Supplementary material:

EC-Net: an Edge-aware Point set Consolidation Network

Submission ID: 817

Overview

There are four sections in this supplementary material.

- Section A presents our training data set.
- Section B presents the network architecture of EC-Net.
- Section C presents more experimental results.
- Section D presents additional results from input point clouds with different number of point samples.

A. Our Training Data set

We collected 12 everyday objects and 24 CAD models in total as our training data set, and manually annotate sharp edges on the models. Figures 1 and 2 show the training models with the annotated edges.

Figure 1. Training models (12 everyday objects) with the annotated edges marked in red.

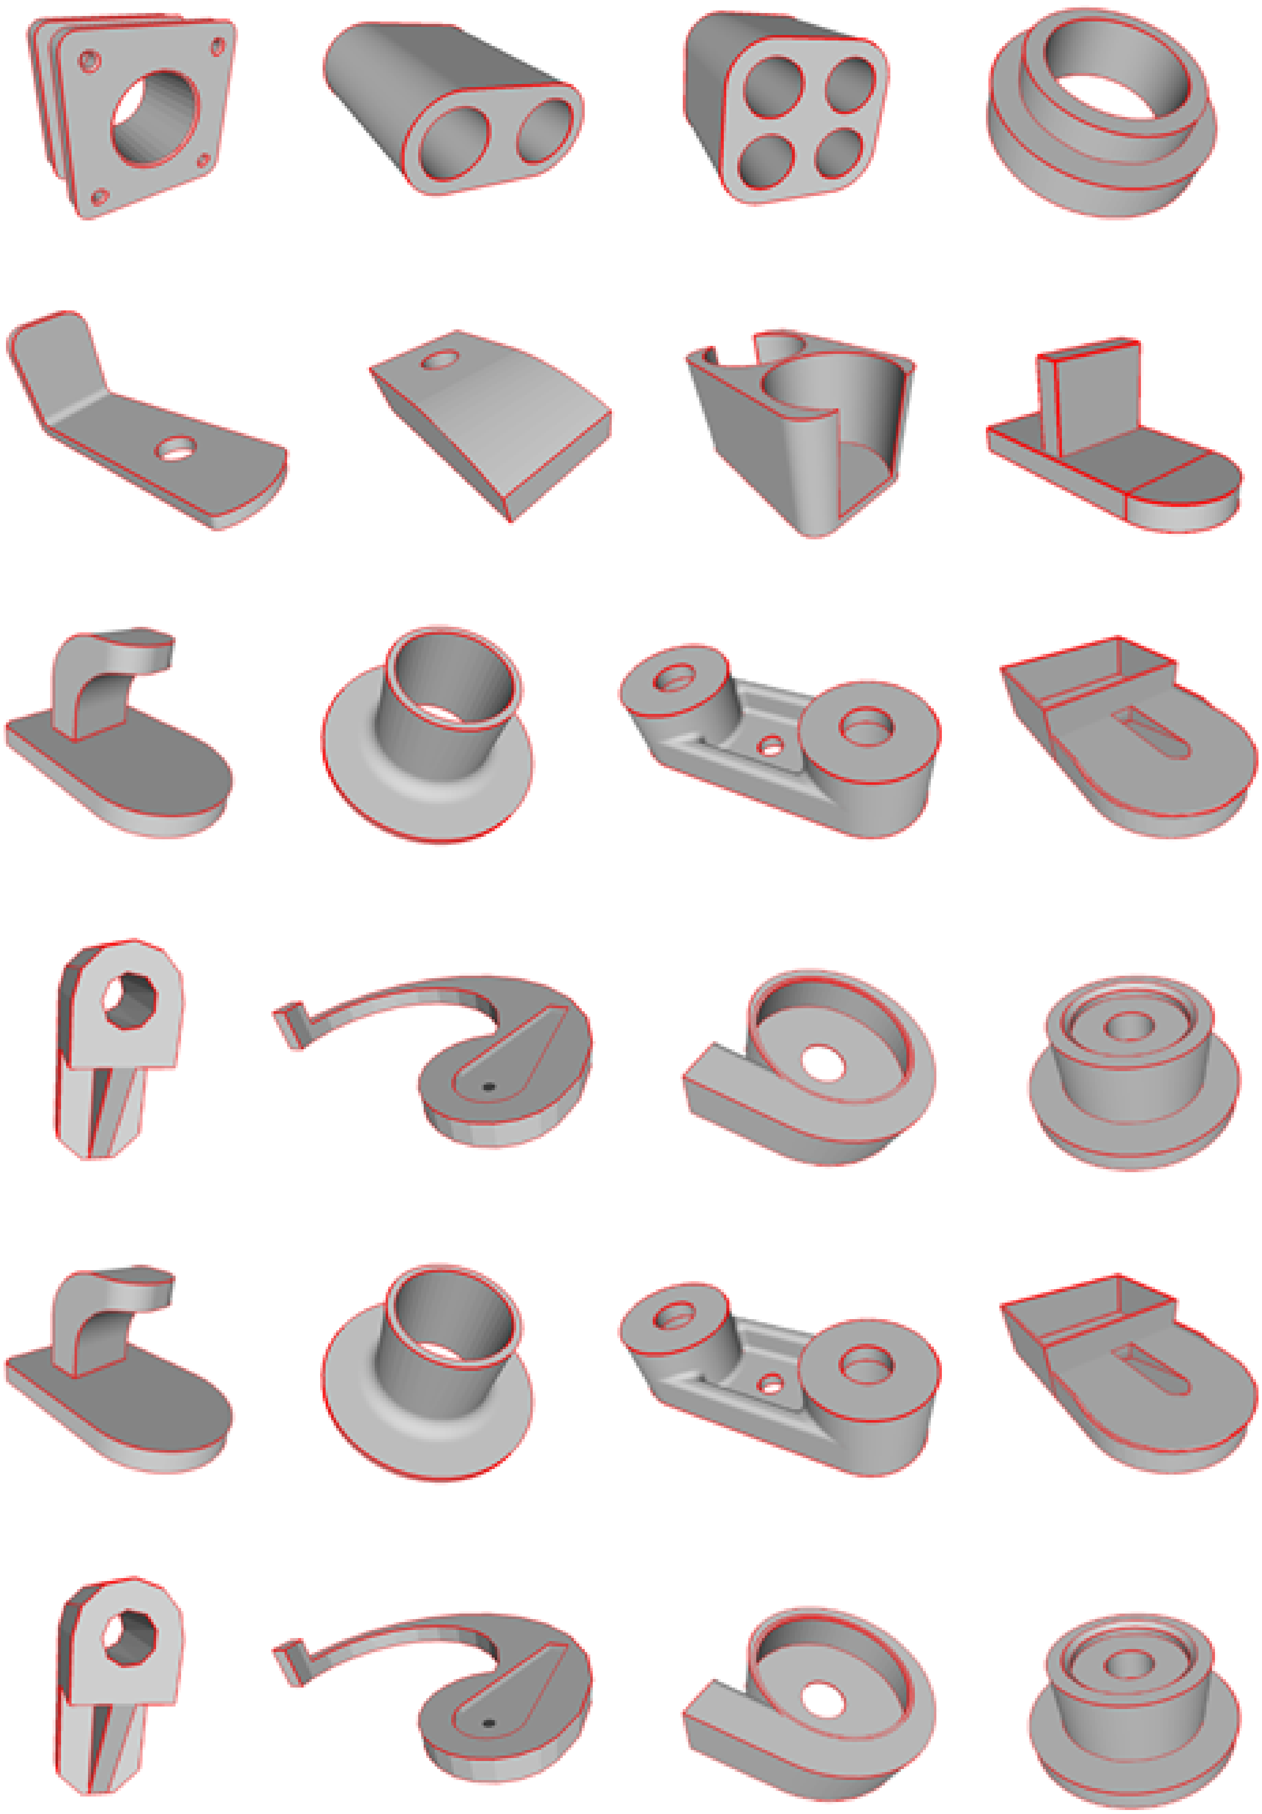

Figure 2. Training models (24 CAD models) with the annotated edges marked in red.

