Machine learning for particle based simulations

Olof Mogren, PhD
RISE Learning Machines Seminars
2020-04-23

What is physics?

Physics (from Ancient Greek: φυσική (ἐπιστήμη), romanized: *physikḗ (epistḗmē)*, lit. 'knowledge of nature', from φύσις *phýsis* 'nature')[1][2][3] is the natural science that studies matter,[4] its motion and behavior through space and time, and the related entities of energy and force.

Wikipedia

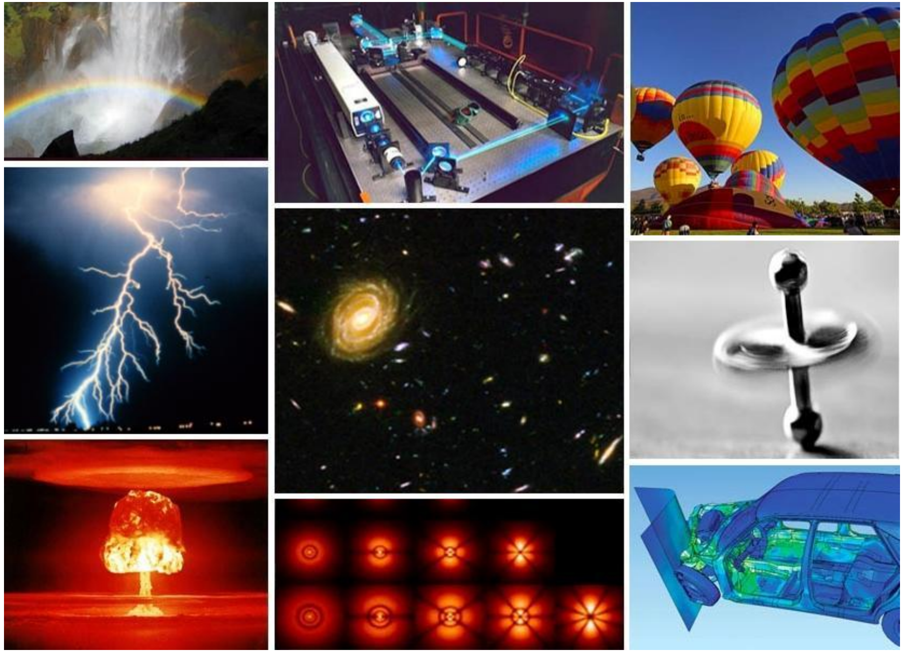

(Today: *classical mechanics*).

What is modelling?

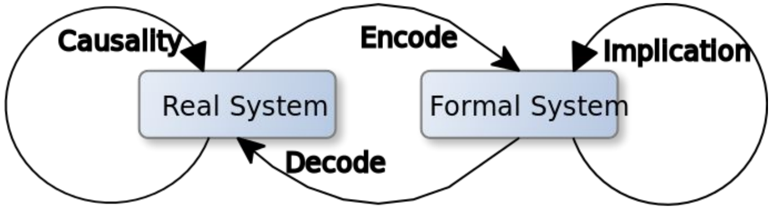

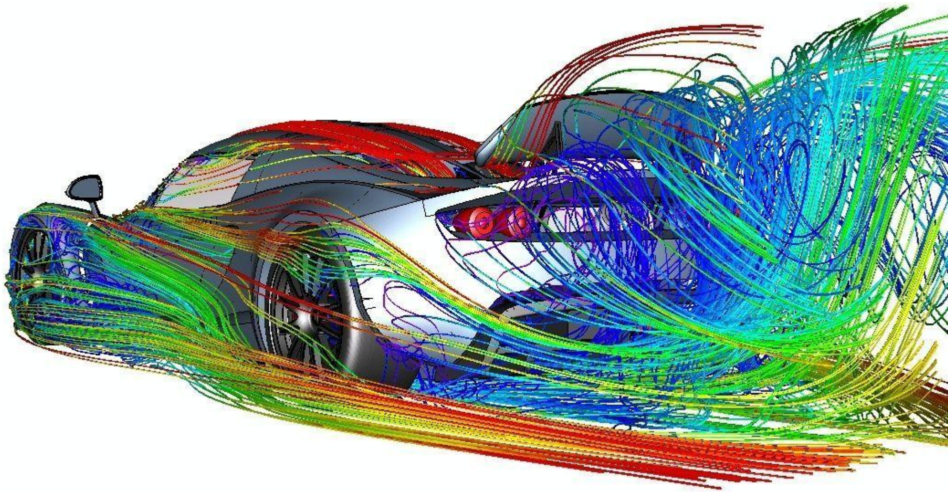

Scientific modelling is a scientific activity, the aim of which is to make a particular part or feature of the world easier to understand, define, quantify, visualize, or simulate by referencing it to existing and usually commonly accepted knowledge.

Wikipedia

What is machine learning?

Machine learning algorithms build a mathematical model based on sample data, known as "training data", in order to make predictions or decisions without being explicitly programmed to do so.

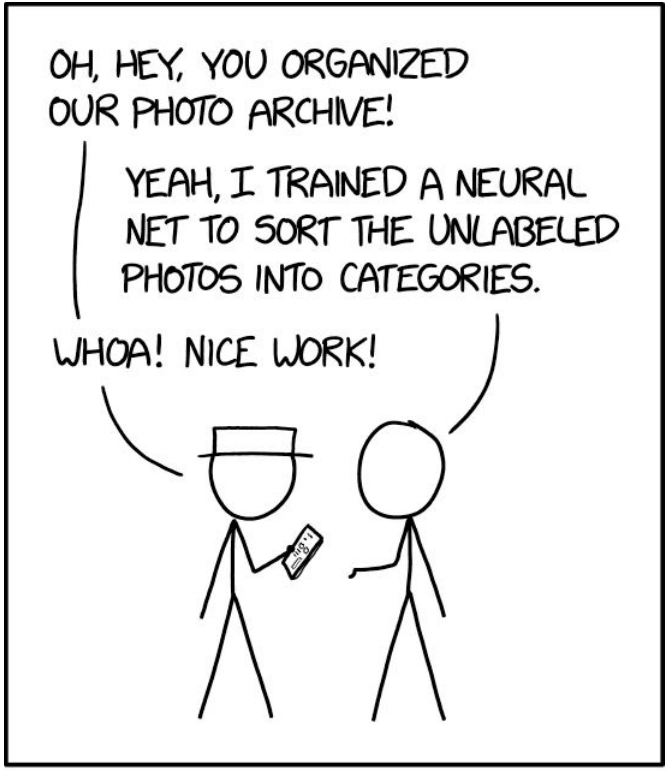

ENGINEERING TIP:
WHEN YOU DO A TASK BY HAND, YOU CAN TECHNICALLY SAY YOU TRAINED A NEURAL NET TO DO IT.

