

VF-Plan: Bridging the Art Gallery Problem and Static LiDAR Scanning with Visibility Field Optimization

B. Xiong^{1†}, L. Zhang¹, R. Huang¹, J. Zhou¹, S.R.U.N. Jafri², B. Wu^{3‡}, F. Li⁴

¹ Wuhan University of Technology, ² NED University of Engineering and Technology
³ Zhejiang University, ⁴ The Advanced Laser Technology Laboratory of Anhui Province




Figure 1: Viewpoint optimization for static LiDAR Scanning. (a) The input floor plan; (b) Optimized viewpoints indicated by green dots, with the Visibility Field displayed as the background image; (c) and (d) Simulated LiDAR points from optimal viewpoints, shown in 2D and 3D respectively, with each scan distinguished by a unique color.

Abstract

Viewpoint planning is critical for efficient 3D data acquisition in applications such as 3D reconstruction, building life-cycle management, navigation, and interior decoration. However, existing methods often neglect key optimization objectives specific to static LiDAR systems, resulting in redundant or disconnected viewpoint networks. The viewpoint planning problem (VPP) extends the classical Art Gallery Problem (AGP) by requiring full coverage, strong registrability, and coherent network connectivity under constrained sensor capabilities. To address these challenges, we introduce a novel Visibility Field (VF) that accurately captures the directional and range-dependent visibility properties of static LiDAR scanners. We further observe that visibility information naturally converges onto a 1D skeleton embedded in the 2D space, enabling significant searching space reduction. Leveraging these insights, we develop a greedy optimization algorithm tailored to the VPP, which constructs a minimal yet fully connected Viewpoint Network (VPN) with low redundancy. Experimental evaluations across diverse indoor and outdoor scenarios confirm the scalability and robustness of our method. Compared to expert-designed VPNs and existing state-of-the-art approaches, our algorithm achieves comparable or fewer viewpoints while significantly enhancing connectivity. In particular, it reduces the weighted average path length by approximately 95%, demonstrating substantial improvements in compactness and structural efficiency. Code is available at <https://github.com/xiongbiaostar/VFPlan>.

CCS Concepts

• Theory of computation → Computational geometry;

1. Introduction

The VPP has emerged as a critical research spot due to the growing use of LiDAR and camera sensors in applications such as

3D reconstruction, autonomous navigation, facility assessment, building life-cycle management, and interior decoration applications [SRR03, MLL23]. VPP focuses on determining the optimal placements and configurations of sensors to achieve specific reconstruction goals, such as maximizing observation coverage and minimizing occlusions, while balancing efficiency and data quality. Both model-free and model-based algorithms for viewpoint and

† Corresponding author: b.xiong@whut.edu.cn

‡ Corresponding author: ustcbjwu@gmail.com

path planning have gained significant attention due to their applicability across various industries, including robotics, construction, and surveillance [MHS*23].

Conceptually, VPP is rooted in the AGP, a classic challenge in computational geometry that seeks the minimal number of “guards” required to cover a given area, typically a polygon, with full 360° visibility [LL86]. AGP is often simplified by constraining guard placements to polygon vertices; however, this simplification does not capture the true computational complexity of the problem. In its generalized form, AGP is $\exists\mathbb{R}$ -complete, meaning it is at least as hard as any NP-complete problem, with no known efficient solutions [CdRdS11, AAM21]. This NP-hard nature has inspired substantial research interest, particularly in constrained versions that make idealized assumptions, such as complete visibility and static scenes [CdRdS11, SRR03].

VPP, however, introduces additional complexities beyond those of AGP. Unlike static guards, sensors in VPP must maintain overlapping fields of view for accurate data registration, accommodate limited and directional visibility ranges, and meet high accuracy requirements for 3D scene reconstruction. These real-world demands, especially in cluttered or partially obstructed environments, render VPP a more challenging NP-hard problem, requiring innovative optimization strategies beyond those used for AGP [SRR03, AAM21].

We address the challenges of static LiDAR coverage planning by introducing a novel greedy algorithm based on the concept of a *Visibility Field (VF)*. The VF encapsulates the physical constraints of LiDAR sensors, including limited sensing range and angle-dependent visibility, enabling a substantial reduction in computational complexity. We formally demonstrate that visibility information is inherently concentrated along a one-dimensional skeleton embedded within the two-dimensional space. This observation allows us to restrict the search space from 2D to 1D by targeting critical structural features such as the medial axis and junctions within the VF. Consequently, we construct a minimal, fully connected network of viewpoints that achieves complete scene coverage with reduced redundancy. Our main contributions are summarized as follows:

1. **Unified Solution to VPP and AGP:** Our method addresses the combined challenges of viewpoint planning and the AGP, achieving efficient coverage with minimal viewpoints and robust network connectivity.
2. **Visibility Field for Dimensionality Reduction:** We introduce a continuous VF tailored to static LiDAR systems, reducing the optimization space from 2D to 1D by leveraging structural elements like medial axis and joints.
3. **Greedy Algorithm for Efficient Coverage:** Our greedy algorithm utilizes the VF’s structural insights to construct a minimal, fully connected VPN, ensuring both robust coverage and network connectivity.

2. Related Work

This section reviews key advancements in sensor network planning, covering both static and mobile sensing, with a focus on recent works in LiDAR and computer vision area. These serve as the

foundation for the viewpoint planning method proposed here, addressing gaps in efficient coverage, connectivity, and optimization complexity.

2.1. Station-Based Scanning

Sensor network design shares similarities with the Planning for Scanning (P4S) problem, observed in applications like surveillance cameras [MADR11, ZYCH13, GCW15], directional sensors [FG09], 5G base stations [SSZ*20, WZL*20], stealth game [XTV14], and station-based LiDAR scanning [ABT21]. For LiDAR, the goal is to cover the entire scene with comprehensive wall scanning while ensuring data quality criteria such as completeness and accuracy [LXHL21, ABT21].

Advancements in LiDAR technology have led to P4S methods using prior information like 2D drawings and 3D models [LXHL21]. These methods have been applied to civil infrastructures [XWYZ25], space frame structures [LXHL21], rebar detection [LKL21, LKS*25], and landslide monitoring [ZCH*23]. In cases where prior information is outdated or unavailable, low-resolution preliminary scans are required, either indoors [QWTW21] or outdoors using UAVs [LLZ*22].

P4S optimizations aim to maximize coverage and overlap while minimizing the number of viewpoints [IGFER22, CSW*25]. Methods like CMA Evolution Strategy [RLGA22] and regular grid sampling [JL19, NB22] increase the likelihood of finding optimal viewpoints but also raise computational complexity. Jia’s approach using grid-based sampling [JL18, JL19, JL22] scores candidate viewpoints based on observed wall segments, though this approach is limited by its reliance on reference objects for registration. Noichl [NLB24] improve Jia’s method into 3D scene for an uniform covering by triangulating 3D scene into meshes and checking their visibility. For scenarios lacking BIM priors, structural elements can be surveyed and detected using low-cost sensors to serve as a preliminary prior for static station planning. A precise scan is then performed in a stop-and-go manner using a static scanner mounted on a robotic dog [CKP*25, HG25].

Our approach addresses these limitations by using a continuous VF model, which reduces optimization from 2D to 1D by focusing on structural convergence points like medial axis and joints, thus minimizing computation and enhancing adaptability.

2.2. Mobile-Based Exploration

Coverage Path Planning (CPP) focuses on computing optimal, collision-free paths for robots to fully cover target areas [TMMA21]. Techniques like Rapidly-exploring Random Trees (RRT) and their variants (RRT* [KF11], RRT# [OF15]) are widely used for autonomous exploration in unknown spaces, including UAVs [MHS*23] and wheeled robots [PSC*23]. RRT-based methods maximize information gain while minimizing travel distance [HNI18, XYH*18, ZYX*21].

Methods utilizing tensor fields for guiding robot movements enable efficient path routing in partially reconstructed indoor scenes [XZY*17, XZD*24]. Recent approaches like the Signed




Figure 2: Pipeline of viewpoint network (VPN) optimization using VF. (a) Input floor plan with walls, doors, and windows; (b) Computed VF encoding the valid observed angle at each pixel; (c) Distance field derived from the VF; (d) Extraction of converging lines (medial axis) and converging points (joints) used as candidate viewpoints; (e) Final optimized VPN ensuring full coverage and strong connectivity.

Distance Field (SDF) and Hamilton-Jacobi skeleton further improve path planning in known environments [NLMC23]. These concepts share similarities with our VF, which leverages converging lines for efficient viewpoint planning.

For known environments, algorithms like the receding horizon "next-best-view" approach [BKA*16] and dual-resolution mapping schemes [CZR*23] balance detailed local mapping with efficient global exploration. For applications in autonomous driving and exploration, recent advances such as "LookOut" [CCS*21] propose diverse multi-future prediction models that enhance autonomous navigation by predicting possible trajectories, balancing safety, and efficiency.