B. Network Architecture of EC-Net

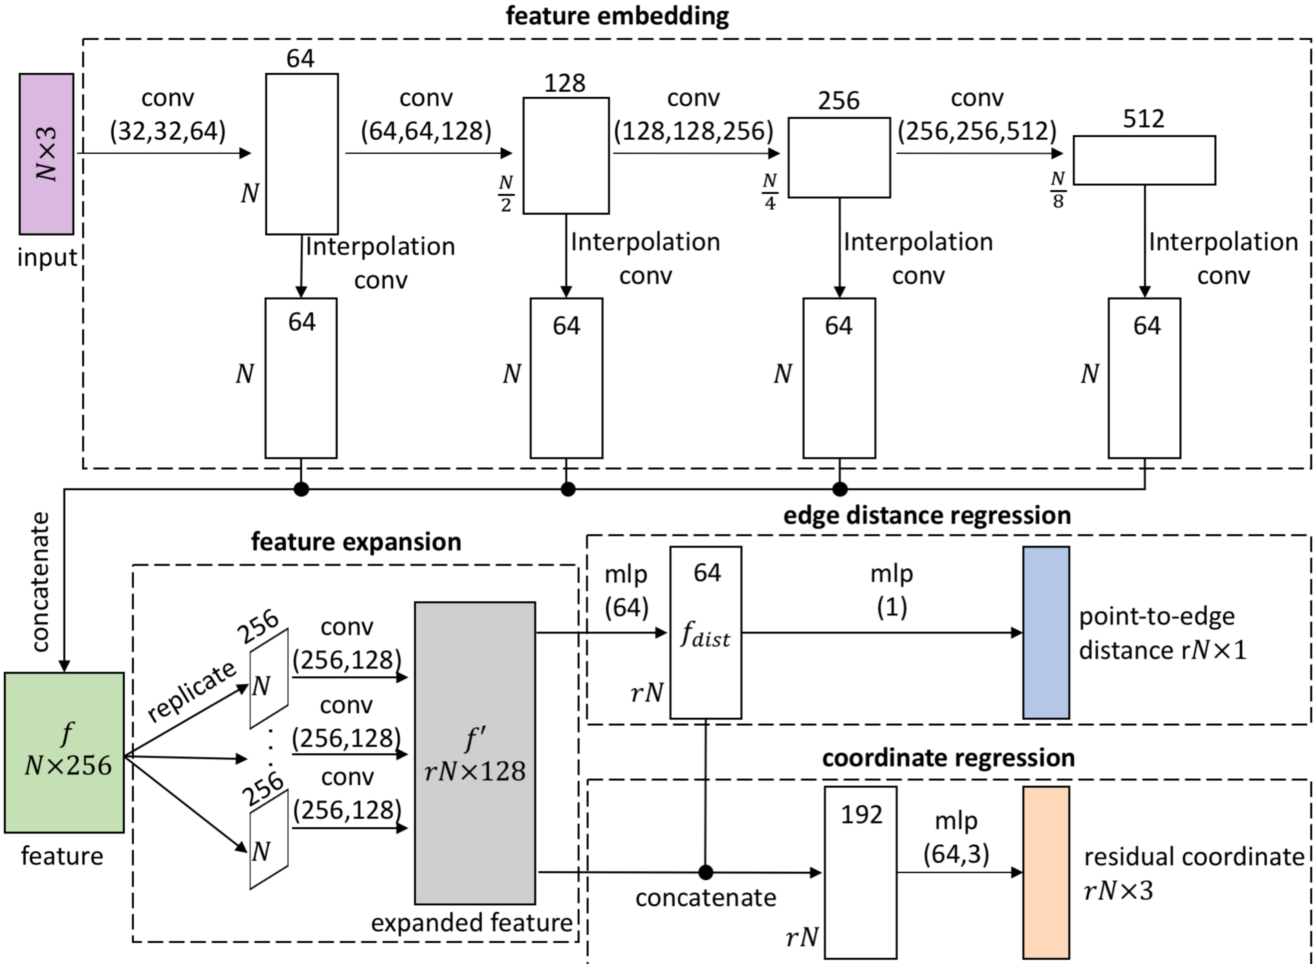

Figure 3. The network architecture of EC-Net.

Figure 3 presents the network architecture of EC-Net; see the explanation below for the details:

- The feature embedding component (see the dashed box on top) is based on PointNet++, where we adopt four levels with grouping radii 0.1, 0.2, 0.4, and 0.6 to extract the local features. The corresponding number of point samples in these four levels are N, $\frac{N}{2}$, $\frac{N}{4}$, and $\frac{N}{8}$, respectively.

- Next, we directly concatenate features from different levels to aggregate the multi-level feature. Specifically, for each level, we use the interpolation in PointNet++ to restore the feature of the level, and then use a convolution to reduce the restored feature to 64 dimensions. After that, we concatenate the restored features from the four levels to form the D = 256 dimensional feature denoted as f (see the green box above).

- In the feature expansion component (see the dashed box on lower left), we follow the feature expansion module in PU-Net. Specifically, we create r copies of feature f, independently apply two separated convolutions with 256 and 128 output channels to each copy, and then concatenate the outputs to form the expand feature denoted as f' (see the grey box above).

- In the edge distance regression component, we first use one fully-connected layer with width 64 to regress f_{dist} from f'. Then, we use another fully-connected layer with width 1 to regress the point-to-edge distance d.
