Towards Automated 3D BIM Reconstruction of Existing Industrial Buildings from Point Cloud Data

Patricia González-Cabaleiro[1], Muataz S. A. Albadri[1], Antonio Fernández[1], Lucía Díaz-Vilariño[1]

[1]CINTECX, Universidade de Vigo, GeoTECH Group, 36310 Vigo, Spain
{patricia.gonzalez.cabaleiro, muatazsafaaabed.albadri, antfdez, lucia}@uvigo.gal

Keywords: 3D Point Clouds, Semantic Segmentation, Industrial Buildings, Profile Recognition, Morphological Processing, Structural Reconstruction.

Abstract

This paper presents a comprehensive methodology for the automated semantic segmentation and 3D reconstruction of industrial building elements, including roof panels, floor, rafters, purlins, and columns, from unstructured point clouds. The proposed approach integrates orientation-based filtering, projection onto characteristic planes, morphological analysis, and optimization-based I-profile fitting to generate accurate 3D models. The workflow begins with point cloud preprocessing, where the data are aligned with the building axes and cleaned of outliers, followed by subdivision into two subsets based on local surface orientation. Binary projections are then processed to extract element contours, while roof slopes and panel inclinations are automatically estimated to guide the reconstruction of rafters and purlins. The method was validated on a real-case study of 930 m² industrial warehouse scanned with a mobile laser scanner, resulting in a raw dataset of seven million points. The segmentation achieved F1-scores above 0.90 for floors, roof panels, rafters, and columns, and 0.75 for purlins. Profile fitting yielded an average width error of 3.8%, confirming the robustness and reliability of the reconstruction across diverse structural components.

1. Introduction

The effective maintenance, and renovation of industrial facilities require a comprehensive understanding of their as-built conditions and their alignment with design documentation, to ensure operational safety, efficiency, and long-term reliability (Azhar, 2011; Son et al., 2013). Building Information Modelling (BIM) has emerged as a powerful solution, providing a semantically rich, object-oriented framework for digitally representing facilities throughout their lifecycle (Toth and Petrie, 2018; Chen and Xue, 2023). The accurate representation of as-built information within BIM environments is fundamental to performance and safety (Rausch and Haas, 2021).

Capturing the geometric reality of these environments typically relies on non-contact sensing technologies such as photogrammetry and laser scanning technologies, which produce dense, high-precision three-dimensional point clouds that serve as a robust geometric foundation for as-built documentation (Son et al., 2013; Wang and Kim, 2019; Yang et al., 2020; Mirzaei et al., 2023). Despite their efficiency in capturing complex geometries, converting point clouds into parametric BIM representations (commonly referred to as Scan-to-BIM) remains a major bottleneck in digital reconstruction, where geometric processing alone can account for up to 90% of total modelling time (Ochmann et al., 2016; Agapaki et al., 2018; Yang et al., 2020; Zhang et al., 2025).

This challenge is particularly acute in industrial facilities, which present far greater complexity than residential or commercial buildings due to their geometric complexity, dense layouts, and the abundance of machinery and equipment (Noichl et al., 2025). Point clouds captured in such environments often exhibit gaps, occlusions and noise, while reflective metallic surfaces further degrade accuracy and complicate object recognition, critical issues where sub-millimetre precision is often required (Anil et al., 2012; Noichl et al., 2025). These difficulties are particularly evident in the reconstruction of structural elements, which account for roughly one-third of all components in industrial facilities (Agapaki et al., 2018). Among these, elongated members with standardized cross-sections, such as I-beams, channels, and circular hollow sections (CHSs), are especially frequent yet remain among the most demanding to reconstruct due to occlusions, interference from surrounding equipment, and their intricate geometries (Agapaki and Brilakis, 2020, 2022).

Driven by these challenges, there is a growing need for frameworks capable of automating the detection and reconstruction of key elements in industrial environments (Smith, 2020; Noichl et al., 2025; Bosché et al., 2015). This study addresses this challenge by proposing a methodology for the automated segmentation, reconstruction, and recognition of key facility elements from point cloud data. The proposed approach integrates geometric reasoning with orientation-based filtering and spatial regularity detection to isolate both structural elements, such as columns, rafters (longitudinal beams supporting the roof panels), and purlins (transverse beams resting on rafters, supporting roof panels), and non-structural elements, including roof panels and floors. By focusing on the precise identification of standardized steel sections, the method provides a reliable foundation for structural evaluation, quality assessment, and the generation of semantically rich as-built BIM, directly addressing a critical automation gap in the digital documentation of industrial facilities (Agapaki et al., 2018; Noichl et al., 2025)

This paper is organized as follows. Section 2 reviews the state of the art in automatic detection and reconstruction of industrial elements from point cloud data. Section 3 describes the proposed methodology for semantic segmentation and 3D reconstruction of key structural and non-structural elements. Section 4 evaluates the method using a real-world industrial case study, assessing segmentation accuracy, profile fitting, and reconstruction performance. Finally, Section 5 summarizes the main findings and future research directions.

