MARIE CURIE ACTIONS

SAGIVTECH

IMVC 2024

Australian National University

TECHNION Israel Institute of Technology

OG-INR:

Octree Guided Unoriented Surface Reconstruction

Chamin Hewa Koneputugodage, Yizhak Ben-Shabat (Itzik), Stephen Gould



Problem statement: Surface Reconstruction

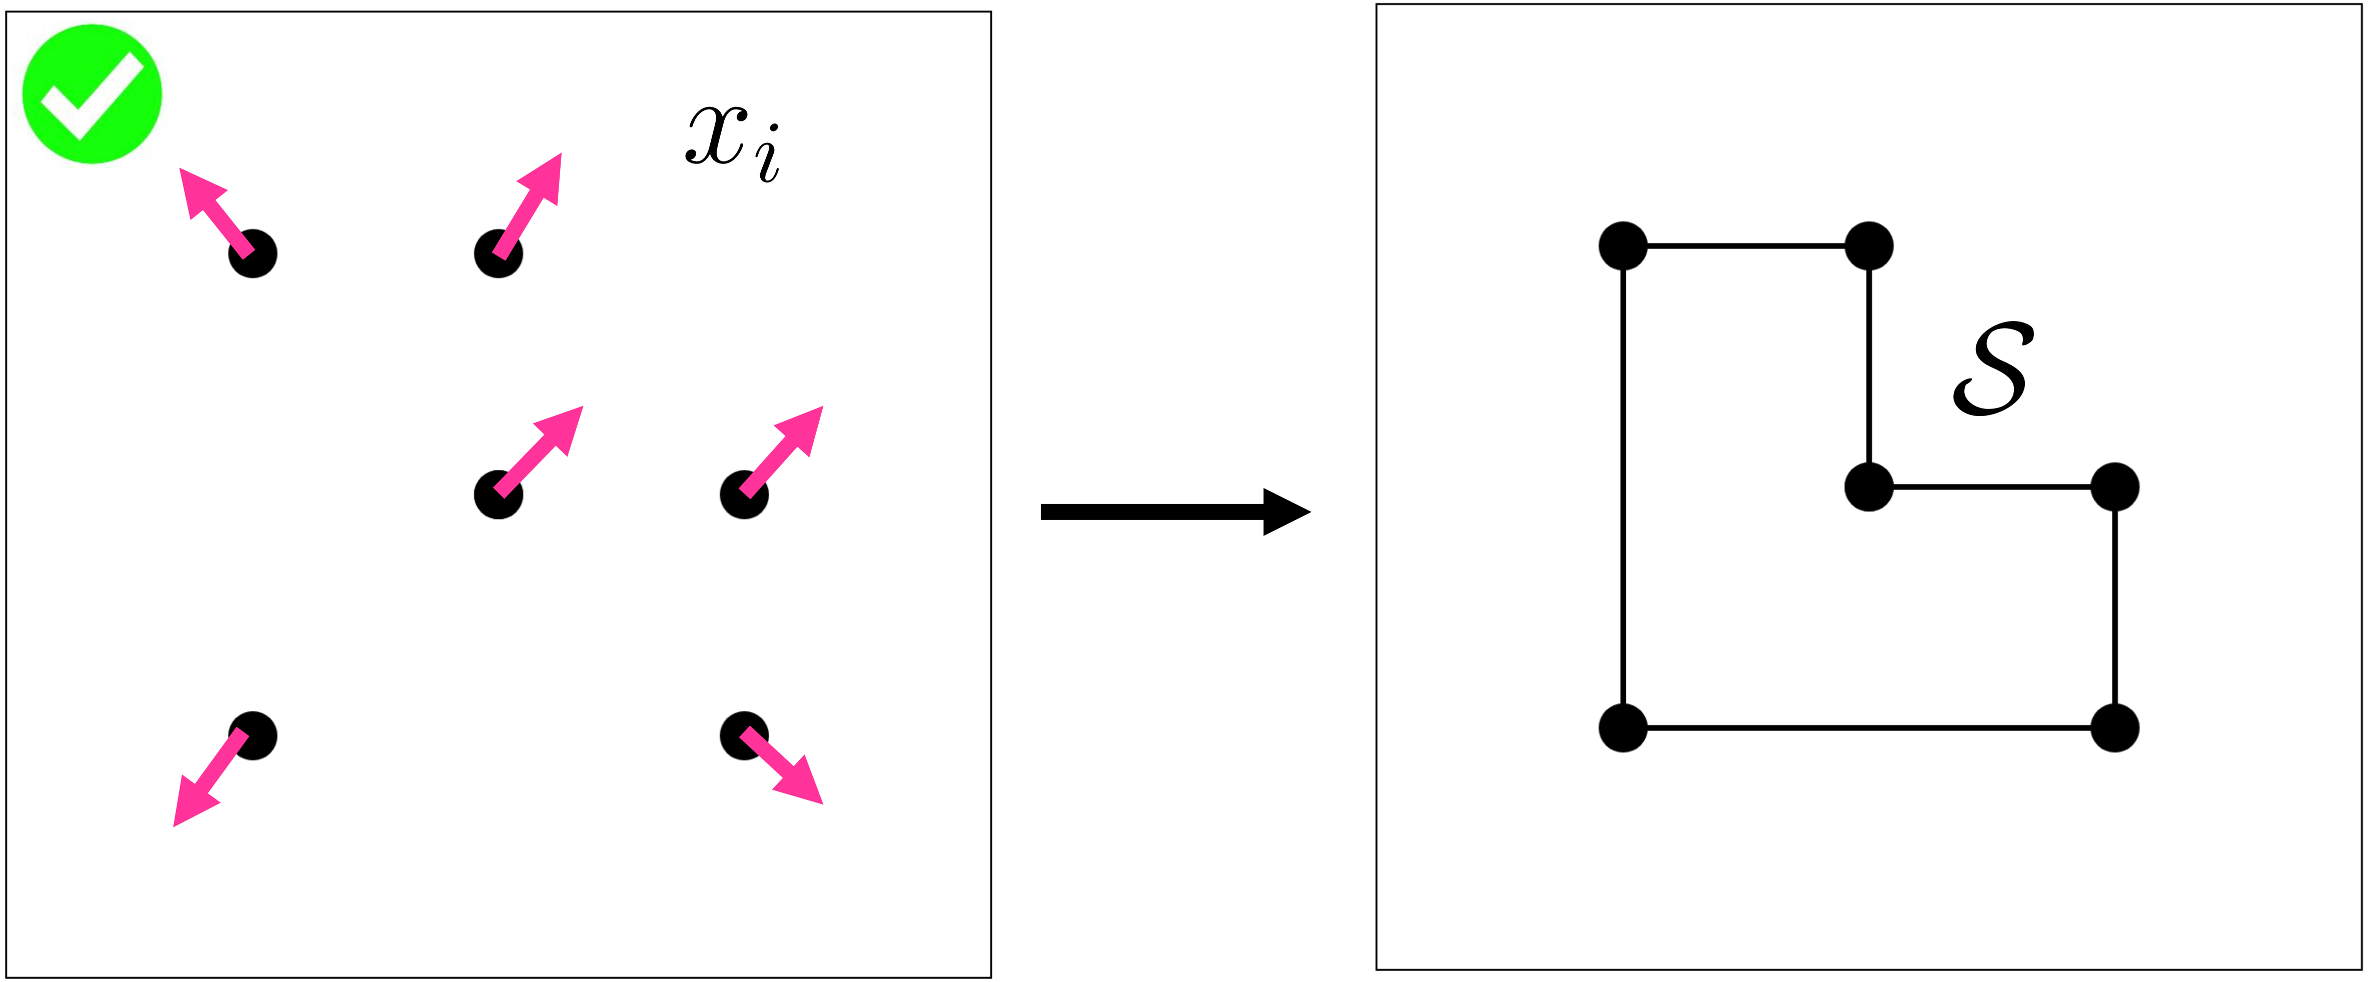

Input

Output

$$\mathcal{S} = \{x \mid \Phi(x) = 0\}$$

Introduction – Implicit shape representation

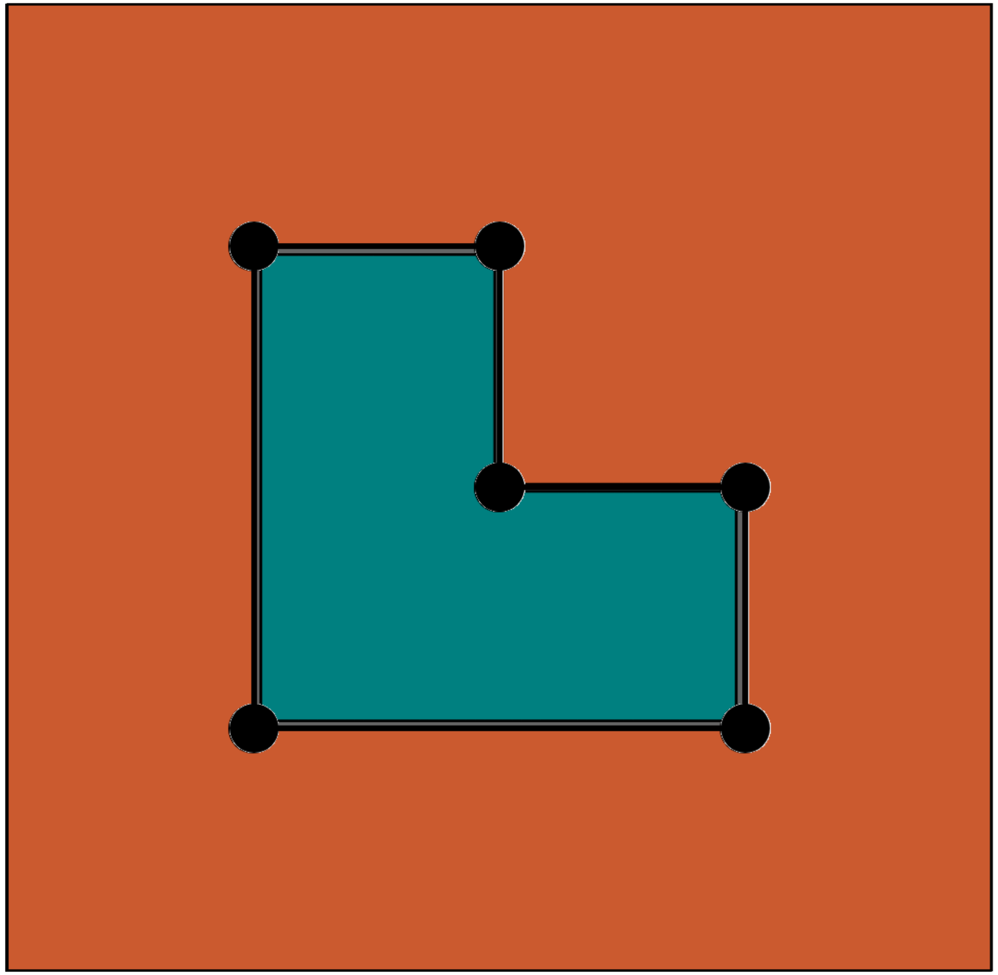

Indicator / Occupancy

$$\Phi_{ind}(x) = \begin{cases} 1 & x \text{ inside} \\ 0 & x \text{ outside} \end{cases}$$