- In the coordinate regression component, we first concatenate the f_{dist} and f' to form a 192 dimensional feature. Then, we use two fully-connected layers with width 64 and 3 to regress the residual point coordinates from the concatenated feature, and output the final 3D coordinates by adding back the original point coordinates.

All the convolution layers and fully-connected layers in the above network are followed by the ReLU operator, except for the last two point-to-edge distance and residual coordinate regression layers.

C. More Experimental Results

C.1. Comparison with other reconstruction methods on benchmark models

Besides the comparison in the main paper, we further evaluate our method, and compare with other surface reconstruction methods on the reconstruction benchmark models [1]. We also include a recent method for large-scale surface reconstruction [2].

Fig. 4 below shows the comparison results on two sharp objects in the benchmark, where (b)-(d) are results downloaded from the benchmark's website related to the best three visual methods presented in the website. The comparison clearly shows that the reconstructions with our consolidation better preserve the sharp edges and have better visual quality even on the noisy models from the benchmark dataset with random error and systematic error.

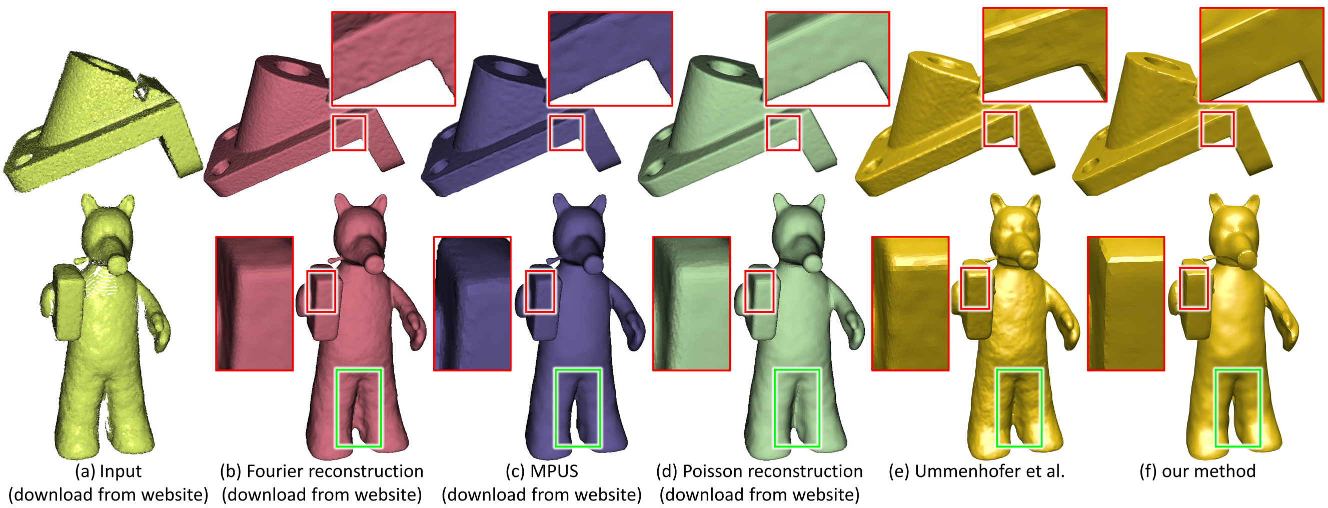

Figure 4. Comparison of different reconstruction methods on benchmark models.

