

BIMStruct3D: A Fully Automated Hybrid Learning Scan-to-BIM Pipeline with Integrated Topology Refinement

Mahdi Chamseddine^{1,2}, Fabian Kaufmann², Marius Schellen², Christian Glock²,
Didier Stricker^{1,2}, and Jason Rambach¹

¹German Research Center for Artificial Intelligence (DFKI), Kaiserslautern, Germany

²RPTU Kaiserslautern-Landau, Kaiserslautern, Germany

Accepted for presentation. This pre-print includes supplementary material that is not part of the official conference proceedings.

Abstract

Automatic generation of Building Information Models (BIM) from building scans is a key challenge in architecture and construction. We present a modular pipeline for generating IFC-compliant BIM from 3D point clouds. The hybrid approach combines learning-based semantic segmentation with topology-aware geometric reconstruction to model structural elements accurately. We propose vIoU, adapting voxel-based overlap evaluation to Scan-to-BIM by enabling holistic, instance-matching-free comparison of reconstructed and ground-truth models. We release the German Hospital dataset (DeKH), including high-resolution point clouds, ground truth BIMs, and semantic annotations. Experiments on DeKH and CV4AEC datasets show significant improvements over a RANSAC-based baseline, demonstrating robustness and scalability.

Introduction

Capturing spatial data of existing built environments can be efficiently achieved using mobile (Kanayama et al., 2025) and terrestrial scanning systems (Chang et al., 2017; Armeni et al., 2017). To leverage such data in construction (Wu et al., 2021) and facility management (Xu et al., 2021), it must be converted into semantic models in accordance with the Building Information Modelling (BIM) methodology. Given the volume of data and the complex steps required to transform raw 3D point clouds into BIM models, full automation of this process is highly desirable. However, this transformation remains a challenge.

Despite the importance of the task, scan-to-BIM is rarely treated as a comprehensive, end-to-end problem with standardized benchmarks; however, some effort is being made through initiatives like the CV4AEC challenges (Armeni et al., 2024). In this work, we introduce BIMStruct3D, one of the first comprehensive frameworks for automating the conversion of noisy and incomplete 3D point clouds into IFC-compatible BIM models. Our BIMStruct3D pipeline employs a hybrid approach that integrates both learning-based and classical point cloud processing techniques. It includes modules for 3D point cloud segmentation, geometric reconstruction and primitive extraction, as well as model post-processing, as illustrated in Figure 1.

BIMStruct3D is designed to process data from multi-storey buildings such as offices, hospitals, and residential complexes, and generate BIM representations of architectural




Figure 1: BIMStruct3D is a hybrid pipeline for generating IFC BIM models from 3D point cloud scans

elements including walls, doors, and columns. Accompanying our framework is `pystruct3d`, a modular library tailored for processing 3D point cloud data of buildings.

To support reproducibility and facilitate benchmarking in the scan-to-BIM domain, we provide a comprehensive evaluation of our method at both modular and end-to-end levels. This includes tests on publicly available datasets and newly acquired data, featuring high-precision, handcrafted ground truth BIMs. We also propose volumetric Intersection over Union (vIoU), which adapts voxel-based IoU evaluation to the Scan-to-BIM domain. Unlike per-object voxel IoU used in 3D shape reconstruction, vIoU operates at the class level without instance matching, making it robust to the fragmented and variably segmented wall representations common in BIM reconstruction.

In real-world scenarios, building point clouds are often incomplete due to factors such as the limitations of scanning technology, human error and environmental conditions like lighting, occlusion, and material reflectivity (e.g. glass surfaces). BIMStruct3D is built to handle such noisy and incomplete data to create usable BIM representations.

In summary, our contributions are as follows:

- BIMStruct3D, a fully automated hybrid pipeline generating IFC BIM models from 3D point clouds using semantic segmentation and geometric reconstruction.
- Specialized algorithms for primitive extraction and BIM model generation.
- vIoU, an adaptation of voxel-based IoU to Scan-to-BIM evaluation, enabling class-level comparison without instance-level matching.
- `pystruct3d`¹, an open-source library for processing 3D building point clouds, released with DeKH², a new dataset including high-resolution scans, semantic annotations, and handcrafted ground truth BIMs.

¹<https://github.com/humantecheu/pystruct3d>

²<https://huggingface.co/datasets/RPTU-FGMB/DeKH>

Related Work

Point Cloud Semantic Segmentation

Semantic segmentation of point clouds is a fundamental computer vision task, currently dominated by deep neural networks. Early approaches include projection-based (Chen et al., 2017), voxel-based (Graham et al., 2018), and point-based networks (Qi et al., 2017), each trading off between computational cost, geometric fidelity, and flexibility. More recent developments have incorporated self-attention mechanisms, demonstrating strong performance in large-scale 3D scene understanding (Wang et al., 2019). Point Transformer by Zhao et al. (2021) advanced point-based segmentation by integrating local self-attention (Vaswani et al., 2017), vector attention (Zhao et al., 2020), and appropriate positional encodings. Subsequent iterations (Wu et al., 2024) introduce architectural improvements leading to better segmentation accuracy and model efficiency.

