ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XII-4/W1-2026
21st International 3D GeoInfo Conference 2026, 28–29 September 2026, Sofia, Bulgaria

Wireframe Extraction of Urban Linear Objects from Aerial Lidar Point Clouds

Haohua Gan[1], Hugo Ledoux[1], Weixiao Gao[1,2]*

[1] Delft University of Technology, The Netherlands – ganhhua01200@outlook.com, h.ledoux@tudelft.nl
[2] Wageningen University & Research, The Netherlands – gaoweixiaocuhk@gmail.com

Keywords: Point Cloud Processing, Wireframe Extraction, RANSAC, Region Growing, Hough Transform

Abstract

Wireframe extraction of urban linear objects from aerial lidar point clouds is important for 3D city modelling, infrastructure inspection, and asset management. Although power-line and pylon reconstruction has been widely studied, reliable type-agnostic wireframe extraction remains difficult because aerial point clouds are sparse, incomplete, and structurally complex. To address this gap, we present an exploratory comparison of four representative strategies—3D RANSAC, 3D–2D RANSAC, Region Growing, and Hough Transform—and propose a pairwise Markov Random Field (MRF) formulation optimised by graph cuts. The methods are evaluated on Dutch aerial lidar data using manually delineated reference wireframes. Results reveal clear trade-offs among the baselines. On the selected difficult pylon case, the proposed method achieves the lowest angular RMSE and qualitatively retains some internal members, but its high unmatched rate indicates that many false-positive edges remain. We identify the remaining challenges of wireframe extraction from sparse point clouds and discuss directions for more robust hybrid solutions. The implementation is publicly available at https://github.com/Ganbusier/final_thesis.

1. Introduction

Three-dimensional city models have become essential for applications such as environmental simulation, infrastructure management, and urban decision-making (Biljecki et al., 2015). While substantial progress has been made in reconstructing dominant urban elements such as buildings and vegetation from point clouds (Huang et al., 2022; Lafarge and Mallet, 2012; Du et al., 2019; Nan and Wonka, 2017), power-line and pylon reconstruction has also received considerable attention. However, most existing methods are tailored to a specific object category, sensing configuration, or parametric model, whereas general-purpose wireframe extraction across power lines and structurally diverse pylons remains less studied. These structures are critical components of urban infrastructure and require accurate geometric representation for inspection, maintenance, and digital city modelling.

Recent advances in aerial lidar have enabled large-scale acquisition of urban point clouds, creating new opportunities for automated reconstruction (see Fig. 1). For urban linear objects, a key intermediate representation is the wireframe model, which describes the structure using 3D line segments and can support subsequent (volumetric) geometric reconstruction. However, extracting reliable wireframes from aerial point clouds remains challenging because the data are often sparse, incomplete, and noisy. As further explaiend in Section 2, existing approaches typically rely on strong prior assumptions about object type, predefined parametric models, or pipelines tailored to specific structures. Although such methods can achieve good performance in constrained cases, they often lack generality and are difficult to transfer across different categories or structural styles of urban linear objects.

To address this gap, this paper investigates wireframe extraction from a general perspective, aiming for methods that are type-agnostic, data-driven, and holistic rather than object-specific. We comparatively evaluate four representative extraction strategies, namely 3D RANSAC, 3D–2D RANSAC, Region Growing, and Hough Transform, and further propose an energy-minimization framework based on a Markov Random Field. The core idea is to infer a structurally consistent wireframe by jointly considering local geometric evidence and global structural coherence.

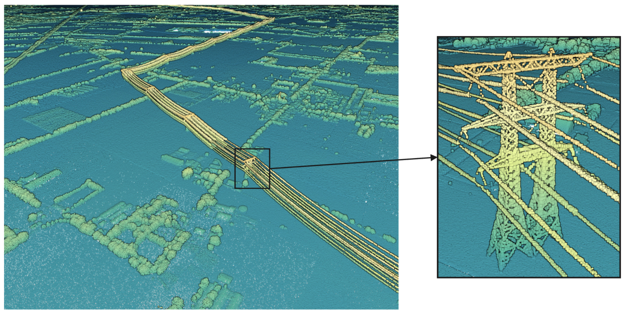

Figure 1. Large-scale aerial lidar point cloud containing man-made urban linear objects (power lines and pylons).

The main contributions of this paper are threefold: (1) we cast wireframe extraction of man-made urban linear objects as a general, type-agnostic research problem; (2) we curate a manually annotated reference dataset and systematically compare representative general-purpose extraction methods on selected aerial lidar cases of power lines and pylons; and (3) we propose and analyse a pairwise-MRF formulation for globally selecting candidate line elements from a dense graph, and identify both its potential and its current failure modes.

2. Related work

Existing research on geometric primitive extraction from point clouds can be broadly grouped into two categories: primitive-detection methods and energy-minimization methods. Primitive-detection methods aim to fit geometric elements

* corresponding author

such as lines, cylinders, or planes directly from point clouds, while energy-minimization methods seek globally consistent solutions by optimizing geometric and structural constraints.

Representative primitive-detection approaches include RANSAC-based fitting, region growing, and Hough Transform. RANSAC and its variants have been widely used to fit cylinders or curve-based models for vegetation, industrial facilities, and power lines (Du et al., 2019; Ghahremani et al., 2021; Jin and Lee, 2019; Guo et al., 2016b). Their core idea is to iteratively generate model hypotheses and identify inliers according to geometric residuals. These methods are effective when the target shape is simple and the data quality is sufficient, but they often rely on accurate initialization, strong geometric assumptions, or object-specific parametric models. Region-growing approaches exploit local continuity, neighbourhood connectivity, and normal consistency to aggregate points into structural components (Qiao et al., 2024). They can better preserve local completeness, but their performance is sensitive to point density, noise, and normal estimation (often difficult to evaluate for a pylon or a power line). Hough-Transform-based methods detect lines through voting in a discretized parameter space (Dalitz et al., 2017; Yan et al., 2019). They do not necessarily require normals and can robustly detect dominant linear structures, but they become computationally expensive and less effective for complex scenes with overlapping or intersecting elements.

