SHARC

*Reference-point-driven **S**pherical **H**armonic **R**epresentation for **C**omplex Shapes*

Panagiotis Sapoutzoglou · George Terzakis · Maria Pateraki

National Technical University of Athens, Greece

{psapoutzoglou,gterzakis,mpateraki}@mail.ntua.gr

pose-lab.github.io/SHARC/

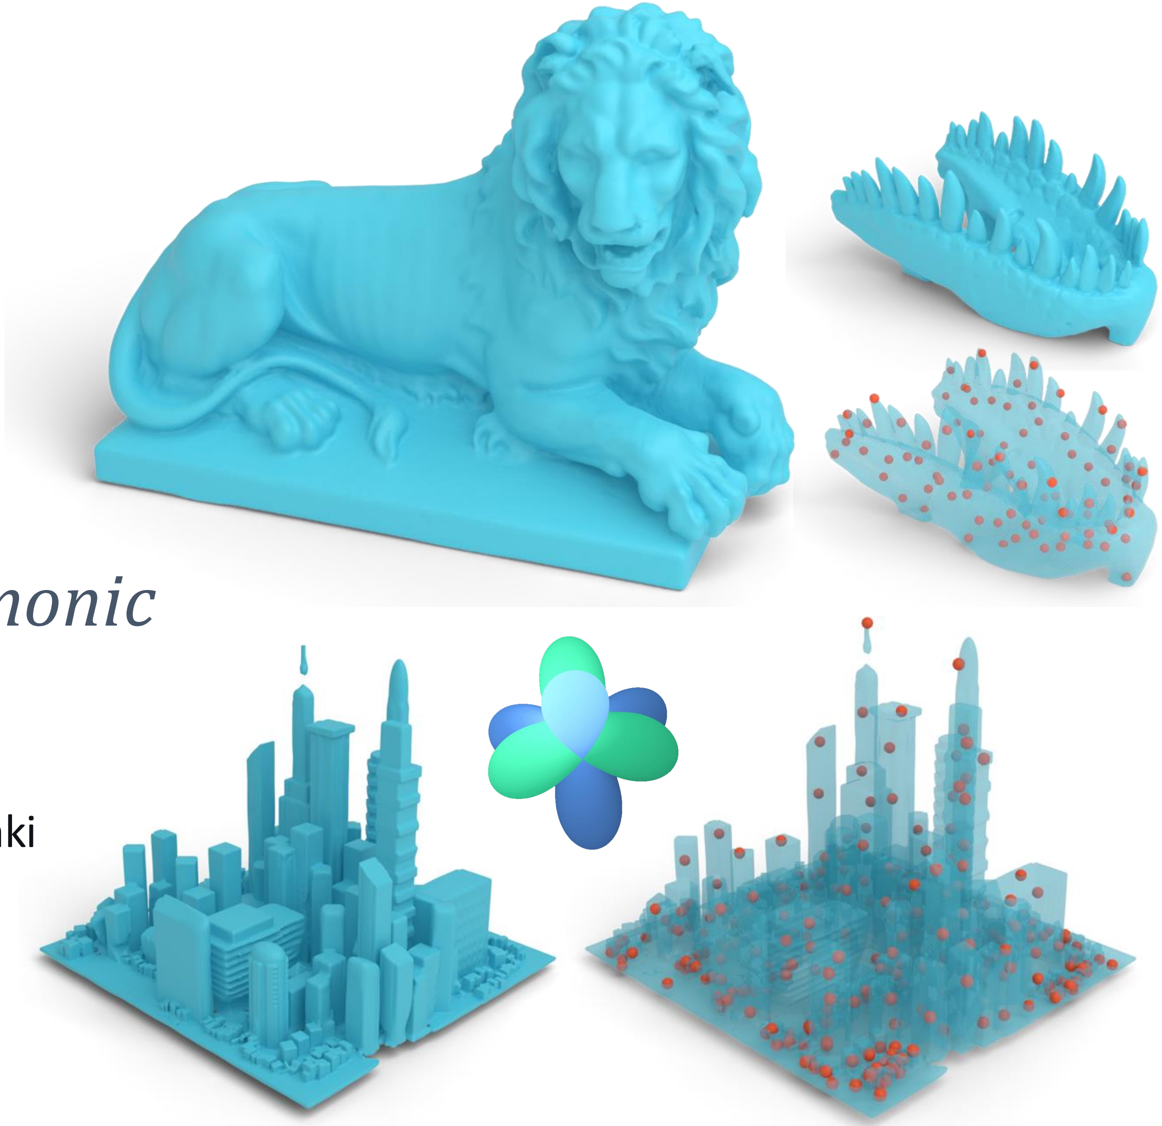

THE PROBLEM

One representation, three demands

Encoding complex 3D geometry that is at once compact, detail-preserving, and valid for any topology remains a central open challenge across graphics, geometry processing, and 3D generative modeling.

1. **Compact**
   small to store, transmit & optimize
2. **High-fidelity**
   fine, high-frequency detail preserved
3. **Genus-agnostic**
   any topology, holes and handles

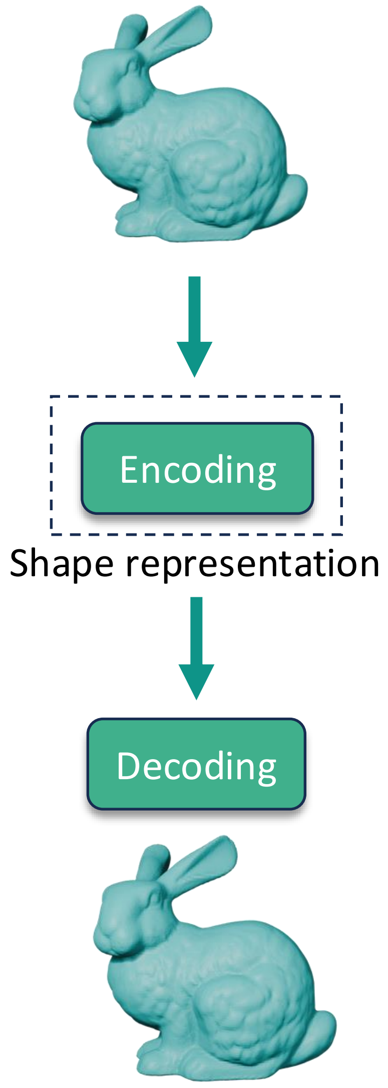

PRIOR WORK

How shapes are represented today

Explicit

Meshes · Point clouds

Detail equals size. Complexity scales with geometry, making them costly to store, transmit and optimize.

- ***Cost grows with detail***

Neural implicit

Occupancy Nets (e.g. DeepSDF[1], 3DShape2VecSet[2])

Compact, continuous fields from MLPs, but parameters are opaque and inference latency is often high.

- ***Require training, hard to interpret***

Medial / skeletal

MAT · CoverageAxis++[3]

A structural scaffold from unions of spheres; simple primitives struggle with complex geometry and boundary noise.

- ***Primitive-limited***

[1] Park, J.J., Florence, P., Straub, J., et al.: Deepsdf: Learning continuous signed distance functions for shape representation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
[2] Zhang, B., Tang, J., Nießner, M., Wonka, P.: 3DShape2VecSet: A 3D shape representation for neural fields and generative diffusion models. ACM Trans. Graph.(2023)
[3] Wang, Z., Dou, Z., Xu, R., Lin, C., et al.: Coverage axis++: Efficient inner point selection for 3d shape skeletonization. Computer Graphics Forum (2024)

Why Spherical Harmonics?

Among compact, analytic representations, spherical harmonics stand out: a handful of coefficients capture smooth geometry, frequency bands separate coarse shape from fine detail, and reconstruction is a closed-form inverse transform.

1. **Natural basis** - Distance fields ***r(ϑ,φ)*** *are functions on the sphere.*
2. **Frequency-structured** - Low degrees carry global shape, high degrees carry fine detail (built-in multiresolution).
3. **Analytic & Training-free** - A deterministic, per-object encoding. No dataset, no network. Closed-form fit and inverse.