3. Problem Definition and Global Optimality

We formalize the VPP for static LiDAR scanning across a collection of two-dimensional polygonal regions $\mathcal{P} = \{P_1, P_2, \dots, P_k\}$, each comprising structural boundaries and possible internal obstacles. The boundaries of these regions are uniformly discretized into a set of equal-length line segments $L = \{l_1, l_2, \dots, l_n\}$, representing the structural elements to be scanned. The objective is to select a set of viewpoints $V = \{v_1, v_2, \dots, v_m\}$ that satisfies two core constraints: *coverage* and *connectivity*.

Coverage: Every segment $l_j \in L$ must be visible from at least one viewpoint $v_i \in V$. We formalize this visibility relationship using a binary coverage table C , where each entry C_{ij} indicates whether v_i observes l_j :

$$C_{ij} = \begin{cases} 1 & \text{if } v_i \text{ sees } l_j, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Connectivity: Let the *overlap ratio* O_{ik} as the fraction of L observed by both viewpoints v_i and v_k . The induced subgraph (S, E_S) is connected under the overlap threshold τ , where

$$E_S = \{(v_i, v_k) \in E \mid v_i, v_k \in S, O_{ik} \geq \tau\}. \quad (2)$$

We aim to find the smallest subset $S \subseteq V$ such that:

$$\min_{S \subseteq V} |S| \quad \text{subject to: (1) and (2)} \quad (3)$$

This NP-hard problem extends the classical set cover problem

by incorporating additional connectivity constraints. We demonstrate that restricting V to the *medial axis* (MA) preserves both complete coverage and network connectivity. As observation data in the VF naturally converge toward the MA and its joint points, placing viewpoints along these skeletal structures ensures comprehensive coverage. This property reduces VPN optimization from a 2D search to a 1D problem constrained to the MA. Consequently, a minimal cover on the MA remains globally optimal. Formal proofs of visibility completeness and global optimality for this reduction are provided in Appendix A.

4. Method

4.1. Overview

Our proposed method, VF-Plan, addresses the NP-hard complexity of the View Planning Problem (VPP) through a greedy strategy centered on the VF. Given a structural 2D floorplan, VF-Plan constructs a VPN that ensures full coverage with a minimal number of viewpoints, while maintaining connectivity and registrability. The selected viewpoints form a fully connected observation network, which can subsequently be used to minimize registration errors via Least Squares Fitting (LSF).

As shown in fig. 2, VF-Plan starts from a floorplan containing structural elements such as walls, columns, doors, and windows. The process includes VF computation, medial axis and joint-point extraction, and candidate viewpoint initialization. Overlap ratios between candidates are then evaluated, and a greedy algorithm selects the minimal subset forming a connected VPN. To accelerate the visibility checks for each viewpoint, we employ a BSP-tree [FKN80] with balanced splitting and no depth limit, starting from the median line at each iteration. Prior to BSP-tree construction and VF computation, the floorplan is preprocessed into closed polygons while preserving all structural elements.

4.2. Visibility Field (VF)

Laser scanning is limited to visible surfaces, leaving occluded areas unobserved. Therefore, visibility and coverage analysis are essential for viewpoint planning, with key objectives of maximizing observation Coverage and registrability. The field of view is cen-




Figure 3: Scan model of a static scanner. The scanner provides valid observations within a 360° angle and a range from r_{min} to r_{max} .

tral to VPN optimization, defining the information captured at each viewpoint and guiding the selection of optimal viewpoints.

A common method for evaluating the information scanned by a LiDAR sensor is based on the length of observed walls [JL22]. Typically, wall length is counted either by ensuring both endpoints are visible or by partitioning walls into uniformly sampled segments within the scanner's range. For 3D cases, visibility is evaluated by counting the number of observed voxels. However, these approaches lack normalization, yielding values that vary with discretization resolution and do not account for differences in distance or orientation to the scanner, impacting computational stability and load without improving coverage accuracy.

Our approach instead uses the valid scanned angle to integrate relevant factors into a single index, uniformly scaled from 0 to 2π . We define a static LiDAR scanner with a 360° field of view in the X-Y plane, a range from r_{min} to r_{max} , and uniform scanning speed, as shown in fig. 3. Only objects within this range are observable. Generally, closer viewpoints collect more information, but excessive proximity can occlude other parts of the field. Scanning is performed in discrete observation angles, with observed information defined as the valid observed angle, ensuring alignment with scanner properties.

To quantify observed data, let Circle_min intersect line segment AB at points C and D , and Circle_max at points E and F . The valid observed line segment is:

$$L_{valid} = \cap(L_{AB}, L_{EF}) - \cap(L_{AB}, L_{CD}) \quad (4)$$

and the corresponding observed angle is:

$$\theta_{valid} = \text{Angle}_O(L_{valid}) \quad (5)$$

Each viewpoint is assigned a View Angle, with the scene represented as a grid where each cell center is a viewpoint. Thus, the scene is represented as a VF, as shown in fig. 4(b), which aligns well with LiDAR properties and expert experience, displaying expected visibility patterns.

4.3. Converging Lines and Joints

In the VF, observation data naturally converge toward the medial axis and its joint points. Placing viewpoints along these skeletal structures guarantees full coverage (section 3, appendix A), reducing VPN optimization from a 2D search to a 1D problem. We extract the medial axis using [SP08] via OpenCV [Bra00] and detect joints by restricting inside room polygons: a skeleton point is




Figure 4: Visibility Field (VF) for static LiDAR scanning. (a, c) VF computed with only one wall (green highlight); (b, d) VF with all walls, revealing the global visibility structure. (a, b) show a rectangular room, and (c, d) depict a flat layout.

marked as a joint candidate if at least three of its eight neighbors are also skeleton points. Connected-component analysis groups candidates, and each cluster's centroid is retained as a representative joint.

For each skeleton ridge R_{ij} between joints J_i and J_j , we compute the overlap ratio O_{ij} . If O_{ij} is below the threshold, a midpoint J_m is inserted to split the ridge, ensuring connectivity. We refer to all joints and added points as *converging points*, and their connecting ridges as *converging lines*. This process iterates until every converging point satisfies the overlap requirement for reliable registration.

4.4. Overlap Ratio




Figure 5: Overlap computation between wall segments from two viewpoints. (a, b) Segments visible from viewpoints A and B. (c) Segments jointly visible, used to measure connectivity.

The overlap ratio quantifies shared visibility between two viewpoints and is key for robust registration. We define five variants based on either segment length or angular alignment, evaluated for their ability to maintain connectivity independent of viewpoint spacing.

As shown in fig. 5, L_a and L_b denote wall segments visible from viewpoints V_a and V_b , with observation angles θ_a and θ_b :

$$\theta_a = \text{Angle}(V_a, L_a) \quad (6)$$

The segments observed by both V_a and V_b are L_{ab} :

$$L_{ab} = L_a \cap L_b \quad (7)$$

Length-based ratios:

$$O^{\text{Min_Len}} = \frac{L_{ab}}{\min(L_a, L_b)} \quad (8)$$

$$O^{\text{Mean_Len}} = \frac{2L_{ab}}{L_a + L_b} \quad (9)$$

$$O^{\text{Union_Len}} = \frac{L_{ab}}{L_a \cup L_b} \quad (10)$$

Angle-based ratios, where θ_a^{ab} and θ_b^{ab} are the angular spans of L_{ab} :

$$O_{\text{Union_Ang}} = \frac{\theta_a^{ab} + \theta_b^{ab}}{\theta_a + \theta_b} \quad (11)$$

$$O_{\text{Mean_Ang}} = \frac{\theta_a^{ab}}{2\theta_a} + \frac{\theta_b^{ab}}{2\theta_b} \quad (12)$$

4.5. Greedy Optimization for VPN Construction

The VF-Plan greedy algorithm constructs an efficient VPN by initializing candidate viewpoints at strategically selected joint points within the VF. For each pair of candidate viewpoints, an Overlap Ratio O is computed to quantify shared visibility. Pairs exceeding a predefined overlap threshold are retained as edges in the VPN, ensuring robust connectivity. To enforce the coverage constraint, we construct a binary Coverage Table C (see eq. (1)) that records whether each viewpoint observes each structural segment.

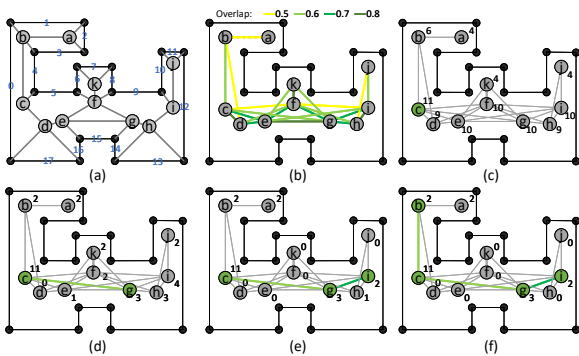


Figure 6: Greedy Optimization of VPN. (a) Floorplan showing candidate viewpoints (labeled a to k) positioned along walls (numbered 0 to 17); (b) Connectivity graph, where edges link viewpoints with overlap ratio O exceeding the threshold criterion; (c) Initial seed viewpoint selection based on maximum wall visibility (numbers near viewpoints indicate the count of observed walls); (d) - (f) Iterative expansion of the seed viewpoint set by adding adjacent viewpoints that provide the highest additional coverage, enhancing overall viewpoint arrangement.