Scan-to-BIM Approaches

To align with the scope of this work, we focus on research targeting fully automated (end-to-end) scan-to-BIM conversion from point clouds. Gourguechon et al. (2022) and Bassier and Vergauwen (2020) distinguish between room-based and wall-based (or structural component-based) reconstruction approaches. The former identifies rooms and then reconstructs surrounding elements, while the latter directly targets structural components such as walls. Unfortunately, to this day, there are no open-source implementations of complete scan-to-BIM pipelines or benchmark datasets for evaluation.

Room-based reconstruction

Room-based methods segment indoor spaces and reconstruct surrounding elements. Representative approaches address plane fitting and wall estimation (Ochmann et al., 2016, 2019), 2D region growing with slab reconstruction (Macher et al., 2017), comprehensive element reconstruction including windows and fixtures (Xiong et al., 2023), hybrid IFC wall reconstruction with topology refinement (Bassier et al., 2018), and learning-based segmentation followed by space decomposition (Tang et al., 2022; Hu et al., 2020). While these methods offer good indoor segmentation, Bassier and Vergauwen (2020) note their limitations in modelling non-room structures and handling clutter. Thus, component-based approaches targeting structural elements have gained traction.

Structural component reconstruction

Component-based reconstruction identifies and models structural elements like walls and columns through clustering, semantic segmentation, and primitive fitting. Bassier and Vergauwen (2020) focused on extracting standard wall objects with topological rules for merging, while follow-up work (Bassier et al., 2020) compared 2D and 3D reconstruction methods. Other approaches combine 2D and 3D processing with point cloud upsampling (Gankhuyag and

Han, 2021), use RANSAC and convex hull-based boundary extraction (Thomson and Boehm, 2015), reconstruct walls and spaces from enclosed rooms (Anagnostopoulos et al., 2016), or segment concrete components via concavity/convexity criteria (Son and Kim, 2017). The closest work to ours is by Kim and Kim (2021), which uses synthetic data to train a segmentation model, followed by instance clustering and planar patch segmentation. Element relationships are modelled with graph networks to infer missing links and extend incomplete elements.

Our approach differs in that we do not rely on synthetic data and instead address several common limitations, including limited generalization and reliance on manual parameter tuning. We integrate semantic segmentation, geometric reconstruction, and topological reasoning into a unified scan-to-BIM pipeline.

Methodology

Our method, outlined in Figure 2, consists of semantic and instance segmentation, geometry reconstruction, topology refinement, and BIM object creation from a 3D point cloud. IFC export is handled via *IfcOpenShell*³, generating a project–site–building–storey hierarchy and assigning reconstructed geometry as corresponding IFC entities (*IfcWall*, *IfcDoor*, *IfcColumn*). During the reconstruction process, we adopt the Manhattan world assumption (Coughlan and Yuille, 1999), which holds for the majority of structural objects in typical building environments.

Initially, levels/storeys are identified by analysing vertical point density distributions, using surfaces like floors and ceilings to detect significant point concentrations (Bassier et al., 2020; Macher et al., 2017). Semantic segmentation using state-of-the-art deep learning models (Wu et al., 2024) serves as the basis for reconstructing walls, doors, and columns. The reconstruction pipeline follows a common structure: instance segmentation, geometry reconstruction, topology refinement, and IFC BIM object creation.

Wall reconstruction is based on axis-aligned point grouping. DBSCAN clustering (Ester et al., 1996) is applied per axis to form spatial clusters, discarding sparse noise as seen in Figure 3. HYSAC, our RANSAC (Schnabel et al., 2007) based plane segmentation algorithm with optimized seed selection, is used to extract wall surfaces. These are enclosed using Horizontal Oriented Bounding Boxes (H-OBB), followed by topological refinements to ensure geometric correctness.

Door points are initially identified via semantic segmentation and then treated as child objects of walls, relying on the reconstructed wall geometry as a starting point. The process first identifies door points within the wall bounding boxes, which typically correspond to frames, linings, or closed door leaves captured during scanning. If such points are found, the bounding box is expanded to include additional door geometry, such as open door leaves, ensuring doors are correctly associated with their respective parent

³<https://ifcopenshell.org/>

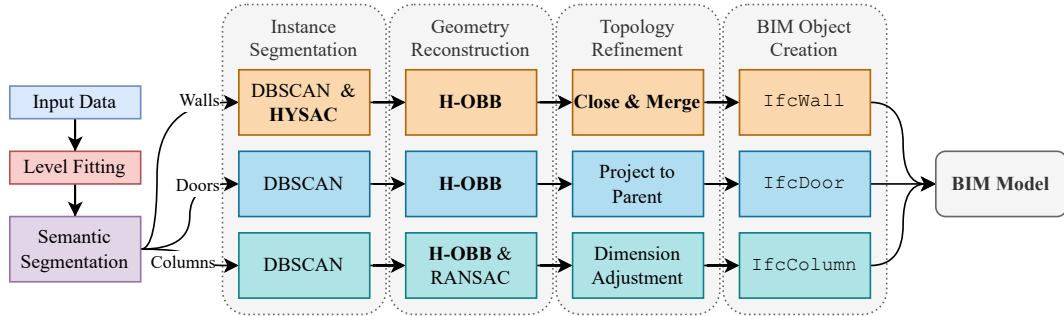


Figure 2: Overview of the reconstruction pipeline showing the different reconstruction stages for walls, doors, and columns. In **bold** are our main contributed components.