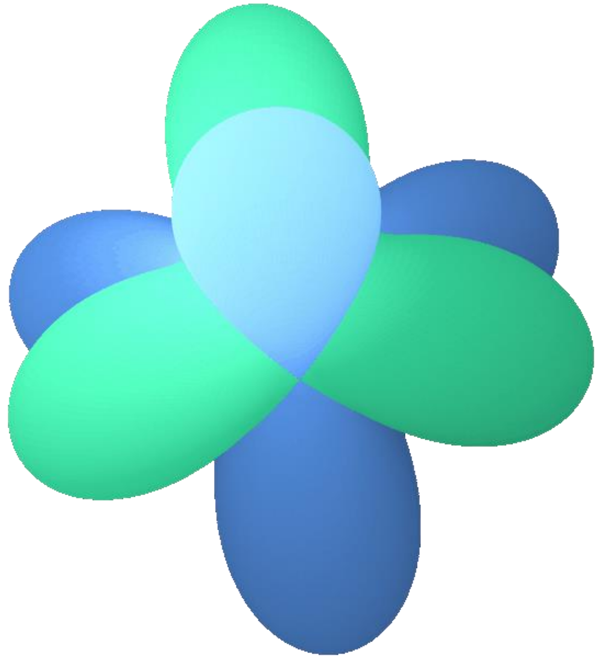

But one Center Isn’t Enough

A radial SH function ***r(ϑ,φ)*** is single-valued per direction so one center can only encode **star-shaped** objects. Any concavity, hole or handle makes the field multi-valued.

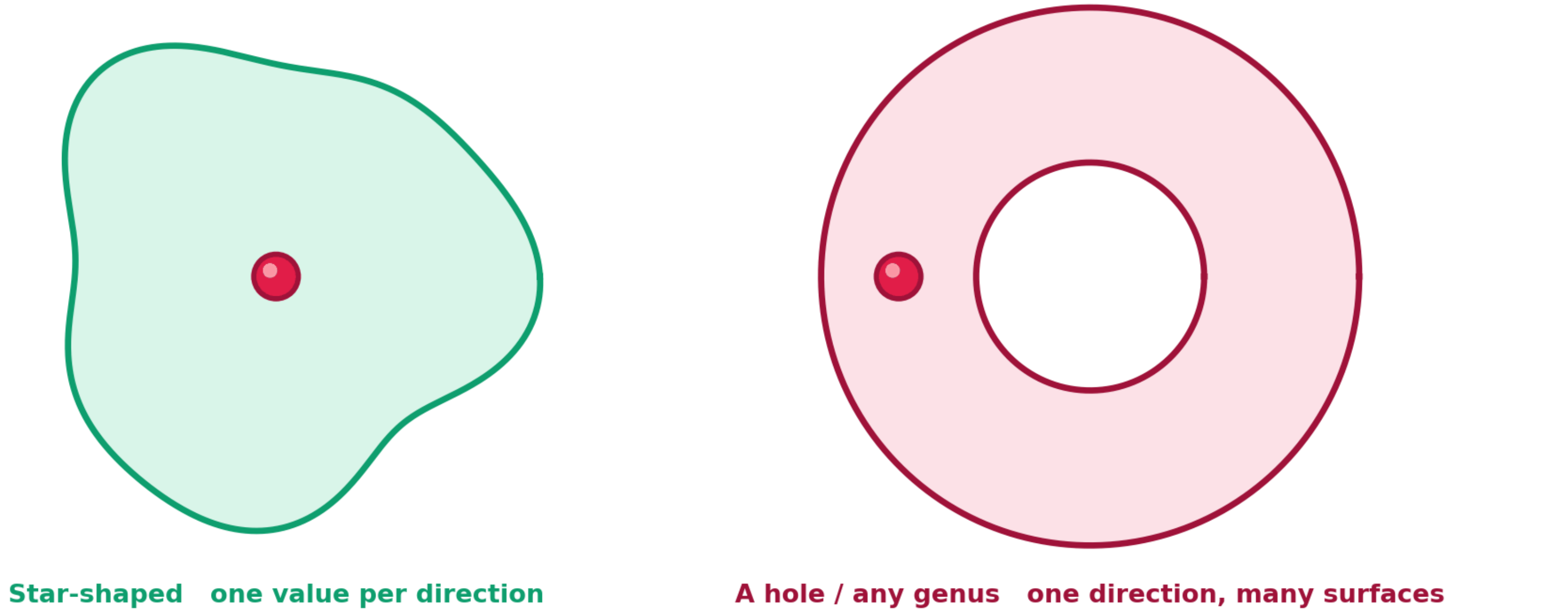

Where Should Reference Points Live?

Local star-shaped patches

Cut the object into many star-shaped sub-parts so SH applies. Segmentation is done using fixed 3D grids into star-shaped sub-parts ignoring the object's natural structure.

MASH[4] (2025) – Exterior anchors

Masked, anchored SH fields. But a multi-view patching strategy with generalized view cones needs many anchors and costly iterative optimization, and blends overlapping patches softening details.

The gap: capture high-frequency detail on any topology, with fewer primitives.

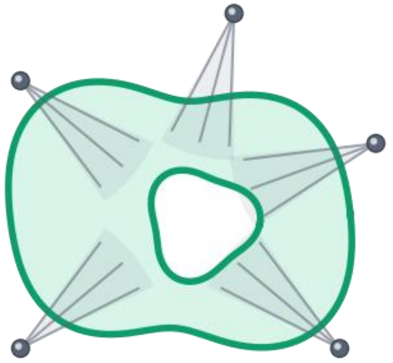

MASH — exterior anchors

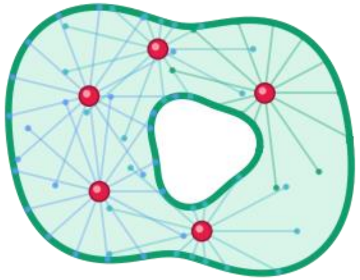

SHARC — interior reference points

[4]Li, C., Xin, Y., Zhou, X., Shamir, A., et al.: Mash: Masked anchored spherical distances for 3d shape representation and generation. ACM SIGGRAPH 1 (2025)

OUR APPROACH

Encode distance fields from multiple interior reference points

Place a sparse set of reference points inside the object volume. From each, encode the radial distance field to the surface with spherical harmonics with visibility modeled explicitly.

1. **Fast** - Encoding and reconstruction run far faster than state-of-the-art methods.
2. **Parsimony & accuracy** - Best reconstruction fidelity using an order-of-magnitude fewer primitives.
3. **Proximity reconstruction** - Nearest-generator filtering resolves overlaps leading to sharper surfaces, no patch blending.

