Uncertainty-Driven 3D CAD-Guided Point Cloud Reconstruction for Hazardous Environment Documentation

Anqi Hu[1], Tomohiro Fukuda[2]

[1,2] Division of Sustainable Energy and Environmental Engineering, Graduate School of Engineering, Osaka University
[1]hu@it.see.eng.osaka-u.ac.jp, 0000-0002-8568-1030
[2]fukuda.tomohiro.see.eng@osaka-u.ac.jp, 0000-0002-4271-4445

This research addresses point cloud reconstruction for architectural documentation in hazardous environments like earthquake-damaged buildings. While AI methods show promise, they lack architectural domain information integration. The proposed approach combines 3D CAD models with deep learning, using model-derived point clouds as ground truth and simulating realistic damage conditions. The methodology employs AdaPoinTr for initial completion, followed by uncertainty-driven CAD-guided fusion that identifies high-uncertainty regions and selectively applies geometric information. Experimental results across four residential case studies show 42.9% improvement in absolute chamfer distance and 66.5% in volume-relative metrics compared to deep learning-only methods. This targeted integration of architectural information with AI enhances reconstruction quality for complex features in uncertain conditions, making it valuable for emergency response, renovation planning, and building assessment in unsafe structures.

Keywords: *Point cloud reconstruction, Computer-aided design, Targeted point cloud repair, Hazardous environment documentation, Deep learning completion.*

INTRODUCTION

Architectural documentation in hazardous environments presents significant challenges for traditional surveying methods. Sites affected by earthquakes, structural instability, chemical contamination, or radiation pose substantial risks to human surveyors, resulting in incomplete documentation that hampers damage assessment, renovation planning, and reconstruction efforts.

LiDAR scanning enables rapid data collection while minimizing human exposure. However, safety concerns restrict scanner placement, creating occluded views that produce incomplete point clouds with significant missing regions.

Deep learning approaches such as AdaPoinTr (Yu et al., 2023) show promise in generating plausible completions but struggle with architectural complexity. Traditional geometry-based methods offer precision but face limitations with extensive missing regions (Berger et al., 2017). Recent benchmarking by Gao et al. (2024) shows methods performing well on object datasets often struggle with building-scale point clouds. 3D CAD models offer valuable architectural information including geometric details and structural relationships. However, in post-disaster scenarios, direct use of CAD models alone is insufficient as they represent the original design state rather than the current damaged condition.

Despite advances in both approaches, a significant research gap exists in methods that can selectively leverage architectural information. Current CAD integration techniques typically apply information uniformly across reconstructions (Avetisyan et al., 2019), failing to distinguish between regions requiring different treatment based on reconstruction uncertainty.

This research proposes an uncertainty-driven approach that selectively integrates CAD model information with deep learning-based point cloud reconstruction. The method begins with AdaPoinTr's initial completion, followed by statistical analysis to identify damage regions where reconstruction uncertainty is high. These regions are then selectively repaired using geometric information from CAD models. The research contributions include:

1. An uncertainty-driven framework identifying high-uncertainty regions in deep learning point cloud reconstructions using statistical distance analysis and DBSCAN clustering
2. A targeted CAD-guided repair approach that selectively integrates CAD model geometry into identified damage regions while preserving well-reconstructed areas
3. A comprehensive evaluation framework assessing reconstruction quality through multiple metrics, including absolute and normalized chamfer distances (CD)
4. Empirical validation across four residential building case studies, demonstrating significant improvements in reconstruction accuracy compared to deep learning-only approaches

This approach advances architectural documentation in hazardous environments, enabling more accurate reconstruction with limited scan data by leveraging existing CAD models to enhance critical areas.

RELATED WORK

This review examines four key areas: point cloud reconstruction, deep learning-based completion, CAD integration, and architectural documentation in hazardous environments.

Point Cloud Reconstruction

Point cloud reconstruction has evolved from geometric approaches to data-driven methods. Conventional techniques like Poisson surface reconstruction (Guennebaud and Gross, 2007) excel with simple geometries but struggle with complex architectural structures. Berger et al. (2017) highlighted limitations with incomplete data. In architectural contexts, specialized approaches have emerged for indoor reconstruction.