C.2. Additional point consolidation and surface reconstruction results

Figures 5 to 7 present additional point consolidation and surface reconstruction results produced with our method.

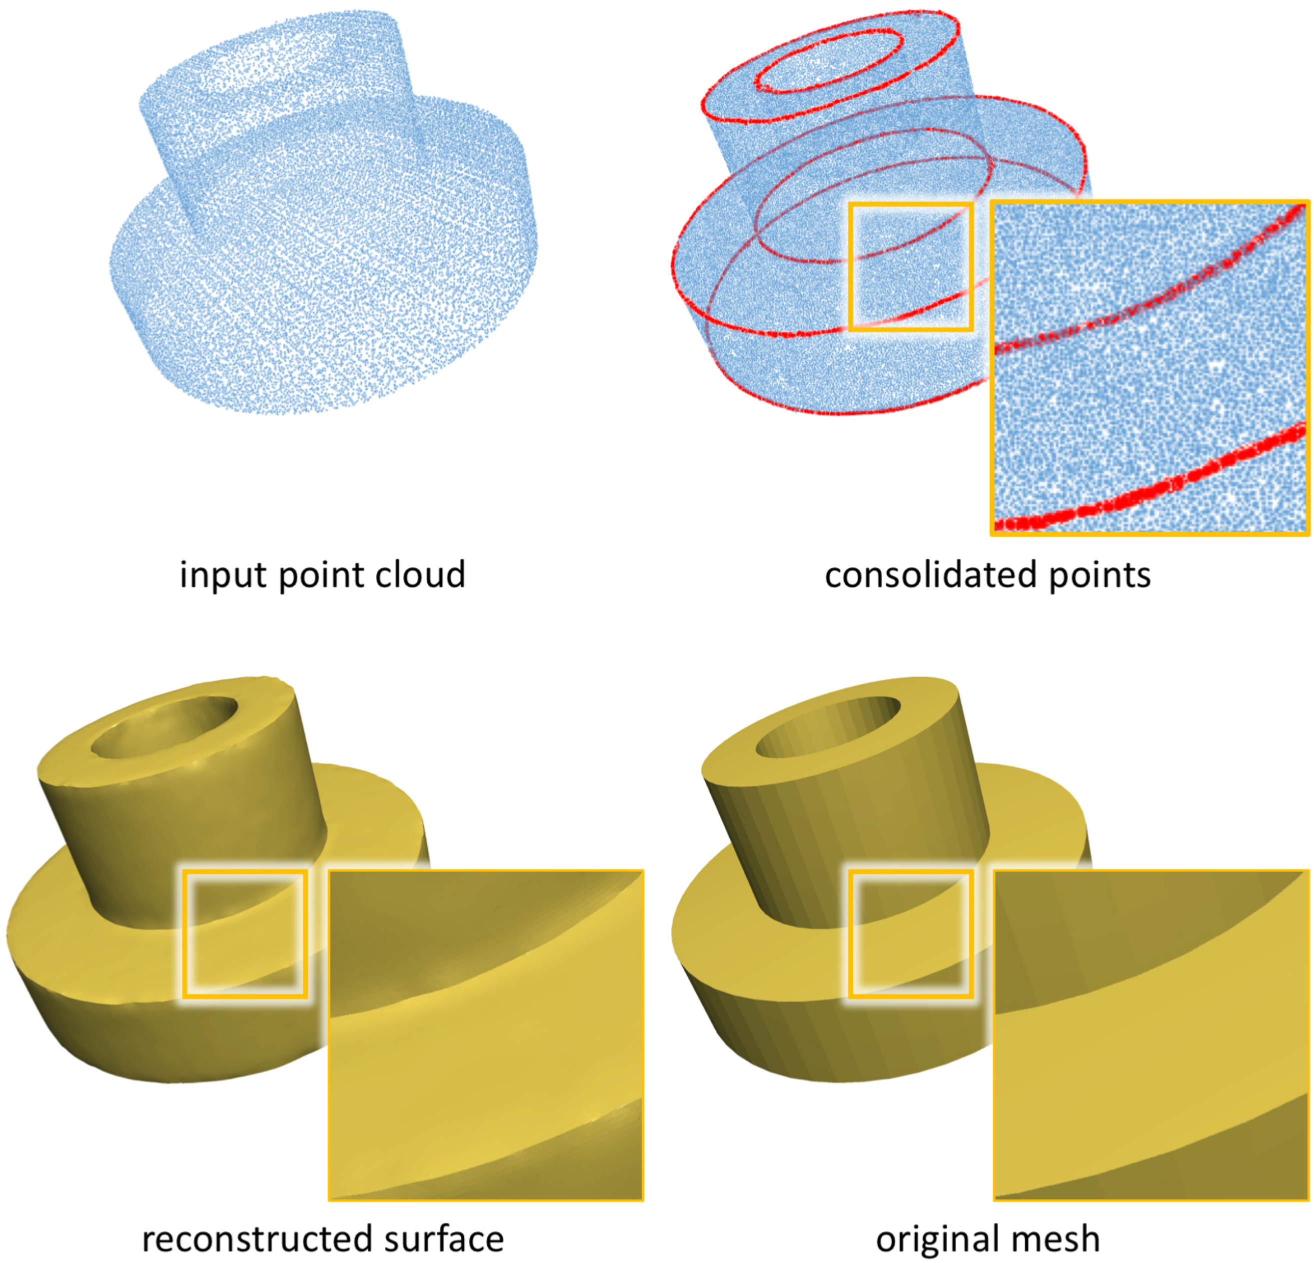

Figure 5. Point consolidation and surface reconstruction results produced with our method.

