Appendix

A. Corruption Functions

We used the benchmark from [3] to introduce 15 common types of corruption occurring in real-world point clouds to evaluate the performance of our 3DD-TTA algorithm. Below is a brief explanation of each corruption function:

Uniform: Each point in the point cloud is perturbed by adding random noise sampled from a uniform distribution within the range of [-0.05, 0.05]. This type of noise introduces a uniform displacement across the points, simulating mild yet consistent distortions in all directions.

Gaussian: Points in the point cloud are randomly perturbed by adding Gaussian noise, where the noise is sampled from a normal distribution with a standard deviation of 0.03. This introduces subtle yet widespread deviations in the spatial positions of the points, simulating real-world sensor noise or environmental interference. Such perturbations can blur fine details and distort surface structures, challenging the model's ability to accurately interpret the point cloud.

Background: Generates 5% of the number of points within the range of -1 to 1 using a uniform distribution and concatenated them with the original point cloud.

Impulse: 10 percent of the points in the point cloud are randomly chosen and perturbed by adding impulse noise with random values of +0.1 and -0.1.

Upsampling: Additional points are generated by duplicating existing points in the point cloud. The newly generated points are perturbed versions of the original points with uniform corruption.

RBF: Point clouds are deformed using the Radial Basis Function [1].

Inverse RBF: The Radial Basis Function [1] and the resulting splines are inverted

Local Density Decrease: Five points are randomly selected from the point cloud as cluster centers, and their 100 nearest neighbors are identified. The density is then decreased by removing three-quarters of the points within these clusters.

Local Density Increase: Five local cluster centers are randomly selected, and their 100 nearest neighbors are identified. These clusters are preserved, while the remaining points are resampled to maintain the original total number of points. As a result, the density of points within these clusters is doubled compared to the rest of the point cloud.

Shear: The point clouds in the xy-plane undergo random transformations through compression or stretching, where the x and y coordinates are scaled by random factors ranging between -0.25 and 0.25. This process introduces geometric distortions that simulate non-uniform scaling in the 2D plane. Such perturbations can affect the overall structure of the point cloud, making it more challenging for models to retain key shape features.

Rotate: The point cloud is randomly rotated in 3D space by an angle between $\pm 15^{\circ}$. This type of corruption introduces subtle but impactful transformations that alter the orientation of the object while preserving its structural integrity. These small rotations are challenging for models to handle, as they must maintain consistent classification performance despite the varying orientations.

Cut-out: Five points are selected as cluster centers, and their 100 nearest neighbors are identified and removed from the point cloud.

FFD: Free Form Distortion (FFD) [2] (denoted as "dist" in tables) with 25 control points with a deformation distance in a range of ±0.5 is applied to the point clouds.

Occlusion: This corruption deletes portions of the point cloud using precomputed meshes and ray tracing from a random camera position [4]. The occluded points are effectively "hidden" or removed, simulating scenarios where parts of the object are not visible due to obstruction. This tests the model's ability to adapt and reconstruct partially visible or occluded objects.

LiDAR: The LiDAR capturing of point clouds is simulated, adding inaccuracies such as occlusion, reflection, and noise.

B. SCD Distance Hyperparameter

We conducted a thorough analysis of the λ parameter, which controls the tightness of the SCD distance (l^{λ}_{CD}) and influences the similarity between the reconstructed and original corrupted latent points. A higher λ value increases this similarity, but choosing the optimal value is critical for achieving effective test-time adaptation. For this analysis, we again focused on background noise, as it requires more denoising steps, making it an ideal case for studying the impact of λ. Figure 6 (left) shows that the optimal value of $\lambda = 0.96$ strikes a balance: lower values fail to guide the diffusion process effectively, while values approaching 1 cause the diffusion process to overfit to the corrupted latent points, including outliers, which hampers the final classification performance. This over-reliance on the corrupted points leads to a degradation in adaptation performance, as the model becomes overly focused on the noise rather than the underlying clean structure. Hence, choosing $\lambda = 0.96$ avoids these extremes, offering a more robust adaptation. The implementation codes will be made publicly available upon the acceptance of this paper. We plan to release them on a dedicated repository to ensure transparency and facilitate further research in this domain.