Deep Learning for Point Cloud Completion

Deep learning has transformed point cloud completion by enabling models to learn complex geometric patterns. Pioneering work by Qi et al. (2017) with PointNet established direct learning on unordered point sets. PCN (Yuan et al., 2019) introduced an encoder-decoder architecture for completion tasks, while PoinTr (Yu et al., 2021) treated completion as a set-to-set translation problem.

AdaPoinTr extends transformer-based approaches with adaptive feature interpolation that better preserves local details. Despite these advances, deep learning faces limitations with architectural documentation as models typically train on object-centric datasets that differ from architectural structures. Gao et al. (2024) demonstrated that methods performing well on object datasets often struggle with building-scale point clouds.

CAD Integration in Reconstruction

CAD models offer valuable structured geometric information for enhancing point cloud reconstruction. Early approaches focused on template matching, where CAD models serve as references. Avetisyan et al. (2019) proposed a framework that simultaneously aligns CAD models to RGB-D scans while reconstructing 3D scenes. However, existing approaches typically apply CAD information uniformly across reconstructions, without differentiating between regions that benefit from deep learning predictions and those requiring CAD guidance.

Architectural Documentation in Hazardous Settings

Documentation in hazardous environments requires remote sensing and robotic systems. Remote techniques minimize human exposure (Lindenbergh and Pietrzyk, 2015) but provide limited interior data. Autonomous scanning robots navigate dangerous areas (Nex and Remondino, 2014) but face coverage constraints.

Combining deep learning with architectural information shows promise. Matrone et al. (2020) demonstrated improved reconstruction accuracy through semantic segmentation. However, methods designed for hazardous environments with unique damage patterns remain limited.

This research addresses this gap by introducing a targeted approach that selectively applies CAD model information to regions where deep learning predictions exhibit high uncertainty, combining strengths of both paradigms while accommodating the specific challenges of hazardous environment documentation.

METHODOLOGY

The uncertainty-driven CAD-guided point cloud reconstruction framework selectively combines deep neural networks' adaptability with CAD models' precision. The method identifies regions with high reconstruction uncertainty in AdaPoinTr predictions and applies targeted CAD-guided repair only to these areas. As shown in Figure 1, the framework consists of four main components: (1) data preparation and hazardous scanning condition simulation, (2) initial point cloud reconstruction using AdaPoinTr, (3) uncertainty estimation through statistical analysis, and (4) targeted CAD-guided fusion of identified regions.

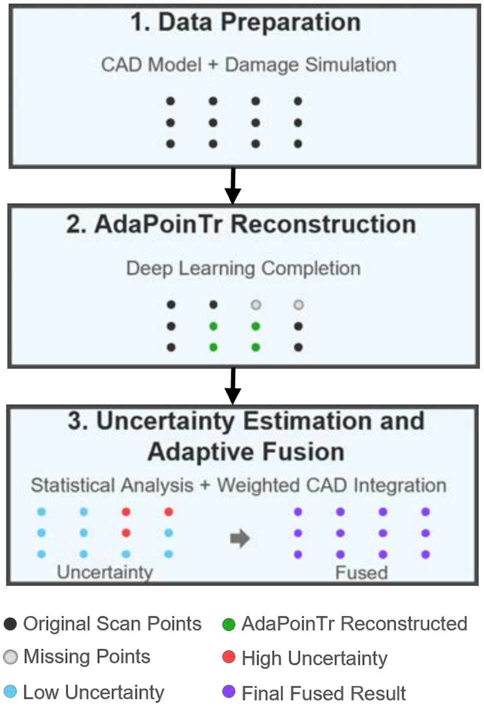

Figure 1
3D CAD-guided point cloud damage detection and repair workflow.

Data Preparation and Simulation

Ground truth point clouds (344,064 points) were generated from high-fidelity 3D CAD models using Open3D's triangle mesh sampling. A physically informed damage simulation algorithm mimics structural damage patterns by selectively removing 40-60% of points from concentrated regions to simulate realistic hazardous documentation conditions. The CAD reference point cloud is generated by uniformly sampling the same number of points as the AdaPoinTr reconstruction to ensure consistent comparison.

Initial Reconstruction using AdaPoinTr

The AdaPoinTr model generates an initial reconstruction through point cloud encoding into feature representation, proxy points generation serving as queries for missing regions, geometry-aware transformer-based completion, and adaptive feature interpolation enhancing detail preservation. The model was trained on ShapeNet objects and Building-PCC data to improve architectural structure performance. While generating plausible completions, AdaPoinTr often struggles with complex architectural features and large missing regions typical in hazardous environment scans.