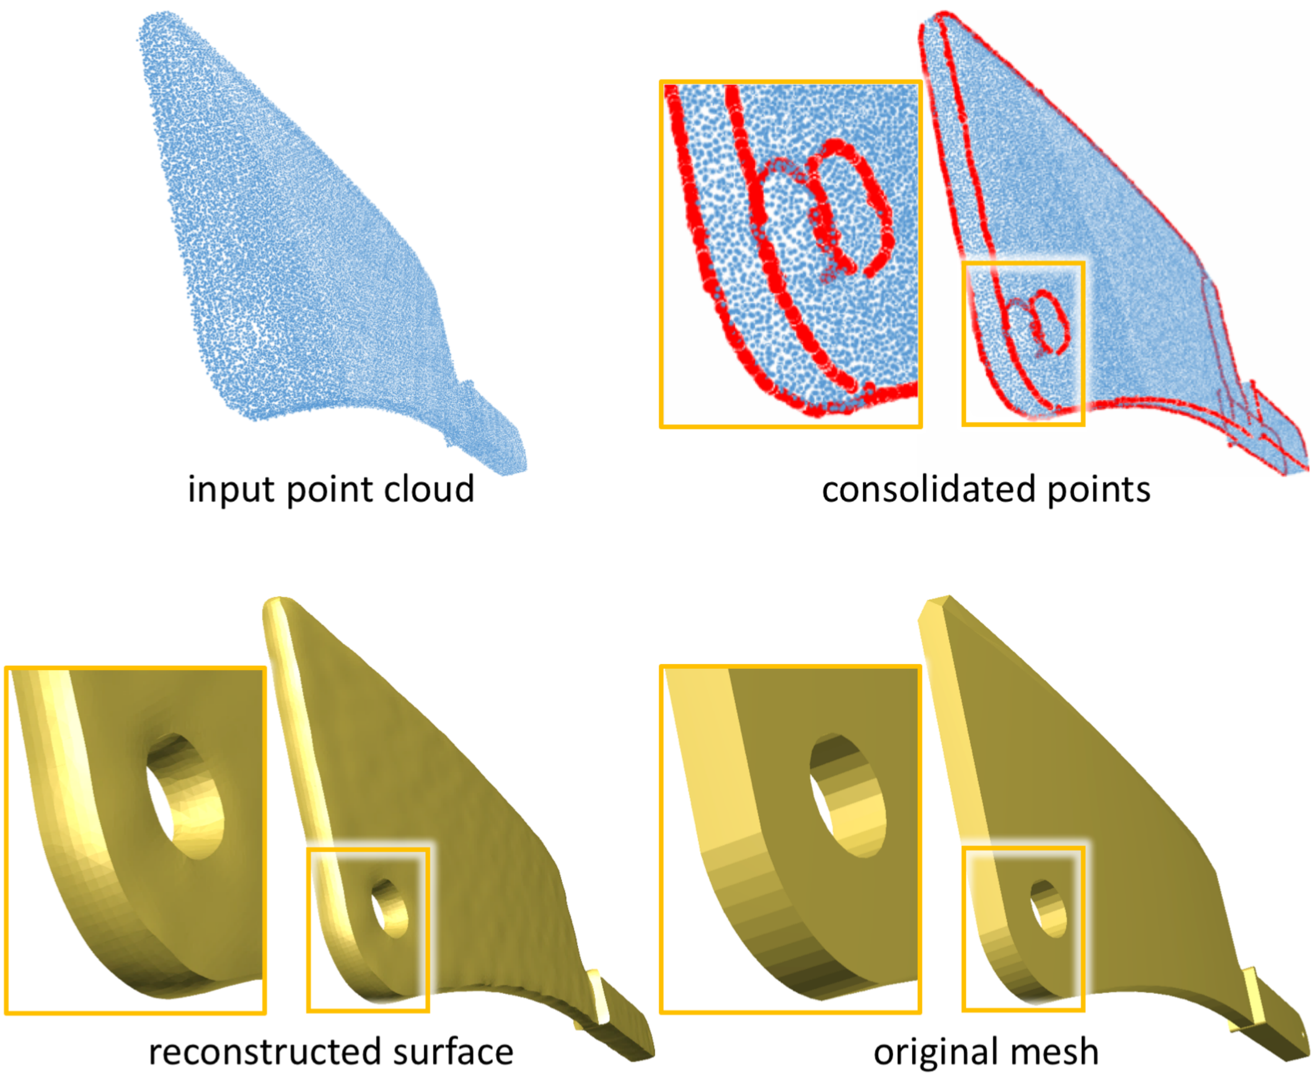

Figure 6. Point consolidation and surface reconstruction results produced with our method.