2. Related Work

The automatic segmentation and reconstruction of steel structural elements from point clouds are critical steps in developing digital twins of industrial facilities. Despite significant advances, this task remains challenging due to the presence of noise, particularly in Mobile Laser Scanning (MLS) data, the diversity of structural shapes, and the occlusions inherent in the scanned environment. Existing methods can be broadly categorized into geometry-driven, learning-based, and integrated parametric reconstruction approaches, each with distinct strengths and limitations when confronted with the complexity of real-world industrial environments.

Early research on point cloud interpretation primarily focused on fitting predefined geometric primitives, such as cylinders, planes, or rectangular profiles, to point cloud data. These model-driven methods commonly rely on RANdom SAmple Consensus (RANSAC) or Hough Transform (Agapaki and Brilakis, 2020; Mirzaei et al., 2023) which have proven effective for simple, repetitive geometries like pipelines, with commercial tools like EdgeWise demonstrating satisfactory performance (Agapaki et al., 2018). Cylinder detection has received particular attention, with reported precision and recall rates of approximately 62% and 70%, respectively (Agapaki et al., 2018; Agapaki and Brilakis, 2020).

Despite their utility, purely geometric methods are highly sensitive to parameter tuning and often require manual preprocessing or prior knowledge of target geometry (Agapaki and Brilakis, 2020). This limitation was clearly illustrated by Anil et al. (2012), who evaluated four manual geometric strategies (point-to-point, edge-distance, plane-intersection, and cross-section tracing) for steel components detection and found that, even under controlled conditions, best-performing methods achieved only 18.75% accuracy for columns and 39.68% for beams when compared to standard AISC sections.

To enhance robustness, later studies integrated geometric reasoning with contextual and structural information. Yang et al. (2020) and (Noichl et al., 2025) developed semi-automated approaches that use PCA, cross-section fitting, and plane intersection techniques to estimate the orientation and location of elongated steel members. Complementing these, the 2D-slice-based technique by (Smith, 2020) improved computational efficiency by performing image-based segmentation on planar slices of the 3D point cloud.

The emergence of deep learning marked a paradigm shift by enabling models to learn discriminative spatial features directly from raw point clouds. Foundational architectures such as PointNet Charles et al. (2017), PointNet++ Qi et al. (2017), KPConv Thomas et al. (2019), and MinkowskiNet Choy et al. (2019) have demonstrated strong performance in 3D semantic segmentation for both synthetic and real-world datasets.

This evolution led to hybrid frameworks that combine deep learning with geometric constraints, achieving notable success in steel profile segmentation. For instance, Agapaki and Brilakis (2020) introduced CLOI-NET, a network trained to segment key industrial steel profiles (C-, L-, O-, and I-sections) from TLS point clouds, reporting up to 80% reductions in manual labour and recall and precision rates of 78.25% and 74.25% for I-beams. Likewise, Mirzaei et al. (2023) proposed an end-to-end inspection framework that integrates geometric definitions, spatial relationships, and a PointNet-based classifier, achieving 94% average classification accuracy on mixed real and synthetic datasets.

Nevertheless, misclassification between visually similar profiles, such as circular and rectangular hollow sections, remains a persistent challenge (Mirzaei et al., 2023). Performance also tends to degrade in the presence of occlusions, noise, and less common profile types such as channels and angles (Agapaki et al., 2018; Smith, 2020). Recent studies have therefore emphasized the incorporation of structural priors and contextual reasoning to improve segmentation reliability under cluttered conditions (Noichl et al., 2025).

Beyond segmentation, research has increasingly focused on converting point clouds into semantically rich, as-built Building Information Models (BIM). Rausch and Haas (2021) proposed a semi-automated BIM updating process that aligns point clouds to a Building Coordinate System (BCS) using surveyed targets or structural features, followed by RANSAC-based segmentation and parameter optimization. This approach allows parametric model updates based on as-built data, including non-rigid deformations, though it still relies heavily on accurate feature extraction. Building on this concept, Noichl et al. (2025) presented a comprehensive Scan-to-BIM pipeline that estimates part orientations and uses procedural modelling to generate IFC-compliant representations from segmented point clouds.

These frameworks underscore a key dependency: robust BIM parametrization is intrinsically limited by the accuracy and completeness of upstream segmentation. Without reliable identification of steel profiles and their spatial relations, the automation of digital twin workflows remains constrained. Parallel efforts have explored optimization-based frameworks for structural refinement and performance-driven modelling. For example, Cucuzza et al. (2024) developed a multi-objective optimization method based on the SPEA-II algorithm to jointly optimize size, shape, and topology, illustrating the growing convergence between automated reconstruction, parametric modelling, and structural performance analysis within digital twin environments.

Despite substantial progress, significant challenges persist. Many existing approaches remain tailored to specific element types or case studies, often requiring manual tuning and lacking generalizability across diverse industrial structures (Agapaki and Brilakis, 2020; Bosché et al., 2015). Even advanced systems such as CLOI-NET (Agapaki and Brilakis, 2020) fall short of achieving fully automated, high-accuracy segmentation, particularly in the presence of complex geometries, occluded components, and less common profile types such as channels and angles (Agapaki et al., 2018; Smith, 2020; Mirzaei et al., 2023). Reliable detection increasingly demands the integration of low-level geometric reasoning with high-level contextual and spatial understanding (Mirzaei et al., 2023; Noichl et al., 2025).