Uncertainty Estimation and Adaptive Fusion

The CAD-derived point cloud and AdaPoinTr reconstruction are aligned through a two-stage registration process involving random sample consensus (RANSAC)-based global registration using fast point feature histograms (FPFH) features followed by iterative closest point refinement for precise alignment. Once aligned, point-wise geometric deviations are computed between the AdaPoinTr reconstruction and CAD reference using k-dimensional tree search to find nearest neighbors. These distances directly indicate reconstruction quality where larger distances suggest areas where AdaPoinTr struggled with accurate geometric recovery. To identify regions requiring CAD guidance, the resulting distance distribution is analyzed statistically with points exceeding a threshold defined as Equation (1):

$$T = M + \alpha \cdot \sigma \qquad (1)$$

where T represents the threshold, M is the median distance, σ is the standard deviation, and α is a distance threshold factor. This statistical approach adapts to different buildings and damage patterns using relative rather than absolute thresholds. Points exceeding this threshold are flagged as high-uncertainty regions requiring targeted repair.

The fusion process operates through adaptive weight calculation where high-uncertainty regions receive greater CAD influence while low-uncertainty regions preserve AdaPoinTr predictions. Identified high-uncertainty regions are expanded using a radius-based approach as shown in Equation (2):

$$R = \bar{d} \cdot \beta \qquad (2)$$

where R is the repair radius, $\bar{d}$ represents the average distance between neighboring points in the uncertain region, and β is a repair radius factor to ensure smooth transitions between repaired and preserved regions. Within these expanded regions, weighted fusion combines points from the CAD-derived cloud and AdaPoinTr reconstruction as shown in Equation (3):

$$P\_f = w\_C \cdot P\_C + w\_A \cdot P\_A \qquad (3)$$

where *P_f* represents the fused point, *P_C* is the point from CAD-derived cloud, *P_A* is the point from AdaPoinTr reconstruction, and weights *w_C* and *w_A* are determined by local uncertainty values. High-uncertainty regions typically receive CAD weights of 0.8-0.9 to ensure architectural accuracy, while transition zones use moderate weighting around 0.6 for CAD and 0.4 for AdaPoinTr to maintain smooth boundaries. Areas outside expanded regions preserve the original AdaPoinTr reconstruction unchanged.

The fused point cloud undergoes comprehensive post-processing including statistical outlier removal to eliminate fusion artifacts, surface normal estimation with consistency orientation, and optional Poisson surface reconstruction for smooth surface generation.

EXPERIMENTAL RESULTS

The experiments were conducted using four residential building models with varied architectural characteristics. Each building model was processed through the complete pipeline: (1) ground truth generation from CAD models, (2) hazardous environment simulation, (3) AdaPoinTr reconstruction, (4) damage region detection, and (5) targeted CAD-guided repair.

The dataset comprises four residential buildings (Buildings 0, 2, 5, and 75) with varying architectural designs, ranging from Building 0's compact L-shaped layout to Building 5's complex multi-room structure, each represented as a colored point cloud with 344,064 points identifying different structural elements. The hazardous environment simulation physically removed 40-60% of points, and the AdaPoinTr model was implemented with a distance threshold factor (α) of 2.5 and repair radius factor (β) of 1.5, determined through ablation studies, with all processing performed on a workstation equipped with an Intel Core i5-11400 CPU, 16GB RAM, and an NVIDIA GeForce RTX 3060 Ti.

Quantitative Results

The performance of the proposed uncertainty-driven CAD-guided reconstruction approach was evaluated using two primary metrics to quantify point cloud reconstruction quality:

Absolute CD: CD is a standard metric for evaluating the similarity between two-point cloud sets, calculated as Equation (4):

$$\begin{aligned} CD(X,Y) &= CD_1 + CD_2 \\ CD_1 &= \frac{1}{|X|}\sum_{x\in X} \min_{y\in Y}\|x-y\|^2 \\ CD_2 &= \frac{1}{|Y|}\sum_{x\in X} \min_{x\in X}\|y-x\|^2 \end{aligned} \qquad (4)$$

where and Y represent the predicted and reference point clouds, $||x\text{-}y||$ denotes the Euclidean distance, and lower CD values indicate higher reconstruction quality.