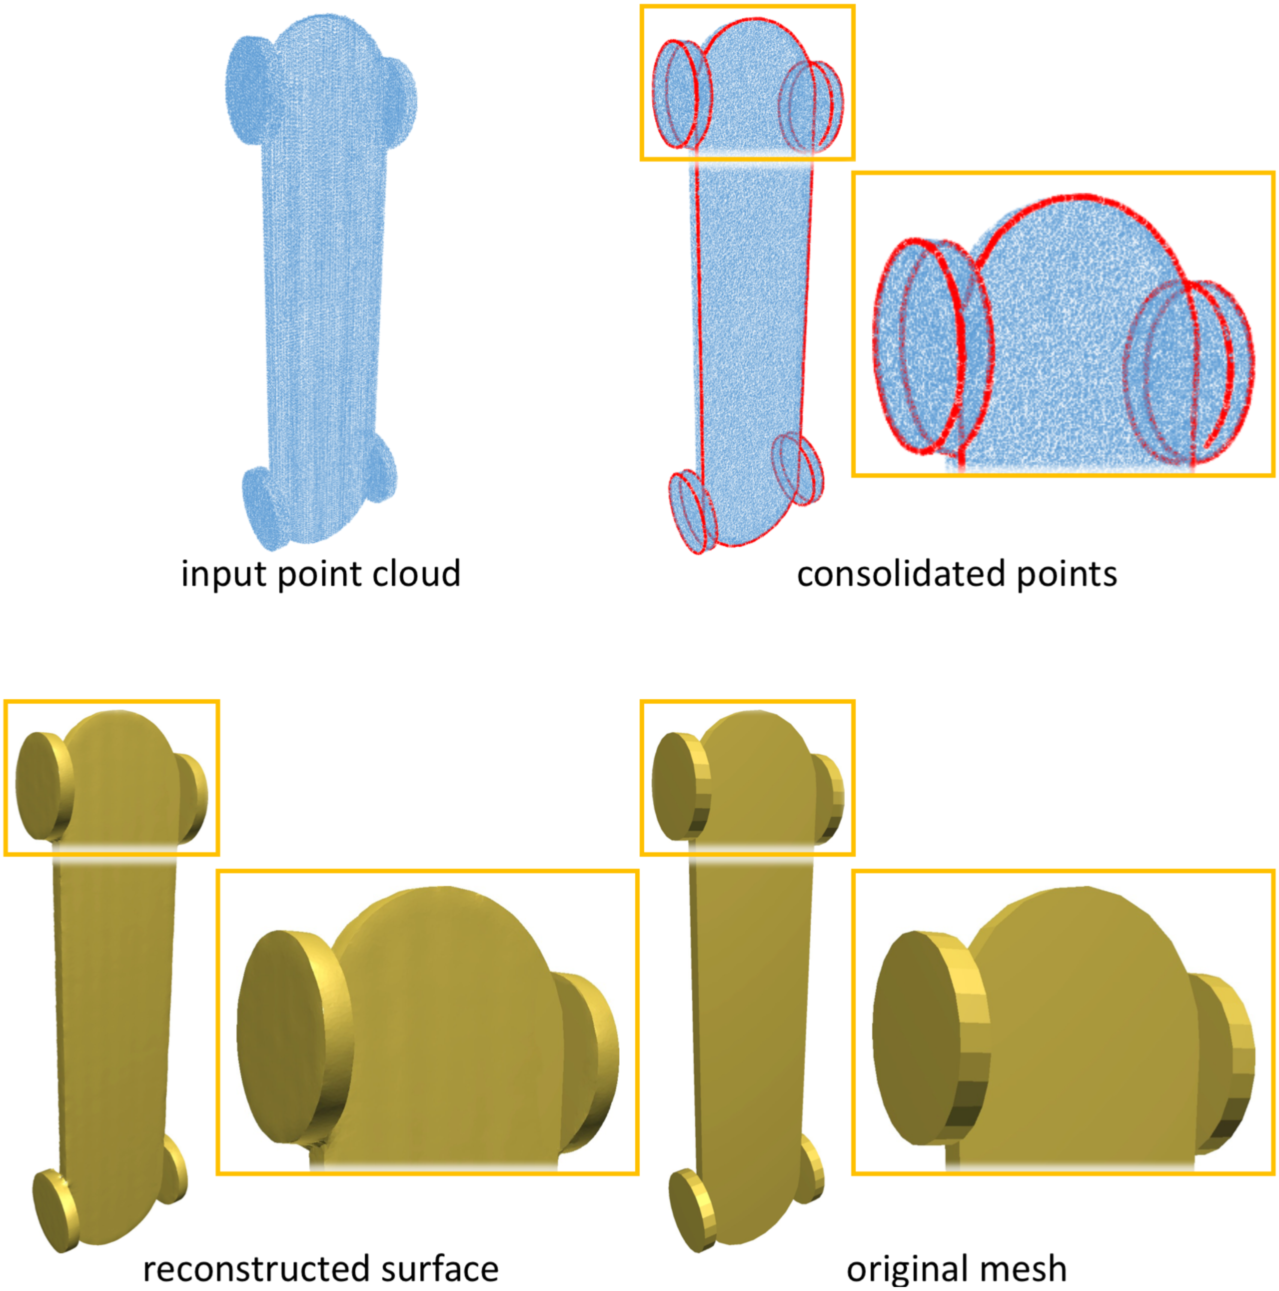

Figure 7. Point consolidation and surface reconstruction results produced with our method.

C.3. Visual comparison of point-to-surface distances

Figure 8 presents visual comparison of point-to-surface distances for results produced with EAR, PU-Net and our EC-Net on three different models (see also Table 1 in the submitted paper). In detail, we color each point based on its minimum distance (proximity) to the original mesh, i.e., the ground truth surface; see the color maps on the right. From these results, we can see that points in results produced with EC-Net are clearly darker, i.e., having lower point-to-surface distances.

