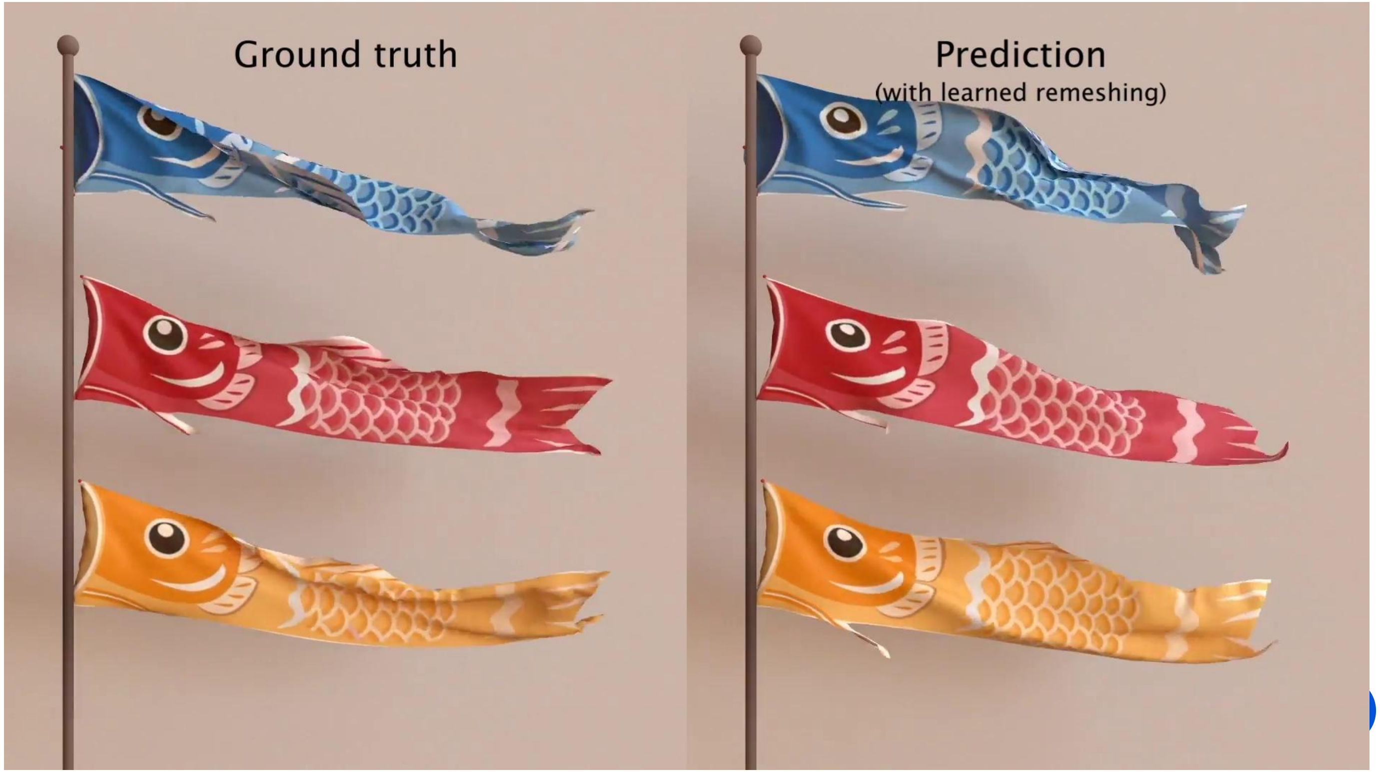

“Learning Mesh-Based Simulation with Graph Networks” (arXiv, under review)
Video page: sites.google.com/view/meshgraphnets

