

Deep Learning on Meshes and Point Clouds

Ruben Wiersma

SGP 2025 Graduate School, Bilbao

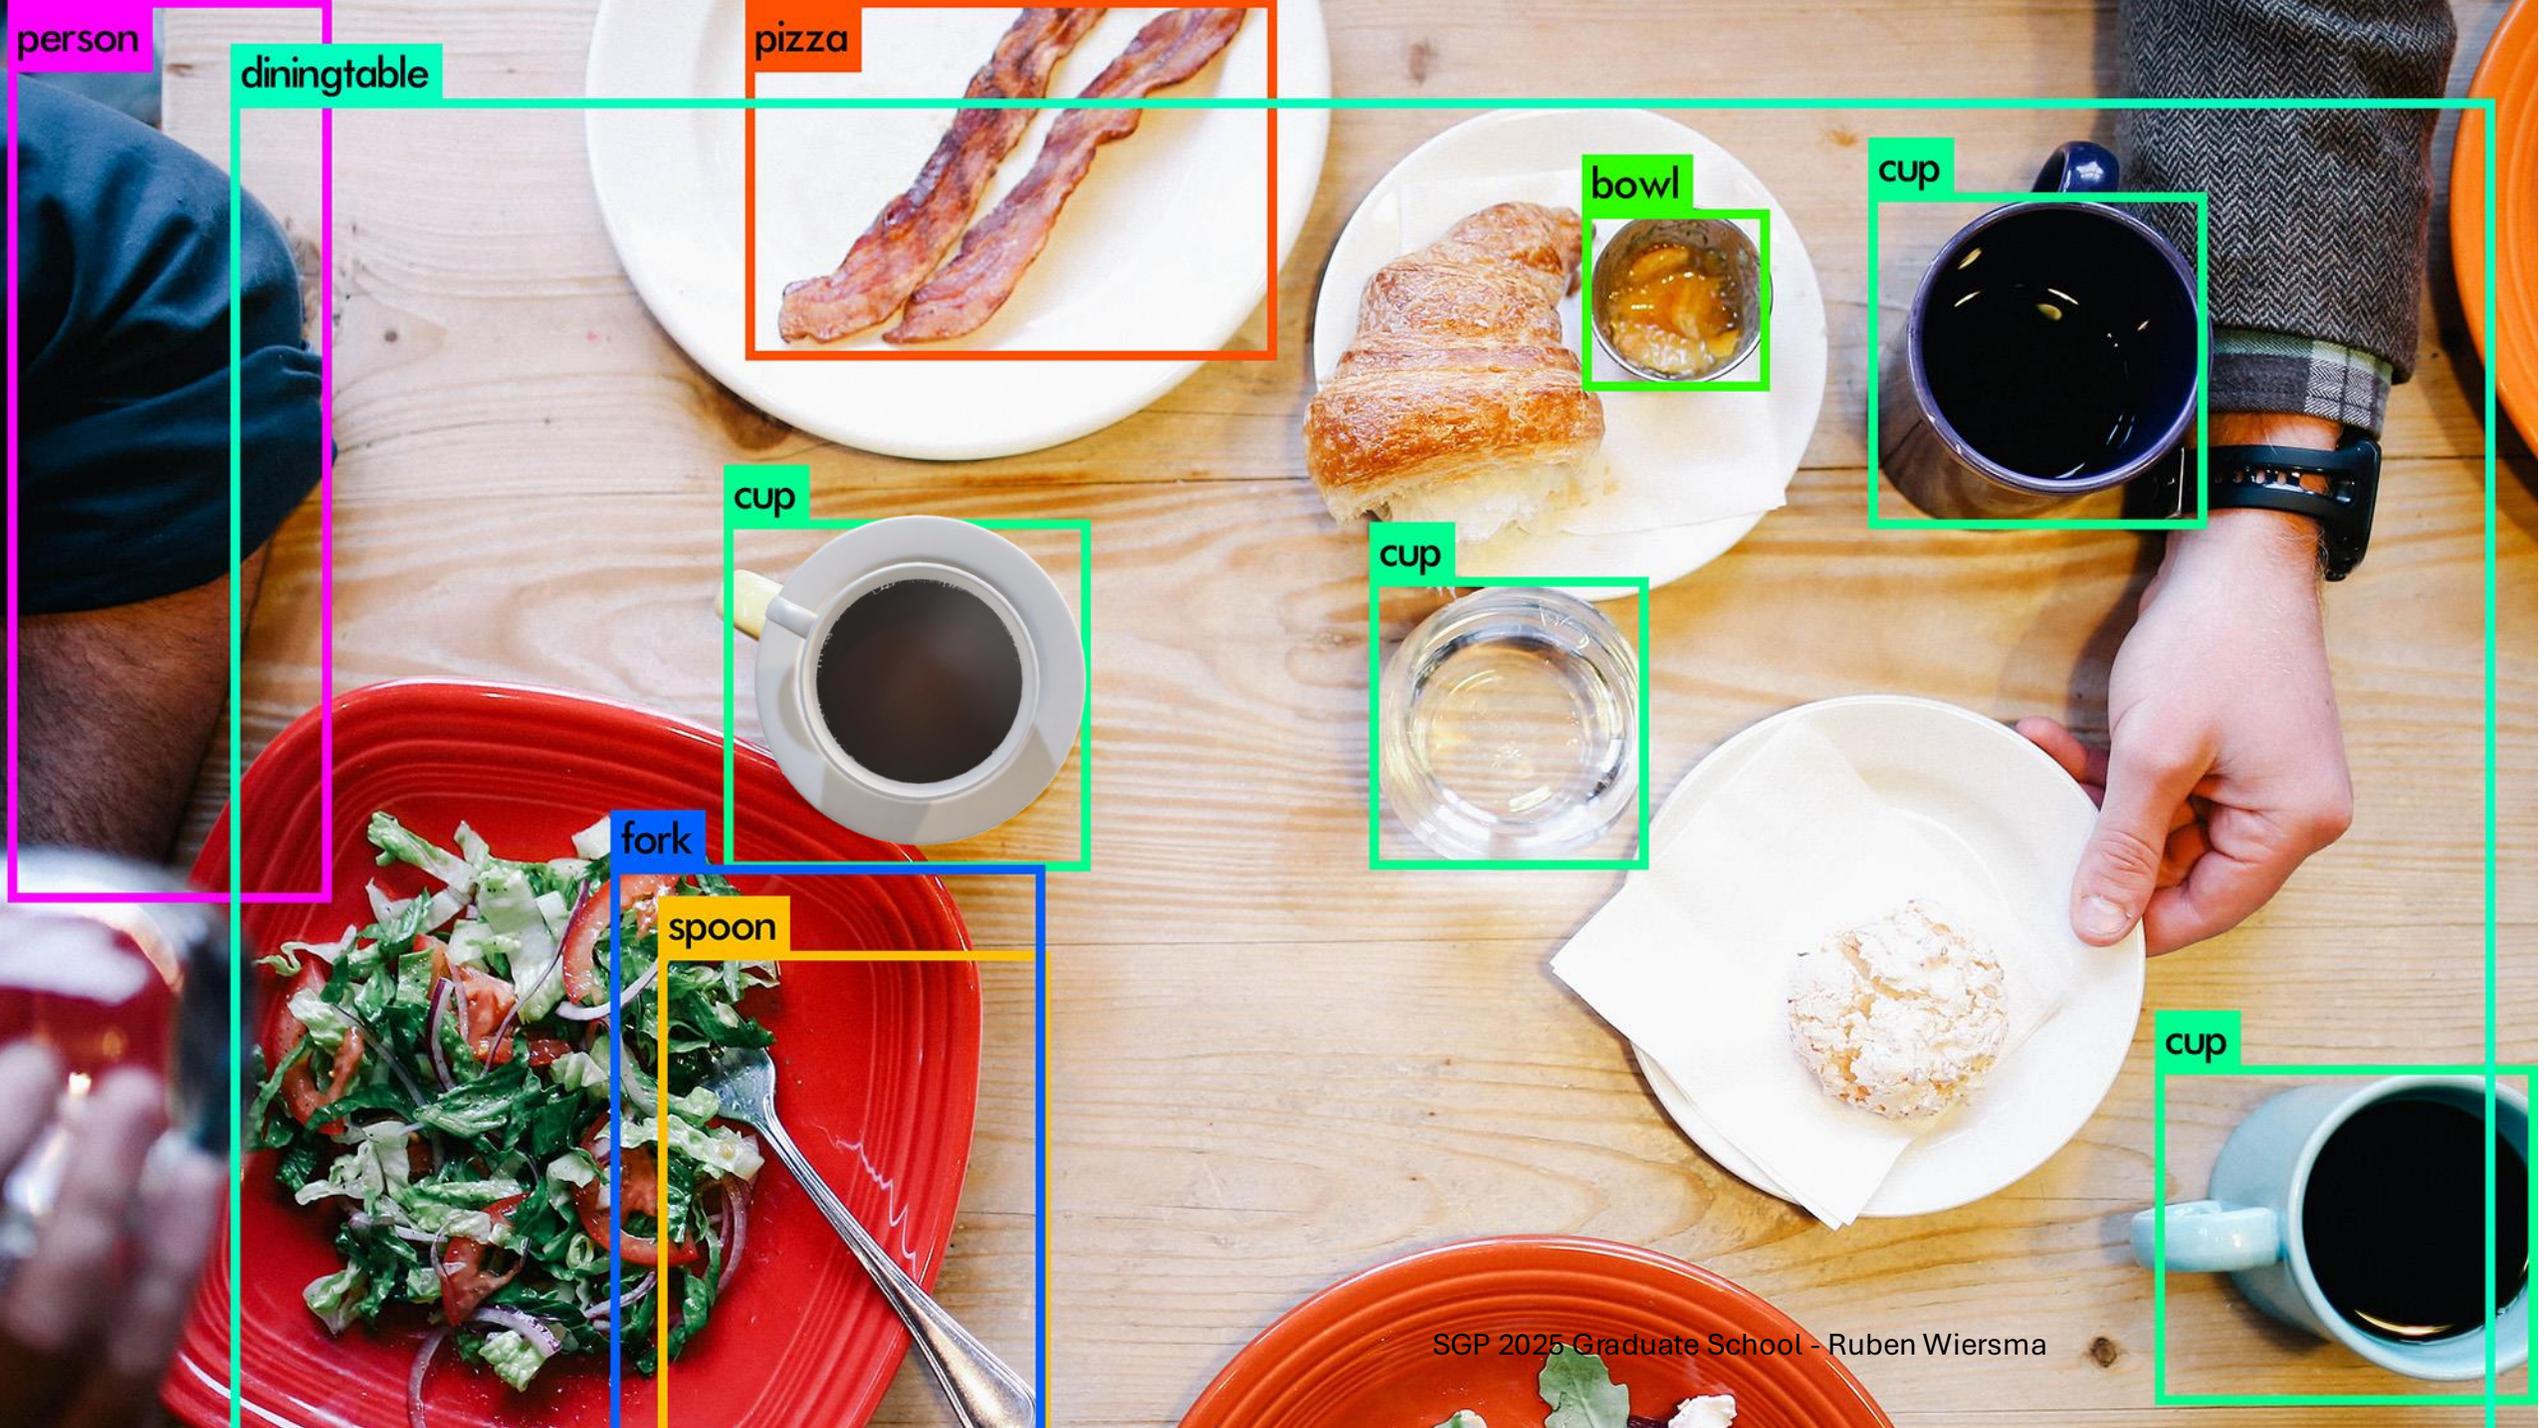
Who are you?

Goal

- Provide a ‘map’ of deep learning on 3D shapes
- Outcome
 - Applying deep learning to 3D tasks
 - Developing deep learning techniques for 3D tasks
- Audience
 - Some familiarity with optimization or machine learning

Why are we interested?