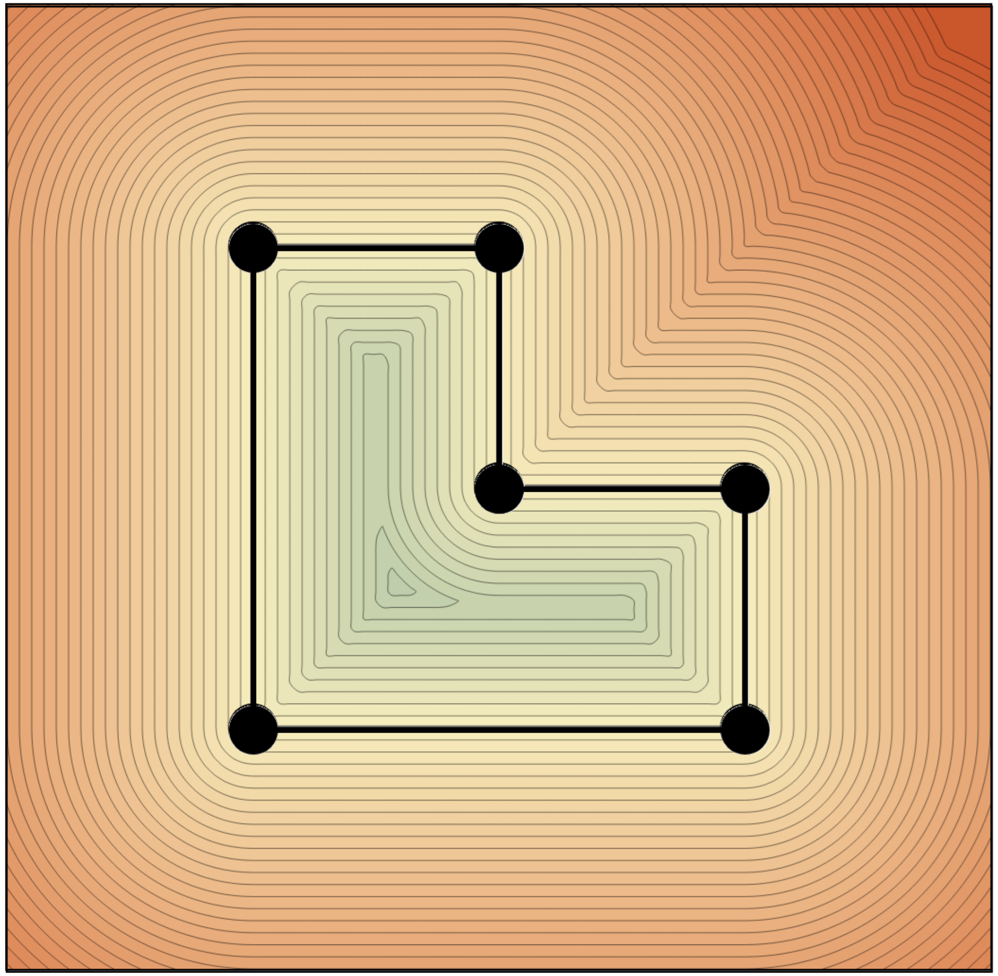

Signed Distance Function

$$\Phi_{sdf}(x) = (-1)^{\Phi_{ind}} \min_{z \in \mathcal{S}} \|x - z\|_2$$

Previous work – NN methods

Can we learn it? Yes!

DeepSDF (Park et al. 2019), Occupancy networks (Mescheder et. al. 2019), Chen & Zhang 2019

Is there a better way to train it? Yes!

SAL, SALD (Atzmon et al. 2019), IGR (Gropp et. al. 2020), PHASE (Lipman 2021) , DiGS (Ben-Shabat et. al. 2022)

Why use ReLUs? Use sinusoidal functions!

SIREN (Sitzmann et al. 2020)

Implicit Neural Representation overview

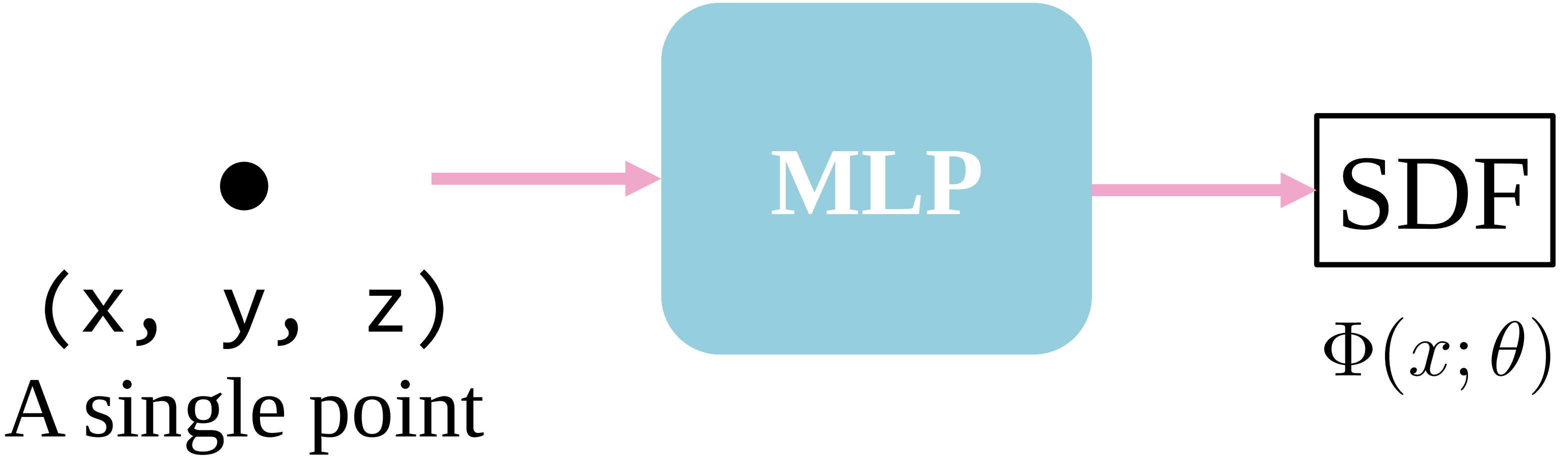

$$\Phi(x;\theta) = \mathbf{w}_n^T \left(\phi_{n-1} \circ \phi_{n-2} \circ \ldots \circ \phi_0\right)(x) + \mathbf{b}_n, \quad x_i \mapsto \phi_i(x_i) = \sin\left(\mathbf{W}_i x_i + \mathbf{b}_i\right)$$

