Neural Network-Based Point Cloud Completion for ICP Registration

Roman Vozivoda[1,*,†], Jan Glaser[1,†] and Marcel Jirina[1,†]

[1]Czech Technical University in Prague, Czech Republic

Abstract

Incomplete 3D point clouds are a common challenge in real-world scanning applications, where occlusions or sensor limitations can lead to missing geometric information. We introduce a neural-network-based point cloud completion framework to reconstruct the missing regions. The framework predicts only the missing points and integrates them with the original input. The impact of the reconstructed geometry on registration performance is evaluated using standard registration metrics under varying levels of point cloud incompleteness. Experimental results demonstrate that the proposed completion strategy substantially improves registration accuracy compared with registration performed directly on the incomplete point clouds. The greatest improvement in fitness and correspondence set size is observed when about 60% of the points are missing, whereas the inlier RMSE degrades at high occlusion levels. These findings indicate that point cloud completion can serve as an effective preprocessing step for improving the accuracy of 3D point cloud registration.

Keywords

3D point cloud, point cloud completion, point cloud registration, ICP, neural network, PointNet, self-attention, encoder-decoder, Chamfer Distance, ShapeNet

1. Introduction

The processing of 3D point clouds is an important topic in computer vision and 3D graphics. Point clouds are often created using laser scanners or depth cameras, and can contain noise or missing parts caused by imperfect scanning conditions. In this paper we use neural networks with an encoder–decoder architecture for the shape completion of point clouds. Our goal is a model that can automatically reconstruct missing areas in a 3D point cloud. We then investigate whether it can improve the quality of point cloud registration.

Deep neural networks have recently become the dominant approach to point cloud completion. The Point Completion Network (PCN) introduced a coarse-to-fine generation strategy in which a sparse global shape estimate is progressively refined into a dense point cloud [1]. Transformer-based approaches such as PoinTr reformulated completion as a set-to-set prediction problem and demonstrated improved recovery of large missing regions through attention mechanisms [2]. Similarly, SnowflakeNet employed hierarchical point generation, progressively refining local geometric structures through multiple upsampling stages [3].

A limitation shared by all of these methods is that they regenerate the whole point cloud, including the regions that are already present in the input. As a result, the original points are replaced by reconstructed approximations, introducing deviations from the acquired sensor data. While such modifications may have limited influence on visual quality, they can negatively affect applications that rely on preserving the exact spatial locations of measured points, such as point cloud registration.

This work addresses the above limitation through a point-preserving completion strategy. Rather than reconstructing the complete shape, the proposed Missing-Point Completion Network (MPCN) predicts only the missing geometric structures and subsequently combines them with the unchanged

ITAT'26: Information Technologies – Applications and Theory, September 25–29, 2026, Biele Karpaty, Slovakia

*Corresponding author.

†These authors contributed equally.

vozivrom@fit.cvut.cz (R. Vozivoda); glaseja1@cvut.cz (J. Glaser); marcel.jirina@fit.cvut.cz (M. Jirina)

0000-0002-6686-1511 (M. Jirina)

© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

published 2026-11-09

CEUR-WS.org/Vol-4323/paper10.pdf

input point cloud. As a result, all original measurements remain preserved while additional points are generated to the point cloud.

The main contributions of this work are as follows:

- a point-preserving completion network that predicts only the missing points while preserving all original measurements;
- an evaluation showing that this completion improves ICP registration.

The remainder of the paper is organized as follows. Section 2 introduces the newly proposed Missing-Point Completion Network (MPCN). Section 3 describes the evaluation methodology. Section 4 presents the experiments. Section 5 discusses the obtained results and relevant aspects. Finally, Section 6 concludes the paper.

2. Proposed Missing-Point Completion Network

The proposed network employs an encoder–decoder architecture illustrated in Fig. 1. The encoder is based on a simplified PointNet feature extractor [4].