As outlined in Algorithm 1 and illustrated in fig. 6, the greedy optimization initiates by selecting the viewpoint with the highest coverage as the seed. The algorithm then iteratively expands by adding neighboring viewpoints that maximize additional coverage, ensuring each new viewpoint significantly contributes to covering line segments while constructing a minimal, fully connected VPN.

Upon completion of the initial search, the algorithm produces a single connected component encompassing all essential viewpoints, with sufficient overlap to facilitate stable connections. The algorithm subsequently enhances network stability by identifying cycles within the VPN graph. For nodes that cannot form cycles (i.e., loose nodes), Converging Points are added to ensure cohesion without redundancy. Additionally, any node with more than three connecting edges from a cycle selects a Converging Joint to establish loops, thereby strengthening network connectivity.

Algorithm 1 Greedy Optimization of VPN

Input: Candidate viewpoints V , line segments L

Output: Minimal set of viewpoints V_{opt} covering all L and forming a connected network

1. Initialize Coverage Table C
2. Identify V_i with maximum C_i , add V_i to Viewpoint Seed Set (VPS)
3. Add adjacent viewpoints of V_i to VPB
4. Remove all L observed by V_i from C
5. While unobserved L exists in C :
 - a. Select V_i in VPB maximizing new coverage C_i
 - b. If multiple options, select V_i with highest overlap O with VPS
 - c. Move V_i to VPS, add adjacencies to VPB
 - d. Remove all L observed by V_i from C

return VPS as the optimized minimal VPN

5. Experiments

In this section, we evaluate the proposed VPN optimization algorithm across multiple scenarios and compare its performance with state-of-the-art methods in both indoor and outdoor environments. Due to the lack of a dedicated benchmark for VPN optimization, we compiled a dataset from established sources, focusing on two primary scene types: indoor and outdoor environments. This dataset provides a robust foundation for evaluating VPN optimization across a range of realistic settings. The pipeline is implemented in C++14 and compiled on Windows 11, running on an Intel Core i9 CPU (24 cores, 2.20 GHz) with 32 GB RAM.

5.1. Dataset

For indoor environments, we selected scenes with varied room layouts to evaluate the algorithm's adaptability across multiple configurations. The Structure3D dataset [ZZL*20] was included for its extensive collection of diverse Asian residential scenes, featuring both complex and compact room types frequently used in 3D modeling and synthetic data generation. VPNs generated from this dataset can also facilitate synthetic LiDAR data production for advancing research in 3D scene understanding. Specific scenes such as 00007, 00009, 00071, 02970, and 02972, discussed in subsequent sections, are drawn from Structure3D. For further comparisons, we incorporated additional scenes from Zeng1 [ZLC*22], Noichl [NL23], Xu [XZY*17], and TUB1 from the ISPRS benchmark [ZZG*24, KTA*21].

For outdoor scenarios, we limited our approach to 2D site plans without considering terrain elevation, distinguishing our method from elevation-based approaches like [CZH*22]. Selected scenes include the CCIT and Crowsnest datasets from the University of Calgary [JL18] and the ECUT dataset from East China University of Technology [CZH*22], which represent open-space environments suitable for VPN optimization.

To evaluate VPN quality, we rely on expert-labeled annotations as standard ground truth data is unavailable. Where possible, we use existing expert-labeled VPNs (e.g., Zeng1, scene1 and TUB1),

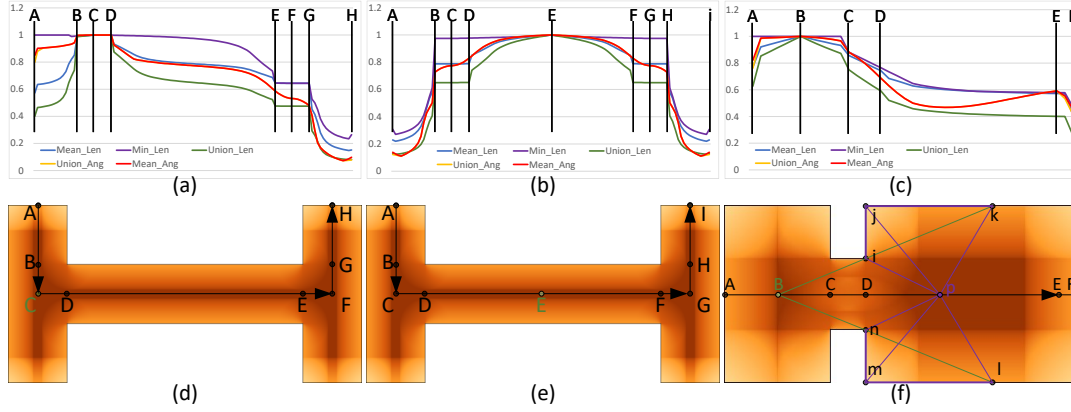


Figure 7: Performance of five proposed overlap ratios. The first row shows overlap ratios between moving and reference viewpoints; the second row illustrates room configurations, with black arrows indicating movement direction and green dots marking reference viewpoints.

and for other cases, independent expert labels were created for comparison with manual designs.

5.2. Performance Metrics

The optimization algorithm aims to construct a minimal, fully connected VPN with reduced redundancy. As most state-of-the-art methods achieve 100% coverage, we focus on two key metrics: the number of optimal viewpoints and network connectivity.

Network connectivity is measured using the Weighted Average Path Length (WAPL), which reflects the compactness of the network [BLM*06]. A shorter WAPL indicates stronger connectivity and a more efficient structure. For a network with weighted edges w , WAPL is calculated as:

$$WAPL = \frac{1}{N(N-1)} \sum_{i \neq j} d_w(i, j) \quad (13)$$

where $d_w(i, j)$ represents the shortest weighted path length between nodes i and j , and N is the total number of nodes.

The edge weight w is defined as

$$w_{ij} = 1 - O_{ij} \quad (14)$$

A larger O_{ij} implies stronger registration-based connectivity, meaning a shorter length between the two nodes i and j .

In cases where baseline methods do not yield a single connected component, some viewpoints may have a small overlap ratio with others. To ensure WAPL remains computable, we assign a large weight (100) to node pairs that lack a connecting path.

5.3. Ablation Study

We evaluate overlap ratio type, scan radius, window inclusion, and, in appendix B, wall partition, VF grid resolution and overlap thresholds. VPN performance is largely insensitive to wall partition length and VF grid resolution, though these parameters affect computation time, while others require tuning for specific scenarios. Default settings for subsequent experiments are: $O^{\text{Mean_Len}}$ for

overlap computation, overlap ratio 0.4 indoors and 0.3 outdoors, $R_{\min} = 0.6$ m indoors and 1.2 m outdoors, $R_{\max} = 30.0$ m indoors and 75.0 m outdoors, wall partition 0.1 m indoors and 1.0 m outdoors, and VF grid resolution 0.02 m indoors and 0.25 m outdoors.

5.3.1. Overlap Ratio Type

We propose five overlap ratios according to the length of scanned walls and according angles to viewpoints in section 4.4. Here we evaluate them in two test indoor scenes, illustrated in fig. 7. Among them, $O^{\text{Min_Len}}$ quickly reaches 100%, showing limited variability along movement paths, as seen from B to H in fig. 7(b) and (e), making it less suitable for dynamic scenes. The angle-based measures, $O^{\text{Union_Ang}}$ and $O^{\text{Mean_Ang}}$, exhibit non-monotonic behavior, particularly along paths from C to E, shown in fig. 7(c) and (f), which limits their usefulness for consistent registration. The $O^{\text{Union_Len}}$ measure tends to be too small, often underestimating connectivity. Therefore, we select $O^{\text{Mean_Len}}$ for subsequent experiments and analyses due to its balanced performance across different scenarios.

5.3.2. Scan Radius

The scan model used in this study operates in the X-Y plane with a 360° field of view, bounded by R_{\min} and R_{\max} , defining near and far blind zones. Typical static scanners, including the FARO M70, Trimble X7, and Leica BLK360, have maximum ranges around 40 meters and minimum ranges of about 1 meter, allowing for comprehensive indoor coverage but requiring precise placement in small rooms and narrow corridors.

For outdoor scenarios, we examine the impact of varying R_{\max} , testing values at 15m, 30m, 45m, and 75m, with R_{\min} fixed at 1m. As illustrated in fig. 8, increasing R_{\max} expands the area of valid placements for viewpoints—defined where the View Angle exceeds 1.0—and reduces the Viewpoint Count (VC) needed for optimized coverage, from 23 to 14, 10, and 8. Larger R_{\max} values are advantageous for open spaces, while smaller values require additional viewpoints for full coverage. These results indicate that our method aligns effectively with real-world scanner specifications and expert guidelines.