Energy-minimization methods provide an alternative by introducing global regularization. For example, Wang et al. (2020) combine RANSAC hypotheses with data, smoothness, and label terms to jointly assign points to primitives, while Guo et al. (2016a) reconstruct pylons by optimizing fidelity and structural terms over a predefined model library. These methods better enforce global consistency, but they are often task-specific and frequently depend on predefined primitive classes or object templates, which limits general applicability.

Taken together, studies on vegetation, industrial facilities, power lines, and pylons reveal a recurring trade-off: RANSAC-based methods depend on model support, region growing depends on reliable local neighbourhoods, and Hough voting becomes costly or ambiguous for complex structures. No generic primitive detector therefore guarantees both high coverage and correct orientation under sparse and noisy sampling. In contrast, we evaluate representative primitive-detection methods within a unified setting and further propose an energy-minimization framework for wireframe extraction. Our method reduces dependence on object-specific assumptions, while its limitations are evaluated explicitly.

3. Method

Our goal is to extract wireframe models, i.e., sets of 3D line segments, from isolated aerial lidar point clouds of man-made urban linear objects. The study contains two complementary components. First, we evaluate whether representative general-purpose primitive-detection methods can recover line structures without object-specific templates. Second, we propose an energy-minimization framework that selects consistent line hypotheses from an over-complete graph. The baseline outputs are not used as input to the proposed method; they provide a common comparison and reveal the failure modes that motivate the graph-based formulation.

3.1 Primitive detection baselines

The purpose of the primitive-detection comparison is to test how far existing generic algorithms can go when applied to aerial point clouds of power lines and pylons. For all baseline methods, the input is an isolated point cloud of one target object, and the output is a set of 3D line segments. These algorithms are treated as baselines rather than new contributions; the modifications below only adapt their outputs and parameters to the present task. They represent three complementary ways of extracting structure: hypothesis fitting, local aggregation, and global voting.

RANSAC-based extraction. RANSAC is included because it is one of the most widely used frameworks for primitive fitting in point clouds. In the original 3D version (Schnabel et al., 2007), the input points are fitted with cylinder primitives, and the axis of each detected cylinder is converted into a 3D line segment. The motivation is that many members of power lines and pylons are elongated and can be locally approximated by cylinders. However, this formulation depends strongly on point normals and on stable radius estimation, both of which become unreliable in sparse aerial lidar. To make the method more suitable for our data, we additionally constrain the minimum and maximum cylinder radius to suppress unrealistic fits.

To further reduce the dependence on normals, we also implement a 3D–2D RANSAC variant. Its motivation is that many urban linear objects exhibit partial planarity: a power line span can often be approximated by a dominant plane, and many pylon members are organized on one or several structural planes. The input remains a 3D point cloud, but the intermediate representation changes. First, 3D RANSAC is used to detect planes; next, the inlier points of each plane are projected to 2D; then, 2D RANSAC is applied to extract line segments; finally, the detected 2D lines are reprojected back to 3D. This design replaces direct fitting of curved surface primitives with a simpler plane-plus-line decomposition, which is better suited to sparse data. In practice, it can recover clearer structural segments than cylinder fitting, but it also inherits a limitation: members that are weakly planar, intersect across planes, or appear only as sparse internal supports may still be missed.

Region growing. Region growing is used to test whether local continuity can compensate for sparse sampling. Its input is again the object point cloud, and its output is a set of cylindrical regions whose axes are interpreted as 3D line segments. The key idea is to start from locally simple seed points and iteratively expand regions by enforcing neighborhood consistency. Because aerial point clouds are sparse, we use a k-nearest-neighbor query instead of a fixed-radius search to ensure that candidate neighbors can still be found in low-density areas. Each region is constrained to remain compatible with a cylinder model, using distance, normal deviation, and minimum region size as acceptance criteria. For our data, this method is useful because it tends to preserve connected linear parts and may overlook fewer points than pure hypothesis fitting. However, its quality depends on the reliability of local normals and neighborhood topology. In cluttered or highly irregular parts of pylons, the grown regions may become fragmented or spatially scattered, which weakens the final wireframe interpretation.

Hough Transform. The Hough Transform is included because it detects dominant lines through global voting rather than local fitting or region expansion. This makes it attractive for line-dominated objects when normals are unreliable. The input is the

3D point cloud, and the output is a set of detected 3D lines represented in Roberts' parameterization (Dalitz et al., 2017). The method discretizes the line parameter space, lets points vote for candidate lines, and iteratively extracts the highest-scoring ones. For our research objects, the method is particularly suitable for prominent, isolated members, because a line can still accumulate votes even when local neighborhoods are weak. At the same time, this advantage comes with a trade-off: when the structure becomes complex, with many intersecting and overlapping elements, the voting space becomes more ambiguous and computationally expensive, which often reduces both efficiency and precision.

Taken together, these baseline methods reveal complementary strengths and weaknesses. RANSAC offers explicit geometric fitting but is sensitive to sparse support and model assumptions; region growing improves continuity but depends on local consistency; Hough voting can recover salient global lines but struggles with complexity. These observations motivate our separate graph-based method, which combines dense hypothesis generation with global structural selection.

