

Article

HiFiRadio: High-Fidelity Radio Map Reconstruction for 3D Real-World Scenes

Ke Liao, Mengyu Ma *, Luo Chen, Yifan Zhang and Ning Jing

College of Electronic Science and Technology, National University of Defense Technology, Changsha 410073, China; liaoke19@nudt.edu.cn (K.L.); luochen@nudt.edu.cn (L.C.)

* Correspondence: mamengyu10@nudt.edu.cn

Abstract

The reconstruction of high-fidelity radio maps is pivotal for wireless network planning but remains challenging due to the tension between physical accuracy and computational efficiency. We propose HiFiRadio, a novel framework that achieves a breakthrough in this balance by integrating centimeter-resolution 3D environmental meshes with semantic-aware propagation modeling. At its core, HiFiRadio introduces a semantic-enhanced 3D indexing structure that efficiently manages complex terrain data, enabling real-time classification of signal paths into line-of-sight, non-line-of-sight, and vegetation-obstructed categories. This classification directly guides a hybrid propagation model, which dynamically applies dedicated loss calculations for buildings and foliage, grounded in physical principles. Extensive experiments demonstrate that HiFiRadio attains an accuracy comparable to commercial ray-tracing tools while being orders of magnitude faster. It also significantly outperforms existing learning-based baselines in both accuracy and scalability, a claim further validated by field measurements. By making high-fidelity, real-time radio map reconstruction practical for large-scale scenes, HiFiRadio establishes a new state of the art with immediate applications in network planning, UAV pathing, and dynamic spectrum access.

Keywords: radio map reconstruction; 3D semantic indexing; ray tracing; path loss modeling; oblique photogrammetry

1. Introduction

The deployment and evolution of next-generation wireless networks (BSG/6G) critically hinge on high-precision sensing and exploitation of the radio environment. Radio maps—representations that characterize the spatio-temporal—spectral distribution of the power spectral density (PSD)—provide indispensable environmental awareness for network planning, dynamic spectrum access, resource allocation, and unmanned aerial vehicle (UAV) path planning [1,2]. Reconstructing high-resolution radio maps that faithfully capture the intricate three-dimensional (3D) characteristics of real-world scenes, however, remains a formidable challenge when only sparse sensor measurements are available.

Model-driven approaches conventionally operate on 2-D planes or invoke strongly simplified 3D terrain abstractions, e.g., assuming a uniform path-loss exponent. These methods rely heavily on empirical propagation models such as the log-distance path-loss model [3]. Their principal limitation is the inability to account for the heterogeneous geometric and material attributes of fine-grained 3D terrain. For instance, signal attenuation through foliage differs fundamentally from that through reinforced-concrete structures;



Academic Editors: Stefano Scanzio, Ramez Daoud, Jetmir Haxhibeqiri and Pedro Santos

Received: 8 November 2025

Revised: 19 December 2025

Accepted: 30 December 2025

Published: 12 January 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

yet traditional models subsume such distinctions into coarse parameters, yielding marked accuracy degradation in topographically complex urban or mountainous regions [4].

Conversely, recent data-driven paradigms leverage machine learning techniques (e.g., CNNs [5], GANs [6]) to learn propagation characteristics directly from data, circumventing explicit modeling. Their efficacy, however, is tightly coupled with the quantity and quality of training data. Models trained in one geographic domain exhibit limited generalizability when deployed in regions with markedly different terrain features, thereby constraining practical applicability. Moreover, the data volume required to span high-fidelity 3D scenes—whether via simulation or in situ acquisition—becomes prohibitive, rendering both training and inference computationally infeasible under real-time constraints.

To transcend these bottlenecks, we posit that high-precision 3D scene models are essential. Recent advances in oblique photogrammetry enable the automated generation of centimeter-resolution 3D meshes that accurately capture architectural facades, windows, and tree canopies [7]. Integrating such detailed 3D data into radio-map reconstruction promises a qualitative leap in fidelity. Yet the sheer granularity of these meshes induces exponential data growth, rendering conventional physics-based simulations (e.g., ray tracing) computationally prohibitive and incompatible with rapid reconstruction.

Addressing this tension, we propose HiFiRadio, a hybrid model- and data-driven radio-map reconstruction framework tailored to high-fidelity 3D scenes. Rather than sacrificing fidelity for tractability, HiFiRadio attacks the computational bottleneck through intelligent data organization and management. Specifically, the framework introduces a spatial data-indexing structure (R-tree) to efficiently store, query, and manipulate massive oblique-photogrammetry meshes. Fine-grained semantic segmentation further differentiates objects—e.g., penetrable vegetation versus impenetrable buildings—enabling physics-aware propagation modeling while achieving orders-of-magnitude acceleration.

Our principal contributions are as follows:

HiFiRadio framework: We present a radio-map reconstruction architecture that synergistically integrates classical propagation models with high-performance 3D spatial indexing. By exploiting the precision gains of detailed terrain data while leveraging spatial indexing for rapid query and computation, the framework achieves extremely high efficiency.

Fine-Grained Semantic 3D Indexing: We devise an R-tree-based indexing scheme for 3D meshes augmented with per-object semantic labels. This segmentation distinguishes between penetrable (vegetation) and impenetrable (buildings) entities, supplying the propagation engine with physically grounded priors.

Semantic-Aware Radio-Map Computation: Guided by the semantic index, the framework dynamically assigns path-loss models to individual segments of the propagation path. Vegetation attenuation models are invoked for foliage, whereas rigorous wall-penetration loss models are applied for building materials, yielding markedly more accurate aggregate path-loss estimates.

Comprehensive Experimental Validation: Using ray-tracing simulations grounded in real-world 3D environments, we benchmark HiFiRadio against leading model- and data-driven baselines, demonstrating superior performance. Furthermore, in situ RF measurements collected at the corresponding physical sites corroborate the framework's operational reliability and efficacy.

2. Related Work

The reconstruction of high-fidelity radio maps lies at the intersection of wireless communications and geoinformation science. Existing research can be broadly categorized into two streams: one focuses on refining propagation models or data algorithms within wire-

less communications, while the other is dedicated to detailed 3D environmental modeling and analysis within geoinformatics. However, both streams exhibit significant limitations when dealing with complex 3D scenes.

2.1. Radio Map Reconstruction in Wireless Communications

Existing radio map reconstruction approaches are broadly categorized as model-driven and data-driven. The COST-231 model [8] simplifies 3D urban scenes into a single equivalent rooftop height for rapid path loss estimation. The WINNER II model [9] categorizes environments into predefined terrain types with fixed correction factors for efficient signal prediction. The 3GPP TR 38.901 [10] model reduces 3D maps into key parameters like height and clutter type for wide-frequency coverage. These model-driven methods overlook fine geometric details, leading to poor accuracy in complex environments. In contrast, data-driven methods learn terrain-to-signal mappings directly from data. RadioUNet [11] treats radio map reconstruction as an image translation task for fast inference but requires high-resolution inputs and generalizes poorly. TiRE-GAN [12] incorporates a radio-depth prior into a generative adversarial network to maintain sharp predictions under sparse sampling. RadioGAT [13] constructs frequency-dependent graphs from sparse measurements and uses graph attention for semi-supervised learning. However, these data-driven approaches require costly high-resolution 3D data and incur high computational loads, making them impractical for large-scale real-time applications.