To address these limitations, our work proposes a robust and generalizable segmentation framework that integrates geometric reasoning with orientation-based filtering and spatial regularity analysis, thereby bridging the gap between raw point cloud data and the automated generation of semantically rich as-built BIM models for complex industrial facilities. Unlike recent deep learning approaches, our framework is non–learning-based by design, emphasizing transparency, reproducibility, and independence from large annotated datasets, which remain scarce for specialized industrial structures. This choice also ensures deterministic outcomes and allows segmentation to be guided directly by the underlying physical and geometric principles of the structure, making it inherently suitable for engineering interpretation and model validation.

3. Method

This section introduces a methodology for the semantic segmentation and 3D reconstruction of fundamental industrial building elements, including roof panels, floor, rafters, purlins and columns (Figure 1), from unstructured 3D point clouds. As illustrated in Figure 2, the approach combines geometric reasoning, projection-based discretization, and profile-fitting techniques to produce accurate 3D reconstructions suitable for structural analysis and digital twin applications.

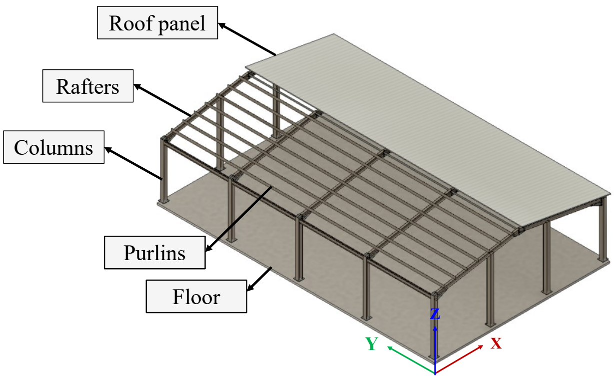

Figure 1. Building elements for segmentation and reconstruction.

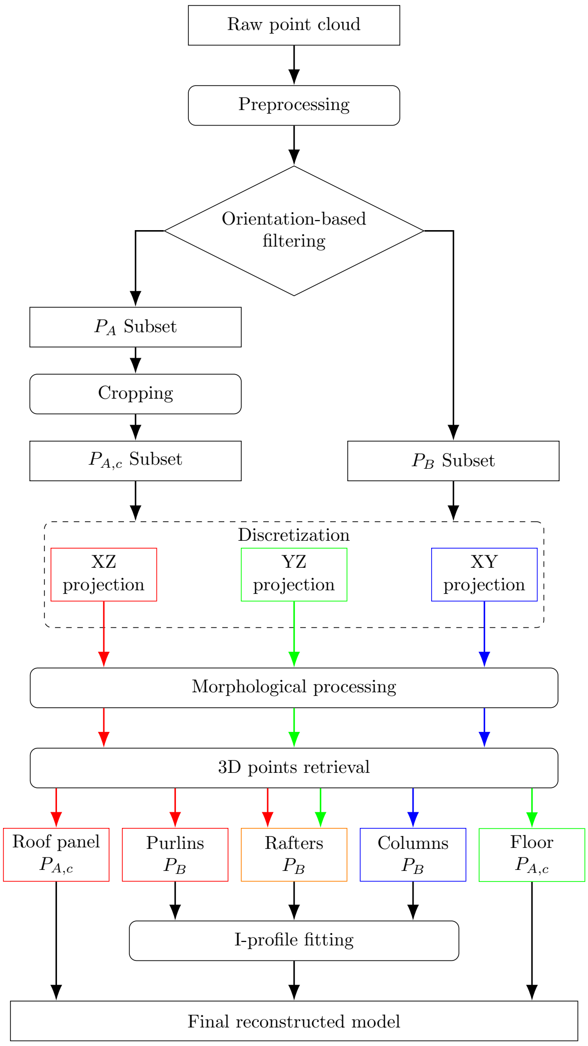

Figure 2. Workflow of the proposed methodology.

3.1 Preprocessing

Prior to segmentation, the point cloud is preprocessed by first aligning it with the principal Cartesian axes (X-axis corresponds to the transverse direction, Y-axis to the longitudinal direction, and Z-axis to the vertical direction) to establish a consistent geometric reference frame. Following the alignment, exterior points not belonging to the studied building are identified by projecting the cloud onto the XY, XZ, and YZ planes and detecting building contours from binary images; points outside the largest contour in each projection are discarded.

3.2 Segmentation

Segmentation begins with an orientation-based filtering of the resulting point cloud $P = \{p_i = (x_i, y_i, z_i) \in \mathbb{R}^3\}_{i=1}^{N}$, where N is the total number of points in the cloud. A normal vector $\mathbf{n}_i$ is estimated for each point p_i using a hybrid KD-tree search with a fixed radius (r) and a maximum number of neighbours (n), capturing the local surface geometry. Each point is then classified by its orientation angle $\theta_i = \arccos\left(|\mathbf{n}_i \cdot \mathbf{v}|\right)$ relative to the vertical axis $\mathbf{v} = [0, 0, 1]^\top$, using a threshold α that separates the cloud into two subsets. Points with $\theta_i > \alpha$ form the A subset (P_A), where the roof panels and the floor are more prominently represented, while points with $\theta_i \leq \alpha$ form the B subset (P_B), which primarily includes columns, rafters, and purlins.

