Google DeepMind

SDF-Sim: Learning rigid-body simulators over implicit shapes for large-scale scenes and vision

Yulia Rubanova, Tatiana Lopez-Guevara, Kelsey Allen, William Whitney,
Kimberly Stachenfeld, Tobias Pfaff

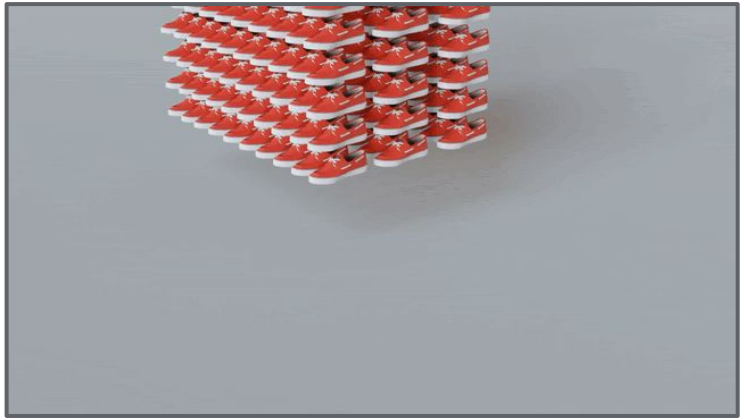

Modeling object interactions

Robotics

Look Before You Leap: Unveiling the Power of GPT-4V in Robotic Vision-Language Planning

Design

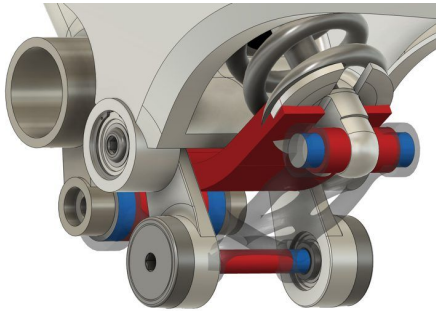

Spatial and physical reasoning

Learning simulators: train from real data

Classic simulators

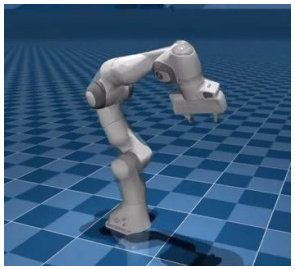

Model rigid interactions with classic simulators?

MuJoCo, Bullet, PhysX, etc.

Problem: sim-to-real gap

Simulations do not exactly match the real data

Learning simulators: train from real data

Classic simulators

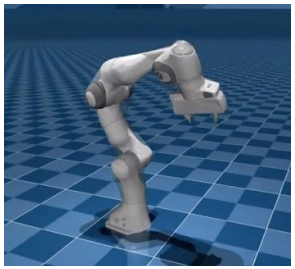

Model rigid interactions with classic simulators?

MuJoCo, Bullet, PhysX, etc.

Problem: sim-to-real gap

Simulations do not exactly match the real data

Learned simulators

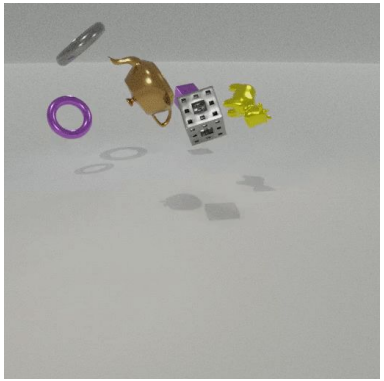

FIGNet - Allen*, Rubanova*
ICLR 2023

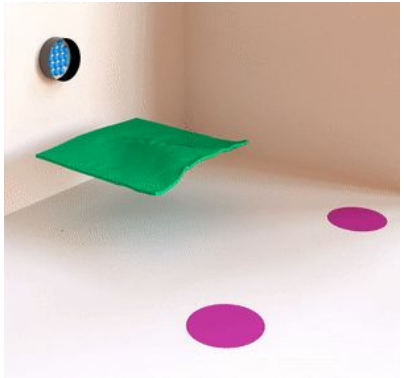

Allen*, Lopez-Guevara*, Stachenfeld*
NeurIPS 2022

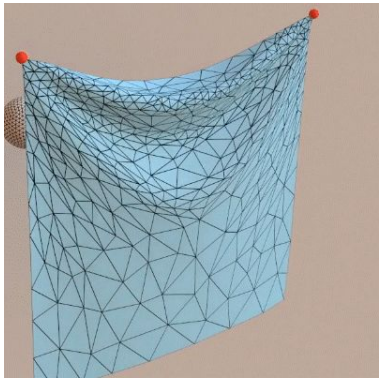

MeshGraphNets - Pfaff*, Fortunato*
ICLR 2021

Allen*, Lopez-Guevara*
CoRL 2022

FIGNet: graph network for rigid-body sim

Allen*, Rubanova* et al. Learning rigid dynamics with face interaction graph networks, ICLR 2023