person

diningtable

pizza

bowl

cup

cup

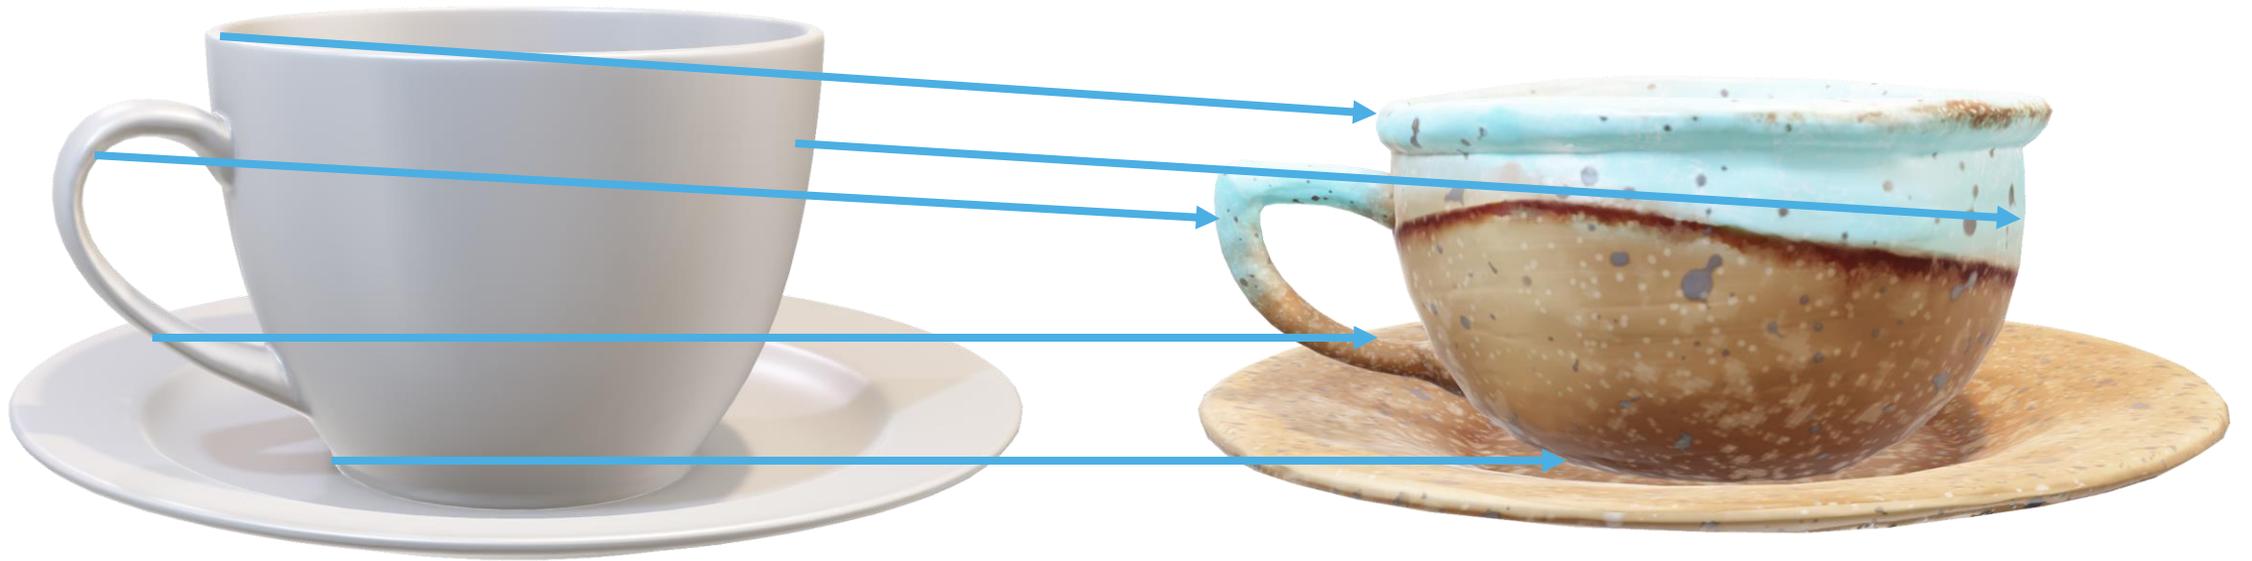
cup

fork

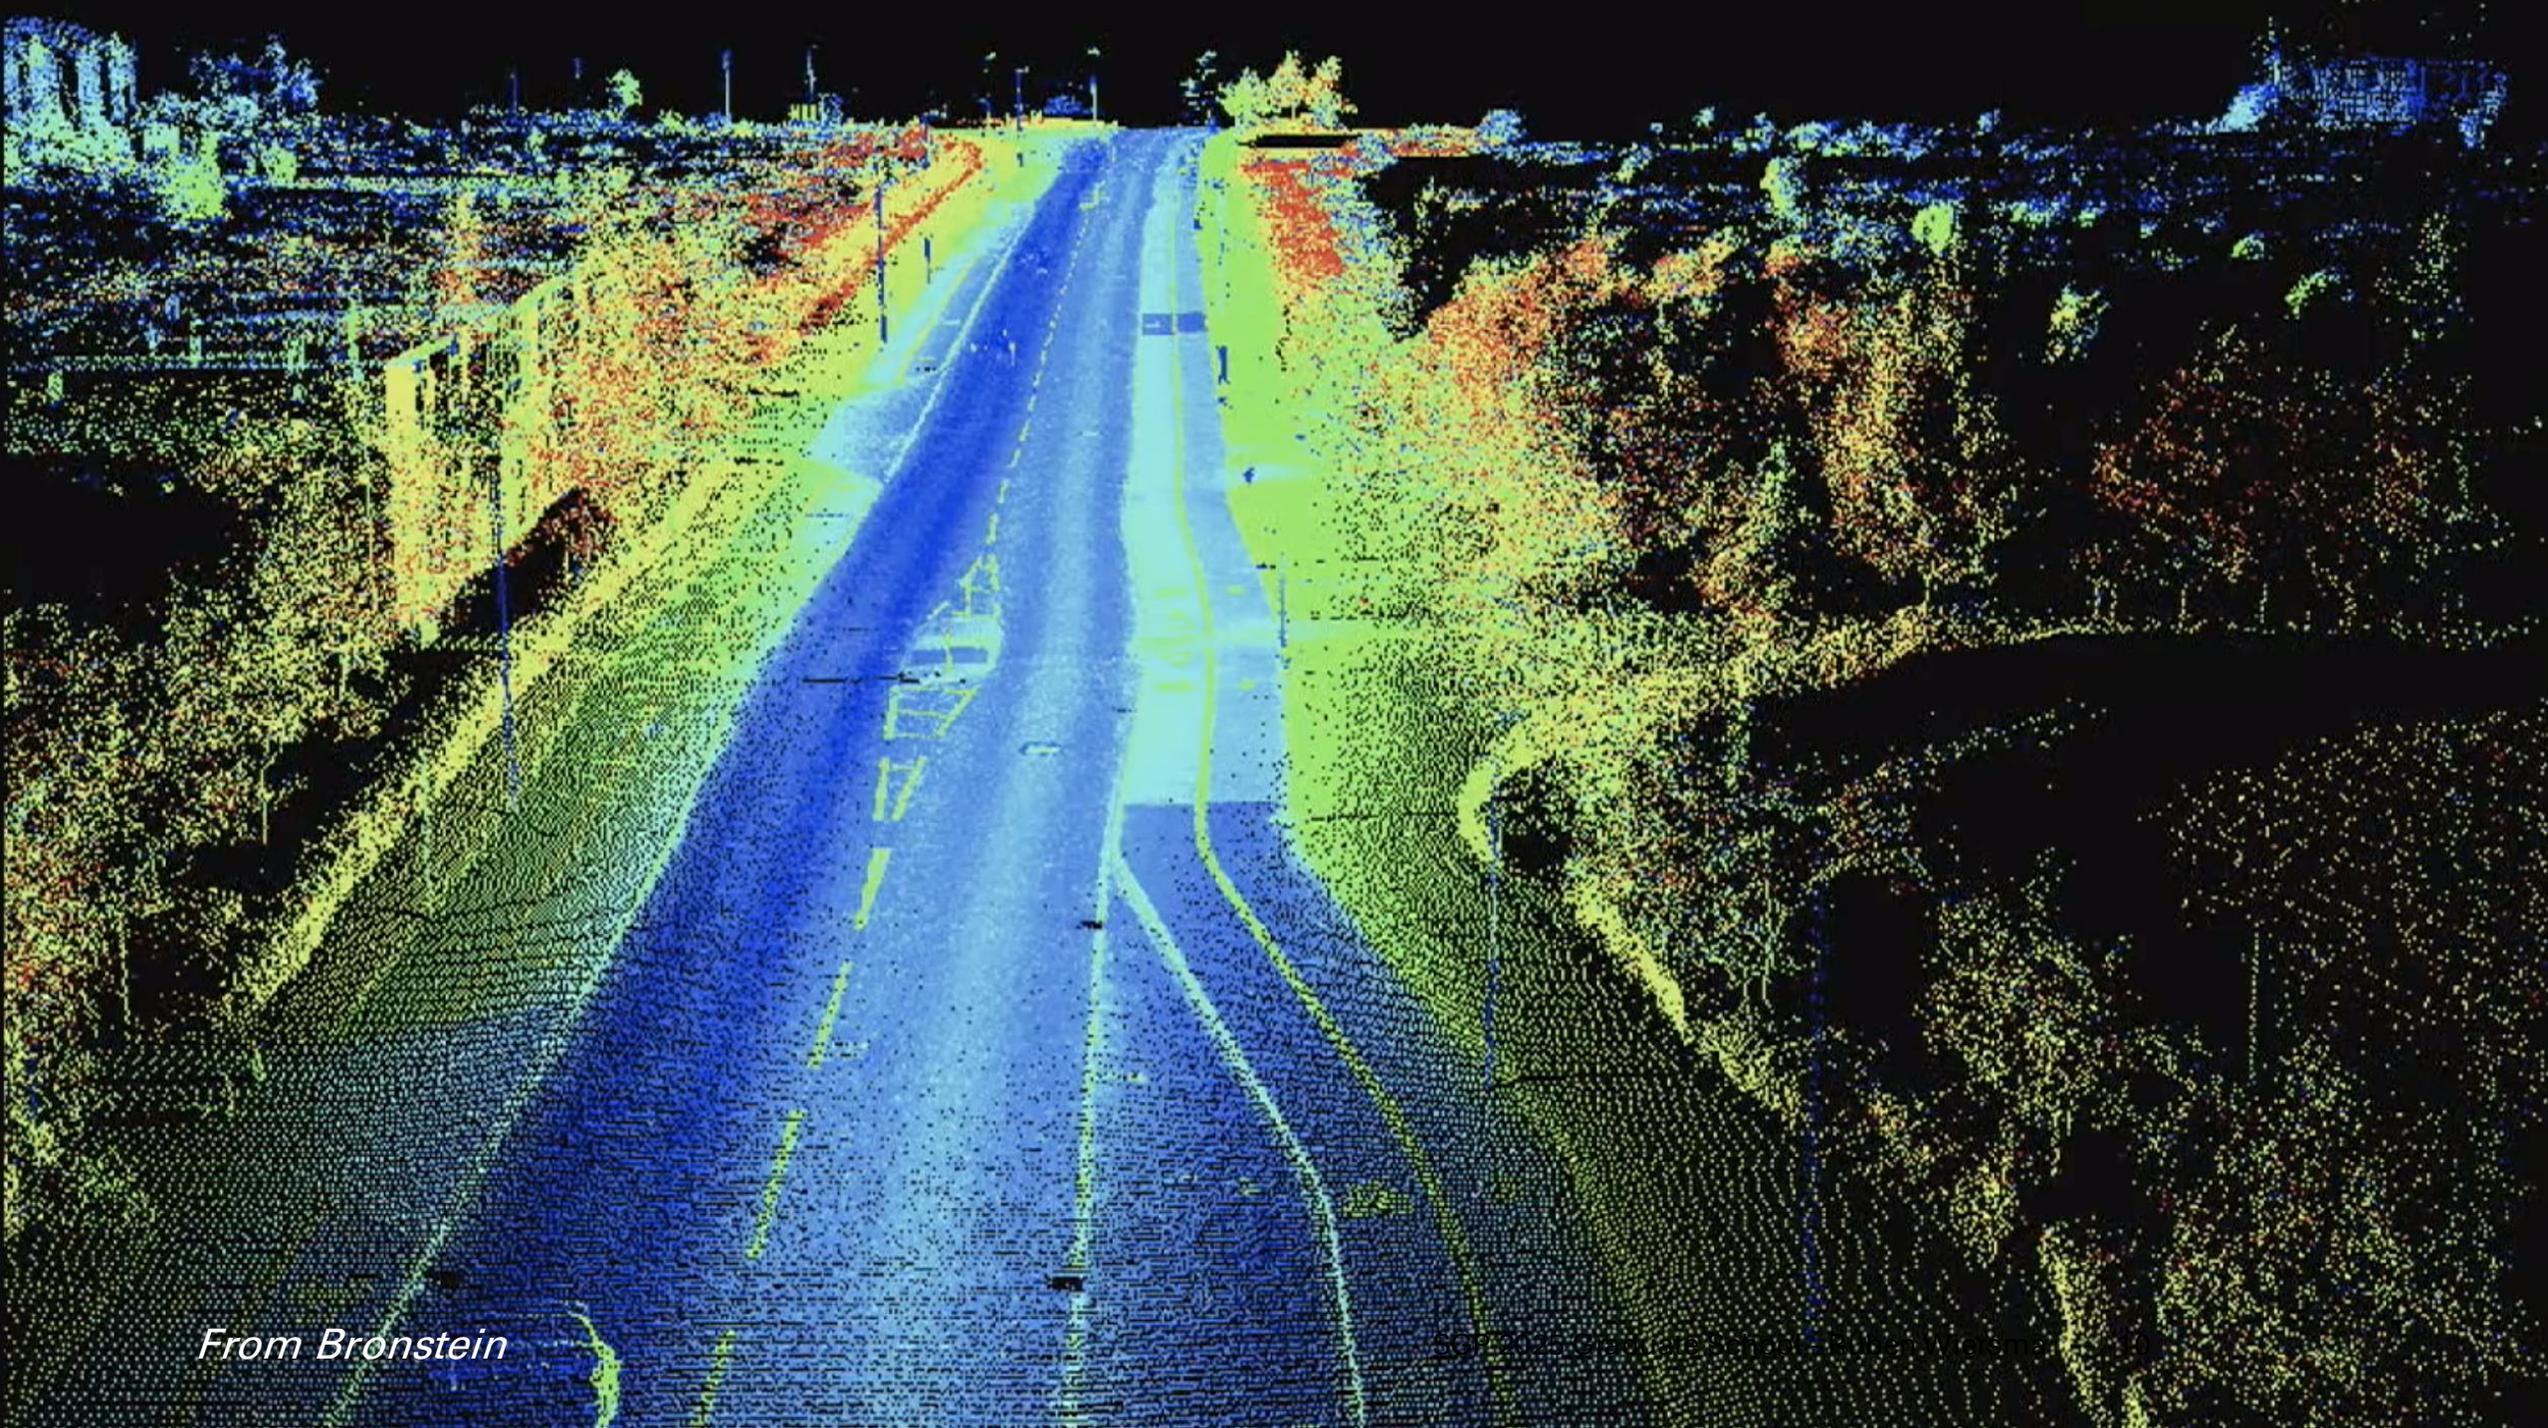
spoon

cup



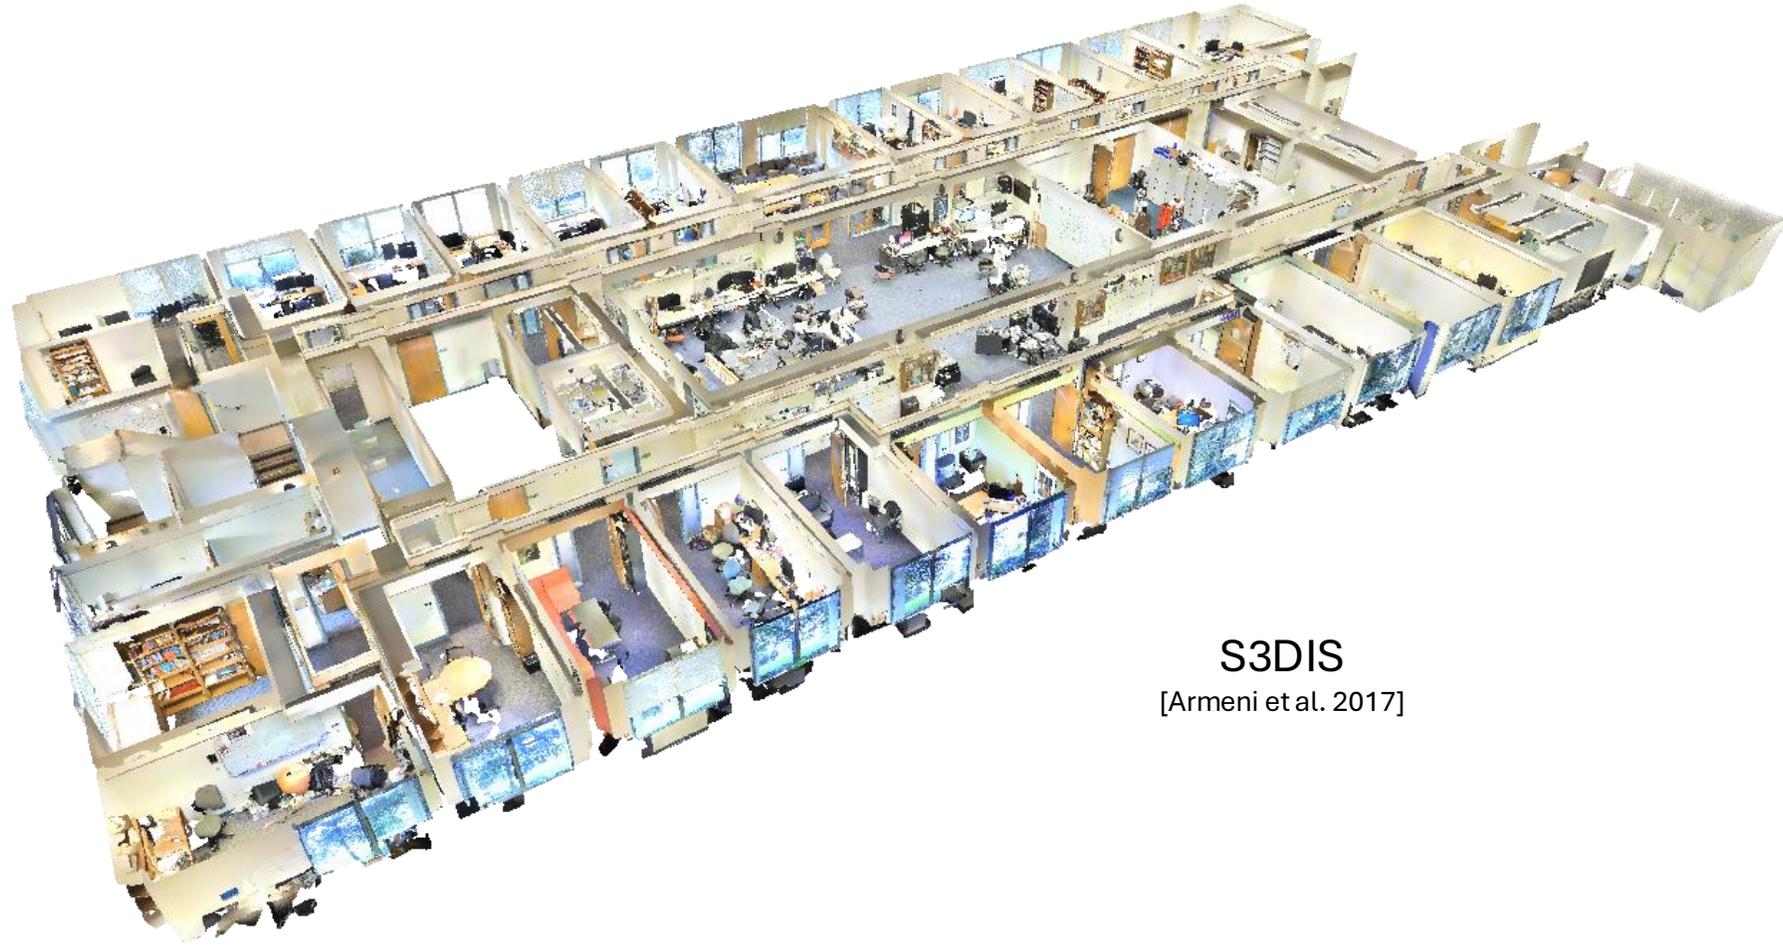




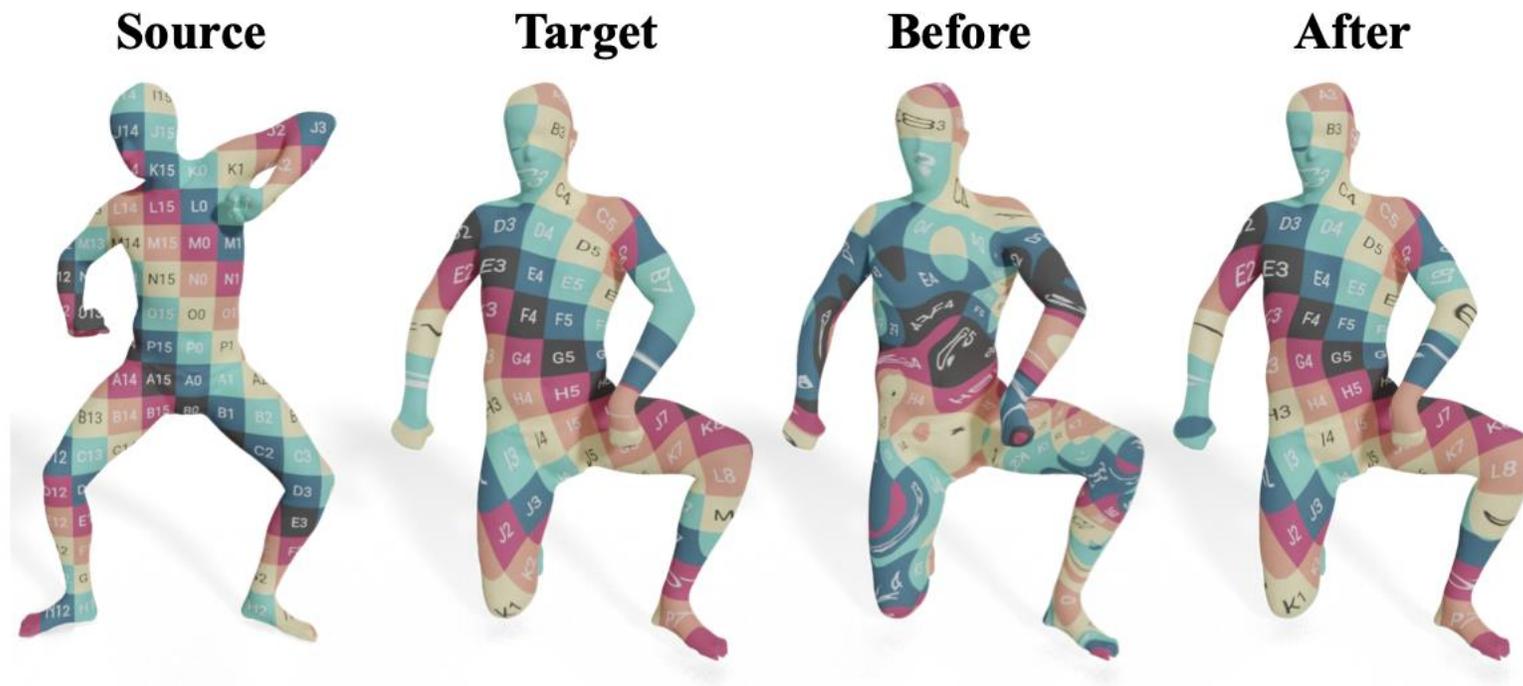


From Bronstein

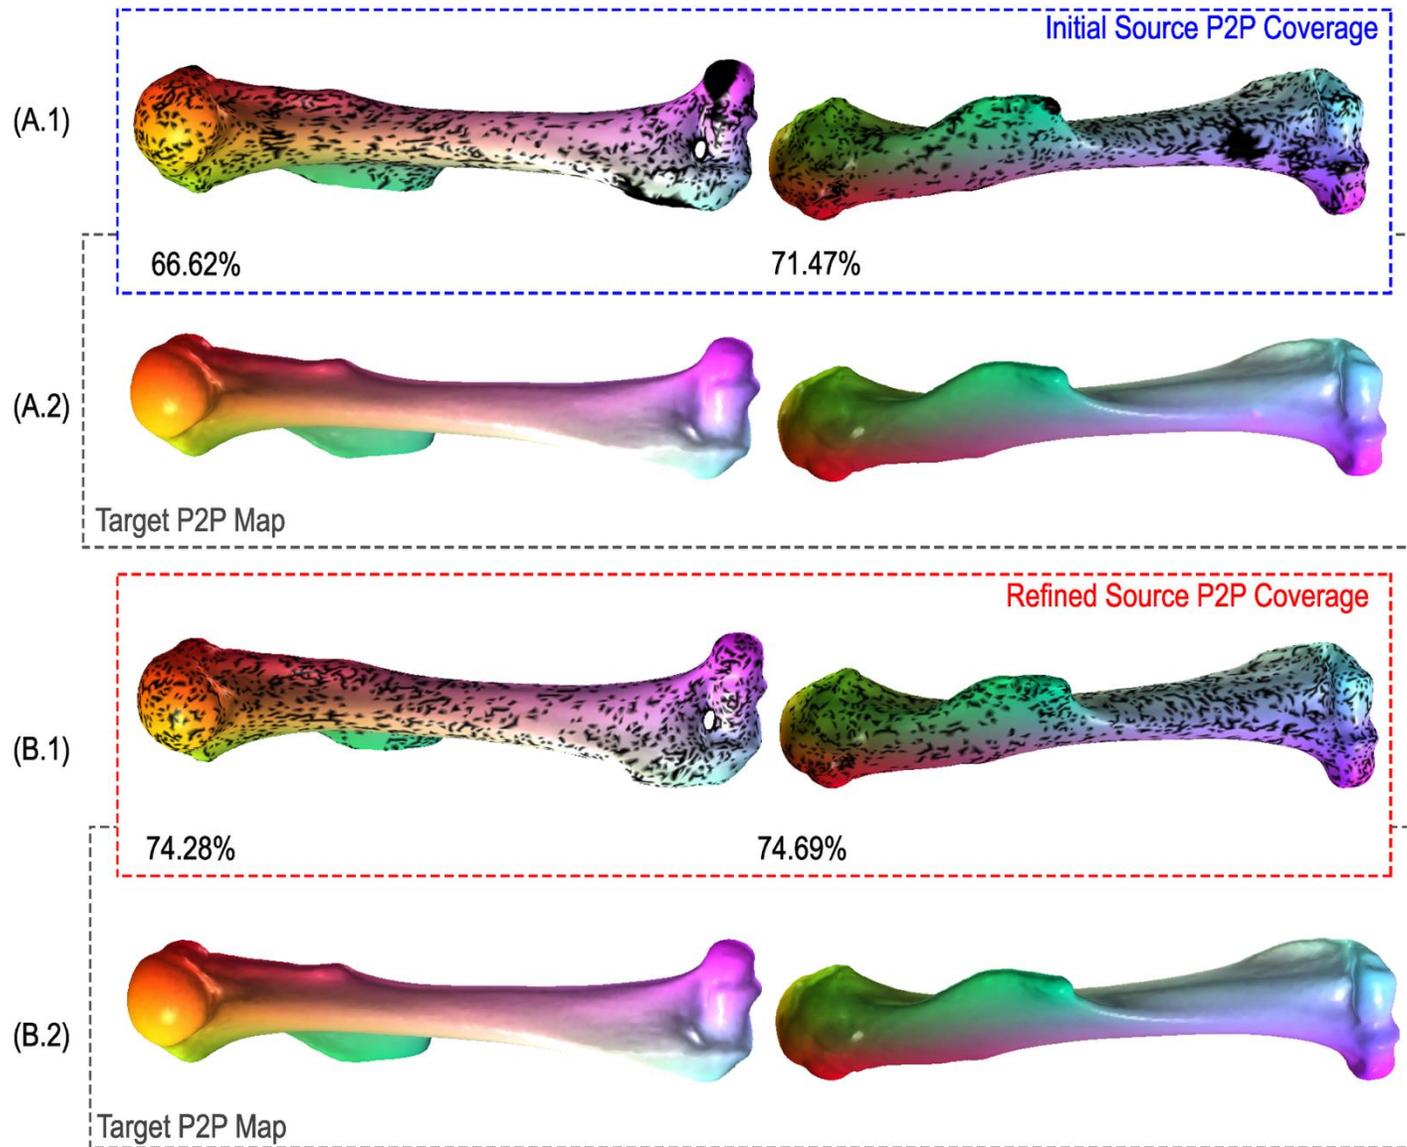
© 2011 Autodesk, Inc.



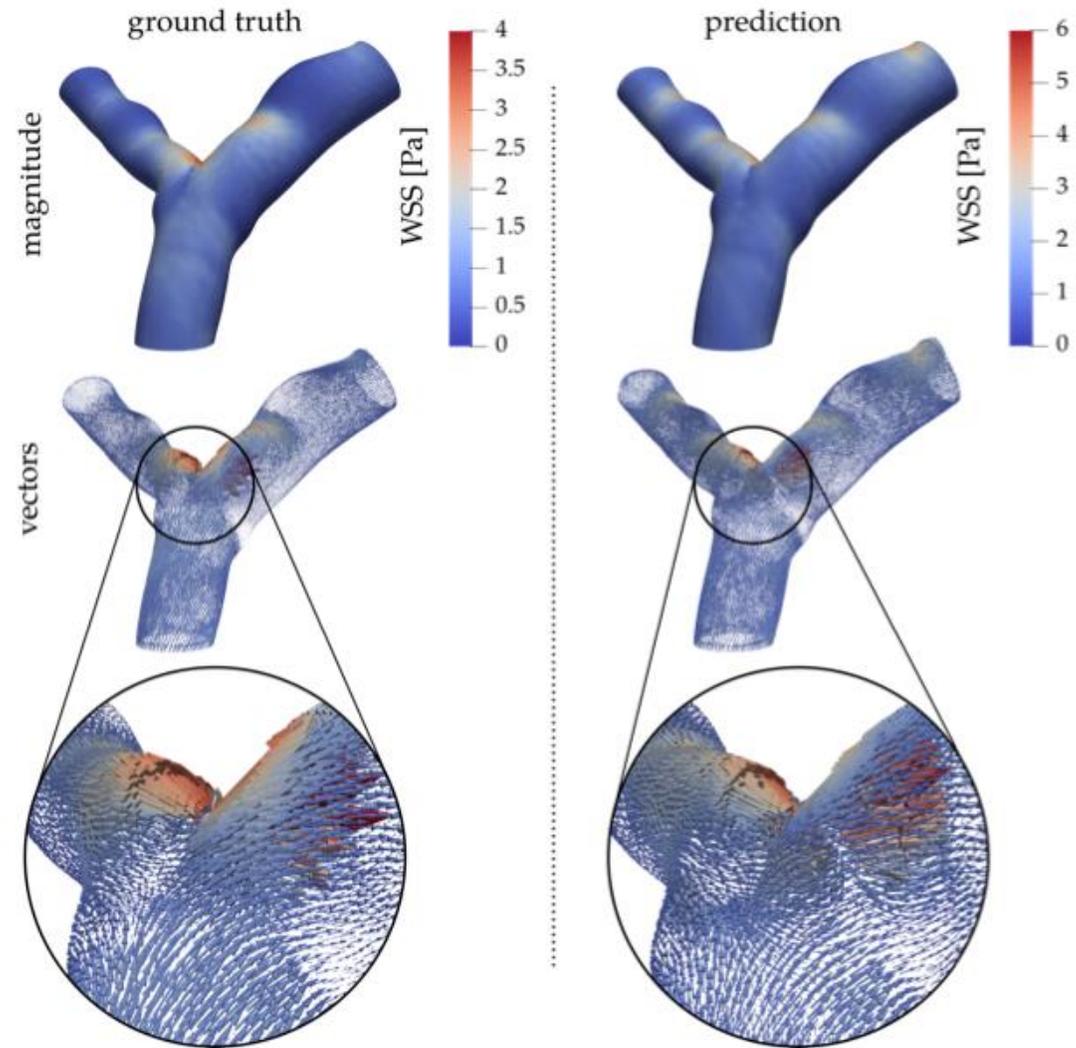
S3DIS
[Armeni et al. 2017]



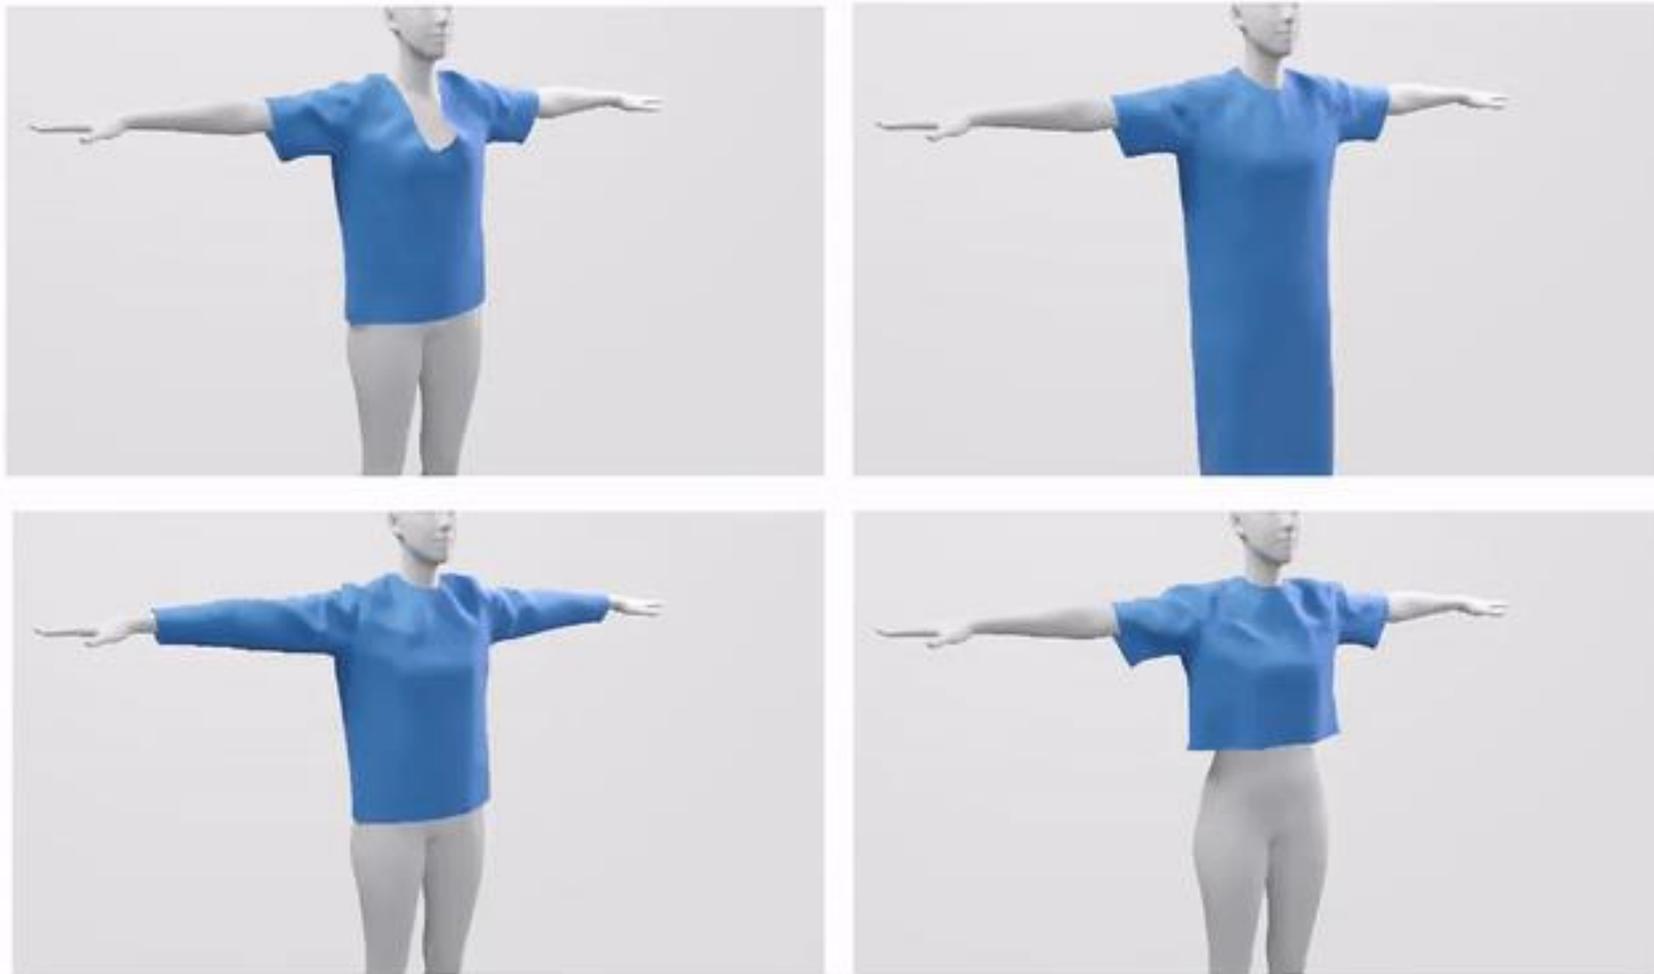
Understanding and Improving Features Learned in Deep Functional Maps, Attaiki and Ovsjanikov (2023)



Automated morphological phenotyping [...], O. Thomas et al. (2023)



Mesh Neural Networks for SE(3)-Equivariant Hemodynamics Estimation of the Artery Wall, Suk et al. (2022)

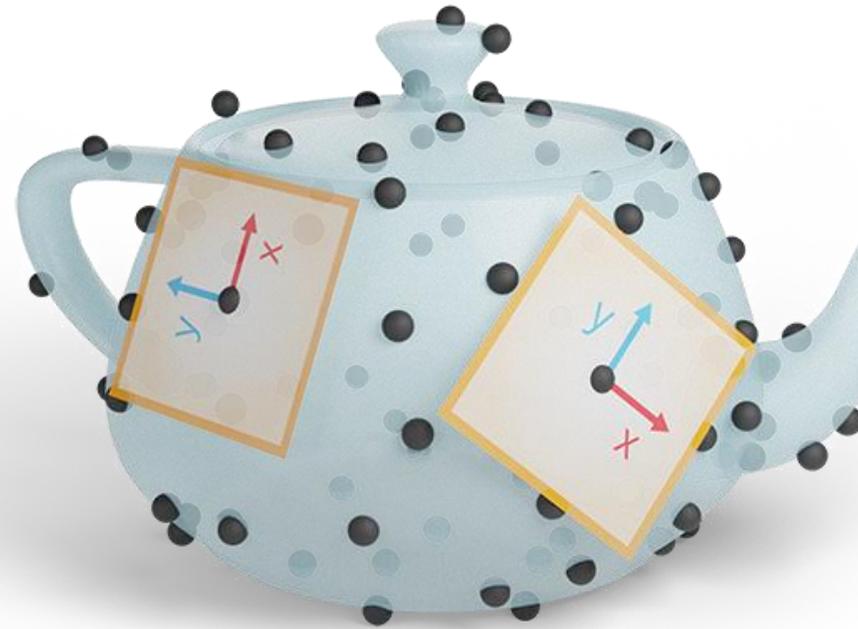
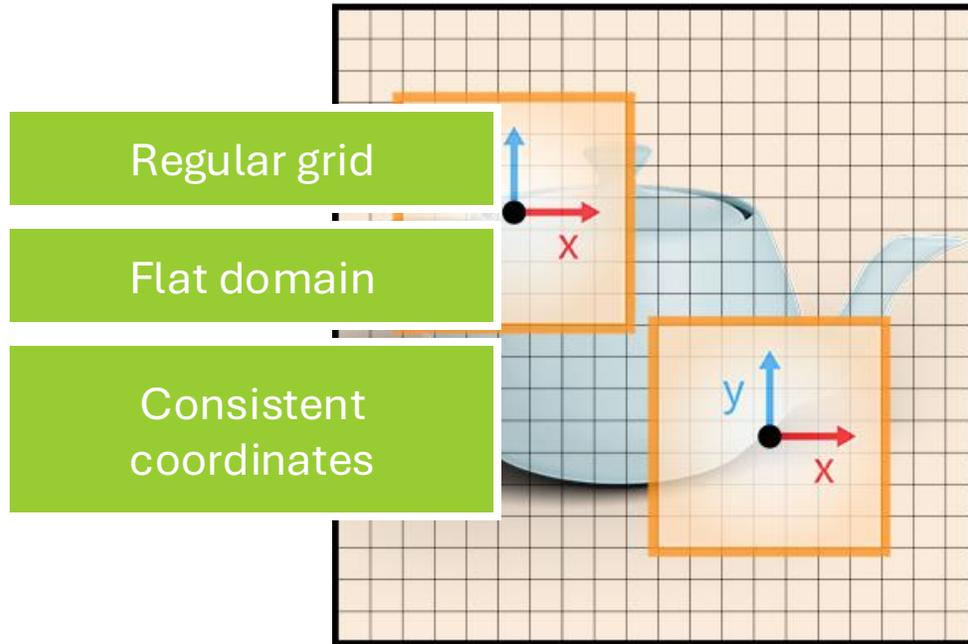


Fully Convolutional Graph Neural Networks for Parametric Virtual Try-On, Vidaurre et al. (2020)

Why are we interested?

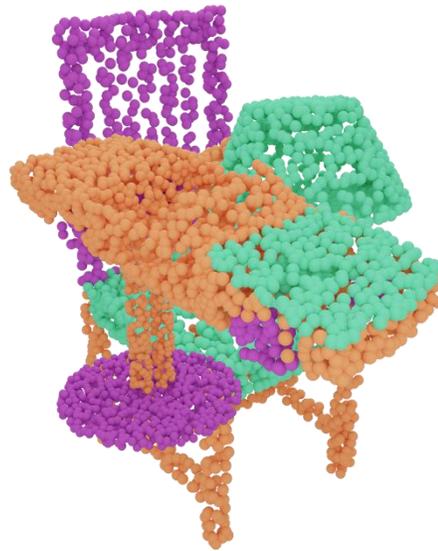
- Classification
- Segmentation
- Registration, correspondence
- Surface reconstruction
- Speeding up classical geometry tasks (i.e., smart lookup table)
- Generative modeling

What's the difference?