Generalization to larger meshes

Training:
2k nodes

Generalization:
>20k nodes

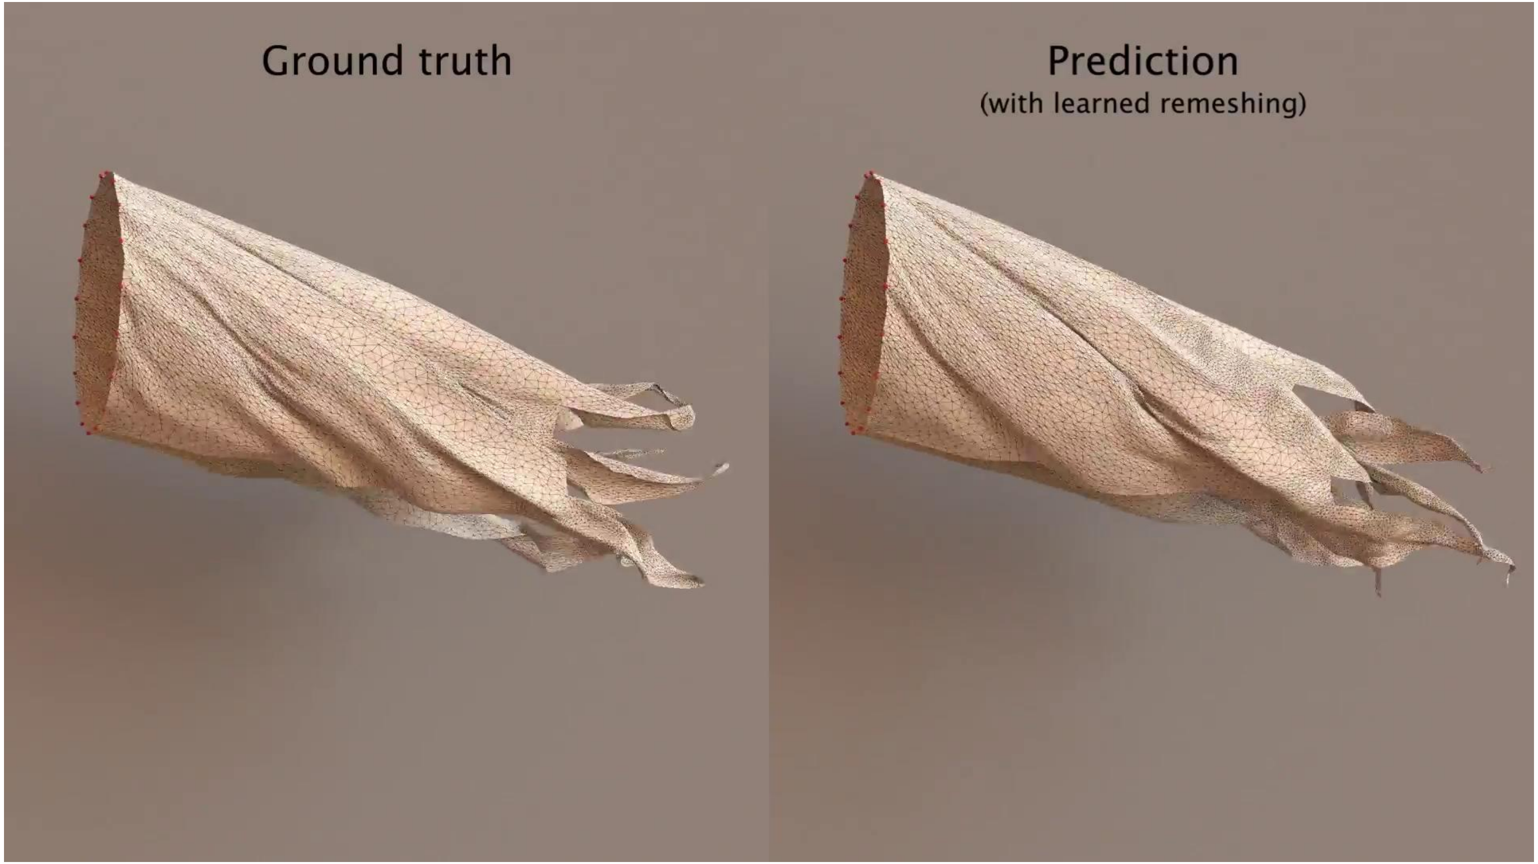

Inverse design - using a forward model

“Move ramps to to get liquid to yellow region”

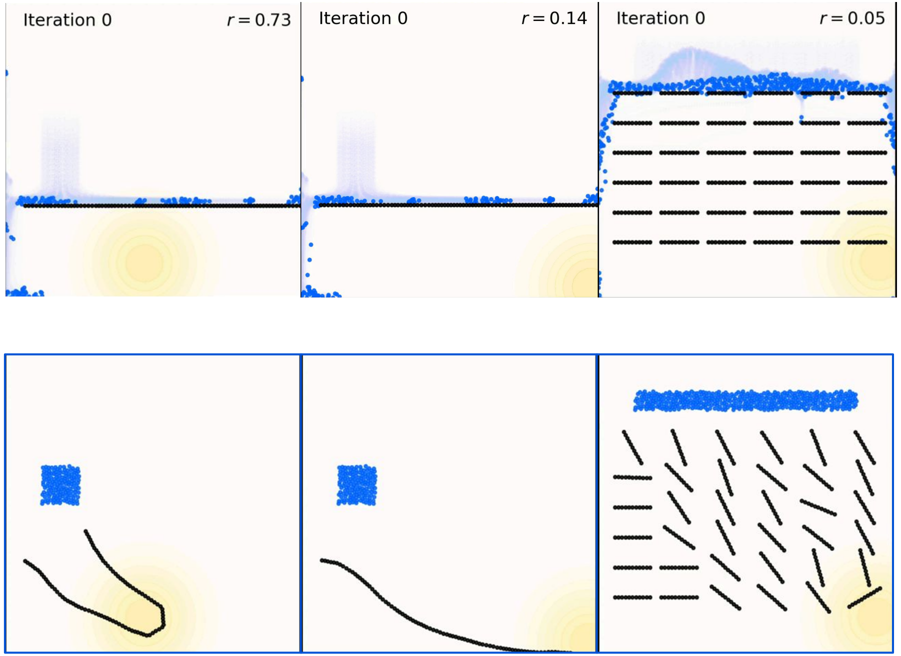

“Move mesh to get particles to purble points”

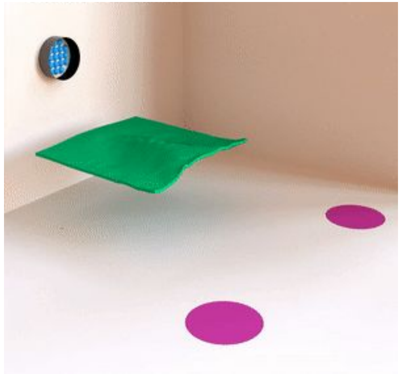

“Maximize lift and minimize drag of the wing”

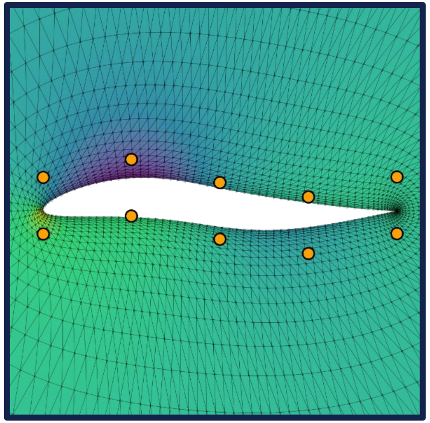

Bonus slide: interpretable graph networks

- "Discovering Symbolic Models from Deep Learning with Inductive Biases"
  Cranmer et al., NeurIPS 2020
- Extract symbolic models from edge and node functions of a GraphNet