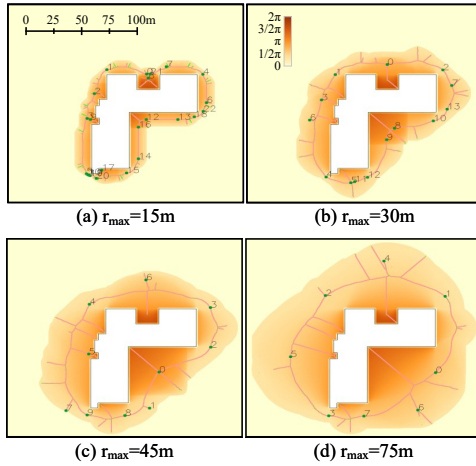


Figure 8: Effect of R_{max} on Viewpoint Placement and Network Efficiency. As R_{max} increases from 15m to 75m, the valid placement area (where View Angle > 1.0) grows, reducing the VC needed for optimized VPN coverage from 23 to 8.

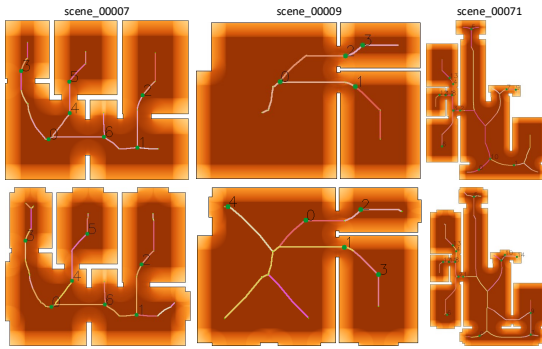


Figure 9: Comparison of VPN structures with and without windows. Including windows (bottom row) adds complexity to the skeleton and slightly increases the VC.

5.3.3. Effect of Windows

Incorporating windows into the model increases the complexity of the skeleton structure by introducing additional joint points, resulting in a more detailed medial axis, as illustrated in fig. 9. However, accurate scanning of windows is often essential in practical applications. To assess the impact of including windows, we evaluate the performance of our algorithm with and without windows considered in the optimization. When windows are included, the VPN algorithm requires only a slight increase in the number of viewpoints to achieve complete coverage, as shown in table 1. On average, the number of viewpoints is approximately 1.4 times the number of rooms without windows, and increases modestly to 1.6 times with windows. These results demonstrate that the algorithm effectively handles the additional geometric detail introduced by windows. For applications requiring precise window coverage, including window frames in the optimization process is recommended.

Scene	00007	00009	00071
Number of Rooms	5	3	9
VC (Without Windows)	7	4	14
VC (With Windows)	7	5	15

Table 1: VC comparison for room coverage with and without windows across different scenes.

5.4. Results and Analysis

Scene	WAPL↓			VC↓			Baseline Method
	Expert	Ours	Baseline	Expert	Ours	Baseline	
Zeng1	0.47	0.53	43.37	12	12	10	Zeng [ZLC*22]
Noichl	0.32	0.38	23.91	8	10	7	Noichl2023 [NB23]
TUB1	1.03	1.00	40.97	28	29	19	Zhai2024 [ZZG*24]
Xu	0.35	0.31	28.69	8	8	7	Xu2017 [XZY*17]
CCIT	0.52	0.47	38.12	8	8	7	Jia2018 [JL18]
Crownsnest	0.50	0.49	34.02	9	8	8	Jia2018 [JL18]
ECUT	0.67	0.76	11.94	58	52	139	Chen2023 [CZH*22]

Table 2: Comparison of VC and WAPL across indoor and outdoor scenes. Lower WAPL indicates more compact network connectivity.

We benchmarked our method against state-of-the-art approaches using datasets from prior studies, representing a range of indoor and outdoor scenes. Indoor datasets include Zeng1 [ZLC*22], Noichl [NB23], and TUB1 from the ISPRS dataset [ZZG*24]. Outdoor datasets consist of CCIT and Crownsnest from the University of Calgary [JL18] and ECUT from East China University of Technology [CZH*22].

Indoor Results. As shown in fig. 10, our method’s viewpoint optimization closely aligns with expert-designed networks, achieving comprehensive coverage and creating compact, interconnected structures. This optimized network design enhances registration accuracy and connectivity, even in complex indoor settings.

Outdoor Results. For outdoor scenes, our method effectively adapts to open spaces, as illustrated in fig. 11. Adjustments to the R_{max} parameter notably improved performance in expansive environments, reducing the required number of viewpoints without compromising coverage.

Table 2 shows that our method matches or surpasses expert-designed networks while significantly outperforming baseline methods. Across all scenes, our VC remains comparable to experts, differing by at most one viewpoint, yet our WAPL is consistently lower in five of seven cases, indicating more compact and cohesive networks. Compared to baselines, the improvement is substantial: in indoor scenes such as Zeng1 and Noichl, WAPL is reduced by 98% and 97% respectively, eliminating the fragmented connectivity seen in baseline outputs (green ellipses in figs. 10 and 11). In challenging outdoor cases, reductions remain high, including 91% in ECUT and 87% in CCIT, demonstrating scalability to large and complex environments. These results confirm that our VF-based optimization preserves expert-level coverage while achieving markedly superior network compactness over existing automated methods.

To further validate the optimized viewpoints, we simulated static LiDAR scanning at each optimized viewpoint using Helios++ [EYW*22]. As shown in fig. 12, the simulated LiDAR point

clouds confirm complete scene coverage, with effective overlap between adjacent scans across both indoor and outdoor environments.

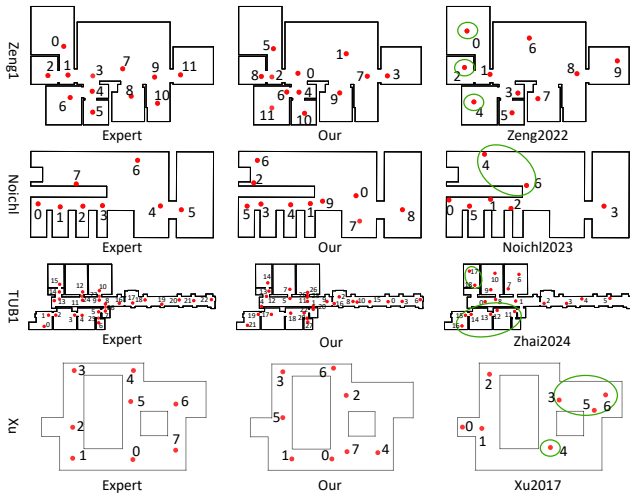


Figure 10: Viewpoint planning results for indoor environments. Red dots represent individual viewpoints, while green ellipses indicate clusters of viewpoints that are not connected to the main network.

6. Conclusions

We presented a novel solution to the VPP and the AGP in static LiDAR scanning on 2D floor plans, leveraging the VF to reduce the optimization complexity from two dimensions to one. This VF-based approach enables the construction of a minimal, fully connected VPN that ensures complete coverage with reduced redundancy, advancing viewpoint planning for static LiDAR applications such as BIM construction. Our method offers a scalable and efficient framework applicable in structured and semi-structured environments, addressing real-world challenges in 3D data collection and autonomous navigation. Our VF framework adapts well to real-world settings with potential obstructions by integrating partial occlusion handling based on initial scans.

The current formulation assumes accurate and complete 2D floor plans, which may not be available in all scenarios. In cases of noisy input or partial scans, the quality of the VF and the resulting VPN may degrade due to inaccuracies in boundary detection or incomplete visibility information. While preliminary layouts can be generated using low-cost SLAM systems or coarse surveys, further robustness to noise, clutter, and missing geometry remains an open challenge. Additionally, the current implementation focuses on planar (2D) environments; extending the approach to handle 3D structures with significant vertical complexity will require additional algorithmic considerations. Future work will explore its application to large-scale LiDAR data collection and synthetic data generation, as well as integration with path planning in both 2D and 3D scenarios for more versatile deployment in complex environments.

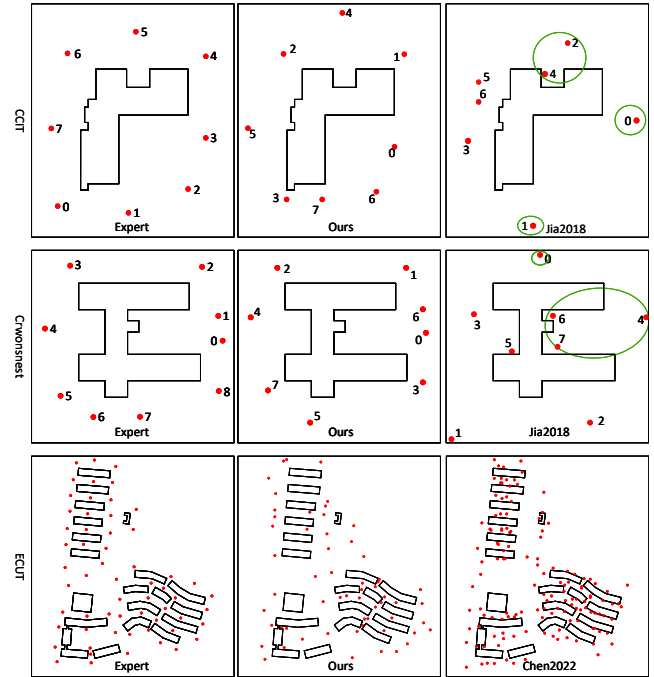


Figure 11: Viewpoint planning results for outdoor environments. Red dots represent individual viewpoints, while green ellipses indicate clusters of viewpoints that are not connected to the main network.