3.2 MRF optimization for wireframe extraction

The proposed method addresses a different question from the baseline algorithms: instead of directly fitting final line segments, can we first generate a dense set of plausible structural hypotheses and then select a consistent subset through global optimization? This design is motivated by two recurrent problems in sparse aerial point clouds of urban linear objects. First, primitive fitting often produces discontinuous segments because neighboring members are detected independently. Second, it may overlook valid structural elements when the local point support is weak. Our method targets both issues by shifting the problem from primitive detection to graph labeling.

Input and output. The input is an isolated point cloud P of one man-made urban linear object. The output is a subset of 3D line segments that forms the extracted wireframe model. Rather than using fitted primitives as initial hypotheses, we build a dense graph directly from the point cloud so that enough candidate line elements are available for optimization.

Graph construction. We first construct an undirected graph $G = \langle V, E \rangle$ from the input point cloud, where $V = \{v_1, \ldots, v_m\}$ are the points and $E = \{e_1, \ldots, e_n\}$ are edges connecting neighboring point pairs. A k-nearest-neighbor graph is used to capture local adjacency, while Delaunay triangulation is additionally introduced to connect gaps that are not bridged by KNN alone. The purpose of this combined graph is to avoid under-connecting sparse structures. We then derive the dual graph $G^* = \langle V^*, E^* \rangle$, where each node in V^* corresponds to an edge in G. In other words, the basic optimization unit is no longer a point but a candidate line element. Wireframe extraction is therefore formulated as a binary labeling problem on G^*: each candidate edge is labeled as either *preserve* or *remove*. Graph cut (Boykov et al., 2001; Kolmogorov and Zabin, 2004) is used to find the final labeling.

MRF formulation and energy design. Each node $x_i \in V^*$ is associated with a binary random variable $y_i \in \{0, 1\}$, where 0 denotes *remove* and 1 denotes *preserve*. Its neighbourhood $\mathcal{N}_i$ is given by adjacent nodes in the dual graph. We assume the pairwise Markov property

$$\Pr(y_i \mid \mathbf{y}_{V^* \setminus i}, \mathcal{P}) = \Pr(y_i \mid \mathbf{y}_{\mathcal{N}_i}, \mathcal{P}), \tag{1}$$

so the posterior distribution can be written in Gibbs form as

$$\Pr(\mathbf{y} \mid \mathcal{P}) = \frac{1}{Z} \exp[-E(\mathbf{y})]. \tag{2}$$

The maximum-a-posteriori labeling (MAP) is therefore obtained by minimizing

$$E(\mathbf{y}) = \sum_i D_i(y_i) + \lambda \sum_{(i,j) \in E^*} w_{ij}\, \mathbf{1}[y_i \neq y_j]. \tag{3}$$

This explicitly defines the MRF: the unary term measures local geometric evidence for each candidate edge, while the pairwise Potts term encourages compatible neighbouring edges to share a label.

For candidate edge x_i, we first compute a score

$$s_i = I_i\, L_i, \tag{4}$$

where I_i is an inlier-support cost and L_i is an edge-length cost. Let $\mathcal{P}_i$ be the points inside the cylinder centred on x_i (Figure 2). The inlier-support cost is

$$I_i = 1 - \frac{\sum_{p_j \in P_i} \rho_{ij}^2}{\sum_{p_j \in P_i} \rho_{ij}}, \tag{5}$$

$$\rho_{ij} = \exp\left(-\frac{d_{ij}^2}{2\sigma^2}\right), \tag{6}$$

where d_{ij} is the perpendicular distance from p_j to x_i. Points closer to the edge receive higher weights; consequently, strong point support produces a lower cost.

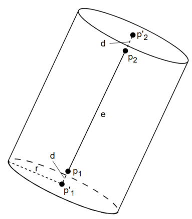

Figure 2. Cylinder constructed around an edge, where d denotes the extension distance applied to both endpoints.

The edge-length cost is

$$L_i = \frac{L_{e_i}}{L_{\mathrm{tol},i}}, \tag{7}$$

where L_{e_i} is the current edge length and $L_{\mathrm{tol},i}$ is accumulated by propagating through the most direction-consistent adjacent edges. Edges embedded in long, nearly collinear chains therefore receive lower costs than isolated short connections. The implementation maps the score to the two unary costs as

$$D_i(y_i) = \begin{cases} 1 - s_i, & y_i = 0 \ (\textit{remove}), \\ s_i, & y_i = 1 \ (\textit{preserve}). \end{cases} \tag{8}$$

For neighbouring candidate edges, the pairwise weight is

$$w_{ij} = \cos^{10}(\alpha_{ij}), \tag{9}$$

where $\alpha_{ij} \in [0, \pi/2]$ is the orientation-invariant angular difference between two neighbouring line elements. The exponent

https://doi.org/10.5194/isprs-annals-XII-4-W1-2026-145-2026

10 is an empirical shape parameter that makes the weight decay rapidly with angular disagreement, so smoothing mainly acts on nearly collinear edges; it is not a mathematically fixed value. Because $w_{ij} \geq 0$ and the label penalty in Eq. (3) is a binary Potts term, the resulting submodular energy is optimised using graph cuts (Boykov et al., 2001; Kolmogorov and Zabin, 2004).

Optimization outcome. After optimization, all candidate edges are partitioned into preserved and removed groups, and the preserved subset forms the extracted wireframe. Compared with the baselines, the framework does not require predefined object templates or explicit part-wise decomposition. It instead uses local evidence and neighbourhood coherence to select from a dense structural graph; whether this improves both completeness and precision is evaluated experimentally.