Wikipedia, xkcd.com

Machine learning for physics

Regression forests

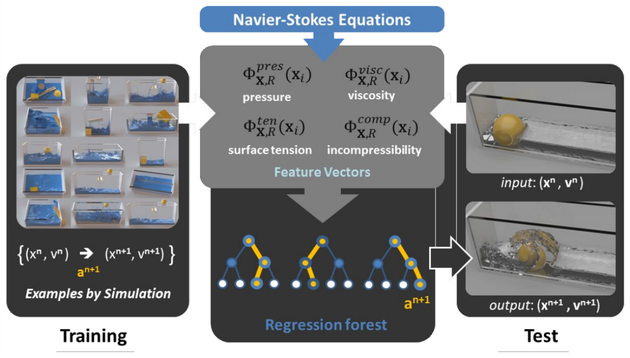

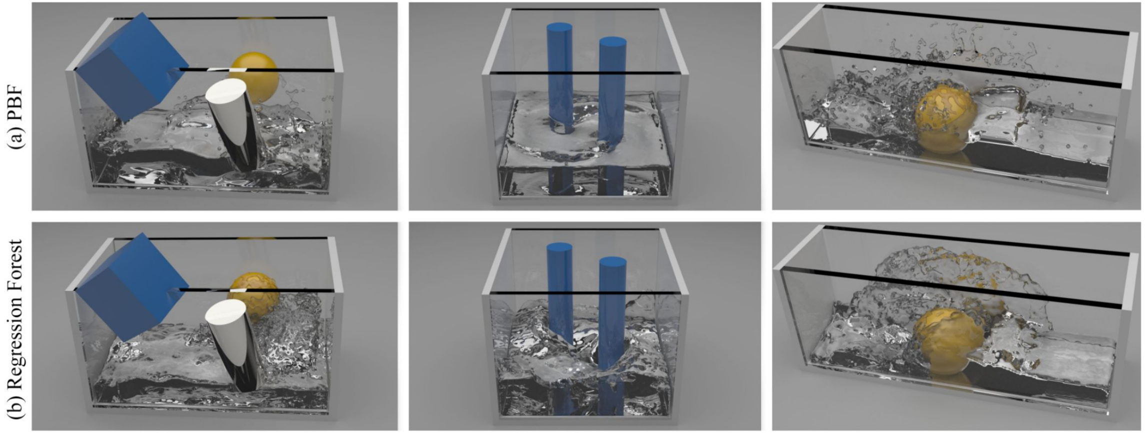

Approximating Navier-Stokes equations on a Lagrangian system (momentum, mass, energy)

Predicting acceleration (individual forces and the incompressibility constraint)

Training data from PBF solver (position based fluids)

Ladicky, et.al. (2015)

Convnets, LSTMS

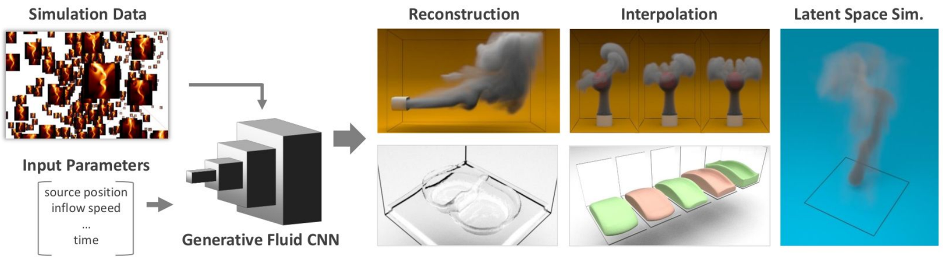

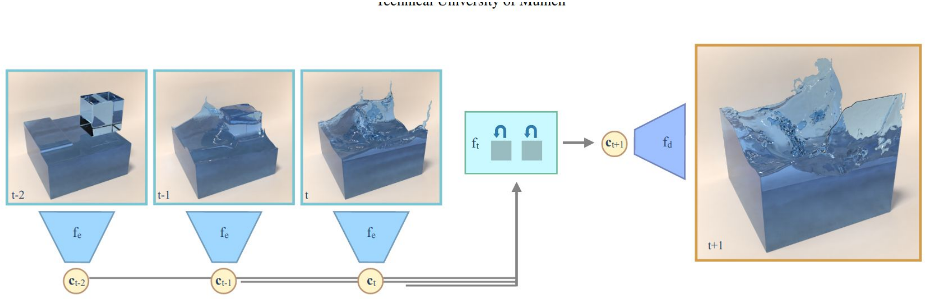

“Deep fluids”

2D, 3D

Turbulent smoke, “gooey” liquids

Kim, et.al. (Eurographics, 2019)

“Latent Space Physics: Towards Learning the Temporal Evolution ofFluid Flow”

LSTM-CNN hybrid

Wiewel, et.al. (arxiv:1802.10123)

Graph networks

Vertex features, ***v***

Edge features, ***e***

Global features, ***u***

Graph in, graph out

1. Message passing phase
2. Read out phase

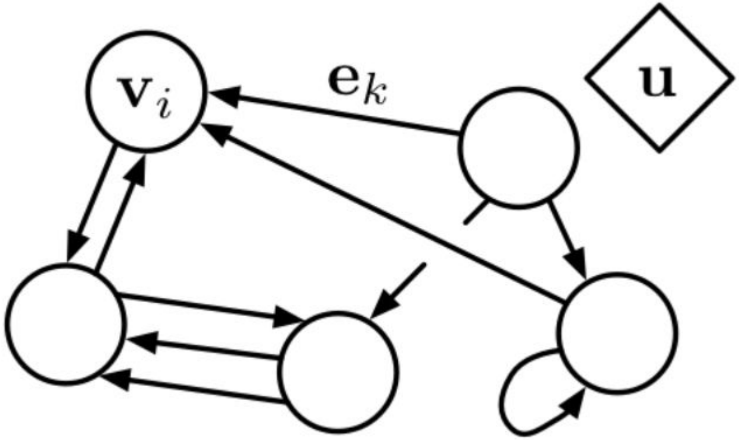

Graph Networks (GNs)

- A GN block is a “graph-to-graph” function approximator
  - The output graph’s structure (number of nodes and edge connectivity) matches the input graph’s
  - The output graph-, node-, and edge-level attributes will be functions of the input graph’s