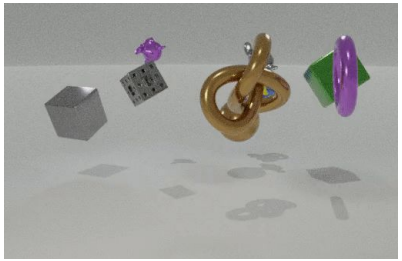

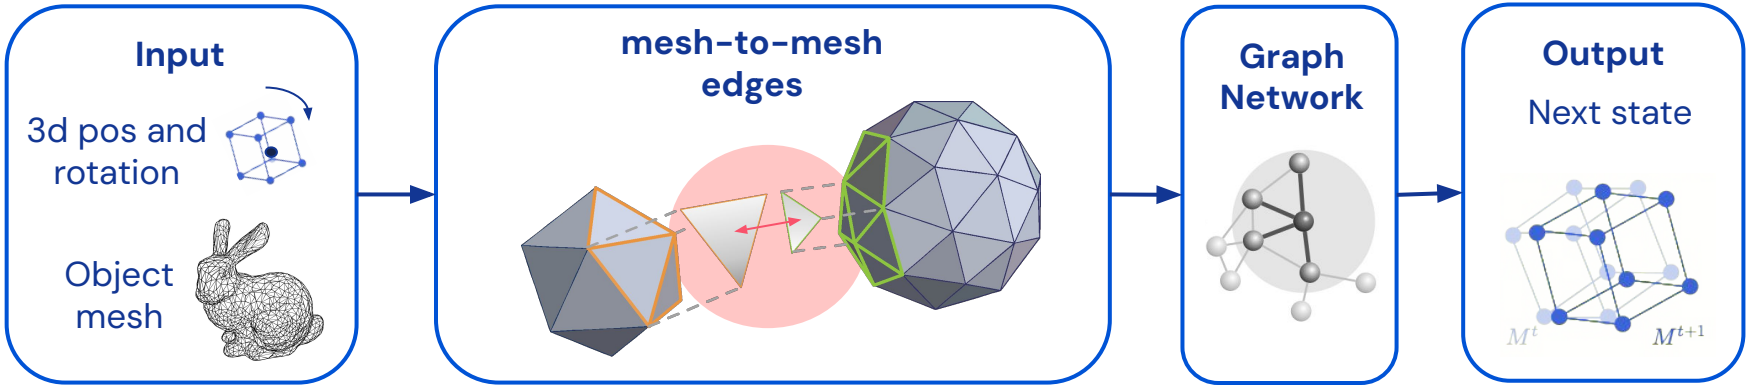

FIGNet's mesh-to-mesh computations are expensive

FIGNet relies on finding pairs of mesh triangles within a certain distance

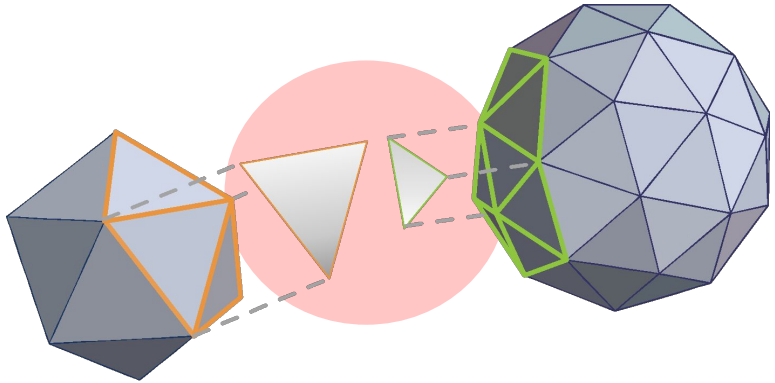

... But typical meshes have thousands of triangles

Expensive and slow!