4. Experiments

4.1 Data, implementation, and evaluation protocol

The experiments were conducted on a manually curated dataset of Dutch power lines and pylons derived from AHN5 aerial lidar data (AHN, 2020). Points with classification value 14 (high-tension objects) were extracted from the original LAZ files and manually isolated in Mapple (Nan, 2021). Manual annotation refers to interactively drawing polylines along point-supported structural members in Rhino 7 and splitting them at junctions into individual 3D line segments; no automatic fitting was used. We therefore refer to these annotations as reference wireframes rather than absolute ground truth. Three cases were selected before the comparative analysis to represent increasing structural complexity:

1. one power line with little directional ambiguity (simple),
2. four neighbouring power lines with repeated parallel structures (medium),
3. a pylon with many intersecting directions and sparsely sampled internal members (difficult).

All methods were implemented mainly in C++ using CGAL, Eigen, Easy3D, and GCO; Hough Transform was adapted from Dalitz et al. (2017). Figure 3 shows the full curated dataset, while quantitative conclusions in this paper are restricted to the three selected cases.

For 3D RANSAC, 3D–2D RANSAC, and Region Growing, parameters were selected using experience-guided grid search. The selection criterion balances coverage and model compactness:

$$\text{score} = \alpha \times L - \beta \times N, \tag{10}$$

where L is the number of leftover points, N is the number of extracted primitives, and $\alpha = 1.0$, $\beta = 0.1$. In practice, 3D RANSAC favoured a small minimum support (5–10), $\epsilon = 0.05$, and large clustering radii; 3D–2D RANSAC used normal-free plane detection followed by 2D line fitting with minimum support 4 and split thresholds of 1.0–1.5; Region Growing worked best with $k = 16$–28, max distance 0.04–0.05, and max angle 20–30°. Hough Transform and the proposed energy-minimization method were tuned qualitatively.

Evaluation was based on 3D line-segment matching against the reference wireframes. For the primitive-detection baselines, we

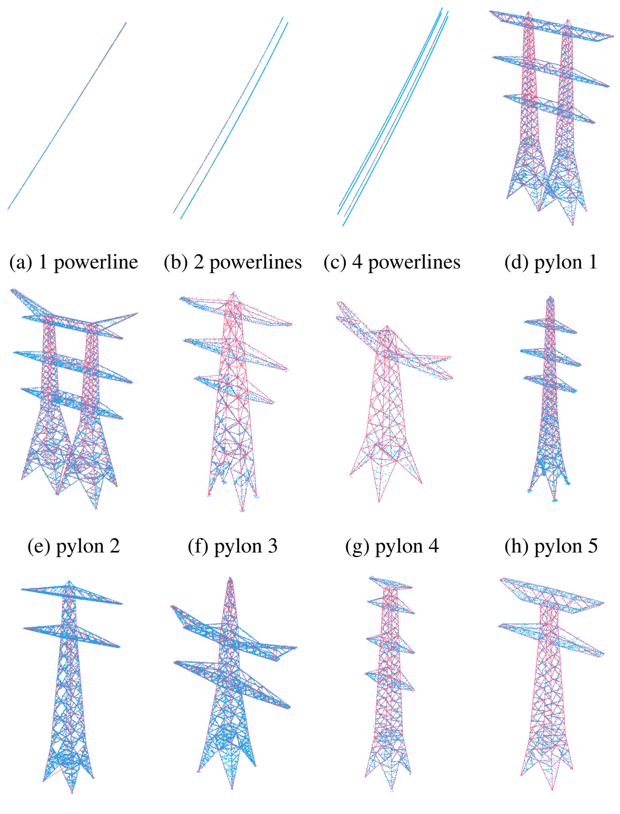

(a) 1 powerline (b) 2 powerlines (c) 4 powerlines (d) pylon 1
(e) pylon 2 (f) pylon 3 (g) pylon 4 (h) pylon 5
(i) pylon 6 (j) pylon 7 (k) pylon 8 (l) pylon 9

Figure 3. Visualization of the curated dataset, showing raw points (blue) and manually annotated reference wireframes (red).

report the leftover-point rate (LoR.); for the proposed graph method, we report the removed-edge rate (RemR.), defined as the proportion of input candidate edges labelled as removed. For all methods, we report the unmatched rate,

$$\text{Unmatched rate} = \frac{M_{\text{unmatched}}}{M_{\text{total}}}, \tag{11}$$

and RMSE of distance and angle deviations,

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_i)^2}. \tag{12}$$

LoR. and RemR. are method-specific diagnostics and are not directly comparable; UnmR. and the RMSE measures are used for cross-method comparison.

4.2 Results and discussion

3D RANSAC. For all test cases, 3D RANSAC was applied with cylinder primitives and fixed radius bounds of 0.01–1.0 m. As summarized in Table 1, the method performs reasonably on the simple and medium cases, with low distance RMSEs of 0.060 m and 0.053 m and small angle RMSEs of 1.313° and 1.169°, respectively. However, its performance degrades markedly on the difficult pylon case, where the leftover rate reaches 42.2%, the unmatched rate 36.1%, and the angle RMSE 62.551°. These results reveal both substantial overlooking and unstable fitting. A main cause of the wrong-fitting problem is the strong dependence of cylinder fitting on point normals. Since no ground-truth normals are available, normals are estimated using KNN, but in

