

ARO-Net: Learning Neural Fields from Anchored Radial Observations

Supplementary Material

Yizhi Wang^{1,2*}, Zeyu Huang^{1*}, Ariel Shamir³, Hui Huang¹, Hao Zhang², Ruizhen Hu^{1†}
¹Shenzhen University ²Simon Fraser University ³Reichman University

1. Anchor Placement

In Figure 1, we demonstrate how anchors are placed by uniform, grid and Fibonacci samplings. In grid sampling, we first choose all grid points $\{(x, y, z) | x, y, z \in \{-0.5, -0.25, 0, 0.25, 0.5\}\}$ inside the unit sphere and then randomly select from remaining grid points to make a total count of 48. Table 1 shows how different anchor placement strategies affect the construction performance on ABC, where Fibonacci sampling is slightly better than grid sampling, and significantly better than uniform sampling. In the early stage of training, Fibonacci sampling exhibits a clear advantage over other two samplings, as shown in Figure 3. In Figure 4, we show the curves in the later stage (from Epoch 100 to 260). We find that grid and uniform sampling narrow their performance gap between Fibonacci sampling in the later stage. However, as shown in the testing loss curve, grid sampling suffers from turbulence of the loss values. In terms of the stability of training process and the overall performance, Fibonacci sampling is superior to other two sampling strategies. We leave more comprehensive investigations on anchor placement as our future work.

2. Adaptive Anchors

As suggested in Table 2, optimizing the anchor positions in the training stage (the first row) cannot improve the performance of ARO-Net on ABC. In the testing phase, we try to adjust the anchor positions for each testing case. Specifically, we fix the network parameters, utilize input point clouds as ground-truth occupancies, and then adjust the anchor positions using Gradient Descent algorithm. Table 2 demonstrates that fine-tuning the anchors for 100 iterations can bring very marginal improvements; 200 iterations will degrade the performance because of over-fitting.

3. Distance and Direction Information

As shown in Table 3, the reconstruction performance of ARO-Net drops significantly when we remove r and $\|r\|$.

*Equal contribution

†Corresponding author. E-mail: ruizhen.hu@gmail.com




Figure 1. Visualization of different anchor placement strategies.

Setting	LFD↓	HD↓	CD↓	EMD↓	IOU↑
Uniform	1.52	2.32	6.29	1.18	8.67
Gird	1.42	2.20	5.63	1.11	8.76
Fibonacci	1.35	2.25	5.46	1.12	8.79

Table 1. How different anchor placement strategies affect the re-construction performance.

Adjusting Stage	LFD↓	HD↓	CD↓	EMD↓	IOU↑
Training	1.60	2.39	5.98	1.333	8.63
Testing (100)	1.58	2.24	5.44	1.098	8.78
Testing (200)	1.80	2.20	5.75	1.314	8.74
Default (Fixed)	1.35	2.25	5.46	1.121	8.79

Table 2. Adjusting anchor positions in the training or testing stage. The numbers behind “testing” denote the iteration numbers of adjusting anchor positions.

Setting	LFD↓	HD↓	CD↓	EMD↓	IOU↑
w/o r and $\ r\ $	1.38	2.38	5.94	1.22	8.62
Full model	1.35	2.25	5.46	1.12	8.79

Table 3. Removing the distance ($\|r\|$) and direction information (r) in ARO-Net.

Since the prediction of ARO-Net is based on engaging r and $\|r\|$ with the radial observations, missing r and $\|r\|$ as network input would make the prediction more difficult.