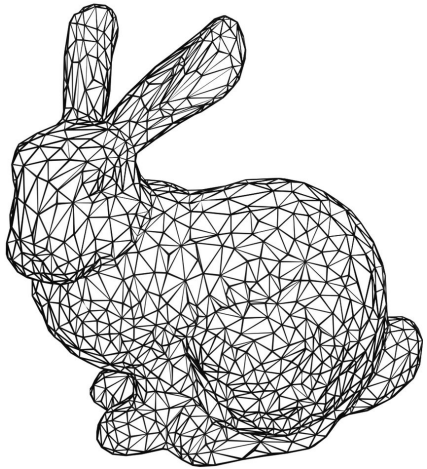

Solution: Signed Distance Functions (SDFs)

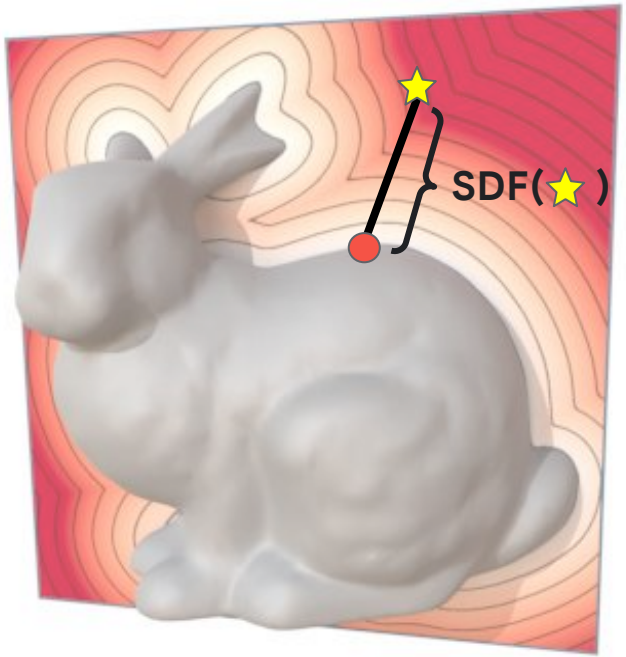

- SDFs define a distance field around the object surface
  distance to object = SDF( )
- How to get an SDF function? Learn it!

  SDF = MLP ( )