Chapter 1: Basics

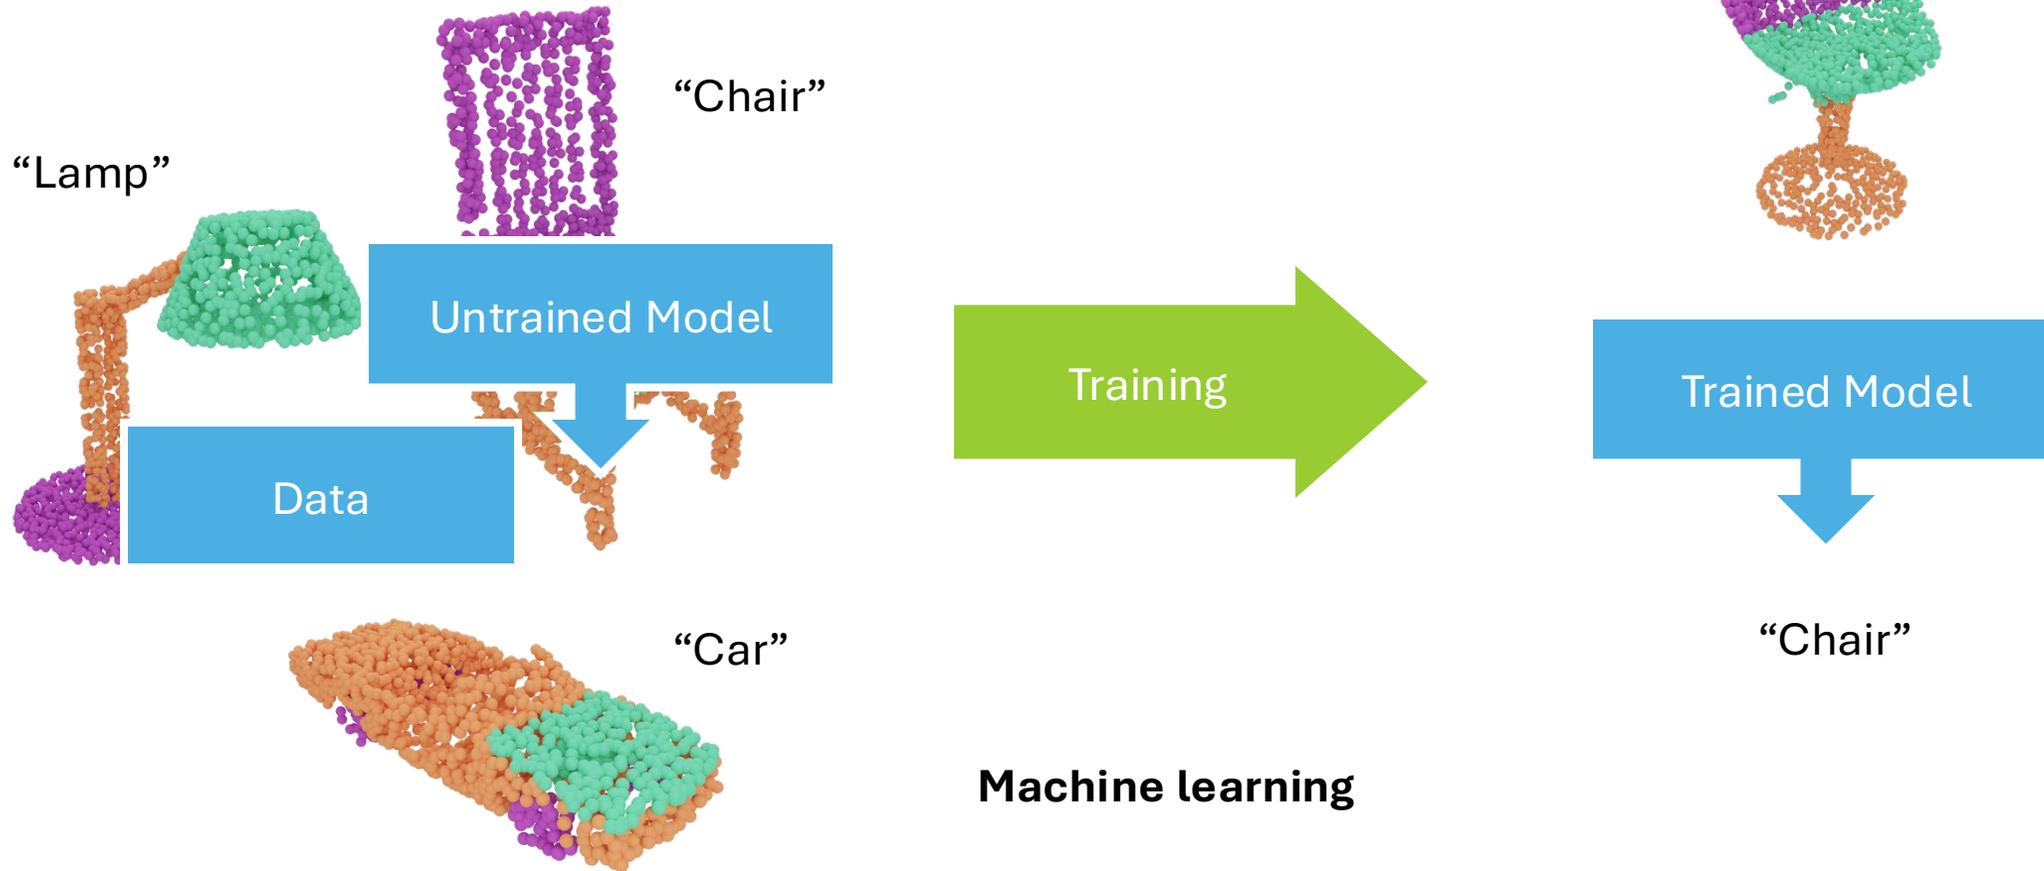
Deep learning is a machine learning method



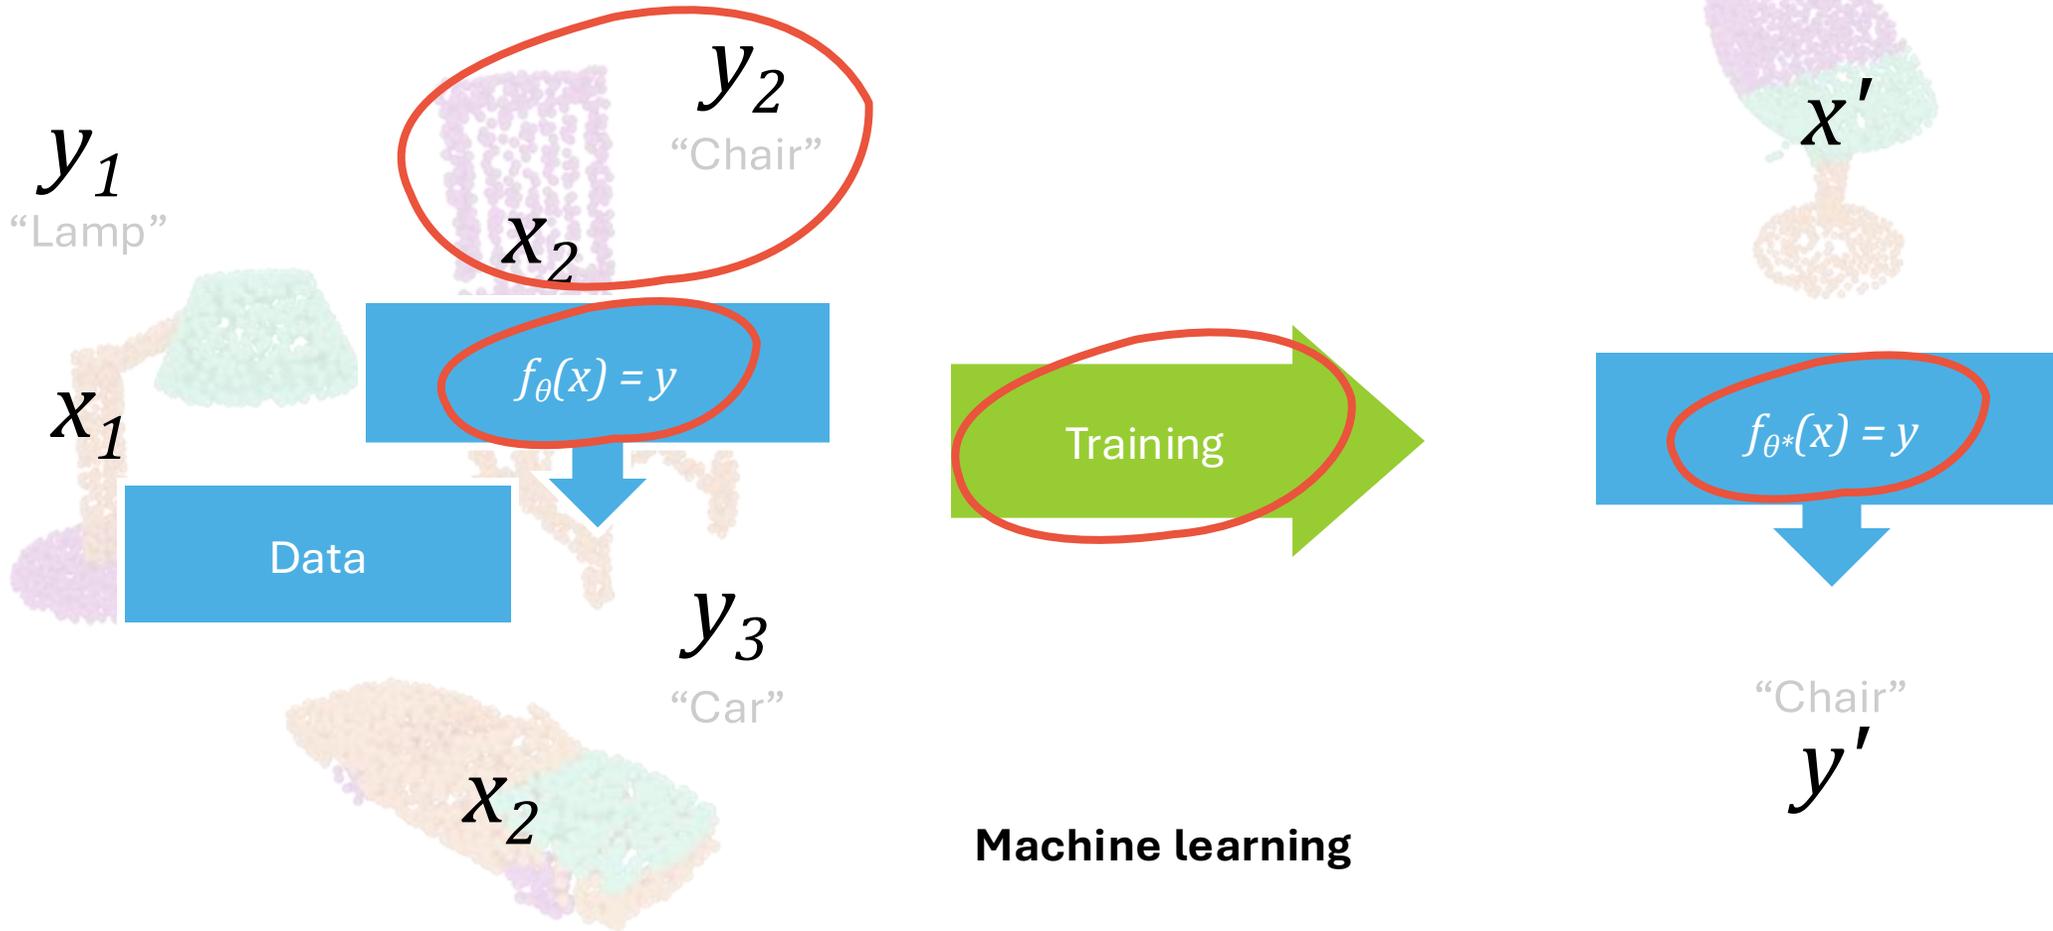
“~~Classic~~”

“Classic” approaches

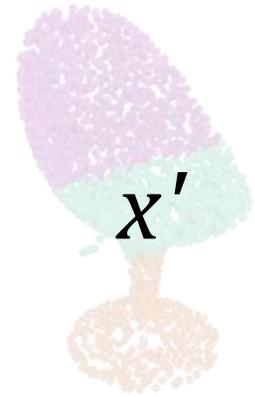
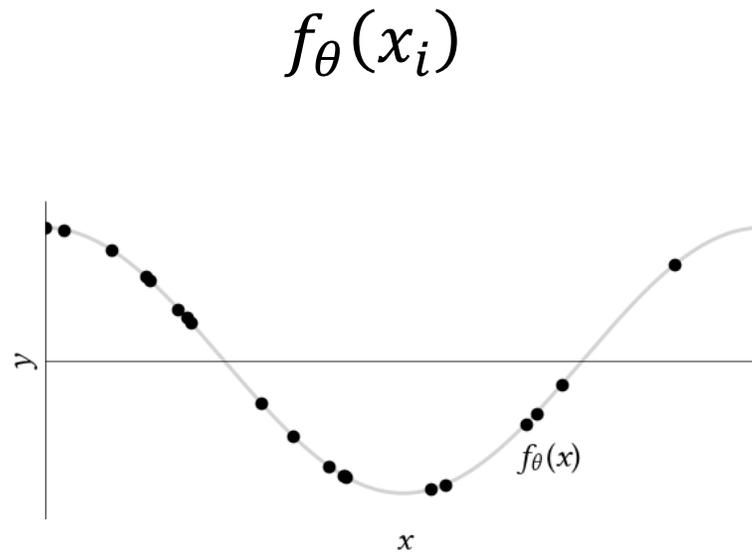
Deep learning is a machine learning method



Deep learning is a machine learning method



Machine learning basics



$f_{\theta^*}(x) = y$

“Chair”
 y'

Deep learning basics

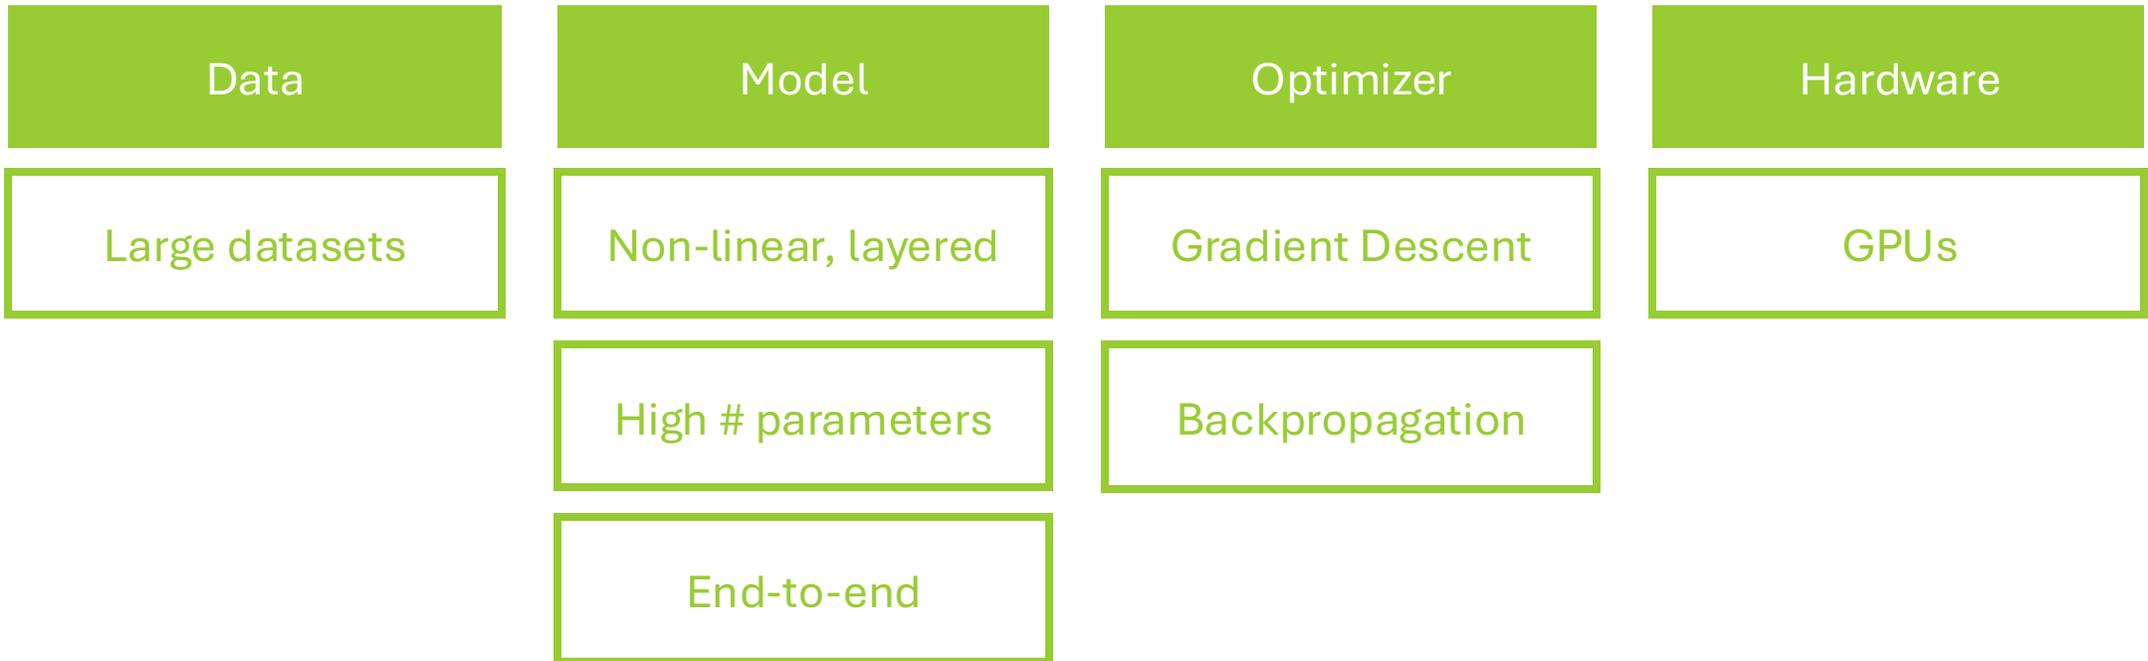
$$\theta^* = \underset{\text{Training}}{\operatorname{argmin}}_{\theta} \sum_i^N \underbrace{\|f_{\theta}(x_i) - y_i\|_2^2}_{\text{Loss}}$$

Model

θ are the unknowns, not x

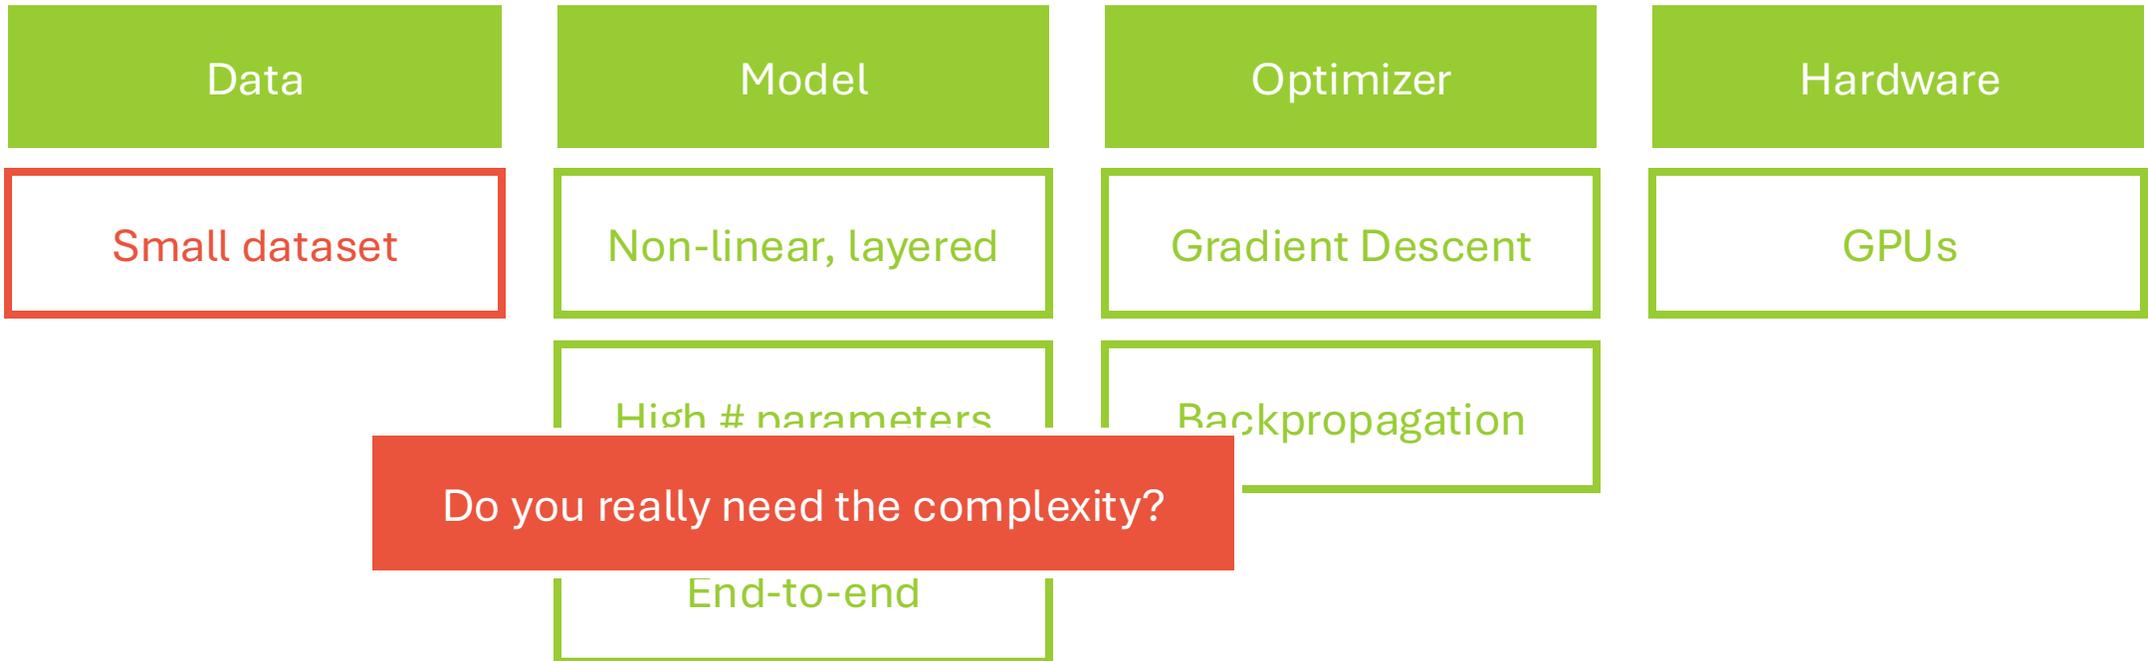
- $f_{\theta}(x)$ is non-linear \rightarrow numerical optimization
- θ is high-dimensional \rightarrow gradient descent (instead of higher order methods)
- N is large \rightarrow stochastic gradient descent (only a few x, y pairs at a time)

Characterization of deep learning



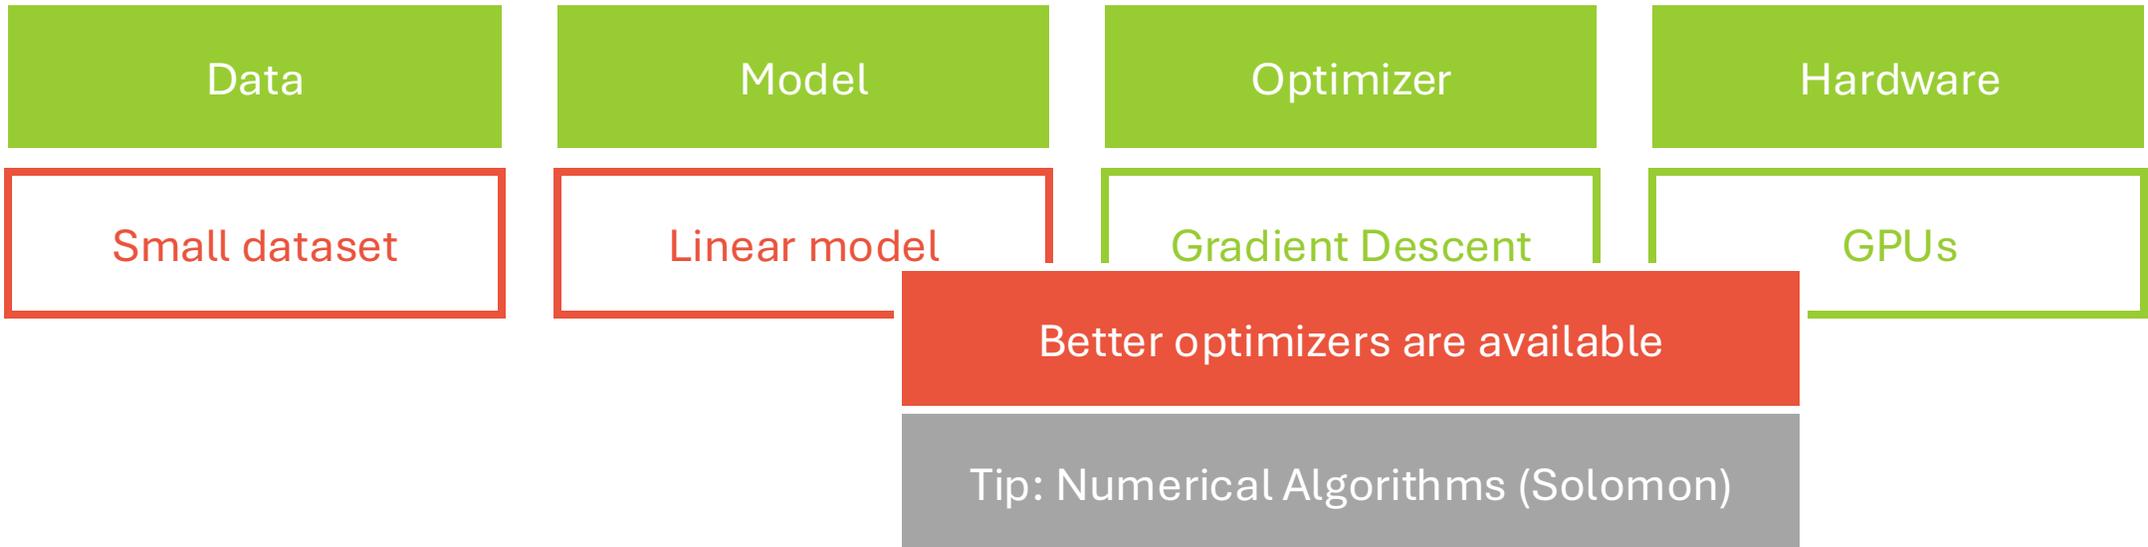
‘Philosophy’: Scaling (data, compute) beats algorithmic complexity

Is deep learning all you need?



‘Philosophy’: Scaling (data, compute) beats algorithmic complexity

Is deep learning all you need?