The network takes a fixed number of $N = 1024$ input points. Point clouds with more points are randomly subsampled to N, and point clouds with fewer points are filled up to N by repeating some of their points. The input partial point cloud, represented as an $N \times 3$ tensor, is first aligned by a T-Net and then processed by a sequence of shared multi-layer perceptrons that progressively lift the three input coordinate channels to a 1024-dimensional per-point feature. A max-pooling operation across all points then aggregates these per-point features into a single 1024-dimensional global feature vector that summarizes the whole shape. From this global vector, a fully connected coarse decoder generates 256 coarse seed points, which are refined by a self-attention module that predicts a per-point displacement. This module is a multi-head self-attention block with 4 heads and an embedding dimension of 128. Finally, an upsampling MLP expands each refined seed into four points and maps them back to the three-dimensional coordinate space, producing the 1024 predicted missing points, which are concatenated with the original 1024-point input to form the 2048-point completed point cloud.

The loss function that is aimed to be minimized is

$$L = L_{\text{fine}} + \alpha\, L_{\text{coarse}} \tag{1}$$

where L_{fine} is the Chamfer distance [5] between the completed point cloud and the ground-truth point cloud, L_{coarse} is the Chamfer distance between the coarse point cloud and the ground-truth point cloud, and α is a weighting coefficient that controls how much the coarse loss affects the total loss.