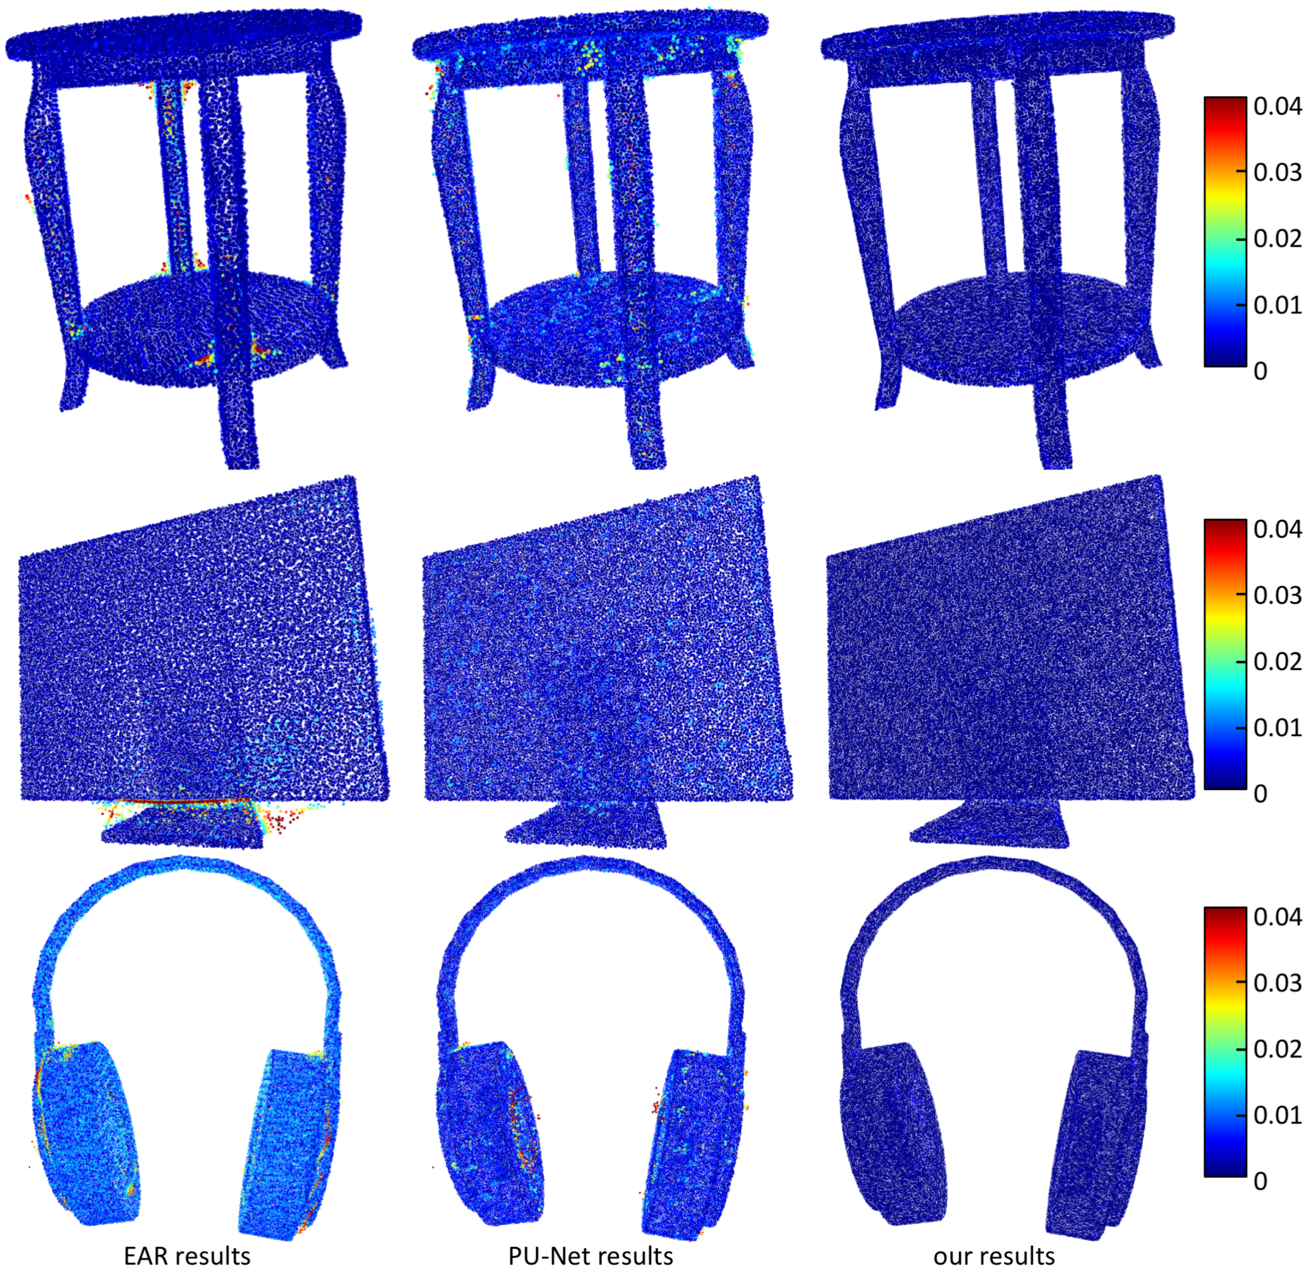

Figure 8. Visual comparison of point-to-surface distances for results produced with EAR, PU-Net and EC-Net. Each point is colored according to its minimum distance from the original (ground truth) mesh surface.