‘Philosophy’: Scaling (data, compute) beats algorithmic complexity

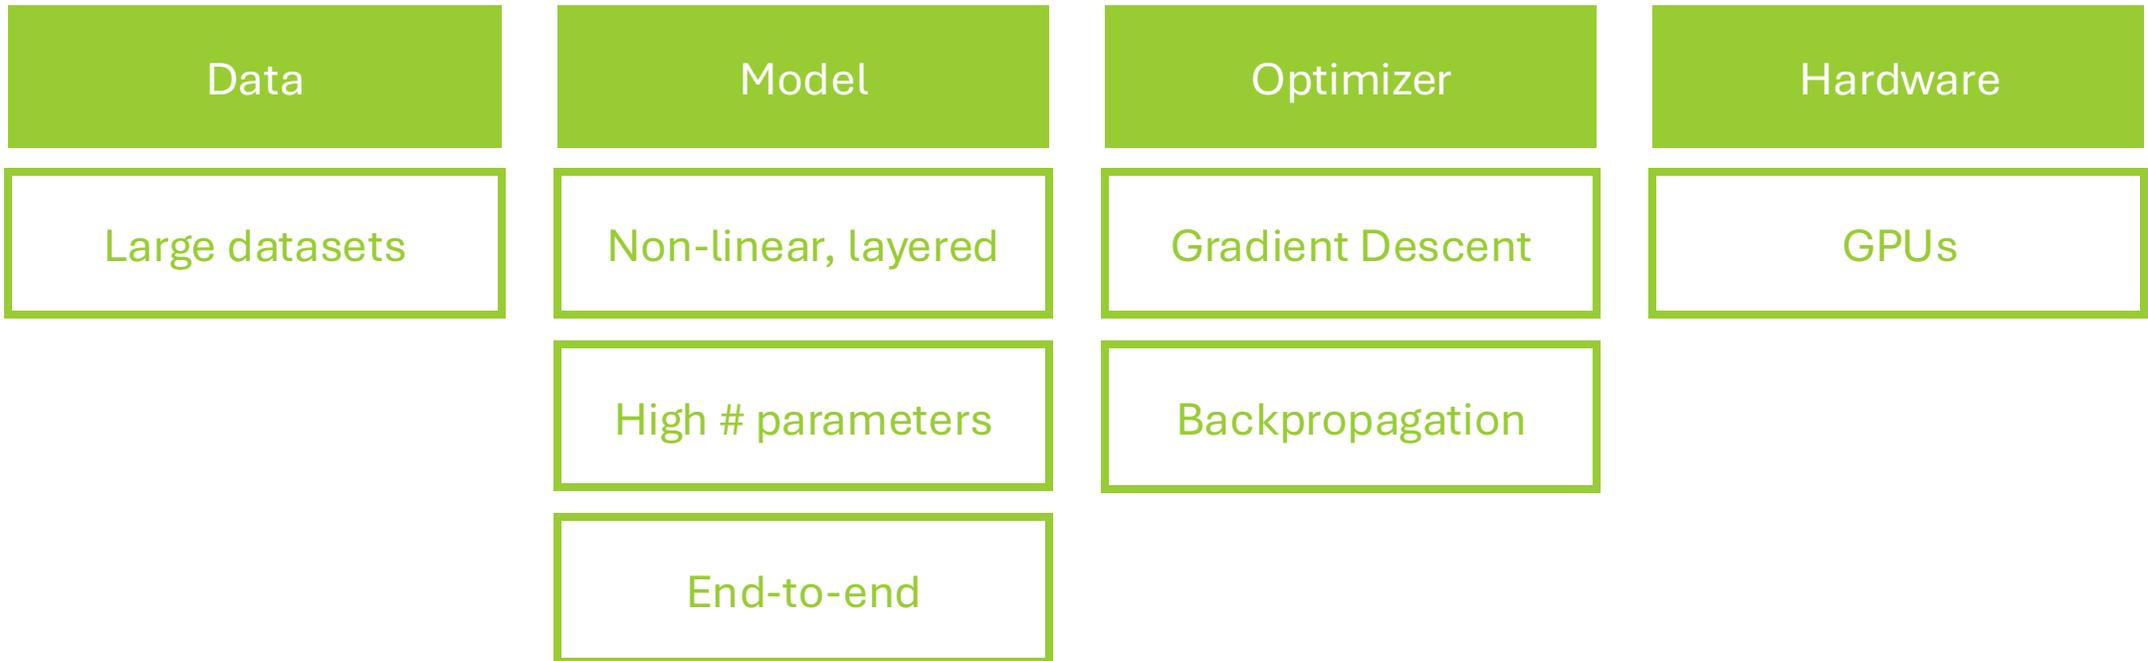
“It’s all about the data” – Alexei Efros

e.g., <https://www.youtube.com/watch?v=M1VHu1d4sGQ>



‘Philosophy’: Scaling (data, compute) beats algorithmic complexity

Characterization of deep learning



‘Philosophy’: Scaling (data, compute) beats algorithmic complexity

Typical models

- MLP – Multi-layer perceptron
 - The Perceptron: A Probabilistic Model For Information Storage And Organization in the Brain – Rosenblatt (**1958**)
- CNN – Convolutional Neural Network
 - Gradient-based learning applied to document recognition – Lecun et al. (**1998**)
- Transformer (Attention)
 - Attention Is All You Need – Vaswani et al. (**2017**), attention was around before that

Multi-Layer Perceptron (geometric perspective)

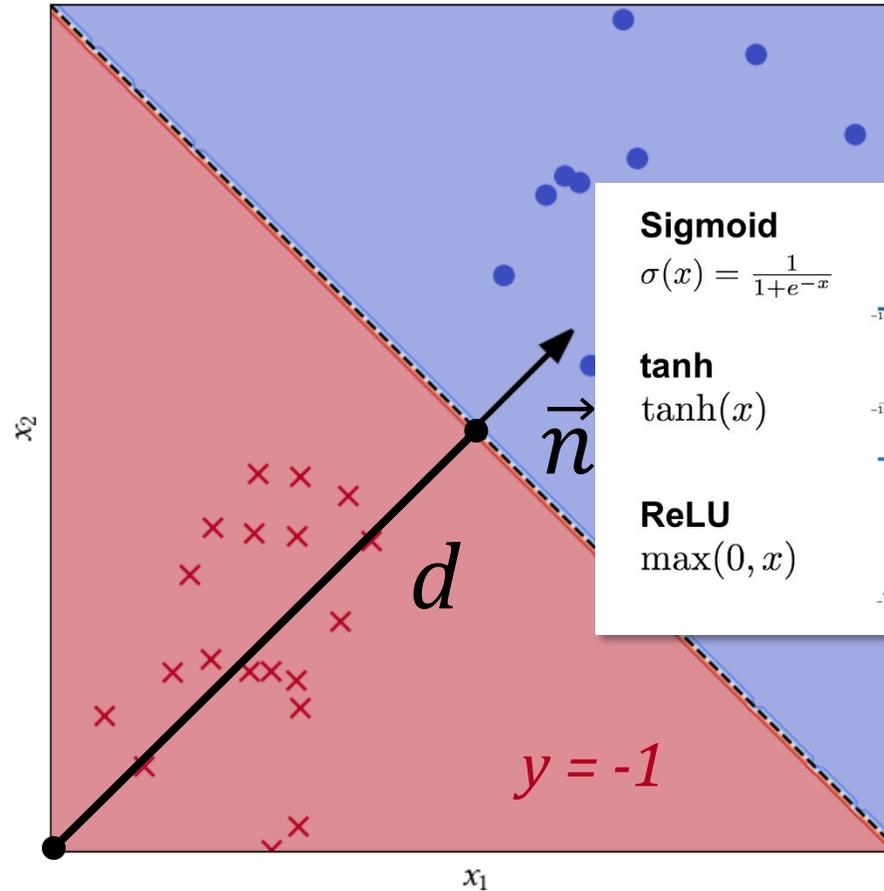
→ Separate these points

- Distance to line/plane

$$\vec{x} \cdot \vec{n} - d$$

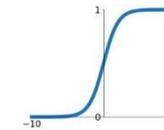
- Step function

$$y = \sigma(\vec{x} \cdot \vec{n} - d)$$



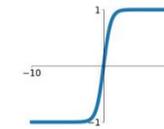
Sigmoid

$$\sigma(x) = \frac{1}{1+e^{-x}}$$



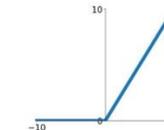
tanh

$$\tanh(x)$$



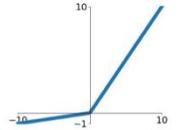
ReLU

$$\max(0, x)$$



Leaky ReLU

$$\max(0.1x, x)$$

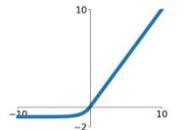


Maxout

$$\max(w_1^T x + b_1, w_2^T x + b_2)$$

ELU

$$\begin{cases} x & x \geq 0 \\ \alpha(e^x - 1) & x < 0 \end{cases}$$



<https://medium.com/@shrutijadon/survey-on-activation-functions-for-deep-learning-9689331ba092>

Multi-Layer Perceptron (geometric perspective)

→ Separate these points

- Distance to line/plane

$$\vec{w}^T \vec{x} + b$$

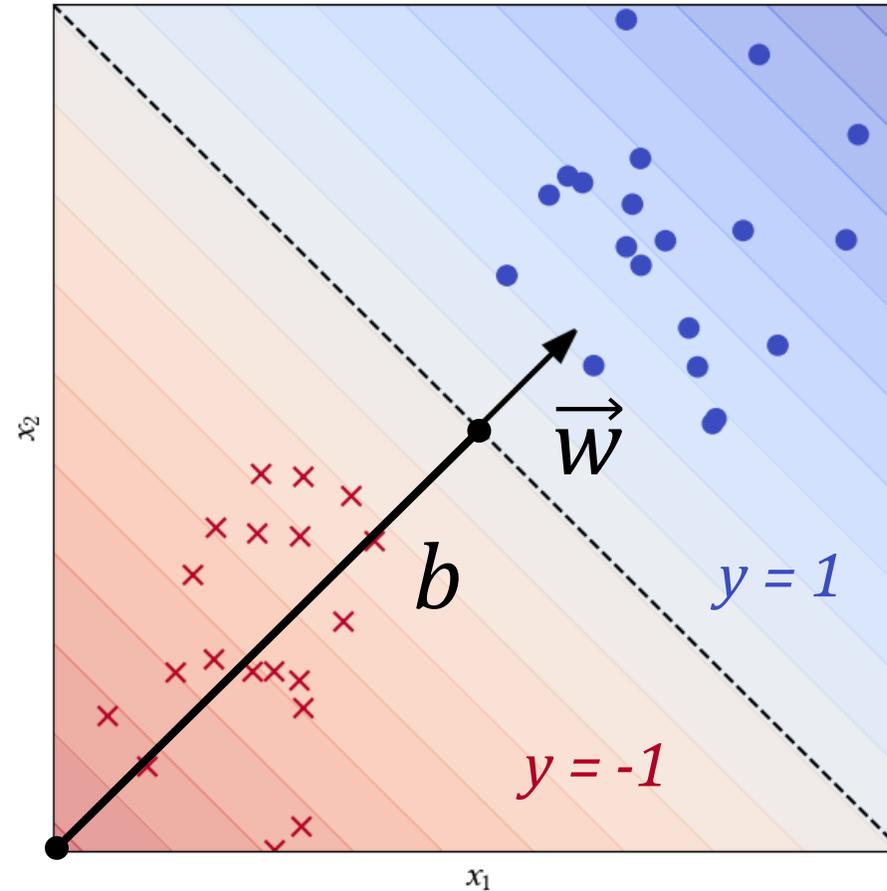
- Step function

$$y = \sigma(\vec{w}^T \vec{x} + b)$$

- Weights, biases

- σ activation function

- Non-linear!



Multi-Layer Perceptron

$$y = \sigma(\vec{w}^T \vec{x} + b)$$

→ Multiple outputs?

“Channels”

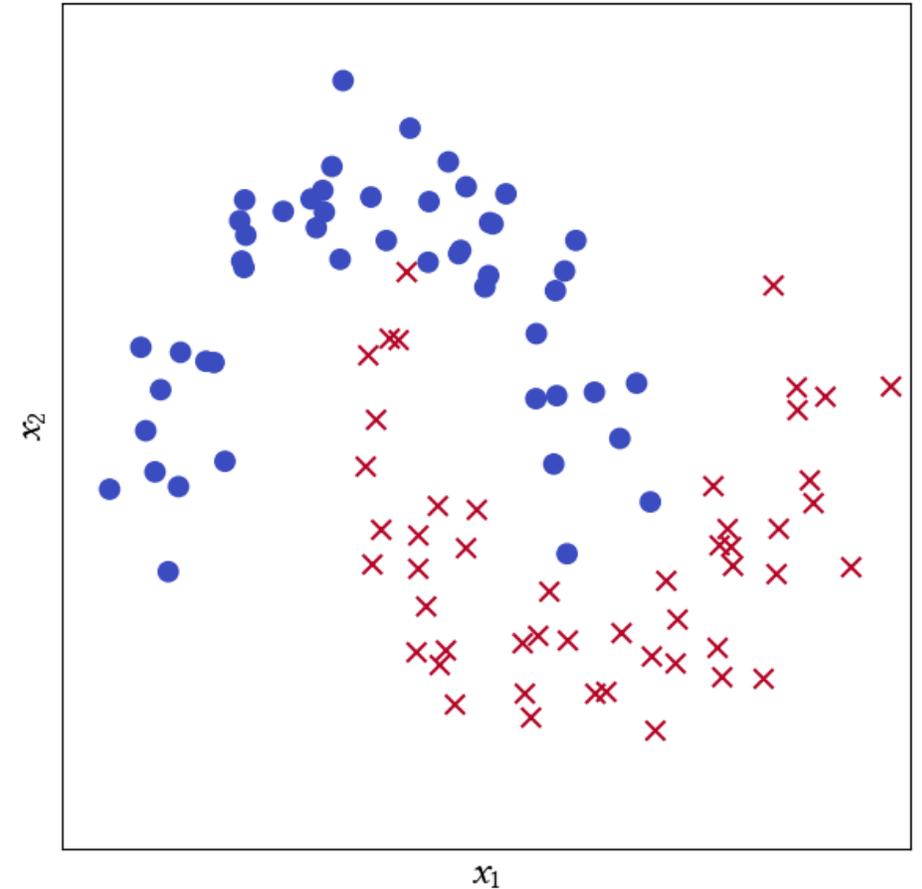
$$\vec{y} = \sigma(W\vec{x} + \vec{b})$$

→ Not linearly separable?

“Layers”

$$\vec{y} = \underbrace{\sigma(W_1\vec{x} + \vec{b}_1)}_{\text{Hidden layers}}$$

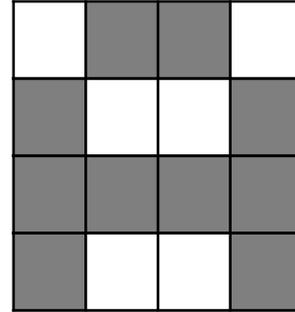
Hidden layers



Multi-Layer Perceptron Notes

- Examples in 2D, but \vec{x} can be of any dimension
- Interpretation: cutting up space with halfspaces
- “Multilayer Feedforward Networks [MLPs] are Universal Approximators”
 - Hornik, Stinchcombe and White (1989)
 - Non-linearity is necessary!
 - Compare, e.g., Fourier transform
- Higher complexity, more risk of overfitting

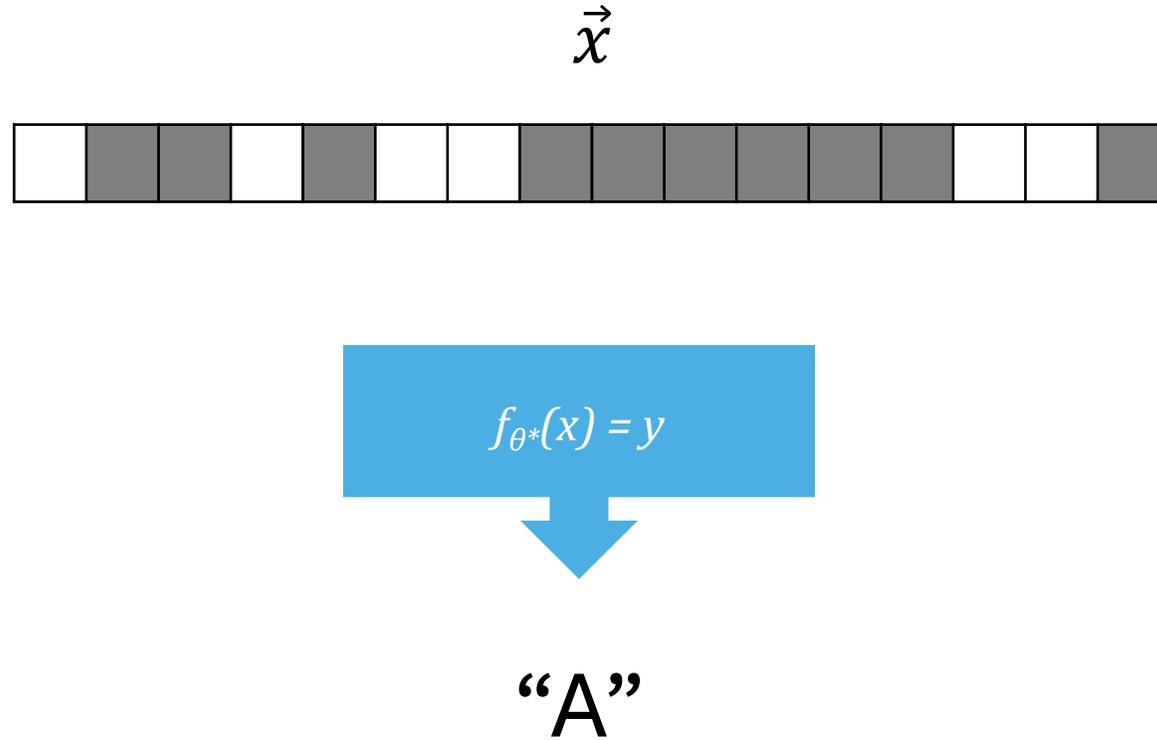
MLP on images



$$f_{\theta^*}(x) = y$$

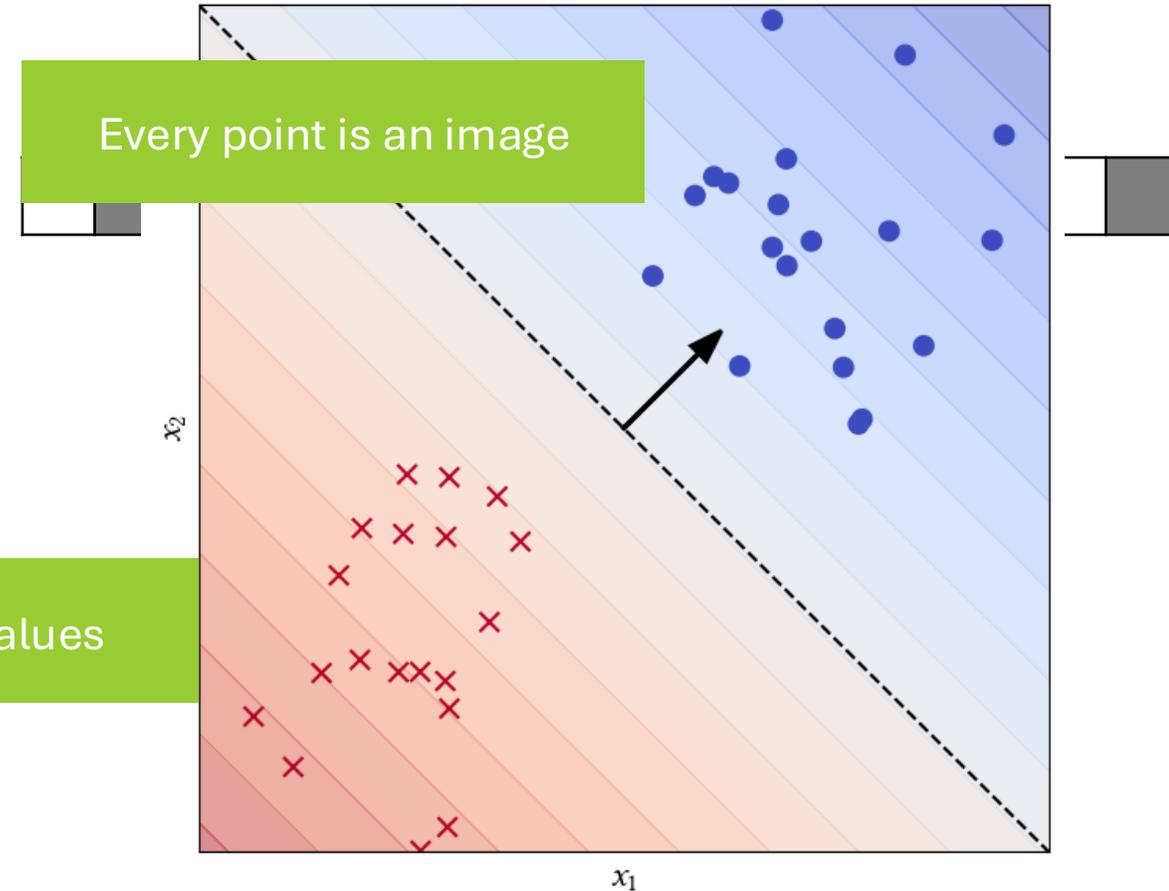
“A”

MLP on images

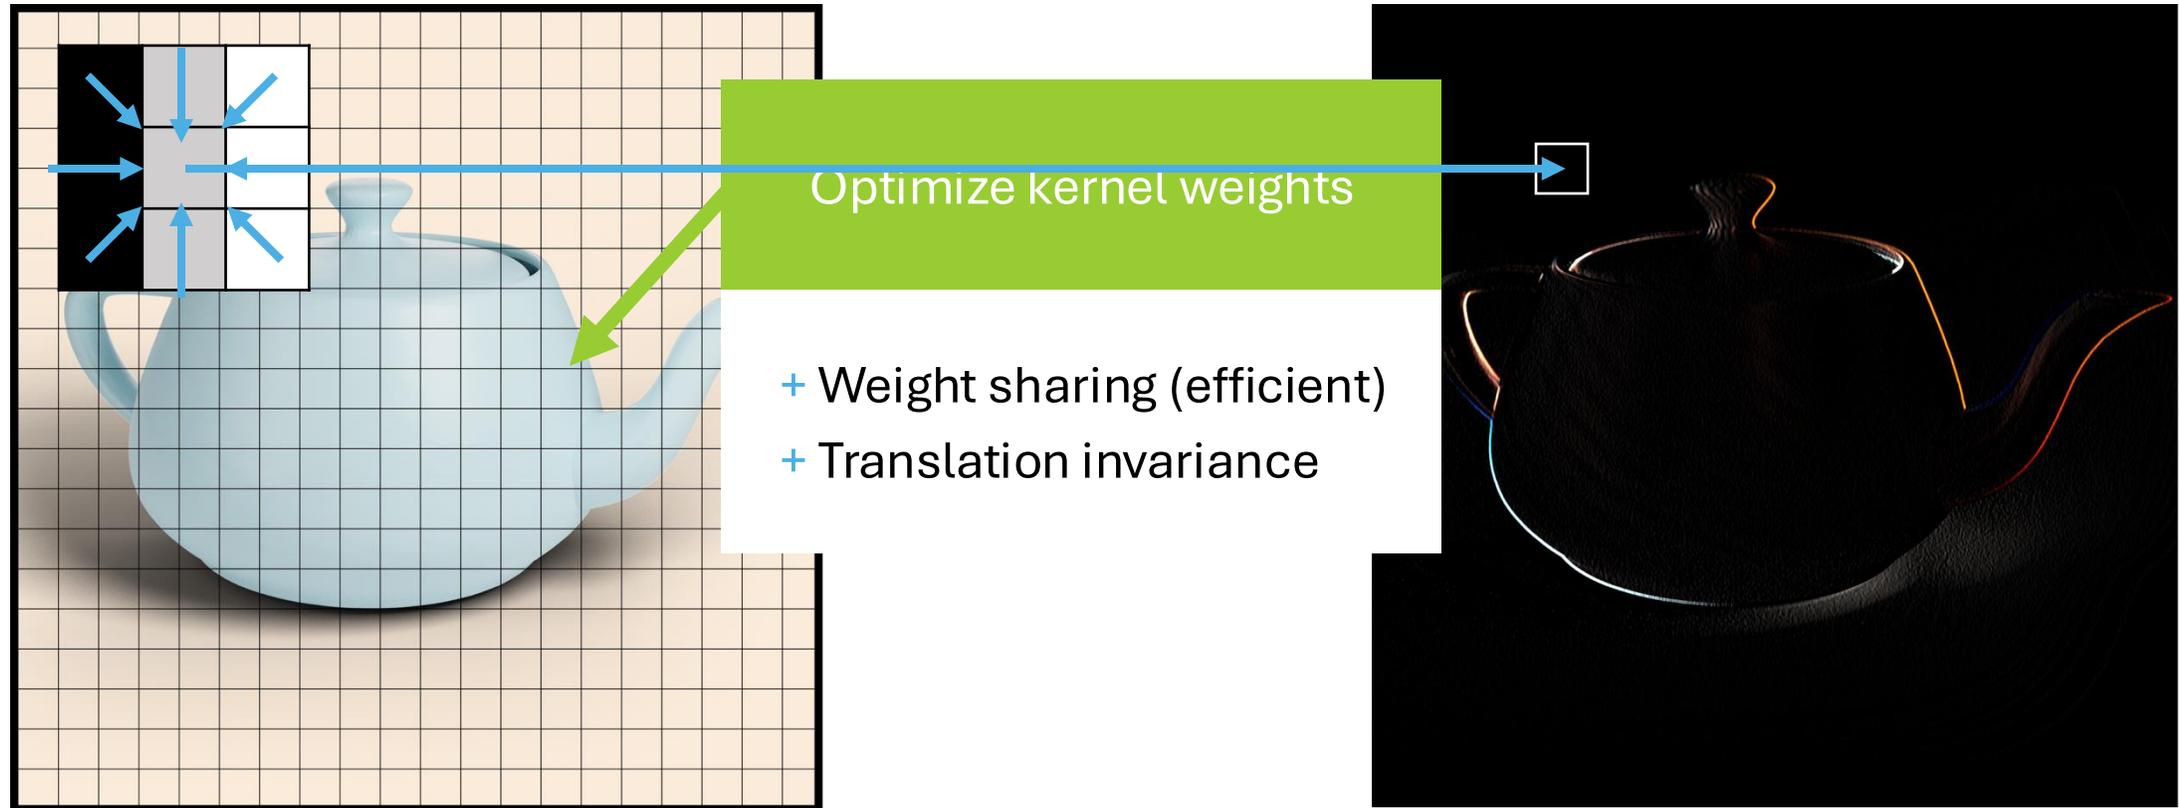


MLP on images

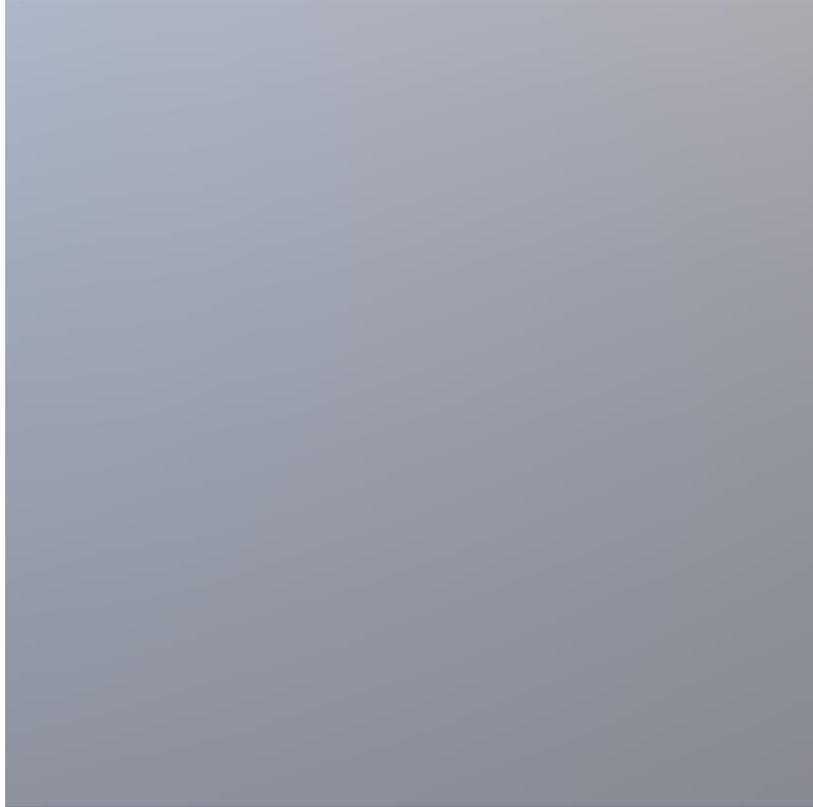
- Stack all pixel values, feed into MLP
- Problem?
 - Not efficient (\vec{x} is of dimension width*height)
 - Patterns can be anywhere in the image



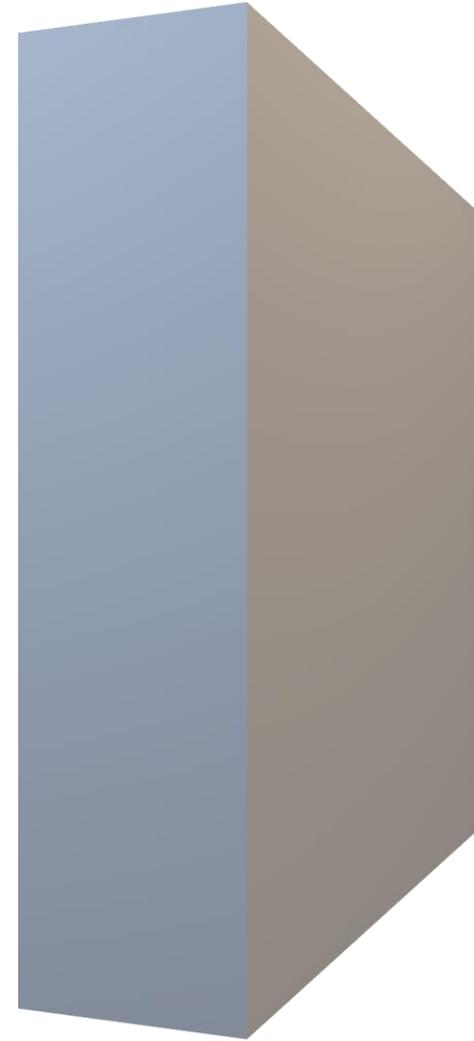
Convolutional Neural Network



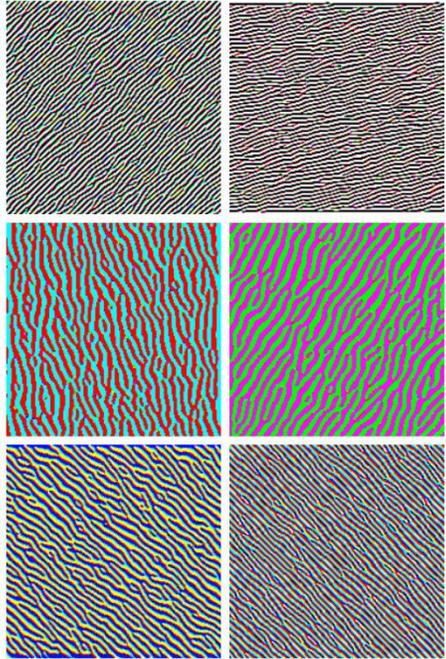
Schematics



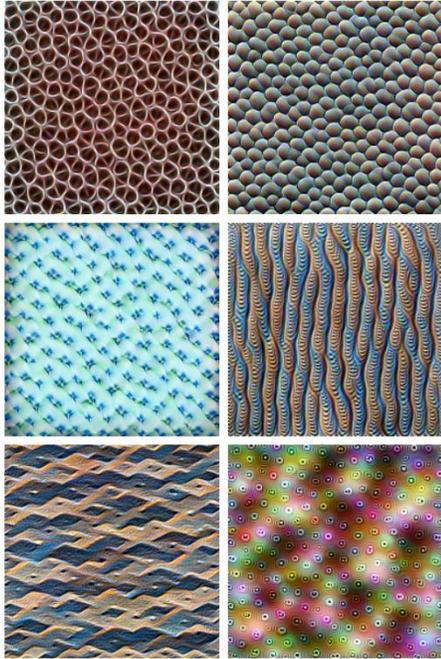
Schematics



Hierarchies of Features in CNNs



Edges (layer conv2d0)