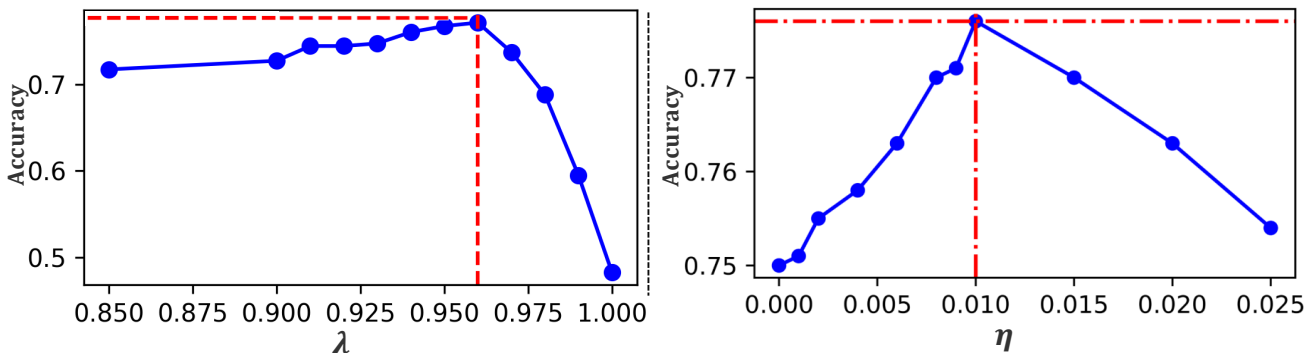

Figure 6. (**left**) Accuracy of the source classifier after adaptation across different values for the SCD hyperparameter (λ). (**right**) Accuracy of the source classifier after adaptation using various shape latent updating coefficients (η).

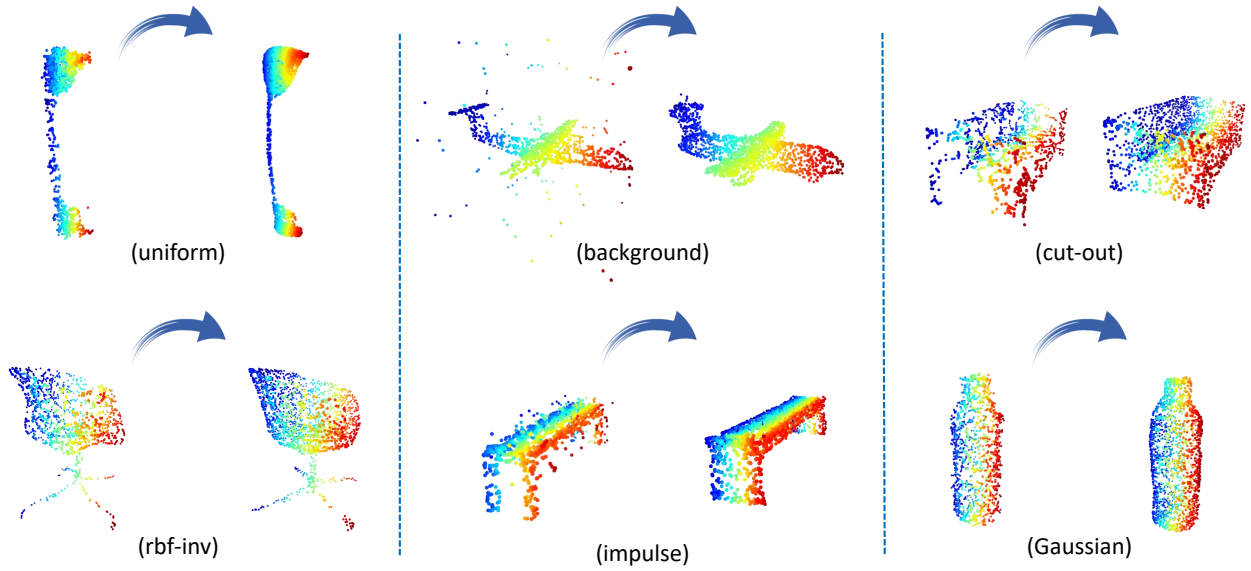

Figure 7. Qualitative assessment of the proposed test-time adaptation across various corruptions shows that the 3DD-TTA model excels in resolving noise corruptions, including uniform, background, impulse, and Gaussian noise. The model can also fill in missing parts within the point cloud and is effective in addressing inverse RBF deformation.

C. Impact of Shape Latent Updating:

We conducted experiments to explore the impact of updating the shape latent via gradient descent during the denoising process on the overall performance of the test-time adaptation (TTA) task. In these experiments, we fixed other hyperparameters and performed a grid search for different values of the shape latent updating coefficient (η). We specifically used background noise as a special case, as it requires more denoising steps, allowing for a more detailed analysis. Figure 6 (right) illustrates the performance of the source classifier across different values of η, with $\eta = 0$ representing no shape latent update. The results indicate that for all values of η greater than zero, there is a noticeable improvement in performance, with $\eta = 0.01$ achieving the highest classification accuracy. Consequently, we set η to 0.01 in all subsequent experiments. This figure further demonstrates that an optimal choice of η can enhance the classification rate by up to 2.5 %.