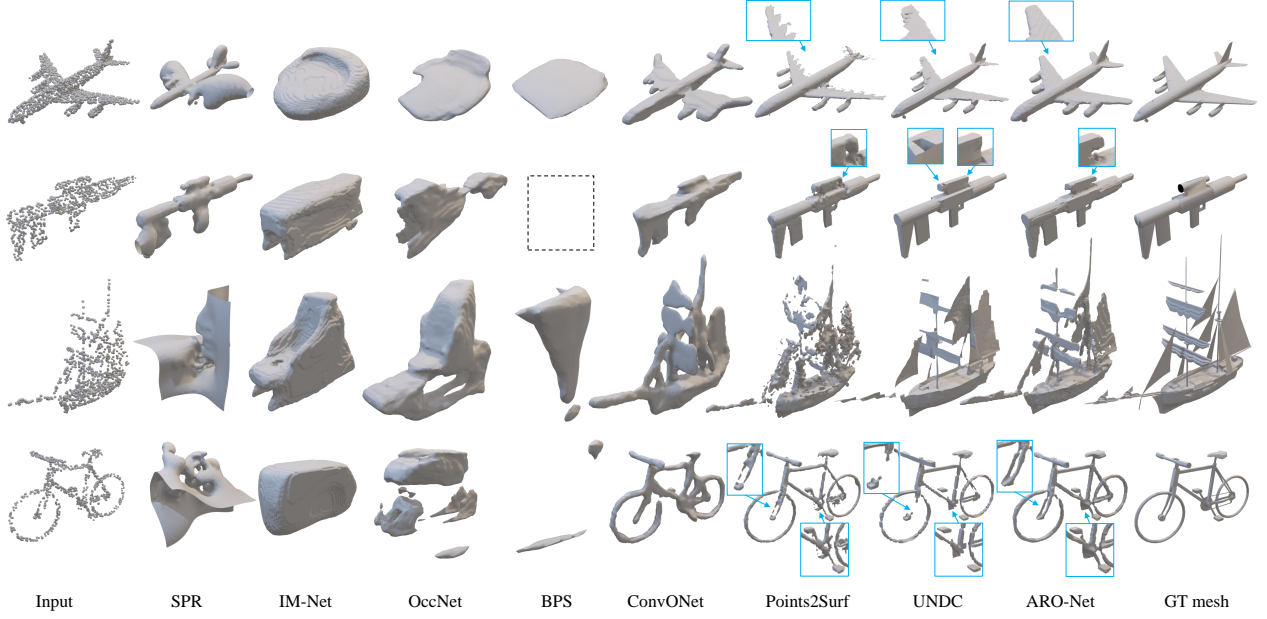


Figure 2. More visual comparisons to state-of-the-art methods on ShapeNet. All methods (except SPR) were trained on chairs in ShapeNet and tested on other categories with zoom-ins highlighting reconstruction artifacts. BPS generated an empty result for the second case.




Figure 3. How different anchor placement strategies affect network performance, where anchor count $m = 48$ for all settings.




Figure 4. The testing loss for different strategies of anchor placement. The loss is calculated by the cross-entropy between predicted probability and GT occupancy.

4. 2D Experiments

In 2D experiments, we replace the PointNet feature f_i with the hit distance of r_i (a ray from anchor a_i to query point) on a given 2D shape S . Specifically, the hit distance (denoted as d_i) is calculated from a_i to the hit point on the boundary pixels of S . Note that it is still non-trivial to reconstruct the 2D shape under this setting because (1) the occupancy of non-boundary points remains unknown and (2) the anchors could observe only partial boundaries because of their count and placement. We send r , $\|r\|$, and d into the network to predict the occupancy. The attention module consists of 3 encoding layers and 4 attention heads. When reconstructing the images of 2D shapes, we utilize the predicted occupancy values to construct the pixel values.

Method	ABC	Chairs	Airplanes
POCO/ARO-Net	1.71/1.35	2.82/1.92	5.11/3.56

Table 4. Comparison between ARO-Net and POCO.

5. More Comparisons

We compare ARO-Net with a more recent work POCO [1] using LFD (\downarrow) in Table 4, where ARO-Net has a big advantage in both ABC and ShapeNet Chairs/Airplanes.

In Figure 2 and Figure 5, we provide more visual compar-

isons on ShapeNet and ABC dataset. In Figure 6, we provide more visual comparisons on one-shape training, and adding Points2Surf [3] and UNDC [2] into comparisons. ARO-Net performs significantly better than others in extremely sparse point clouds (such as the ship and bicycle in Figure 2). ARO-Net also shows great ability in preserving local details such as the holes in 1st and 4th examples in Figure 5 and the ears of horse in Figure 6. Compared to UNDC, ARO-Net can avoid generating undesired holes on the surfaces. Compared to Points2Surf and ConvONet, ARO-Net can reconstruct better details, smoother surfaces and sharper features. Among all of these comparisons across a variety of shapes, ARO-Net achieves the best overall performance.

References

- [1] Alexandre Boulch and Renaud Marlet. Poco: Point convolution for surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6302–6314, 2022.
- [2] Zhiqin Chen, Andrea Tagliasacchi, Thomas Funkhouser, and Hao Zhang. Neural dual contouring. *ACM Trans. Graph.*, 41(4), 2022.
- [3] Philipp Erler, Paul Guerrero, Stefan Ohrhallinger, Niloy J. Mitra, and Michael Wimmer. Points2Surf: Learning implicit surfaces from point clouds. In *Eur. Conf. Comput. Vis.*, pages 108–124, 2020.




Figure 5. More visual comparisons to state-of-the-art methods on ABC with zoom-ins highlighting reconstruction artifacts.




Figure 6. More visual comparisons on the results obtained by training only on the “fertility” model with rotation and scaling.