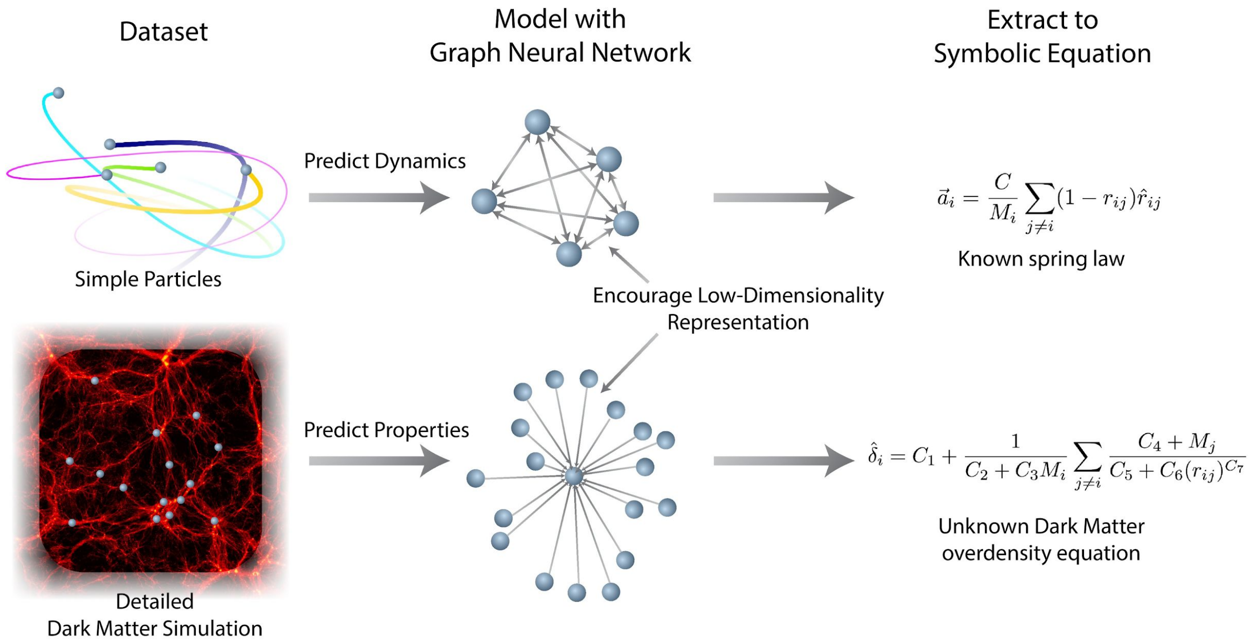

“Learned Coarse Models for Efficient Turbulence Simulation” (ICLR 2022)

*Stachenfeld at al.**

Arxiv: arxiv.org/abs/2112.15275

Video page: sites.google.com/corp/view/learned-turbulence-simulators

Engineering

Forecasting

Science

Classical numerical solvers are powerful but computationally expensive

Can fully-learned simulators capture complex, chaotic turbulence accurately at low-resolutions?