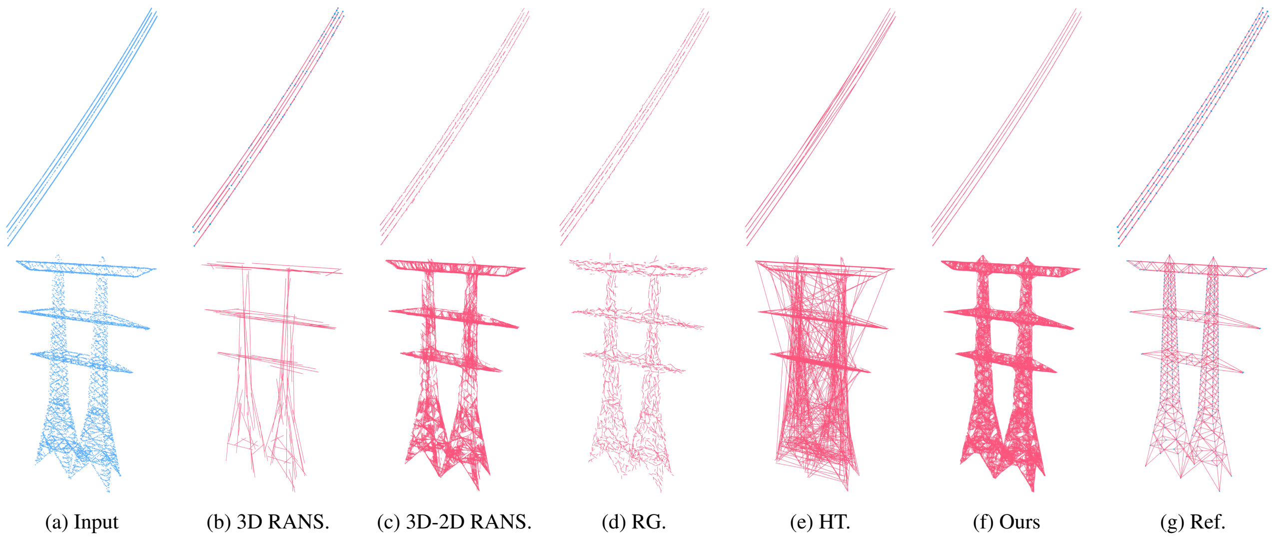

Figure 4. Qualitative comparison of wireframe extraction results on the medium (*top*: four power lines) and difficult (*bottom*: pylon) test cases. From left to right: input point cloud, 3D RANSAC (3D RANS.), 3D–2D RANSAC (3D–2D RANS.), Region Growing (RG.), Hough Transform (HT.), the proposed method (Ours), and the manually annotated reference wireframe (Ref.).

sparse aerial lidar their quality varies with local point density and is often noisy, as illustrated in Figure 5. This makes the fitted cylinder axes unstable and can lead to large directional errors, especially in complex scenes. Qualitative examples in Figure 4(b) further show that RANSAC also favors densely sampled outer members over sparsely sampled internal structures, which aggravates the overlooking problem.

Data	LoR. (%)	UnmR. (%)	D-RMSE (m)	A-RMSE (°)
3a	14.7	18.2	0.060	1.313
3c	9.7	25.8	0.053	1.169
3d	42.2	36.1	0.042	62.551

Table 1. Summary of 3D RANSAC results on the three representative test cases (LoR.: leftover rate; UnmR.: unmatched rate; D-RMSE: distance RMSE; A-RMSE: angle RMSE).

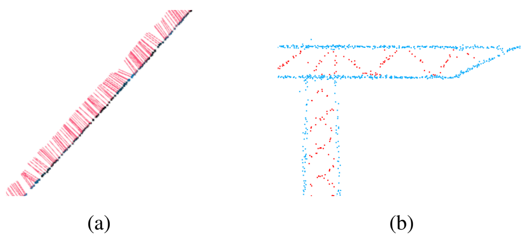

Figure 5. (a) Noisy KNN normal estimation ($k = 16$); (b) imbalance between internal and external pylon structures.

3D–2D RANSAC. The 3D–2D RANSAC variant reduces the reliance on point normals by replacing direct 3D cylinder fitting with plane detection followed by 2D line extraction. As summarized in Table 2, it achieves the best overall performance on the simple case, with only 2.1% unmatched segments, a distance RMSE of 0.040 m, and an angle RMSE of 0.921°. More importantly, for the difficult pylon case, it markedly improves coverage compared with 3D RANSAC, reducing the leftover rate from 42.2% to 13.4% and the unmatched rate from 36.1% to 22.5%, while also lowering the angle RMSE from 62.551° to 55.628°. However, the medium case remains unstable, with a high unmatched rate of 47.5%. Qualitative examples in Figure 4(c) suggest that although the 2D projection increases detectability, projection ambiguity and reprojection errors still lead to misaligned segments.

Data	LoR. (%)	UnmR. (%)	D-RMSE (m)	A-RMSE (°)
3a	24.9	2.1	0.040	0.921
3c	20.4	47.5	0.065	1.571
3d	13.4	22.5	0.044	55.628

Table 2. Summary of 3D–2D RANSAC results on the three representative test cases. LoR., UnmR., D-RMSE, and A-RMSE denote leftover rate, unmatched rate, distance RMSE, and angle RMSE, respectively.

Region Growing. Region Growing achieves the strongest coverage among the primitive-detection baselines, with consistently low leftover rates of 1.1%, 3.2%, and 2.0% for the simple, medium, and difficult cases, respectively (Table 3). However, this advantage is offset by poor structural fidelity. The unmatched rate increases from 34.6% to 70.5%, and the angle RMSE rises sharply from 2.822° to 57.304°, indicating severe directional instability, especially for the pylon. By contrast, the distance RMSE remains relatively stable (0.055–0.060 m), suggesting that many fitted segments stay spatially close to the target but are oriented incorrectly. Qualitative comparisons in Figure 4(d) and the failure cases in Figure 6 show that noisy normal estimation and relaxed region validity criteria tend to generate scattered or wrongly oriented cylinders. Region Growing therefore reveals a clear trade-off between low overlooking and reliable wireframe extraction.