Figure 3: (Left) Wall direction filtering shown in red and green and (Right) DBSCAN clustering for one of the directions in multiple colours.

walls. Our two-step approach reduces errors caused by open doors adjacent to perpendicular walls.

Column points are clustered using DBSCAN (Ester et al., 1996) and filtered by minimum point thresholds to reduce artifacts. Shape classification is achieved through curvature analysis based on eigenvalue decomposition from k-d tree neighbourhood searches. Curvature distributions are modelled with Maximum Likelihood Estimation to distinguish round (low standard deviation) from square or rectangular columns (high standard deviation). We then use RANSAC cylinder fitting for round columns and H-OBB fitting square/rectangular columns.

Semantic Segmentation

Semantic segmentation is essential for producing accurate IFC models, as it forms the foundation for all subsequent steps. High-quality segmentation ensures later geometry reconstruction stages are based on reliable object labels.

We use Point Transformer V3 by Wu et al. (2024), trained on a combined dataset from three public datasets: S3DIS (Armeni et al., 2016), Structured3D (Zheng et al., 2020), and ScanNet (Dai et al., 2017). Training on a joint dataset prevents the model from over-fitting to the data and sensor used in one dataset. To further improve generalization, additional data augmentation techniques were used, such as removing color information to simulate scans from grayscale 3D sensors. This is particularly relevant because the CV4AEC dataset used for the evaluation contains colored and grayscale point clouds. This segmentation

pipeline ensures robust performance across diverse scan types and scene complexities, enabling reliable downstream reconstruction of BIM elements.

Topology aware geometry reconstruction

The reconstruction pipeline leverages both state-of-the-art techniques and our novel contributions. The following sections describe our key innovations in wall reconstruction, oriented bounding box fitting, and topology correction.

HYSAC wall reconstruction

The RANSAC-based plane fitting methods (Schnabel et al., 2007) are highly sensitive to parameter settings like the minimum number of inliers and distance thresholds. Improper parameter tuning often results in under-segmentation, where a single plane, instead of two, is incorrectly fitted to an entire wall instance.

To address this, we propose Hypothesis-based Sample Consensus (HYSAC), a modified RANSAC approach that improves seed point selection. Instead of sampling seed points randomly, HYSAC benefits from the point density distribution perpendicular to the primary wall orientation. Inspired by Armeni et al. (2016), we compute a 1D density histogram along the axis orthogonal to the wall surface and select seed points from regions of highest density. Planes are fitted using Singular Value Decomposition (SVD), which provides a least-squares optimal plane fit robust to point cloud noise, and accepted based on a minimum inlier ratio. Algorithm 1 shows a pseudo code implementation of HYSAC.

As shown in Figure 4, HYSAC reliably detects multiple planes within a wall cluster, distinguishing between valid wall surfaces and outliers. The resulting inlier points are grouped and forwarded for further geometric processing.




Figure 4: HYSAC plane fitting. Result of plane fitting (plane points coloured, outliers grey) with corresponding histogram of point distribution.

Algorithm 1: HYSAC plane fitting pseudocode.

```
input :wall_points # points of a wall cluster
input :n_bins      # number of histogram bins
input :n_seeds     # number of seed points
input :min_points  # minimum points per plane
output:planes     # equations of found planes
output:inliers    # points belonging to the planes

1 def histogram_seeds(wall_points, n_seeds):
    # generate the histogram perpendicular to the
    # wall direction
2     hist = histogram(cluster_points, n_bins)
3     peak = find_peaks(hist)
    # select random points from peaks
4     seed_points = random_sample(peak, n_seeds)
5     return seed_points
6 end

7 def hysac_plane(wall_points, n_seeds):
8     points = wall_points
    # loop until all planes are found
9     while number_of(points) ≥ min_points:
10        seeds = histogram_seeds(points, n_seeds)
11        plane, inliers = fit_plane_svd(seeds)
12        if number_of(inliers) ≥ min_points:
            # remove inliers of the plane from the
            # wall points
13            points = remove(points, inliers)
14            yield plane, inliers
15        end
16    end
17 end

18 planes, inliers = hysac_plane(wall_points, n_seeds)
19 return planes, inliers
```

Horizontal Oriented Bounding Boxes Fitting

To approximate the geometry of wall instances, we fit Horizontal Oriented Bounding Boxes (H-OBB) to the previously segmented plane clusters. Given that architectural structures typically conform to horizontal and vertical alignments, we restrict the bounding boxes to axis-aligned orientations within the XY-plane.

The H-OBB fitting procedure follows a 2D projection strategy. All segmented points are projected onto the XY-plane, and a convex hull is computed. The convex hull is restricted to a quadrilateral, and each edge is evaluated as a potential candidate for the bounding box length. We rotate each edge to align with the X-axis, compute the minimum-area axis-aligned bounding box, and then select the configuration with the smallest area. The inverse rotation is applied to return the bounding box to its original orientation. This process is visualized in Figure 5.