2.2. Three-Dimensional Scene Modeling and Spatial Analysis in Geoinformatics

The reconstruction of high-fidelity radio maps relies on precise 3D representation of the physical environment, for which geoinformation science provides a crucial technological foundation. In 3D scene modeling, the analysis of spatial relationships is key to understanding signal propagation paths, and viewshed analysis, as a core analytical tool in this field, is specifically designed to determine inter-visibility between points in 3D space, providing a geometric foundation for identifying line-of-sight and non-line-of-sight conditions in signal propagation.

In 3D scene modeling, techniques such as oblique photogrammetry enable the generation of centimeter-resolution 3D meshes [7], accurately capturing detailed features like building facades and vegetation canopies. Nevertheless, these detailed geometric models lack descriptions of electromagnetic properties. Although semantic modeling standards like CityGML [14] introduce object-level categorical information, their level of detail typically remains at the “building” level, insufficient to support per-triangle physical property annotation required for electromagnetic propagation modeling.

In spatial analysis techniques, ray tracing serves as the gold standard for 3D electromagnetic field simulation, capable of accurately modeling signal propagation paths, but its computational overhead is prohibitively high, generally limiting its application to simplified or coarse terrain data. To maintain accuracy while improving efficiency, the geoinformatics field has developed a series of optimized methods. R-tree and its variants [15] are widely used to manage massive 3D datasets, supporting efficient spatial queries through hierarchical spatial indexing. In viewshed analysis, several efficient algorithms have been proposed, including GPU-parallel real-time viewshed analysis [16], multi-point viewshed analysis [17], TIN surface decomposition-based viewshed analysis [18], and projection-based viewshed analysis with 3D refinement [19,20]. These methods have achieved significant efficiency improvements in determining inter-visibility between points, but their results are typically confined to binary “visible” or “not visible” judgments. They have not been integrated with complex signal propagation

physics models (e.g., penetration loss, diffraction loss), thus preventing their direct use in calculating continuous signal attenuation.

Table 1 systematically contrasts HiFiRadio with existing state-of-the-art methods. Traditional model-driven approaches (e.g., COST-231, 3GPP TR 38.901) rely on highly simplified environmental parameters and cannot capture fine-grained geometric and material variations. Data-driven methods (e.g., RadioUNet, RadioGAT) learn directly from data but suffer from limited generalizability and high computational costs during training and inference. In contrast, HiFiRadio introduces a semantic-enhanced 3D spatial index and a physics-aware path classification mechanism, achieving near ray-tracing accuracy and data-driven efficiency while maintaining physical interpretability. This enables a breakthrough in the longstanding accuracy–efficiency trade-off in radio map reconstruction.

In summary, although the geoinformatics field provides high-precision environmental data foundations and efficient spatial analysis tools, existing detailed terrain data has not been deeply integrated with efficient spatial analysis methods possessing physical awareness.

Table 1. Comparison of different radio map reconstruction methods.

Category	Representative Methods	Environmental Modeling Granularity	Physical Mechanism	Computational Efficiency	Key Limitations
Model-Driven	COST-231, 3GPP TR 38.901	2D/simplified 3D parameters	Empirical formulas	High	Low accuracy in complex 3D scenes
Data-Driven	RadioUNet, RadioGAT	Depends on input resolution	Black-box learning	Medium (high training cost)	Poor generalization, requires massive labeled data
Physics Simulation	Wireless Insite (Ray Tracing)	Centimeter-level 3D mesh	Full physics (reflection, diffraction, transmission)	Very low	Computationally prohibitive for real-time/large-scale use
Hybrid (Proposed)	HiFiRadio	Cm-level 3D mesh + semantic labels	Path classification + dedicated physical models	High (real-time feasible)	Depends on segmentation accuracy; no multipath reflection modeling

3. Methodology

In complex 3D urban environments, intricate topographic and morphologic features exert a decisive influence on radio-wave propagation. Empirical models that parameterize these terrain characteristics inevitably yield coarse approximations and are therefore inadequate for high-precision electromagnetic situation assessment. By leveraging oblique photogrammetry-based 3D modeling, we treat the acquired city-scale geometry not merely as a visualization backdrop, but as a primary input to the construction of the electromagnetic spectrum map.

Within this framework in Figure 1, the high-fidelity 3D mesh is first processed through semantic segmentation to classify each triangular face into categories (e.g., ‘building’, ‘vegetation’). These labels are then binarized into a penetrability attribute (impenetrable vs. penetrable), thereby enriching the raw geometry with electromagnetic semantics. This semantically enriched mesh is subsequently organized into a specialized 3D R-tree index, which efficiently manages the massive dataset and enables real-time spatial queries coupled with physical properties.

This semantic-enhanced index serves as the computational backbone for our rapid path-category classification algorithm. For any signal trajectory between a transmitter and a receiver, the framework queries the index to identify all intersecting obstacles and analyzes their penetrability attributes. Based on this, it decomposes the path into one of three mutually exclusive classes: line-of-sight (LOS), non-line-of-sight (NLOS), and obstructed line-of-sight (OLOS). Each identified propagation class is subsequently associated with a dedicated, physics-grounded channel model for path-loss prediction, thereby enabling the accurate and efficient generation of electromagnetic spectrum maps in realistic 3D urban settings.




Figure 1. System overview diagram of HiFiRadio.

3.1. Semantic-Enhanced 3D R-Tree Indexing

The foundation of our framework is a high-fidelity 3D mesh model of the environment, typically obtained through oblique photogrammetry. These models, often stored in formats like OBJ or PLY, offer centimeter-level geometric detail, accurately capturing features such as building facades, windows, and vegetation canopies, which are crucial for precise propagation modeling. However, the raw geometric detail alone is insufficient, as different materials exhibit vastly different electromagnetic properties. For instance, the penetration loss through vegetation, while non-negligible, is typically significantly lower (e.g., 10 dB) than through reinforced concrete structures (e.g., 20 dB). Ignoring this distinction leads to substantial errors in path-loss estimation.

To incorporate this physical prior, we first perform semantic segmentation on the input 3D mesh. We employ the MeshCNN [21] model, which is pre-trained on the ShapeNet dataset and fine-tuned on a custom dataset of urban oblique photogrammetry meshes labeled with ‘building’, ‘vegetation’, ‘ground’, and ‘other’ categories. The training set comprises 50 urban scenes with over 10 million labeled triangles. The model achieves an overall segmentation accuracy of 92.3% and per-class IoU of 89.1% for ‘building’ and 85.7% for ‘vegetation’ on our validation set. Although misclassification (e.g., labeling a glass facade as vegetation) could affect path loss estimation, our experiments in Section 4.5 demonstrate that the framework covers robust-to-moderate segmentation errors, with an average RMSE increase of less than 1.2 dB under simulated label noise of 10%. This preprocessing step classifies each triangular face in the mesh into categories such as ‘building’, ‘vegetation’, ‘ground’, etc. For the core purpose of propagation path classification and loss calculation, these labels are subsequently binarized into a penetrability attribute: impenetrable (e.g., buildings, terrain, encoded as ‘0’) and penetrable (e.g., foliage, encoded as ‘1’). As shown in Figure 2, the green triangles are assigned an attribute of 1 (representing vegetation), while the yellow triangles are assigned an attribute of 0 (representing

buildings). This binary classification provides a critical physically grounded prior for the propagation engine.

Managing and performing efficient spatial queries on the resulting massive, semantically enriched mesh (containing often hundreds of millions of triangles) is computationally challenging. To overcome this, we design a semantic-enhanced 3D R-tree spatial index. The core idea is to hierarchically organize the numerous triangular faces using their Minimum Bounding Rectangles (MBRs), enabling efficient range queries and nearest-neighbor searches necessary for ray-path analysis.