Volume-relative CD (ppm): To enable fair comparisons between buildings of different sizes, the volume-relative CD is introduced in Equation (5), measured in parts per million (ppm):

$$Vol - Rel\ CD = \frac{Absolute\ CD}{D} 10^6 \qquad (5)$$

where D is the characteristic dimension of the model (typically the diagonal length of the model's bounding box), and this metric eliminates the influence of a building's absolute size. The performance of the proposed uncertainty-driven CAD-guided reconstruction approach was evaluated against the AdaPoinTr-only reconstruction using these metrics, with results summarized in Table 1.

Table 1
Quantitative comparison of key metrics

Build-ing	Absolute CD		Volume-Relative (ppm)	
	Ada-PoinTr	CAD-guide	Ada-PoinTr	CAD-guide
0	0.0279	0.0162	421.6	143.7
2	0.0243	0.0158	386.4	154.8
5	0.0318	0.0172	512.9	168.5
75	0.0288	0.0153	467.2	132.6
Ave-rage	0.0282	0.0161	447.0	149.9
Impro-vement	42.90%		66.5%	

The results demonstrate that the proposed CAD-guided repair approach significantly outperforms the AdaPoinTr-only method across all metrics and all building models.

The proposed method achieves consistent performance with an average improvement of 42.9% in absolute CD, with improvements observed for all building models.

Regarding volumetric accuracy, the volume-relative metric shows the most dramatic improvement (66.5% on average), indicating that the method particularly excels at preserving the volumetric integrity of buildings.

Building-specific analysis reveals that Building 5, with its complex multi-room design, shows the largest absolute improvement in CD (from 0.0318 to 0.0172), while Building 75 with its more regular structure shows the best volume-relative performance (reduction from 467.2 to 132.6 ppm).

Overall effectiveness is demonstrated by the consistent improvements across different building types, showing the robustness of the approach for architectural point cloud repair in hazardous environments. These results confirm that integrating architectural domain information (through CAD models) with deep learning techniques significantly enhances point cloud reconstruction quality for architectural documentation in hazardous environments.

Qualitative Analysis

Visual comparison between the ground truth, incomplete point cloud, AdaPoinTr reconstruction, and CAD-guided repair results reveals significant quality differences. The qualitative analysis reveals several key observations, as illustrated in Figure 2.

1. Structural Integrity: The AdaPoinTr reconstruction (third column) often produces smoothed approximations of damaged regions, visible in the irregular color transitions in the heat-map visualization. The CAD-guided method (fourth column) shows more distinct color boundaries, accurately preserving structural elements like sharp edges and corners with clearer depth stratification.
2. Feature Preservation: Window and door openings, which are critical architectural features, appear as consistent blue regions (indicating depth) in the CAD-guided approach. AdaPoinTr occasionally fills these openings or distorts their shape, shown by inconsistent coloration across openings.
3. Wall Alignment: The CAD-guided method maintains proper wall alignment and perpendicularity, visible in the structured color gradients along vertical surfaces. AdaPoinTr sometimes produces walls with slight angular deviations, indicated by uneven color transitions along what should be straight edges.
4. Damage Region Identification: The incomplete point cloud (second column) clearly shows the simulated damage patterns, missing substantial structural elements compared to the ground truth, particularly interior walls, partial exterior walls, and floor areas. The heat-map visualization in the CAD-guided results reveals how these areas are effectively targeted and repaired with appropriate depth information.
5. Edge Cases: In severe damage scenarios, such as the large missing areas in Buildings 2 and 75, AdaPoinTr generates plausible but incorrect completions with inconsistent heat-map patterns. The CAD-guided method, guided by CAD information, recovers the correct structural form and spatial relationships even in these challenging cases, shown by more regular and architecturally consistent color transitions.