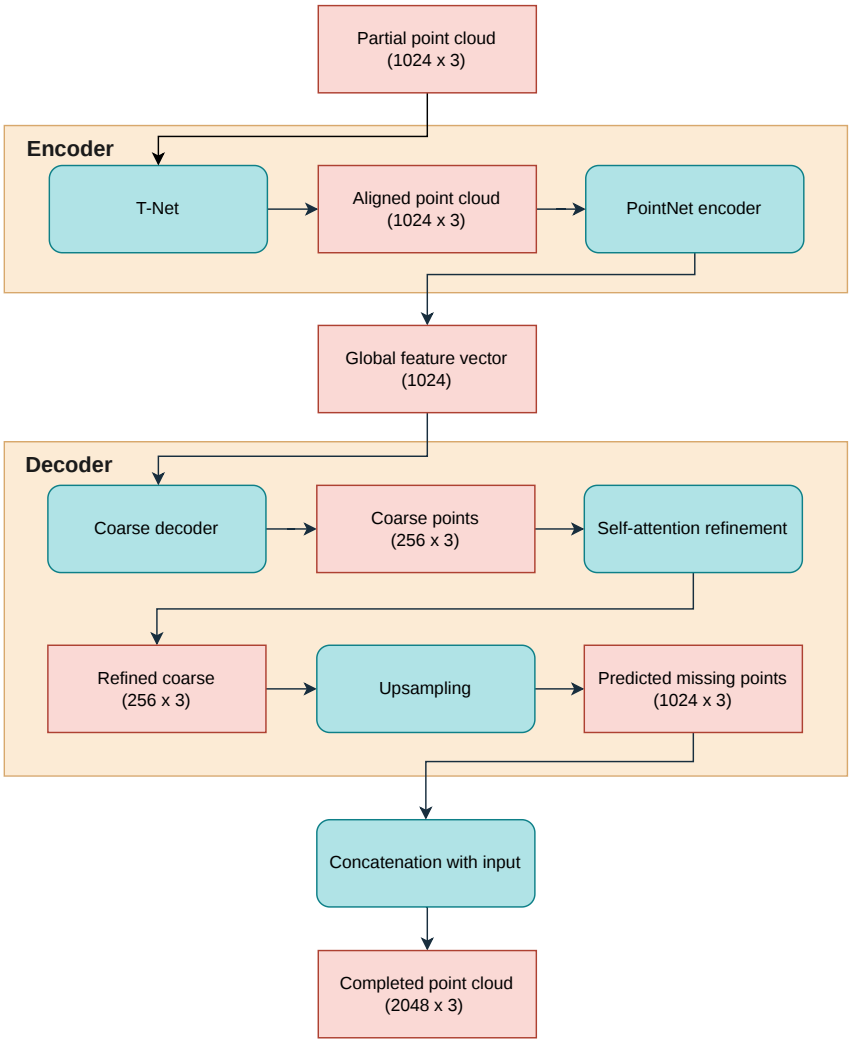


Figure 1: Architecture of the proposed input-preserving completion network. The partial input is first aligned by a T-Net and encoded by a PointNet-style encoder into a 1 024-dimensional global feature vector. The decoder predicts a coarse point cloud, refines it with self-attention, and upsamples it.

3. Evaluation methodology

3.1. Simulating missing points

A real depth camera or LiDAR observes an object from a single viewpoint, so it captures only the surfaces that face the sensor, while the opposite side of the object stays hidden. To reproduce this effect, a viewpoint-based point removal operator $\Phi(\cdot)$ is introduced. Let

$$\mathbf{v} \in \mathbb{R}^3, \qquad \|\mathbf{v}\| = 1 \tag{2}$$

denote a randomly generated unit viewpoint vector and let

$$P = \{\mathbf{p}_i\}_{i=1}^{N}, \qquad \mathbf{p}_i \in \mathbb{R}^3 \tag{3}$$

represent the complete point cloud.

For each point $\mathbf{p}_i$, a projection score is computed as the projection onto the viewpoint direction

$$s_i = \mathbf{p}_i^\top \mathbf{v}, \tag{4}$$

where larger values of s_i correspond to points that are more closely aligned with the selected viewpoint direction. Before the operator is applied, the point cloud is normalised to a unit sphere by centring it at its centroid and scaling it so that its farthest point lies at unit distance from the origin. The score s_i is then computed on this normalised, origin-centred cloud, so that removing the points with the largest projections discards every point lying beyond a plane orthogonal to $\mathbf{v}$, reproducing the flat, half-space occlusion boundary of a real depth sensor.

Let

$$\mathcal{S}_K = \{s_i \mid s_i \text{ belongs to the } K \text{ largest projection values}\} \tag{5}$$

denote the set of the K largest projection scores.

The viewpoint-based point removal operator is then defined as

$$\Phi(P, \mathbf{v}, K) = P \setminus \{\mathbf{p}_i \mid s_i \in \mathcal{S}_K\}. \tag{6}$$

The resulting partial point cloud is obtained as

$$P_{\text{partial}} = \Phi(P, \mathbf{v}, K). \tag{7}$$

Consequently, points located in the direction of the selected viewpoint are removed, which results in a missing region in the point cloud, that resembles an occlusion caused by limited sensor visibility.

3.2. Registration metrics

The influence of point cloud completion on registration performance is evaluated using three commonly adopted registration metrics: fitness, inlier RMSE, and correspondence set size. These measures quantify the degree of overlap between two registered point clouds, the geometric consistency of the established correspondences, and the number of matched point pairs, respectively.

Let

$$P = \{\mathbf{p}_i\}_{i=1}^{N} \tag{8}$$

and

$$Q = \{\mathbf{q}_j\}_{j=1}^{M} \tag{9}$$

denote the source and target point clouds after registration, where N and M are the numbers of points in the source and target clouds, respectively. Let

$$C = \{(\mathbf{p}, \mathbf{q})\} \tag{10}$$

represent the set of inlier correspondences determined within a predefined distance threshold.

The fitness metric evaluates the fraction of source points that participate in valid correspondences and can therefore be interpreted as a measure of overlap between the registered point clouds [6, 7]. It is defined as

$$\text{fitness} = \frac{|C|}{|P|}. \tag{11}$$

Larger fitness values indicate a higher degree of alignment and overlap between the source and target point clouds.