Topology Correction

Even with high-quality plane and bounding box fitting, errors in wall topology often persist, typically caused by 3D scanning noise. We implement a topology-aware correction stage to ensure consistent and realistic wall geometries. Three main correction operations are introduced:




Figure 5: H-OBB fitting procedure on 2D projected data. Left: 2D projected input data with convex hull. Centre: Axis-aligned bounding box fitted to each edge rotated parallel to the X-axis. Right: Inverse rotation to minimal bounding box, final result.

- **Intersection Correction:** Perpendicular bounding boxes that intersect must be clipped or extended to form clean corner connections.
- **Merging:** Collinear bounding boxes with adjacent or overlapping baselines must be merged into longer walls to reduce redundancy.
- **Redundancy Removal:** Smaller boxes fully enclosed by larger ones are treated as artifacts and removed.

Corrections are based on geometric analysis (adjacency, parallelism, orthogonality) as well as domain knowledge. Perpendicular pairs are found by analysing centrelines of all bounding boxes and computing intersection points. If the endpoint of a bounding box lies within a certain distance r of an intersection, the box is clipped or extended accordingly as shown in Figure 6a.




Figure 6: Wall bounding box refinement. (a) Extending and clipping bounding boxes. (b) Merging bounding boxes.

For merging, bounding boxes with endpoints within a distance threshold are joined by averaging the height and Z-coordinates of the endpoints while preserving horizontal alignment as shown in Figure 6b. Intersection correction and merging procedures are applied iteratively, as each operation may introduce new merge or correction candidates. The distance thresholds for these operations (e.g., intersection radius r , merge distance) are set once per project based on domain knowledge such as typical wall widths and corridor dimensions.

Finally, to refine door geometry, we project door bounding boxes into their parent wall to ensure alignment as seen in Figure 7a. For oversized doors, which may arise from

clustering errors, we apply a width threshold and split them into multiple instances. Figure 7b shows the corrected doors after offsets were applied to ensure a realistic spacing between the door instances.

These algorithms yield a geometrically and topologically coherent representation of walls and doors, forming a robust foundation for accurate IFC models.




Figure 7: Bounding box refinement. (a) projecting the bounding box into the parent geometry. (b) splitting bounding boxes too wide.

Evaluation

The primary method for assessing reconstruction accuracy is to compare reconstructed BIMs against ground truth models. While simple cuboid representations can be directly compared, walls and doors may be reconstructed as more complex geometries, making standard geometric comparison difficult. Therefore, we use several evaluation metrics to quantify reconstruction accuracy.

Metrics

The main metric used is the 3D-Intersection over Union (3D-IoU). However, due to its limitations in handling small spatial misalignments, we introduce the Volumetric Intersection over Union (vIoU) as an alternative.

3D-Intersection over Union

The 3D-Intersection over Union (3D-IoU) is calculated as the ratio of the volume of the intersection to the volume of the union of the reconstructed and ground-truth bounding boxes. This metric captures both shape and spatial alignment, providing a compact representation of how well two volumes overlap. However, 3D-IoU is highly sensitive to minor deviations in size or alignment: because volume scales cubically with linear dimensions, even a small positional offset disproportionately reduces the intersection-to-union ratio. Furthermore, 3D-IoU requires instance-level assignment between predicted and ground-truth elements, which becomes ambiguous when walls are fragmented or merged during reconstruction.

Volumetric Intersection over Union

To address these limitations, we propose Volumetric Intersection over Union (vIoU), computed on a voxel grid. Voxel-based IoU is well established in 3D object reconstruction and occupancy prediction Agnew et al. (2021), where it evaluates per-object shape accuracy. However, its direct application to Scan-to-BIM is non-trivial, as standard per-object voxel IoU inherits the same matching dependency

as 3D-IoU. Our vIoU addresses this by operating at the class level: all reconstructed geometry of a given type (e.g., walls) and all corresponding ground-truth geometry are voxelized jointly, and overlap is computed holistically without instance assignment. This is conceptually analogous to how semantic segmentation IoU evaluates per-class predictions in 2D, but applied to 3D volumetric BIM evaluation. We use a voxel size of 5 cm, corresponding to a tolerance of ± 2.5 cm, which we consider reasonable for the structural elements addressed here. The 3D space is discretized, and a voxel is marked as occupied if its centroid lies inside a bounding box. For both ground truth and reconstructed geometries, we count all occupied voxels and compute vIoU as the ratio of their intersection and union. An example is illustrated in Figure 8.

This voxel-based approach simplifies comparison by avoiding complex geometric overlap calculations, exhibits more stable behavior under small misalignments (as voxel resolution acts as a tolerance threshold), and eliminates the need for instance-level bounding box matching. The latter is particularly important for walls, which may appear as a single bounding box, be split into segments, or be fragmented due to incomplete scans. Unlike 3D-IoU, which would assign ground truth to only one predicted instance, vIoU evaluates overlap holistically at the voxel level.