Model: Architecture

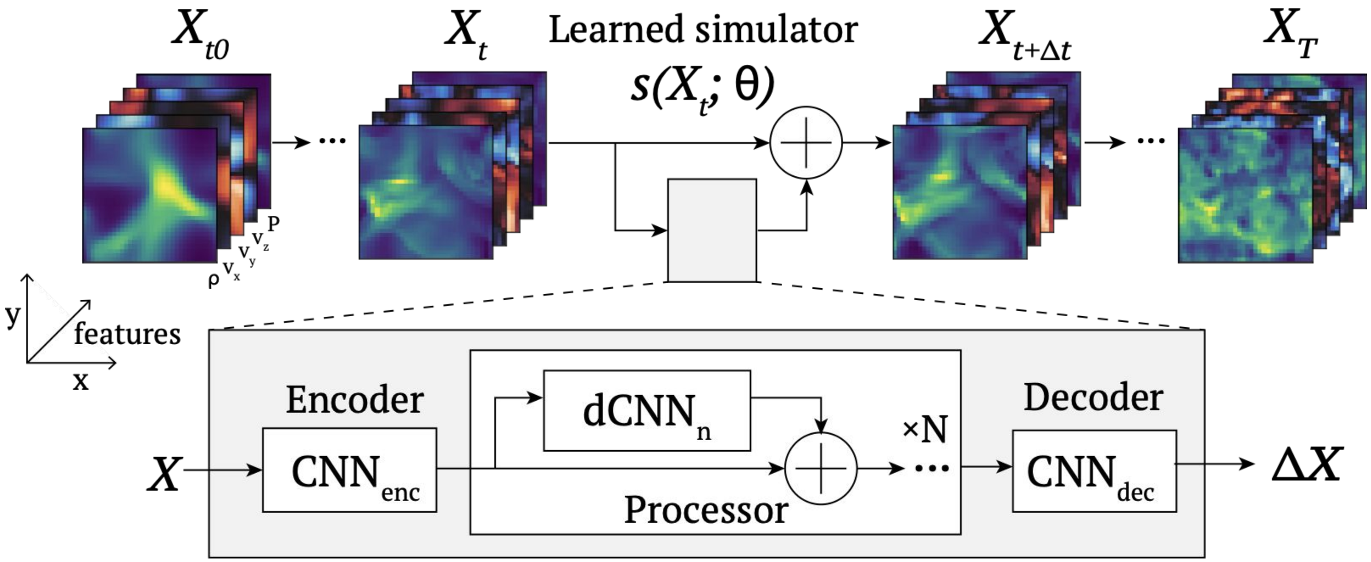

One model → 4 different domains

Domain Generality

1D Kuramoto-Sivashinsky (KS) Equation

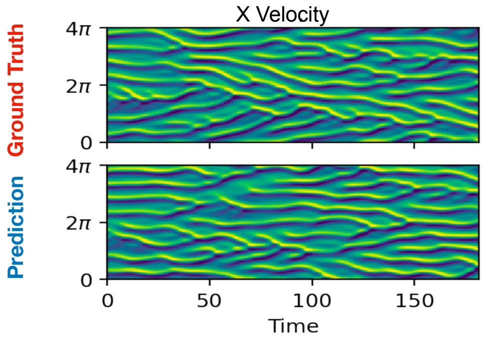

3D Uniform Compressible Decaying Turbulence

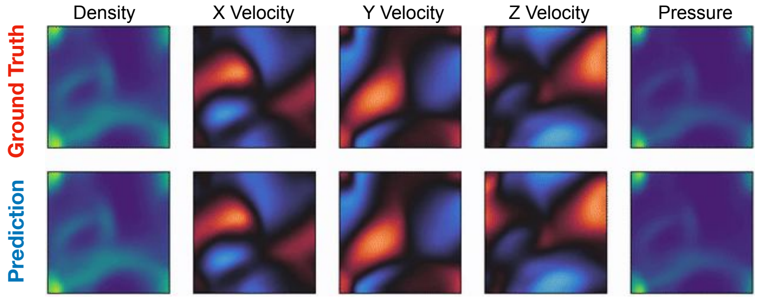

2D Incompressible Turbulence

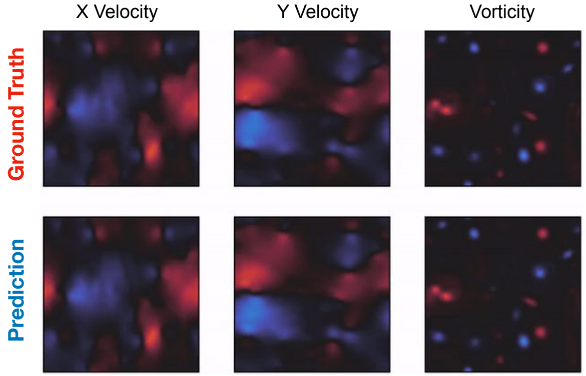

3D Mixing Layer Turbulence with Radiative Cooling

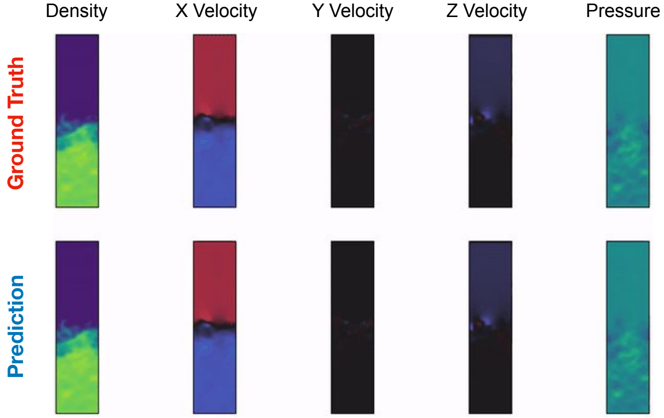

Spatial Coarsening

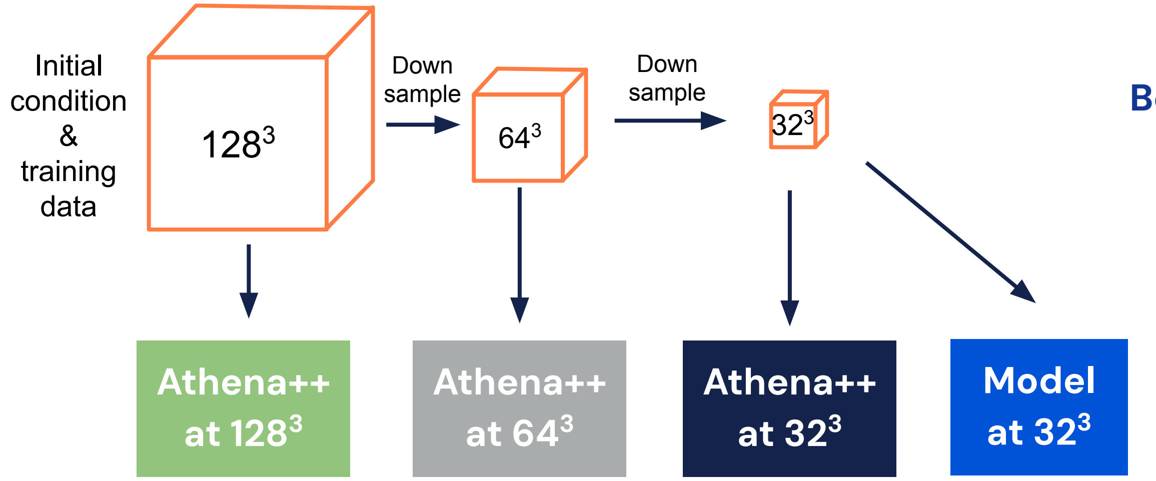

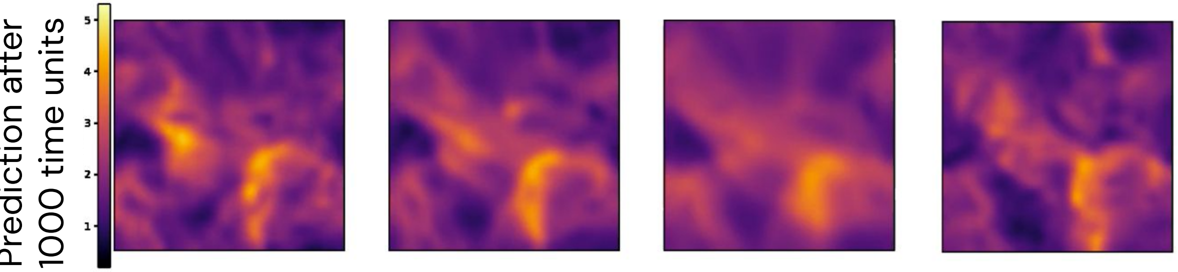

Better spectrum than Athena at 64^3

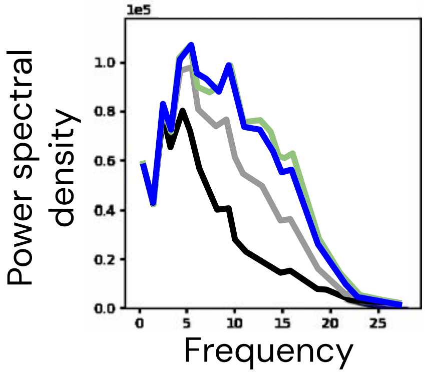

Temporal Coarsening

Model can be trained to work on a large range of timesteps

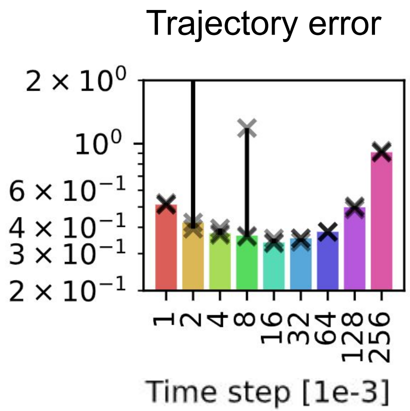

Ground Truth **Learned model timestep**

1 2 4 8 16 32 64 128 256

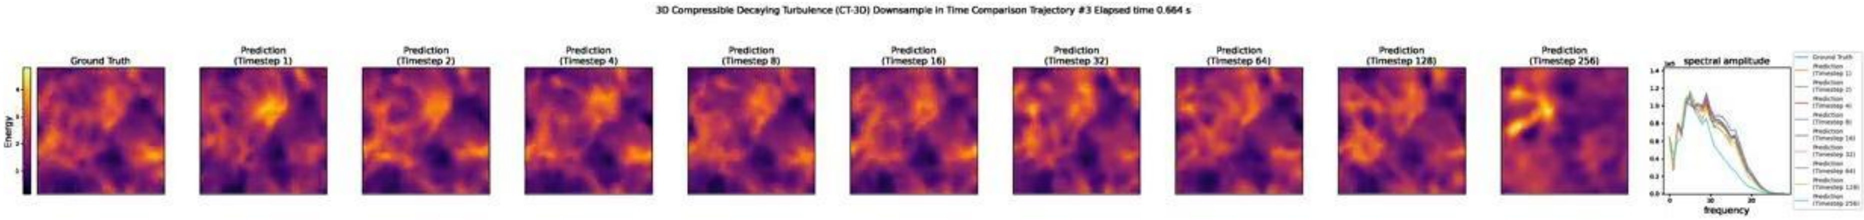