D. More Reconstruction Results

Figures 7 and 8 present additional examples of corrupted point clouds and their adapted versions produced by the proposed 3DD-TTA method. The model demonstrates effectiveness in reconstructing the point clouds by resolving inconsistencies in density and eliminating noise. While the method is not flawless in handling all transformation-based corruptions, it shows the ability to address certain deformations, such as RBF-inv and shear, to a notable degree.

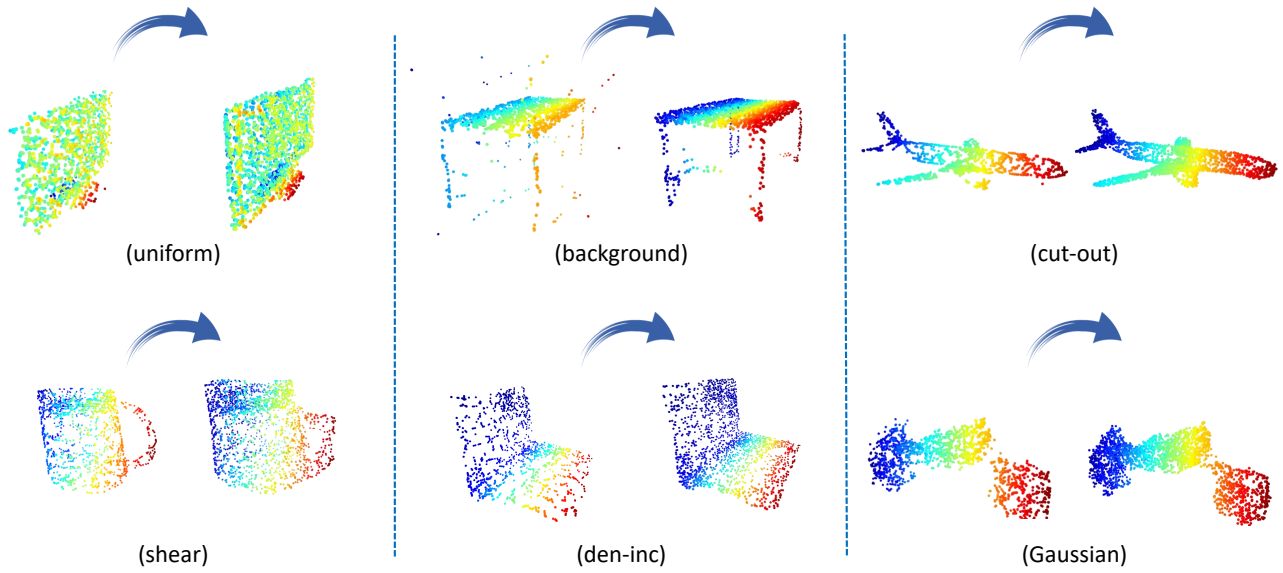

Figure 8. More examples for the qualitative assessment of the proposed test-time adaptation across various corruptions.

It is worth highlighting that this reconstruction is achieved within a training-free pipeline, performed entirely as test-time adaptation, which significantly enhances its practical value. Unlike other point cloud reconstruction and denoising methods that rely on supervised training with paired corrupted and ground-truth point clouds, our approach adapts the input point clouds without requiring access to any paired training data, making it well-suited for real-world applications where labeled data is scarce or unavailable. This ability to generalize without explicit training offers a unique advantage in dynamic and unpredictable environments.

References

[1] Davide Forti and Gianluigi Rozza. Efficient geometrical parametrisation techniques of interfaces for reduced-order modelling: application to fluid–structure interaction coupling problems. *International Journal of Computational Fluid Dynamics*, 28(3-4):158–169, 2014.

[2] Thomas W Sederberg and Scott R Parry. Free-form deformation of solid geometric models. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, pages 151–160, 1986.

[3] Jiachen Sun, Qingzhao Zhang, Bhavya Kailkhura, Zhiding Yu, Chaowei Xiao, and Z Morley Mao. Benchmarking robustness of 3d point cloud recognition against common corruptions. *arXiv preprint arXiv:2201.12296*, 2022.

[4] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3d: A modern library for 3d data processing. *arXiv preprint arXiv:1801.09847*, 2018.