Textures (layer mixed3a)



Patterns (layer mixed4a)



Parts (layers mixed4b & mixed4c)

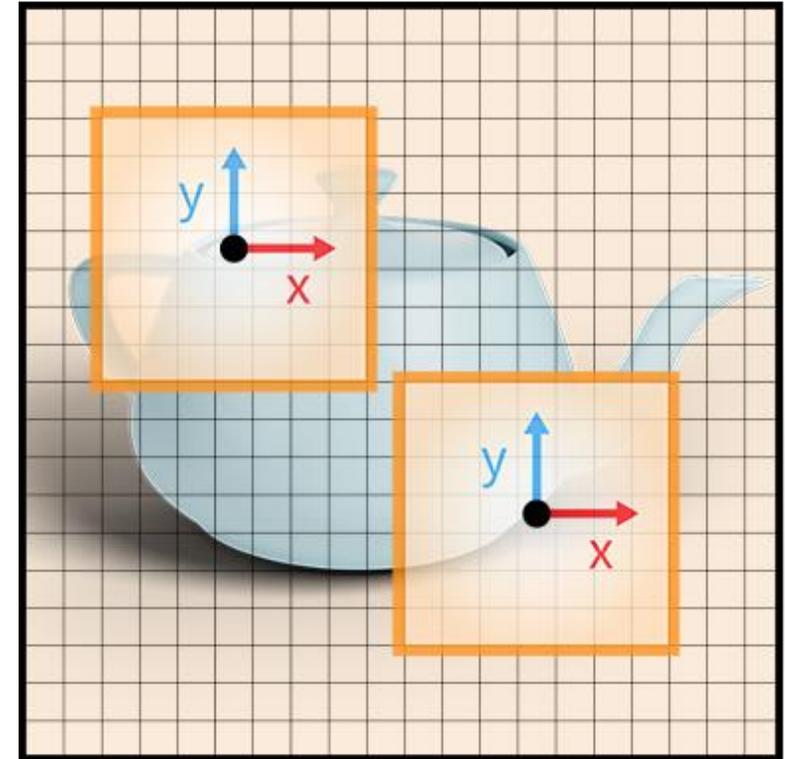


Objects (layers mixed4d & mixed4e)

Feature Visualization, Olah, Mordvintsev, Schubert (2017) – <https://distill.pub/2017/feature-visualization/>

Convolutional Neural Networks Notes

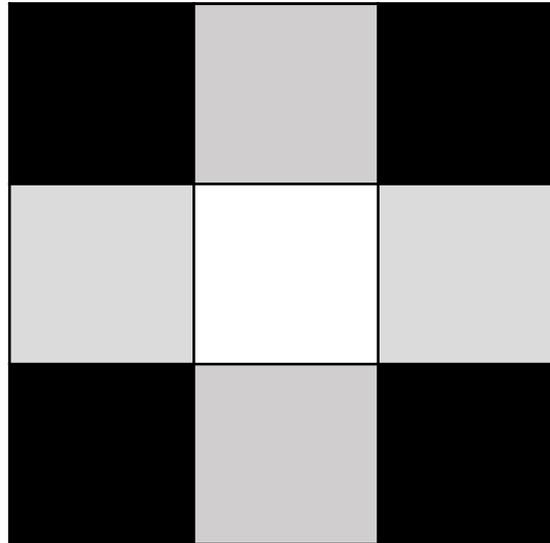
- Interpretation: MLP on patches
- Consistent pixel grid helps
 - Orientation
 - Maps to hardware well
- Downside (?) Template is fixed



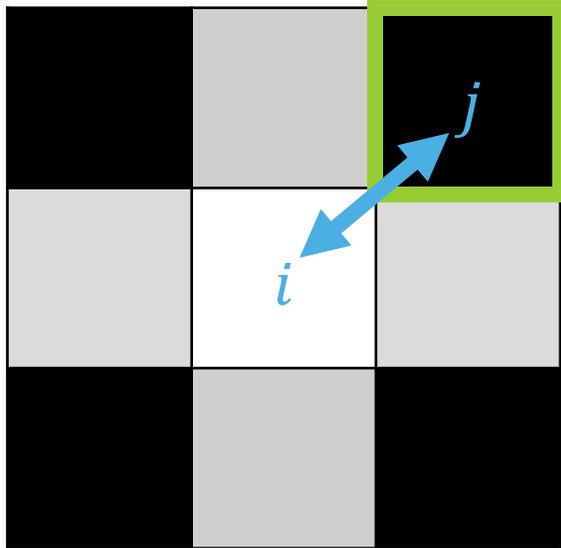
Limitation of Convolution Kernels

Kernel values based on location

Can we adapt to the signal?



Limitation of Convolution Kernels



Gaussian kernel g

$$x'_i = \sum_{j \in \mathcal{N}_i} \underbrace{g(\|p_j - p_i\|)} x_j$$

Based on distance

Denoising with Gaussian filter vs. Bilateral filter

$$x'_i = \sum_{j \in \mathcal{N}_i} g(\|p_j - p_i\|) x_j$$



Isotropic blurring

$$x'_i = \sum_{j=0}^N f(\|x_j - x_i\|) g(\|p_j - p_i\|) x_j$$

Look at pixel value as well



Preserves features

Self-attention (Transformers)

$$x'_i = \sum_{j=0}^N f(\|x_j - x_i\|) g(\|p_j - p_i\|) x_j$$

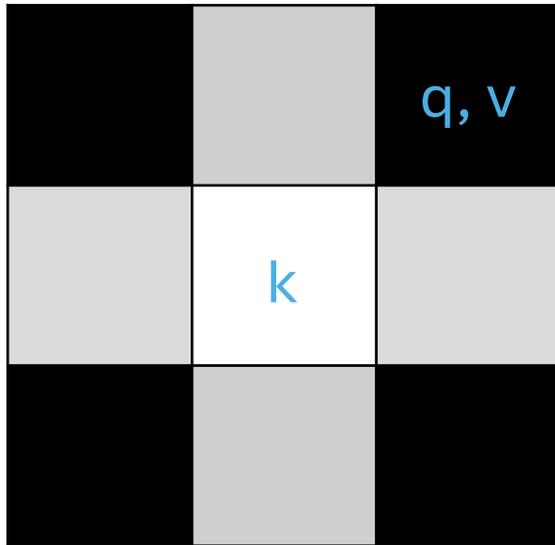
- Replace distance with cosine similarity

$$x'_i = \sum_{j=0}^N x_j^T x_i x_j$$

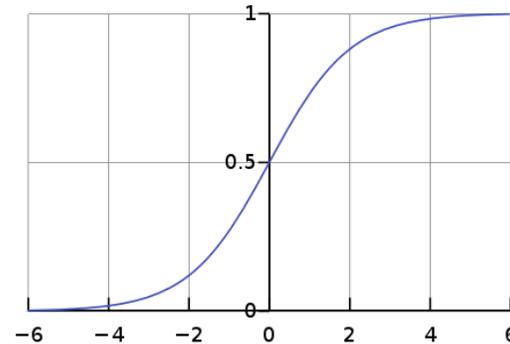
- Add flexibility with weight matrices and make sure weight ‘behaves’

$$x'_i = \sum_{j=0}^N \text{softmax} \left((W_q x_j)^T W_k x_i \right) W_v x_j$$

Self-attention (Transformers)



$$x'_i = \sum_{j=0}^N \underbrace{\text{softmax}} \left(\underbrace{(W_q x_j)^T}_{\text{Query}} \underbrace{W_k x_i}_{\text{Key}} \right) \underbrace{W_v x_j}_{\text{Value}}$$



Self-attention (Transformers)

$$x'_i = \sum_{j=0}^N \text{softmax} \left((W_q x_j)^T W_k x_i \right) W_v x_j$$

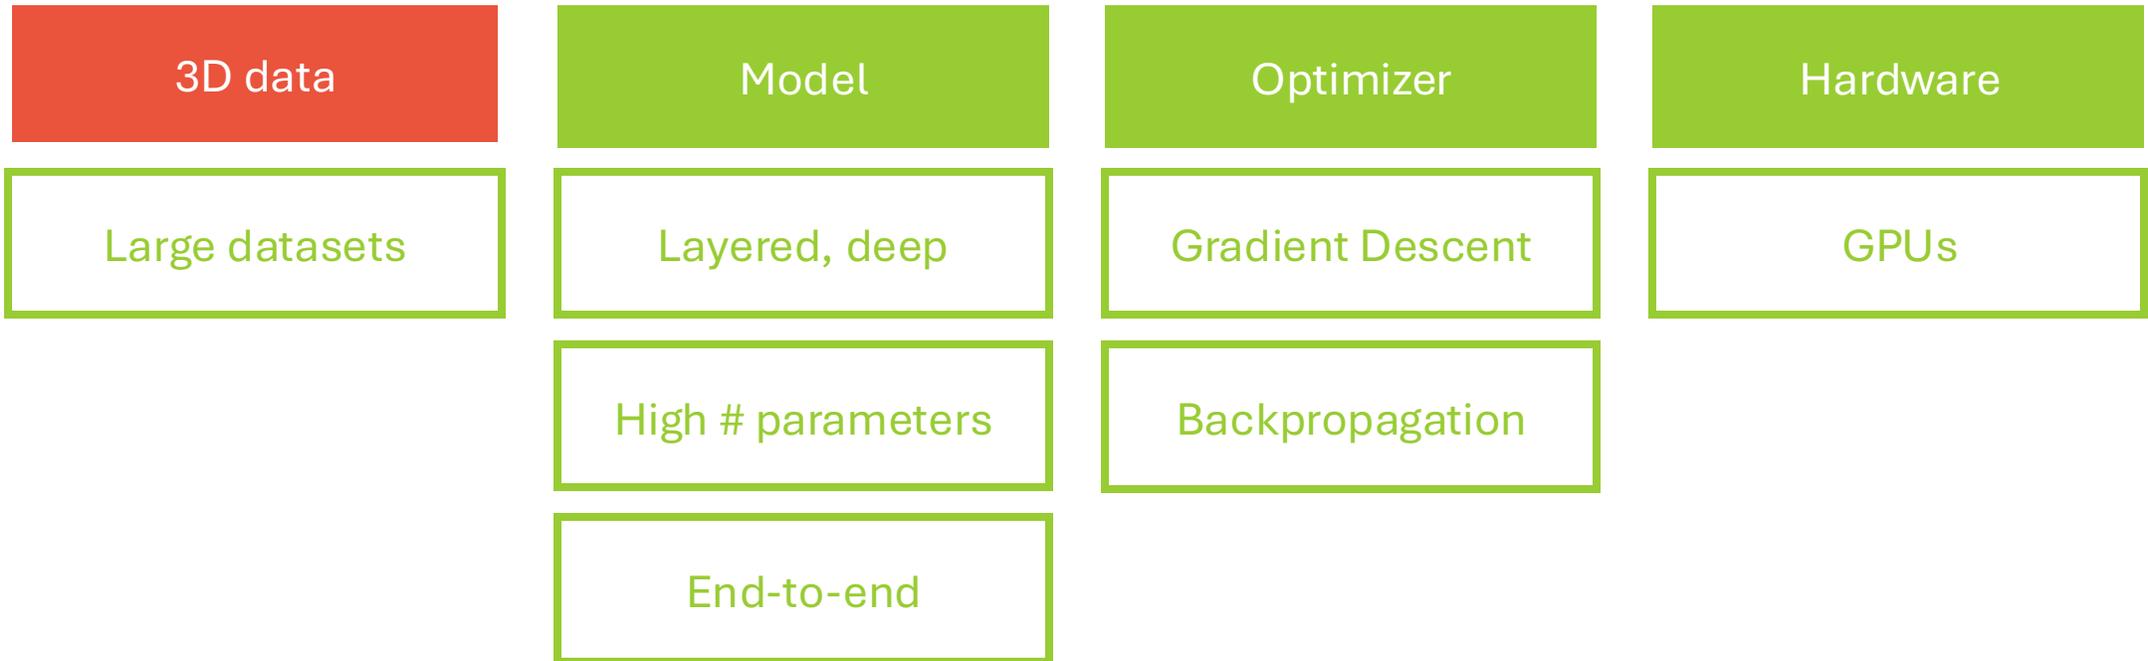
- Transformers combine these blocks
- Positional encoding based on $p_j - p_i$
- **Pro:** Highly flexible, very effective (LLMs, Vision Transformers, etc.)
- **Pro:** ‘Global’ connections (vs. U-Net)
- **Con:** Computationally expensive (compared to CNN)
 - Solutions: work at coarse scale OR only apply locally

Summary

- **Multi-Layer Perceptrons (MLP)**
 - Linear combination, followed by non-linearity, repeated
 - Basic building block of Neural Networks
- **Convolutional Neural Network (CNN)**
 - Learn local kernel, convolve
 - Weight sharing, translation invariance
 - ‘Constructs’ localized features of increasing abstraction
- **Transformers**
 - Learn weights based on feature similarity
 - Highly expressive, but expensive; current SOTA for many tasks

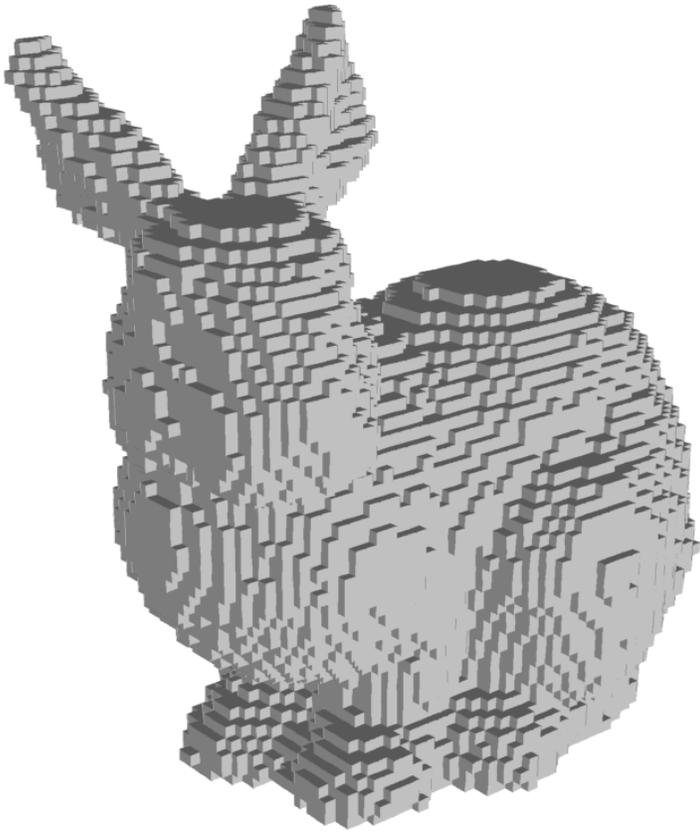
Chapter 2: 3D Data

Deep learning for 3D data

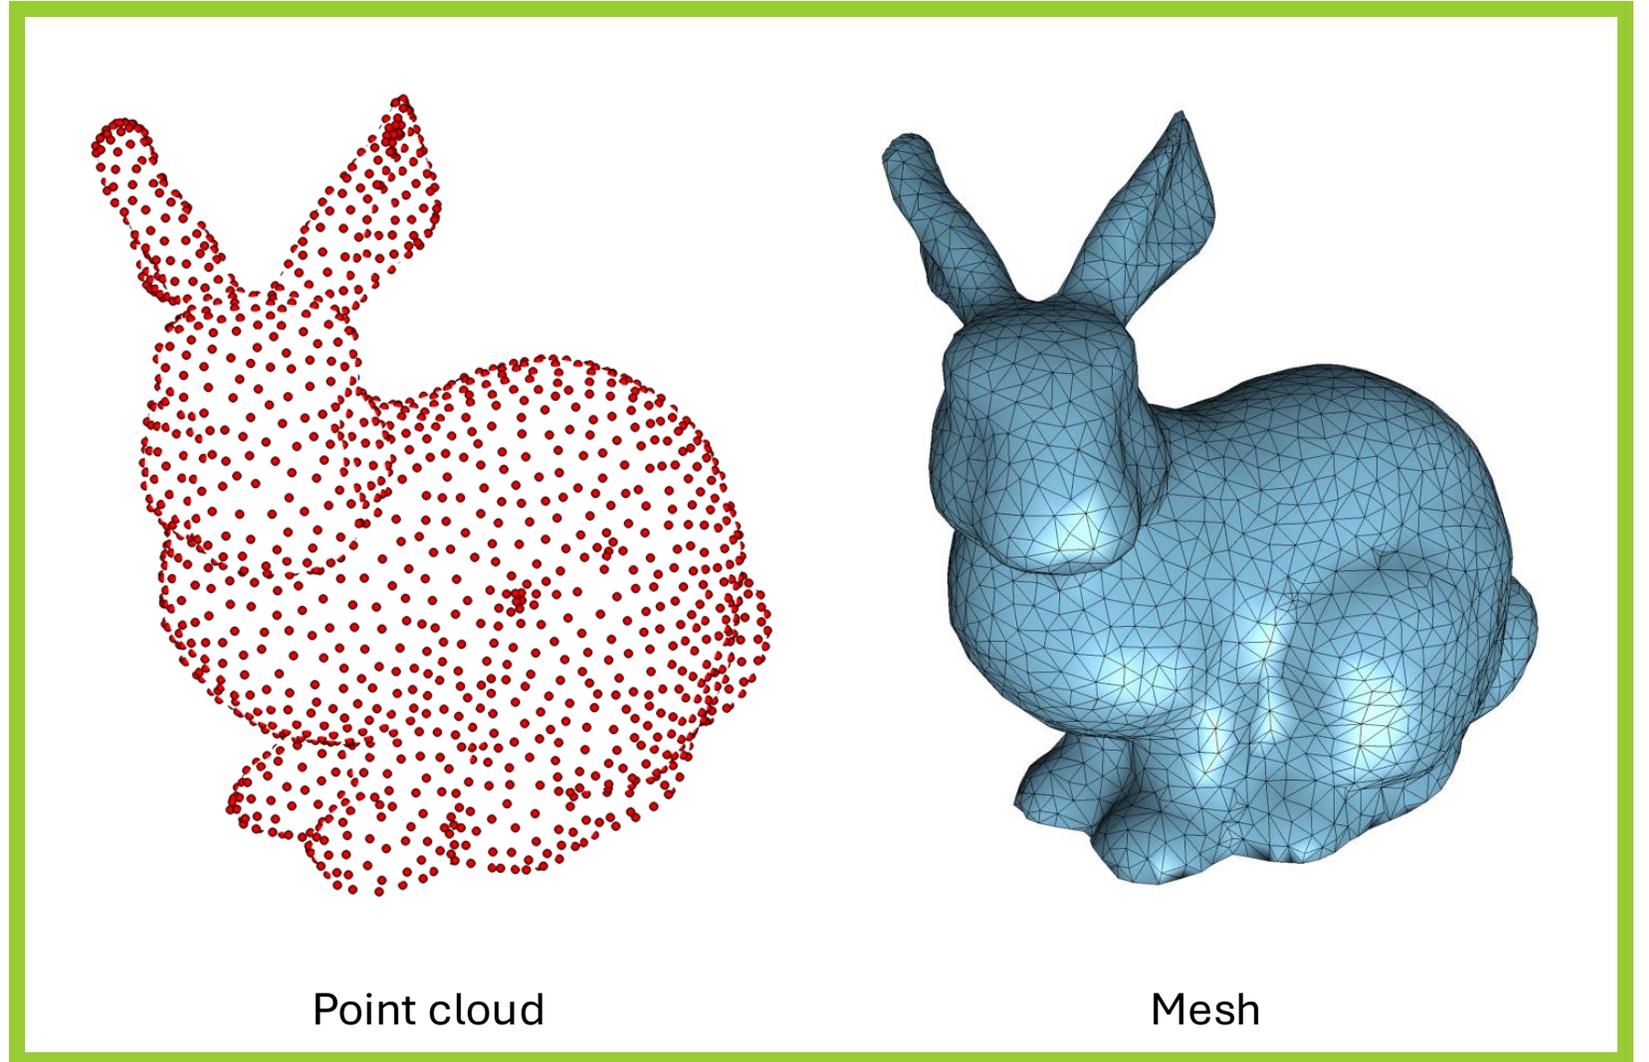


‘Philosophy’: Scaling (data, compute) beats algorithmic complexity

3D Data



Voxels



Point cloud

Mesh

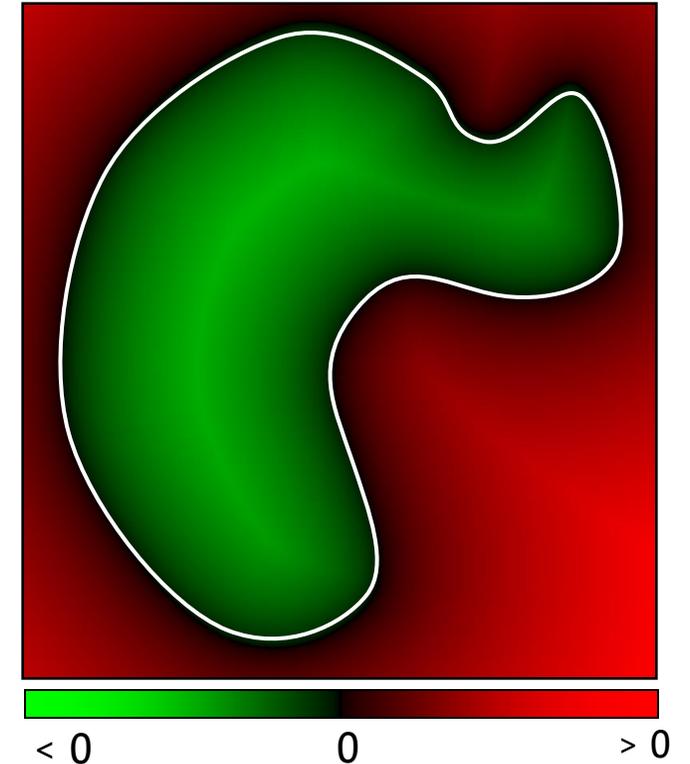
What about implicit functions?

$$f: \mathbb{R}^3 \rightarrow \mathbb{R}$$

Value > 0 outside shape, < 0 inside

In deep learning

- Use MLP to represent f
- Rarely used as input to MLP
 - Convert to voxels/mesh/point cloud

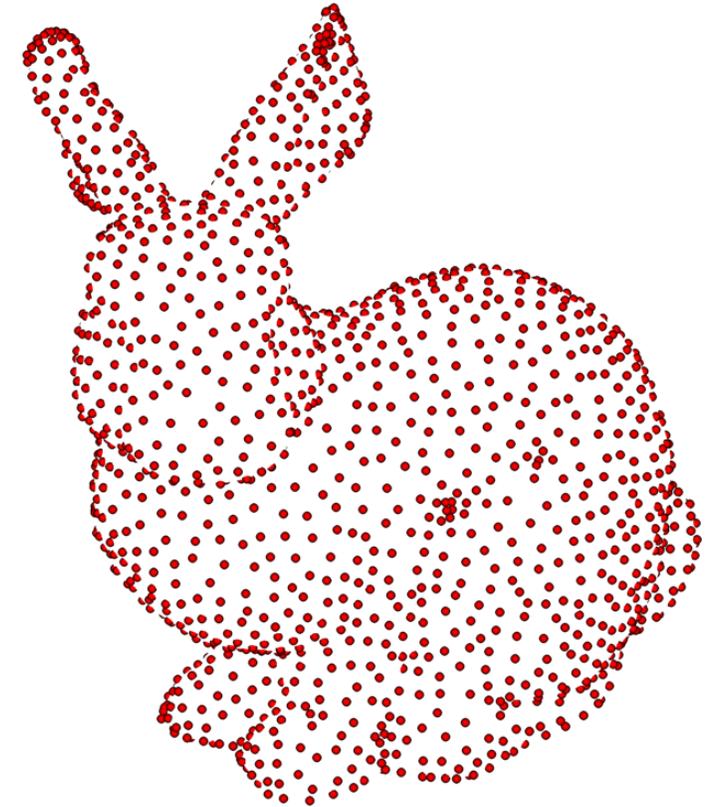


Based on slides by Olga Sorkine-Hornung

Point Cloud

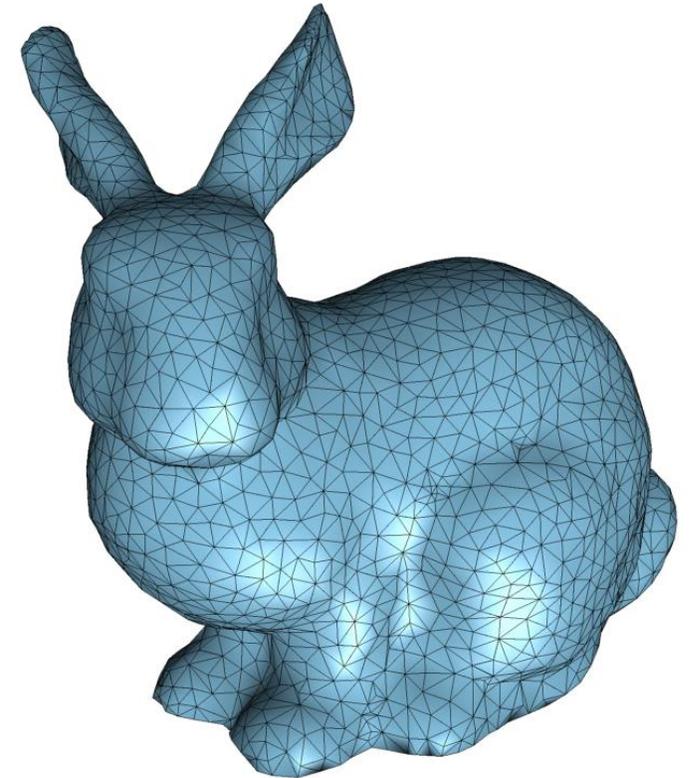
- Set of points in d-dimensional space

$$P \in \mathbb{R}^3$$



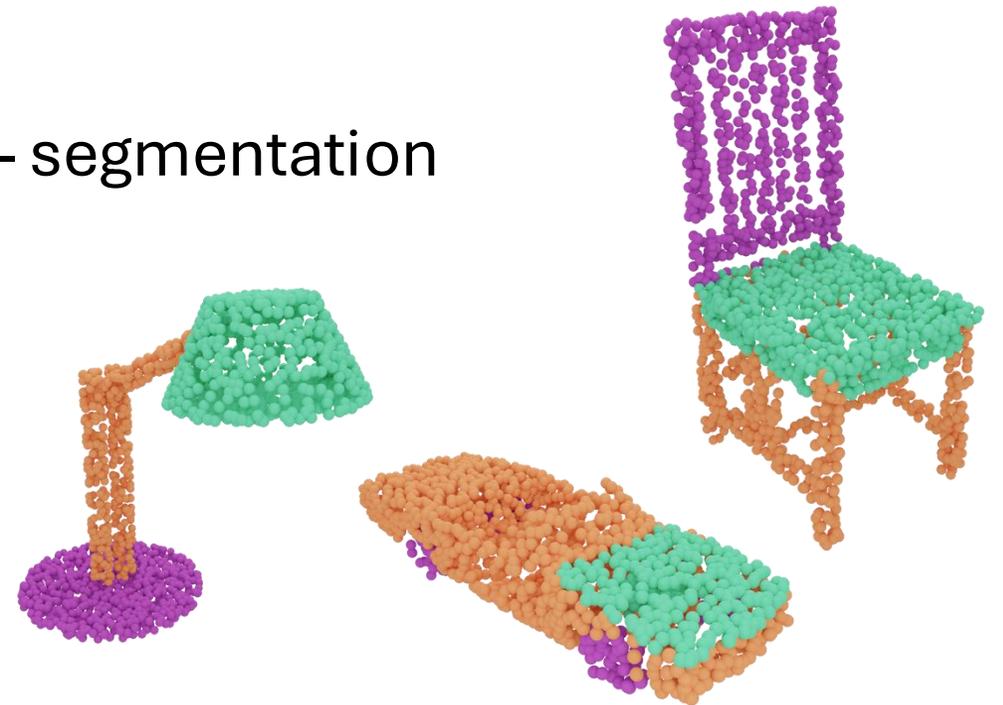
Mesh

- ‘Complements’ point cloud
 - connectivity and surface given by V , E , F
- Geometry can be on points, edges, etc.