Data	LoR. (%)	UnmR. (%)	D-RMSE (m)	A-RMSE (°)
3a	1.1	34.6	0.055	2.822
3c	3.2	50.0	0.060	8.835
3d	2.0	70.5	0.055	57.304

Table 3. Summary of Region Growing results on the three representative test cases. LoR., UnmR., D-RMSE, and A-RMSE denote leftover rate, unmatched rate, distance RMSE, and angle RMSE, respectively.

Hough Transform. Hough Transform performs best on the simple case, yielding a low leftover rate of 3.0%, an unmatched

https://doi.org/10.5194/isprs-annals-XII-4-W1-2026-145-2026 | © Author(s) 2026. CC BY 4.0 License.

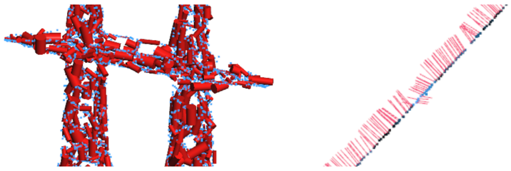

(a) (b)

Figure 6. Typical Region Growing failure cases: wrong cylinder directions and inconsistent grouping caused by noisy normals.

rate of 7.7%, a distance RMSE of 0.056 m, and an angle RMSE of 2.055° (Table 4). For the medium case, the leftover rate remains low (1.0%), but the unmatched rate rises to 39.7%, indicating that many detected segments are not well aligned with the reference wireframe. On the difficult pylon case, the method achieves the lowest distance RMSE among all baselines (0.036 m) and a relatively low unmatched rate (9.9%), but its angle RMSE increases sharply to 61.223°, revealing substantial directional errors. As shown in Figure 4(e) and the failure cases in Figure 7, the method benefits from not requiring normals, yet discretization in Hough space and vote-dominated selection still lead to unstable orientation estimates in complex scenes.

Data	LoR. (%)	UnmR. (%)	D-RMSE (m)	A-RMSE (°)
3a	3.0	7.7	0.056	2.055
3c	1.0	39.7	0.057	1.106
3d	23.2	9.9	0.036	61.223

Table 4. Summary of Hough Transform results on the three representative test cases. LoR., UnmR., D-RMSE, and A-RMSE denote leftover rate, unmatched rate, distance RMSE, and angle RMSE, respectively.

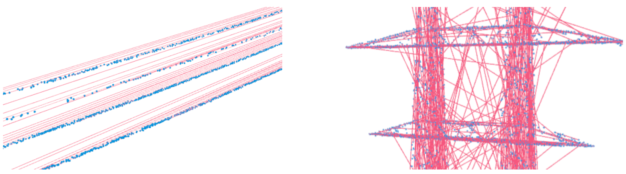

(a) 3c (b) 3d

Figure 7. Typical wrong-fitting results of Hough Transform.

Energy minimization. The proposed energy-minimization framework was evaluated using the formulation introduced in Section 3.2. The inlier-probability term was parameterized by σ^2 derived from Eq. (6), the initial angular threshold for the edge-length term was set to 10°, and the smoothness weight was initialized as $\lambda = 0.1$. Unlike the primitive-detection baselines, this method does not directly fit a small set of final primitives. Instead, it starts from a dense graph of candidate edges and selects a structurally consistent subset through global optimization. This design is intended to alleviate the overlooking problem observed in the previous methods, especially for complex objects with sparsely sampled internal members.

As summarized in Table 5, the proposed method has distance RMSEs between 0.050 and 0.063 m. On the difficult pylon case, it attains the lowest angle RMSE among the evaluated methods (50.313°), compared with 62.551° for 3D RANSAC, 55.628° for 3D–2D RANSAC, 57.304° for Region Growing, and 61.223° for Hough Transform. However, its unmatched rate is 42.0%, substantially higher than the 9.9% of Hough Transform. The result therefore indicates lower angular error among matched structures, not overall superiority.

Data	RemR. (%)	UnmR. (%)	D-RMSE (m)	A-RMSE (°)
3a	3.0	14.1	0.055	9.273
3c	7.7	50.5	0.063	18.331
3d	25.7	42.0	0.050	50.313

Table 5. Summary of energy minimization results on the three representative test cases. RemR., UnmR., D-RMSE, and A-RMSE denote removed-edge rate, unmatched rate, distance RMSE, and angle RMSE, respectively.

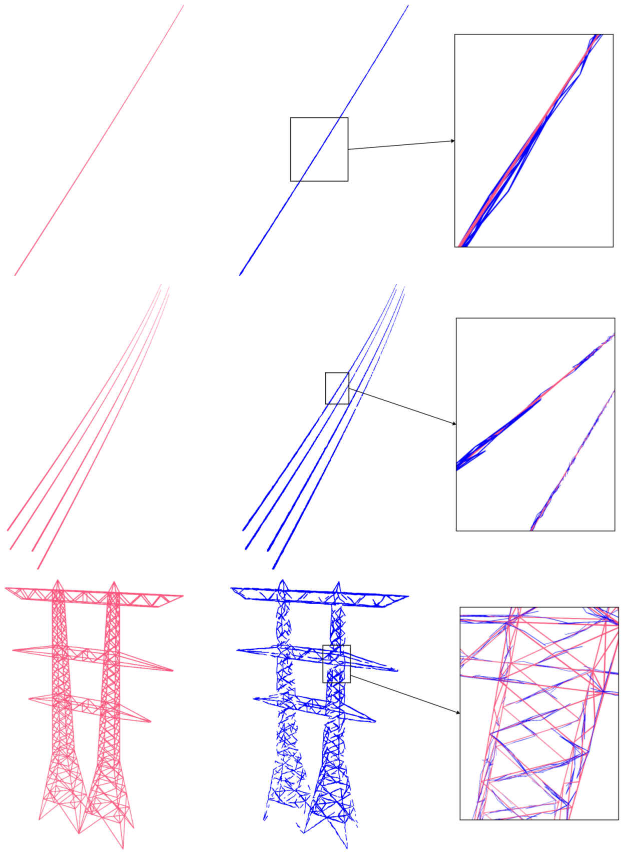

Figure 8. Edge-matching results of the proposed method for the simple power-line (*top*), medium four-power-line (*middle*), and difficult pylon (*bottom*) cases. Reference wireframes are shown in red and preserved candidate edges in blue; the right column shows enlarged local comparisons.