SDF-Sim: replace meshes with learned SDFs

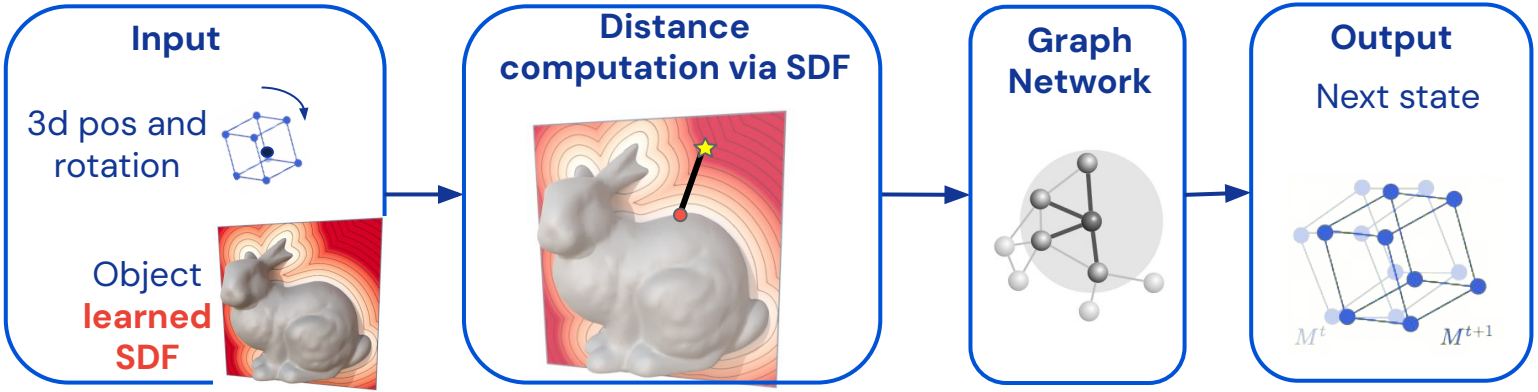

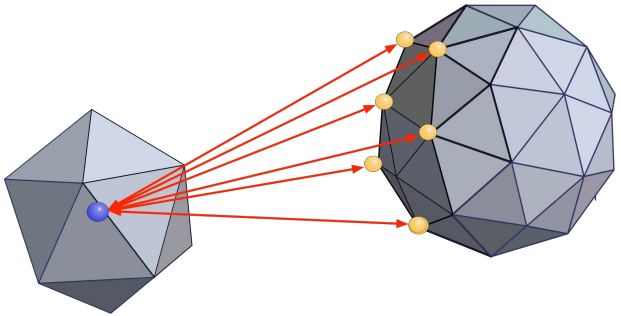

New construction of graph edges
connect nodes directly to the object center

O(n) edges instead of O(n^2) edges -> less memory and runtime required

n – number of nodes

SDF-Sim: scale to 50x larger scenes

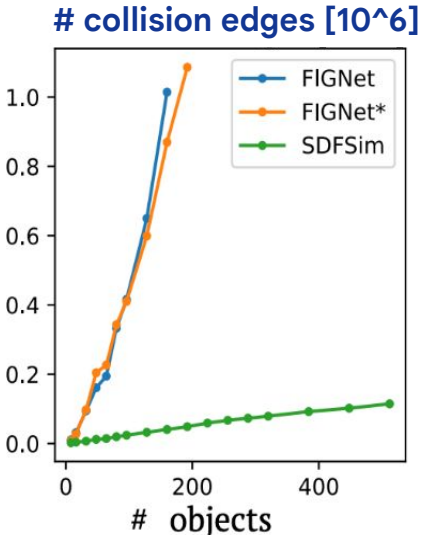

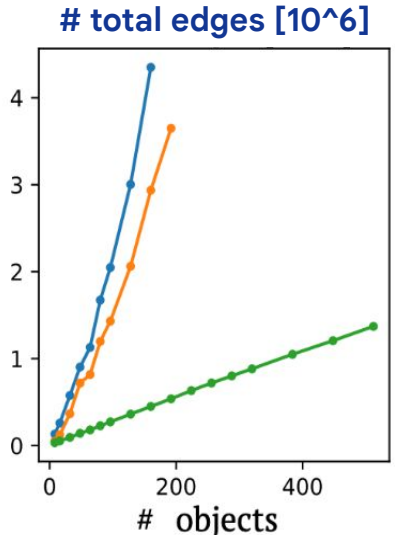

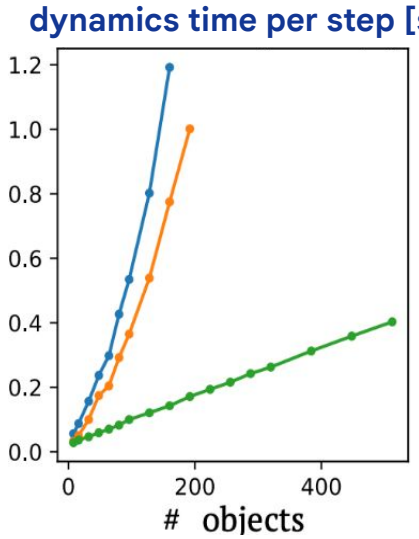

Two orders of magnitude smaller memory and runtime!