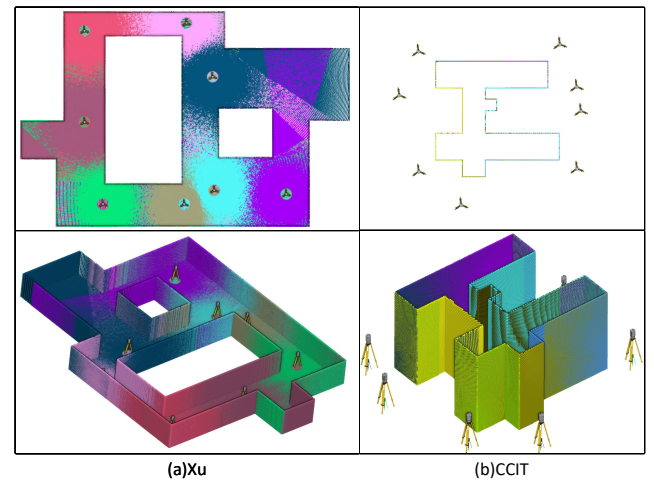


Figure 12: Simulated LiDAR points from optimal viewpoints. Each scan is uniquely colored. Top row shows top views, and bottom row shows oblique views with tripods indicating viewpoints.

References

- [AAM21] ABRAHAMSEN M., ADAMASZEK A., MILTZOW T.: The art gallery problem is $\exists\mathbb{R}$ -complete. *Journal of the ACM* 69, 1 (2021), 1–70. doi:10.1145/3486655. 2
- [ABT21] ARYAN A., BOSCHÉ F., TANG P.: Planning for terrestrial laser scanning in construction: A review. *Automation in Construction* 125 (2021), 103551. 2
- [BKA*16] BIRCHER A., KAMEL M., ALEXIS K., OLEYNIKOVA H., SIEGWART R.: Receding horizon "next-best-view" planner for 3d exploration. In *2016 IEEE International Conference on Robotics and Automation (ICRA)* (2016), IEEE, pp. 1462–1468. doi:10.1109/ICRA.2016.7487276. 3
- [BLM*06] BOCCALETTI S., LATORA V., MORENO Y., CHAVEZ M., HWANG D.-U.: Complex networks: Structure and dynamics. *Physics reports* 424, 4-5 (2006), 175–308. 6
- [Blu67] BLUM H.: A transformation for extracting new descriptions of shape. *Models for the perception of speech and visual form* (1967), 362–380. 11
- [Bra00] BRADSKI G.: The opencv library. *Dr. Dobbs' Journal: Software Tools for the Professional Programmer* 25, 11 (2000), 120–123. 4
- [CCS*21] CUI A., CASAS S., SADAT A., LIAO R., URTASUN R.: Look-out: Diverse multi-future prediction and planning for self-driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021), pp. 16107–16116. 3
- [CdRdS11] COUTO M. C., DE REZENDE P. J., DE SOUZA C. C.: An exact algorithm for minimizing vertex guards on art galleries. *International Transactions in Operational Research* 18, 4 (2011), 425–448. doi:10.1111/j.1475-3995.2011.00800.x. 2
- [Chv79] CHVATAL V.: A greedy heuristic for the set-covering problem. *Mathematics of operations research* 4, 3 (1979), 233–235. 11
- [CKP*25] CHUNG D., KIM J., PAIK S., IM S., KIM H.: Automated system of scaffold point cloud data acquisition using a robot dog. *Automation in Construction* 170 (2025), 105944. 2
- [CSW*25] CHEN Z., SONG C., WANG B., TAO X., ZHANG X., LIN F., CHENG J. C.: Automated reality capture for indoor inspection using bim and a multi-sensor quadruped robot. *Automation in Construction* 170 (2025), 105930. 2
- [CZH*22] CHEN Z., ZHANG W., HUANG R., DONG Z., CHEN C., JIANG L., WANG H.: 3d model-based terrestrial laser scanning (tls) observation network planning for large-scale building facades. *Automation in Construction* 144 (2022), 104594. 5, 7
- [CZR*23] CAO C., ZHU H., REN Z., CHOSSET H., ZHANG J.: Representation granularity enables time-efficient autonomous exploration in large, complex worlds. *Science Robotics* 8, 80 (2023), eadf0970. doi:10.1126/scirobotics.adf0970. 3
- [EYW*22] ESMORÍS A. M., YERMO M., WEISER H., WINIWARTER L., HÖFLE B., RIVERA F. F.: Virtual lidar simulation as a high performance computing challenge: Toward hpc helios++. *IEEE Access* 10 (2022), 105052–105073. URL: <https://ieeexplore.ieee.org/document/9906068>, doi:https://doi.org/10.1109/ACCESS.2022.3211072. 7
- [FG09] FUSCO G., GUPTA H.: Selection and orientation of directional sensors for coverage maximization. In *2009 6th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks* (2009), IEEE, pp. 1–9. doi:10.1109/SAHCN.2009.5168945. 2
- [FKN80] FUCHS H., KEDEM Z. M., NAYLOR B. F.: On visible surface generation by a priori tree structures. In *Proceedings of the 7th annual conference on Computer graphics and interactive techniques* (1980), pp. 124–133. 3
- [GCW15] GHANEM B., CAO Y., WONKA P.: Designing camera networks by convex quadratic programming. *Computer Graphics Forum* 34, 2 (2015), 77–88. doi:10.1111/cgf.12544. 2
- [HG25] HU D., GAN V. J.: Semantic navigation for automated robotic inspection and indoor environment quality monitoring. *Automation in Construction* 170 (2025), 105949. 2
- [Hnh18] HEPP B., NIESSNER M., HILLIGES O.: Plan3d: Viewpoint and trajectory optimization for aerial multi-view stereo reconstruction. *ACM Transactions on Graphics (TOG)* 38, 1 (2018), 1–17. doi:10.1145/3272127. 2
- [IGFER22] IBRAHIM A., GOLPARVAR-FARD M., EL-RAYES K.: Metrics and methods for evaluating model-driven reality capture plans. *Computer-Aided Civil and Infrastructure Engineering* 37, 1 (2022), 55–72. doi:10.1111/mice.12733. 2
- [JL18] JIA F., LICHTI D. D.: An efficient, hierarchical viewpoint planning strategy for terrestrial laser scanner networks. In *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* (2018), vol. 4, pp. 137–144. doi:10.5194/isprs-annals-IV-2-137-2018. 2, 5, 7
- [JL19] JIA F., LICHTI D. D.: A model-based design system for terrestrial laser scanning networks in complex sites. *Remote Sensing* 11, 15 (2019), 1749. doi:10.3390/rs11151749. 2
- [JL22] JIA F., LICHTI D. D.: A practical algorithm for the viewpoint planning of terrestrial laser scanners. *Geomatics* 2, 2 (2022), 181–196. doi:10.3390/geomatics2020013. 2, 4
- [KF11] KARAMAN S., FRAZZOLI E.: Incremental sampling-based algorithms for optimal motion planning. *Robotics Science and Systems (RSS)* (2011). doi:10.15607/rss.2011.v7.021. 2
- [KTA*21] KHOSHELHAM K., TRAN H., ACHARYA D., VILARIÑO L. D., KANG Z., DALYOT S.: Results of the isprs benchmark on indoor modelling. *ISPRS Open Journal of Photogrammetry and Remote Sensing* 2 (2021), 100008. 5
- [LKL21] LI F., KIM M.-K., LEE D.-E.: Geometrical model-based scan planning approach for the classification of rebar diameters. *Automation in Construction* 130 (2021), 103848. doi:10.1016/j.autcon.2021.103848. 2
- [LKS*25] LI F., KIM M.-K., SIM S.-H., CHI H.-L., LEE D.-E.: Full-scale application of dimensional quality assessment on precast slabs: A scan planning approach. *Measurement* 242 (2025), 115850. 2
- [LL86] LEE D., LIN A.: Computational complexity of art gallery problems. *IEEE Transactions on Information Theory* 32, 2 (1986), 276–282. doi:10.1109/TIT.1986.1057160. 2
- [LLZ*22] LI D., LIU J., ZENG Y., CHENG G., DONG B., CHEN Y. F.: 3d model-based scan planning for space frame structures considering site conditions. *Automation in Construction* 140 (2022), 104363. doi:10.1016/j.autcon.2022.104363. 2
- [LXHL21] LIU J., XU D., HYYPPÄ J., LIANG Y.: A survey of applications with combined bim and 3d laser scanning in the life cycle of buildings. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 14 (2021), 5627–5637. 2
- [MADR11] MORSLEY Y., AOUF N., DJOUADI M. S., RICHARDSON M.: Particle swarm optimization-inspired probability algorithm for optimal camera network placement. *IEEE Sensors Journal* 12, 5 (2011), 1402–1412. doi:10.1109/JSEN.2011.2171706. 2
- [MHS*23] MABOUDI M., HOMAËI M., SONG S., MALIHI S., SAADATSERESHT M., GERKE M.: A review on viewpoints and path planning for uav-based 3d reconstruction. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 16 (2023), 5026–5048. doi:10.1109/JSTARS.2023.3269826. 2
- [MLL23] MA Z., LIU Y., LI J.: Review on automated quality inspection of precast concrete components. *Automation in Construction* 150 (2023), 104828. 1
- [NB22] NOICHL F., BORRMANN A.: Towards multicriterial scan planning in complex 3d environments. In *International Conference on Computing in Civil and Building Engineering* (2022), Springer, pp. 223–235. 2