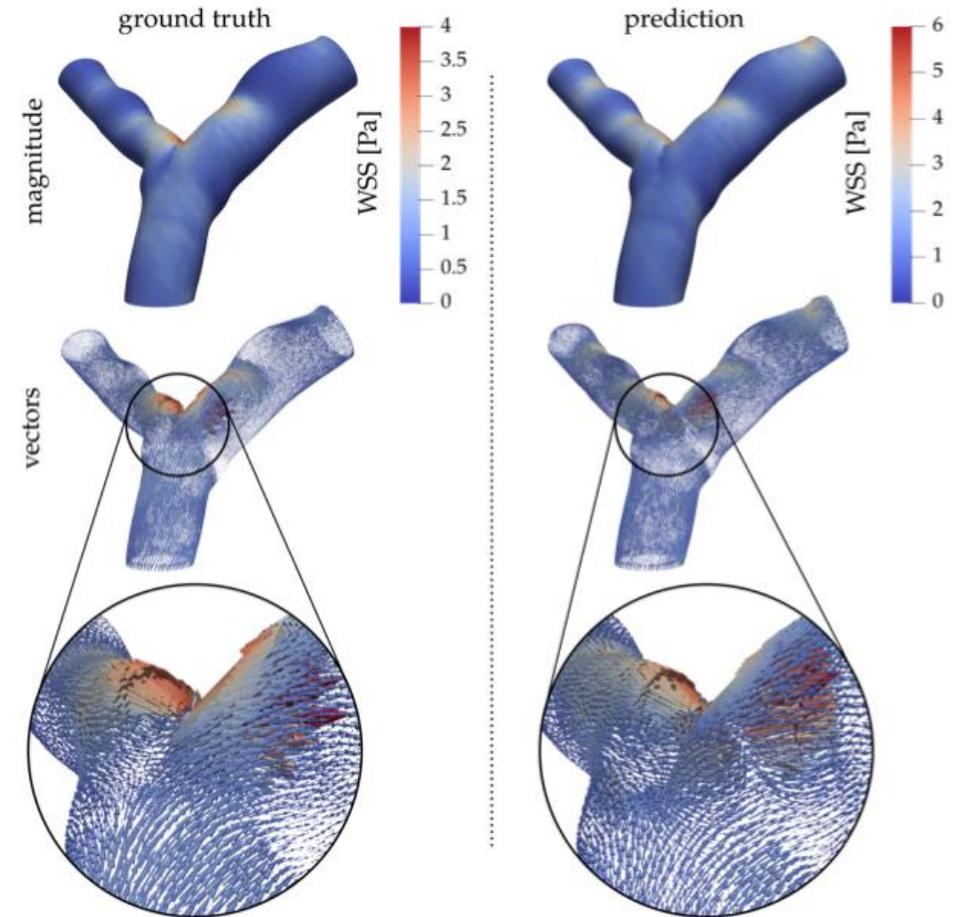
Datasets

- **ShapeNet** (meshes, point clouds) – segmentation, classification
- **ModelNet40** (meshes) – classification
- **ScanNet (v2, v3)** (RGB-D + reconstructions) – segmentation
- **S3DIS** (point clouds) – segmentation
- **Thingy10k** (meshes) – n/a
- **ABC** (meshes) – n/a
- Consider ethics, copyright!
 - E.g., <https://huggingface.co/datasets/allenai/objaverse/discussions/18>



More exotic inputs and outputs: simulation

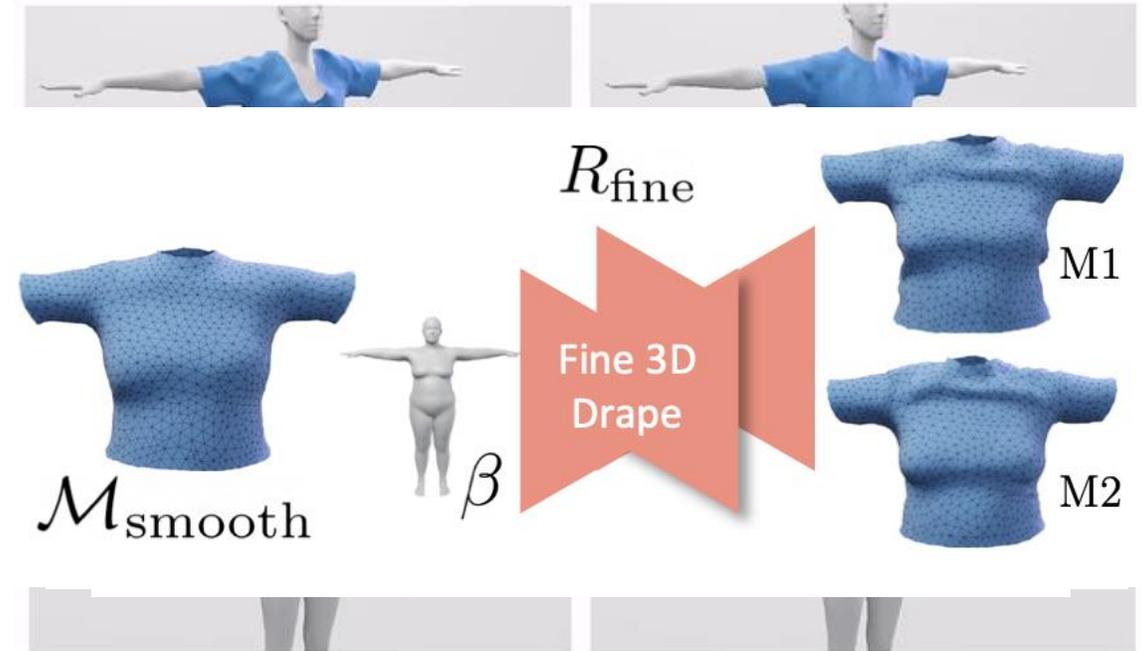
- **Input:** positions, distance to loops
- **Output:** tangent vector at each vertex



Mesh Neural Networks for SE(3)-Equivariant Hemodynamics Estimation of the Artery Wall, Suk et al. (2022)

More exotic inputs and outputs: simulation

- Stage 1
 - **Input:** Garment parameters
 - **Output:** Mesh positions on mean shape
- Optimize topology
- Stage 2
 - **Input:** Mesh positions, target shape parameters
 - **Output:** Smooth mesh positions
- Stage 3
 - **Input:** Smooth mesh positions
 - **Output:** Fine mesh positions (wrinkles, etc.)



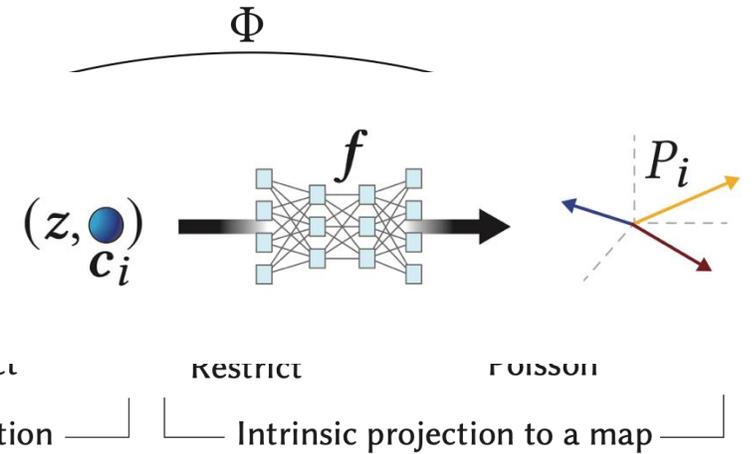
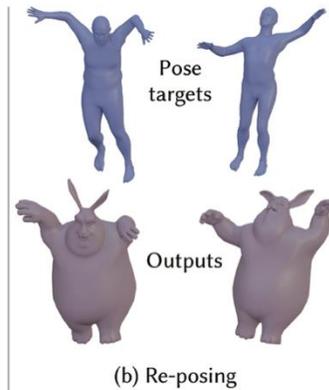
Fully Convolutional Graph Neural Networks for Parametric Virtual Try-On, Vidaurre et al. (2020)

More exotic inputs and outputs: Jacobians

- Neural Jacobian Fields
 - **Input:** Per-triangle centroid + global code
 - **Output:** Jacobian matrix
 - **Model:** MLP

- Post-processing

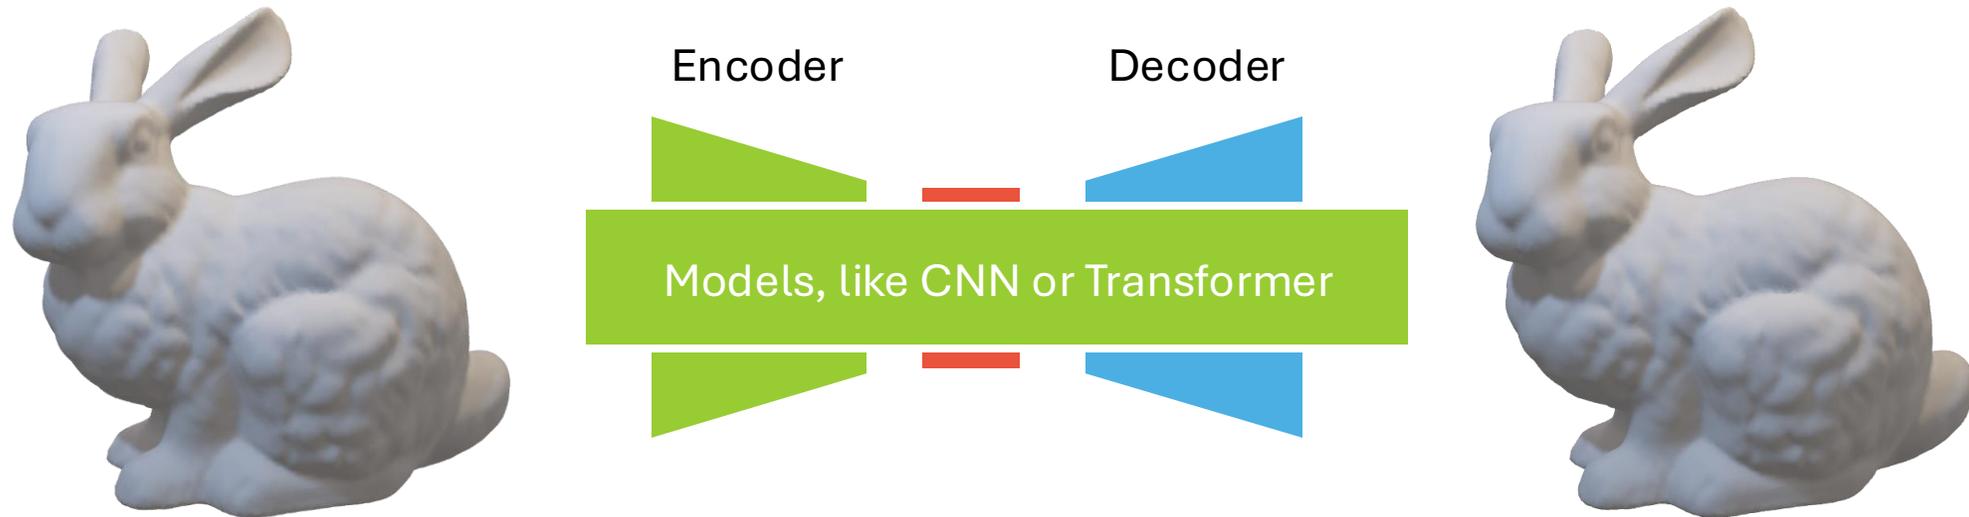
- Restrict Jacobian
- Solve Poisson



Neural Jacobian Fields: Learning Intrinsic Mappings of Arbitrary Meshes, Aigerman et al. (2022)

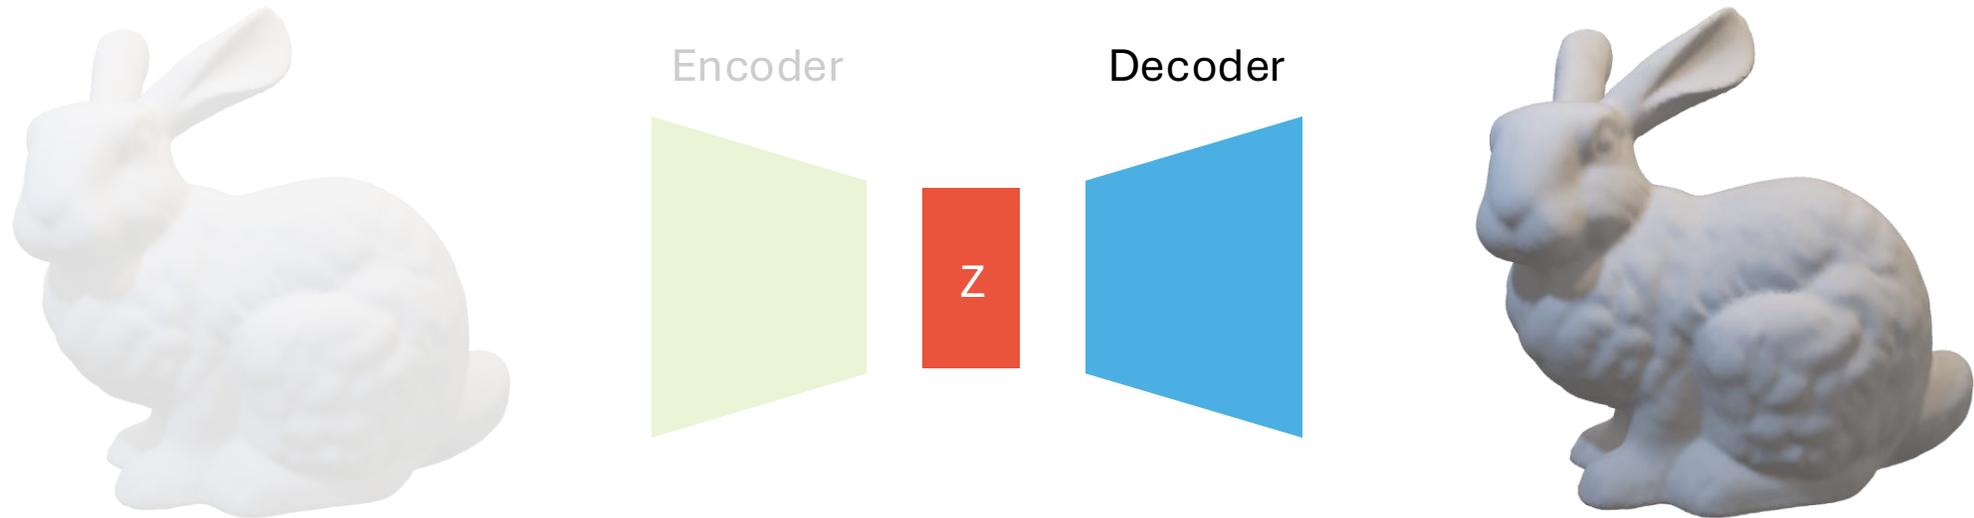
What if I have data, but no labels?

- Unsupervised learning
- For example: auto-encoders



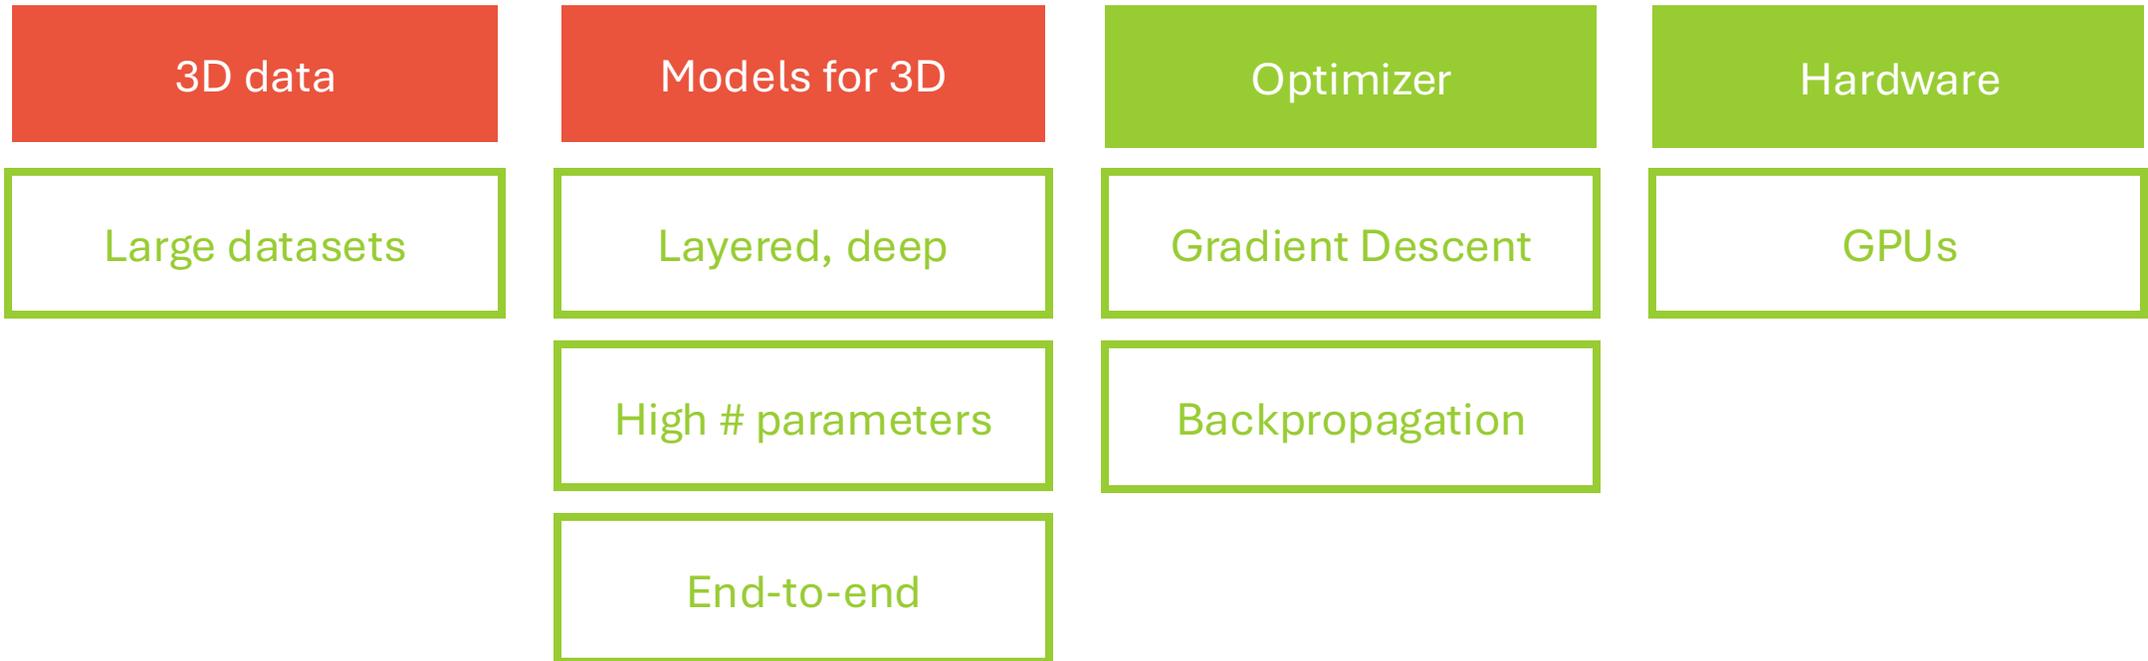
What if I have data, but no labels?

- Unsupervised learning
- For example: auto-encoders



Chapter 3: 3D Deep Learning Models

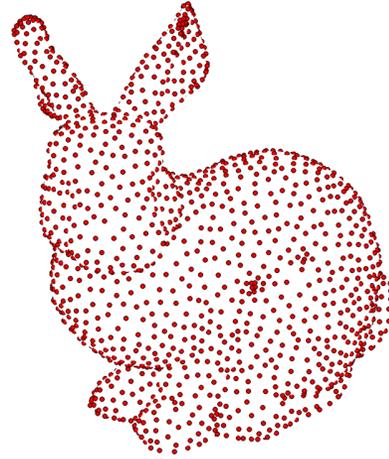
Deep learning for 3D data



‘Philosophy’: Scaling (data, compute) beats algorithmic complexity

Let's start simple

MLP



Stack all points?

Which order?

MLP

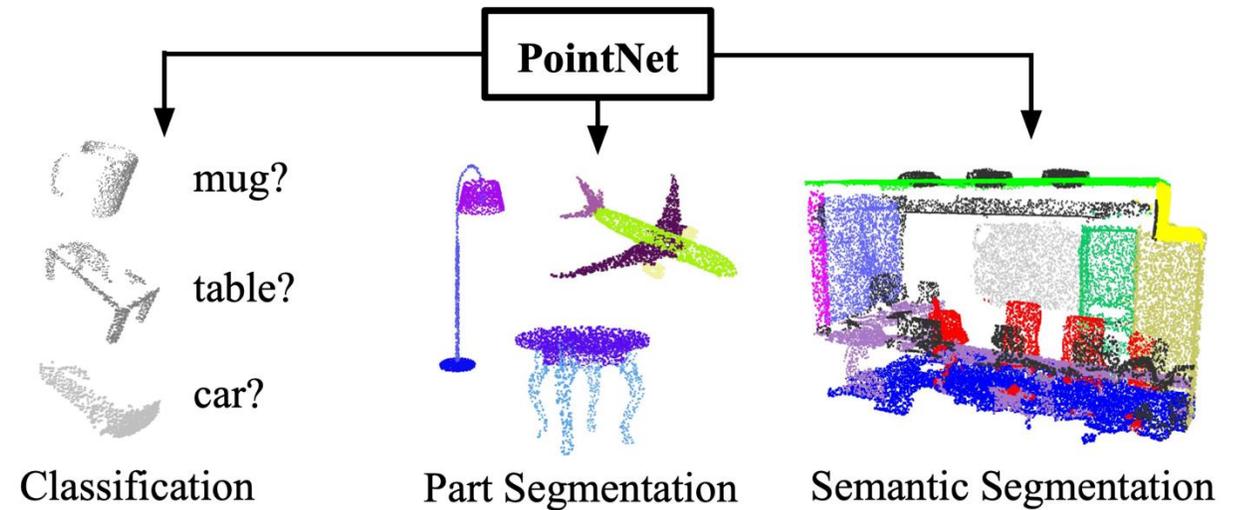
“Bunny”

PointNet

- MLP on each point
- Maximum over all points

$$y = \max_{j \in P} f_{\theta}(x_j)$$

Cannot learn from neighborhoods



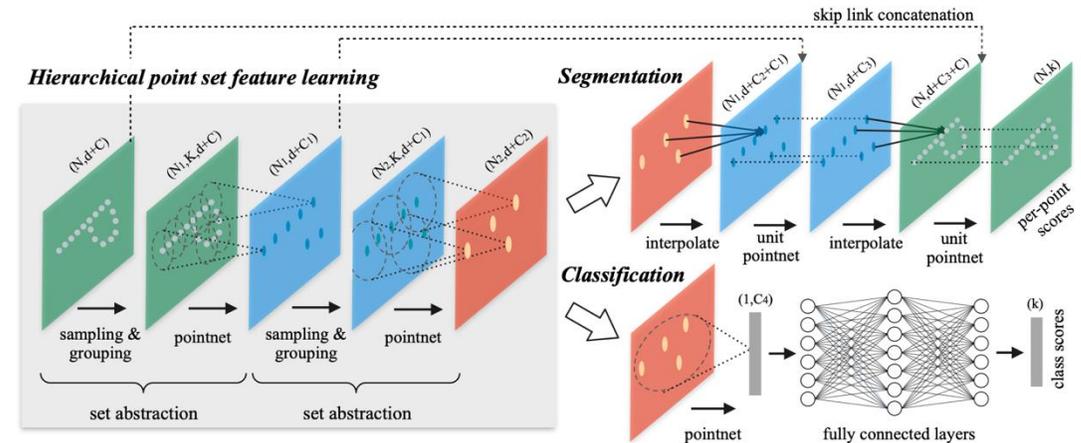
PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation, Qi, Su, Mo, Guibas (2016)

PointNet++

- MLP on each point
- Maximum over neighborhood
 - kNN – homogeneous
 - Radius – more robust to sampling
 - Geodesic/Euclidean?

$$x'_i = \max_{j \in N_i} f_\theta(x_j, p_j - p_i)$$

- Hierarchies with maximum pooling (like CNN)

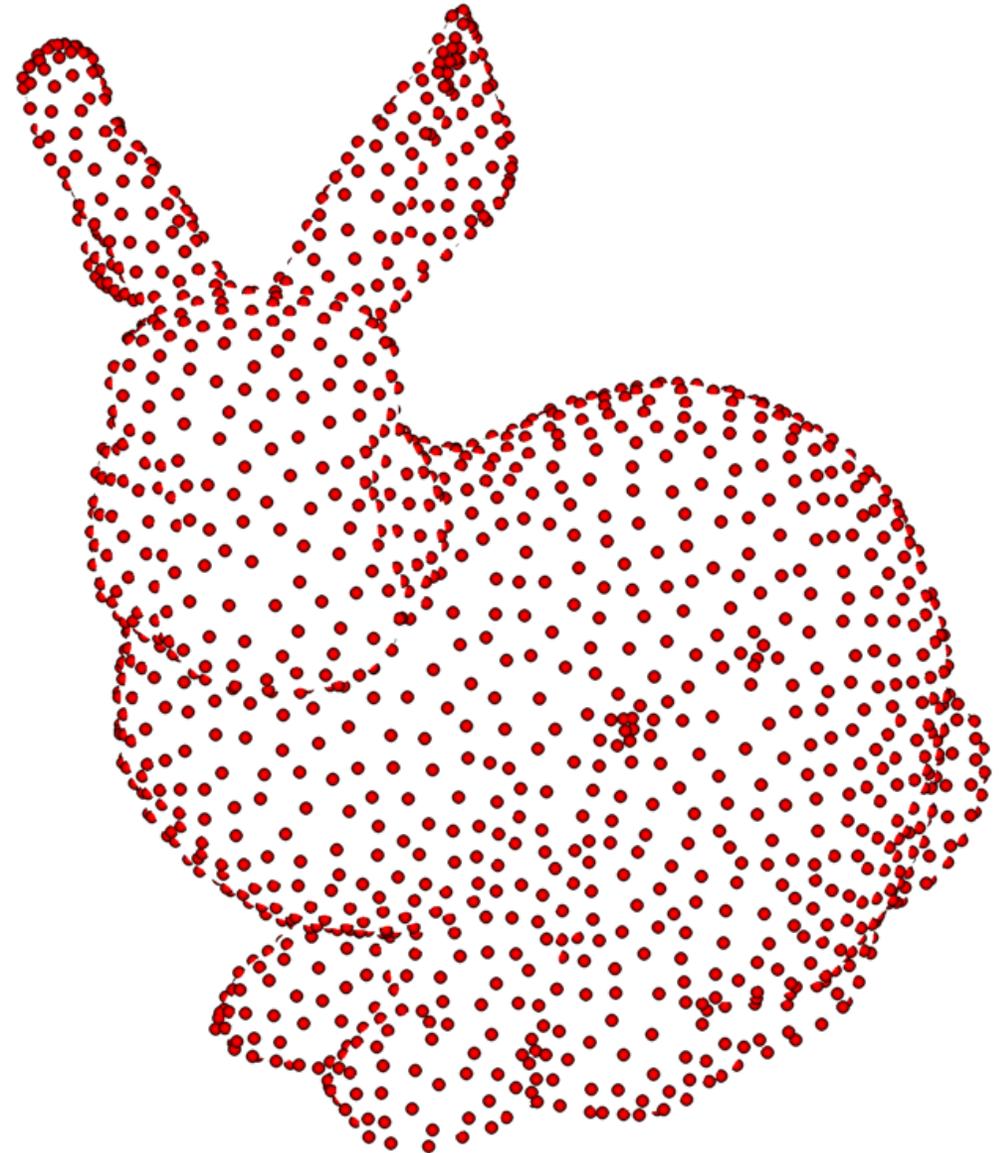


PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space, Qi, Yi, Su, Guibas (2017)