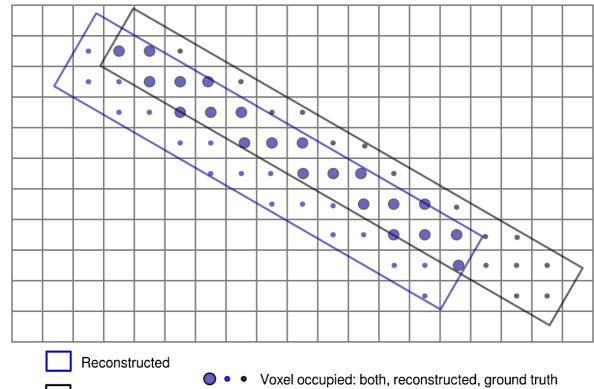


Figure 8: Volumetric IoU computation. Each voxel is classified as belonging to the reconstructed or ground-truth bounding box if its centroid lies within the corresponding geometry.

Results

We evaluate our reconstruction pipeline on two datasets: Our newly proposed German Hospital dataset (DeKH) and the CV4AEC Scan-to-BIM Challenge dataset. While DeKH features largely empty interiors, CV4AEC is fully furnished featuring a lot of clutter and complex occlusions. Furthermore, an ablation study is presented to investigate the impact of individual components in our pipeline and assess the effect of semantic segmentation accuracy.

CV4AEC

The CV4AEC dataset is more challenging due to its cluttered, incomplete scans and presence of interior furniture and occlusions. Six scans from the test set are used for evaluation, with ground-truth BIMs manually created from the input point clouds. Results are presented in Table 1.

Table 1: Results of the evaluation on the CV4AEC dataset test scenes. The vIoU is calculated for a 5 cm voxel size.

Metrics	Office_1	Office_2	Office_3	Office_4	Parking_1	Parking_2	Mean %	Segmentation IoU
Wall 3D-IoU [-]	0.347	0.358	0.261	0.297	0.280	0.318	31.0	0.672
Wall vIoU [-]	0.437	0.383	0.259	0.302	0.386	0.420	36.5	
Door 3D-IoU [-]	0.388	0.452	0.157	0.182	0.192	0.000	22.9	0.586
Door vIoU [-]	0.367	0.407	0.150	0.149	0.247	0.000	22.0	
Column 3D-IoU[-]	0.382	0.400	0.011	0.033	0.197	0.105	18.8	0.293
Column vIoU [-]	0.404	0.547	0.025	0.057	0.250	0.125	23.5	
Mean 3D-IoU [-]	0.372	0.403	0.143	0.171	0.223	0.141	24.2	0.517
Mean vIoU [-]	0.403	0.446	0.145	0.169	0.294	0.182	27.3	

In the first two scenes, (Office_1 and Office_2) with standard wall and door configurations, our pipeline performs best due to alignment with our assumptions. Scenes Office_3 and Office_4 contain open spaces, cubicles, and complex facade elements, making the reconstruction more challenging. The last two scenes (Parking_1 and Parking_2) are from underground parking structures. Although some improvement in column detection is observed, results remain lower than in the office scenes. Notably, no doors are detected in (Parking_2), likely due to closed doors during scanning and segmentation failure suggesting that a dedicated detection algorithm for doors might be beneficial. The segmentation IoU presented in the last column gives an indication of the challenge these types of data are posing to state-of-the-art point cloud segmentation methods.

German Hospital Dataset (DeKH)

Our DeKH dataset is a new public dataset including four scans from three buildings in an unused, largely unfurnished hospital building in Germany. Building **A**, over 100 years old, features structural walls, central corridors, and a symmetrical layout. Building **B** includes an Intensive Care Unit (ICU) area with integrated equipment, while building **C** contains empty surgical rooms with built-in furniture. In addition to the point clouds annotated following a construction ontology-based guideline (Kaufmann et al., 2023), we provide the manually created BIMs for all of the scans. The results of the reconstruction are presented in Table 2.

Table 2: Results of the evaluation on the DeKH dataset. The vIoU is calculated for a voxel size of 5 cm.

Metrics	A		B	C	Mean %
	1st floor	2nd floor	ICU	Surgery	
Wall 3D-IoU [-]	0.395	0.219	0.270	0.328	30.3
Wall vIoU [-]	0.553	0.463	0.388	0.451	46.4
Door 3D-IoU [-]	0.417	0.104	0.193	0.223	23.4
Door vIoU [-]	0.402	0.103	0.219	0.236	24.0
Mean 3D-IoU[-]	0.406	0.162	0.232	0.275	26.9
Mean vIoU [-]	0.478	0.283	0.304	0.344	35.2

Performance in the **A** building’s second floor is lower than the first floor. The pipeline faced challenges in reconstructing smaller wall compartments. The door reconstruction had issues due to sparse segmentation. In the **B** ICU building, performance improves. Door reconstruction is accurate but scored poorly due to a mismatch in door representations between the ground truth of thin door leaves and the recon-

structed geometry including frame and lining. In building **C**, some furniture surfaces were misclassified as walls, and some open doors lead to double reconstructions.