Before further analysis, the A subset P_A is cropped along the Z-axis using percentile thresholds, r_p for roofs and f_p for the floor, to remove irrelevant points and reduce complexity in the next steps. The resulting cropped subset ($P_{B,c}$), together with the B subset (P_B), are then projected onto the corresponding planes and discretized into a 2D grid via $\Pi : \mathbb{R}^3 \longrightarrow \mathbb{Z}^2$, with pixel size S defining the resolution. During discretization, each grid cell within the element's bounding box is assigned a value of 1 (white) if occupied by points and 0 (black) if empty, producing a binary representation for morphological processing. At the same time, the complete preprocessed point cloud (P) is projected onto each plane using the same parameters, generating reference binary images that enable consistent 3D point retrieval for all elements.

To facilitate element recognition, the projection planes are chosen based on the geometric characteristics and typical orientation of the elements, so that their shapes are clearly distinguishable. The cropped A subset ($P_{A,c}$) is projected onto the XZ-plane to detect roof panels and onto the YZ-plane to locate the floor. The B subset (P_B) is projected onto the XZ-plane to identify purlins. For rafters, two projections are performed sequentially: first onto the YZ-plane to determine their positions, and then onto the XZ-plane to refine their longitudinal extent. A third projection, performed later once the roof inclination angle (γ) is estimated, involves rotating the YZ-plane by this angle to align the rafter cross-sections with the plane, enabling their accurate geometric characterization. Columns, which are vertical structural elements, are identified using search regions defined by the previously detected rafters. In the longitudinal direction (Y-axis), each rafter extreme defines a search area, while for the transverse direction (X-axis), the total width of the structure determines the number of search areas (two symmetric areas if less than 20 m). Within each search area, a vertical region of interest (ROI) of height h is selected in P_B, corresponding to a section of the column, and then projected onto the XY-plane. These ROI projections generate the binary images that will be used for column extraction.

Mathematical morphology is then applied to all binary images generated from the aforementioned projections to extract element contours. A closing operation is used with kernel size k_s and a specified number of iterations (ℓ). For both $P_{A,c}$ images, only the largest contour is retained. For P_B images, contours are filtered based on area limits defined a_{min} and a_{max} to isolate relevant elements. Purlins are extracted directly. Rafters contours are first localized in the YZ-plane, and contour extraction is refined in the XZ-plane within these regions. Column contours are extracted from the XY-plane images of each previously identified ROI. Since columns are often attached to walls, the largest contour in each morphological image is vertically split, and the resulting sections are evaluated. When one side of the contour is significantly larger, corresponding to wall contact, the contour is trimmed according to the shorter section, ensuring accurate column extraction.

Finally the interior of each detected contour is used as a mask on the binary image of the complete preprocessed point cloud, allowing the corresponding 3D points to be retrieved and enabling precise segmentation and reconstruction of each element. The sequence of operations, from binary images (B), through morphological processing (D) to the final mask used for 3D point retrieval (M) is illustrated in Figure 3.

Once the main elements are segmented, additional geometric attributes can be derived directly from the corresponding point clouds. For the roof, its inclination angle, γ is estimated by fitting multiple planes to the cloud using RANSAC. The normal vectors of these planes define their local slopes and the median of these angles represents the overall roof pitch. This information is crucial for reconstructing rafters and purlins accurately. For rafters, it defines the longitudinal inclination of the span, while for purlins, it determines the transverse rotation of their cross-section required to align with the roof slope. Moreover, knowing the roof slope enables the identification of profiles in inclined beams and ensures that cross-sections are generated perpendicular to their true axis.

3.3 I-profile fitting

With the point clouds of the different elements segmented, the next step involves identifying the structural profiles for subsequent reconstruction. To this end, an *optimization-based fitting* algorithm was developed.

Prior to fitting, each point cloud is oriented such that the cross-section plane is perpendicular to the element's longitudinal axis, ensuring an accurate representation of the profile geometry. Each cluster of points is fitted to an idealized European I-beam (IPE) cross-section, a standardised hot-rolled steel section with an I-shaped geometry commonly used in structural frameworks. Two configurations were considered: vertical, composed of two vertical flanges and a horizontal web, or horizontal, composed of two horizontal flanges and a vertical web. The algorithm minimizes the sum of squared distances between the observed points and the nearest segment of the corresponding IPE outline:

$$\mathcal{L} = \sum_{i:d_i<d_{\max}} (\min\{d_{i,\text{flanges}},\, d_{i,\text{web}}\})^2$$

where d_i = minimum distance from each point to the closest segment (flanges or web),
$d_{i,\text{flanges}}$, $d_{i,\text{web}}$ = distances from point i to the flange and web regions, respectively,
$d_{\max}$ = maximum distance considered for the fit
$\mathcal{L}$ = squared minimum distances summed over points with $d_i < d_{\max}$.

The combination of configuration and nominal IPE size that yields the smallest value of $\mathcal{L}$ is selected as the best-fitting profile. To ensure physically feasible profiles, the algorithm enforces bounds that keep all points within the section's bounding box and linear constraints that maintain realistic spacing between flanges and the web. These conditions guarantee positive dimensions, proper aspect ratios, and consistency with standard steel IPE-profiles, enabling reliable classification and reconstruction of each structural element.