The geometric accuracy of the established correspondences (Eq. 10) is quantified by the inlier RMSE [6, 8], defined as

$$\mathrm{RMSE}_{\mathrm{inlier}} = \sqrt{\frac{1}{|C|} \sum_{(\mathbf{p},\mathbf{q})\in C} \|\mathbf{p}-\mathbf{q}\|^2}. \tag{12}$$

Lower RMSE values correspond to smaller distances between the corresponding points and therefore indicate more accurate registration.

The correspondence set size is given by

$$N_C = |C|, \tag{13}$$

and denotes the total number of inlier correspondences. Larger correspondence sets generally reflect increased overlap and provide a stronger geometric basis for registration.

The computation of all three metrics depends on an evaluation distance threshold used to determine inlier correspondences. Only point pairs whose Euclidean distance falls below the specified threshold are included in the set C. In all experiments the threshold is set to 0.05. All point clouds are normalized to the unit sphere, so this value corresponds to $5\,\%$ of the radius of the object.

4. Experimental analysis

The experiments were performed on eight object categories from the ShapeNet [9] dataset, as shown in Fig. 2. These are the eight categories of the completion benchmark introduced by PCN [1]. The eight categories contain 31 774 models in total. Each category was split into $70\,\%$ training, $15\,\%$ validation and $15\,\%$ test data, which gives 22 243 training models.

In all experiments, the model was trained for up to 100 epochs in mini-batches of 64 full point clouds and evaluated on the validation set after every epoch. The network was trained with the Adam optimizer, using an initial learning rate of 10^{-3} that was halved every 30 epochs. The whole network has approximately 4.87 million trainable parameters. Training was performed on two NVIDIA Tesla T4 GPUs and continued until the convergence of the validation loss, as seen in Fig. 3. For the coefficient α in Eq. (1), the value 0.1 was chosen by a grid search over $\{1.0, 0.1, 0.01\}$.

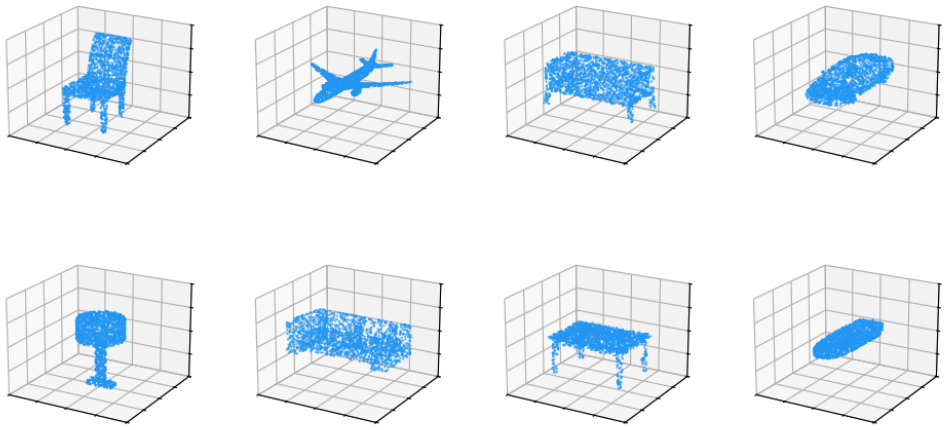


Figure 2: Representative point cloud from each of the eight test categories, chair, airplane, cabinet, car, lamp, sofa, table and watercraft [9].