For comparison the timestep of Athena++ at 128^3 resolution is ~0.125 → **Timestep increase: 8x – 2000x**

“Learning skillful medium-range global weather forecasting” (Science 2023)

Lam at al.,

www.science.org/doi/10.1126/science.adi2336

Goal: Transform global operational weather forecasting with ML

- More **accurate** and more **efficient** forecasts made by models learned **directly from data**.
- **Improve downstream applications** that rely on weather forecasts for the benefit of humanity.

Scientific objective: Predict **whole-Earth’s** surface & atmospheric weather, **10 days ahead**, at **30 km resolution**

Surface E-W wind

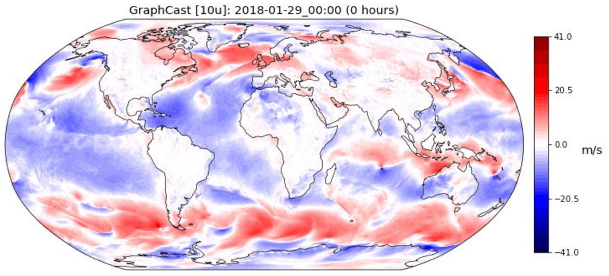

Surface temperature

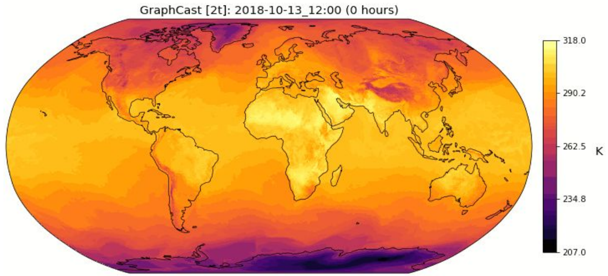

Temperature @ 500 hPa

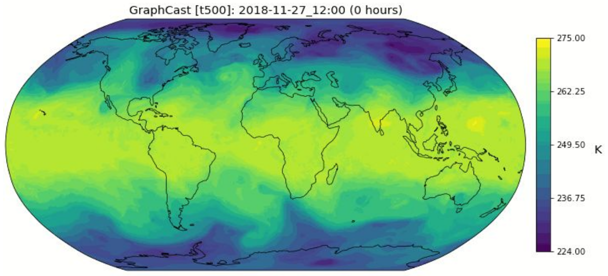

Difficulty and scale: A single atmospheric state is ~1GB worth of data.