The construction algorithm for our Semantic R-Tree is as follows:

Input Preparation: For each triangular face i in the semantically segmented mesh, we form a data entry $e_i = (MBR_i, attr_i)$. Here, MBR_i is the 3D axis-aligned bounding box enclosing the triangle, computed from the minimum and maximum x , y , and z coordinates of its three vertices. $attr_i$ is a compact data structure storing the triangle's vertex coordinates and its binary penetrability attribute.

Leaf Node Formation: The algorithm starts by grouping a predefined number M (e.g., $M = 50$) of these data entries into leaf nodes. The MBR of a leaf node is the tightest bounding box that contains the MBRs of all its constituent data entries.

Hierarchical Clustering: These leaf nodes are then treated as objects themselves, grouped into higher-level nodes based on the proximity and overlap of their MBRs. This process continues recursively until all entries are contained under a single root node. The grouping at each level aims to minimize the overlap and area of the resulting node MBRs, following the typical R*-tree insertion strategy which employs a combination of least area enlargement and least overlap enlargement criteria for choosing subtrees during insertion.

Node Structure: Each node (both internal and leaf) in the tree contains: (a) a list of child nodes (for internal nodes) or data entries (for leaf nodes); (b) the MBR that covers all objects in its subtree; and (c) for leaf nodes, pointers to the actual geometric and semantic data of the triangles (stored externally for efficiency, with the $attr_i$ in the leaf entry serving as a key or storing minimal essential data).




Figure 2. Illustration of the R-tree architecture. Green triangles represent vegetation (penetrability attribute 1), yellow triangles represent buildings (penetrability attribute 0).

The construction algorithm for our Semantic R-Tree is detailed in Algorithm 1.

Complexity Analysis: The construction of the R-tree index has a time complexity of $O(n \log n)$, where n is the number of triangles. Each ray query exhibits an average time complexity of $O(\log n + k)$, where k is the number of actually intersected triangles—which is significantly lower than the $O(n)$ required for brute-force traversal. **Parameter Selection:** The leaf node capacity $M = 50$ was chosen empirically to balance tree depth and query

performance. The R*-tree insertion strategy is adopted to minimize bounding box overlap and area, thereby optimizing query efficiency.

Algorithm 1 Constructing the semantic-enhanced 3D R-tree index.

Input: A set of triangular meshes F , where each mesh f_i is associated with its vertex coordinates and a penetrability attribute $a_i \in \{0, 1\}$.

Output: A fine-grained semantic 3D R-tree index model.

```

1 for all triangular mesh  $f$  in  $F$  do
2   Calculate the MBR of  $f$  and form a data entry  $e = (MBR, attr)$ , where  $attr$  stores the
   vertex data and  $a_i$ .
3   Find the leaf node  $N$  in the current tree whose MBR requires the minimum expansion
   to incorporate  $e$ .MBR.
4   if not exceedsThreshold( $N, e$ ) then
5     // Check if adding  $e$  exceeds node capacity  $M$ .
6     Insert entry  $e$  into node  $N$ .
7     adjustAABB( $N$ ) // Adjust the MBR of node  $N$  to encompass all its entries.
8   else
9      $(N, N_1) \leftarrow \text{splitNode}(N, e)$  // Split the overfull node  $N$  into two nodes  $N$  and  $N_1$ .
10    while  $N$  is not the root do
11       $L \leftarrow \text{findParentNode}(N)$  //  $L$  is the parent node of  $N$ .
12      adjustAABB( $L$ ) // Adjust the MBR of parent node  $L$  to fit its children.
13      if exceedsMaxStorageConstraint( $L$ ) then
14        // Check if  $L$  is now overfull due to the split.
15         $(L, L_1) \leftarrow \text{splitNode}(L, N_1)$  // Split the overfull parent node  $L$ .
16         $N \leftarrow L$ 
17         $N_1 \leftarrow L_1$ 
18 return root // The root of the constructed R-tree.
```

Given the real-world characteristics of oblique photogrammetry data—where triangles are not uniformly distributed but cluster densely on building surfaces and vegetation—the R-tree effectively groups geographically proximate and semantically similar objects (due to the spatial coherence of semantic labels). This structure allows the propagation engine to rapidly retrieve all potential obstacles (triangles) intersecting a given radio propagation path (e.g., a ray between transmitter and receiver) by traversing the tree from the root, pruning entire branches whose MBR does not intersect the query volume. The concomitant retrieval of the penetrability attribute for each intersecting triangle enables immediate, physics-aware decision-making during path loss calculation.

3.2. Signal Occlusion Detection Algorithm Based on Index

Upon establishing the semantic 3D R-tree index, the next critical step is to accurately and efficiently determine the signal occlusion between a propagation path and the detailed 3D terrain. This occlusion detection is fundamental for classifying propagation paths and calculating aggregate path loss. To address this, we designed a two-stage occlusion detection algorithm that leverages the index structure for rapid culling of non-occluding geometry, followed by precise occlusion validation.

The algorithm operates as follows: For a given signal path—a ray from a Transmitter (Tx) to a Receiver (Rx)—the R-tree index is first traversed to retrieve all leaf nodes (and their contained triangular terrain facets) whose MBRs are intersected by the ray. This broad-phase occlusion culling efficiently eliminates large portions of the scene that do not occlude the signal. For each terrain facet within these candidate MBRs, a second, precise narrow-phase occlusion validation is performed to verify if the signal path is actually occluded (blocked) by the facet itself. A confirmed occlusion indicates that the signal path is obstructed by the terrain at that specific location. The two-stage process is detailed below.

Part I: Broad-Phase—Signal Occlusion Culling via MBR Testing

The first part of the algorithm performs a rapid screening to determine whether a signal path warrants further detailed occlusion checking with the terrain facets inside an MBR. We employ the efficient Slab [22] Method for this purpose. The core idea is to treat the AABB as the intersection of three pairs of parallel planes (slabs). The signal path's intersection intervals with each pair of slabs are computed independently, and the path is a candidate for occlusion if the overlap of these intervals is non-empty.

Algorithm description:

As shown in Figure 3, given a signal path defined by its origin point O (the transmitter) and a normalized direction vector D (towards the receiver), and an AABB defined by its minimum and maximum extents $(B_{\min}^x, B_{\min}^y, B_{\min}^z)$ and $(B_{\max}^x, B_{\max}^y, B_{\max}^z)$:




Figure 3. Slab Method for broad-phase occlusion culling.

For each axis $i \in \{x, y, z\}$:

Calculate the distances to the two slabs along axis i :

$$t_{i\min} = \frac{B_{\min}^i - O^i}{D^i}, \quad t_{i\max} = \frac{B_{\max}^i - O^i}{D^i} \quad (1)$$

To ensure numerical stability and handle rays pointing in the negative direction, the near and far intersection parameters for each axis are defined as follows:

$$t_{inear} = \min(t_{i\min}, t_{i\max}), \quad t_{ifar} = \max(t_{i\min}, t_{i\max}) \quad (2)$$

Find the overall entry and exit points of the signal path with the union of the slabs:

$$t_{\text{enter}} = \max(t_{xnear}, t_{ynear}, t_{znear}) \\ t_{\text{exit}} = \min(t_{xfar}, t_{yfar}, t_{zfar}) \quad (3)$$

Occlusion Culling Condition: The signal path is a candidate for occlusion by the contents of the AABB if and only if

$$t_{\text{enter}} \leq t_{\text{exit}} \text{ and } t_{\text{enter}} \geq 0 \quad (4)$$

The first condition ensures the intervals overlap, meaning the signal path enters the box before it exits. The second condition ensures the potential occlusion occurs in the positive direction of the path (i.e., not behind the transmitter).