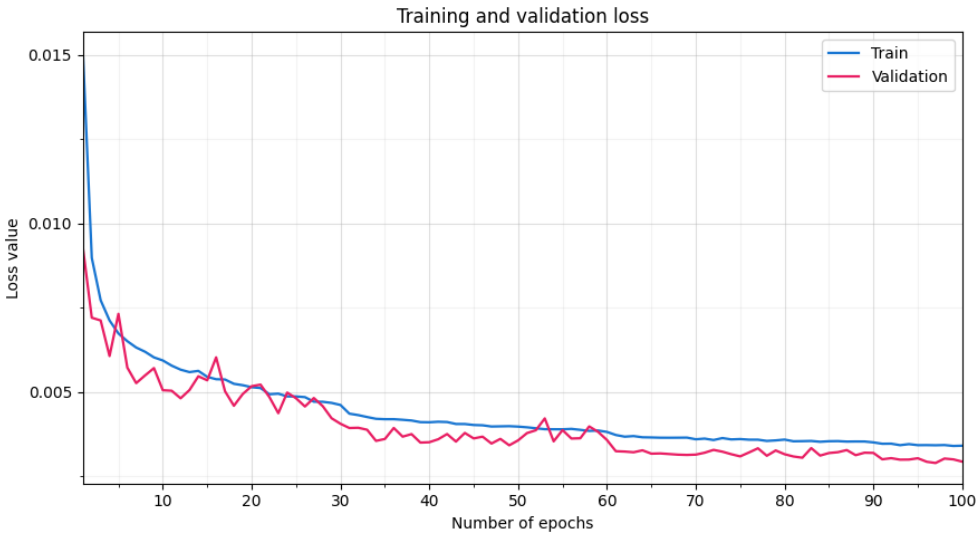


Figure 3: Training and validation loss over epochs for the proposed MPCN on the ShapeNet dataset [9].

To create incomplete point clouds, points were removed with the viewpoint-based removal operator defined in Eq. (6). The same fraction of points is removed from the source and the target point cloud, but the viewpoint vector $\mathbf{v}$ is drawn independently for each of them, so that the removed regions differ between the two point clouds. Gaussian noise with a standard deviation of 0.05 is then added to each point cloud to approximate sensor noise, and each point cloud is resampled to 1024 points.

We evaluated the performance of ICP [10] in two different settings using the same source–target pairs. In the first setting, both point clouds are incomplete. In the second setting, both point clouds are first completed by the proposed completion network.

The results presented in Table 1 are averaged over 800 samples (100 samples per category across 8 test categories) from the ShapeNet dataset [9].

Table 1
Comparison of registration of partial point clouds and point clouds completed by the MPCN, in terms of fitness, inlier RMSE, and correspondence set size, reported as mean $\pm$ standard deviation over 800 samples per occlusion ratio (100 samples per category across 8 test categories). The points were removed using the viewpoint-based point removal operator. Higher fitness and correspondence set size are better, lower RMSE is better.

	Partial point clouds			Completed point clouds		
Occ. ratio (%)	Fitness	RMSE	Corresp. size	Fitness	RMSE	Corresp. size
20	0.612 ± 0.149	0.0332 ± 0.0028	627 ± 153	0.807 ± 0.100	0.0314 ± 0.0031	$1\,652 \pm 205$
40	0.527 ± 0.183	0.0324 ± 0.0029	540 ± 187	0.787 ± 0.108	0.0315 ± 0.0032	$1\,611 \pm 221$
60	0.368 ± 0.251	0.0307 ± 0.0078	376 ± 257	0.685 ± 0.143	0.0320 ± 0.0031	$1\,403 \pm 293$
80	0.261 ± 0.277	0.0217 ± 0.0147	268 ± 284	0.391 ± 0.230	0.0324 ± 0.0031	800 ± 471

Table 2
Relative change in registration after completion, in terms of fitness, inlier RMSE, and correspondence set size, averaged over 800 samples per occlusion ratio. The values are changes with respect to the partial point clouds, so a positive value always means that registration was better on the completed point clouds.

Occ. ratio (%)	Δ Fitness (%)	Δ RMSE (%)	Δ Corresp. size (%)
20	31.80	5.44	163.59
40	49.23	3.00	198.43
60	86.65	-4.31	273.21
80	49.54	-49.50	199.03

We have also performed evaluation of registration accuracy on per category, in terms of fitness, inlier RMSE and correspondence set size. The points were removed using the viewpoint-based removal operator (see Eq. (6)). The results are shown in Table 3.