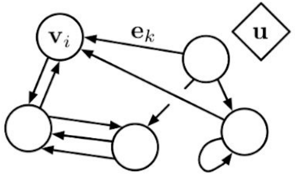

Graph Network
(a type of Graph Neural Network)
Battaglia et al., 2018, arXiv

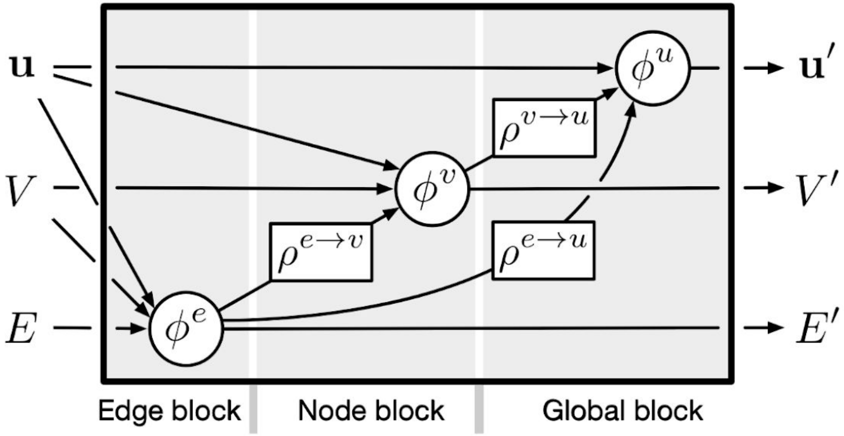

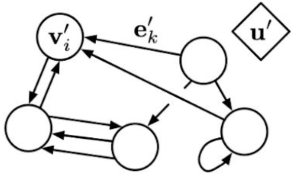

Peter Battaglia

General graph network processing pipeline

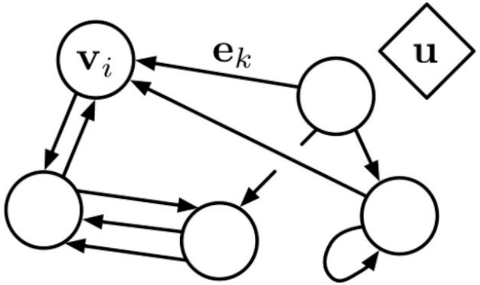

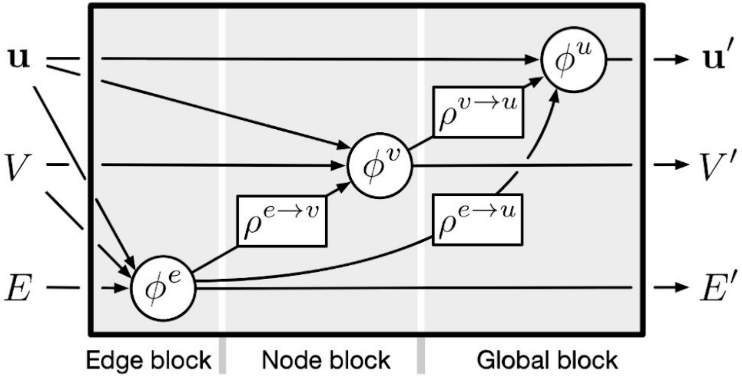

Edge block

For each edge, $\mathbf{e}_k, \mathbf{v}_{s_k}, \mathbf{v}_{r_k}, \mathbf{u}$,
are passed to an "edge-wise function":

$$\mathbf{e}'_k \leftarrow \phi^e\left(\mathbf{e}_k, \mathbf{v}_{r_k}, \mathbf{v}_{s_k}, \mathbf{u}\right)$$

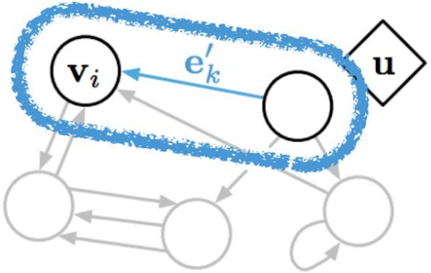

Node block

For each node, $\bar{\mathbf{e}}'_i, \mathbf{v}_i, \mathbf{u}$, are passed to a "node-wise function":

$$\mathbf{v}'_i \leftarrow \phi^v\left(\bar{\mathbf{e}}'_i, \mathbf{v}_i, \mathbf{u}\right)$$

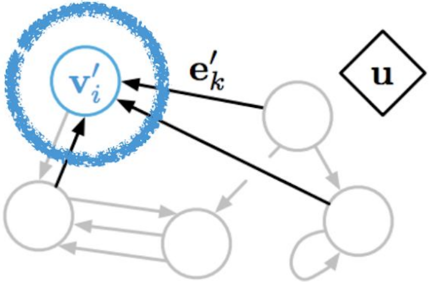

Global block

Across the graph, $\bar{\mathbf{e}}', \bar{\mathbf{v}}', \mathbf{u}$, are passed to a "global function":

$$\mathbf{u}' \leftarrow \phi^u\left(\bar{\mathbf{e}}', \bar{\mathbf{v}}', \mathbf{u}\right)$$

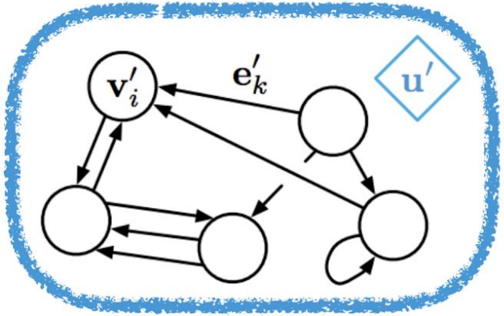

Peter Battaglia

Composing GN blocks

The GN's graph-to-graph interface promotes stacking GN blocks, passing one GN's output to another GN as input