This method is computationally inexpensive, requiring only comparisons and divisions, making it ideal for quickly culling non-occluding nodes during the R-tree traversal.

Part II: Narrow-phase—precise signal occlusion validation with terrain facets

If a signal path is identified as a candidate for occlusion by an MBR (broad-phase), a precise check is required for each terrain facet within that MBR to validate the occlusion. For this narrow-phase test, we utilize the fast and widely adopted Moller-Trumbore [23] algorithm.

Algorithm description:

Given a signal path defined by origin O (transmitter) and direction D , and a terrain facet defined by its three vertices V_0, V_1, V_2 :

Find two edge vectors of the triangle:

$$\begin{aligned} e_1 &= V_1 - V_0 \\ e_2 &= V_2 - V_0 \end{aligned} \quad (5)$$

Compute the determinant which determines whether the signal path is parallel to the facet:

$$\begin{aligned} P &= D \times e_2 \\ \det &= e_1 \cdot P \end{aligned} \quad (6)$$

Calculate the vector from the facet vertex V_0 to the transmitter O :

$$T = O - V_0 \quad (7)$$

Compute the first barycentric coordinate u :

$$u = (T \cdot P) / \det \quad (8)$$

If $u < 0$ or $u > 1$, the point of occlusion lies outside the facet.

Compute the second barycentric coordinate v :

$$\begin{aligned} Q &= T \times e_1 \\ v &= (D \cdot Q) / \det \end{aligned} \quad (9)$$

If $v < 0$ or $u + v > 1$, the point of occlusion lies outside the facet.

Occlusion Condition: The signal path is occluded by the terrain facet if and only if all the following conditions are met:

$$\begin{aligned} 0 &\leq u \leq 1, \\ 0 &\leq v \leq 1, \\ u + v &\leq 1 \end{aligned} \quad (10)$$

The distance t from the transmitter to the occlusion point is given by

$$t = (e_2 \cdot Q) / \det \quad (11)$$

This distance t must also be positive for a valid occlusion ahead of the transmitter.

Input and Output for Both Stages:

Input: A signal path (defined by origin O and direction D) and a bounding volume (AABB for Part I, a terrain facet for Part II).

Output: A Boolean result indicating whether occlusion is detected (True/False). For Part II, the barycentric coordinates and occlusion distance t are also valuable outputs for subsequent path loss calculations (e.g., penetration depth).

By combining the efficient occlusion culling of the Slab Method with the precise occlusion validation of the Moller-Trumbore algorithm, our framework achieves high-performance signal occlusion detection, which is paramount for real-time radio map reconstruction in high-fidelity 3D scenes.

3.3. Radio Map Construction Based on Multi-Path Classification

Leveraging the semantic 3D terrain index and the efficient signal occlusion detection method, we can accurately classify the propagation path of a signal in complex terrain into three distinct categories: Line-of-Sight (LOS), Non-Line-of-Sight (NLOS), and Obstructed-Line-of-Sight (OLOS). As shown in Figure 4. This classification is the cornerstone for applying appropriate and physically grounded path loss models. The core classification logic is formalized in Algorithm 2.

Based on the classification result from Algorithm 2, a dedicated path loss model is applied for each category to calculate the total loss PL (in dB) for the link.

Algorithm 2 Path category classification

Input: A ray defined by its origin T_x (transmitter) and direction D (towards receiver R_x); the semantic 3D R-tree index RTree

Output: Path category category; A list of intersected triangles obstacles; The total vegetation penetration depth veg_depth (for OLOS)

```

1 obstacles  $\leftarrow$  [] // Initialize an empty list to store intersected triangles
2 veg_depth  $\leftarrow$  0.0
3 candidate_triangles  $\leftarrow$  RTree.rayQuery( $T_x, D$ ) // Query the R-tree to get all triangles
  potentially occluding the signal path
4 Step 1: Check for any occlusion
5 if candidate_triangles is empty then
6   category  $\leftarrow$  "LOS" // No obstacles found
7   return category, obstacles, veg_depth
8 Step 2: Check semantics of obstacles
9 for all triangle tri in candidate_triangles do
10  if tri.penetration == 0 then // Impenetrable (e.g., building) found
11    category  $\leftarrow$  "NLOS"
12    obstacles.append(tri)
13    return category, obstacles, veg_depth // NLOS dominates, return early
14  else
15    obstacles.append(tri) // Add penetrable obstacle (vegetation)
16 Step 3: If only penetrable obstacles were found
17 if obstacles is not empty then
18   veg_depth  $\leftarrow$  calculateVegetationDepth( $T_x, D, obstacles$ ) // See Algorithm 3
19   category  $\leftarrow$  "OLOS"
20 return category, obstacles, veg_depth

```




Figure 4. Illustration of propagation path categories (e.g., LOS, NLOS, OLOS).

1. Line-of-Sight (LOS) Path Loss

For unobstructed paths, the signal attenuation is modeled by the Free-Space Path Loss (FSPL) model, which predicts the loss incurred by the spreading of the wavefront in free space. The formula is given by

$$PL_{LOS} = 20 \log_{10}(d) + 20 \log_{10}(f) - 147.55 \quad (12)$$

where d is the distance between the transmitter and receiver in meters, and f is the frequency in Hz. This model serves as the fundamental baseline for signal attenuation.

2. Non-Line-of-Sight (NLOS) Path Loss

When the path is blocked by impenetrable obstacles like buildings, we adopt the empirically derived 3GPP Urban Macro (UMa) NLOS path loss model from TR 38.901 [10]. This model accounts for additional losses due to diffraction and reflection over rooftops and around structures.

$$\begin{aligned} PL_{NLOS} &= \max(PL_{LOS}, PL'_{NLOS}) \\ PL'_{NLOS} &= 36.7 \log_{10}(d) + 22.7 + 26 \log_{10}(f) \\ &\quad - 0.3(h_{UT} - 1.5) \end{aligned} \quad (13)$$

where d is the 3D distance (in meters), f is the carrier frequency in GHz, and h_{UT} is the user terminal height in meters. The max operation ensures the NLOS loss is never less than the free-space loss.

3. Obstructed Line-of-Sight (OLOS) Path Loss

OLOS paths are obstructed solely by penetrable vegetation. We model the total path loss as the sum of the free-space loss and an additional vegetation attenuation loss (L_{veg}) based on the ITU-R P.833 [24] recommendation. The vegetation loss is a function of the carrier frequency f and the total depth d_{veg} (in meters) that the signal travels through the foliage.

$$\begin{aligned} PL_{OLOS} &= PL_{LOS} + L_{veg} \\ L_{veg} &= 0.2 \cdot f^{0.3} \cdot d_{veg}^{0.6} \end{aligned} \quad (14)$$

where f is the frequency in MHz.

The key to calculating L_{veg} is accurately determining the total vegetation penetration depth d_{veg} . We designed Algorithm 3 for this purpose, leveraging the semantic index to identify and sum the penetration segments.