Table 3
Per-category comparison of registration of partial point clouds and point clouds completed by the MPCN, in terms of fitness, inlier RMSE, and correspondence set size, averaged over 100 samples per category, for viewpoint-based point removal. Higher fitness and correspondence set size are better, lower RMSE is better.

		Partial point cloud			Completed point cloud		
Category	Occ. ratio (%)	Fitness	RMSE	Corresp. size	Fitness	RMSE	Corresp. size
Chair	20	0.5848	0.0340	599	0.7744	0.0324	1 586
	40	0.5061	0.0332	518	0.7598	0.0322	1 556
	60	0.3418	0.0325	350	0.6628	0.0321	1 357
	80	0.1747	0.0219	179	0.3395	0.0331	695
Airplane	20	0.7435	0.0309	761	0.8761	0.0289	1 794
	40	0.6236	0.0300	639	0.8554	0.0288	1 752
	60	0.4333	0.0304	444	0.7723	0.0295	1 582
	80	0.3321	0.0231	340	0.4997	0.0314	1 024
Cabinet	20	0.4542	0.0357	465	0.7172	0.0338	1 469
	40	0.4068	0.0349	416	0.6993	0.0340	1 432
	60	0.2707	0.0324	277	0.5542	0.0347	1 135
	80	0.1541	0.0204	158	0.3187	0.0338	653
Car	20	0.5387	0.0347	552	0.7898	0.0331	1 618
	40	0.4758	0.0340	487	0.7813	0.0332	1 600
	60	0.2481	0.0303	254	0.6125	0.0344	1 254
	80	0.2530	0.0223	259	0.3112	0.0336	637
Lamp	20	0.7259	0.0303	743	0.8609	0.0281	1 763
	40	0.6007	0.0296	615	0.8187	0.0284	1 677
	60	0.4773	0.0269	489	0.7352	0.0292	1 506
	80	0.3953	0.0223	405	0.5472	0.0294	1 121
Sofa	20	0.5238	0.0349	536	0.7428	0.0337	1 521
	40	0.4516	0.0342	462	0.7272	0.0337	1 489
	60	0.3244	0.0315	332	0.6281	0.0338	1 286
	80	0.2285	0.0218	234	0.3220	0.0334	659
Table	20	0.6009	0.0339	615	0.8021	0.0320	1 643
	40	0.5031	0.0331	515	0.7822	0.0318	1 602
	60	0.3499	0.0323	358	0.6884	0.0321	1 410
	80	0.2630	0.0225	269	0.3884	0.0327	796
Watercraft	20	0.7368	0.0315	754	0.8902	0.0297	1 823
	40	0.6300	0.0308	645	0.8768	0.0298	1 796
	60	0.4373	0.0295	448	0.7829	0.0306	1 603
	80	0.3157	0.0213	323	0.4645	0.0320	951

Table 4
Per-category relative change in fitness, inlier RMSE, and correspondence set size after completion using viewpoint-based point removal, averaged over 100 samples per category. The values are signed so that a positive number always indicates an improvement after completion, i.e. a higher fitness, a higher correspondence set size, or a lower inlier RMSE than on the partial point clouds. Negative values therefore indicate that the metric became worse after completion.

Category	Occ. ratio (%)	Δ Fitness (%)	Δ RMSE (%)	Δ Corresp. size (%)
Chair	20	32.35	4.79	164.67
	40	50.09	3.09	200.14
	60	93.89	1.15	287.74
	80	93.48	-50.94	286.99
Airplane	20	17.85	6.67	135.70
	40	37.16	3.84	174.29
	60	78.26	3.02	256.63
	80	50.59	-36.08	201.15
Cabinet	20	57.95	5.13	215.89
	40	71.93	2.73	243.89
	60	104.84	-7.10	309.70
	80	106.29	-66.36	312.55
Car	20	46.58	4.46	193.11
	40	64.19	2.26	228.41
	60	146.76	-13.47	393.39
	80	23.40	-50.39	146.91
Lamp	20	18.60	7.43	137.17
	40	36.36	3.78	172.72
	60	54.01	-8.42	207.98
	80	38.46	-32.15	176.92
Sofa	20	41.77	3.58	183.52
	40	61.03	1.50	221.99
	60	93.83	-7.22	287.77
	80	40.83	-53.07	181.79
Table	20	33.40	5.69	166.81
	40	55.44	3.78	210.87
	60	96.77	0.47	293.55
	80	47.45	-45.16	194.91
Watercraft	20	20.82	5.76	141.64
	40	39.17	3.41	178.34
	60	79.03	-3.97	258.06
	80	46.86	-49.59	193.75