By comparing four building models of different complexities, the consistency of the method in handling different types of structural damage can be observed. Especially in complex structures like Building 5, CAD-guided repair can accurately reconstruct interior partition walls and openings with clear depth differentiation,

Figure 2
Comparison of ground truth, incomplete point cloud, AdaPoinTr reconstruction, and CAD-guided repair results for four building models, visualized using a height-based heat map (blue to red indicating increasing elevation from ground level)

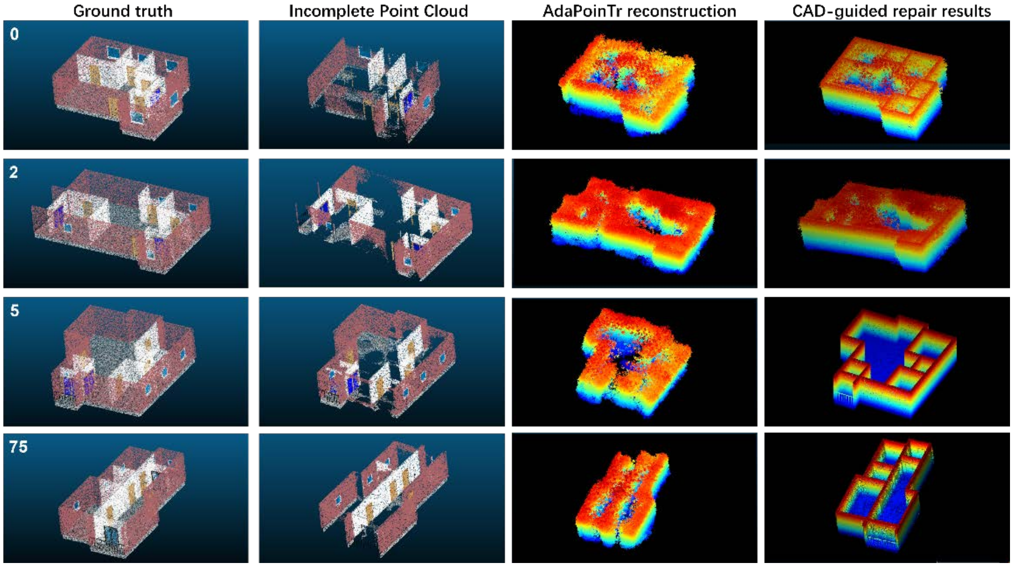

whereas AdaPoinTr produces blurred results with less defined depth boundaries in these areas. To illustrate the effectiveness of the uncertainty-driven approach, Figure 3 provides a zoomed-in comparison of a complex wall section in Building 5. As shown in Figure 3, the ground truth (a) displays clear geometric structure and properly aligned surfaces. The AdaPoinTr reconstruction (b) exhibits geometric distortion and irregular surface quality, particularly at structural junctions. After applying the statistical thresholding method (Equation 1) to identify high-uncertainty regions, the targeted CAD-guided fusion (c) successfully restores structural accuracy while preserving detail in well-reconstructed areas. This focused comparison helps explain why Building 5 shows the largest absolute improvement in CD (from

Figure 3
Structural detail comparison showing CAD-guided fusion improvement for complex architectural features

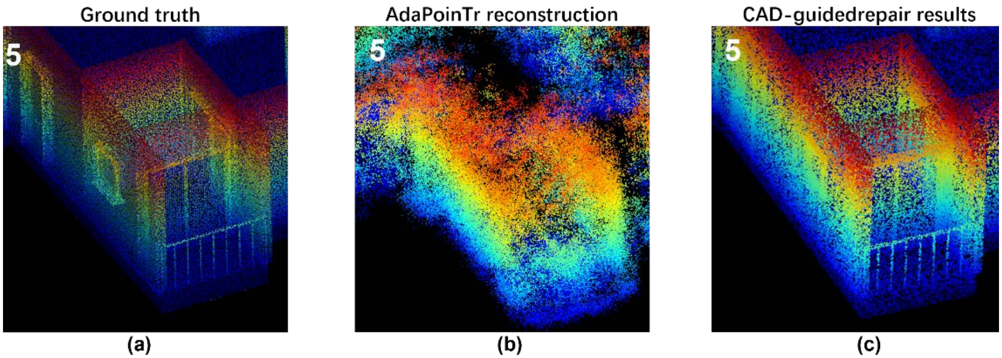

0.0318 to 0.0172 as shown in Table 1), demonstrating the effectiveness of selective CAD integration for complex architectural structures.

This uncertainty-driven approach automatically identifies problematic reconstruction regions and applies targeted CAD guidance without manual threshold tuning, making it adaptable to various damage patterns and building types while ensuring accurate reconstruction of genuinely damaged building areas suitable for safety-critical applications.