Algorithm 3 calculateVegetationDepth: Computation of Total Vegetation Penetration Depth

Input: Ray origin T_x , direction D ; List of intersected (vegetation) triangles obstacles

Output: Total vegetation penetration depth total_depth

```

1 intersection_points ← []
2 for all triangle tri in obstacles do
3   point ← computeRayTriangleIntersection( $T_x$ ,  $D$ , tri) // Using Möller-Trumbore
4 sorted_points ← sort(intersection_points, key=lambda p: distance( $T_x$ , p)) // Sort all
   intersection points by their distance from the transmitter  $T_x$ 
5 total_depth ← 0.0
6 for i ← 0 to length(sorted_points) - 1 do
7    $p_1$  ← sorted_points[i]
8    $p_2$  ← sorted_points[i + 1] // For the last point, if the number is odd, it is ignored,
   as it represents an exit without a subsequent entry.
9   segment_depth ← distance( $p_1$ ,  $p_2$ )
10 total_depth ← total_depth + segment_depth
11 return total_depth
```

Finally, the received power P_r (dBm) at any receiver location is computed as follows:

$$P_r = P_t - PL \quad (15)$$

where P_t is the transmitted power (dBm). By executing this process for a dense grid of receiver locations, a high-fidelity, physics-aware radio map of the entire area is reconstructed.

4. Experimental Validation

4.1. Accuracy Comparison of Different Methods

To quantitatively evaluate the performance of the proposed HiFiRadio framework, we conducted a comparative experiment against leading model-driven and data-driven baselines using a realistic urban environment.

Dataset and Experimental Setup:

The experiment was conducted on a $250\text{ m} \times 250\text{ m}$ area extracted from the open-source SUM_Helsinki dataset, as illustrated in Figure 5. This dataset provides a centimeter-resolution 3D mesh with precise pre-computed semantic segmentation labels (e.g., building, vegetation, ground), making it ideally suited for our method. A transmitter (Tx) was placed at coordinates (100, 100, 80) m with a transmission power of 40 dBm and a frequency of 3 GHz. The objective was to reconstruct the received signal power across a $200\text{ m} \times 200\text{ m}$ area centered on the Tx, specifically at a receiver height of 1.5 m, simulating a typical handheld device level. The ground truth was generated by importing the detailed 3D mesh and its semantics into the commercial ray-tracing simulator Wireless Insite, with identical Tx parameters, and exporting the simulation result at 1.5 m height.




Figure 5. Semantic segmentation results on SUM_Helsinki dataset.

Compared Methods:

RadioGAT: A state-of-the-art data-driven method using Graph Attention Networks.
 RadioUNet: A prominent deep learning approach based on U-Net for radiomap estimation.
 COST-231 Hata Model: A widely used empirical model-driven method for urban areas.
 HiFiRadio (Ours): The proposed hybrid framework.

Evaluation Metrics:

We employed two metrics for a comprehensive evaluation:

Root Mean Square Error (RMSE): Measures the average magnitude of the power estimation error in dB.

Correct Detection Zone Ratio (CDZR)/False Alarm Zone Ratio (FAZR): A threshold-based metric. A prediction is considered correct if its absolute error is within 7 dB. CDZR is the ratio of the number of correctly predicted pixels to the total number of pixels. FAZR is the ratio of the number of incorrectly predicted pixels to the total number of pixels. A higher CDZR and a lower FAZR indicate better performance.

Results and Analysis:

The reconstructed radio maps and the ground truth are visually presented in Figure 6. Qualitatively, our method most closely resembles the ground truth, accurately capturing

the complex shadowing effects caused by buildings and the subtle attenuation through vegetation canopies.

The COST-231 Hata model produces a smooth, radially symmetric pattern that completely fails to capture the intricate interactions with the 3D geometry. It is unable to model NLOS conditions behind buildings or OLOS conditions within parks, leading to massive errors, as confirmed by its high RMSE and poor CDZR/FAZR.

The RadioUNet model, trained primarily on building footprints and lacking explicit 3D awareness, struggles with the height disparity between the Tx (80 m) and the receiver plane (1.5 m). It fails to accurately reconstruct the deep shadows cast by tall buildings, resulting in significant underestimation of path loss in those areas.

The RadioGAT model, which processes the scene in partitioned blocks, exhibits visible artifacts and power discontinuities at the boundaries between these blocks. This tiling effect is a direct result of its limited receptive field and leads to unrealistic predictions, which is reflected in its suboptimal metrics.

In contrast, HiFiRadio explicitly models the 3D propagation paths by leveraging the detailed geometry and semantics for precise signal occlusion detection. It correctly identifies LOS, NLOS (behind buildings), and OLOS (through vegetation) conditions, applying physically grounded models for each. This results in a highly accurate and spatially consistent reconstruction.

The quantitative results, summarized in Figure 6, strongly support the visual observations. Our method achieves the lowest RMSE and the highest CDZR, significantly outperforming all other baselines. This demonstrates the critical importance of incorporating fine-grained 3D semantic geometry into the radio map reconstruction process for achieving high fidelity.




Figure 6. Cont.




Figure 6. Signal strength heatmaps and comparison results from various methods.

4.2. Efficiency Comparison of Different Methods

While accuracy is paramount, computational efficiency is equally critical for the practical deployment of radio map reconstruction systems, especially in large-scale scenarios. This subsection evaluates the time efficiency of the proposed HiFiRadio framework against other methods across varying reconstruction scales.

Dataset and experimental setup:

To assess scalability, we selected a larger $500 \text{ m} \times 500 \text{ m}$ area from the SUM_Helsinki dataset. A transmitter was placed at the center (250, 250, 80) m with the same parameters as those listed in Section 4.1 (3 GHz, 40 dBm). We measured the total computation time required by each method to reconstruct the received signal power map at 1.5m height for five increasingly large areas centered on the T_x : $50 \text{ m} \times 50 \text{ m}$, $100 \text{ m} \times 100 \text{ m}$, $150 \text{ m} \times 150 \text{ m}$, $200 \text{ m} \times 200 \text{ m}$, and $250 \text{ m} \times 250 \text{ m}$.

Compared Methods:

Wireless Insite (Ray Tracing): This high-fidelity commercial simulator is used as the accuracy benchmark in Section 4.1. It represents the upper bound of accuracy but also of computational cost.

RadioUNet (NET) & RadioGAT (GAT): Representative data-driven deep learning methods.

HiFiRadio (Ours): The proposed framework.


The COST-231 Hata model was excluded from this comparison due to its negligible computational cost and significantly inferior accuracy, which does not constitute a meaningful baseline for efficiency in this context.

Results and Analysis:

The detailed runtime results for all methods are presented in Table 2. The data is also visualized in Figure 7 to illustrate the scaling trends more intuitively.

Table 2. Runtime comparison (in seconds) for reconstructing radio maps of different sizes.

Area (m ²)	50 × 50	100 × 100	150 × 150	200 × 200	250 × 250
Wireless Insite	2355	8760	18,677	37,587	54,331
RadioGAT (GAT)	15.61	37.13	61.90	93.58	145.37
RadioUNet (NET)	0.78	3.18	7.10	12.65	20.33
HiFiRadio (Ours)	0.047	0.186	0.391	0.643	0.953

**Figure 7.** Execution time of various approaches across different problem scales.

The results lead to several key conclusions:

Prohibitive Cost of Ray Tracing: As expected, the precision of Wireless Insite comes at an enormous computational expense. Its runtime grows rapidly with the area size, taking over 15 h to simulate the 250 m × 250 m area. This renders it utterly impractical for real-time applications, despite its high accuracy.

Superior Efficiency of HiFiRadio: Our proposed HiFiRadio framework demonstrates exceptional computational efficiency. It is over 16 times faster than RadioUNet and over 150 times faster than RadioGAT at the 250 m × 250 m scale. Most notably, it is nearly 60,000 times faster than the ray-tracing benchmark for the same area. This immense speedup is a direct result of our hybrid approach, which replaces millions of computationally intensive physics-based ray-triangle intersection calculations (in pure ray tracing) with an intelligent combination of rapid spatial indexing and simplified—yet physics-aware—propagation models.