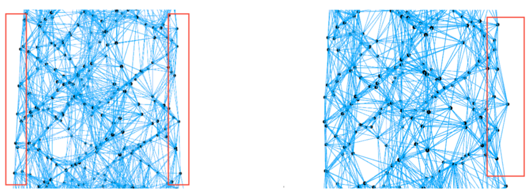

(a) Failure of the inlier-probability term on external corner edges.

(b) Failure of the edge-length term on non-collinear outer members.

Figure 9. Typical failure cases of the data term on external pylon structures.

Figure 8 further illustrates this behaviour. For the simple power-line case, the method preserves the dominant linear structure with an unmatched rate of 14.1%. For the medium case, the main geometry remains recognizable, but many redundant edges produce an unmatched rate of 50.5%. In the difficult pylon case, the qualitative comparison in Figure 4(f) shows

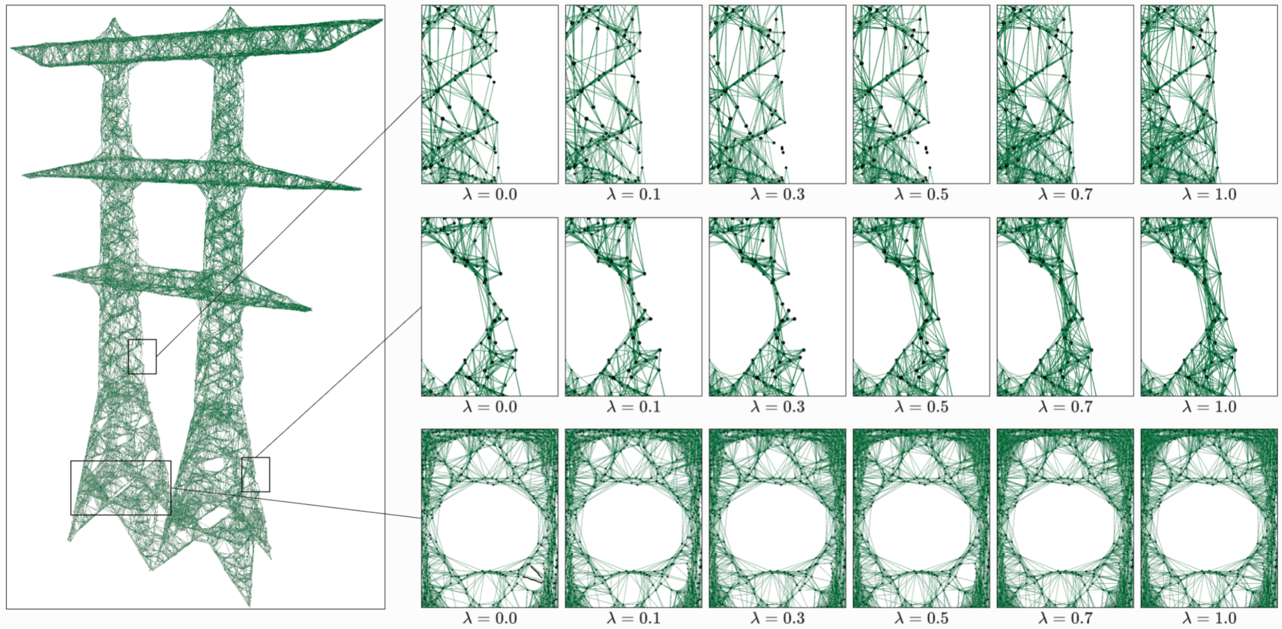

Figure 10. Effect of the smoothness weight λ on the optimization result, illustrating limited smoothing, incorrect smoothing, and spatially varying smoothing behavior.

that some internal members missed by the baselines are retained. This observation suggests that dense graph hypotheses can reduce overlooking, although the current metrics do not yet demonstrate a consistent completeness advantage.

At the same time, the results also reveal the current limitation of the framework, namely insufficient discrimination in the data term. A considerable number of noisy or structurally implausible edges are still retained, especially in the medium and difficult cases, resulting in unmatched rates of 50.5% and 42.0%, respectively. This indicates that the optimization is effective at preserving coverage, but not yet selective enough to suppress false positives.

A closer inspection shows that the main limitation of the proposed framework lies in the current design of the data term. As introduced in Section 3.2, the data term combines an inlier-probability term and an edge-length term. For outer pylon members, however, both terms can be unreliable. First, the inlier-probability term searches supporting points within a cylinder constructed around each edge. For corner-like external members, this cylinder may contain very few or even no nearby points, leading to underestimated support (Figure 9). Second, the edge-length term favors edges that can be propagated through nearly collinear neighbors. External members often meet adjacent edges at larger angles, so their accumulated length remains small and the corresponding penalty becomes overly large. As a result, structurally valid outer edges may be removed during optimization.

This limitation also explains the inconsistent behavior of the smoothness term. In principle, the smoothness term should suppress local noise by encouraging neighboring edges with similar directions to share the same label. In practice, its effect strongly depends on the correctness of the data term. Figure 10 shows three typical cases during λ tuning: little visible change when most edges are already labeled as preserved, incorrect smoothing that removes valid structural edges, and different smoothing thresholds across local regions. These patterns indicate that smoothing alone cannot compensate for systematic misclassification in the unary term.