Shared GN core

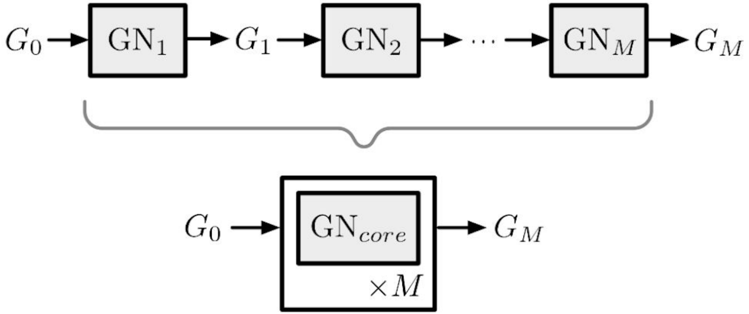

Encode-process-decode

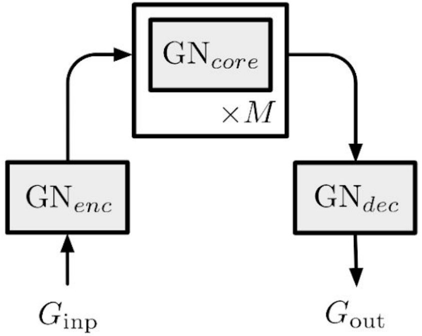

Recurrent GN architecture

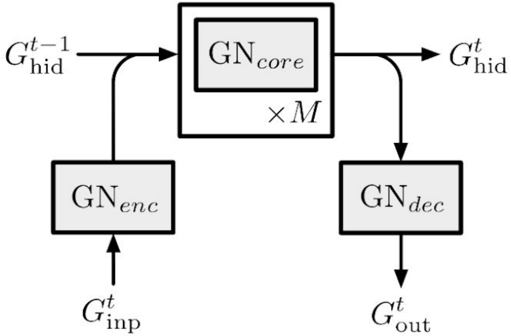

Battaglia et al., 2018, arXiv

The Graph Neural Network: History

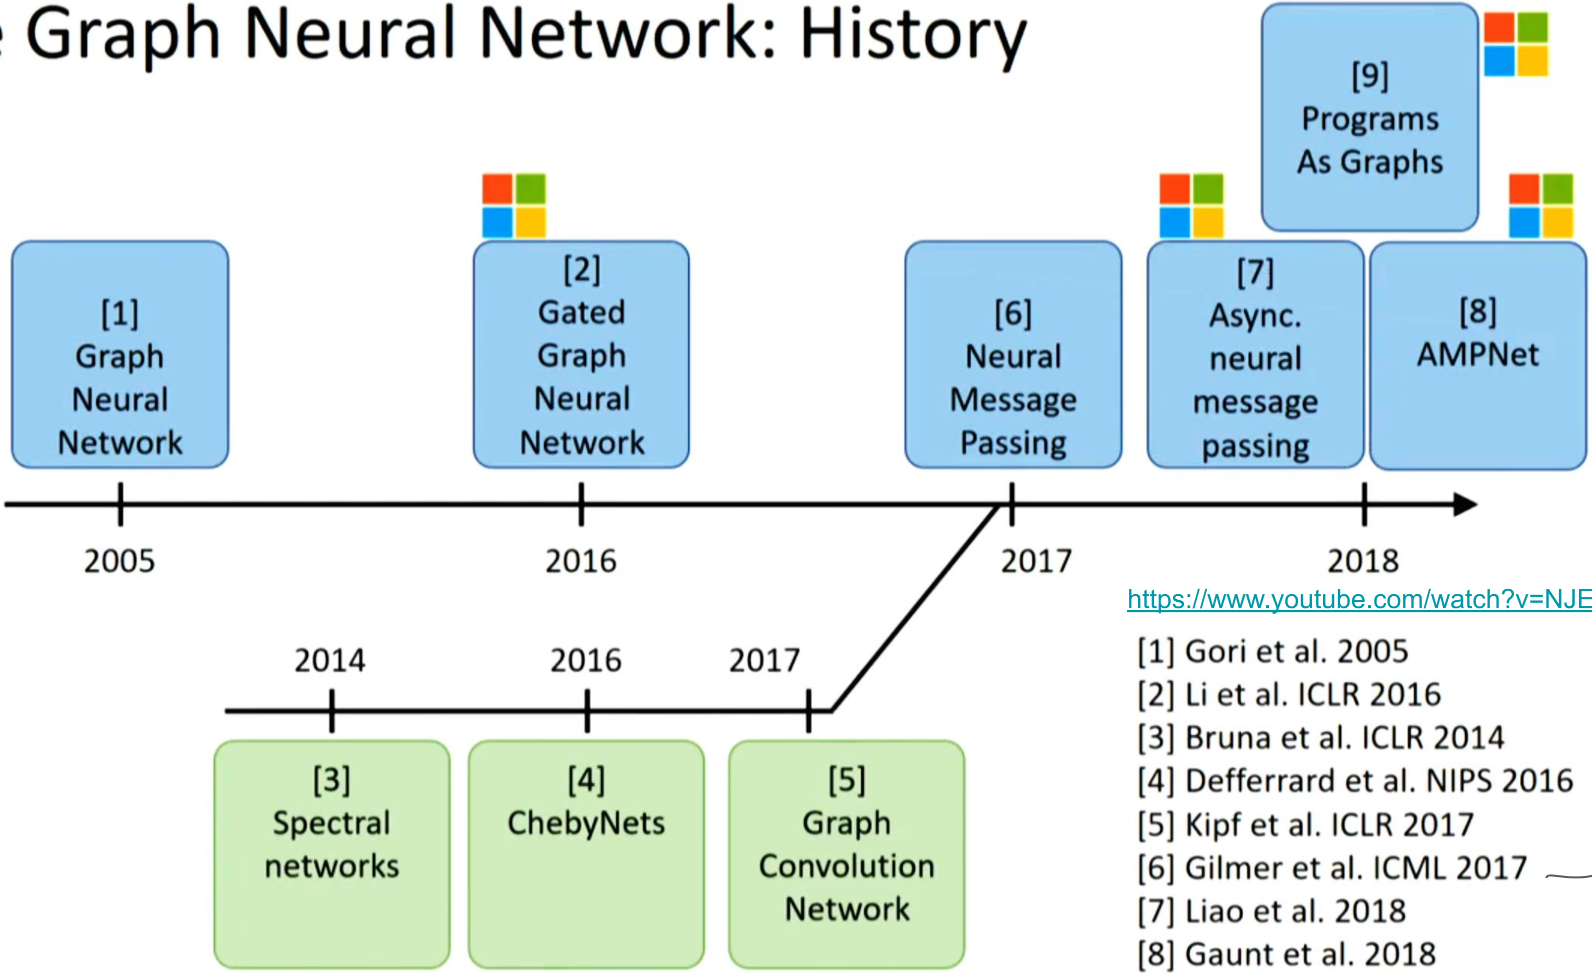

https://www.youtube.com/watch?v=NJEb5sqjv2w

[1] Gori et al. 2005
[2] Li et al. ICLR 2016
[3] Bruna et al. ICLR 2014
[4] Defferrard et al. NIPS 2016
[5] Kipf et al. ICLR 2017
[6] Gilmer et al. ICML 2017
[7] Liao et al. 2018
[8] Gaunt et al. 2018
[9] Allamanis et al. 2018

Alex Gaunt, Microsoft Research

Interaction Network: Learning simulation as message-passing

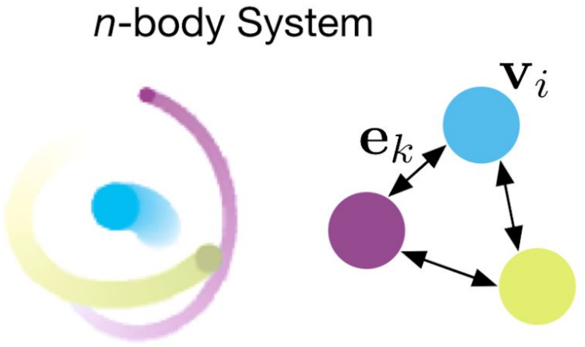

Battaglia et al., 2016, NeurIPS

Interaction Network: Learning simulation as message-passing

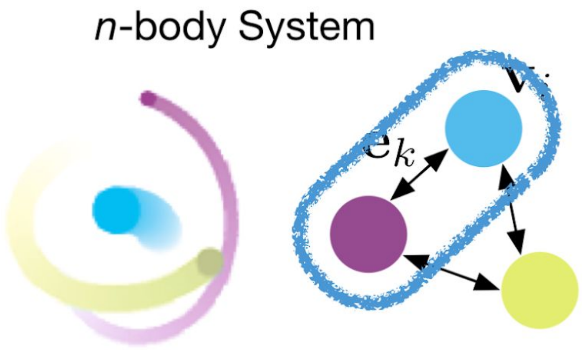

Edge function

$$\mathbf{e}'_k \leftarrow \phi^e(\mathbf{e}_k, \mathbf{v}_{r_k}, \mathbf{v}_{s_k})$$

- Compute "message" from node and edge attributes associated with an edge

Battaglia et al., 2016, NeurIPS

Interaction Network: Learning simulation as message-passing

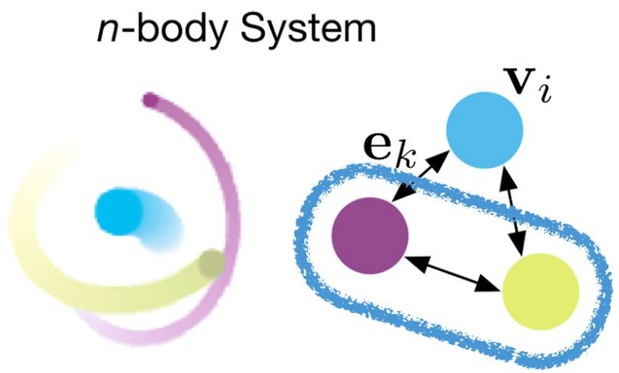

Edge function

$$\mathbf{e}'_k \leftarrow \phi^e(\mathbf{e}_k, \mathbf{v}_{r_k}, \mathbf{v}_{s_k})$$