Favorable Scaling: The runtime of HiFiRadio scales approximately linearly with the number of receiver points (i.e., the area), as evidenced by the data. This predictable and favorable scaling behavior is crucial for applications requiring the reconstruction of city-scale radio maps.

Accuracy–Efficiency Trade-off: Crucially, this dramatic efficiency gain, as shown in Table 1 and Figure 7, is achieved without sacrificing accuracy. As demonstrated in Section 4.1, HiFiRadio significantly outperforms the data-driven methods (GAT, NET) in reconstruction fidelity and approaches the accuracy of ray tracing. Therefore, our framework successfully breaks the long-standing trade-off between accuracy and efficiency in radio map reconstruction.

In summary, the HiFiRadio framework establishes a new state of the art by achieving high-fidelity reconstruction that is orders of magnitude faster than both precise physical simulations and contemporary data-driven deep learning models, making it viable for large-scale, practical deployment.

4.3. Parallel Performance

To further demonstrate the practicality and scalability of the HiFiRadio framework for large-scale scenarios, we evaluated its parallel computing performance. The core computation in our method—performing ray queries against the indexed 3D mesh for a large number of receiver points—is highly parallelizable, as each point can be processed independently.

Dataset and experimental setup:

This experiment was conducted on Dataset-1 (the 250 m × 250 m area from SUM_Helsinki, comprising approximately 340,000 triangular faces, introduced in Section 4.1). We implemented a parallel version of the HiFiRadio reconstruction pipeline using OpenMP, focusing on the embarrassingly parallel task of computing received power for each point on a grid.

We varied two key parameters:

Number of Threads: 1, 2, 4, 8, 16, 32, 64.

Reconstruction Scale (Number of Receiver Points): 50 × 50 (2500 points), 100 × 100 (10,000 points), and 200 × 200 (40,000 points).

The transmitter location and parameters remained consistent with previous experiments. For each configuration, we recorded the total runtime. The Speedup Ratio(η), defined as $\eta = T_{\text{serial}} / T_{\text{parallel}}$, where T_{serial} is the runtime with 1 thread and T_{parallel} is the runtime with N threads, was calculated to quantify parallel efficiency.

Results and Analysis:

The runtime and calculated speedup for all configurations are presented in Table 3. The scaling behavior is further visualized in Figure 8.

Table 3. Parallel runtime (in seconds) and speedup ratio (in parentheses) for different receiver grid sizes and thread counts.

Number of Threads	50 × 50 (2500 Points)	100 × 100 (10,000 Points)	200 × 200 (40,000 Points)
1	0.61 (1.00)	2.47 (1.00)	9.35 (1.00)
2	0.35 (1.74)	1.42 (1.74)	5.31 (1.76)
4	0.19 (3.21)	0.81 (3.05)	3.05 (3.07)
8	0.10 (6.10)	0.34 (7.26)	1.38 (6.78)
16	0.05 (12.20)	0.17 (14.53)	0.64 (14.61)
32	0.03 (20.33)	0.11 (22.45)	0.41 (22.80)
64	0.05 (12.20)	0.19 (13.00)	0.81 (11.54)

The results lead to the following conclusions:

Near-Linear Scaling for Moderate Thread Counts: For up to 32 threads, the HiFiRadio framework demonstrates excellent parallel scalability. The speedup ratio increases almost linearly with the number of threads across all problem sizes. This is because the computational workload for each receiver point is significant enough to outweigh the overhead of parallel task scheduling, making efficient use of the available cores.

Performance Degradation Beyond Hardware Support: When the thread count (64) exceeds the number of available physical hardware threads (32 threads on our 16-core/32-thread test system), performance deteriorates. The runtime increases and the speedup ratio drops significantly (e.g., from 22.8× to 11.5× for the 200 × 200 case). This is due to the increased overhead of context switching and contention for shared resources among an excessive number of logical threads, confirming that optimal performance is achieved when the thread count aligns with the hardware’s parallel capabilities.




Figure 8. Execution time vs. parallel threads and problem size.

Efficiency for Large-Scale Computation: The parallel efficiency is more pronounced for larger receiver grids (e.g., 200×200). The absolute time reduction is greatest for these cases, and the speedup remains stable until the hardware limit is hit. This underscores the framework's particular suitability for accelerating the reconstruction of large-scale, high-resolution radio maps, reducing the computation time from nearly 10 s serially to under half a second in parallel—a critical improvement for near-real-time applications.

In conclusion, the HiFiRadio framework is not only accurate and efficient but also effectively leverages modern multi-core architectures through parallelization. This design choice future-proofs the method, ensuring it can capitalize on the increasing core counts of future processors to tackle even larger scenes, such as Dataset-2 (the $500 \text{ m} \times 500 \text{ m}$ area with ~ 1.45 million faces from Section 4.2), with continued high performance.

4.4. Computational Performance on Large-Scale Datasets

A critical requirement for a practical radio map reconstruction framework is its ability to handle increasingly large and detailed 3D environmental models without a catastrophic increase in computational cost. This subsection evaluates the scalability of the HiFiRadio framework with respect to the size and complexity of the input 3D mesh.

Dataset and experimental setup:

We employed four datasets of vastly different scales to stress-test the framework:

Dataset-1: The $250 \text{ m} \times 250 \text{ m}$ area from SUM_Helsinki, comprising ~ 0.34 million triangular faces.

Dataset-2: The $500 \text{ m} \times 500 \text{ m}$ area from SUM_Helsinki, comprising ~ 1.45 million triangular faces (introduced in Section 4.2).

Dataset-3: A large urban area comprising ~ 3.14 million triangular faces.

Dataset-4: An extensive city-scale area comprising ~ 16.03 million triangular faces.

For all experiments, the parallel thread count was fixed to 16 to leverage the performance gains established in Section 4.3. We measured the runtime required to reconstruct received power maps of three different resolutions (50×50 , 100×100 , and 200×200 receiver points) on each dataset. The transmitter was placed at a central location within each scene.

Results and analysis:

The runtime results are summarized in Table 4. The data clearly demonstrates the framework's robustness against scene complexity.


Table 4. Runtime (in seconds) for different receiver grid sizes across datasets of increasing mesh complexity (16 threads).

Number of Faces	~0.34 M (Dataset-1)	~1.45 M (Dataset-2)	~3.14 M (Dataset-3)	~16.03 M (Dataset-4)
50 × 50 (2500)	0.05	0.07	0.08	0.11
100 × 100 (10,000)	0.17	0.23	0.32	0.40
200 × 200 (40,000)	0.64	0.89	1.25	1.53

The key observation from Table 3 and Figure 9 is the sub-linear growth of runtime as the number of triangular faces in the mesh increases by nearly 50 times (from 0.34 M to 16.03 M). For instance, when reconstructing a 200 × 200 map,

The mesh size increases by a factor of ~ 47 .

The runtime only increases by a factor of ~ 2.4 (from 0.64 s to 1.53 s).

**Figure 9.** Execution time across datasets of varying sizes and computational complexities.

This highly favorable scaling is a direct consequence of the spatial indexing structure (R-tree) at the core of our framework. The R-tree organizes the massive mesh data in a hierarchical manner. During a ray query, the algorithm does not need to check every single triangle in the scene. Instead, it efficiently traverses the tree, quickly pruning away large branches of the hierarchy whose bounding volumes do not intersect the ray. This reduces the average time complexity of a single ray query from $O(n)$ (checking all n triangles) to approximately $O(\log n)$.