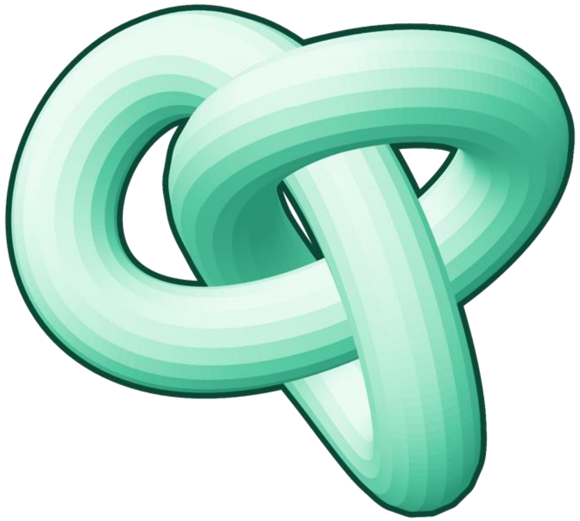

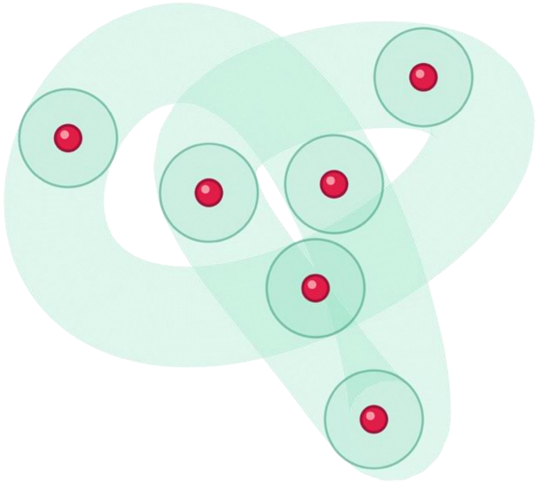

METHOD

The SHARC pipeline

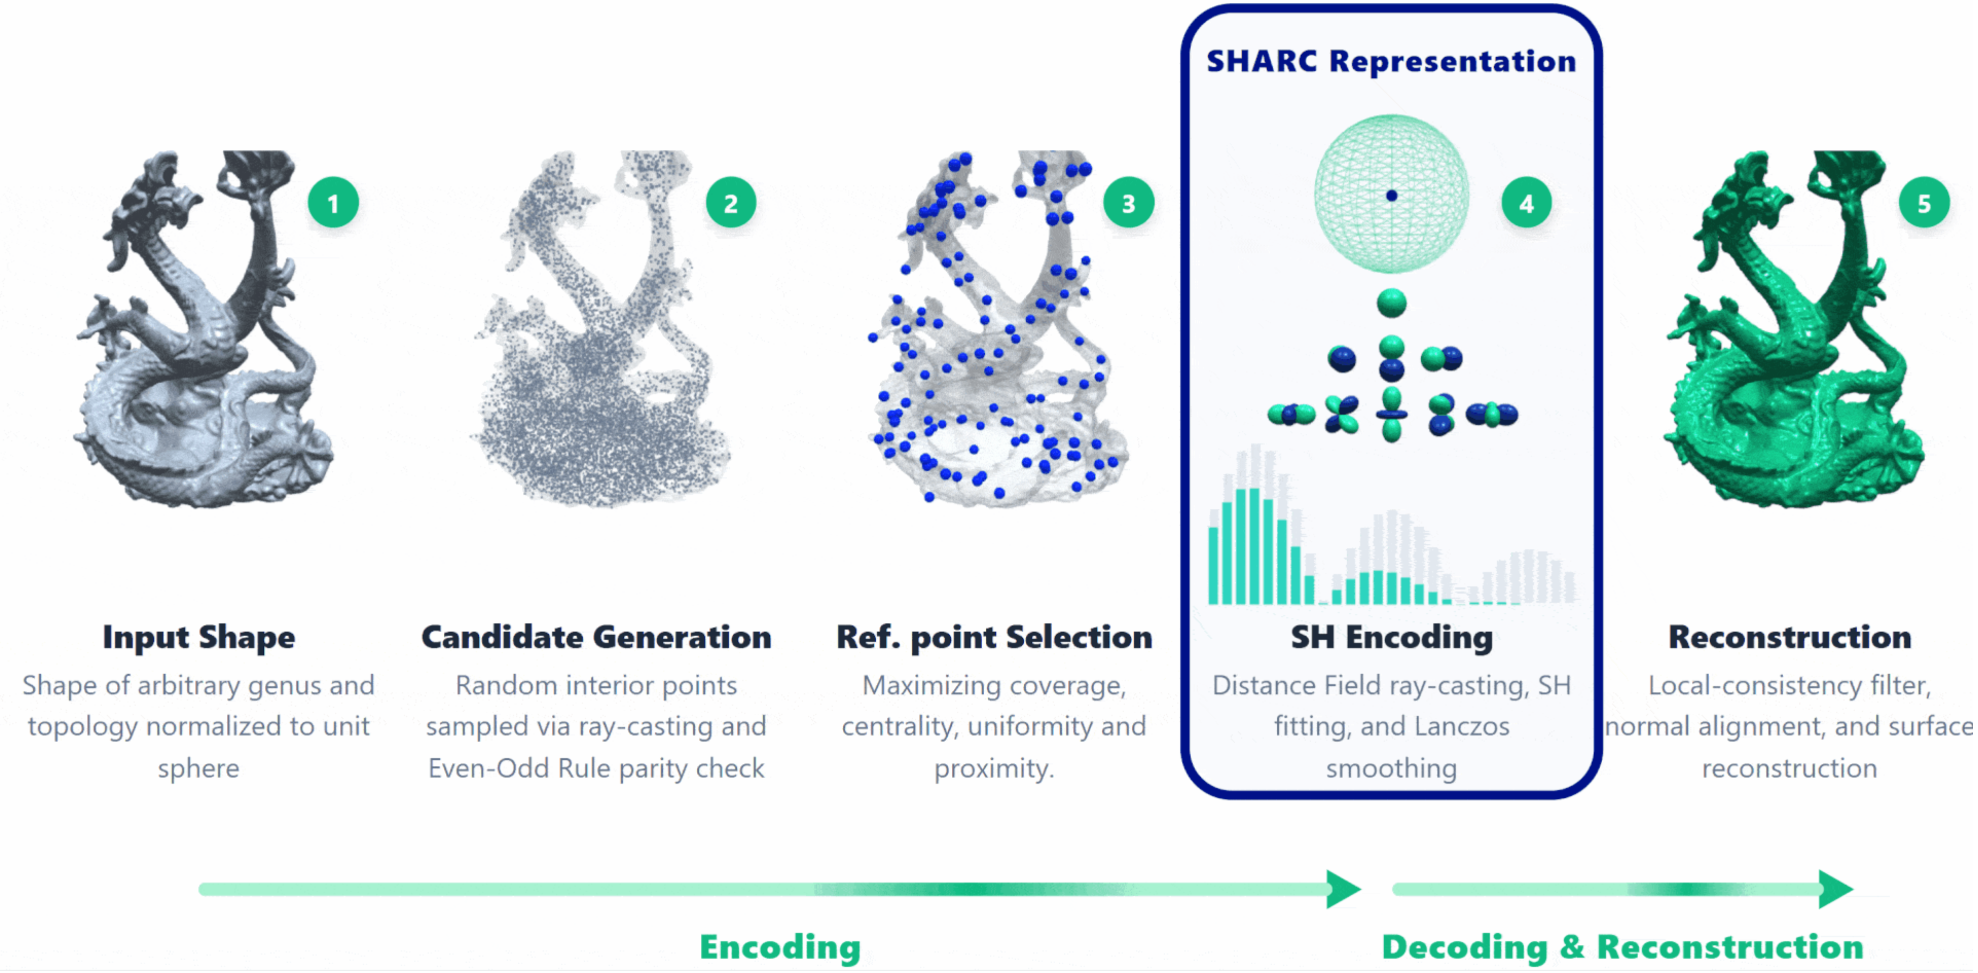

Choosing where to look from

Candidate pool (C). Sample random points inside the object's volume.

Surface points (S). Uniformly sample surface points.

Greedy selection. At each step add the candidate (c ∈ C) that maximizes a standardized score:

H(c) = w·Coverage + w·Centrality + w·Uniformity

- **Coverage** — visibility × proximity to surface samples

$$E_{cov}(\mathbf{c}) = \frac{1}{|\mathcal{S}|}\sum_{\mathbf{s}\in\mathcal{S}} \mathrm{Vis}(\mathbf{s};\mathbf{c})\,\mathrm{Prox}(\mathbf{s};\mathbf{c})$$

- **Centrality** — near the medial axis (inverse medial radius)

$$E_{cen}(\mathbf{c}) = -\frac{1}{\min_{\mathbf{s}\in\mathcal{S}} \|\mathbf{c}-\mathbf{s}\|}$$

- **Uniformity** — spread out from already-chosen anchors

$$E_{uni}(\mathbf{c}) = \min_{\mathbf{p}\in\mathcal{C}_t} \|\mathbf{c}-\mathbf{p}\|$$

A surface point s ∈ S is considered covered if ||c − s|| ≤ τ and s is visible from c.