3.4 Reconstruction

Once the point cloud is segmented and all structural elements have undergone I-profile fitting, the reconstruction focuses on determining their 3D positions and orientations. The final reconstructed model is represented using an IFC schema. For roof panels, rafters, and purlins, the previously identified roof planes are used to divide each element along the building's midpoint into left and right sections, illustrated in purple and magenta in Figure 4, with each section analysed separately.

For the roof panels, the slope of each panel is determined by fitting a plane to the 3D points using PCA and extracting the normal vector $\mathbf{n}$ (purple and magenta arrows in Figure 4), where the Z-component indicates the upward direction and the X and Y components define the horizontal tilt. The bounding box corners (orange crosses on Figure 4) are adjusted according to the slope, and the heights of the section centroids along the slope direction are compared to assign the correct positive or negative inclination. The floor is reconstructed in a similar manner by computing the bounding box of its point cloud, identifying the minimum and maximum coordinates along X, Y, and Z, and extracting the four corner points in the XY-plane (Green crosses on Figure 4), with the Z-coordinate representing the floor elevation and the vertical extent defining its thickness.

For each structural element, the centroid of its corresponding cluster within the complete point cloud is used to define its nominal location, while a normal vector is assigned to indicate its longitudinal (span) direction, perpendicular to the element's cross-section. The nominal positions of columns, purlins and rafters are illustrated in Figure 4. Column centroids are marked with blue crosses and extend vertically along the blue vector $\mathbf{n}_c$. Rafters are not represented by crosses, as their centroids coincide with the column locations along the Y-axis; instead,

https://doi.org/10.5194/isprs-archives-XLIX-B2-2026-1237-2026 |

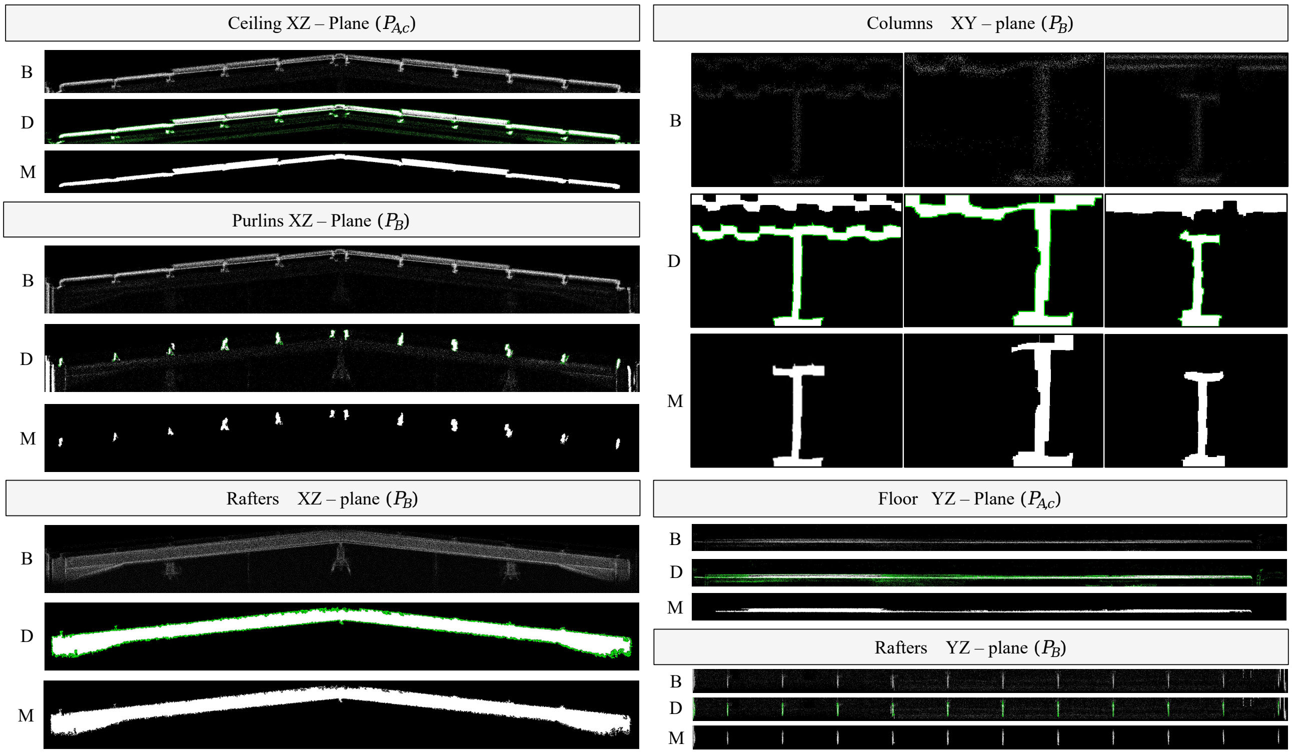

Figure 3. Morphological processing workflow. B indicates the binary images obtained from the projections, D the dilated images showing the processed contours, and M the final masks used for 3D point retrieval. Each row corresponds to a different element type.

an example of their longitudinal extend is shown with the green vector $\mathbf{n}_r$, aligned with the roof slope. Purlins extend horizontally along the yellow vector $\mathbf{n}_p$ and are rotated according to the roof inclination angle γ, referenced in their local coordinate system (X', Z').