SDF-Sim: scale to 50x larger scenes

Training: <10 objects

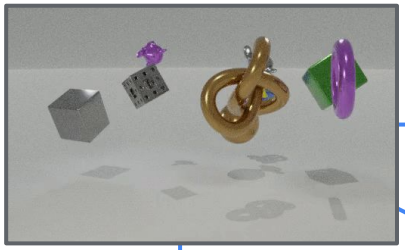

270 objects, 384k nodes, 200 timesteps

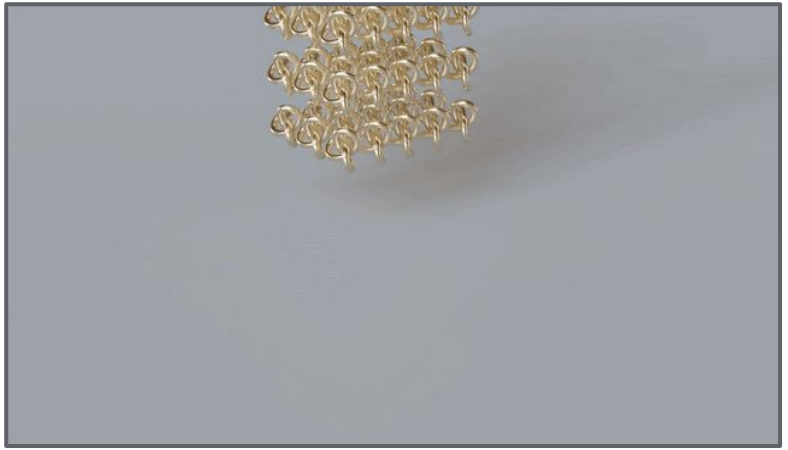

300 objects, 851k nodes, 200 timesteps

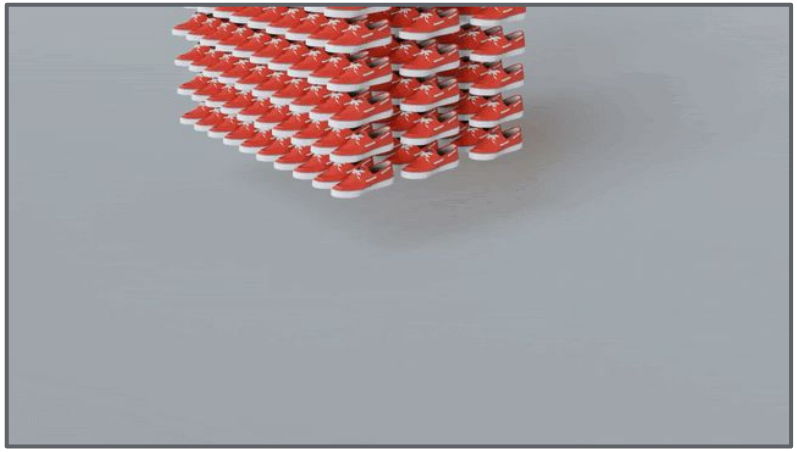

380 objects, 1.1M nodes, 400 timesteps

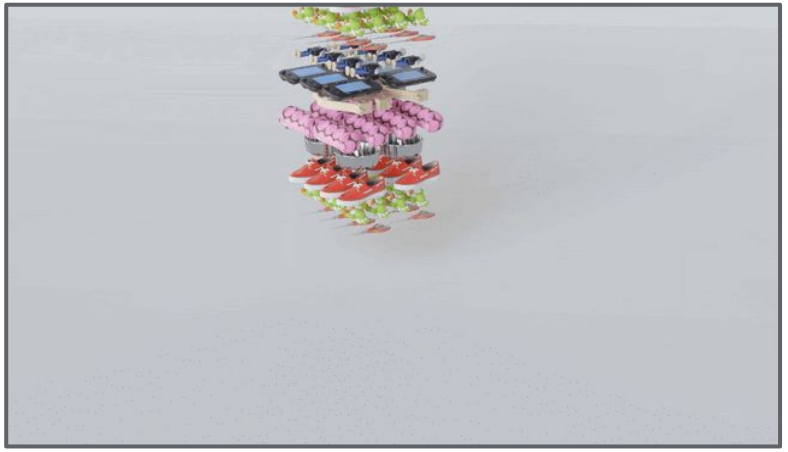

More advantages of learned SDFs

Learned SDF is memory-compact

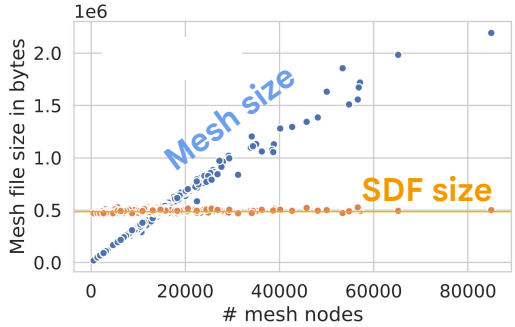

Small MLP SDFs are enough

Reconstruction quality is negligible between SDF MLP sizes

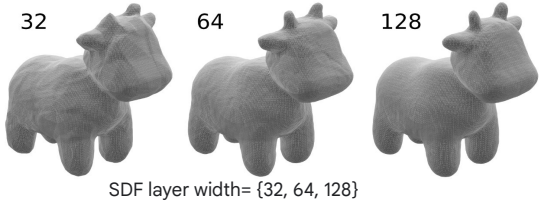

SDF layer width= {32, 64, 128}

Flexible computation

Sample desired number of object nodes from SDF

More nodes = more accurate

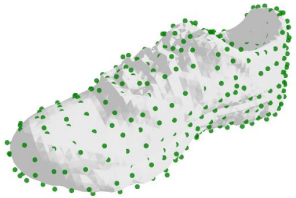

Fewer nodes = cheaper

Run simulation directly on 3D reconstructions (e.g. NeRF)

Google DeepMind

SDF-Sim: Learning rigid-body simulators over implicit shapes for large-scale scenes and vision


Yulia Rubanova


Tatiana López


Kelsey Allen


Will Whitney


Kim Stachenfeld


Tobias Pfaff

Poster session: Fri 13 Dec 11 a.m. PST — 2 p.m. PST

Oral presentation: Oral Session 5A Fri 13 Dec 10 - 10:20 a.m. PST

https://arxiv.org/abs/2405.14045