Therefore, the total computational cost for reconstructing a radio map with k receiver points scales roughly as $O(k \log n)$, rather than $O(kn)$. This logarithmic scaling with respect to scene complexity (n) is the reason why the runtime increase is minimal even when moving from a medium-sized urban model (Dataset-1) to a massive city-scale model (Dataset-4).

The experimental results confirm that the HiFiRadio framework is not only efficient but also highly scalable. Its performance degrades gracefully with increasing input size, thanks to the intelligent data organization provided by the semantic 3D R-tree index. This property is indispensable for practical applications, as it ensures that the framework can be deployed for city-scale planning and optimization using the highly detailed 3D models that are becoming increasingly available.

4.5. Validation with Field Measurements

To further validate the accuracy of our proposed HiFiRadio framework in a real-world scenario, we conducted a field measurement campaign. Simulation-based validation, while valuable, can be influenced by model imperfections. In situ measurements provide the most credible ground truth for evaluating radio map reconstruction performance.

Field measurement setup:

The experiment was conducted near the Hunan Aerospace Hospital in Changsha, China. We first captured the 3D geometry of a 200 m $\times$ 200 m area using an unmanned aerial vehicle (UAV) equipped with an oblique photogrammetry camera. The resulting 3D mesh was semantically segmented using the MeshCNN model to classify objects (e.g., buildings, vegetation) and subsequently used to build the semantic 3D R-tree index, as described in Section 3.1.

A signal transmitter was placed at coordinates (40, 40, 8) m within this area, operating at 3 GHz with a transmission power of 40 dBm. A handheld spectrum analyzer was used as a mobile receiver to measure the received signal power (RSS) along a pre-defined path, as shown in Figure 10. A total of 4000 measurement points, along with their precise coordinates, were collected.




Figure 10. Setup and Measured Results. The route-planning diagram (a) shows the measurement path; the field-measured heatmap (b) shows the received signal strength (RSS) along the path. Regions A, B, and C correspond to LOS, NLOS, and OLOS conditions, respectively.

Results and Analysis:

The measurement path and the corresponding recorded RSS values are visualized in Figure 10, respectively. The RSS distribution along the path can be clearly divided into three distinct intervals, which correspond remarkably well to the three propagation categories defined by our framework:

Region A (green): Corresponds to a Line-of-Sight (LOS) area with minimal obstruction.

Region B (red): Corresponds to a Non-Line-of-Sight (NLOS) area blocked by large buildings.

Region C (yellow): Corresponds to an Obstructed-Line-of-Sight (OLOS) area where the signal penetrates through vegetation, experiencing attenuated but not completely blocked propagation.

This strong qualitative agreement between the physical environment, the measured signal strength, and our path categorization logic provides compelling evidence for the soundness of our methodological foundation.

Quantitative comparison:

To quantitatively assess the accuracy, we generated the predicted radio maps for the entire area using both our HiFiRadio framework and the ray-tracing software Wireless Insite (version 3.3.2), using the same transmitter parameters and a receiver height of 1.5 m. We then compared the predictions against the field measurements. For a balanced evaluation, we randomly selected 500 sample points from each of the three propagation regions (LOS, NLOS, OLOS) identified in Figure 10) and calculated the root mean square error (RMSE). The results are summarized in Table 5.

Table 5. RMSE (dB) comparison between predictions and field measurements in different propagation regions.

Method	LOSRegion	NLOSRegion	OLOSRegion
HiFiRadio (Ours)	4.17	9.31	5.75
Wireless Insite	1.91	5.65	4.95

Discussion and Conclusions:

The results in Table 5 lead to two key insights:

Inherent Limitations of Modeling: Even the computationally intensive ray-tracing method (Wireless Insite) exhibits non-negligible errors when compared to real-world measurements. This highlights the inherent challenge of perfectly modeling complex electromagnetic propagation phenomena, which are influenced by countless factors in the real environment that are difficult to fully capture (e.g., exact material properties, surface roughness, small-scale scatterers).

Favorable Accuracy–Efficiency Trade-off: Although HiFiRadio’s accuracy, particularly in complex NLOS and OLOS conditions, is lower than that of the dedicated ray-tracing simulator, the margin is acceptable. Crucially, this minor compromise in accuracy is made in exchange for a computational speedup of four to five orders of magnitude (as demonstrated in Section 4.2). Our method achieves RMSE values within the same order of magnitude as the high-fidelity benchmark.

This experiment confirms that the HiFiRadio framework successfully strikes an effective balance between efficiency and accuracy. It provides predictions that are sufficiently reliable for numerous practical applications (e.g., network planning, UAV pathing) while being computationally feasible for real-time operation—a feat that pure ray-tracing methods cannot accomplish.

5. Limitations and Future Work

Despite its promising performance, HiFiRadio has several limitations that point to fruitful directions for future research.

1. **Dependency on Semantic Segmentation:** The accuracy of HiFiRadio relies on the quality of semantic segmentation. Misclassification (e.g., glass facades labeled as vegetation) could lead to incorrect path categorization. Future work could explore more robust segmentation networks or multi-sensor fusion (e.g., LiDAR + RGB) to improve label accuracy.

2. **Simplified Physical Modeling:** For NLOS paths, HiFiRadio currently employs the empirical 3GPP model without explicitly modeling multipath reflections or diffraction. While this ensures real-time performance, it may underestimate received power in dense urban canyons. Incorporating simplified ray tracing (e.g., first-order reflections) could enhance NLOS modeling fidelity.

3. **Binary Material Representation:** Materials are currently classified as either penetrable or impenetrable, ignoring frequency-dependent permittivity and conductivity. Future versions could integrate a material library with electromagnetic properties for different surface types.

4. **Dynamic Environment Modeling:** The current framework assumes a static environment. Extending HiFiRadio to handle dynamic objects (e.g., moving vehicles, seasonal foliage changes) would enable time-varying radio map reconstruction.

5. **Multi-frequency and Multi-antenna Support:** Experiments were conducted for a single frequency and omnidirectional antenna. Future work will extend HiFiRadio to multi-band and MIMO systems for 5G/6G network planning.

6. **GPU Acceleration Comparison:** The current parallel implementation uses CPU-based OpenMP. While significant speedup is demonstrated, a comparison with GPU-accelerated ray tracing would provide a more balanced evaluation. A GPU version of HiFiRadio is planned for future work.

6. Conclusions

In this paper, we have presented HiFiRadio, a hybrid model- and data-driven framework for high-fidelity radio-map reconstruction in 3D scenes. By synergistically integrating detailed oblique-photogrammetry-derived 3D meshes with a semantic-aware spatial indexing structure for efficient signal occlusion detection, HiFiRadio effectively bridges the gap between physical accuracy and computational tractability, achieving extremely high efficiency.

Extensive experimental validation demonstrates that HiFiRadio significantly outperforms both classical model-driven methods and contemporary data-driven approaches in terms of reconstruction accuracy, while achieving a computational speedup of several orders of magnitude compared to high-fidelity ray-tracing simulations. Field measurements further corroborate the practical viability of HiFiRadio, showing close alignment between predicted and actual signal propagation patterns across diverse terrain types.