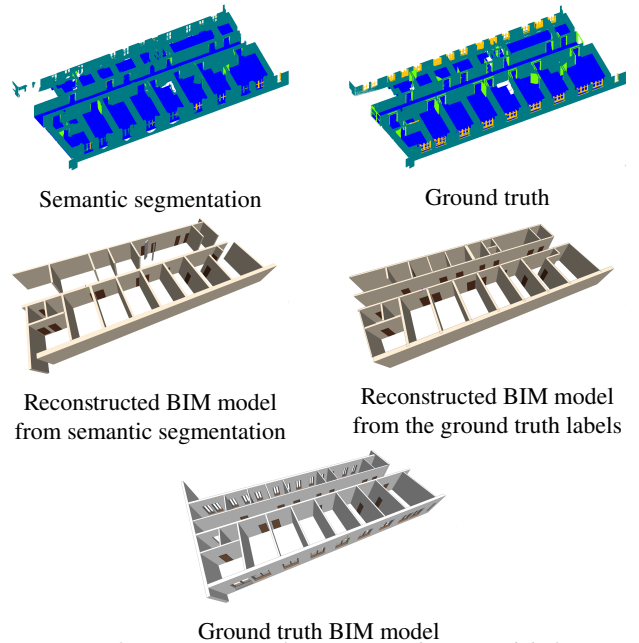


Figure 9: Comparison of reconstructed BIM models from semantic segmentation vs ground truth in DeKH–building A.

The ablation study in Table 3 presents three configurations: **Baseline**, which follows the reconstruction approach of Kaufmann et al. (2022) using RANSAC for primitive fitting and H-OBB estimation without topology correction; **Ours**, representing the full pipeline with HYSAC, H-OBB fitting, and topology refinement; and **GT labels**, which applies the same method as **Ours** but uses ground truth semantic labels to isolate the effect of segmentation quality. The results confirm our pipeline significantly improves wall and door reconstruction accuracy compared to the baseline. Notably, even with ground truth labels, wall reconstruction does not improve. Our approach achieves a 4.9% higher bounding box IoU while volumetric IoU for walls shows only a slight advantage with ground truth labels. The door reconstruction improves considerably with ground truth segmentation. These results highlight that, while our method can tolerate some incomplete segmentation especially for large structures like walls, accurate semantic segmentation remains crucial for smaller elements like doors. The qualitative results in Figure 9 also support these findings.

Table 3: Ablation study on the first floor of the DeKH A building.

Metrics	Baseline	Ours	GT Labels
Wall 3D-IoU [-]	0.371	0.395	0.346
Wall vIoU [-]	0.466	0.553	0.571
Door 3D-IoU [-]	0.322	0.417	0.519
Door vIoU [-]	0.391	0.402	0.521
Mean 3D-IoU[-]	0.346	0.406	0.432
Mean vIoU [-]	0.429	0.478	0.546

Discussion and Conclusion

This work presents a robust and modular pipeline that advances the state of scan-to-BIM automation, generating IFC-compliant BIM models from 3D point clouds. The proposed pipeline achieves competitive reconstruction accuracy across diverse scenes, demonstrating its practical viability and scalability. We contribute to the topic with methodological advances and a new benchmark dataset.

Our approach delivers accurate reconstructions of structural elements using a hybrid methodology of learning-based and geometric processing techniques. Evaluation across two datasets confirms the pipeline’s robustness and effectiveness, particularly for walls and doors, where reconstruction quality exceeds a RANSAC-based baseline. Even with incomplete or noisy data, the system maintained reliable performance. Direct comparison with existing scan-to-BIM methods is currently not feasible due to the absence of open-source implementations and standardized evaluation protocols. The CV4AEC challenge provides the closest available shared benchmark, where the proposed method was evaluated competitively. The release of DeKH aims to further support reproducible comparison in future work.

A major contribution is the introduction of the **DeKH** benchmark dataset. The release of DeKH, with high-resolution scans, ground truth BIMs, and semantic annotations, provides a new benchmark for scan-to-BIM.

Additionally, we propose vIoU, which adapts voxel-based IoU to class-level evaluation without instance matching, providing more robust comparison in scenes with misaligned, fragmented, or variably segmented elements.

To strengthen the pipeline, future efforts should expand beyond the current object classes toward more functional BIMs. The reconstruction algorithms are adaptable and could extend to elements such as windows, stairs, and mechanical, electrical, and plumbing (MEP) systems.

The topology refinement procedures showed sensitivity to parameter selection. For instance, these operations can introduce inaccuracies when small segments are wrongly merged. Thus, adaptive, context-aware refinement strategies or learned alternatives could address this vulnerability. Our pipeline adopts the Manhattan world assumption, which constrains reconstructed geometry to predominantly orthogonal orientations. Consequently, it is best suited to buildings with rectilinear layouts, such as offices, hospitals, and residential complexes, which constitute the majority of typical building stock. Structures with significant non-orthogonal elements (such as curved walls, angled façades, or irregular floor plans) would not be handled correctly,

as the axis-aligned clustering would either discard or misassign their points, and the bounding box fitting would produce poor approximations. Relaxing this assumption, for instance through data-driven orientation estimation or by supporting polygonal wall primitives, is an important direction for future work.

Looking ahead, extending the pipeline to support a broader range of object classes and incorporating subcomponent-level reasoning will be important steps toward producing complete and functionally rich BIMs. In parallel, research into end-to-end learning approaches that can infer BIM semantics and geometry directly from raw point cloud data and other sources will help move the field forward.