- Compute “message” from node and edge attributes associated with an edge

Battaglia et al., 2016, NeurIPS

Interaction Network: Learning simulation as message-passing

n-body System

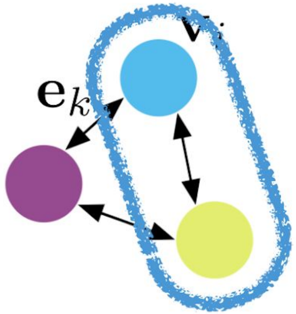

Edge function

$$\mathbf{e}'_k \leftarrow \phi^e(\mathbf{e}_k, \mathbf{v}_{r_k}, \mathbf{v}_{s_k})$$

- Compute “message” from node and edge attributes associated with an edge

Battaglia et al., 2016, NeurIPS

Interaction Network: Learning simulation as message-passing

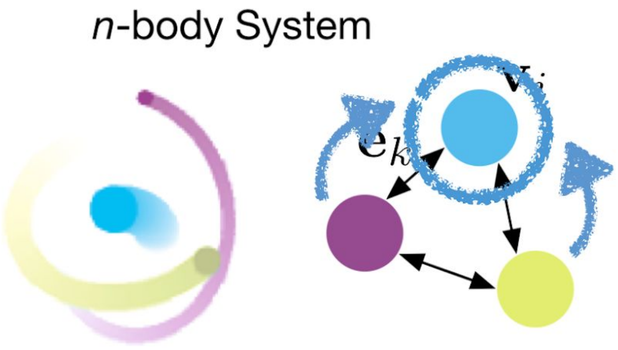

Edge function

$$\mathbf{e}'_k \leftarrow \phi^e(\mathbf{e}_k, \mathbf{v}_{r_k}, \mathbf{v}_{s_k})$$

- Compute "message" from node and edge attributes associated with an edge

Message aggregation

$$\bar{\mathbf{e}}'_i \leftarrow \sum_{r_k=i} \mathbf{e}'_k$$

Node function

$$\mathbf{v}'_i \leftarrow \phi^v(\bar{\mathbf{e}}'_i, \mathbf{v}_i, \mathbf{u})$$

- Update node info from previous node state and aggregated "messages"

Battaglia et al., 2016, NeurIPS

Interaction Network: Learning simulation as message-passing

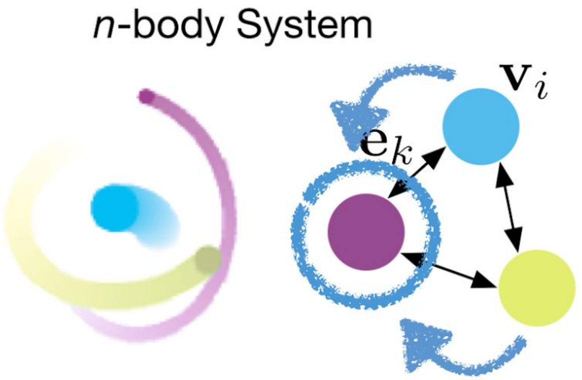

Edge function

$$\mathbf{e}'_k \leftarrow \phi^e(\mathbf{e}_k, \mathbf{v}_{r_k}, \mathbf{v}_{s_k})$$

- Compute "message" from node and edge attributes associated with an edge

Message aggregation

$$\bar{\mathbf{e}}'_i \leftarrow \sum_{r_k=i} \mathbf{e}'_k$$

Node function

$$\mathbf{v}'_i \leftarrow \phi^v(\bar{\mathbf{e}}'_i, \mathbf{v}_i, \mathbf{u})$$

- Update node info from previous node state and aggregated "messages"

Battaglia et al., 2016, NeurIPS

Interaction Network: Predicting potential energy

n-body System

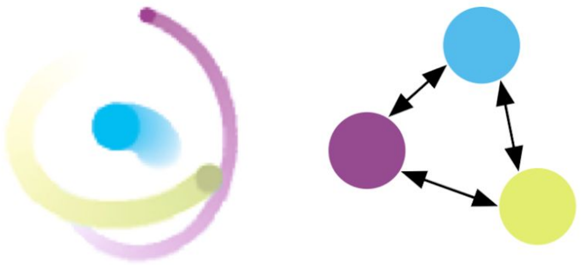

Node aggregation and global function

$$\bar{\mathbf{v}}' \leftarrow \sum_i \mathbf{v}'_i$$

$$\mathbf{u}' \leftarrow \phi^u(\bar{\mathbf{v}}')$$

- Rather than making node-wise predictions, node updates can be used to make global predictions.

Trained to predict system's potential energy

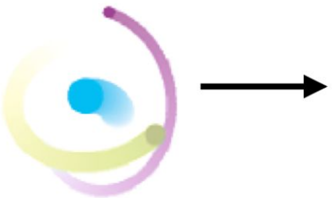

P.E.

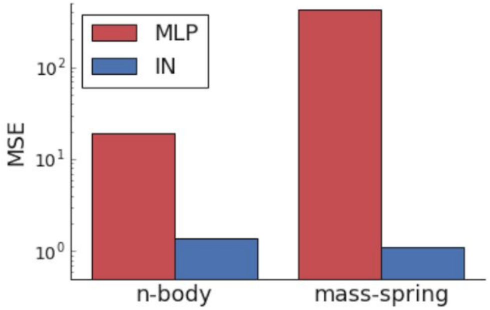

Battaglia et al., 2016, NeurIPS

Interaction networks

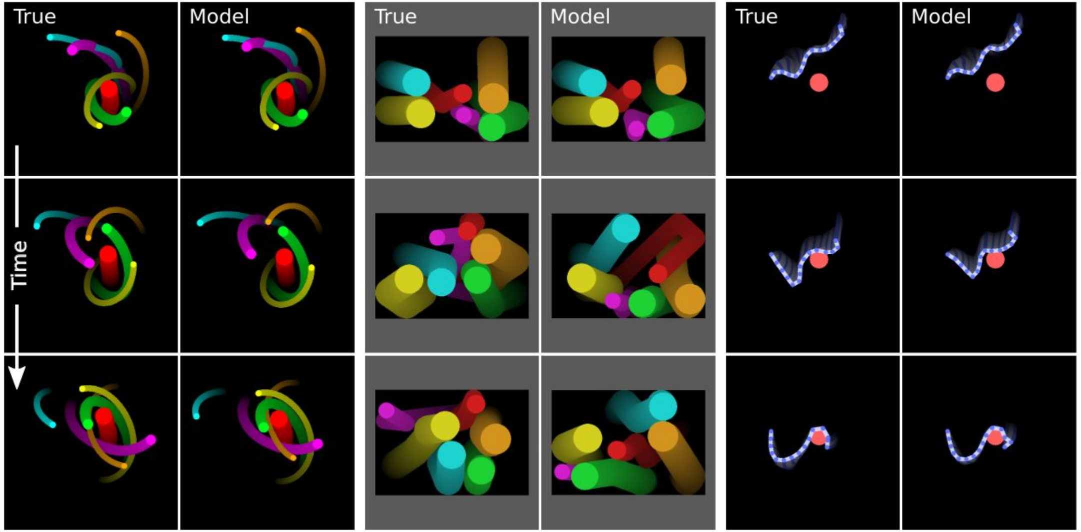

Learning interactions and trajectories from simulations

Simulated data: n-body systems; balls bouncing in a box; and strings composed of springs that collide with rigid objects