5. Discussion

Table 2 summarises the relative changes after completion. Point cloud completion improves fitness and correspondence set size at every occlusion ratio, and the largest improvement is reached at 60 % occlusion, where fitness increases by 86.7 % and the correspondence set size by 273.2 %. The inlier RMSE behaves differently: it improves by 5.4 % at 20 % occlusion and degrades by 49.5 % at 80 %. Tables 3 and 4 show that the same pattern holds for every category individually. At 80 % occlusion the relative improvement drops and the absolute fitness becomes too low for a reliable point cloud registration. As shown in Table 4, cabinet has the largest improvement in ICP metrics after completion, because it has a simple box-like geometry that is easy for the network to complete. Lamp has the smallest improvement, because it has complex, distinctive shapes which are harder for the network to reconstruct accurately. Examples of completion results obtained for cabinet and lamp objects are presented in Fig. 4 and Fig. 5, respectively. The completion quality varies across object categories and appears to be strongly influenced by geometric complexity. Cabinet objects are characterized by regular structures and dominant planar surfaces, enabling more reliable reconstruction of missing regions. In contrast, lamp objects frequently contain thin components and intricate geometric details that are more difficult to reproduce from the partial point clouds only.

This trend is consistent with the registration experiments, where categories exhibiting higher completion quality also achieve greater registration improvements. The cabinet category therefore benefits more substantially from completion-assisted registration than the lamp category, as seen in Table 4.

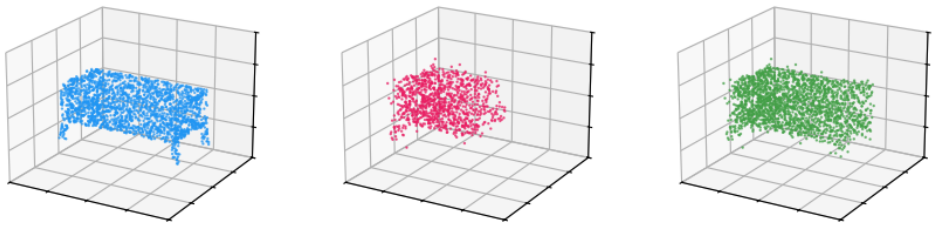


Figure 4: A cabinet point cloud from the ShapeNet dataset [9]. a) ground-truth point cloud, b) partial point cloud, c) point cloud completed by the MPCN.

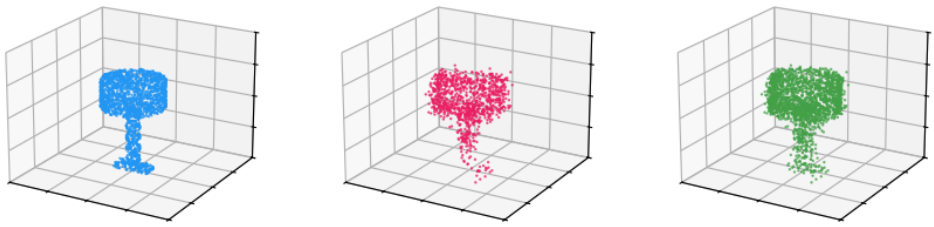


Figure 5: A lamp point cloud from the ShapeNet dataset [9]. a) ground-truth point cloud, b) partial point cloud, c) point cloud completed by the MPCN.