- [NB23] NOICHL F., BORRMANN A.: Automated deterministic model-based indoor scan planning. In *ECPPM 2022-eWork and eBusiness in Architecture, Engineering and Construction 2022*. CRC Press, 2023, pp. 559–566. [5](#), [7](#)
- [NLB24] NOICHL F., LICHTI D. D., BORRMANN A.: Automating adaptive scan planning for static laser scanning in complex 3d environments. *Automation in Construction* 165 (2024), 105511. [2](#)
- [NLMC23] NOËL T., LEHUGER A., MARCHAND E., CHAUMETTE F.: Skeleton disk-graph roadmap: A sparse deterministic roadmap for safe 2d navigation and exploration. *IEEE Robotics and Automation Letters* 9, 1 (2023), 555–562. [doi:10.1109/LRA.2023.3236339](#). [3](#)
- [OF15] OTTE M., FRAZZOLI E.: Rt^2 : Real-time motion planning/replanning for environments with unpredictable obstacles. In *Algorithmic foundations of robotics XI: selected contributions of the eleventh international workshop on the algorithmic foundations of robotics* (2015), Springer, pp. 461–478. [2](#)
- [PSC*23] PLACED J. A., STRADER J., CARRILLO H., ATANASOV N., INDELMAN V., CARLONE L., CASTELLANOS J. A.: A survey on active simultaneous localization and mapping: State of the art and new frontiers. *IEEE Transactions on Robotics* 39, 3 (2023), 1686–1705. [2](#)
- [QWTW21] QIU Q., WANG M., TANG X., WANG Q.: Scan planning for existing buildings without bim based on user-defined data quality requirements and genetic algorithm. *Automation in Construction* 130 (2021), 103841. [2](#)
- [RLGA22] ROUGERON G., LE GARREC J., ANDRIOT C.: Optimal positioning of terrestrial lidar scanner stations in complex 3d environments with a multiobjective optimization method based on gpu simulations. *ISPRS Journal of Photogrammetry and Remote Sensing* 193 (2022), 60–76. [2](#)
- [SP08] SIDDIQI K., PIZER S.: *Medial representations: mathematics, algorithms and applications*, vol. 37. Springer Science & Business Media, 2008. [4](#)
- [SRR03] SCOTT W. R., ROTH G., RIVEST J.-F.: View planning for automated three-dimensional object reconstruction and inspection. *ACM Computing Surveys (CSUR)* 35, 1 (2003), 64–96. [1](#), [2](#)
- [SSZ*20] SHI L., SHI D., ZHANG X., MEUNIER B., ZHANG H., WANG Z., VLADIMIRESCU A., LI W., ZHANG Y., COSMAS J., ET AL.: 5g internet of radio light positioning system for indoor broadcasting service. *IEEE Transactions on Broadcasting* 66, 2 (2020), 534–544. [2](#)
- [TMMA21] TAN C. S., MOHD-MOKHTAR R., ARSHAD M. R.: A comprehensive review of coverage path planning in robotics using classical and heuristic algorithms. *IEEE Access* 9 (2021), 119310–119342. [2](#)
- [WZL*20] WANG Q., ZHAO X., LV Z., MA X., ZHANG R., LIN Y.: Optimizing the ultra-dense 5g base stations in urban outdoor areas: Coupling gis and heuristic optimization. *Sustainable Cities and Society* 63 (2020), 102445. [2](#)
- [XTV14] XU Q., TREMBLAY J., VERBRUGGE C.: Generative methods for guard and camera placement in stealth games. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* (2014), vol. 10, pp. 87–93. [2](#)
- [XWYZ25] XU Y., WANG Y., YANG J., ZHANG J.: Two-stage terrestrial laser scan planning framework for geometric measurement of civil infrastructures. *Measurement* 242 (2025), 115785. [2](#)
- [XYH*18] XIE K., YANG H., HUANG S., LISCHINSKI D., CHRISTIE M., XU K., GONG M., COHEN-OR D., HUANG H.: Creating and chaining camera moves for quadrotor videography. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–13. [2](#)
- [XZD*24] XI Y., ZHU C., DUAN Y., YI R., ZHENG L., HE H., XU K.: Thp: Tensor-field-driven hierarchical path planning for autonomous scene exploration with depth sensors. *Computational Visual Media* (2024), 1–15. [2](#)
- [XZY*17] XU K., ZHENG L., YAN Z., YAN G., ZHANG E., NIESSNER M., DEUSSEN O., COHEN-OR D., HUANG H.: Autonomous reconstruction of unknown indoor scenes guided by time-varying tensor fields. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 1–15. [2](#), [5](#), [7](#)
- [ZCH*23] ZHANG W., CHEN Z., HUANG R., DONG Z., JIANG L., XIA Y., CHEN B., WANG H.: Multiobjective optimization-based terrestrial laser scanning layout planning for landslide monitoring. *IEEE Transactions on Geoscience and Remote Sensing* 61 (2023), 1–15. [2](#)
- [ZLC*22] ZENG Y., LIU J., CAO Q., WU Z., CHEN B., LI D., CHENG G.: Optimal planning of indoor laser scans based on continuous optimization. *Automation in Construction* 143 (2022), 104552. [5](#), [7](#)
- [ZYCH13] ZHAO J., YOSHIDA R., CHEUNG S.-C. S., HAWS D.: Approximate techniques in solving optimal camera placement problems. *International Journal of Distributed Sensor Networks* 9, 11 (2013), 241913. [2](#)
- [ZYG*21] ZHANG H., YAO Y., XIE K., FU C.-W., ZHANG H., HUANG H.: Continuous aerial path planning for 3d urban scene reconstruction. *ACM Trans. Graph.* 40, 6 (2021), 225–1. [2](#)
- [ZYG*24] ZHAI R., ZOU J., GAN V. J., HAN X., WANG Y., ZHAO Y.: Semantic enrichment of bim with indoorgml for quadruped robot navigation and automated 3d scanning. *Automation in Construction* 166 (2024), 105605. [5](#), [7](#)
- [ZZL*20] ZHENG J., ZHANG J., LI J., TANG R., GAO S., ZHOU Z.: Structured3d: A large photo-realistic dataset for structured 3d modeling. In *Proceedings of The European Conference on Computer Vision (ECCV)* (2020). [5](#)

Supplementary Materials

Appendix A: Theoretical Foundations, Medial Axis Convergence, and Greedy Optimality

The key idea of VF-Plan is that visibility information concentrates on a one dimensional skeleton of the polygonal domain. In the visibility field, visibility propagates from the boundary and accumulates on the medial axis, which reduces the search from two dimensions to one. We formalize this convergence, prove that restricting candidate viewpoints to the medial axis preserves global optimality, and present quantitative performance guarantees for the greedy selection strategy.

Skeleton Completeness and Visibility Propagation

Medial Axis (MA). The medial axis of a polygon P is the locus of points that are equidistant to at least two boundary elements. In Blum’s grassfire transform [Blu67] the boundary ignites at time zero and inward fronts meet on the medial axis, which explains the concentration of visibility information on the skeleton.

Theorem 1 (Skeleton Completeness) Let $p \in P$ see a set of boundary segments $L_p \subseteq L$. There exists a finite set $\{q_\ell\} \subseteq \text{MA}(P)$ such that

$$L_p \subseteq \bigcup_{\ell \in L_p} \text{Vis}(q_\ell). \quad (15)$$

Hence any interior coverage provided by p can be reproduced by viewpoints on the medial axis.

Sketch of Proof. By the grassfire view [Blu67] visibility from each boundary segment propagates inward until it meets $\text{MA}(P)$. For every $l_j \in L_p$ there exists $q_j \in \text{MA}(P)$ that sees l_j . Therefore $\text{Vis}(p) \subseteq \bigcup_{l_j \in \text{Vis}(p)} \text{Vis}(q_j)$. Figure 13 illustrates the construction. Restricting candidates to $\text{MA}(P)$ does not reduce achievable coverage.

Global Minimal Cover and Connectivity on the Skeleton

By theorem 1 any feasible interior solution transfers to the medial axis. In the eq. (3) restricting V to $\text{MA}(P)$ including branch and joint nodes therefore captures all feasible solutions.

Theorem 2 (Global Minimal Cover and Connectivity) Let $S^* \subseteq \text{MA}(P)$ be a minimal solution to eq. (3) that covers all $l_j \in L$ and satisfies the required overlap-based connectivity. Then S^* is irreducible and no smaller configuration in P can achieve the same objectives. Therefore S^* is globally optimal.