GraphCast v1.0: A learned simulator based on GNNs

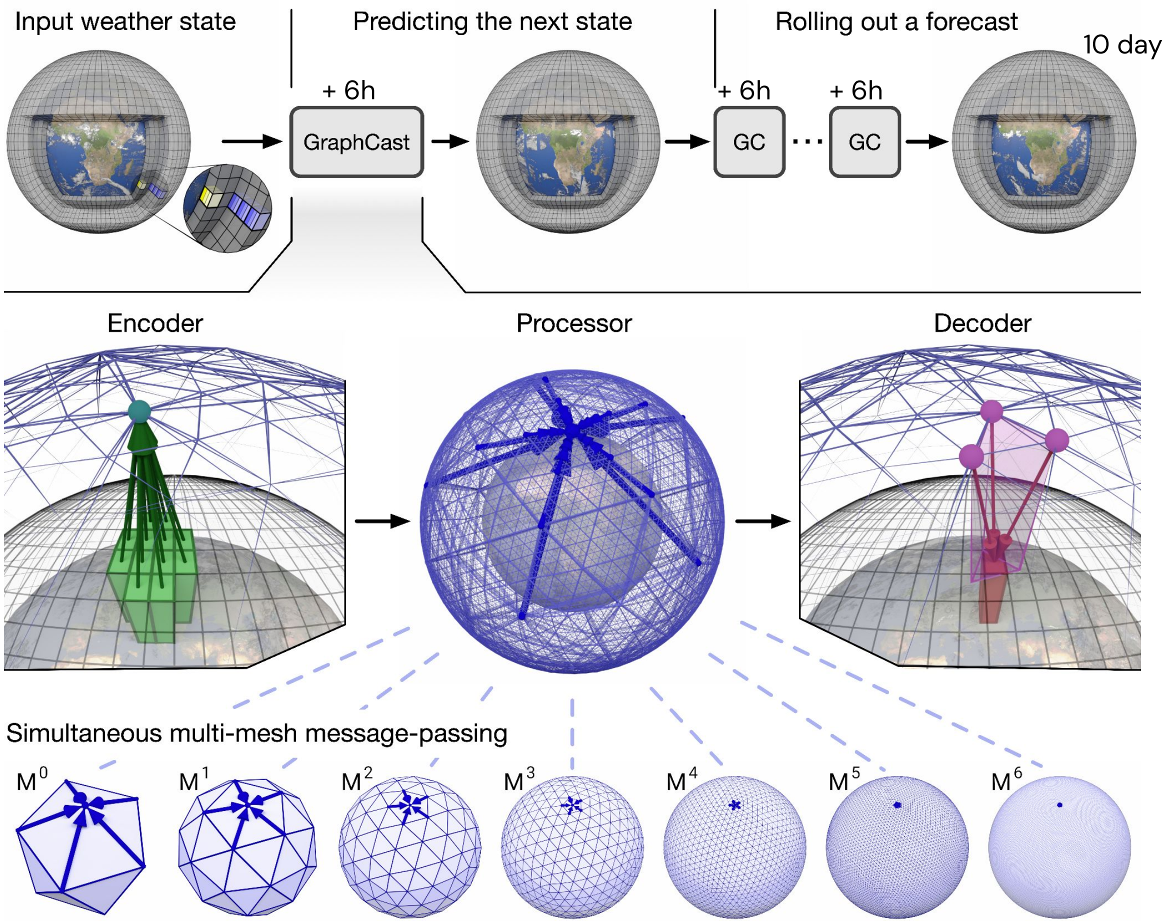

- GraphCast takes a weather state as input, and is applied iteratively to roll out a forecast.
- Only **26M** model parameters!

GraphCast v1.0 vs operational forecasts: RMSE

HRES:
Best operational model available (by ECMWF)