Implicit Neural Representation overview

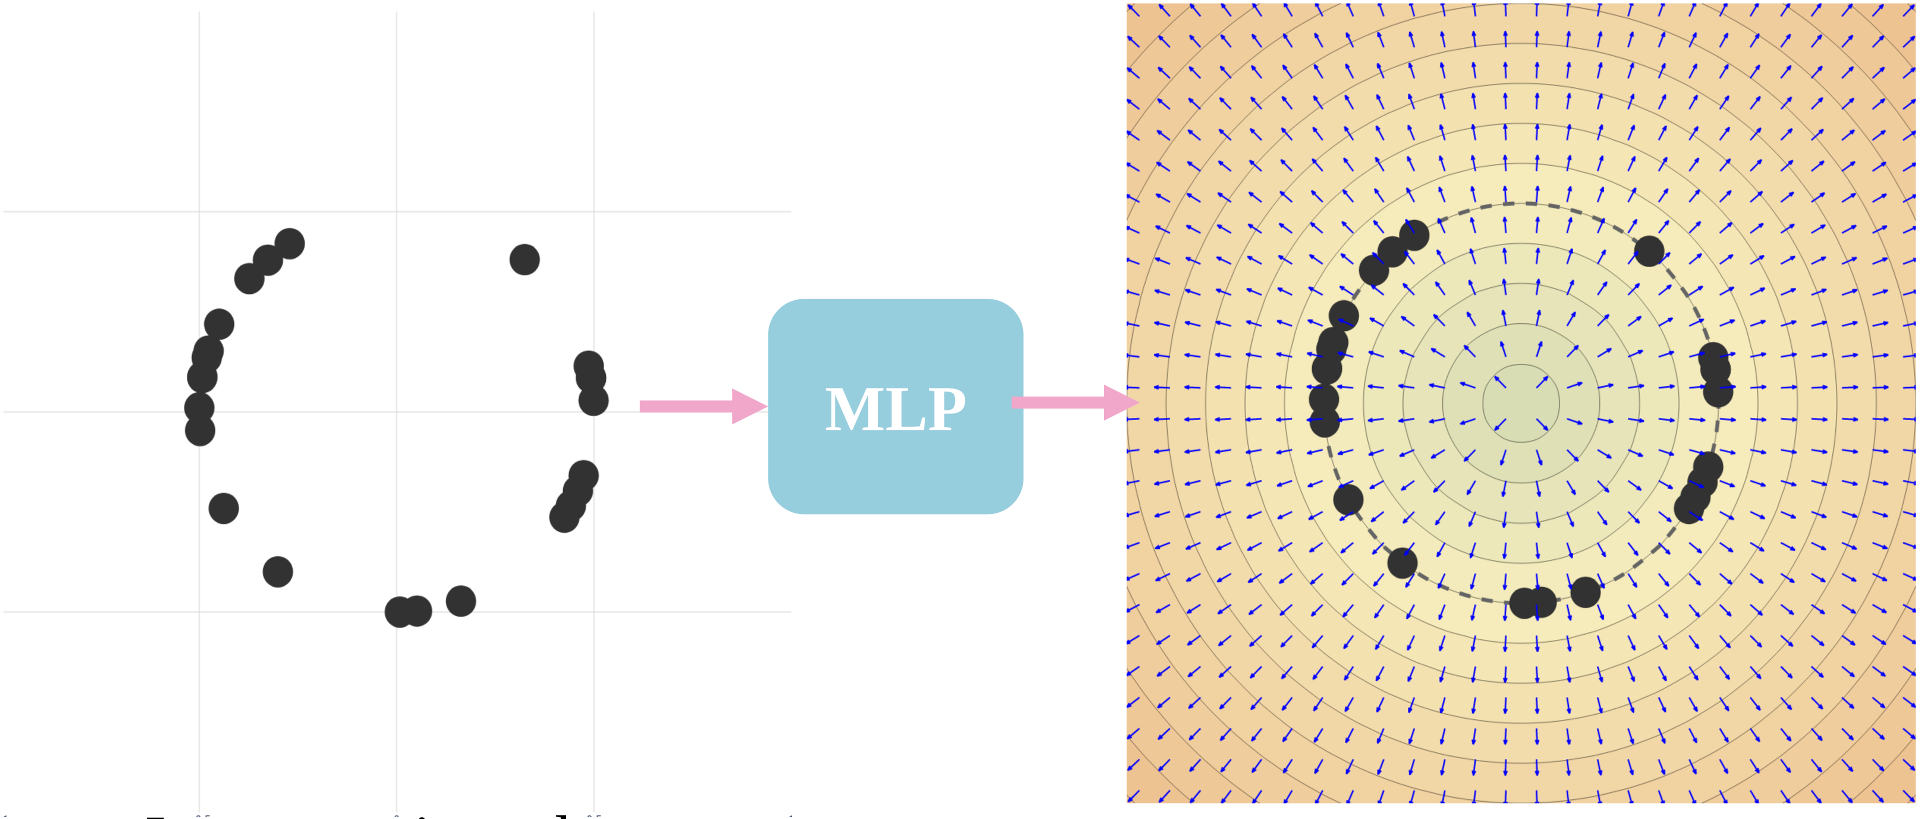

Input unoriented
Point Cloud

SDF, Gradients

$$\Phi(x;\theta), \nabla\Phi(x;\theta)$$

Can we prevent “ghost geometries”?

Spoiler alert: YES!!!
(by figuring out inside / outside regions?)