Graph neural networks: Relation encoder MLP, object encoder MLP

Generalize to larger systems

Train on single-step, predict using roll-outs

Output: x,y velocity

Objects: n-body objects, balls, walls, points masses that represented string elements
Object state: dynamic state component (e.g., position and velocity) and a static attribute component (e.g., mass, size, shape)

Relations: e.g., gravitational attraction, collisions, springs

Battaglia, et.al. (NeurIPS, 2016)

Learning object interactions using video

Learning physics from video

2D data from simulator, drawn on top of images from CIFAR-10

Visual interaction network: 1.Visual encoder →2.Dynamics prediction →3.State decoder

(1.Convnet →2.MLP→3.Output layer)

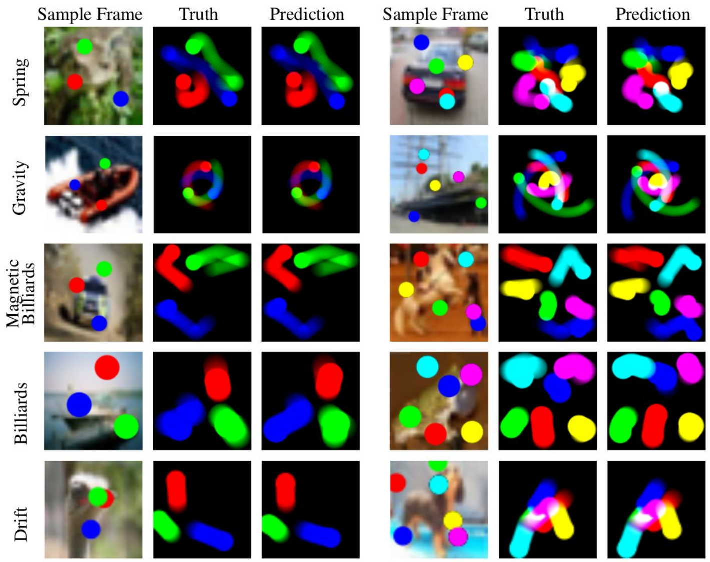

Watters, et.al. (NeurIPS, 2017)

FFWD: Graph neural networks for physics simulations

3D Data from CFD simulators

Several kinds of particles, e.g. liquid, goop, sand, solid.

Generalize to new initial conditions, more particles, many timesteps.

Input particle state: position, 5 previous velocities, static material properties (e.g., water, sand, goop, rigid, boundary particle)

Edged: added to particle pairs at a connectivity radius < *R*

Output: particle acceleration (simulate using Euler integration)

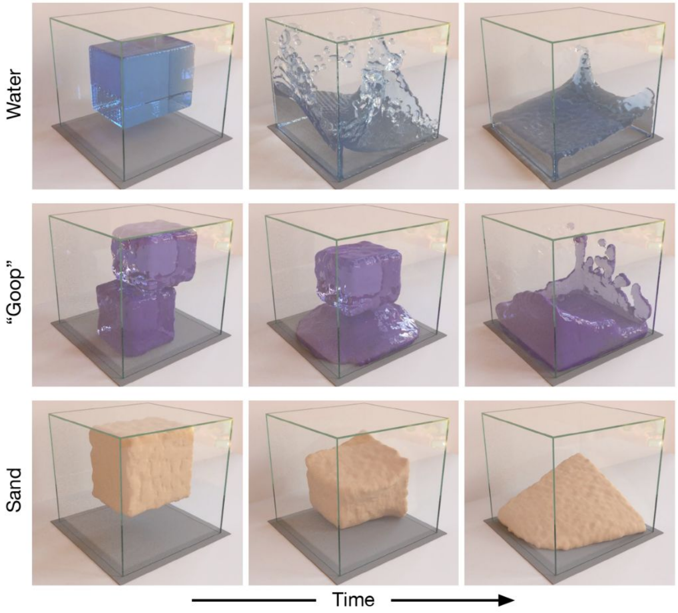

Sanchez-Gonzalez, et.al. (arxiv:2002.09405)

Generated training data

Simulators used:

- BOXBATH: Flex (position-based dynamics method)
- WATER-3D: SPlisHSPlasH (SPH-based; strict volume preservation)
- Other: Taichi-MPM engine

Training:

- ?k particles
- 1 timestep

Test:

- 1k-85k particles (up to 43x training size)
- 5000 timesteps

Experiments

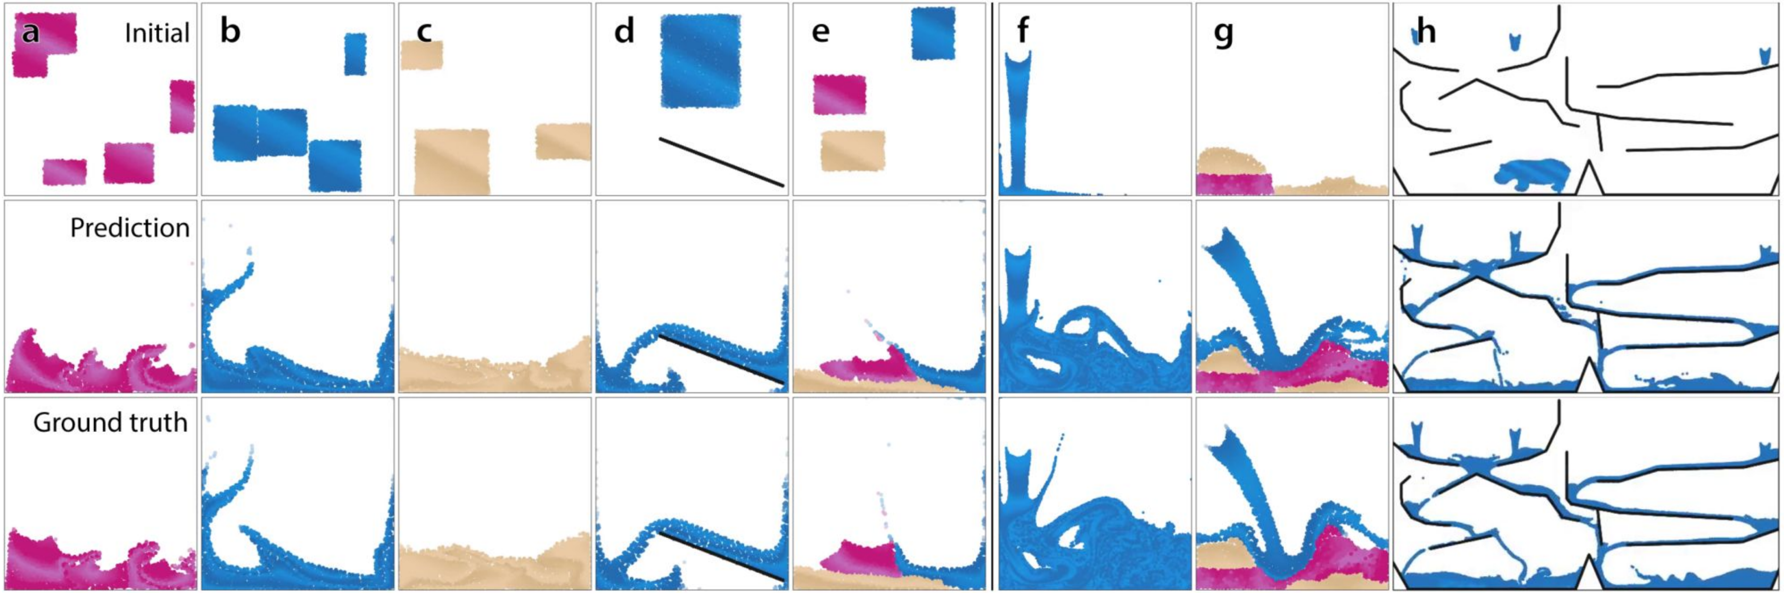

Simulations

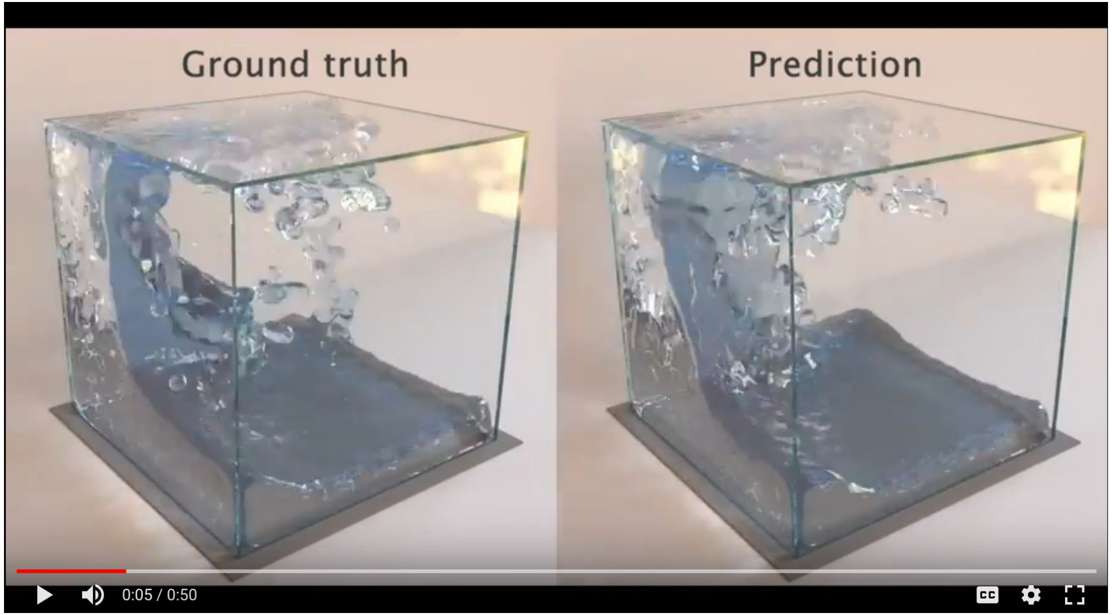

https://sites.google.com/view/learning-to-simulate/home#h.p_hjnaJ6k8y0wo

Ablations

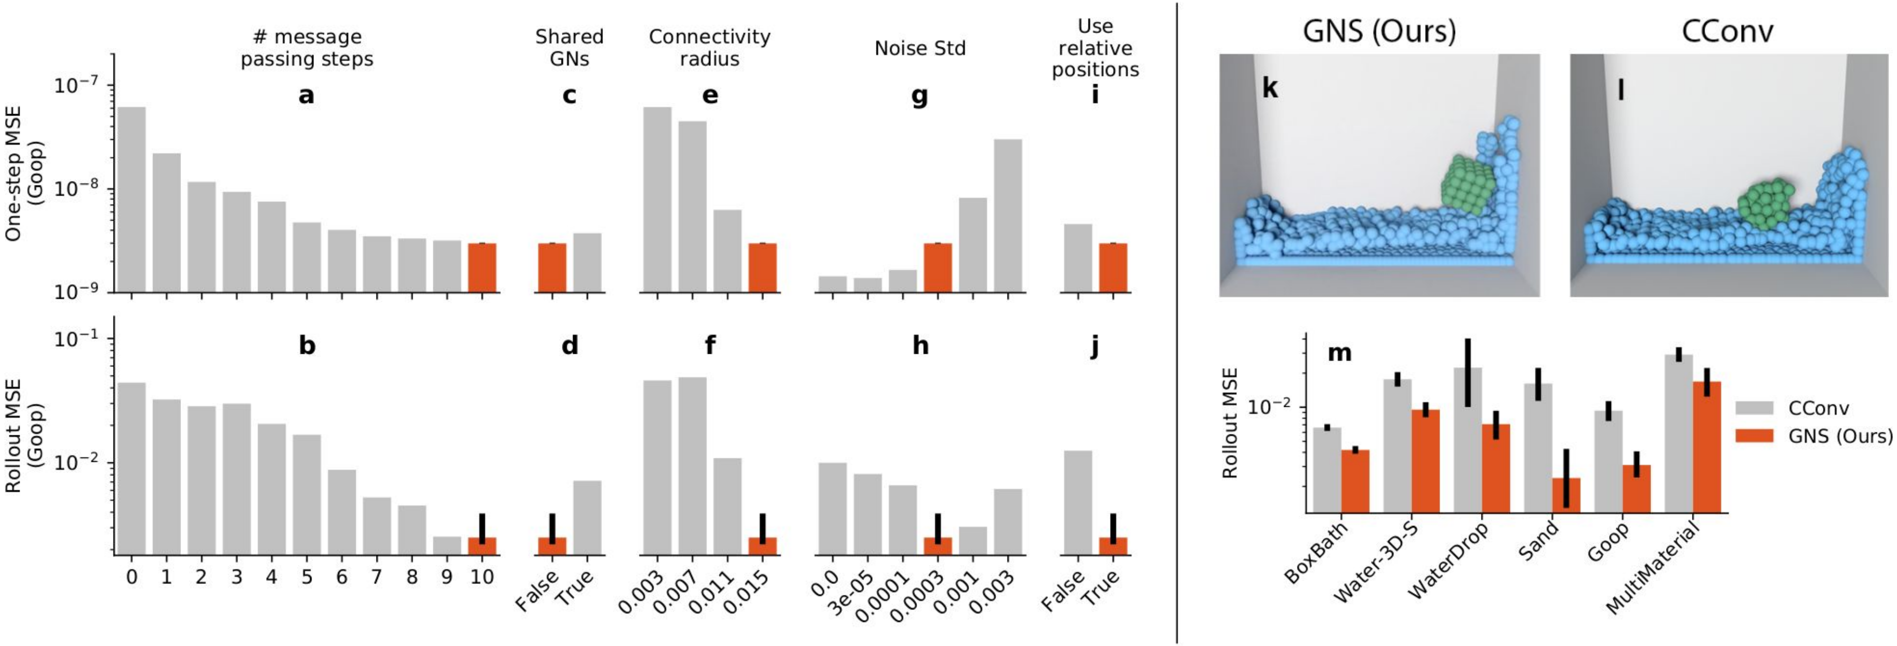

message passing steps
Shared GNs
Connectivity radius
Noise Std
Use relative positions
a
c
e
g
i
b
d
f
h
j
One-step MSE (Goop)
Rollout MSE (Goop)
0 1 2 3 4 5 6 7 8 9 10
False True
0.003 0.007 0.011 0.015
0.0 3e-05 0.0001 0.0003 0.001 0.003
False True
GNS (Ours)
CConv
k
l
m
Rollout MSE
CConv
GNS (Ours)
BoxBath
Water-3D-S
WaterDrop
Sand
Goop
MultiMaterial