Example: Geopotential at 500 hPa

Absolute performance

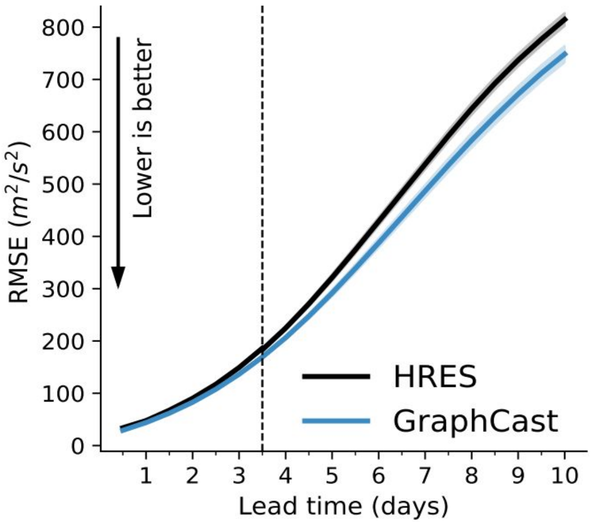

Relative performance

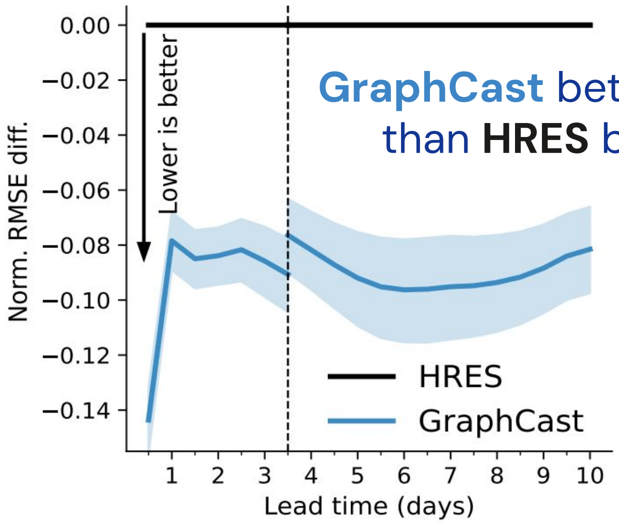

GraphCast better RMSE than **HRES** by 10%

GraphCast better than **HRES** on **99.7%** of all targets below the stratosphere

GraphCast (1 PufferFish) **>100x faster** than **HRES** (on a supercomputer)

GraphCast v1.0 vs operational forecasts: Hurricane tracking

Hurricane Maria (2017):

- Most intense storm of 2017.
- Worst storm to ever hit Dominica, Saint Croix, Puerto Rico.
- ~$100B in damage. 3rd costliest storm on record.

HRES

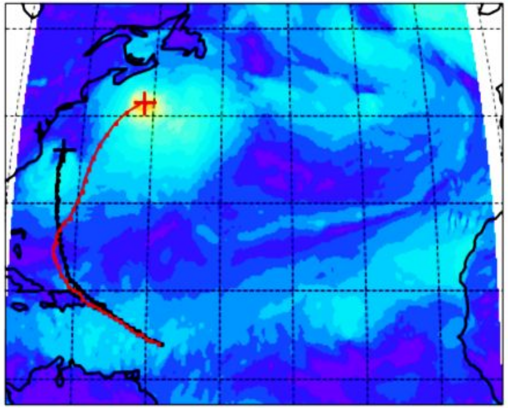

GraphCast

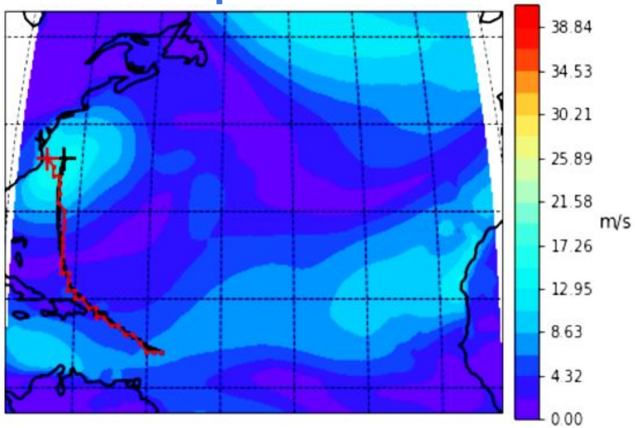

All hurricanes from 2018-2021

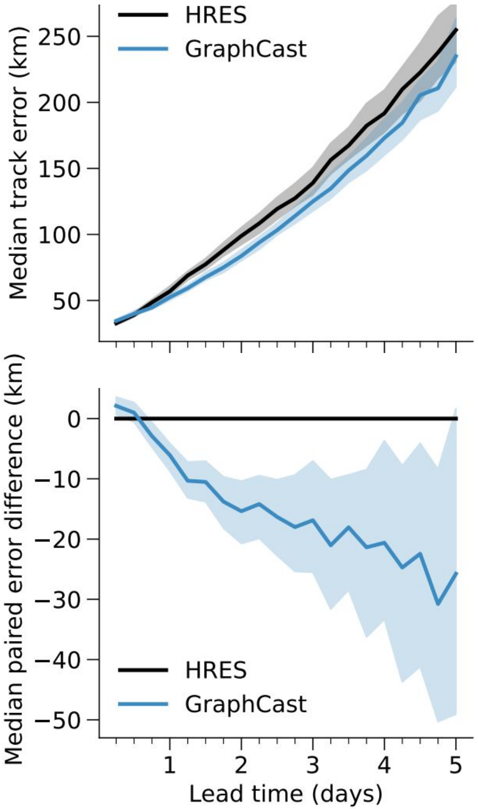

Severe weather

Extreme heat

Predict when surface temperature will reach top 2% extremes
- GraphCast dominates at long lead times.

Atmospheric rivers

'Rivers in the sky' transporting water vapor from tropics.
- GraphCast has ~24h advantage in predicting IVT

Credit: Mark Ross, [Scientific American]

Performance in extreme events is consistent with RMSE results.

GraphCast available to the world!