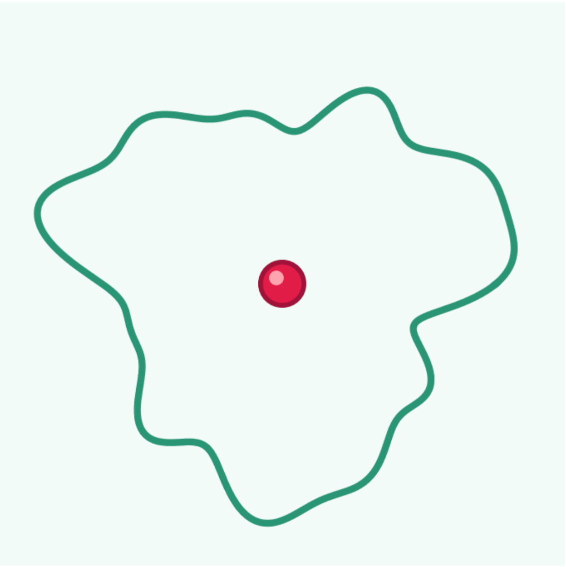

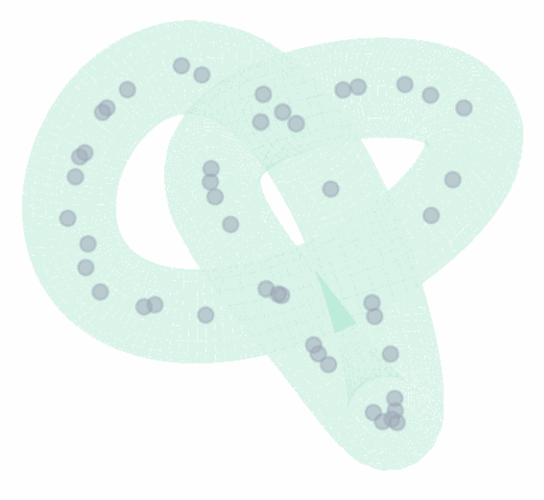

Encoding the distance fields with SH

1

Sample the field

Ray-cast the radial distance r(θ,ϕ) from each reference point to the surface.

2

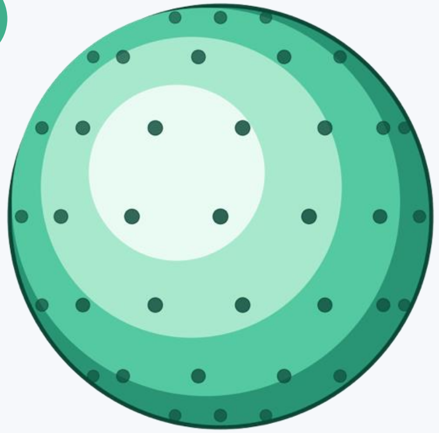

HEALPix grid

Equal-area sampling on the sphere, no polar distortion, uniform sampling.

3

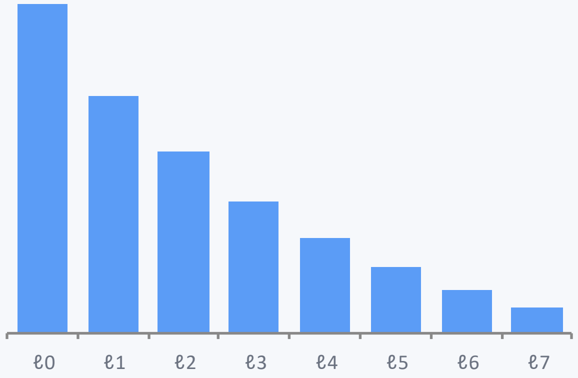

Fit SH — FSHT

Fast Spherical Harmonic Transform up to bandwidth L. We use L = 64.

From coefficients back to surface

1. **Inverse transform** Reconstruct each reference point's distance field from its SH coefficients → dense point cloud.

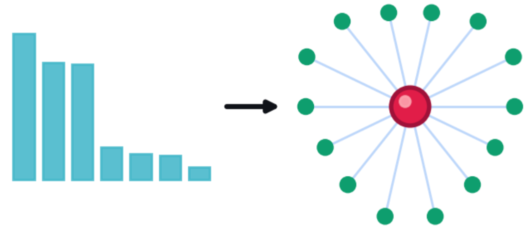

2. **Lanczos smoothing** A σ-window attenuates high frequencies — ringing gone, global shape kept.

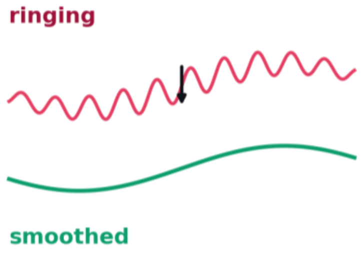

3. **Local generator selection** Enforce locality — each decoded point is kept only by its nearest reference point. This resolves overlaps, sharper than blending.

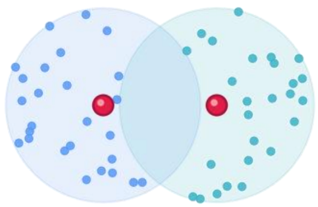

4. **Normals for free** Each point's normal is just its outward ray direction. No costly normal-orientation step before meshing.

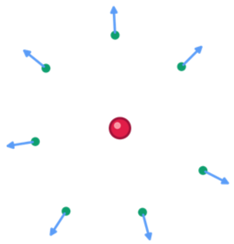

Best fidelity with far fewer primitives

Table 1. Quantitative comparison on the **PSB**, **Stanford**, and **Thingi10k** datasets. We report the primitive count $|C|$, Chamfer Distance (d_{CD}), and Hausdorff Distance components $\overrightarrow{\epsilon}$, $\overleftarrow{\epsilon}$, $\overleftrightarrow{\epsilon}$. All metrics are expressed as a percentage of the object's diameter. Best results are highlighted in **bold**. Detailed per-object results are provided in the Suppl. Sec. D.

Method	PSB					Stanford					Thingi10k				
	$\|C\|$	d_{CD}	$\overrightarrow{\epsilon}$	$\overleftarrow{\epsilon}$	$\overleftrightarrow{\epsilon}$	$\|C\|$	d_{CD}	$\overrightarrow{\epsilon}$	$\overleftarrow{\epsilon}$	$\overleftrightarrow{\epsilon}$	$\|C\|$	d_{CD}	$\overrightarrow{\epsilon}$	$\overleftarrow{\epsilon}$	$\overleftrightarrow{\epsilon}$
Raw MA	30k	0.56	2.43	2.25	2.46	83.7k	0.61	3.26	2.26	3.27	53.3k	2.23	7.34	6.84	7.42
CoverageAxis++	310	0.62	2.86	2.20	2.87	510	0.74	3.30	2.46	3.37	654*	1.59	5.53	5.26	5.83
MASH	200	0.31	1.02	1.53	1.64	400	0.38	1.39	1.33	1.49	400	0.52	1.77	2.74	2.99
Ours	**46**	**0.28**	**0.73**	**0.62**	**0.74**	**77**	**0.33**	**1.07**	**0.89**	**1.08**	**118**	**0.39**	**1.72**	**1.28**	**1.78**

Table 2. Quantitative comparison across four ShapeNet Categories (Airplane, Chair, Lamp, Sofa). We evaluate 100 models per category.

Category	Raw MA					CoverageAxis++					MASH					Ours				
	$\|C\|$	d_{CD}	$\overrightarrow{\epsilon}$	$\overleftarrow{\epsilon}$	$\overleftrightarrow{\epsilon}$	$\|C\|$	d_{CD}	$\overrightarrow{\epsilon}$	$\overleftarrow{\epsilon}$	$\overleftrightarrow{\epsilon}$	$\|C\|$	d_{CD}	$\overrightarrow{\epsilon}$	$\overleftarrow{\epsilon}$	$\overleftrightarrow{\epsilon}$	$\|C\|$	d_{CD}	$\overrightarrow{\epsilon}$	$\overleftarrow{\epsilon}$	$\overleftrightarrow{\epsilon}$
Airplane	82k	0.30	**0.90**	0.58	**0.91**	199	0.49	2.79	1.55	2.80	200	0.27	1.21	1.47	1.63	**46**	**0.23**	0.92	**0.53**	0.97
Chair	83k	0.44	1.11	0.98	1.15	200	0.72	3.49	2.34	3.52	200	0.46	1.47	2.67	2.72	**99**	**0.35**	**0.83**	**0.89**	**1.02**
Lamp	71k	0.32	2.20	**0.65**	**2.22**	167	0.54	3.79	1.63	3.87	200	0.34	**1.51**	2.14	2.33	**59**	**0.28**	2.37	0.75	2.60
Sofa	85k	0.44	1.65	**0.96**	1.68	200	0.79	4.41	2.61	4.41	200	0.43	1.63	1.96	2.16	**116**	**0.36**	**0.86**	1.06	**1.19**