Taking a step back

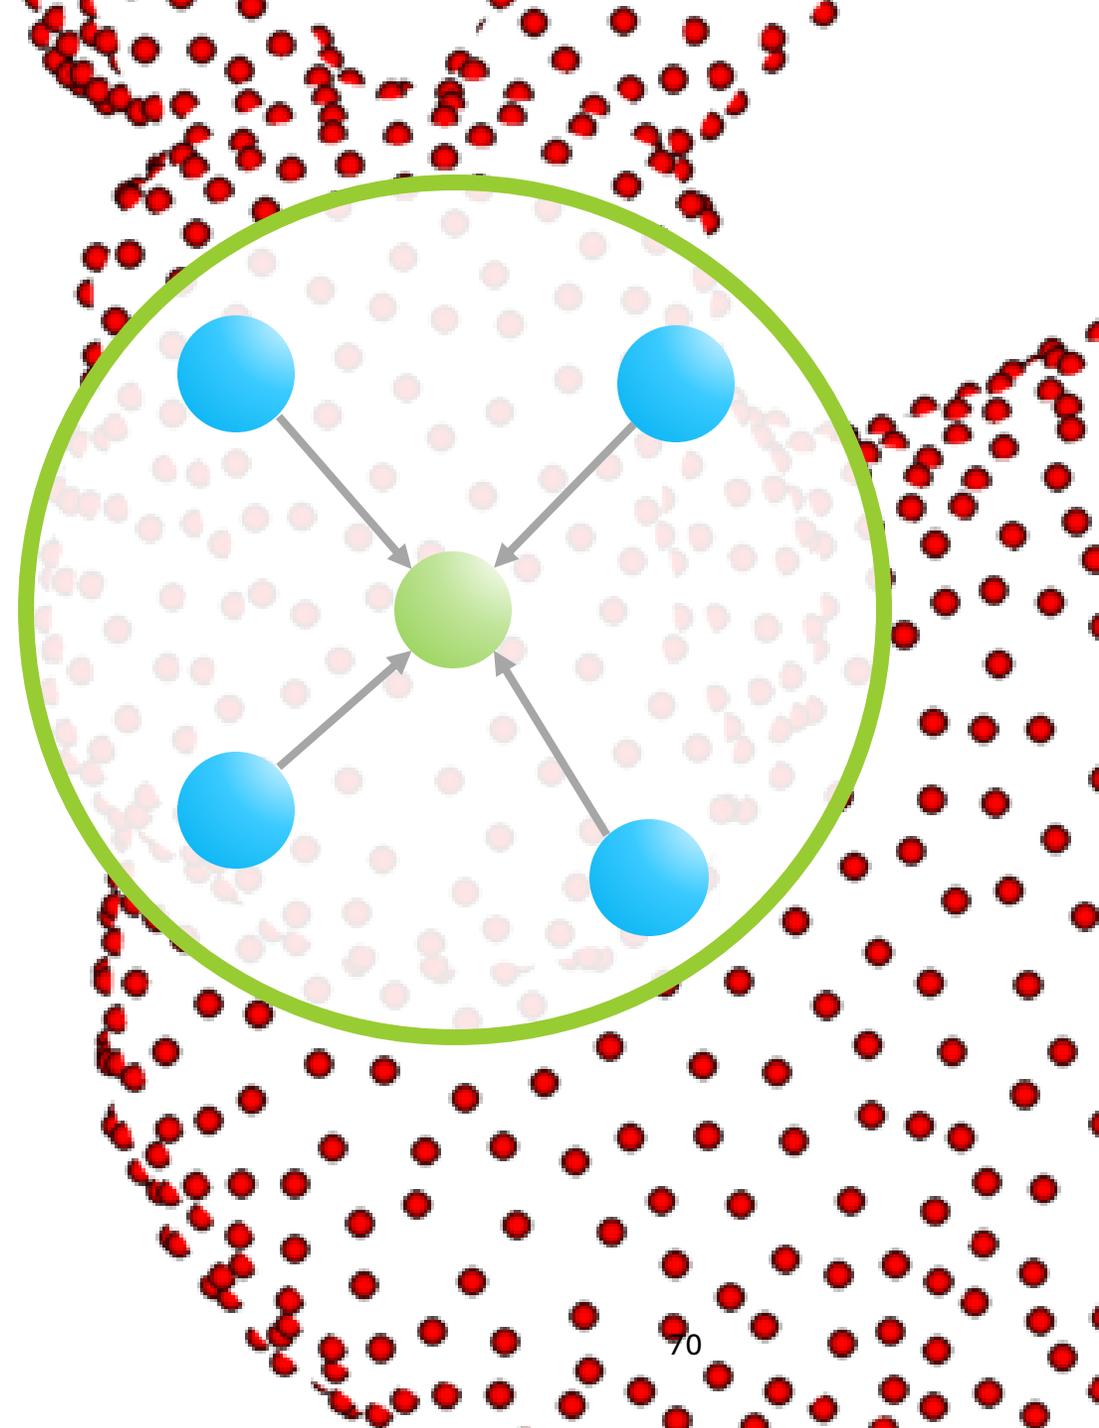
Surface as a graph

- Vertices or points are nodes



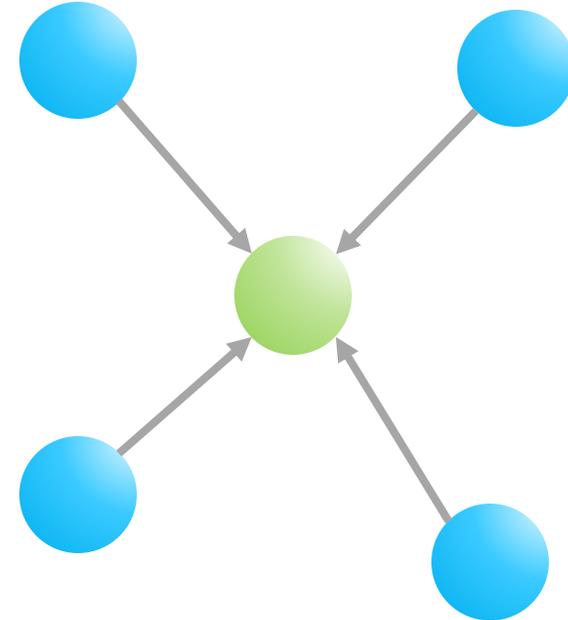
Surface as a graph

- Vertices or points are nodes
- Edges connect nearby points (radius graph, k-NN graph)
- For meshes: use mesh edges
- What can meshes help with?
 - Geodesic neighborhoods (we know connectivity)
 - Encode geometry (e.g., MeshCNN)



Message passing

1. Compute 'message' on each node
2. Aggregate messages over edges



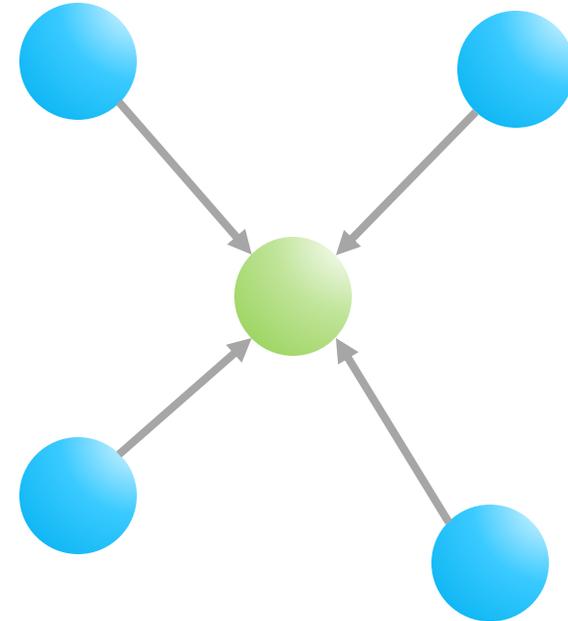
Neural Message Passing for Quantum Chemistry, Gilmer, Schoenholz, Riley, Vinayals, Dahl (2017)

https://pytorch-geometric.readthedocs.io/en/latest/tutorial/create_gnn.html

PointNet++ as Message Passing

1. **Message:** MLP on features + relative location
2. **Passing:** Maximum over neighbors

$$x'_i = \max_{j \in N_i} f_\theta(x_j, p_j - p_i)$$

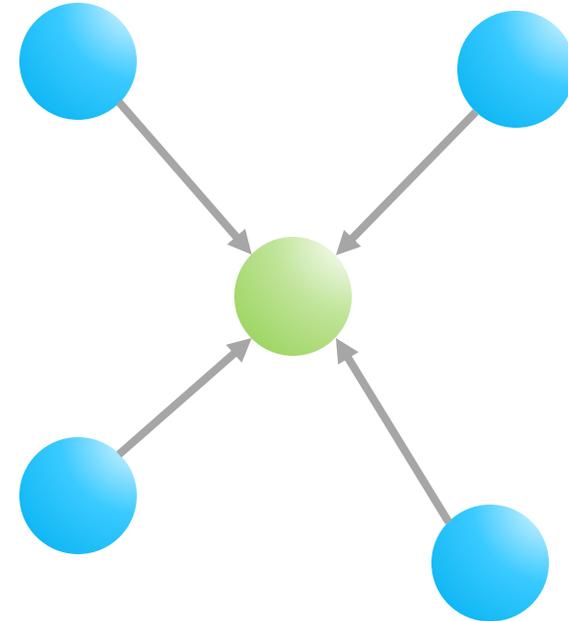


PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space, Qi, Yi, Su, Guibas (2017)

EdgeConv as Message Passing

1. **Message:** MLP on 'relative features' (edges)
2. **Passing:** Maximum over neighbors

$$x'_i = \max_{j \in N_i} f_{\theta}(x_i, x_j - x_i)$$

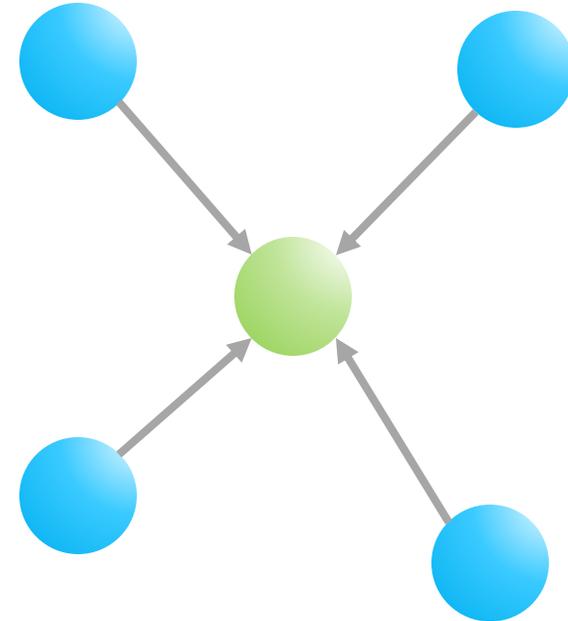


Dynamic Graph CNN for Learning on Point Clouds, Wang, Sun, Liu, Sarma, Bronstein, Solomon (2019)

GCN as Message Passing

1. **Message:** Linear transformation of features
2. **Passing:** Weight by degree, average

$$x'_i = \sigma(W_0 x_i + \sum_{j \in N_i} \frac{1}{c_{ij}} W_1 x_j)$$



Graph Convolutional Networks, Kipf, Welling (2016)

Laplacian in GCN and EdgeConv

- Laplacian: Sum of second derivatives
 - Discrete setting: Difference to average of neighbors

- GCN: graph Laplacian

$$x'_i = \sigma(W_0 x_i + \sum_{j \in N_i} \frac{1}{c_{ij}} W_1 x_j)$$

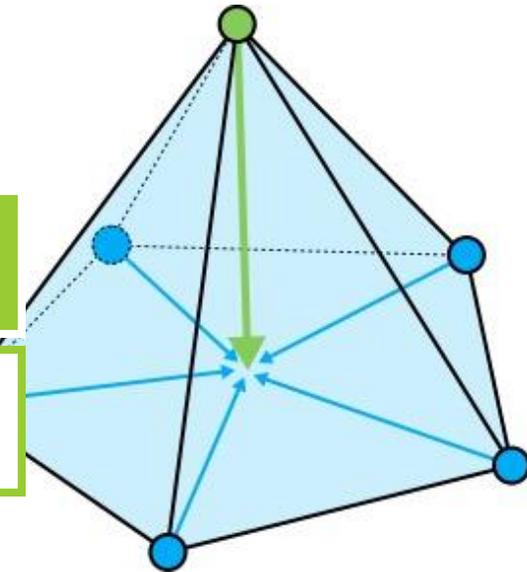
- EdgeConv

$$x'_i = \max_{j \in N_i} f_\theta(x_i, x_j - x_i)$$

Learning to Diffuse

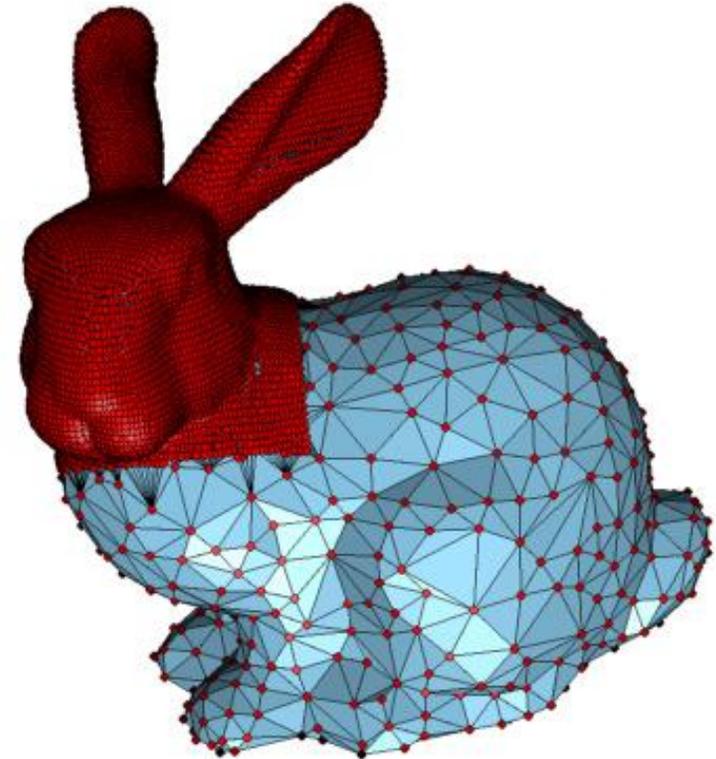
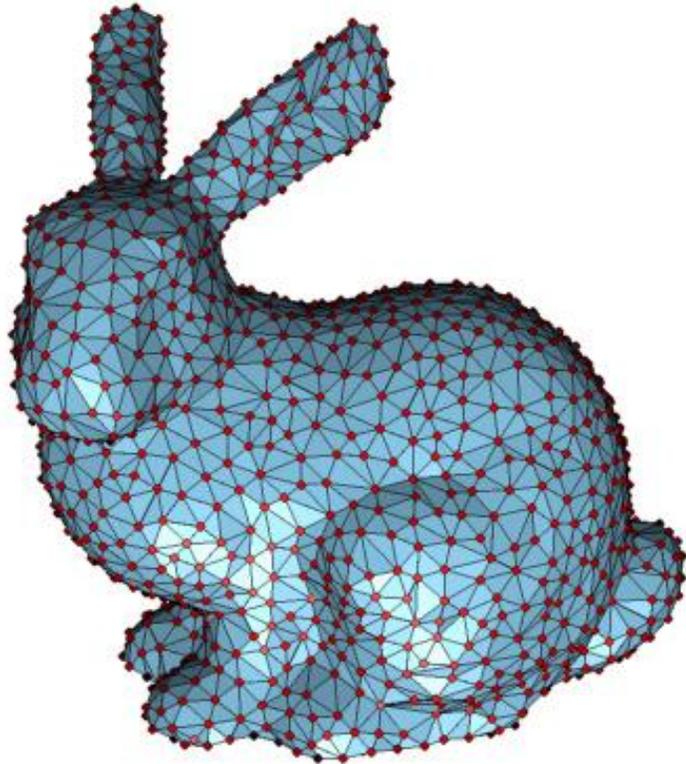
Smoothly

'Sharply'



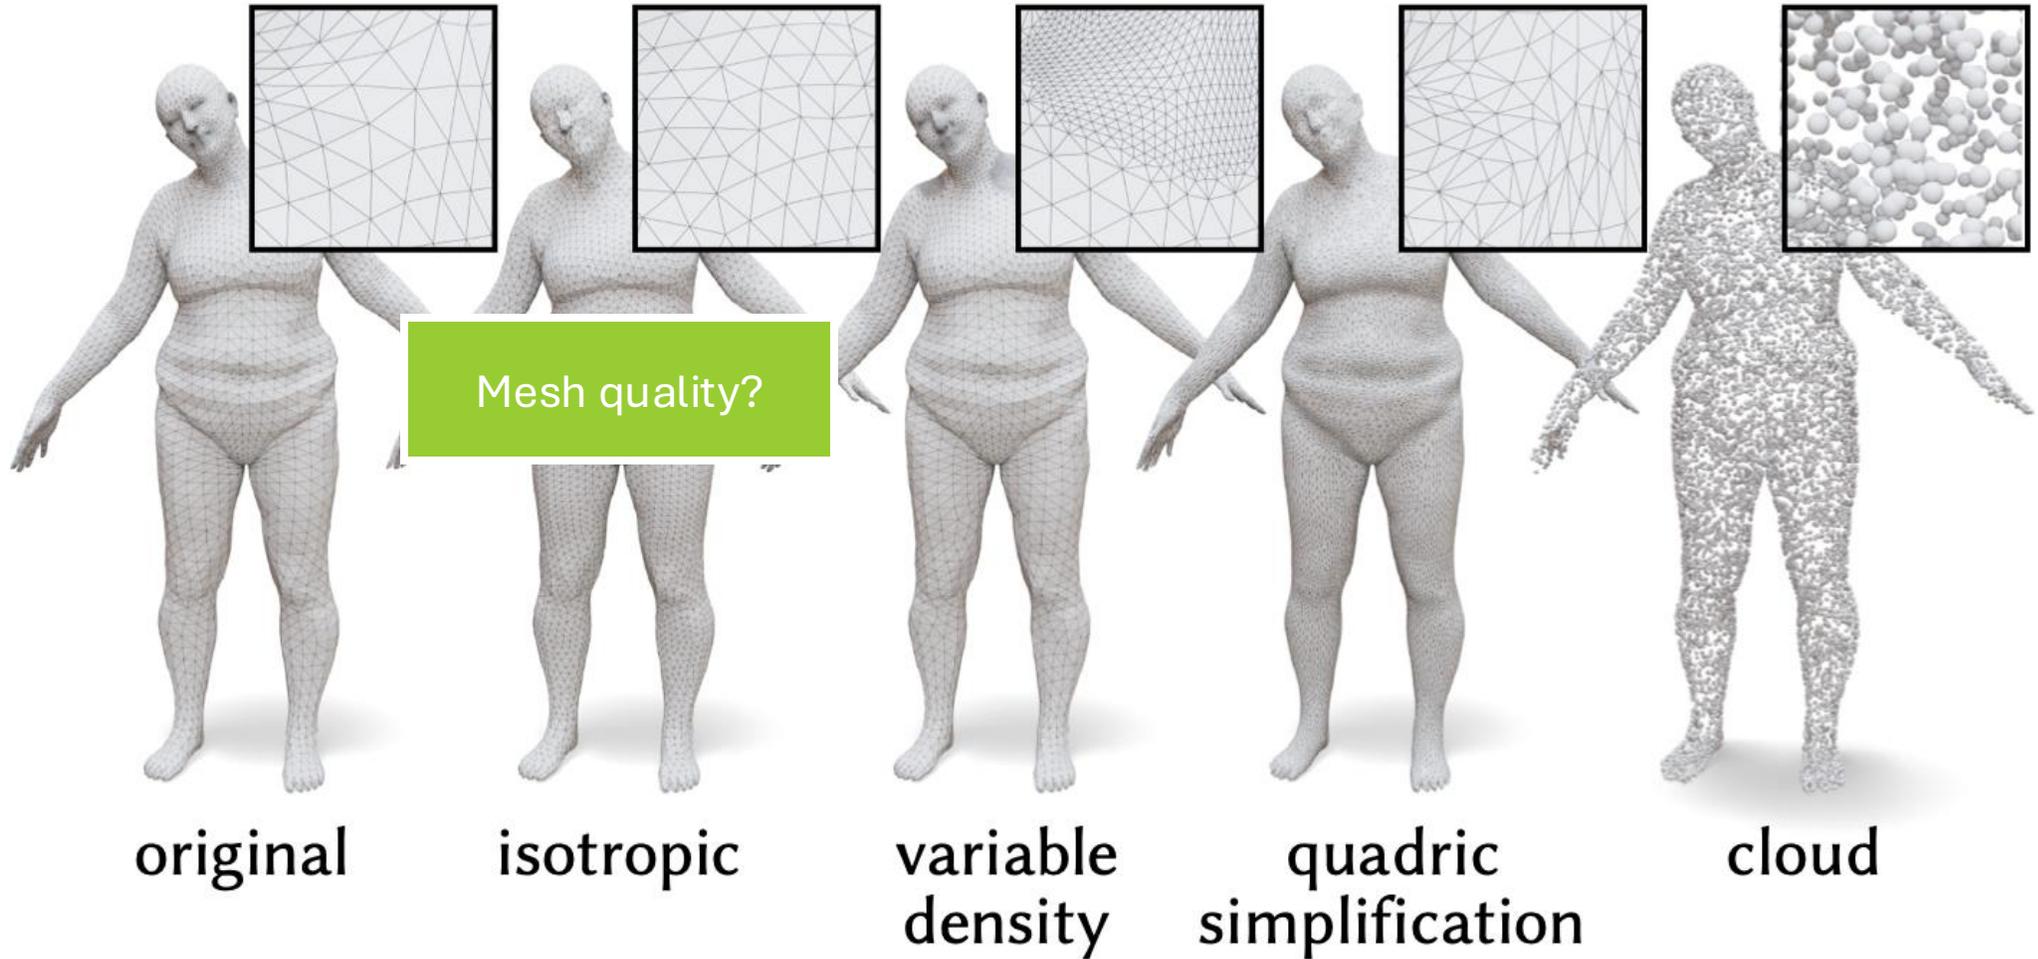
More geometry

Why geometry?



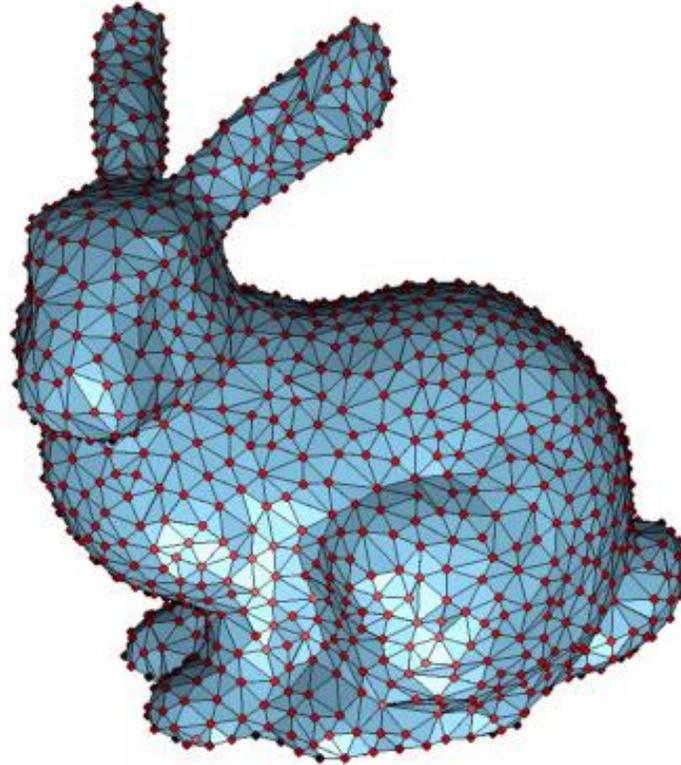
Based on slides by Klaus Hildebrandt

Why geometry?



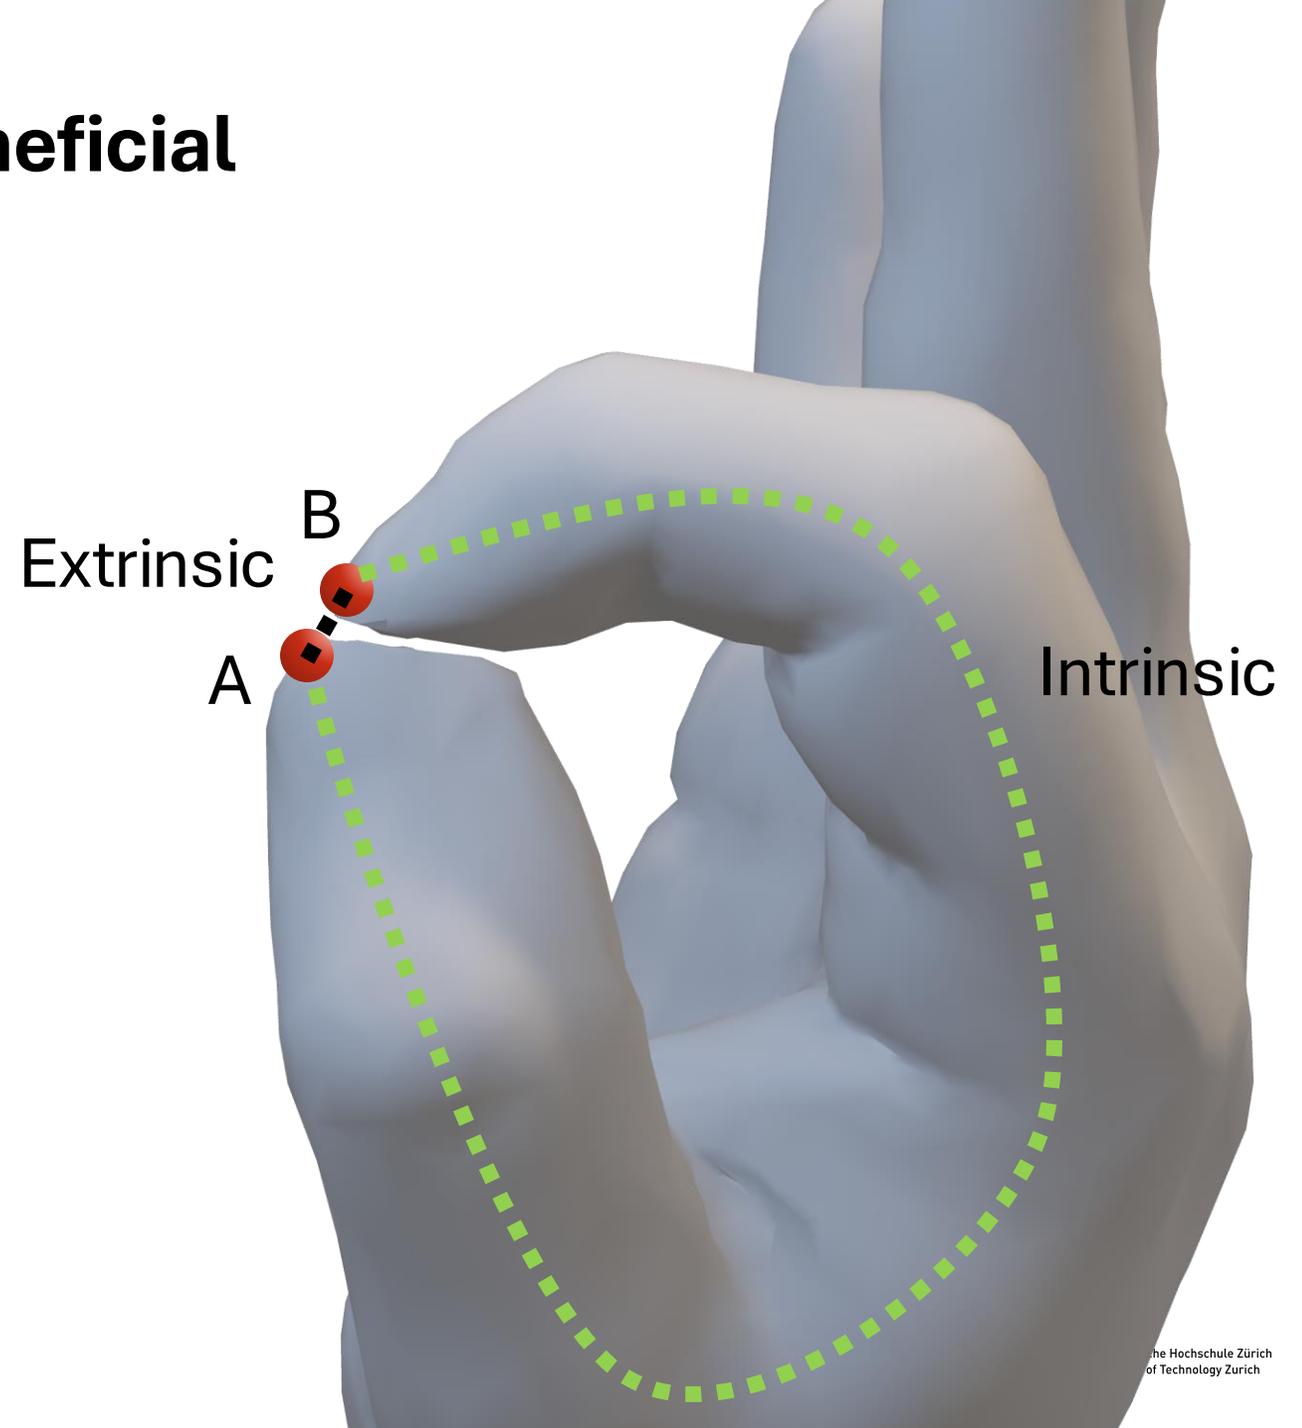
Why geometry?