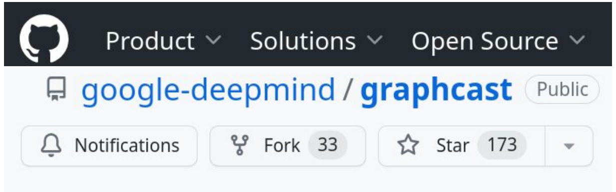

ECMWF running GraphCast online

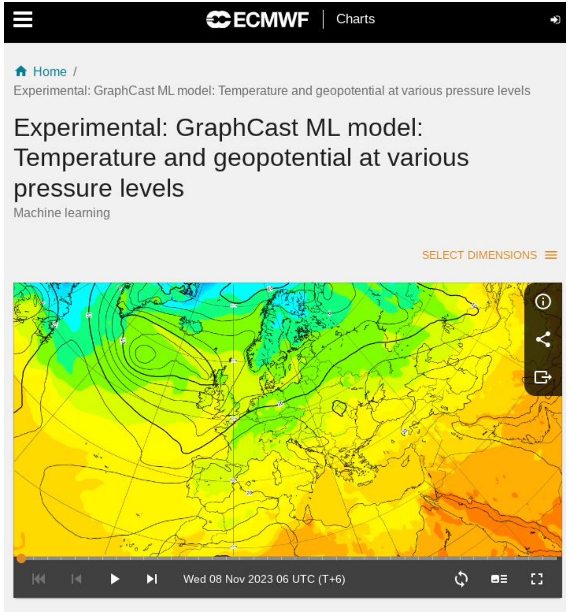

Traction in the community via **real-time tests**

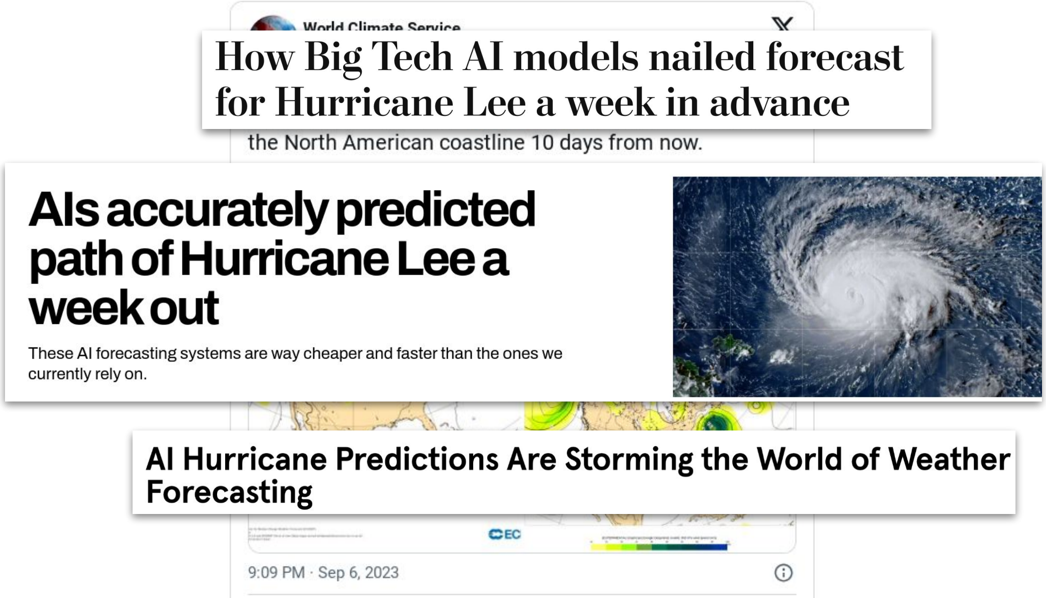

“GenCast: Diffusion-based ensemble forecasting for medium-range weather” (2023)

Price at al.,

arxiv.org/abs/2312.15796

Probabilistic version of GraphCast:

- Train as **diffusion model**

What do model predictions look like?

Region of interest

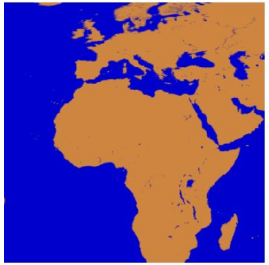

Example ERA5 state

2018-09-20 6pm

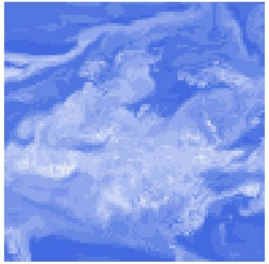

Specific humidity at 700 hPa

	GraphCast	GenCast Samples	GenCast Sample Mean
12h earlier forecast			
5d earlier forecast			
10d earlier forecast			

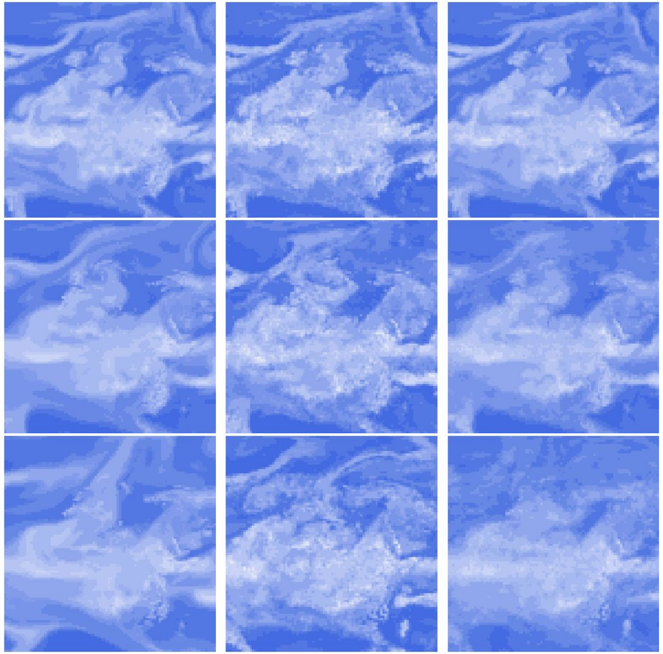

Uncertainty as blurring!

Each sample is sharp!

Still captures uncertainty!

Conclusions

- AI for physical simulation is **here to stay**!

Thanks for your attention!

Question time?