Key Idea – Construct Octree, label it, and train INR

1. Constructs an **octree** and label **inside and outside**

+

2. Optimizes an **INR** guided by the **octree's labels**

=

Better & Faster surface reconstruction

Octree Guided INR - Approach

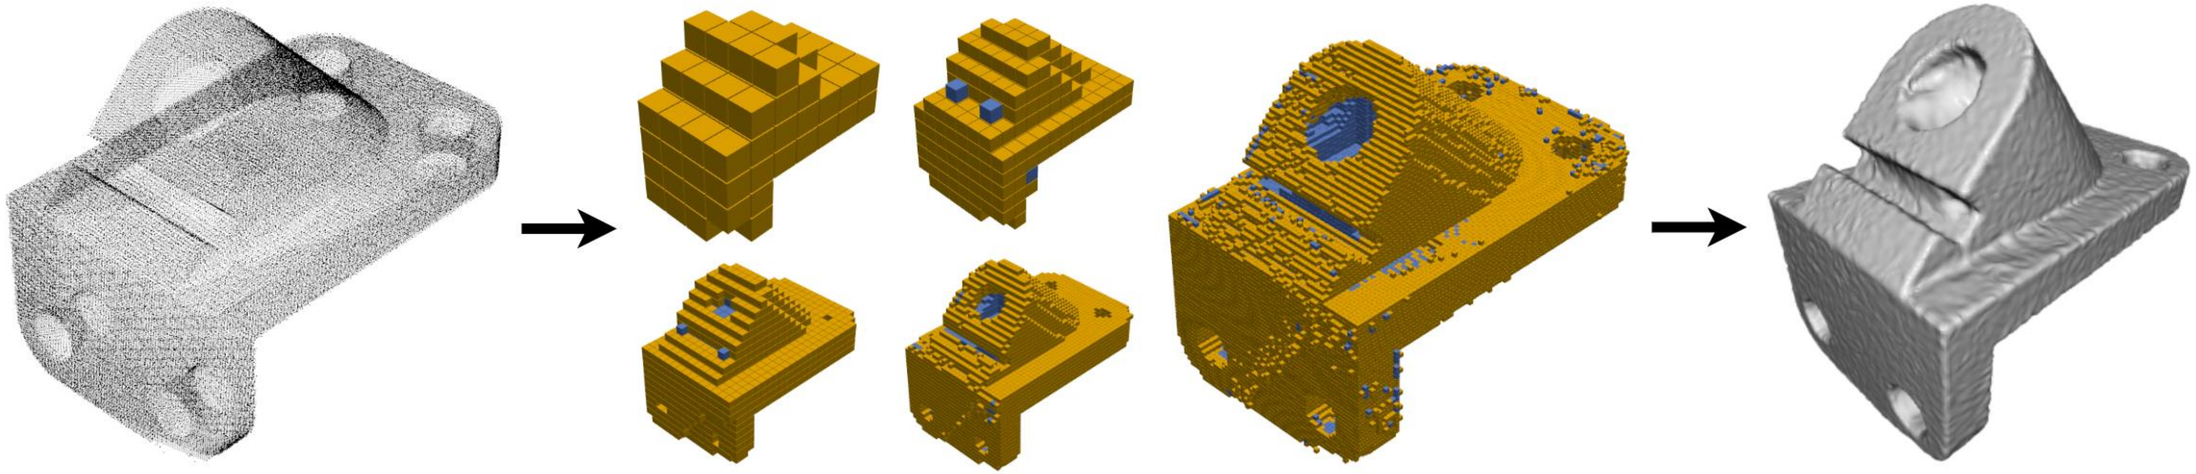

Raw unoriented point cloud

- Build octree to initial depth d_i (surface nodes yellow)
- Label other nodes as inside (blue) or outside (transparent)
- Then continue expanding surface leaves and labelling till final depth d_f

- Use labels as supervision for an INR to learn the shape continuously

Using this as supervision for INR training

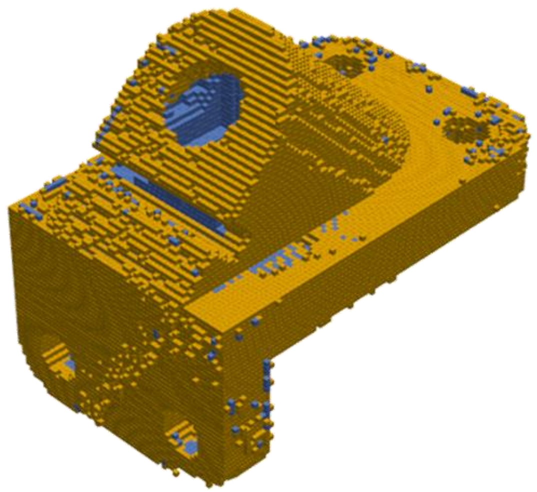

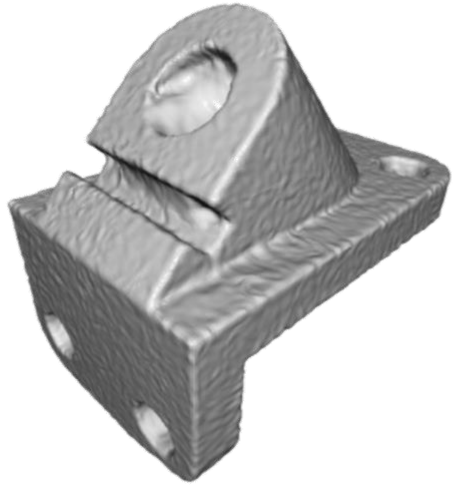

$$\mathcal{L}(\Phi) = \lambda_1 \mathcal{L}_1 + \lambda_2 \mathcal{L}_2 + \boxed{\lambda_3 \mathcal{L}_3 + \lambda_4 \mathcal{L}_4}$$

Standard INR Losses

$$\mathcal{L}_1 = \sum_{x \in \chi} |\Phi(x, \theta)|$$

$$\mathcal{L}_2 = \int_{\mathcal{D}} \left| \|\nabla_x \Phi(x, \theta)\|_2 - 1 \right| \mathrm{d}x$$

Octree Guidance

Approximate SDF Loss

$$\mathcal{L}_3 = \int_{\mathcal{D}} |\Phi(x, \theta) - \tilde{d}_s(x)| \, \mathrm{d}x$$

(Pred SDF should match approx. SDF: sign from octree, distance is to χ)

Octree Label Loss

$$\mathcal{L}_4 = \sum_{\ell \in \mathcal{L}} \left(\sum_{x \in S_\ell} \left[(-1)^{(1-y'_\ell)} \Phi(x, \theta) \right]_{>0} \right)$$