6. Conclusions

This paper investigated the use of neural point cloud completion as a preprocessing step for point cloud registration. A point-preserving completion network was proposed that predicts only the missing regions of a partial point cloud while preserving all originally observed points. Unlike conventional

completion approaches that reconstruct the entire shape, the proposed method maintains the exact sensor measurements and augments them with geometrically plausible predictions in occluded regions.

Experimental evaluation on ShapeNet data demonstrated that the completed point clouds consistently improved ICP registration accuracy compared to registration performed directly on the incomplete point clouds. Increased fitness and larger correspondence sets were observed at all evaluated occlusion levels, with the largest improvement at $60\,\%$ occlusion, whereas the inlier RMSE improved only at low and moderate occlusion and degraded once most of the object was missing. We conclude that the completion of point clouds is beneficial for point set registration.

Acknowledgments

This work was supported in part by the Student Grant Competition (SGS) of Czech Technical University in Prague, project No. SGS26/186/OHK3/3T/18.

Declaration on Generative AI

During the preparation of this work, the authors used Claude and GPT-5 in order to: Grammar and spelling check, paraphrase and reword. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] W. Yuan, T. Khot, D. Held, C. Mertz, M. Hebert, PCN: Point completion network, in: 2018 International Conference on 3D Vision (3DV), 2018, pp. 728–737. URL: https://doi.org/10.1109/3DV.2018.00088. doi:10.1109/3DV.2018.00088.

[2] X. Yu, Y. Rao, Z. Wang, Z. Liu, J. Lu, J. Zhou, PoinTr: Diverse point cloud completion with geometry-aware transformers, in: 2021 IEEE/CVF International Conference on Computer Vision (ICCV), 2021, pp. 12478–12487. URL: https://doi.org/10.1109/ICCV48922.2021.01227. doi:10.1109/ICCV48922.2021.01227.

[3] P. Xiang, X. Wen, Y.-S. Liu, Y.-P. Cao, P. Wan, W. Zheng, Z. Han, SnowflakeNet: Point cloud completion by snowflake point deconvolution with skip-transformer, in: 2021 IEEE/CVF International Conference on Computer Vision (ICCV), 2021, pp. 5479–5489. URL: https://doi.org/10.1109/ICCV48922.2021.00545. doi:10.1109/ICCV48922.2021.00545.

[4] C. R. Qi, H. Su, K. Mo, L. J. Guibas, PointNet: Deep learning on point sets for 3D classification and segmentation, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 77–85. URL: https://doi.org/10.1109/CVPR.2017.16. doi:10.1109/CVPR.2017.16.

[5] H. Fan, H. Su, L. Guibas, A point set generation network for 3D object reconstruction from a single image, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 2463–2471. URL: https://doi.org/10.1109/CVPR.2017.264. doi:10.1109/CVPR.2017.264.

[6] Q.-Y. Zhou, J. Park, V. Koltun, Open3D: A modern library for 3D data processing, arXiv:1801.09847, 2018. URL: https://arxiv.org/abs/1801.09847.

[7] Open3D, ICP registration — open3d documentation, online, ???? URL: http://www.open3d.org/docs/release/tutorial/pipelines/icp_registration.html, [Online; accessed 31-October-2023].

[8] J. Baek, J. Park, S. Cho, C. Lee, 3d global localization in the underground mine environment using mobile lidar mapping and point cloud registration, Sensors 22 (2022) 2873.

[9] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, F. Yu, ShapeNet: An information-rich 3D model repository, arXiv:1512.03012, 2015. URL: https://arxiv.org/abs/1512.03012.

[10] P. J. Besl, N. D. McKay, A method for registration of 3-D shapes, IEEE Transactions on Pattern Analysis and Machine Intelligence 14 (1992) 239–256. URL: https://doi.org/10.1109/34.121791. doi:10.1109/34.121791.