Implication and heuristic search. No point outside $\text{MA}(P)$ improves coverage or connectivity over a minimal skeletal solution. Solving eq. (3) on $\text{MA}(P)$ yields a configuration that is optimal for the whole polygon. The problem remains NP-hard, yet a greedy strategy is effective in practice and admits quantitative guarantees for the number of viewpoints.

Greedy Selection and Suboptimality Bounds

Let $n = |L|$ be the number of discretized boundary segments. Ignoring connectivity, selecting viewpoints to cover L is a set cover instance. The classic greedy rule achieves an approximation factor $H_n \leq 1 + \ln n$ for the number of selected sets [Chv79]. Our solver first runs greedy coverage on medial-axis candidates and then augments connectivity on the induced viewpoint graph whose edge weights are $w_{ij} = 1 - O_{ij}$ where O_{ij} is the overlap ratio. Connecting the greedy terminals with a Steiner tree in this graph adds only connector viewpoints along medial-axis paths. Using a constant-factor approximation for Steiner tree on graphs, the number of added connectors is within a constant factor of the optimum connector count. Combining the two stages gives

$$|S_{\text{greedy}}| \leq H_n \cdot \text{OPT}_{\text{cover}} + \alpha \text{OPT}_{\text{conn}}, \quad (16)$$

where α is a constant and $\text{OPT}_{\text{cover}}$ and OPT_{conn} are the minimal numbers of coverage and connector viewpoints on $\text{MA}(P)$. When connectivity is not stricter than coverage, $\text{OPT}_{\text{conn}} \leq \text{OPT}_{\text{cover}}$, which yields the compact bound

$$|S_{\text{greedy}}| \leq (H_n + \alpha) \text{OPT}, \quad (17)$$

with OPT the minimal skeletal solution of eq. (3). Empirically our solutions match or improve expert VC in five of seven scenes and reduce WAPL by up to 98% relative to automated baselines while remaining within one viewpoint of expert designs in most cases, which corroborates the tightness of the bound in practice.

Appendix B: More Ablation Study

Effect of Wall Partition

To evaluate the impact of wall partition length, wall boundaries are divided into segments of varying lengths including 1.0 m, 0.5 m, 0.1 m, and 0.01 m, as illustrated in fig. 14. Experimental results show that the partition length does not affect the final optimization outcome, due to the continuous nature of the VF employed in our method. The resolution of partitioning influences only the order in which viewpoints are selected. Coarser partitions, corresponding to longer segments or no partitioning, tend to place initial viewpoints at major structural intersections. A coarse partition reduces computational cost while maintaining effective coverage. Based on this observation, we set the default partition length to 0.1 m for indoor scenes and 1.0 m for outdoor scenes in all subsequent experiments.

Effect of Overlap Ratio Threshold

We analyze the influence of overlap ratio (OR) on VPN optimization using the indoor Scene_00001 dataset. The OR defines the minimum visible segment proportion shared between two viewpoints for them to be considered connected. Experiments were conducted for OR values ranging from 0.3 to 0.7, with all other parameters fixed: $R_{\min} = 0.6$ m, $R_{\max} = 30.0$ m, resolution 0.02 m, and wall partition length 0.2 m. Results in table 3 and figs. 15 and 16 show that increasing OR generally enhances network connectivity but often requires more viewpoints to maintain full coverage. For example, raising OR from 0.3 to 0.7 increases the number of viewpoints from 10 to 18 (an 80% increase) while WAPL rises from 0.38 to 0.44. Moderate OR values, such as 0.4, achieve a balanced




Figure 13: Illustration of skeleton completeness. (a) A simple polygon (black boundary) with its medial axis (gray arcs). (b) Coverage by an interior viewpoint p can be replicated by a single skeletal viewpoint q lying on the medial axis (MA). (c) In the general case, coverage from an interior viewpoint p may require multiple skeletal viewpoints (here, m and n) to fully reproduce p 's visibility. Thus, no coverage capability is lost by restricting candidate viewpoints to the MA. For clarity, boundary segments commonly observed by both p and the MA viewpoints (q , or m and n) are not drawn.

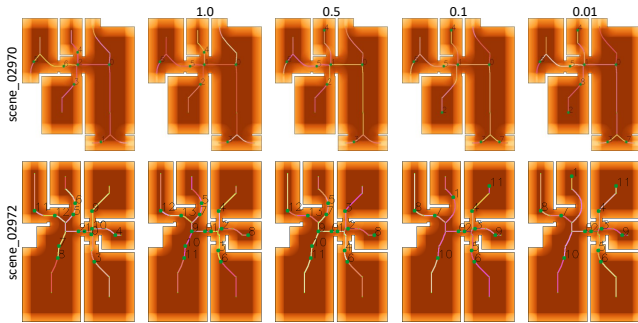


Figure 14: Effects of wall partitions. From left to right: unpartitioned walls and walls partitioned into segments of 1.0m, 0.5m, 0.1m, and 0.01m.

trade-off with only 9 viewpoints, a WAPL of 0.40, and stable runtime (≈ 0.74 s across all settings). Very high OR values (≥ 0.6) lead to sharp viewpoint growth without proportional gains in WAPL, indicating diminishing returns in connectivity improvement. These observations suggest that an OR in the range 0.3–0.5 is sufficient for most scanning tasks, providing robust connectivity while keeping network size compact.

OR	0.3	0.4	0.5	0.6	0.7
VC	10	9	11	13	18
WAPL	0.38	0.40	0.45	0.45	0.44
Time (s)	0.72	0.75	0.73	0.74	0.74

Table 3: Effect of overlap ratio (OR) on Scene_00001.

Effect of VF Grid Resolution

We examine the impact of VF grid resolution on VPN optimization using the indoor Zeng_SceneA and outdoor CCIT datasets. For Zeng_SceneA, parameters are $R_{\min} = 0.6$ m, $R_{\max} = 30.0$ m,

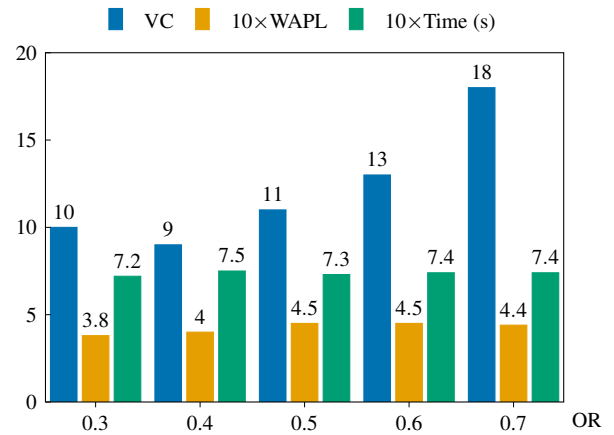


Figure 15: Effect of overlap ratio (OR) on Scene_00001. WAPL and runtime are multiplied by 10 for scale comparability.

overlap ratio 0.4, and wall partition 0.2 m. For CCIT, parameters are $R_{\min} = 1.2$ m, $R_{\max} = 75.0$ m, overlap ratio 0.3, and wall partition 1.0 m. All other settings remain fixed, and only the grid resolution is varied. As shown in table 5 and figs. 17 to 19, resolution changes have minimal influence on VPN count or WAPL, with variations under 5% for both datasets. In Zeng_SceneA, VPN count ranges from 10 to 12 and WAPL from 0.52 to 0.46 as resolution decreases from 0.01 m to 0.05 m, while runtime drops from 1.57 min to 0.03 min. In CCIT, VPN count varies from 8 to 9 and WAPL from 0.43 to 0.44 as resolution decreases from 0.15 m to 0.35 m, with runtime falling from 0.21 min to 0.03 min. Runtime scales approximately quadratically with grid resolution due to the growth in visibility evaluations. These results confirm that the VF-based approach depends primarily on the topological structure of the floor-plan, allowing the use of lower resolutions that capture necessary geometric detail while significantly improving runtime.




Figure 16: Effect of overlap ratio (OR) on viewpoint planning in Scene_00001. Increasing OR from 0.3 to 0.7 improves connectivity but raises viewpoint count (VC), showing a trade-off between coverage and compactness, while WAPL and runtime remain stable.

Table 4: Effect of resolution on Viewpoint count, WAPL, and runtime for indoor (Zeng_SceneA) and outdoor (Outdoor_CCIT) cases.

	Zeng_SceneA					Outdoor_CCIT				
Resolution (m)	0.01	0.02	0.03	0.04	0.05	0.15	0.20	0.25	0.30	0.35
VC	11	10	11	11	12	8	8	8	8	9
WAPL	0.52	0.52	0.49	0.48	0.46	0.43	0.45	0.43	0.43	0.44
Time (min)	1.57	0.16	0.07	0.04	0.03	0.21	0.08	0.05	0.03	0.03

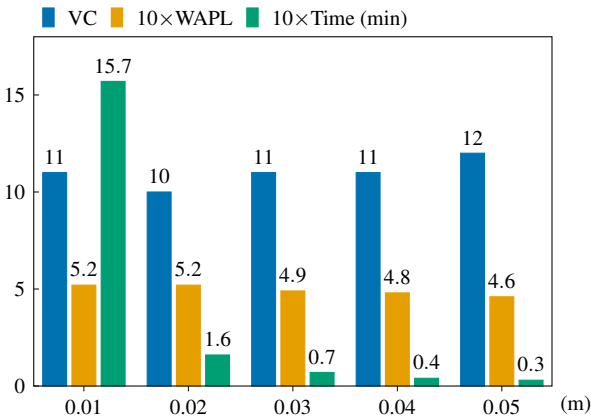


Figure 17: Effect of resolution on Viewpoint count (VC), WAPL, and runtime for indoor (Zeng_SceneA).