(Sample in each leaf, penalize if pred sign ≠ octree sign)

OG-INR Qualitative Results

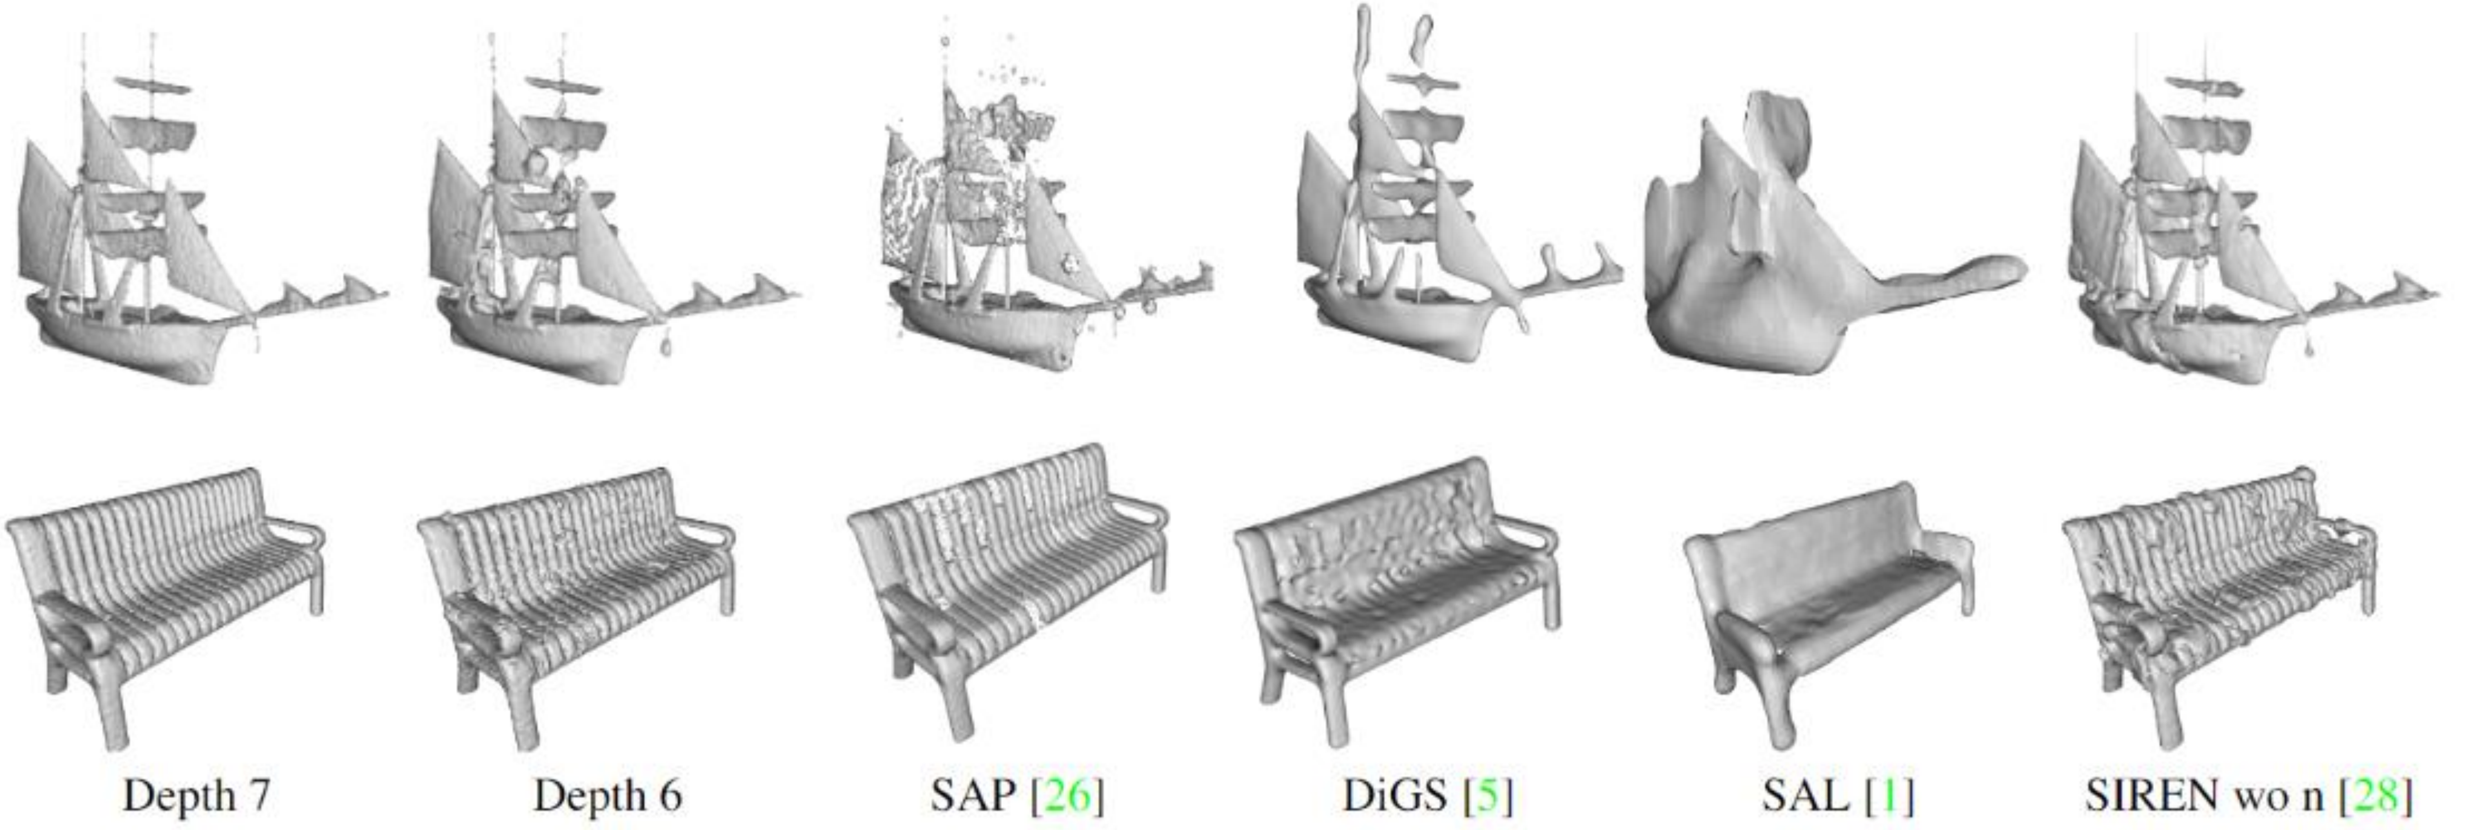

OG-INR Qualitative Results

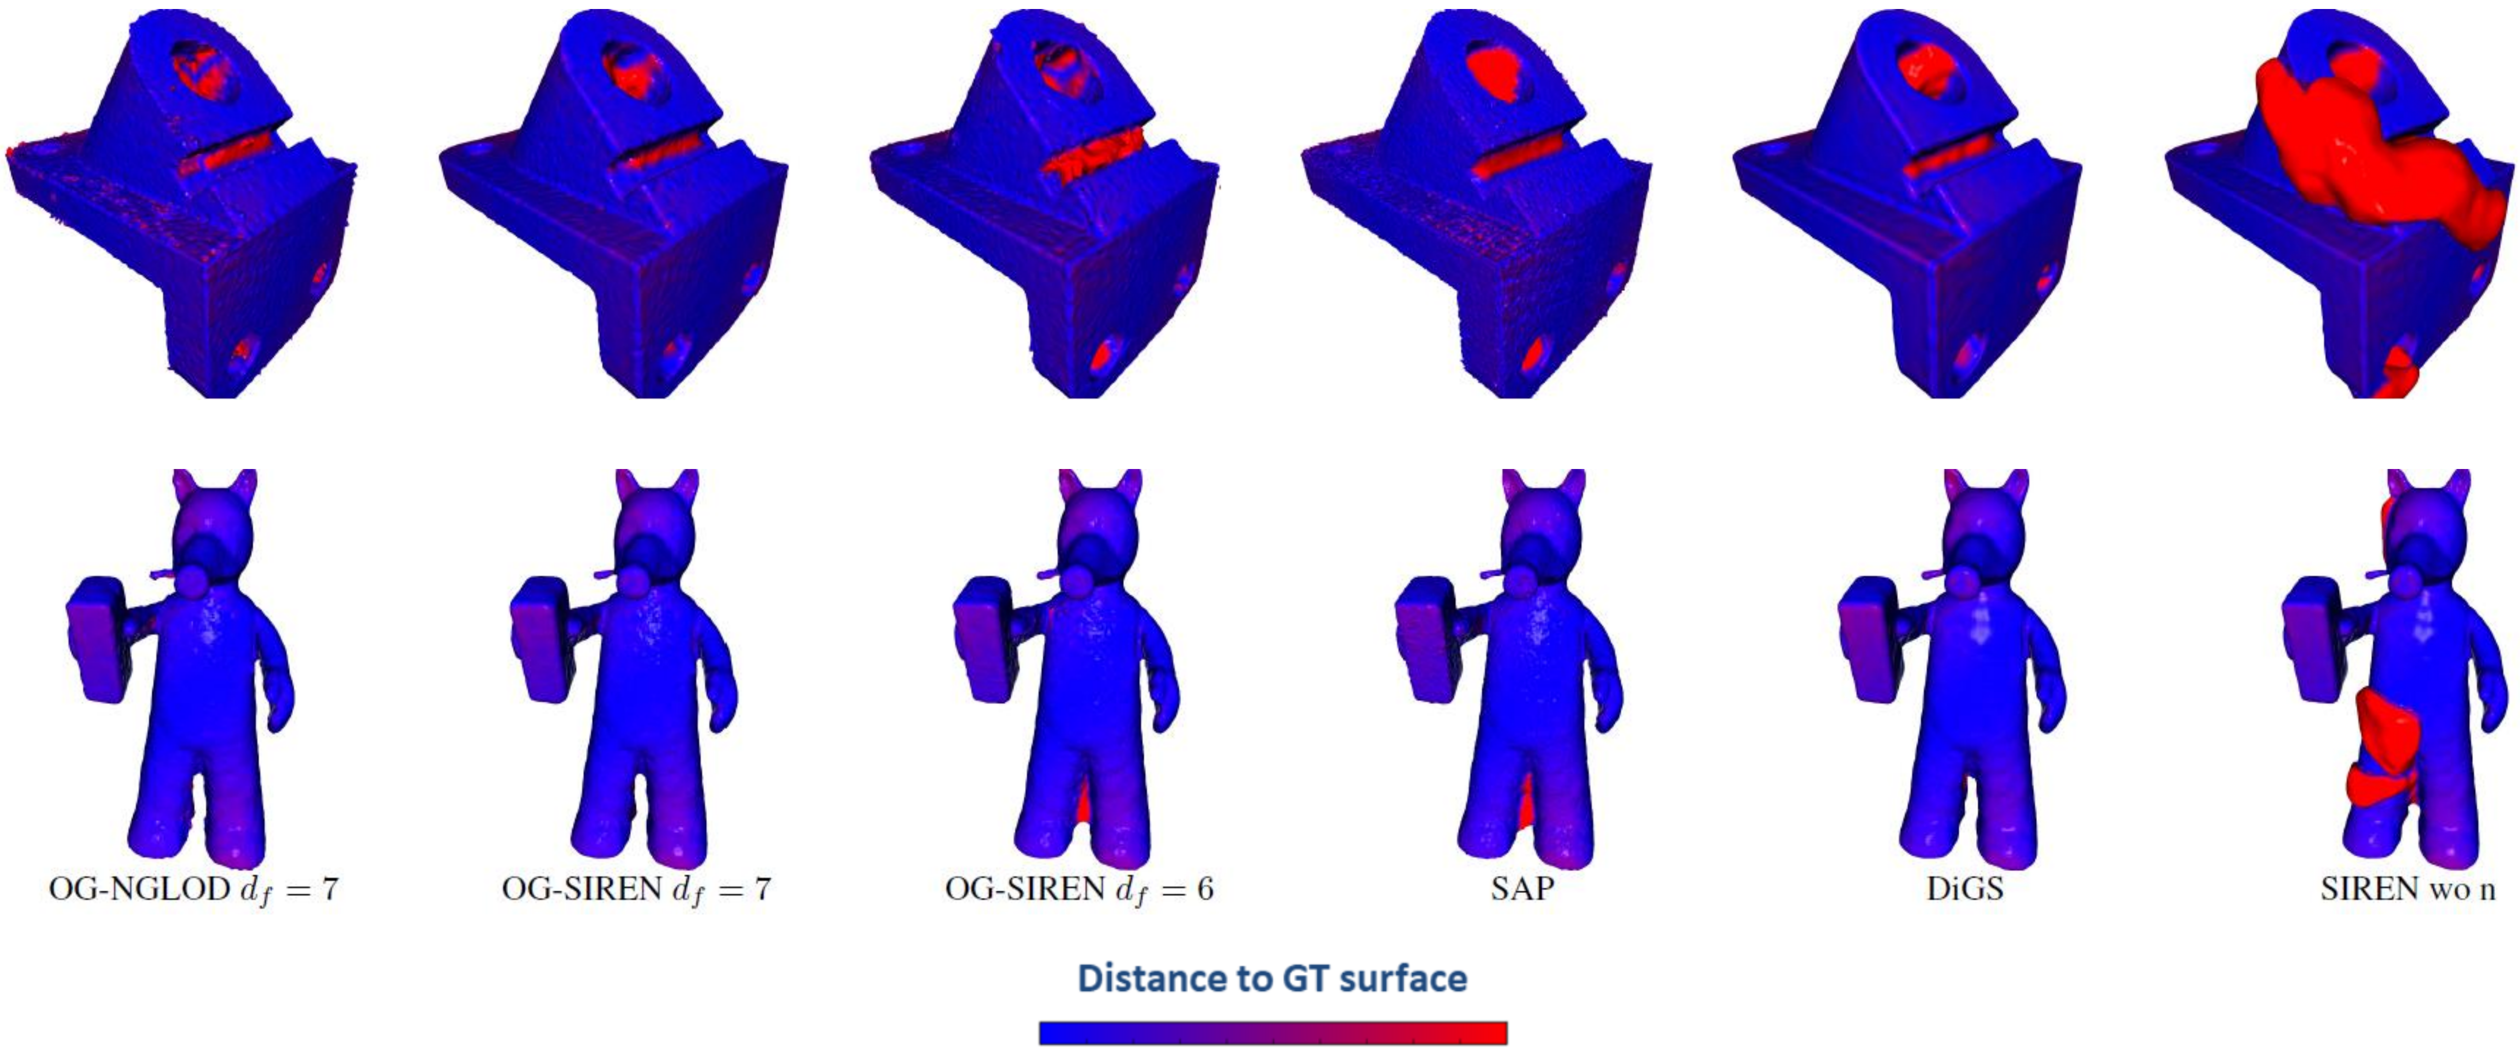

Octree guidance gives a significant speedup

We get a **5.5x speedup** on INR training over SIREN, making an overall speedup of 3.9x in the full method

Method	Parameters	Time per iter. (s)	Num iters	Time (s)	Speed Up
N Est. +SPSR	-	-	-	42 (13 + 29)	$12\times (\cdot + 18\times)$
SIREN (wo n)	264K	0.052	10000	520	$1\times$
SAL	2.1M	0.175	10000	1750	$0.3\times$
DiGS	264K	0.120	10000	1200	$0.4\times$
SAP	120K	-	3200	330	$1.6\times$
Our OG-SIREN	264K	0.158	600	135 (40 + 95)	$3.9\times (\cdot + 5.5\times)$