Acknowledgements

This research was funded by the European Union as part of the projects: HumanTech (Grant Agreement 101058236) and ShieldBOT (Grant Agreement 101235093).

References

- Agnew, W., Xie, C., Walsman, A., Murad, O., Wang, Y., Domingos, P., and Srinivasa, S. (2021). Amodal 3d reconstruction for robotic manipulation via stability and connectivity. In Conference on robot learning. PMLR.
- Anagnostopoulos, I., Belsky, M., and Brilakis, I. (2016). Object boundaries and room detection in as-is bim models from point cloud data. In ICCCBCE.
- Armeni, I., Che, E., Fischer, M., Hall, D., Jung, J., Li, F., Olsen, M., Pollefeys, M., Turkan, Y., Rastiveis, H., and et al. (2024). Computer vision in the built environment.
- Armeni, I., Sax, S., Zamir, A. R., and Savarese, S. (2017). Joint 2d-3d-semantic data for indoor scene understanding. arXiv:1702.01105.
- Armeni, I., Sener, O., Zamir, A. R., Jiang, H., Brilakis, I., Fischer, M., and Savarese, S. (2016). 3d semantic parsing of large-scale indoor spaces. In CVPR. Ieee.
- Bassier, M., Klein, R., Van Genechten, B., and Vergauwen, M. (2018). Ifc wall reconstruction from unstructured point clouds. ISPRS Annals.
- Bassier, M. and Vergauwen, M. (2020). Unsupervised reconstruction of building information modeling wall objects from point cloud data. Automation in Construction.
- Bassier, M., Yousefzadeh, M., and Vergauwen, M. (2020). Comparison of 2d and 3d wall reconstruction algorithms from point cloud data for as-built bim. ITcon.
- Chang, A., Dai, A., Funkhouser, T., Halber, M., Niebner, M., Savva, M., Song, S., Zeng, A., and Zhang, Y. (2017). Matterport3d: Learning from rgb-d data in indoor environments. In 3DV. IEEE.

- Chen, X., Ma, H., Wan, J., Li, B., and Xia, T. (2017). Multi-view 3d object detection network for autonomous driving. In CVPR.
- Coughlan, J. M. and Yuille, A. L. (1999). Manhattan world: Compass direction from a single image by bayesian inference. In ICCV.
- Dai, A., Chang, A. X., Savva, M., Halber, M., Funkhouser, T., and Nießner, M. (2017). Scannet: Richly-annotated 3d reconstructions of indoor scenes. In CVPR.
- Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In KDD. AAAI Press.
- Gankhuyag, U. and Han, J.-H. (2021). Automatic bim indoor modelling from unstructured point clouds using a convolutional neural network. *Intelligent Automation & Soft Computing*.
- Gourguechon, C., Macher, H., and Landes, T. (2022). Automation of as-built bim creation from point cloud: An overview of research works focused on indoor environment. *ISPRS Annals*.
- Graham, B., Engelcke, M., and van der Maaten, L. (2018). 3d semantic segmentation with submanifold sparse convolutional networks. In CVPR. Ieee.
- Hu, Q., Yang, B., Xie, L., Rosa, S., Guo, Y., Wang, Z., Trigoni, N., and Markham, A. (2020). Randla-net: Efficient semantic segmentation of large-scale point clouds. In CVPR. Ieee.
- Kanayama, H., Chamseddine, M., Guttikonda, S., Okumura, S., Yokota, S., Stricker, D., and Rambach, J. (2025). Tof-360-a panoramic time-of-flight rgb-d dataset for single capture indoor semantic 3d reconstruction. In CVPRW.
- Kaufmann, F., Chamseddine, M., Guttikonda, S., Glock, C., Stricker, D., and Rambach, J. (2023). Ontology-based semantic labeling for rgb-d and point cloud datasets. In EC3. European Council on Computing in Construction.
- Kaufmann, F., Glock, C., and Tschickardt, T. (2022). Scalebim: Introducing a scalable modular framework to transfer point clouds into semantically rich building information models. In EC3. European Council on Computing in Construction.
- Kim, H. and Kim, C. (2021). 3d as-built modeling from incomplete point clouds using connectivity relations. *Automation in Construction*.
- Macher, H., Landes, T., and Grussenmeyer, P. (2017). From point clouds to building information models: 3d semi-automatic reconstruction of indoors of existing buildings. *Applied Sciences*.
- Ochmann, S., Vock, R., and Klein, R. (2019). Automatic reconstruction of fully volumetric 3d building models from oriented point clouds. *ISPRS Annals*.
- Ochmann, S., Vock, R., Wessel, R., and Klein, R. (2016). Automatic reconstruction of parametric building models from indoor point clouds. *Computers & Graphics*.
- Qi, C. R., Su, H., Mo, K., and Guibas, L. J. (2017). Pointnet: Deep learning on point sets for 3d classification and segmentation. In CVPR. Ieee.
- Schnabel, R., Wahl, R., and Klein, R. (2007). Efficient ransac for point-cloud shape detection. *Computer Graphics Forum*.
- Son, H. and Kim, C. (2017). Semantic as-built 3d modeling of structural elements of buildings based on local concavity and convexity. *Advanced Engineering Informatics*.
- Tang, S., Li, X., Zheng, X., Wu, B., Wang, W., and Zhang, Y. (2022). Bim generation from 3d point clouds by combining 3d deep learning and improved morphological approach. *Automation in Construction*.
- Thomson, C. and Boehm, J. (2015). Automatic geometry generation from point clouds for bim. *Remote Sensing*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *NeurIPS*.
- Wang, Y., Sun, Y., Liu, Z., Sarma, S. E., Bronstein, M. M., and Solomon, J. M. (2019). Dynamic graph cnn for learning on point clouds. *ACM ToG*.
- Wu, C., Yuan, Y., Tang, Y., and Tian, B. (2021). Application of terrestrial laser scanning (tls) in the architecture, engineering and construction (aec) industry. *Sensors*.
- Wu, X., Jiang, L., Wang, P.-S., Liu, Z., Liu, X., Qiao, Y., Ouyang, W., He, T., and Zhao, H. (2024). Point transformer v3: Simpler faster stronger. In CVPR.
- Xiong, B., Jin, Y., Li, F., Chen, Y., Zou, Y., and Zhou, Z. (2023). Knowledge-driven inference for automatic reconstruction of indoor detailed as-built bims from laser scanning data. *Automation in Construction*.
- Xu, J., Chen, K., Xue, F., and Lu, W. (2021). 3d point cloud data enabled facility management: A critical review. In CRIOCM. Springer.
- Zhao, H., Jia, J., and Koltun, V. (2020). Exploring self-attention for image recognition. In CVPR.
- Zhao, H., Jiang, L., Jia, J., Torr, P., and Koltun, V. (2021). Point transformer. In ICCV. Ieee.
- Zheng, J., Zhang, J., Li, J., Tang, R., Gao, S., and Zhou, Z. (2020). Structured3d: A large photo-realistic dataset for structured 3d modeling. In ECCV. Springer.