- Invariances



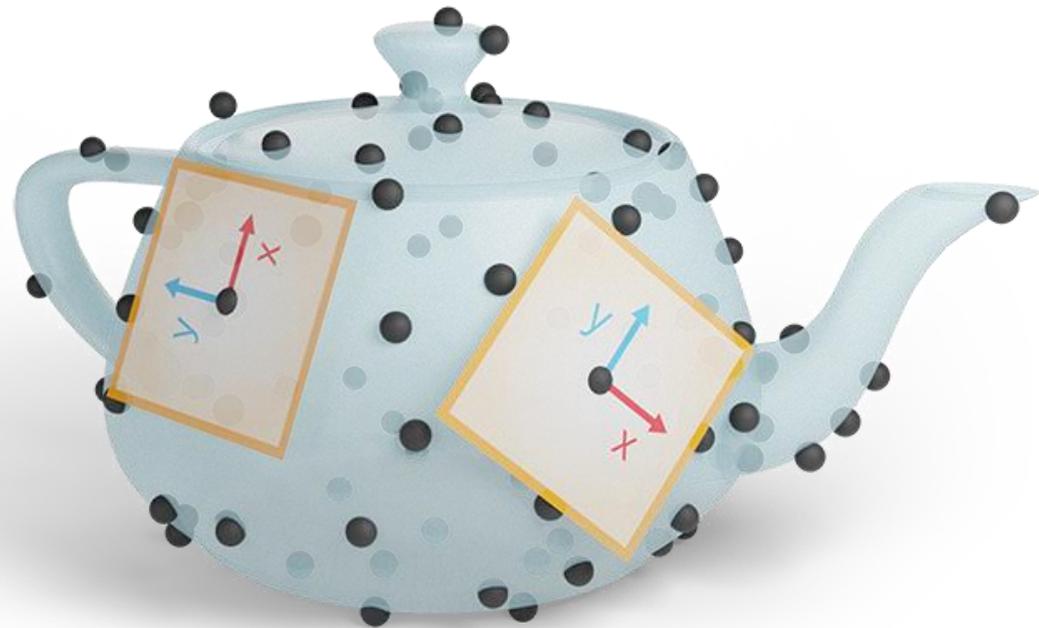
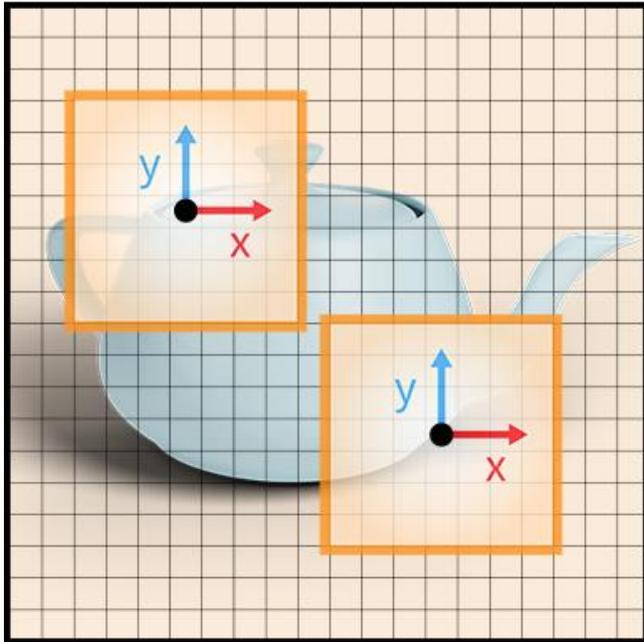
Intrinsic operations can be beneficial

- + Robust to isometric deformations
- + 2D instead of 3D
- + No/less distortion or occlusion



Why Geometry?

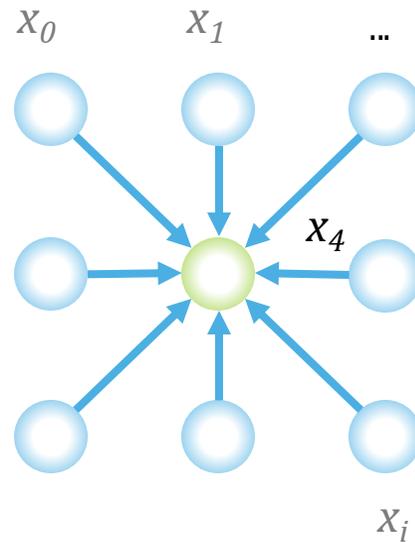
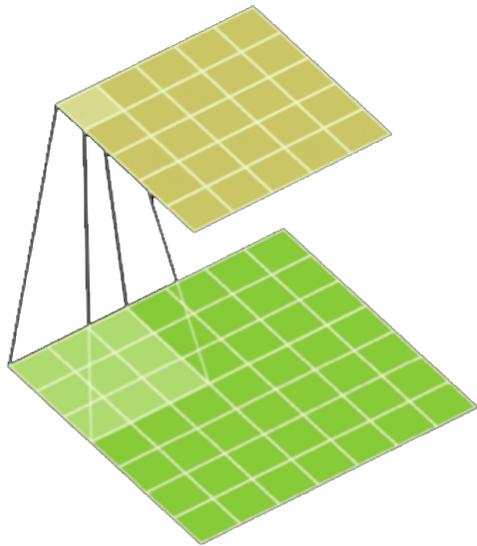
- Invariances within an invariance



Back to CNNs

CNN on a Graph

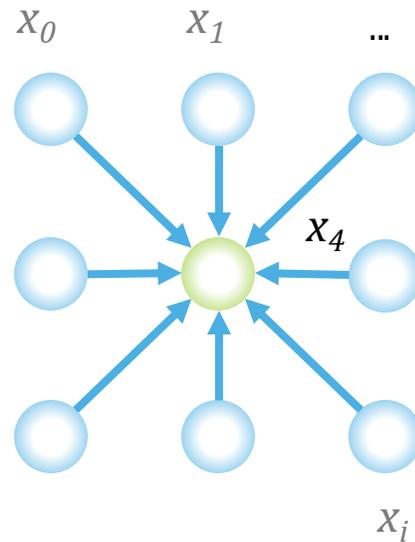
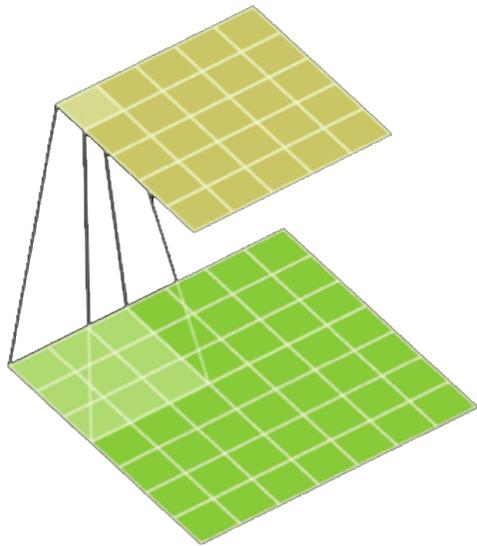
Single CNN layer with 3x3 filter



$$x'_4 = \sigma\left(\sum_i W_i x_i\right)$$

CNN on a Graph

Single CNN layer with 3x3 filter



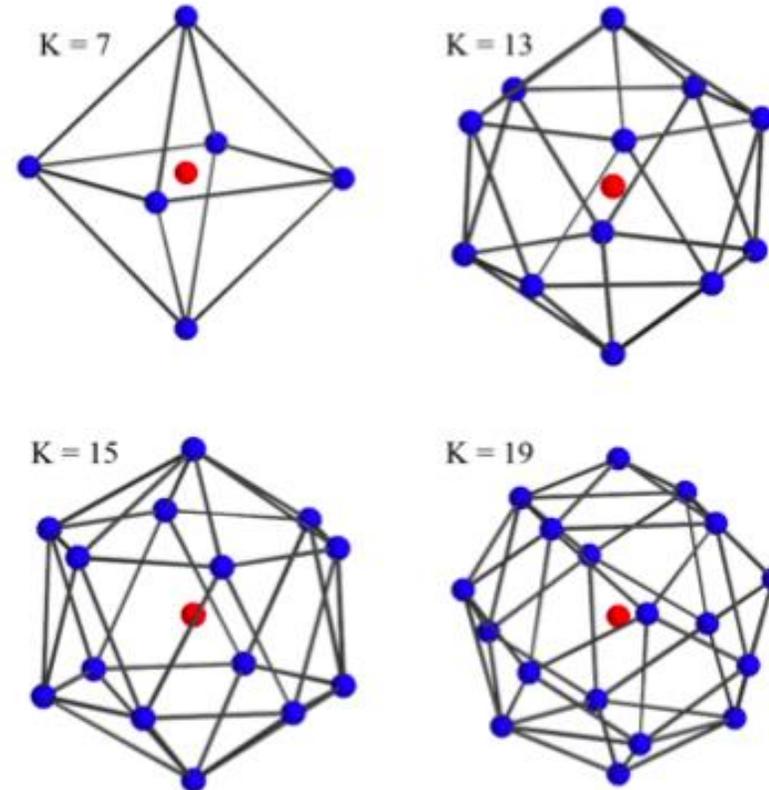
$$x'_4 = \sigma\left(\sum_i \boxed{W_i} x_i\right)$$

Compare with GCN

$$x'_i = \sigma\left(W_0 x_i + \sum_{j \in N_i} \frac{1}{c_{ij}} \boxed{W_1} x_j\right)$$

CNNs for 3D

- Graph- and point based
 - GCN, PointNet++, EdgeConv
- 3D kernel (extrinsic)
 - KPConv, MinkowskiNet, SSCN

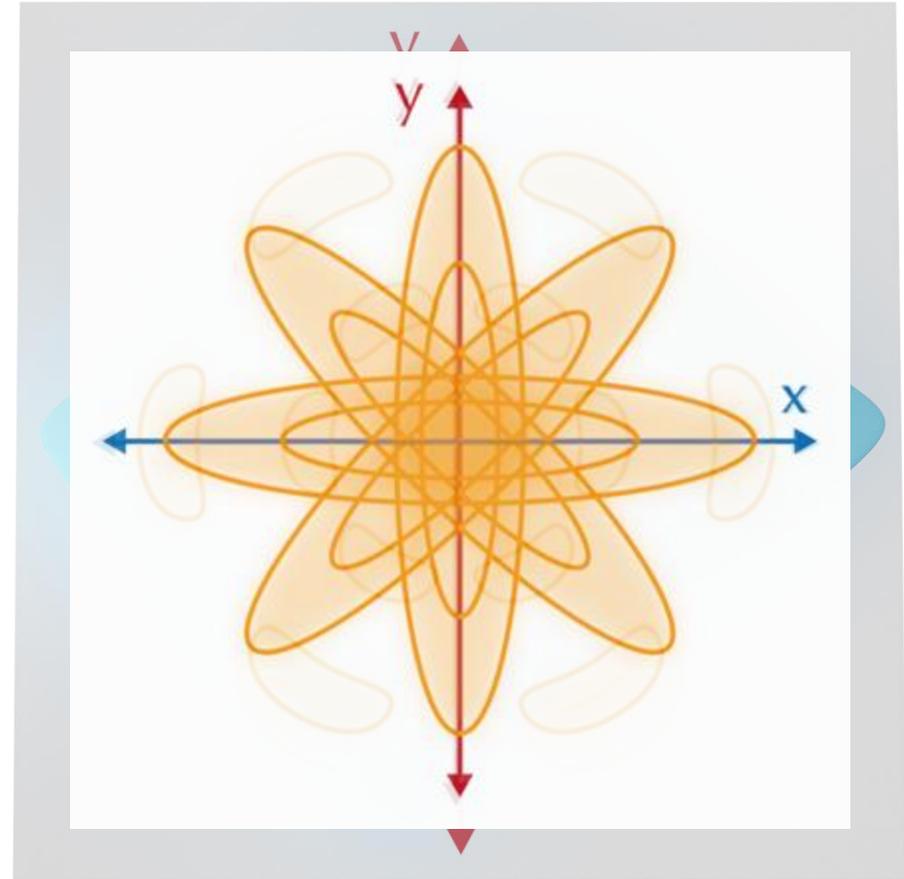


KPConv

[Thomas et al. 2019]

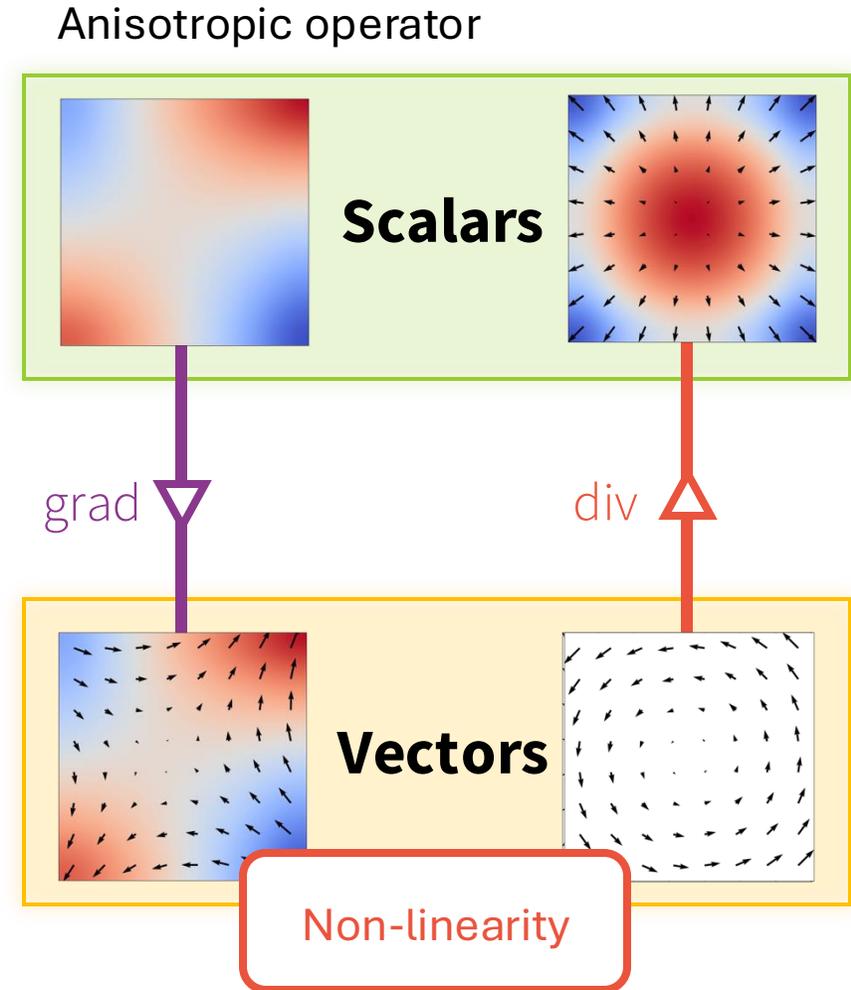
CNNs for 3D

- Graph- and point based
 - GCN, PointNet++, EdgeConv
- 3D kernel (extrinsic)
 - KPConv, MinkowskiNet, SSCN
- 2D kernels *on* surfaces (intrinsic)
 - GCNN, ACNN, MoNet, MDGCNN, HSN



CNNs for 3D

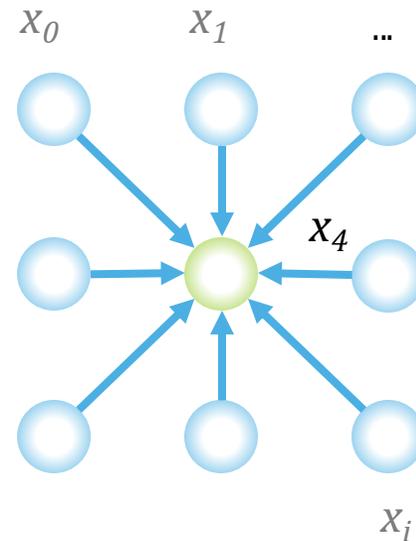
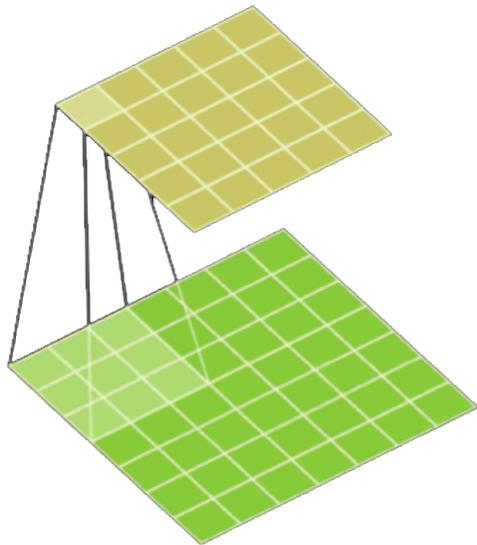
- Graph- and point based
 - GCN, PointNet++, EdgeConv
- 3D kernel (extrinsic)
 - KPConv, MinkowskiNet, SSCN
- 2D kernels *on* surfaces (intrinsic)
 - GCNN, ACNN, MoNet, MDGCNN, HSN
- Operator-based (e.g., Laplacian)
 - DeltaConv



Transformers

Attention on a Graph

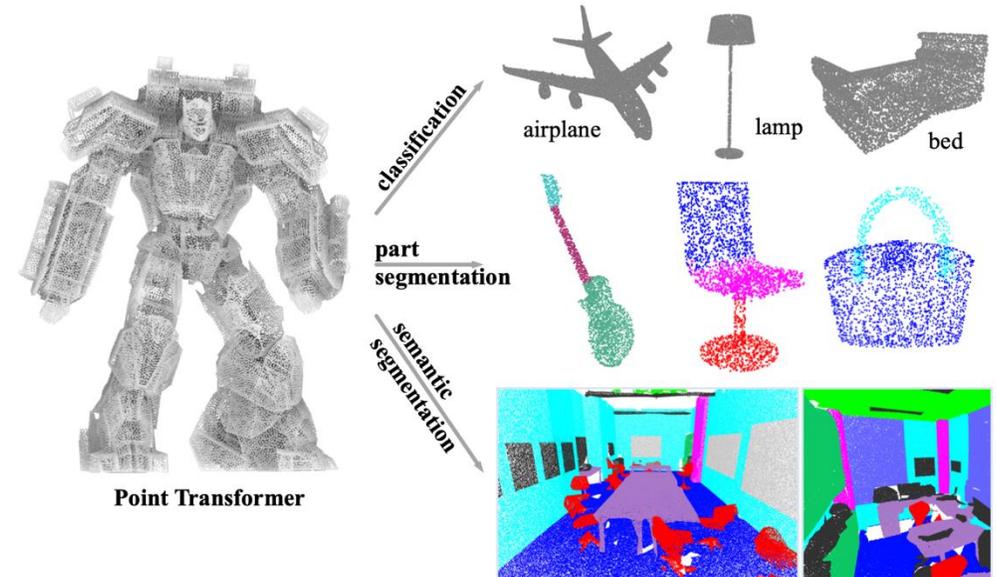
Single CNN layer with 3x3 filter



$$x'_i = \sum_{j=0}^N \text{softmax} \left((W_q x_j)^T W_k x_i \right) W_v x_j$$

Transformers for 3D

- PointTransformer - Zhao et al. 2021
- Challenge: how to reduce cost?
 - Apply on local neighborhoods
 - Apply on serialized point clouds (space-filling curves, v3 2024)
- Challenge: positional encoding
 - PointNet++-like positional encoding, $MLP(p_j - p_i)$
 - Rotation invariance?



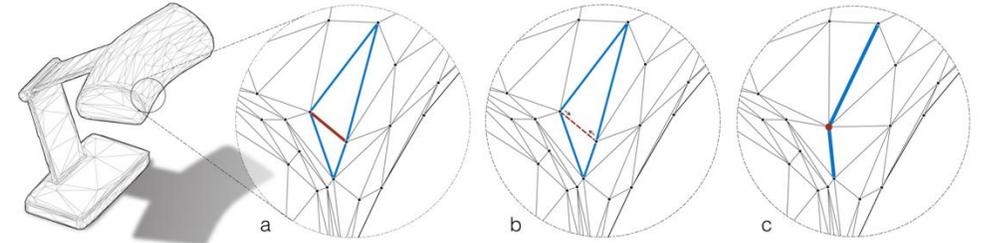
Current state-of-the-art

- Many approaches, *but* PointNet++ still seems to work quite well
 - Due to tasks/benchmarks?
- Other options often seen
 - DGCNN (EdgeConv), PointTransformer v3, Sparse Voxel Convolutions
- Which one to use depends on your task
 - Do you expect the network to understand curvature? Maybe not PointNet++/DGCNN?
 - Do you want efficiency, simplicity? – Try the simple networks first.

Methods you should know, but we didn't cover

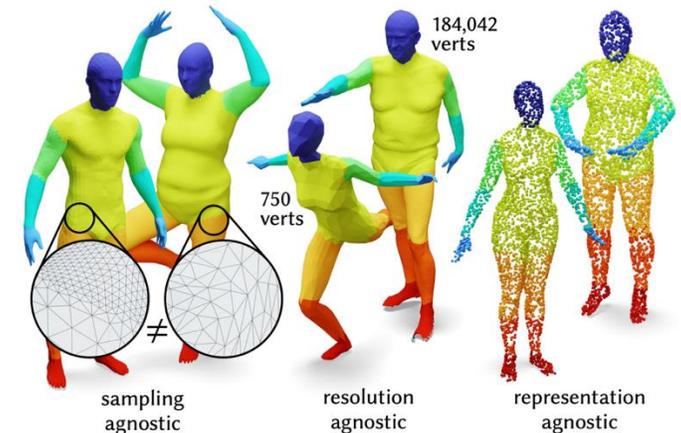
- **MeshCNN** (Hanocka et al. 2019)

- Operates on mesh edges
- Geometric features
 - dihedral angle, two inner angles and two edge-length ratios for each face



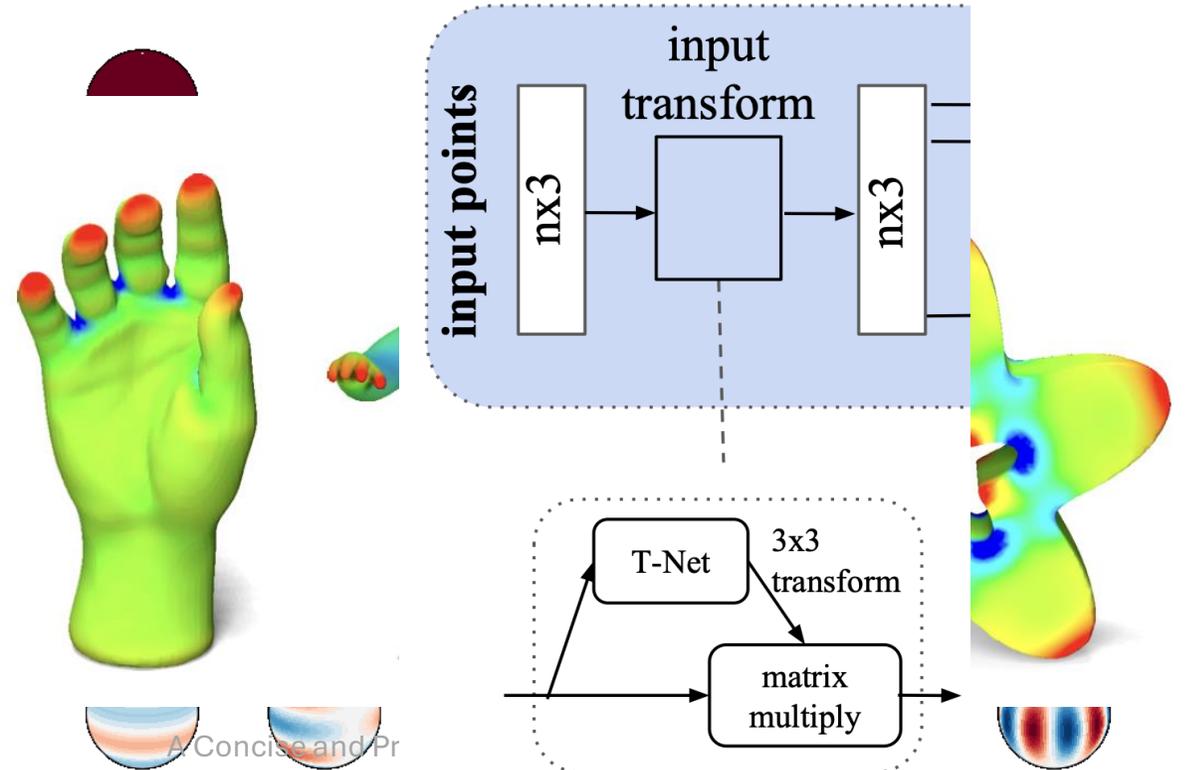
- **DiffusionNet** (Sharp et al. 2022)

- Learning to diffuse
- Accelerated in frequency domain (Eigenvectors of Laplacian)
- Robust to discretizations (e.g., mesh \rightarrow point cloud)



Rotation invariance, translation invariance?

- Data augmentation
 - Random rotation/translation
- ‘Learn’ the transformation
 - Spatial transformer, e.g., in PointNet
- Invariant input features
 - Heat-Kernel Signatures (DiffusionNet)
 - MeshCNN – Edges, intrinsic features
- Encode features in local frames
 - Point difference, normals (translation-invariant)
 - Distances (translation-rotation-invariant)
- Rotation-equivariance
 - E.g., Tensor Field Networks (Thomas et al. 2018), Spherical Harmonics (Poulenard et al. 2019)



A Concise and Practical Introduction to Spherical Harmonics
Based on Heat Diffusion, Gao, Ovsjanikov, Guibas (2009)
PointNet: Deep Learning on Point Sets for 3D Classification
and Segmentation, Qi, Su, Mo, Guibas (2016)
github.com/nvidia/torch-harmonics

Chapter 4: Hands-on with PyTorch Geometric

PyTorch Geometric

<https://pytorch-geometric.readthedocs.io/en/latest/>

- Implementation of message-passing paradigm
- Supports many convolution types
- Many helpful utilities

- Alternatives
 - Kaolin, PyTorch3D, GraphGym



Setting up the environment

Dataset, loading and visualization

Model

Training

Deep Learning on Meshes and Point Clouds

Ruben Wiersma

Graduate School - SGP 2025, Bilbao