D. Surface Reconstruction Results from Point Clouds of Varying Number of Points

To study the ability of our method to handle input point clouds of varying number of point samples, we virtually scan the same 3D model to produce four point clouds of around 8k, 16k, 32k, and 64k point samples; see the top row in Figure 9. Then, we apply our method to each point cloud and produce a surface reconstruction result accordingly; see the bottom row in the figure. From the results, we can observe that for the input point cloud with only $\sim$8k points, the point distribution is severely non-uniform and inhomogeneous, with some of the points being very close to one another. For such a severe case, our method can still help reconstruct the surface and edges, yet the surface may not be very smooth.

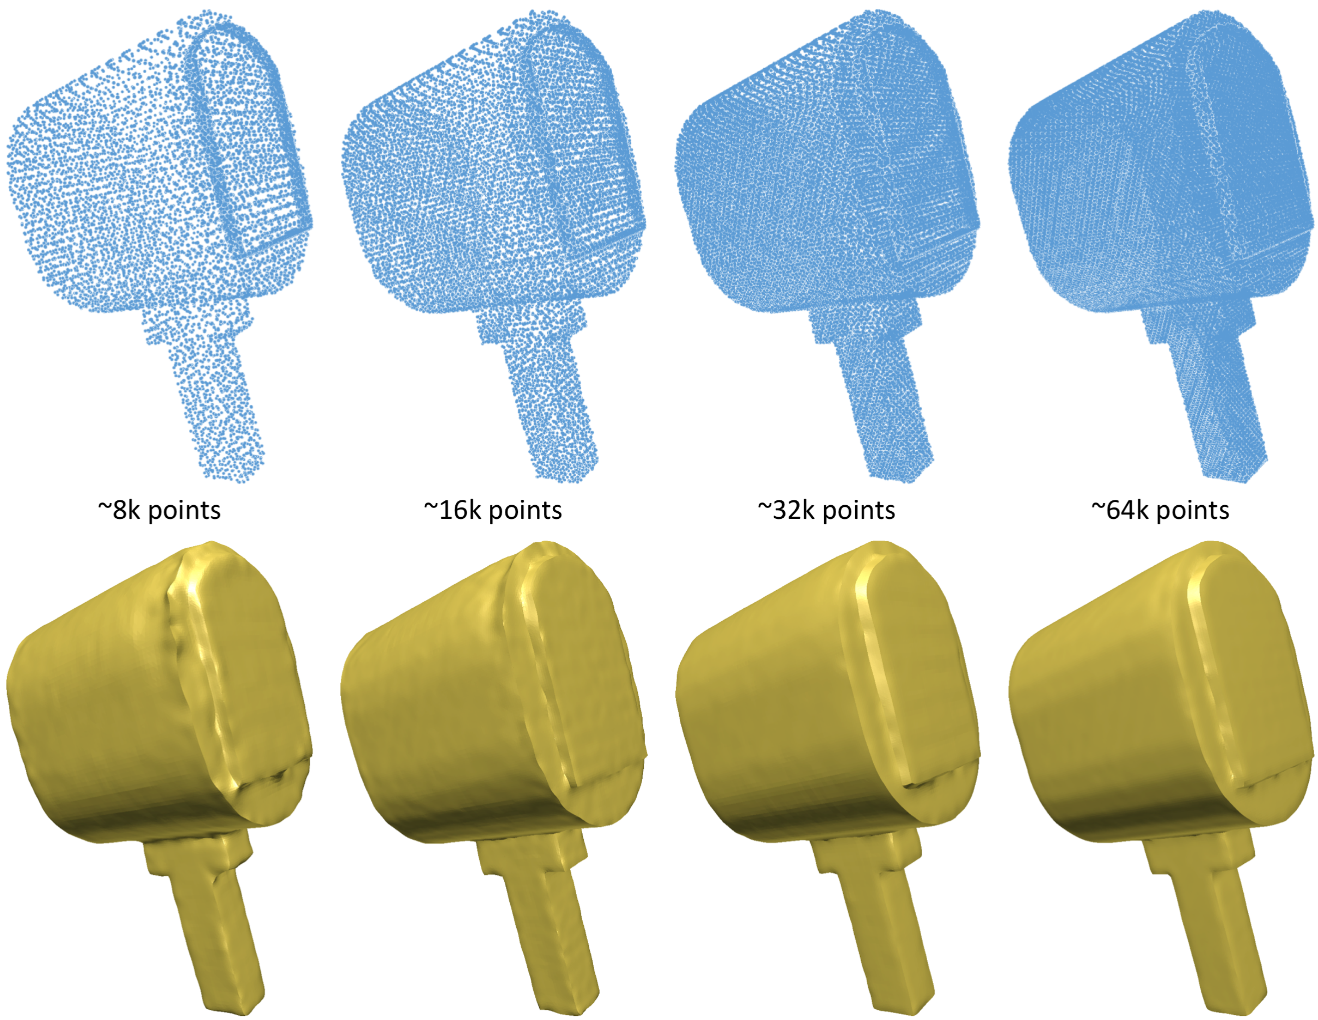

Figure 9. Surface reconstruction results (bottom) produced from point clouds of different number of points (top).

References

[1] Berger, M., Levine, J.A., Nonato, L.G., et al.: A benchmark for surface reconstruction. ACM Trans. on Graphics **32**(2) (2013) 20 6

[2] Ummenhofer, B., Brox, T.: Global, dense multiscale reconstruction for a billion points. Int. J. Comp. Vision (2017) 1–13 6