$|C| \approx 46–120$

primitives — vs 30k–84k for Raw MA

Lowest dCD

on every benchmark dataset

Up to 4× fewer

primitives than MASH, better accuracy

[5]Shilane, P., Min, P., Kazhdan, M., Funkhouser, T.: The princeton shape benchmark.In: Proc. IEEE Shape Modeling Applications (2004)
[6]University, S.: The Stanford 3d scanning repository. http://graphics.stanford.edu/data/3Dscanrep/ (2014), accessed: 2026-01-07
[7]Zhou, Q., Jacobson, A.: Thingi10k: A dataset of 10,000 3d-printing models. arXiv:1605.04797 (2016)
[8]Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., et al.: Shapenet: Aninformation-rich 3d model repository. arXiv (2015)

Best fidelity with far fewer primitives

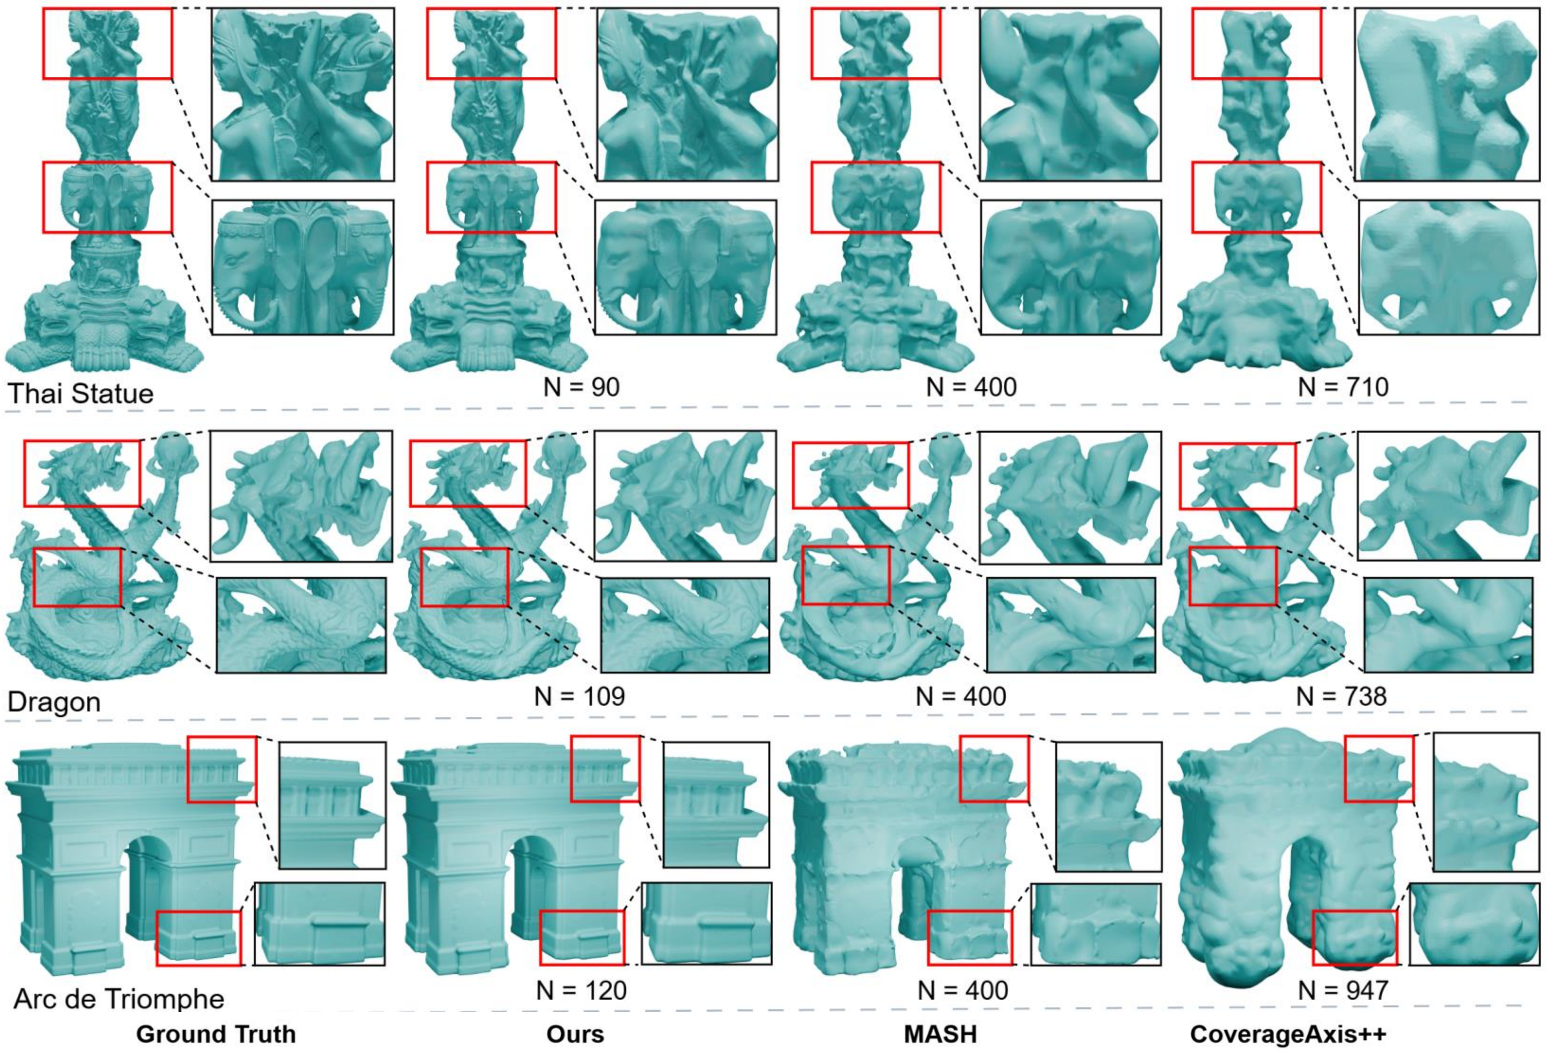

Best fidelity with far fewer primitives

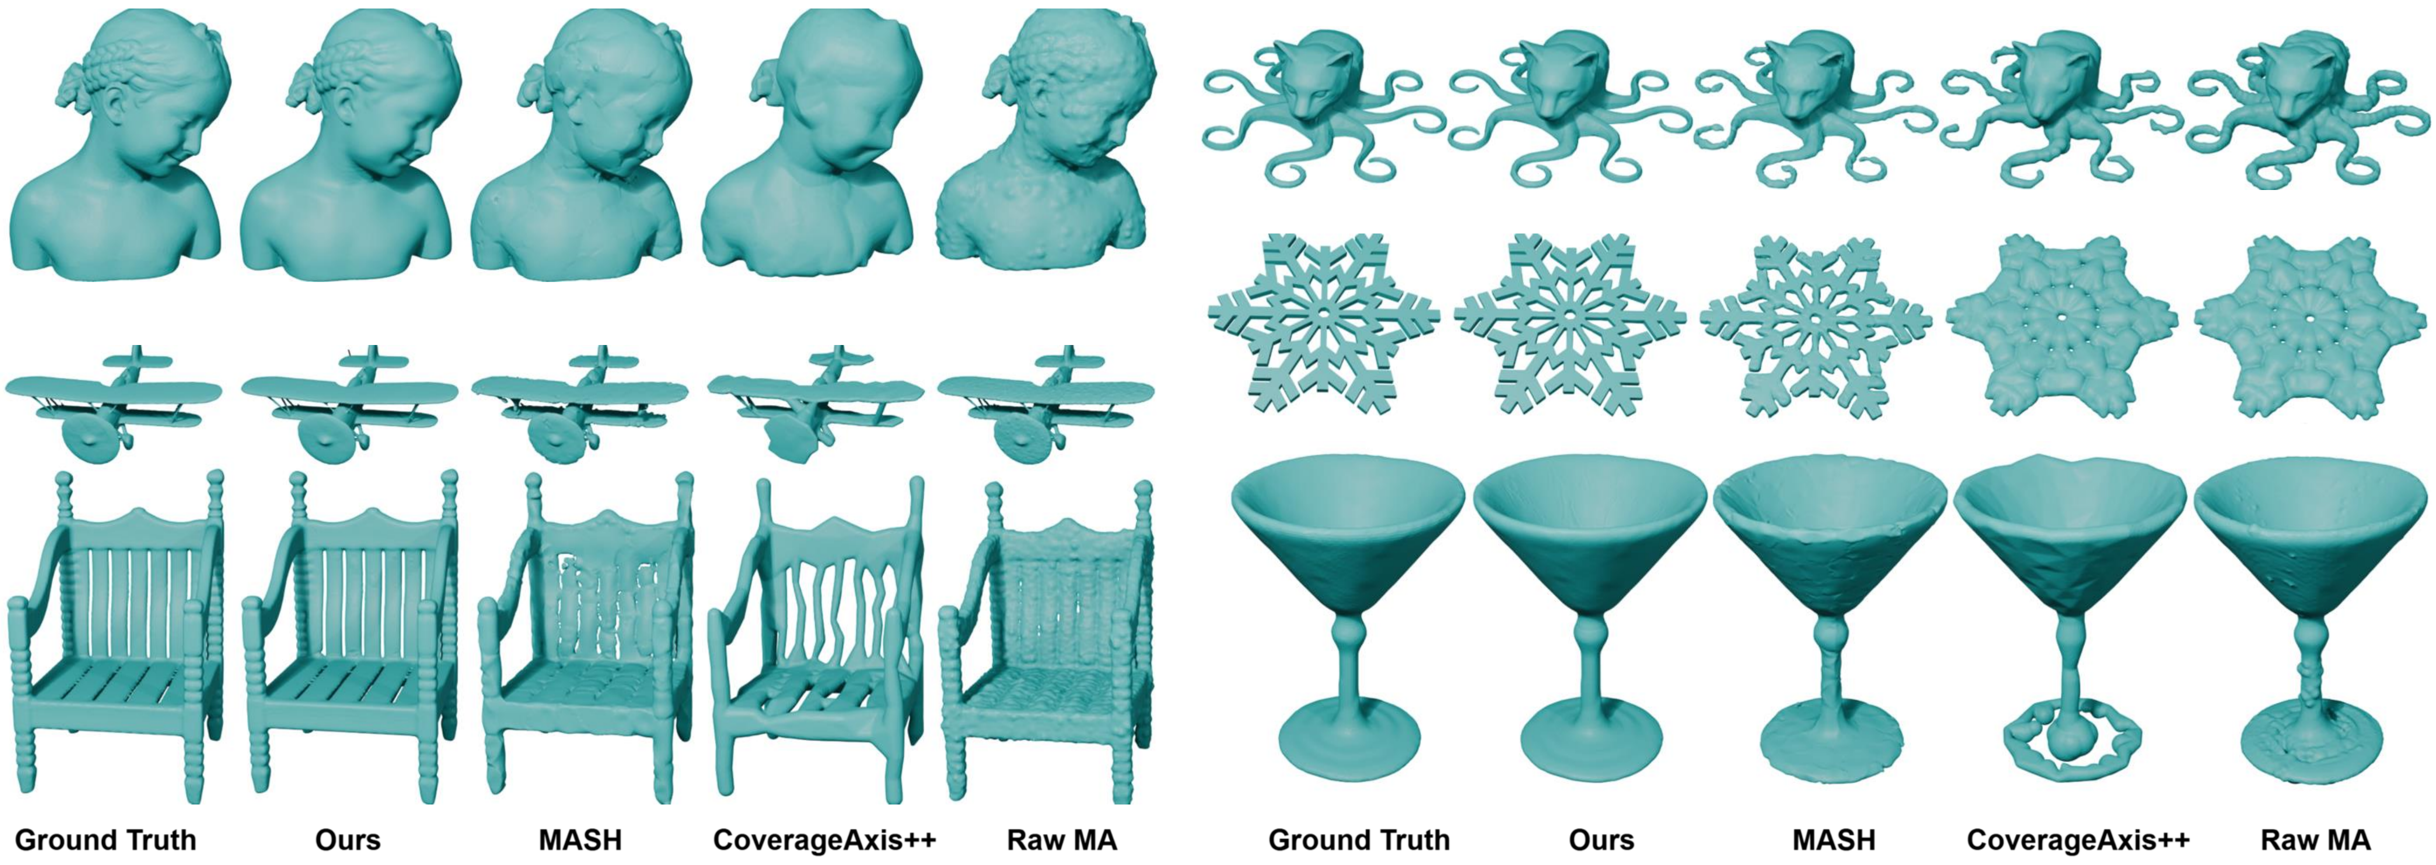

Fastest to encode, high compression ratios

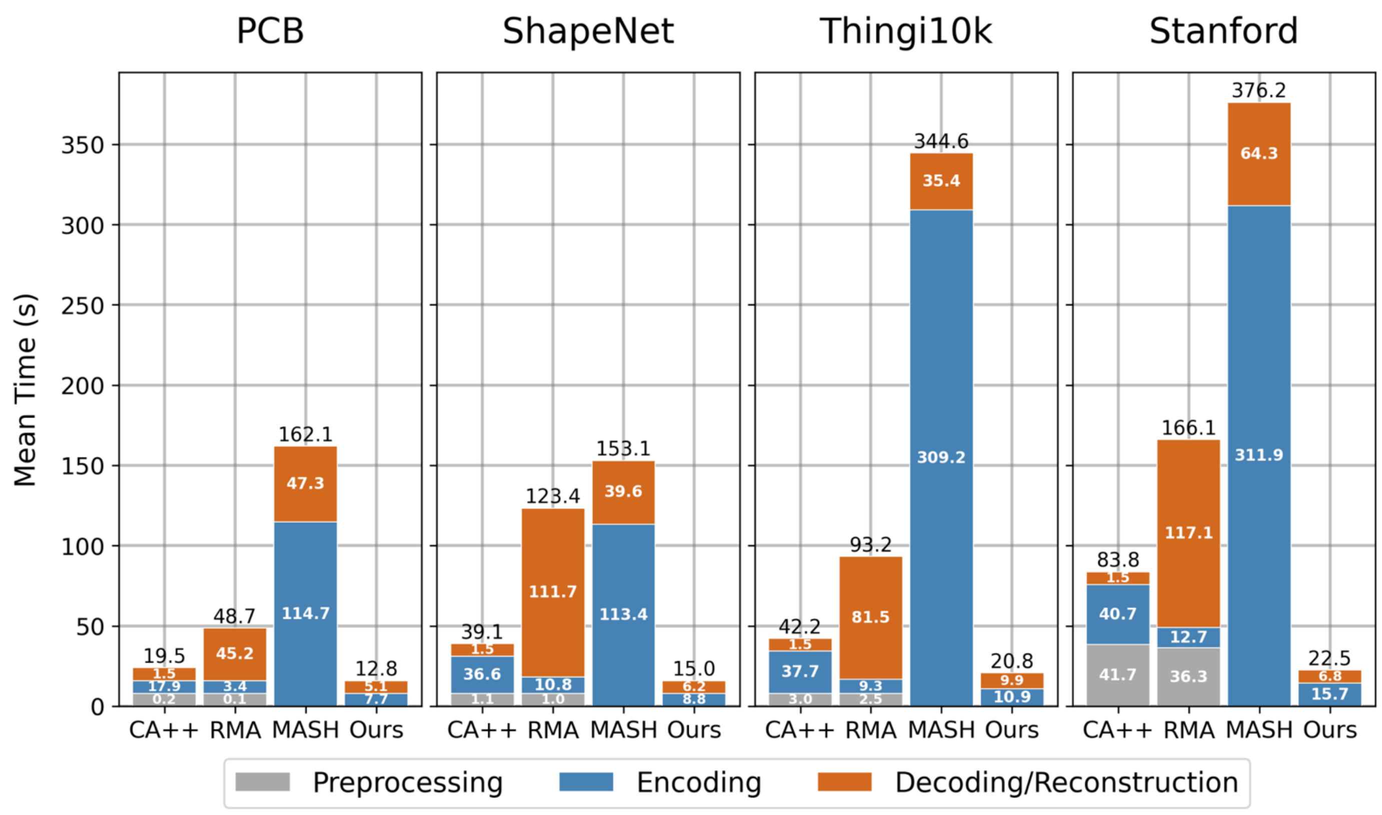

~8 s

to fit a 10M-face Thai Statue

121×

smaller: 181 MB → 1.49 MB (dCD = 0.28% loss)