Inferring system properties from simulations

Inverse modelling, infer system given data

Graph neural networks

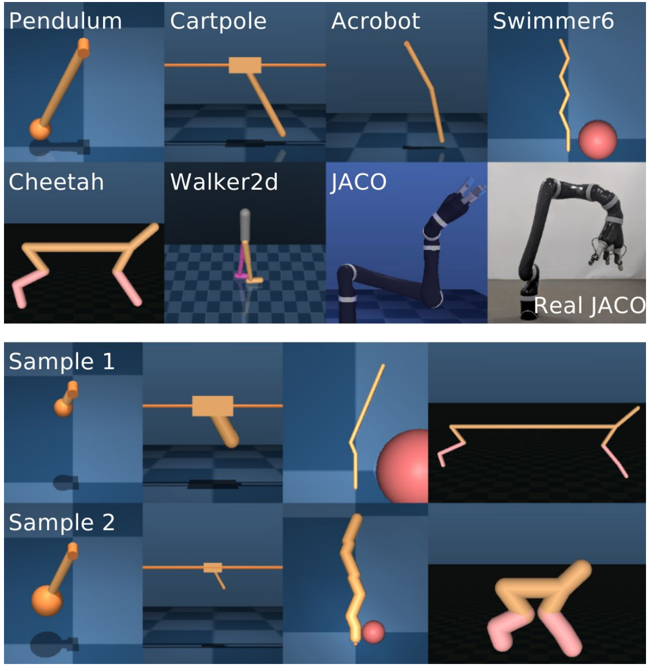

Sanchez-Gonzalez, et.al. (ICML, 2018)

MeshGraphNets

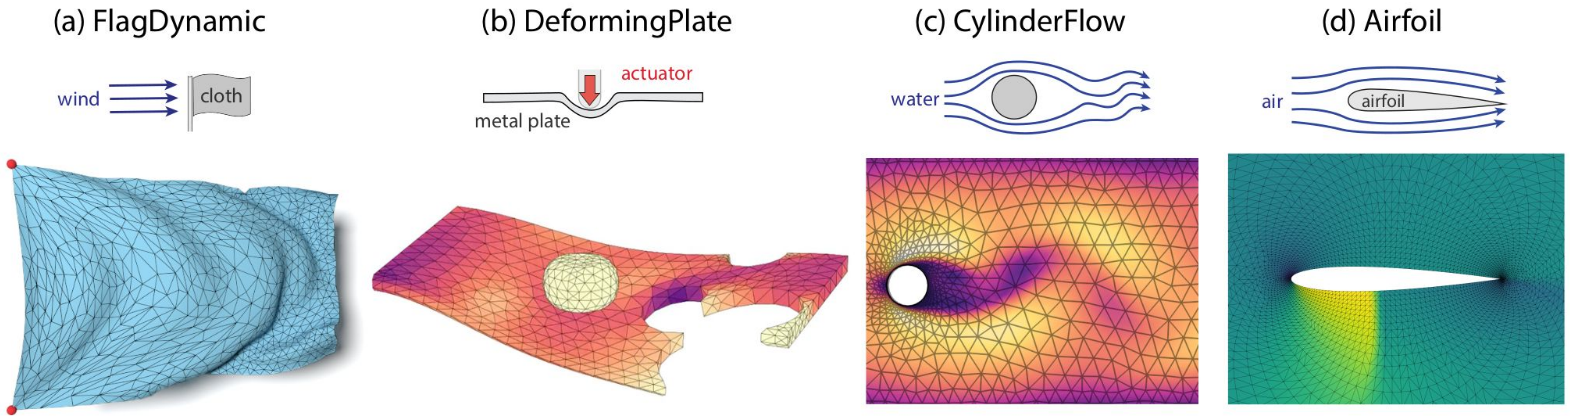

- Mesh graph
- 1-2 orders of magnitude faster

Pfaff, et.al., Learning Mesh-Based Simulation with Graph Networks, arxiv: arXiv:2010.03409

References

Mentioned today:

- *Ladicky, et.al., Data-driven Fluid Simulations using Regression Forests, ACM Transactions on Graphics 2015*
- *Battaglia, et.al., Interaction networks for learning about objects, relations and physics, NeurIPS 2016*
- *Kim, et.al., Deep Fluids: A Generative Network for Parameterized Fluid Simulations, arxiv:1806.02071*
- *Sanchez-Gonzalez, et.al., Learning to Simulate Complex Physics with Graph Networks, arxiv:2002.09405*
- *Pfaff, et.al., Learning Mesh-Based Simulation with Graph Networks, arxiv: arXiv:2010.03409*

Videos:

- *Alex Gaunt, GNNs: https://www.youtube.com/watch?v=cWIeTMklzNg&t=707s*
- *Gori, Message passing NNs: https://www.youtube.com/watch?v=NJEb5sqjv2w*
- *Peter Battaglia, Learning structured models of physics: https://www.youtube.com/watch?v=RwrzKtnSwrw*
- *Learning to simulate complex physics: https://sites.google.com/view/learning-to-simulate/home#h.p_hjnaJ6k8y0wo*

Barely or not mentioned today:

- *Gori, et.al., A New Model for Learning in Graph Domains, International Joint Conference on Neural Networks 2005*
- *Scarselli, et.al., The graph neural network model, Transactions on Neural Networks 2009*
- *Duvenaud, et.al., Convolutional networks on graphs for learning molecular fingerprints. arXiv preprint arXiv:1509.09292*
- *Agrawal, et.al., 2016, Learning to Poke by Poking: Experiential Learning of Intuitive Physics, arxiv:1606.07419*
- *Tompson, et.al., Accelerating Eulerian Fluid Simulation With Convolutional Networks, arxiv:1607.03597*
- *Gilmer, et.al., Neural Message Passing for Quantum Chemistry, arxiv:1704.01212*
- *Tassa, et.al., DeepMind Control Suite, arxiv:1801.00690*
- *Sanchez-Gonzalez, et.al., Hamiltonian graph networks with ode integrators, ICML 2018*
- *Sanchez-Gonzalez, et.al., Graph Networks as Learnable Physics Engines for Inference and Control, ICML 2018, arxiv:1806.01242*
- *Battaglia, et.al., Relational inductive biases, deep learning, and graph networks, arxiv:1806.01261*
- *Wiewel, et.al., Latent-space Physics: Towards Learning the Temporal Evolution of Fluid Flow, arXiv:1802.10123*