

Points2Surf

Learning Implicit Surfaces from Point Clouds

Supplementary Material

Anonymous ECCV submission

Paper ID 2965

S.1 Overview

In this supplementary material, we provide evaluations on a new dataset with more organic shapes (THING10K [1]), consisting of both quantitative (Section S.2) and qualitative (Section S.3) comparisons, and additional qualitative comparisons on the ABC dataset (Section S.4).

S.2 Quantitative Comparison on Thingi10k

We show a quantitative comparison of reconstruction errors on the THING10K dataset in Table S1. The Chamfer distance between reconstructed and ground truth surfaces *averaged* over all shapes in a dataset is shown. Both the absolute value of the error multiplied by 100 (abs.), and the error relative to Point2Surf (rel.) are shown to facilitate the comparison. Our method consistently performs better than the baselines, due to its strong and generalizable prior. Note that POINTS2SURF was *not* retrained on the Thingi10k – this shows generalization results.

Table S1. Quantitative comparison of reconstruction errors on the Thingi10k dataset. Note that none of the methods was *not* retrained on the Thingi10k in order to test generalization to new data.

	DeepSDF		AtlasNet		SPR		POINTS2SURF	
	abs.	rel.	abs.	rel.	abs.	rel.	abs.	rel.
THING10K no-noise	9.16	6.48	5.29	3.74	1.78	1.26	1.41	1.00
THING10K med-noise	8.83	5.99	5.19	3.52	1.81	1.23	1.47	1.00
THING10K max-noise	12.28	4.68	4.90	1.87	3.23	1.23	2.62	1.00
THING10K sparse	9.56	4.54	5.64	2.68	2.35	1.12	2.11	1.00
THING10K dense	8.35	6.19	5.02	3.72	1.57	1.16	1.35	1.00
average	9.64	5.58	5.21	3.11	2.15	1.20	1.79	1.00