Our OG-NGLOD	68.7M	0.173	300	**92 (40 + 52)**	$\mathbf{5.7\times (\cdot + 10\times)}$

Contribution

- Propose **OG-INR**, which
  1. Constructs an **octree** and labels **inside and outside**
  2. Optimizes an **INR** guided by the **octree's labelling**
- Propose an **energy function** over the octree and provide an **efficient move-making algorithm** that explores many possible labelling options.
- This **avoids many local minima** that SGD gets stuck in

Talking Papers Podcast

More episodes coming out soon!

https://talking.papers.podcast.itzikbs.com/

Acknowledgements

- The European Union’s Horizon 2020 research and innovation programme under the **Marie Sklodowska-Curie grant** agreement No 893465

- **NVIDIA** hardware grant

OG-INR:

Octree Guided Unoriented Surface Reconstruction

Chamin Hewa Koneputugodage, Yizhak Ben-Shabat (Itzik), Stephen Gould

https://chumbyte.github.io/OG-INR-Site/

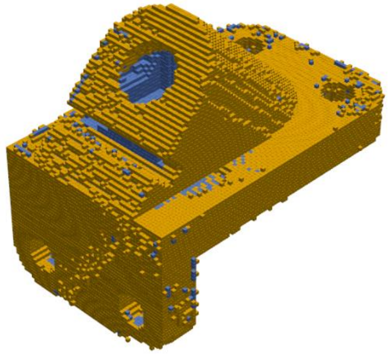

@sitzikbs

sitzikbs

Itzikbs.com

Octree Guided INR – Move making algorithm

Initialize -> Repeatedly grow, and make moves that minimize

$$E(\boldsymbol{y}) = \sum_{j \in \mathcal{L}_s} E_j^{SP}(\boldsymbol{y}) + \lambda \sum_{i,j \in \mathcal{L}_{ns}} E_{i,j}^{MSC}(\boldsymbol{y})$$

Surface Property + **Minimal Surface Constraint**

Surface nodes should be close to an inside and an outside node

Surface area between inside (**Blue**) and outside (**Transparent**) nodes should be minimised