4. Results

To demonstrate the applicability of the proposed method, a real-world industrial warehouse with a floor area of 930 m^2, containing furniture, equipment, and other sources of clutter, was analysed as a case study. The building was scanned using a BLK2GO mobile laser scanner, with a range of 25 m, noise level ± 3 mm, and accuracy ± 10 mm, resulting a point cloud of 7 million points that included the floor and roof panels, along with 12 columns, 12 rafters and 12 purlins.

The raw point cloud was first aligned with the main axes and reduced to 2.5 million after removing outliers, using a pixel size of 0.2 m as part of the preprocessing step. The point cloud was then divided into two subsets through orientation-based filtering. Surface normals were estimated using a 0.1 m search radius and up to 30 neighbours, classifying points with normal deviations below 45 ° from the vertical axis as B, and the reminder as A.

Element-specific parameters were tuned to account for variations in scale and point density. For non-structural elements, including ceiling and floor segmentation, a projection pixel size of 0.0015 m/px and a morphological kernel of 2 pixels, applied over 2 iterations, were used, with percentile thresholds of 40% and 25%, respectively. For structural elements, column extraction employed localized regions of interest with 2 m lateral margins and 3 m height, using a projection pixels size of

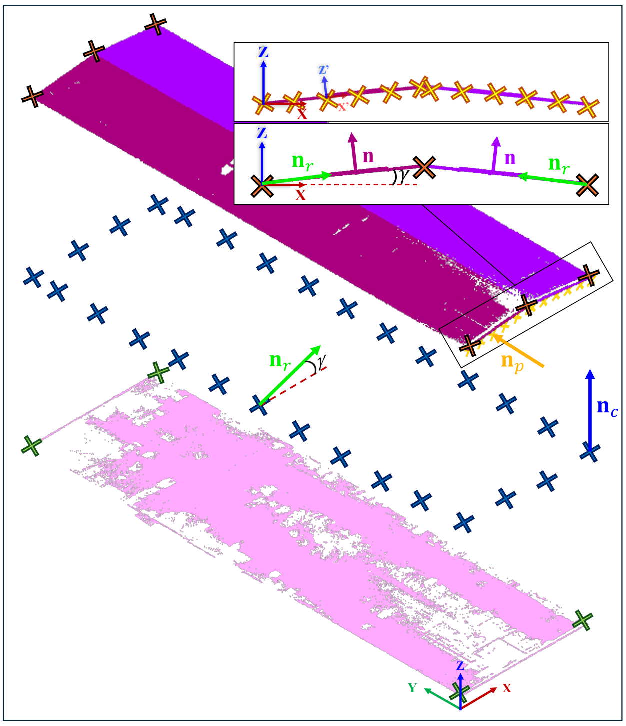

Figure 4. Key elements for model construction. Column centroids (blue crosses), purlin centroids (yellow crosses), roof panels corners (orange crosses), and floor corners (green crosses) are shown, with the rafter orientation indicated by the $\mathbf{n}_r$ vector.

0.001 m/px. Contour filtering was performed based on area an point density, applying morphological kernel of 7–10 pixels over up to 5 iterations. For rafters and purlins, projections were performed on the XZ- and YZ-planes using pixel sizes of 0.005 and 0.0015 m/px, respectively, with kernel size of 2 pixels and 1–2 iterations. Area limits between 0.01 to 0.05 m^2, were set to ensure accurate isolation of structural components despite varying occlusions and clutter. The fitting of IPE profiles for structural elements was performed using geometric parameters from the IPE catalogue, with a maximum deviation tolerance of $d_{\max} = 0.125$ m to constraint the optimization and ensure realistic cross-sectional dimensions.

The method effectively extracted roof panels, floors, rafters, purlins, and columns using *orientation-based filtering*, projection onto characteristic planes, and morphological analysis in less than one minute. For 3D reconstruction, IPE-profiles were matched and fitted using the *optimization-based* algorithm. Segmentation performance was high across most elements, with F1-scores exceeding 0.90 for floors, roof panels, rafters, and columns (Table 1). Purlins exhibited slightly lower performance (F1=0.75), mainly due to partial occlusions and lower point density. Overall accuracy, computed across all points of all elements, was 0.92, reflecting the method’s effectiveness on the complete dataset. The majority of the detected structural elements were accurately matched to standard profiles: IPE 360 for 88.9% of columns, IPE 330 for 94.4% of rafters, and IPE 200 for 64.3% of purlins. The fitted profiles showed an average width error of 3.8% with a standard deviation of 5.9%, confirming consistent and reliable fitting across all structural components. Figure 5 illustrates the method and the resulting segmentation and fitting outcomes.

Element	**Precision**	**Recall**	**F1-score**
Floor	**0.99**	**1.00**	**1.00**
Roof panels	**0.94**	**0.94**	**0.94**
Purlins	0.73	0.77	0.75
Rafters	**0.90**	**0.97**	**0.93**
Columns	**0.91**	**0.98**	**0.94**

Table 1. Segmentation results for the case study. Values ≥ 0.90 are highlighted in bold.