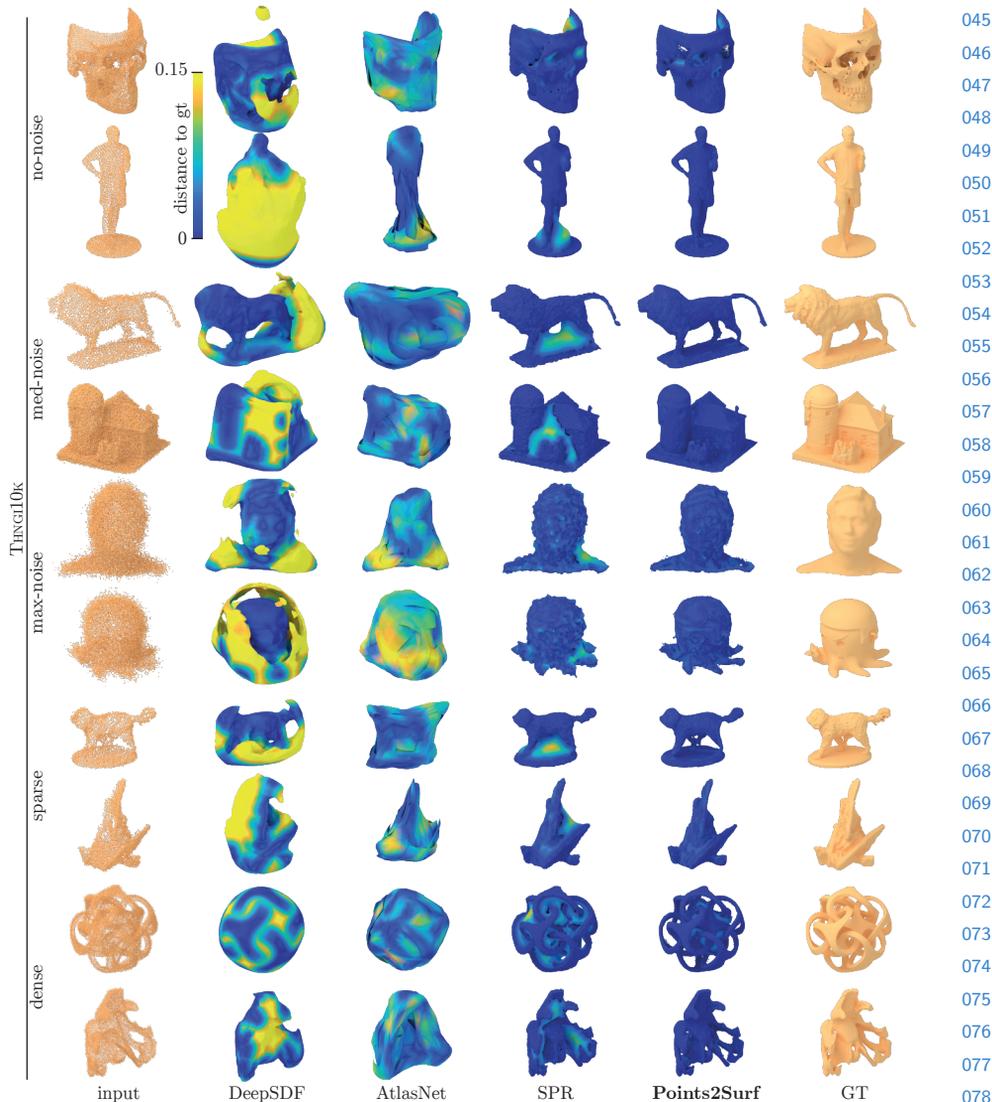


Fig. S1. Qualitative comparison of surface reconstructions on the THING10K dataset. We evaluate two examples from each dataset variant with each method. Colors show the distance of the reconstructed surface to the ground truth surface, shades of blue indicating low error.

S.3 Qualitative Results on Thingi10k

We evaluate POINTS2SURF on the THING10K dataset [1]. We take 100 meshes that are tagged with ‘scan’ or ‘sculpture’. These objects are mostly animals, humans and faces, many of them realistic, some artistic. We create the same dataset variants as in the FAMOUS dataset: a version without noise (THING10K

no-noise), a version with a medium noise strength $0.01L$ (THING10K *med-noise*), and a version with maximum amount of noise $0.05L$ (THING10K *max-noise*). Additionally we create sparser and denser point clouds by varying the number of scans: a variant with 5 scans instead of 10 (THING10K *sparse*), and a version with 30 scans (THING10K *dense*), both with a medium noise strength of $0.01L$. Some qualitative results are shown in Figure S1.

S.4 Additional Qualitative Results on ABC

We provide additional qualitative comparisons on the ABC dataset in Figure S2. The evaluation setup is the same as for Figure 4 in the paper.

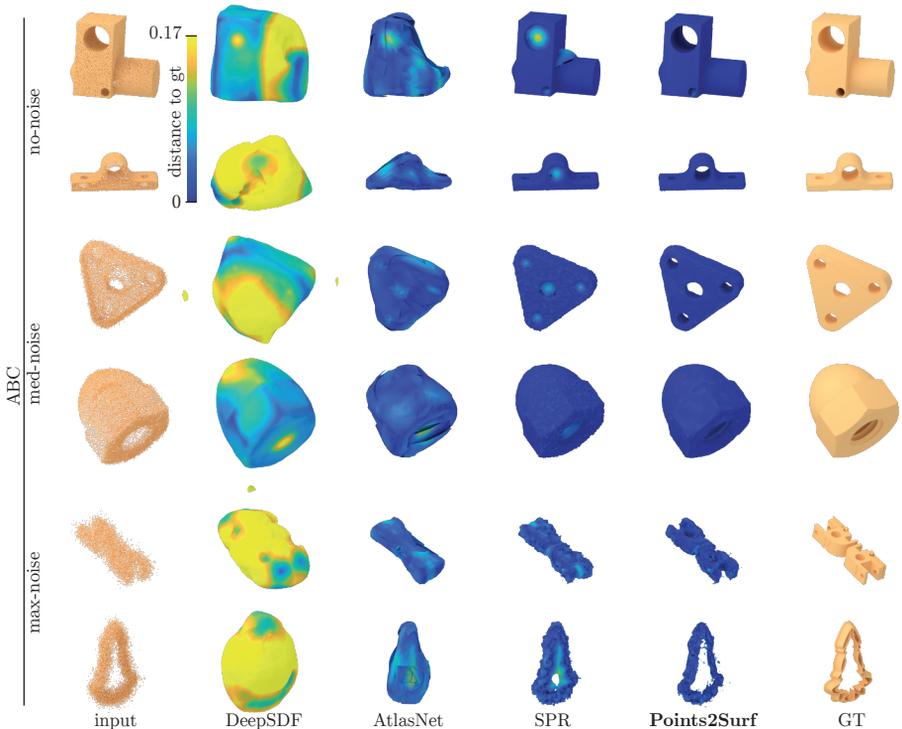


Fig. S2. Additional qualitative comparison of surface reconstructions on the ABC dataset. We evaluate two examples from each dataset variant with each method. Colors show the distance of the reconstructed surface to the ground truth surface.

References

1. Zhou, Q., Jacobson, A.: Thing10k: A dataset of 10,000 3d-printing models. arXiv preprint arXiv:1605.04797 (2016)