Supplementary Material: BIMStruct3D: A Fully Automated Hybrid Learning Scan-to-BIM Pipeline with Integrated Topology Refinement

Mahdi Chamseddine^{1,2}, Fabian Kaufmann², Marius Schellen², Christian Glock²,
 Didier Stricker^{1,2}, and Jason Rambach¹

¹German Research Center for Artificial Intelligence (DFKI), Kaiserslautern, Germany

²RPTU Kaiserslautern-Landau, Kaiserslautern, Germany

CV4AEC Scene Names

The mapping of the scene names used in Table 1 is provided in Table 4 for reference and potential comparison of results.

Table 4: CV4AEC dataset (Armeni et al., 2024) scene name mapping.

Short Name	CV4AEC Scene
Office_1	08_ShortOffice_01_F1
Office_2	08_ShortOffice_01_F2
Office_3	11_MedOffice_05_F2
Office_4	11_MedOffice_05_F4
Parking_1	25_Parking_01_F1
Parking_2	25_Parking_01_F2

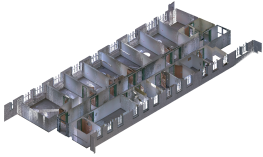
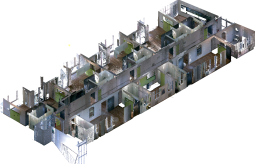
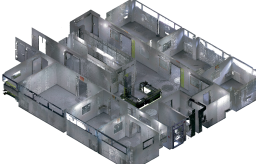
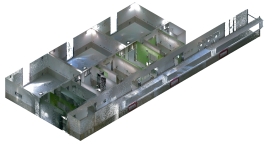
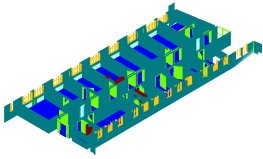
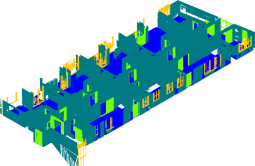
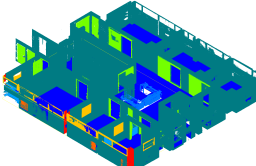
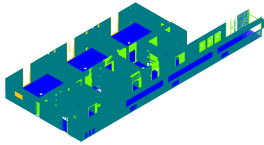
The CV4AEC challenge comprises of two challenges, a 2D and 3D one. Here we use the test data of the 3D challenge for evaluation. The training data of the challenge was used for training the point cloud segmentation model.

DeKH Dataset

The German hospital dataset was recorded in an empty hospital with various types of facilities and rooms: offices, toilets, reception, stairs, operation rooms, and more. It boasts a variety of window and door shapes as well as some embedded furniture. For annotation, the ontology based segmentation guideline by Kaufmann et al. (2023) was used. It is a construction specific joint guideline for images and point clouds.

The dataset is divided into four areas shown in Table 5 and providing high quality semantic segmentation labels as well as manually created BIM ground truth models. The quality of the scans and labels, as well as the variety of the buildings and rooms make this dataset an important contribution to the scan-to-BIM research and the construction community.

Table 5: Overview of the DeKH dataset, including colored point clouds, semantic labels, and ground truth BIM models

	Building A - 1st floor	Building A - 2nd floor	Building B - ICU	Building C - Surgery
PCD RGB				
PCD labeled				
GT BIM	