< 20 s

even at |C| = 200 primitives

Ablation studies on design choices

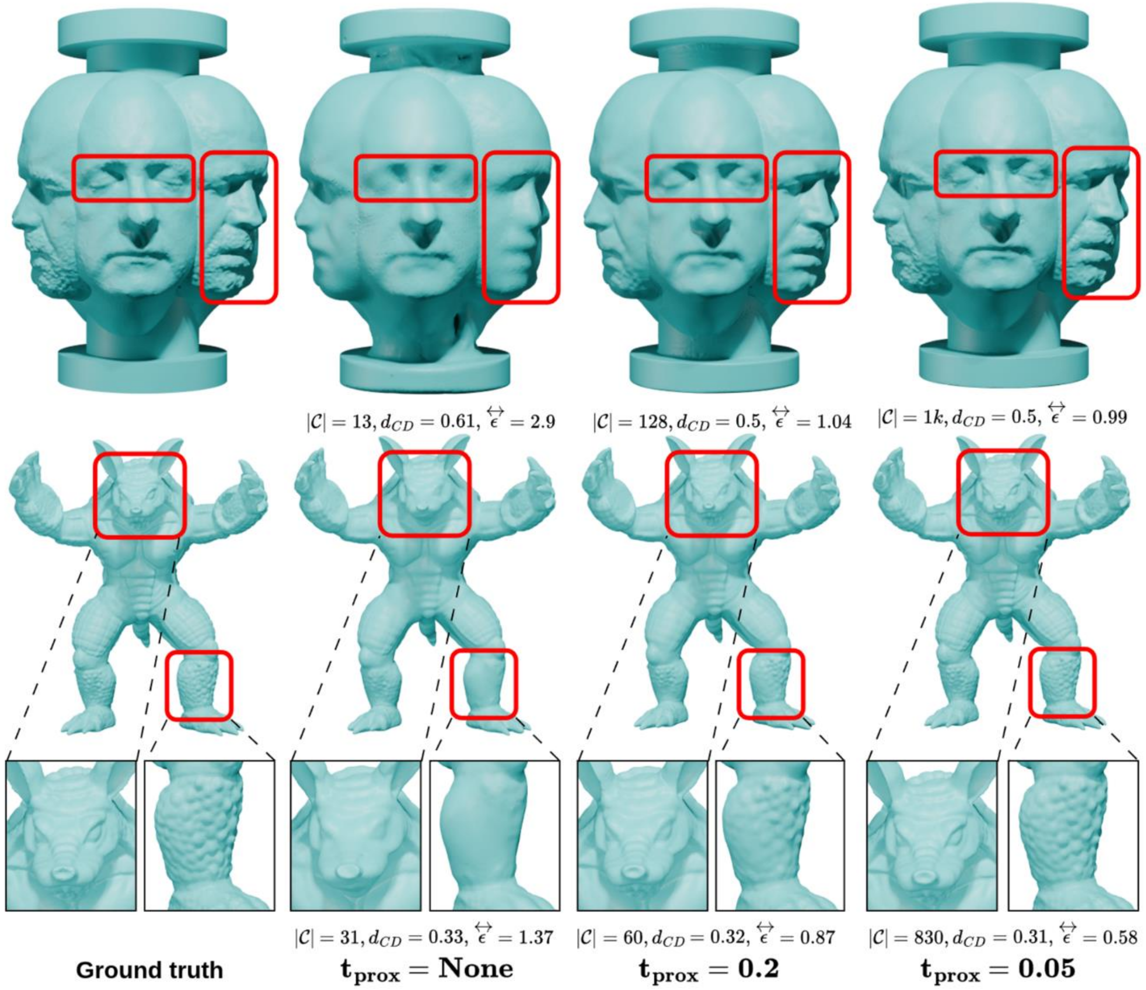

Proximity τ_{prox}

Trades sparsity for locality + detail.

None	|C| = 31 · over-smoothed
0.2	|C| = 60 · sweet spot
0.05	|C| = 830 · redundant

Ablation studies on design choices

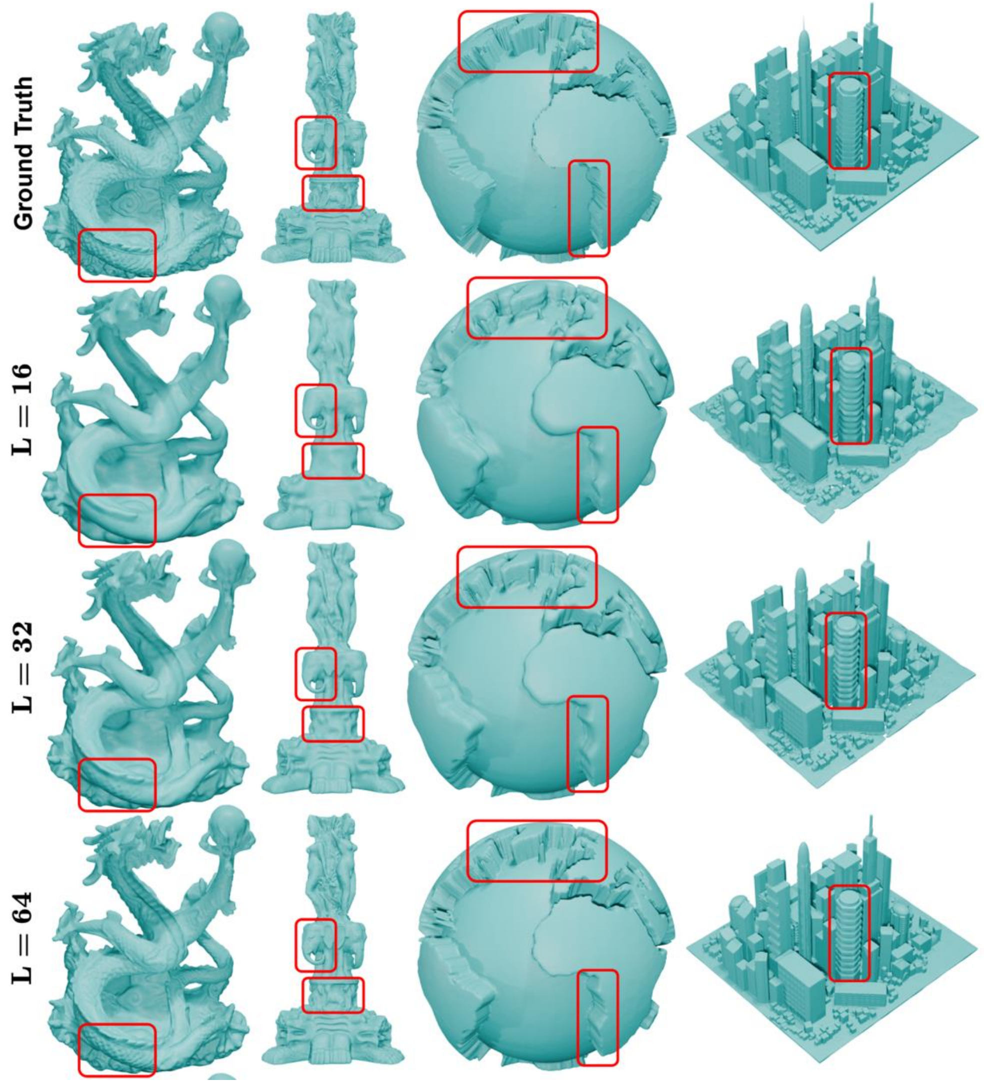

Bandwidth L

Controls high-frequency fidelity.