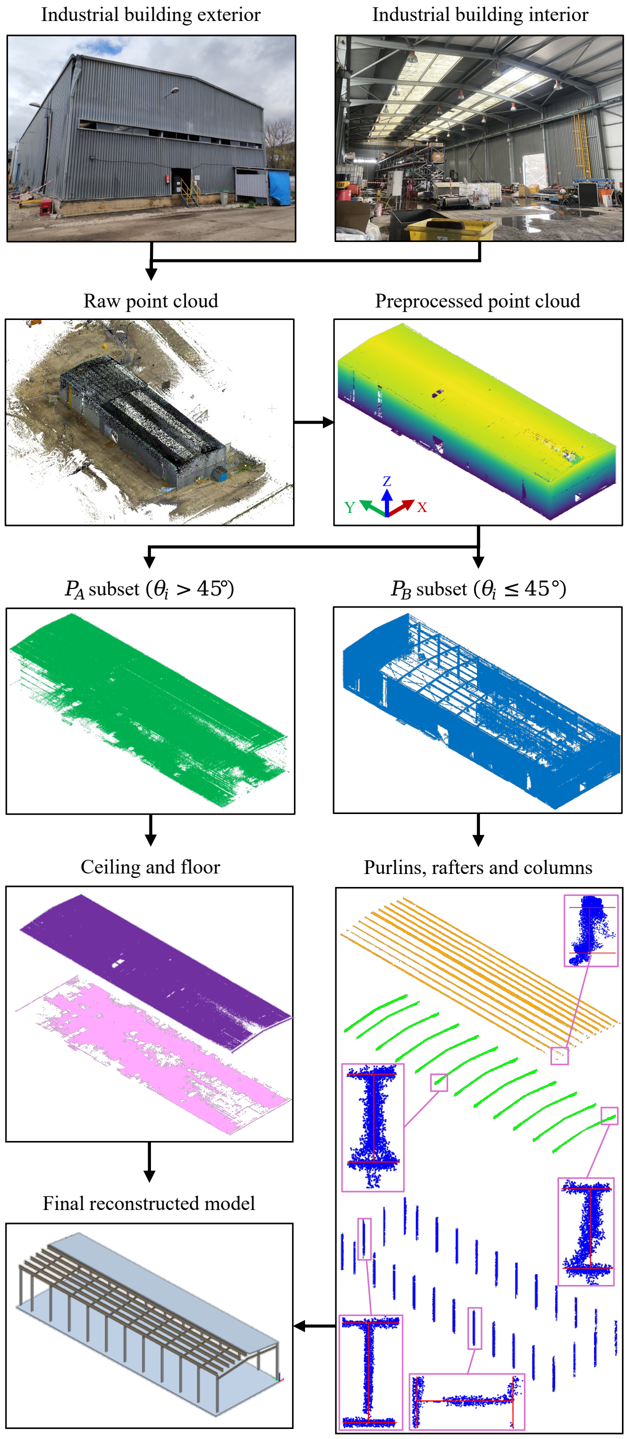

Figure 5. Workflow results from the raw point cloud to the final reconstructed model, including preprocessing, subset division into A ($\theta_i > \alpha$) and B ($\theta_i \leq \alpha$) subsets, detection of non-structural elements (roof panels and floor), structural elements (purlins, rafters, and columns) with profile recognition.

5. Conclusions

This work presents a robust methodology for the semantic segmentation and 3D reconstruction of industrial building elements from unstructured point clouds. The method accurately identifies and reconstructs roof panels, floors, rafters, purlins, and columns, even in the presence of clutter such as furniture and equipment. High segmentation performance, with F1-scores above 0.90 for most elements, and reliable IPE-profile fitting demonstrate the method’s suitability for producing 3D models for analysis and inspection. The workflow efficiently handles large-scale point clouds, providing geometrically and structurally meaningful information for each element.

The quality of the input point cloud significantly influences segmentation accuracy and profile recognition. Variations in sensor type, scanning resolution, noise levels, and occlusions can affect the detection of smaller elements or partially obscured components. In this case study, for example, purlins were less reliably segmented and recognized due to limited point density and partial occlusions. Future studies will investigate the sensitivity of the method to these factors, assessing how changes in point cloud density, noise, or coverage impact the precision of both segmentation and 3D reconstruction.

The methodology can also be extended to include additional structural elements, such as crane beams, rod bracings or girts, as well as and non-structural elements like walls, improving the comprehensiveness of reconstructed digital models. Future work will explore incorporating a profile library within the optimization framework to enable adaptive fitting of alternative cross-sections, such as rectangular hollow sections (RHS), circular hollow sections (CHS), and channels (UPE), extending the method’s applicability to a broader range of structural systems and materials used in industrial buildings.

Building on these methodological extensions, automating IFC model generation and integrating the results with BIM platforms would facilitate interoperability and streamline workflows for structural documentation. Similarly, improving the robustness of the method for complex layouts and partially occluded point clouds would broaden its applicability across diverse real-world industrial scenarios.