Balanced Perspective on Performance: Although HiFiRadio demonstrates a favorable trade-off between accuracy and efficiency, it is important to acknowledge its current simplifying assumptions. The framework's accuracy is contingent upon the quality of semantic segmentation, and its physical modeling for NLOS conditions relies on empirical rather than deterministic propagation models. Furthermore, the binary classification of materials overlooks nuanced electromagnetic properties. Despite these limitations, HiFiRadio achieves a level of accuracy sufficient for many practical applications—such as network planning, dynamic spectrum access, and UAV path optimization—while operating at speeds that enable near-real-time, city-scale reconstruction, a capability that pure ray-tracing methods cannot provide.

Looking forward, we envision several promising directions for future work, including the integration of dynamic environmental factors (e.g., moving vehicles, seasonal foliage changes), support for multi-frequency and multi-antenna systems, and the extension to time-varying radio maps. Once its current limitations are addressed, particularly through enhanced physical modeling and more robust semantic understanding, the HiFiRadio framework can evolve into an even more powerful tool. It establishes a solid foundation for next-generation radio-environment sensing and exploitation, paving the way toward more intelligent and adaptive wireless networks.

Author Contributions: Conceptualization, K.L., M.M. and L.C.; methodology, K.L. and M.M.; software, K.L. and Y.Z.; validation, K.L., M.M. and L.C.; formal analysis, K.L. and M.M.; investigation, K.L., M.M. and N.J.; resources, M.M. and N.J.; data curation, K.L. and Y.Z.; writing—original draft preparation, K.L.; writing—review and editing, M.M., L.C., Y.Z. and N.J.; visualization, K.L.;

supervision, M.M. and N.J.; project administration, M.M.; funding acquisition, M.M. and N.J. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Zhang, S.; Zhang, R. Radio Map-Based 3D Path Planning for Cellular-Connected UAV. *IEEE Trans. Wireless Commun.* **2021**, *20*, 1975–1989. [\[CrossRef\]](#)
2. Chen, J.; Yatnalli, U.; Gesbert, D. Learning Radio Maps for UAV-Aided Wireless Networks: A Segmented Regression Approach. In Proceedings of the IEEE International Conference on Communications (ICC), Paris, France, 21–25 May 2017; pp. 1–6.
3. Goldsmith, A. *Wireless Communications*; Cambridge University Press: Cambridge, UK, 2005.
4. Romero, D.; Kim, S.J. Radio Map Estimation: A Data-Driven Approach to Spectrum Cartography. *IEEE Signal Process. Mag.* **2022**, *39*, 53–72. [\[CrossRef\]](#)
5. Alzubaidi, L.; Zhang, J.; Fadhel, A.M.; Lu, Z.; Al-Shamma, M.S.; Yang, M.; Zhang, J. Review of Deep Learning: Concepts, CNN Architectures, Challenges, Applications, Future Directions. *J. Big Data* **2021**, *8*, 53. [\[CrossRef\]](#) [\[PubMed\]](#)
6. Creswell, A.; White, T.; Dumoulin, V.; Arulkumaran, K.; Sengupta, B.; Bharath, A.A. Generative Adversarial Networks: An Overview. *IEEE Signal Process. Mag.* **2018**, *35*, 53–65. [\[CrossRef\]](#)
7. Hu, B.; Huang, X.; Wang, M.; Zhang, L. Automated 3D Building Reconstruction from Oblique Photogrammetry Imagery Using Deep Learning. *IEEE Trans. Geosci. Remote Sens.* **2020**, *58*, 4889–4902.
8. Hata, M. Empirical Formula for Propagation Loss in Land Mobile Radio Services. *IEEE Trans. Veh. Technol.* **1980**, *29*, 317–325. [\[CrossRef\]](#)
9. Kyösti, P.; Holtzman, J.; Tufvesson, F.; Benjebbour, A. WINNER II Channel Models for Future Radio Systems. *EURASIP J. Wirel. Commun. Netw.* **2011**, *2011*, 1–14.
10. 3GPP. *Study on Channel Model for Frequencies from 0.5 to 100 GHz*; Technical Report TR 38.901; 3GPP: Sophia Antipolis, France, 2020.
11. Levie, R.; Yapar, Ç.; Kutyniok, G.; Caire, G. RadioUNet: Fast Radio Map Estimation with CNNs. *IEEE Trans. Wireless Commun.* **2021**, *20*, 4001–4015. [\[CrossRef\]](#)
12. Zhang, S.; Wijesinghe, A.; Ding, Z. RME-GAN: A Learning Framework for Radio Map Estimation Based on Conditional Generative Adversarial Network. *IEEE Internet Things J.* **2023**, *10*, 18016–18027. [\[CrossRef\]](#)
13. Li, X.; Zhang, S.; Li, H.; Li, X.; Xu, H.; Mei, H.; Zhu, G.; Qi, N.; Xiao, M.; Xu, L. RadioGAT: A Joint Model-Based and Data-Driven Framework for Multi-Band Radiomap Reconstruction via Graph Attention Networks. *IEEE Trans. Wireless Commun.* **2024**, *23*, 17777–17792. [\[CrossRef\]](#)
14. Gröger, G.; Plümer, L.; Kolbe, T.H. The CityGML Standard—An Introduction to the Virtual 3D City Model Standard. In Proceedings of the Proc. Int. Symp. 3D Geoinf. (3DGI), Berlin, Germany, 29–30 November 2012; pp. 1–10.
15. Guttman, A. R-Trees: A Dynamic Index Structure for Spatial Searching. In Proceedings of the ACM SIGMOD, Boston, MA, USA, 18–21 June 1984; pp. 47–57.
16. Feng, W.; Gang, W.; Deji, P.; Yuan, L.; Yang, L.; Wang, H. A Parallel Algorithm for Viewshed Analysis in Three-Dimensional Digital Earth. *Comput. Geosci.* **2015**, *75*, 57–65. [\[CrossRef\]](#)
17. Chen, Y.; Chen, J. A Parallel Multipoint Viewshed Analysis Method for Urban 3D Building Scenes. *Trans. GIS* **2021**, *25*, 2010–2028. [\[CrossRef\]](#)
18. Orlof, J.; Ozimek, P. TIN Surface and Radial Viewshed Determination Algorithm Parallelisation on Multiple Computing Machines. *Symmetry* **2021**, *13*, 424. [\[CrossRef\]](#)
19. Jung, J.; Olsen, M.J.; Hurwitz, D.S.; Kashani, A.; Buker, K. 3D Virtual Intersection Sight Distance Analysis Using LiDAR Data. *Transp. Res. Part C* **2018**, *86*, 563–579. [\[CrossRef\]](#)
20. Parent, J.R.; Lei-Parent, Q. Rapid Viewshed Analyses: A Case Study with Visibilities Limited by Trees and Buildings. *Appl. Geogr.* **2023**, *154*, 102942. [\[CrossRef\]](#)
21. Hanocka, R.; Hertz, A.; Fish, N.; Giryas, R.; Fleishman, S.; Cohen-Or, D. MeshCNN: A Network with an Edge. *ACM Trans. Graph.* **2019**, *38*, 1–12. [\[CrossRef\]](#)
22. Preparata, F.P.; Shamos, M.I. *Computational Geometry: An Introduction*, 2nd ed.; Springer: Berlin/Heidelberg, Germany, 2012.

23. Möller, T.; Trumbore, B. Fast, Minimum Storage Ray/Triangle Intersection. In Proceedings of the ACM SIGGRAPH Courses, Los Angeles, CA, USA, 31 July–4 August 2005; p. 7-es.
24. ITU-R. *Attenuation in Vegetation*; Technical Report Recommendation P.833; ITU-R: Geneva, Switzerland, 2020.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.