16 / 32	over-smoothed detail
64	recovers details
128	marginal gain, larger memory footprint

Ablation studies on design choices

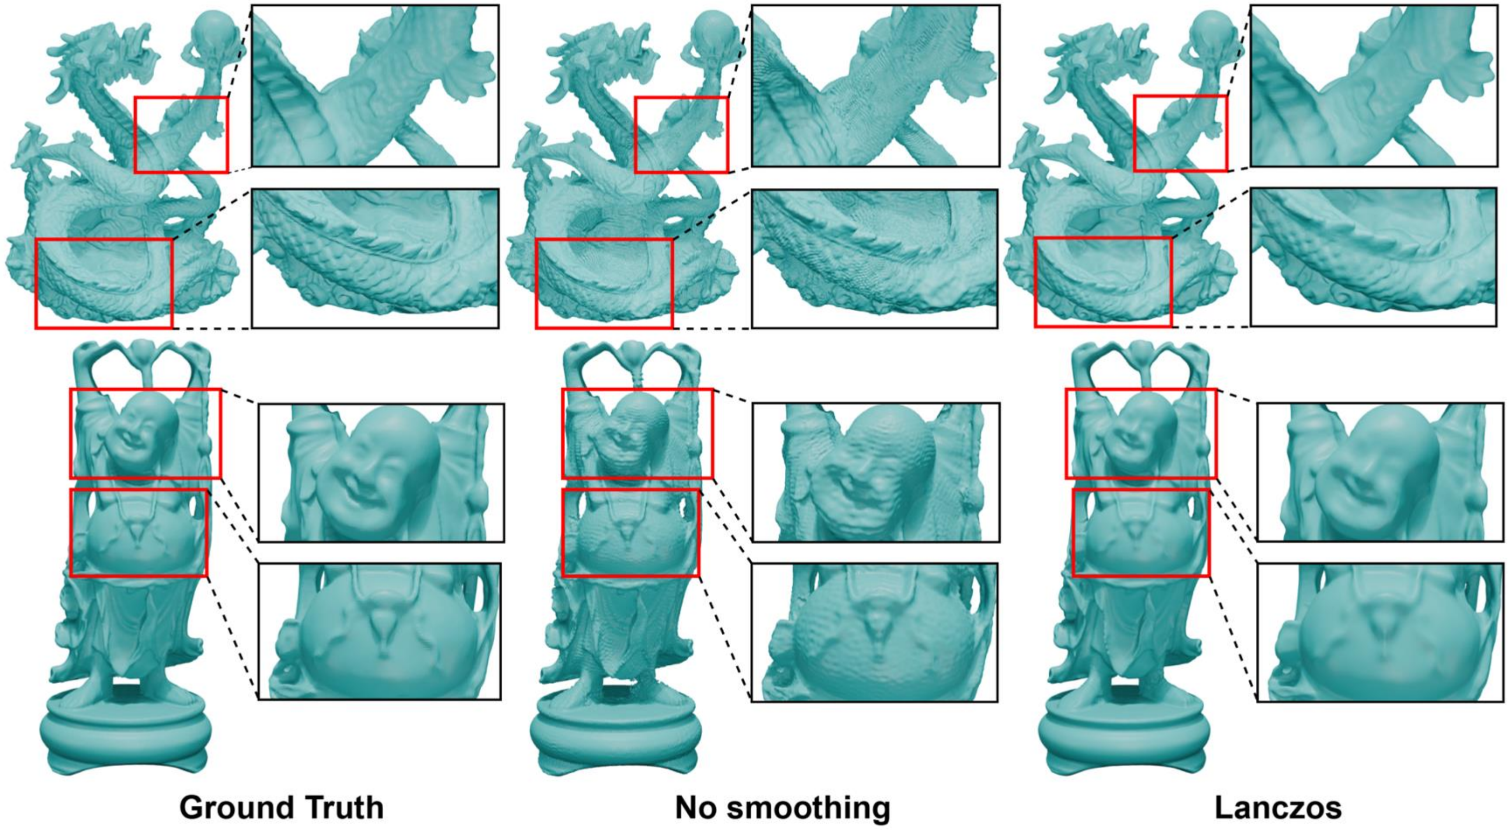

Lanczos smoothing

Mitigates spectral ringing.

Off	ringing artifacts
On	clean surface, detail kept

preserves geometry

CONCLUSION & FUTURE WORK

SHARC

Sparse interior reference points + spherical-harmonic distance fields deliver state-of-the-art accuracy — faster, more scalable, with a tiny footprint, for any topology.

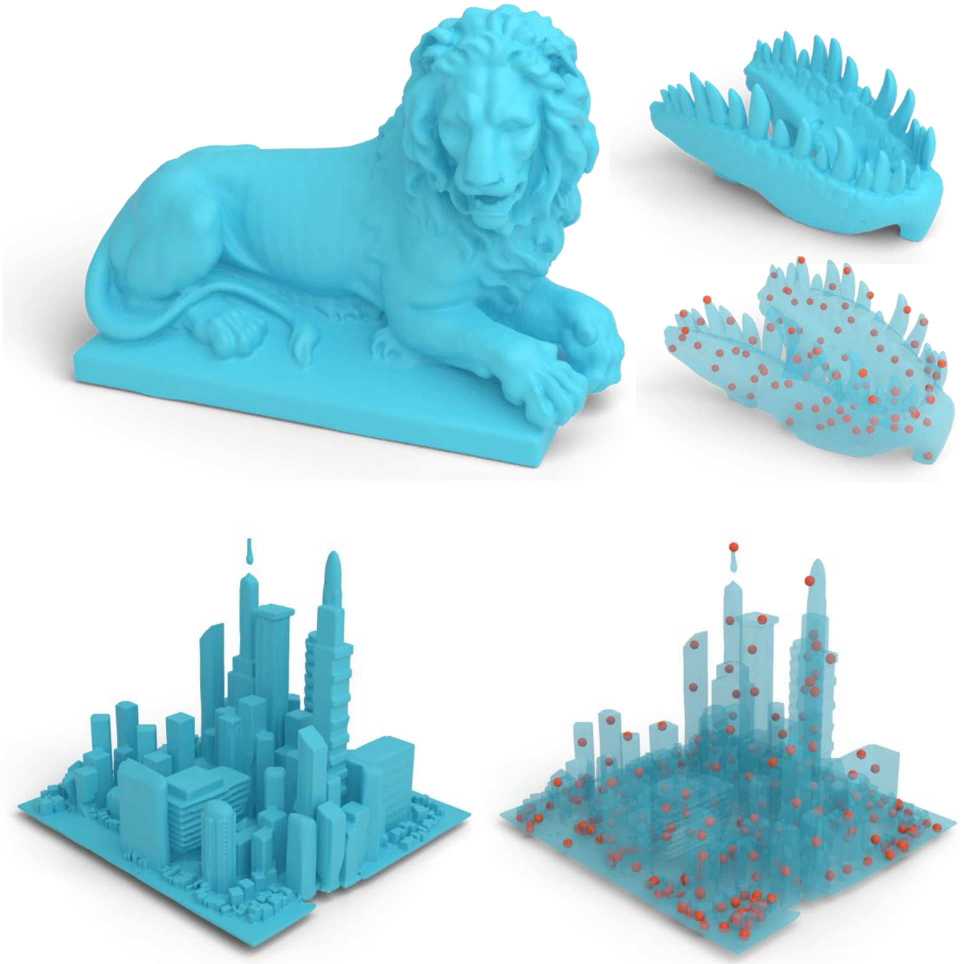

KEY CONTRIBUTIONS

- Multiple interior reference points – arbitrary topology representation
- SOTA fidelity with 4× fewer primitives than prior SH methods
- No training, fast per-object encoding, compression

FUTURE WORK

- 3D shape generation
- UV-free texture (RGB / radiance)
- Differentiable point selection

SHARC

pose-lab.github.io/SHARC/

Thank You!

Project page

Panagiotis Sapoutzoglou · George Terzakis · Maria Pateraki

National Technical University of Athens, Greece

{psapoutzoglou,gterzakis,mpateraki}@mail.ntua.gr