Ablation Studies

To evaluate parameter contributions, ablation studies were conducted on distance threshold factor α and repair radius factor β.

The distance threshold α controls damage detection sensitivity. Testing values from 1.5 to 4.0 (Figure 4) showed that lower thresholds (α = 1.5) over-identify damage regions (~9.55% of points), causing excessive CAD influence.

Higher thresholds (α = 4.0) are conservative, identifying only ~1.48% as damaged but finding coherent regions via DBSCAN clustering. The optimal α = 2.5 balances accurate damage identification while preserving successful AdaPoinTr reconstructions.

The repair radius β determines repair extent beyond damage regions (Figure 5).

Smaller radii (β = 1.0) create abrupt transitions with visible boundaries, while larger radii (β ≥ 2.0) risk over-modifying well-performing areas. The optimal β = 1.5 ensures smooth transitions while respecting appropriate AdaPoinTr reconstructions.

Results demonstrate that repair radius significantly impacts final CD and Hausdorff distances, with appropriate expansion ensuring seamless fusion between repair regions and surrounding point cloud.

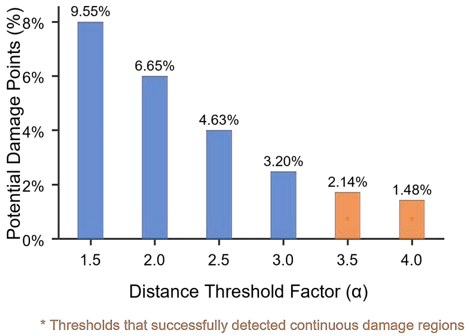

Figure 4
Effect of distance threshold factor (α) on damage detection

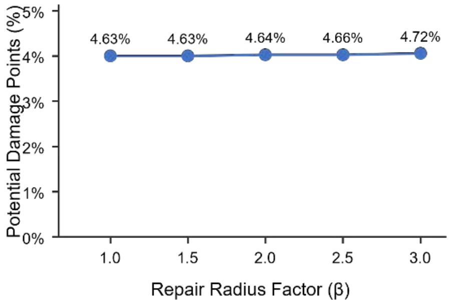

Figure 5
Effect of repair radius factor (β) on reconstruction quality

DISCUSSION

This section discusses the broader implications of the experimental findings, considering both the technical significance and potential applications of uncertainty-driven CAD-guided point cloud reconstruction in hazardous environments. The results not only demonstrate quantitative improvements but also reveal insights into the integration of domain information with deep learning approaches for architectural documentation.

Key Findings and Implications

The experimental results demonstrate that the uncertainty-driven CAD-guided point cloud reconstruction approach significantly outperforms pure deep learning methods for architectural reconstruction in hazardous environments. Several key findings emerge with important implications:

First, the substantial performance gap between the targeted approach and global CAD fusion confirms that selective integration is more effective than uniform application of CAD information. This challenges traditional paradigms and suggests computational resources should focus on high-uncertainty regions.

Second, the consistent improvement across buildings with varied scales and architectural styles indicates the approach addresses fundamental limitations in deep learning-based reconstruction rather than exploiting specific geometries, suggesting broader applicability beyond tested residential buildings.

Third, the statistical distance analysis combined with DBSCAN clustering effectively localizes high-uncertainty regions in deep learning predictions, suggesting reconstruction uncertainty can be reliably detected through geometric analysis without requiring model uncertainty estimates, with potential applications in active scanning guidance.

Limitations and Future Work

Despite promising results, several limitations remain. The current approach assumes the availability of CAD models that reasonably match the documented buildings. While common for modern structures, historical buildings or those with significant undocumented modifications may lack accurate CAD representations. Future work could explore adapting generic architectural templates to specific buildings based on available scan data.

The current implementation processes buildings sequentially, with acceptable processing time for offline documentation (3-5 minutes per building). Real-time applications would require optimization through GPU acceleration and parallel processing of building sections.

Additionally, the approach focuses on geometric repair without explicitly considering semantic information. Future work could integrate semantic segmentation to apply different repair strategies based on element types—windows might benefit from template-based repair while walls could use planar constraints.

Broader Impact

The uncertainty-driven CAD-guided point cloud reconstruction approach has potential impacts beyond technical improvements. By reducing required scan coverage for accurate documentation, this approach could significantly decrease human
exposure time in hazardous environments. This reduction in required scan data could translate to proportional reductions in on-site time, with important safety benefits in scenarios like post-earthquake documentation or contaminated site assessment.