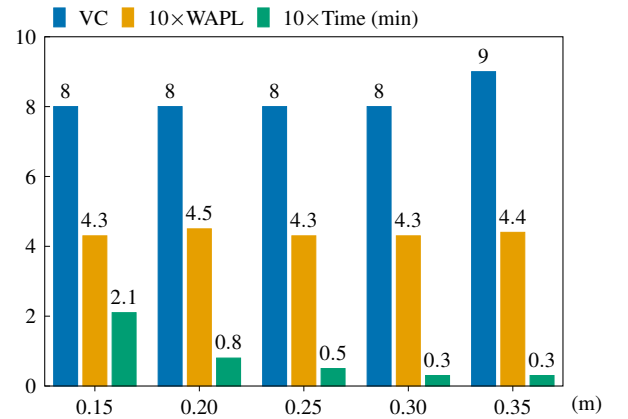


Figure 18: Effect of resolution on Viewpoint count (VC), WAPL, and runtime for outdoor (CCIT).

Time and Space Complexity Analysis

We evaluate the computational complexity of VF-based VPN optimization in a large-scale outdoor town scenario. To assess scalability, the original dataset is cropped into subregions with approximately 200, 400, 600, 800, and 1000 viewpoints (VC). The full town comprises 637 buildings within a block of $976\text{m} \times 893\text{m}$, totaling $80,549\text{m}^2$ of built area with an average footprint of 126m^2 . All parameters are fixed across experiments: VF resolution 2.0m, minimum scan radius 0.6m, maximum scan radius 600m, wall density 5m, and overlap ratio 0.3. This setup enables controlled scaling of VC while preserving realistic geometric complexity (table 5 and figs. 20 and 21).

Table 5 shows that memory usage grows nearly linearly with

VC, from 265MB at VC=200 to 5703MB at VC=1000 (a 21.5 \times increase). On a logarithmic scale, $\log_{10}(\text{RAM})$ increases by only 1.34 across the five settings, confirming mild scaling. By contrast, runtime grows superlinearly: from 2.68min at VC=200 to 1505.36min (≈ 25 hrs) at VC=1000, corresponding to a 561 \times increase. The logarithmic slope of $\log_{10}(\text{Runtime})$ averages 0.55 per VC-doubling, consistent with near-quadratic growth due to pairwise visibility evaluations.

Comparative scaling further highlights runtime as the primary bottleneck. For example, increasing VC from 400 to 800 multiplies memory by 3.2 \times (1018 to 3268 MB) but runtime by 31 \times (30.8 to 967.5 min). Despite this, memory remains within typical workstation limits (<6 GB), while runtime quickly becomes prohibitive at

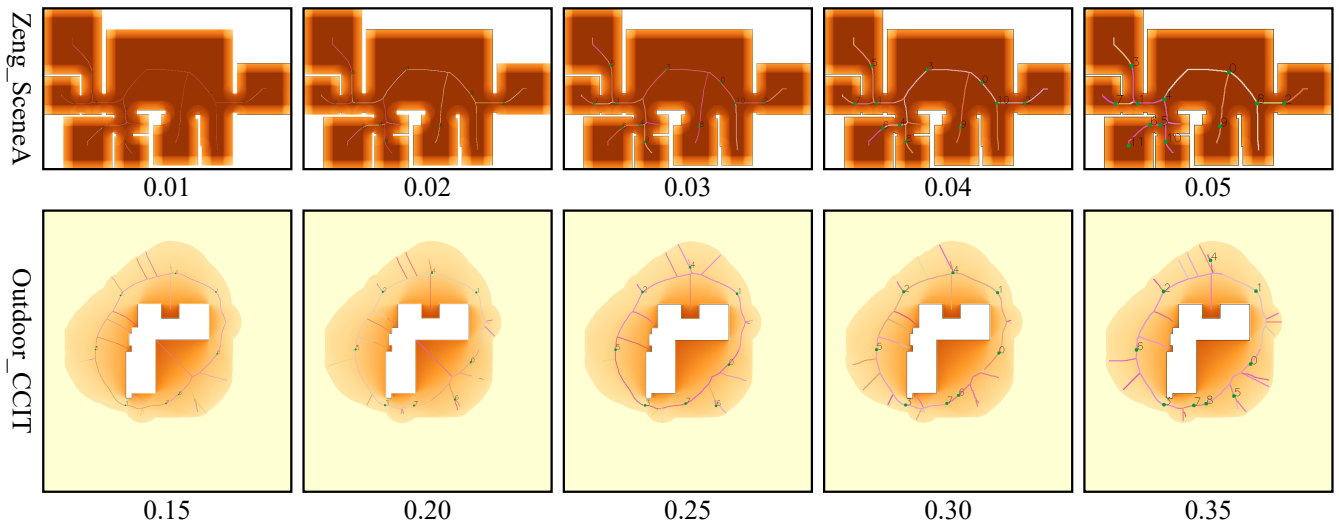


Figure 19: Effect of resolution on VF-based viewpoint planning. The top row shows results for the indoor Zeng_SceneA case with resolutions ranging from 0.01 m to 0.05 m, and the bottom row shows results for the outdoor CCIT case with resolutions ranging from 0.15 m to 0.35 m. The VF's topological nature results in low sensitivity to resolution changes in both scenarios.

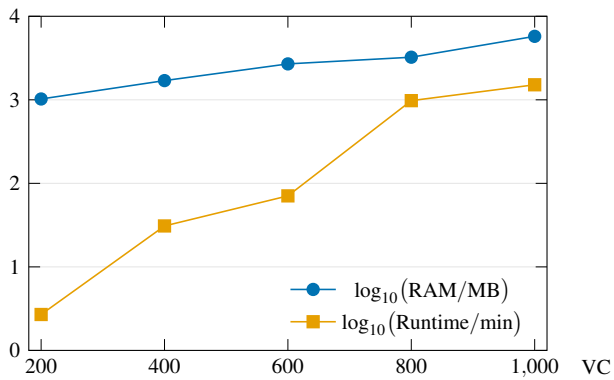
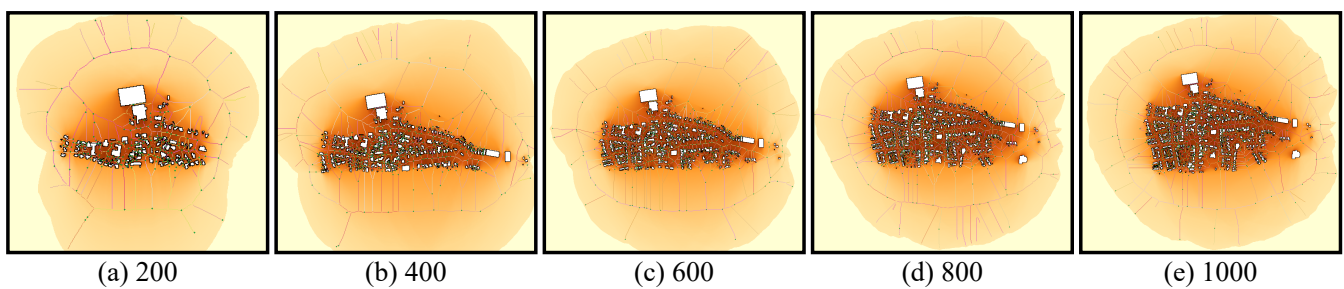


Figure 20: Log-scaled memory and runtime vs. viewpoint count (VC) for the town dataset. Both $\log_{10}(\text{RAM/MB})$ and $\log_{10}(\text{Runtime/min})$ increase with VC, highlighting time and space complexity trends.

$VC \geq 1000$. These findings demonstrate that VF-based optimization is tractable at city-block scales but requires algorithmic acceleration—such as hierarchical visibility pruning, parallel BSP-tree queries, GPU acceleration, or distributed processing—for deployment in very large regions.

Table 5: Time and space complexity analysis for the town dataset with cropped regions.

VC	VC (Actual)	RAM [MB]	\log_{10} (RAM/MB)	Runtime [min]	\log_{10} (Runtime/min)	Building Count	Block Area [m × m]	Total Building Area [m ²]	Avg Area Building [m ²]
200	197	265	2.42	2.68	0.43	132	520 × 292	22059	167
400	409	1018	3.01	30.84	1.49	274	965 × 366	39359	144
600	553	1683	3.23	71.02	1.85	379	976 × 441	49428	130
800	757	3268	3.51	967.48	2.99	529	976 × 588	67448	128
1000	954	5703	3.76	1505.36	3.18	637	976 × 893	80549	126

**Figure 21:** Optimized VPNs for the town dataset. Panels (a) to (e) show cropped regions for viewpoint counts VC = 200, 400, 600, 800, 1000 respectively.