Acknowledgments

This work was partially supported by project ED431F 2022/08 funded by Xunta de Galicia, Spain- GAIN and by the projects PCI2022-132943 and CNS2022-135730, funded by MCIN/AEI /10.13039/501100011033 and by European Union NextGenerationEU/PRTR. L. Díaz-Vilariño acknowledges the grant RYC2020-029193-I funded by MCIN/AEI/10.13039/ 501100011033 and FSE ’El FSE invierte en tu futuro’ and P. González-Cabaleiro acknowledges the Grant PRE2024-PRAZA 10 for the training of predoctoral researchers, associated with the project PID2024-157803OB-I00 funded by MICIU/AEI/10.13039/501100011033. The statements made herein are solely the responsibility of the authors.

References

Agapaki, E., Brilakis, I., 2020. CLOI-NET: Class segmentation of industrial facilities point cloud datasets. *Advanced Engineering Informatics*, 45, 101121.

Agapaki, E., Brilakis, I., 2022. Geometric digital twinning of industrial facilities: Retrieval of industrial shapes. arXiv preprint arXiv:2202.04834.

Agapaki, E., Miatt, G., Brilakis, I., 2018. Prioritizing object types for modelling existing industrial facilities. *Automation in Construction*, 96, 211-223.

Anil, E., Sunnam, R., Akinci, B., 2012. Challenges of identifying steel sections for the generation of as-is BIMs from laser scan data. *Gerontechnology*, 11.

Azhar, S., 2011. Building Information Modeling (BIM): Trends, Benefits, Risks, and Challenges for the AEC Industry. *Leadership and Management in Engineering*, 11(3), 241-252.

Bosché, F., Ahmed, M., Turkan, Y., Haas, C. T., Haas, R., 2015. The value of integrating Scan-to-BIM and Scan-vs-BIM techniques for construction monitoring using laser scanning and BIM: The case of cylindrical MEP components. *Automation in Construction*, 49, 201-213.

Charles, R. Q., Su, H., Kaichun, M., Guibas, L. J., 2017. PointNet: Deep learning on point sets for 3D classification and segmentation. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 77–85.

Chen, S.-H., Xue, F., 2023. Automatic BIM detailing using deep features of 3D views. *Automation in Construction*, 148, 104780.

Choy, C., Gwak, J., Savarese, S., 2019. 4D spatio-temporal convnets: Minkowski convolutional neural networks. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3070–3079.

Cucuzza, R., Aloisio, A., Rad, M. M., Domaneschi, M., 2024. Constructability-based design approach for steel structures: From truss beams to real-world inspired industrial buildings. *Automation in Construction*, 166, 105630.

Mirzaei, K., Arashpour, M., Asadi, E., Masoumi, H., Mahdiyar, A., Gonzalez, V., 2023. End-to-end point cloud-based segmentation of building members for automating dimensional quality control. *Advanced Engineering Informatics*, 55, 101878.

Noichl, F., Borrmann, A., Pan, Y., 2025. Automated Steel Structure Model Reconstruction through Point Cloud Instance Segmentation and Parametric Shape Fitting. *Journal of Information Technology in Construction*, 30, 1099-1122.

Ochmann, S., Vock, R., Wessel, R., Klein, R., 2016. Automatic reconstruction of parametric building models from indoor point clouds. *Computers & Graphics*, 54, 94-103.

Qi, C. R., Yi, L., Su, H., Guibas, L. J., 2017. PointNet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in Neural Information Processing Systems (NeurIPS)*, 5099–5108.

Rausch, C., Haas, C., 2021. Automated shape and pose updating of building information model elements from 3D point clouds. *Automation in Construction*, 124, 103561.

Smith, A. G., 2020. Extraction of structural component geometries in point clouds of metal buildings. Master’s thesis, Virginia Polytechnic Institute and State University, Blacksburg, Virginia.

Son, H., Kim, C., Kim, C., 2013. Knowledge-based approach for 3D reconstruction of as-built industrial plant models from laser-scan data. F. Hassani, O. Moselhi, C. Haas (eds), *Proceedings of the 30th International Symposium on Automation and Robotics in Construction and Mining (ISARC 2013): Building the Future in Automation and Robotics*, International Association for Automation and Robotics in Construction (IAARC), Montreal, Canada, 885–893.

Thomas, H., Qi, C. R., Deschaud, J.-E., Marcotegui, B., Goulette, F., Guibas, L., 2019. KPConv: Flexible and deformable convolution for point clouds. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 6410–6419.

Toth, C., Petrie, G., 2018. Chapter 2: Terrestrial laser scanners. *Topographic Laser Ranging and Scanning*, CRC Press, Boca Raton, FL, 21–50.

Wang, Q., Kim, M.-K., 2019. Applications of 3D point cloud data in the construction industry: A fifteen-year review from 2004 to 2018. *Advanced Engineering Informatics*, 39, 306-319.

Yang, L., Cheng, J. C., Wang, Q., 2020. Semi-automated generation of parametric BIM for steel structures based on terrestrial laser scanning data. *Automation in Construction*, 112, 103037.

Zhang, Y., Nie, G., Wang, D., 2025. Structural health monitoring based on three-dimensional point cloud technology: A systematic review. *Results in Engineering*, 27, 106552.