Overall, the proposed method shows potential for retaining sparsely sampled internal members, but the present evidence is qualitative and its false-positive rate remains high. Its performance is constrained by the limited discriminative power of the data term. Future improvements should therefore focus on more reliable edge-support measures and stronger structural priors, enabling the optimization to distinguish true wireframe elements from noisy graph connections more reliably.

5. Conclusion

In this paper, we investigated wireframe extraction from aerial lidar point clouds of power lines and pylons without relying on object-specific templates, strong priors, or part-wise handcrafted pipelines. We evaluated four representative primitive-detection strategies and introduced a pairwise-MRF formulation for selecting candidate line elements from a dense graph. The experiments show that current general-purpose methods can recover useful 3D line segments, but all suffer from overlooking valid members or producing wrongly fitted segments. The proposed method obtains the lowest angular RMSE on the selected difficult pylon case and qualitatively retains some internal members, but its high unmatched rate shows that many false-positive edges remain.

The contributions are threefold. First, we curate a manually annotated reference dataset of urban linear objects. Second, we compare representative general-purpose methods under a common evaluation protocol and identify their recurring failure modes. Third, we formulate candidate-edge selection as a binary pairwise MRF optimised by graph cuts and analyse the limitations of its current unary and smoothness terms.

The study remains exploratory: quantitative evaluation is limited to three representative cases, and no object-specific reconstruction method is included as a comparator. Future work should therefore evaluate the full dataset, include domain-specific baselines where applicable, and adopt metrics that separately quantify reference-line coverage and false-positive length. Methodologically, more discriminative edge-support measures, stronger structural priors, and hybrid pipelines combining primitive detection with graph optimization are promising directions for improving both completeness and precision.

References

AHN, 2020. Actueel Hoogtebestand Nederland: Dataroom page. https://www.ahn.nl/dataroom.

Biljecki, F., Stoter, J., Ledoux, H., Zlatanova, S., Çöltekin, A., 2015. Applications of 3D City Models: State of the Art Review. *ISPRS International Journal of Geo-Information*, 4(4), 2842–2889.

Boykov, Y., Veksler, O., Zabih, R., 2001. Fast Approximate Energy Minimization via Graph Cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11), 1222–1239.

Dalitz, C., Schramke, T., Jeltsch, M., 2017. Iterative Hough Transform for Line Detection in 3D Point Clouds. *Image Processing On Line*, 7, 184–196.

Du, S., Lindenbergh, R., Ledoux, H., Stoter, J., Nan, L., 2019. AdTree: Accurate, Detailed, and Automatic Modelling of Laser-Scanned Trees. *Remote Sensing*, 11(18), 2074.

Ghahremani, M., Williams, K., Corke, F., Tiddeman, B., Liu, Y., Wang, X., Doonan, J. H., 2021. Direct and Accurate Feature Extraction from 3D Point Clouds of Plants Using RANSAC. *Computers and Electronics in Agriculture*, 187, 106240.

Guo, B., Huang, X., Li, Q., Zhang, F., Zhu, J., Wang, C., 2016a. A Stochastic Geometry Method for Pylon Reconstruction from Airborne LiDAR Data. *Remote Sensing*, 8(3), 243.

Guo, B., Li, Q., Huang, X., Wang, C., 2016b. An Improved Method for Power-Line Reconstruction from Point Cloud Data. *Remote Sensing*, 8(1), 36.

Huang, J., Stoter, J., Peters, R., Nan, L., 2022. City3D: Large-Scale Building Reconstruction from Airborne LiDAR Point Clouds. *Remote Sensing*, 14(9), 2254.

Jin, Y.-H., Lee, W.-H., 2019. Fast Cylinder Shape Matching Using Random Sample Consensus in Large Scale Point Cloud. *Applied Sciences*, 9(5), 974.

Kolmogorov, V., Zabin, R., 2004. What Energy Functions Can Be Minimized via Graph Cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2), 147–159.

Lafarge, F., Mallet, C., 2012. Creating Large-Scale City Models from 3D-Point Clouds: A Robust Approach with Hybrid Representation. *International Journal of Computer Vision*, 99(1), 69–85.

Nan, L., 2021. Easy3D: a lightweight, easy-to-use, and efficient C++ library for processing and rendering 3D data. *Journal of Open Source Software*, 6(64), 3255.

Nan, L., Wonka, P., 2017. PolyFit: Polygonal Surface Reconstruction from Point Clouds. *2017 IEEE International Conference on Computer Vision (ICCV)*, IEEE, Venice, 2372–2380.

Qiao, Y., Xi, X., Wang, C., Du, M., Nie, S., Liu, W., Fan, H., 2024. A Framework for Automated Reconstruction of Communication Towers from Terrestrial Laser Scanning Point Clouds. *International Journal of Digital Earth*, 17(1), 2366431.

Schnabel, R., Wahl, R., Klein, R., 2007. Efficient RANSAC for Point-Cloud Shape Detection. *Computer Graphics Forum*, 26(2), 214–226.

Wang, L., Yan, B., Duan, F., Lu, K., 2020. Energy Minimisation-Based Multi-Class Multi-Instance Geometric Primitives Extraction from 3D Point Clouds. *IET Image Processing*, 14(12), 2660–2667.

Yan, B., Xu, N., Zhao, W.-B., Xu, L.-P., 2019. A Three-Dimensional Hough Transform-Based Track-Before-Detect Technique for Detecting Extended Targets in Strong Clutter Backgrounds. *Sensors*, 19(4).