The ability to generate accurate reconstructions from limited scans makes architectural documentation more accessible in challenging contexts like disaster-affected remote areas, conflict zones with limited access, or economically constrained regions where extensive scanning equipment may be unavailable.

CONCLUSION

This paper presented an uncertainty-driven CAD-guided point cloud reconstruction approach for architectural documentation in hazardous environments. By selectively integrating CAD information based on uncertainty localization, the method significantly outperforms pure deep learning approaches, with improvements exceeding 40% in CD metrics across four residential buildings. The method focuses computational resources on high-uncertainty regions, enabling more accurate documentation with limited scan data while preserving the actual current condition captured in scans.

This selective integration strategy is particularly valuable for environments with high radiation levels, such as nuclear facilities, where minimizing human exposure time is critical. Rather than restoring structures to their original CAD model state, our approach enhances only areas with insufficient data, making it suitable for emergency response, renovation planning, and structural assessment in unsafe facilities. Future work will address historical structure adaptation, computational efficiency, and validation with real-world damaged building scans from various hazardous environments.

ACKNOWLEDGMENTS

This work was supported by JAEA Nuclear Energy S&T and Human Resource Development Project Grant Number JPJA24P24020673. The authors would like to thank the following members for their support of this study: Mugita, Y., Nishimoto, R., Ogura, A., Shimizu, H., Shirahase, T., and Tsurunaga, S.

REFERENCES

Avetisyan, A., Dahnert, M., Dai, A., Savva, M., Chang, A.X. and NieBner, M. (2019) 'Scan2CAD: Learning CAD Model Alignment in RGB-D Scans', in: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Presented at the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Long Beach, CA, USA, pp. 2609–2618.

Berger, M., Tagliasacchi, A., Seversky, L.M., Alliez, P., Guennebaud, G., Levine, J.A., Sharf, A. and Silva, C.T. (2017) 'A Survey of Surface Reconstruction from Point Clouds', *Computer Graphics Forum 36*, pp. 301–329.

Gao, W., Peters, R. and Stoter, J. (2024) 'Building-PCC: Building Point Cloud Completion Benchmarks', *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* X-4-W5-2024, pp. 179–186.

Guennebaud, G. and Gross, M. (2007) 'Algebraic point set surfaces', *ACM Trans. Graph.* 26, 23. https://doi.org/10.1145/1276377.1276406

Lindenbergh, R. and Pietrzyk, P. (2015). 'Change detection and deformation analysis using static and mobile laser scanning', *Appl Geomat* 7, pp. 65–74.

Matrone, F., Lingua, A., Pierdicca, R., Malinverni, E.S., Paolanti, M., Grilli, E., Remondino, F., Murtiyoso, A. and Landes, T. (2020) 'A benchmark for large-scale heritage point cloud semantic segmentation', *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci. XLIII-B2-2020*, 1419–1426. https://doi.org/10.5194/isprs-archives-XLIII-B2-2020-1419-2020

Nex, F. and Remondino, F. (2014) 'UAV for 3D mapping applications: a review', *Appl Geomat 6*, 1–15. https://doi.org/10.1007/s12518-013-0120-x

Qi, C.R., Su, H., Mo, K. and Guibas, L.J. (2017) 'PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation'. https://doi.org/10.48550/arXiv.1612.00593

Yu, X., Rao, Y., Wang, Z., Liu, Z., Lu, J. and Zhou, J. (2021) 'PoinTr: Diverse Point Cloud Completion with Geometry-Aware Transformers', in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. Presented at the 2021 IEEE/CVF International Conference on Computer Vision (ICCV), IEEE, Montreal, QC, Canada, pp. 12478–12487. https://doi.org/10.1109/ICCV48922.2021.01227

Yu, X., Rao, Y., Wang, Z., Lu, J. and Zhou, J. (2023) 'AdaPoinTr: Diverse Point Cloud Completion With Adaptive Geometry-Aware Transformers', *IEEE Transactions on Pattern Analysis and Machine Intelligence 45*, 14114–14130.

Yuan, W., Khot, T., Held, D., Mertz, C. and Hebert, M. (2019) *PCN: Point Completion Network*.