Wireless Radiation Field Reconstruction with Random Initialized 3D Gaussians

ZEZHOU YAO, Beijing Institute of Technology, China
CHUNHUI DUAN*, Beijing Institute of Technology, China
RUOXING WANG, Beijing Institute of Technology, China
FAN LI, Beijing Institute of Technology, China

Modern communication systems demand precise wireless channel modeling to underpin the rapid advancement of high-rate, low-latency network technologies. Wireless radiation field reconstruction based on neural networks has recently emerged as a promising research direction, attracting significant attention from the research community due to its potential for accurate scene-aware channel representation. However, existing solutions heavily rely on high-quality point clouds with geometric priors for scene construction. This reliance imposes substantial limitations, primarily stemming from the imperfection of such data and the high costs associated with its acquisition in complex real-world settings. To mitigate this challenge, we propose WRF-RAIN, a generic and effective framework capable of reconstructing robust wireless radiation fields from imperfect, low-quality input data. To this end, WRF-RAIN innovatively introduces a flexible initialization scheme that trains a 3D Gaussian splatting model for wireless channel modeling using only randomly initialized point clouds, which eliminates reliance on prior geometric information and thus fundamentally relaxes the stringent data requirements inherent to prior methods. Furthermore, we develop a progressive low-pass filtering mechanism and an adaptive density control algorithm to adjust point cloud distribution, jointly guiding optimization from coarse to fine. These designs effectively compensate for the absence of geometric priors caused by random initialization and balance global scene coherence with local detail preservation. Extensive evaluations demonstrate that our framework outperforms state-of-the-art methods under sparse or noisy data conditions, offering a scalable, real-time solution for next-generation networks and various applications.

CCS Concepts: • **Human-centered computing** → **Ubiquitous and mobile computing theory, concepts and paradigms**; • **Networks** → *Mobile networks*.

Additional Key Words and Phrases: Wireless radiation field reconstruction, 3D Gaussian splatting, Random initialization

ACM Reference Format:
Zezhou Yao, Chunhui Duan, Ruoxing Wang, and Fan Li. 2026. Wireless Radiation Field Reconstruction with Random Initialized 3D Gaussians. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 10, 3, Article 185 (September 2026), 27 pages. https://doi.org/10.1145/3832013

1 INTRODUCTION

Modern communication systems are undergoing transformative evolution, driven by the rapid deployment of 5G and the imminent advent of 6G networks [2]. As these technologies pave the way for an increasingly interconnected world, the demand for more efficient and reliable wireless communication is growing exponentially.

*Corresponding author.

Authors' Contact Information: Zezhou Yao, School of Computer Science and Technology, Beijing Institute of Technology, Beijing, China, yaozz2026@163.com; Chunhui Duan (corresponding author), School of Computer Science and Technology, Beijing Institute of Technology, Beijing, China, duanch@bit.edu.cn; Ruoxing Wang, School of Computer Science and Technology, Beijing Institute of Technology, Beijing, China, wangrx@bit.edu.cn; Fan Li, School of Computer Science and Technology, Beijing Institute of Technology, Beijing, China, fli@bit.edu.cn.



This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2026 Copyright held by the owner/author(s).
ACM 2474-9567/2026/9-ART185
https://doi.org/10.1145/3832013

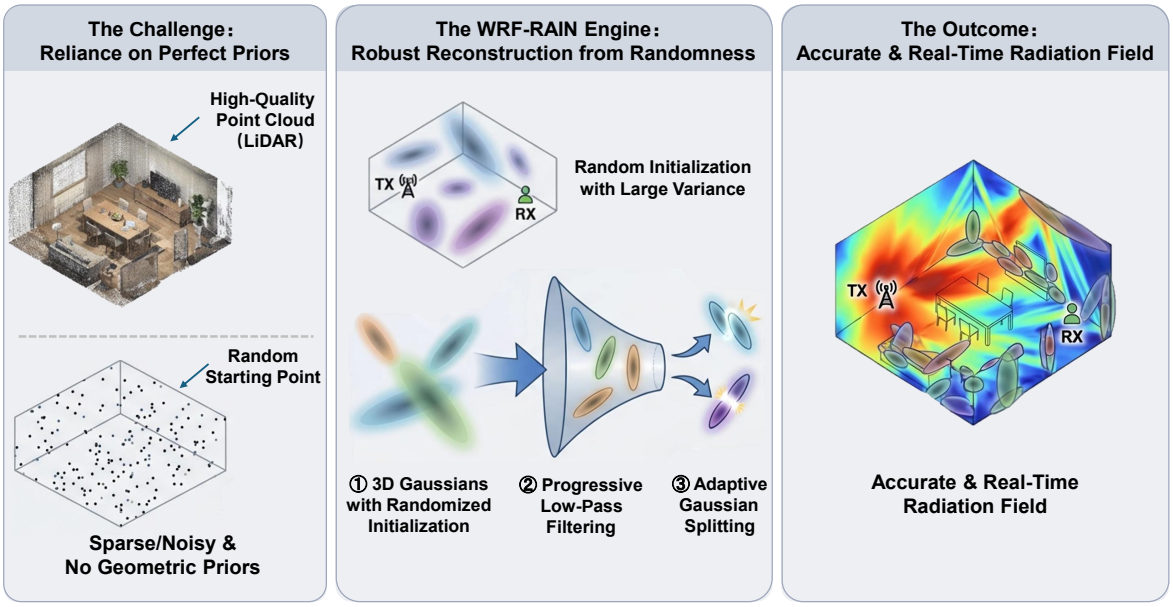


Fig. 1. A visual comparison between traditional methods that rely on precise data and the proposed WRF-RAIN engine, which robustly reconstructs accurate 3D scenes in real-time from sparse, noisy data.

The core of optimizing these systems is the ability to model and understand the intricate nature of wireless signals. Consequently, wireless signal modeling has become indispensable for the design, analysis, and optimization of communication systems, as it characterizes how radio waves propagate, interact with the environment, and ultimately influence system performance [25].

As communication networks grow denser, more complex, and increasingly bandwidth-intensive, the modeling of wireless channels has become markedly more sophisticated. Although Maxwell's equations fully describe electromagnetic wave propagation [3], real-world wireless environments introduce significant challenges, including high variability, severe multipath effects, and dynamic environmental changes. In practice, wireless signal propagation is frequently affected by obstacles, leading to reflection, scattering, diffraction, and absorption, which are fundamental phenomena extensively studied in classical wireless propagation theory [14, 21, 23, 37]. These factors necessitate highly accurate yet computationally efficient modeling approaches capable of supporting large-scale and real-time prediction, particularly in the context of next-generation wireless systems [2, 22, 27].

To address these challenges, traditional wireless signal modeling approaches, such as deterministic ray tracing [5, 37], and probabilistic or statistical models [26], have been widely explored. Ray tracing techniques explicitly simulate the physical propagation paths of radio waves as they interact with surrounding environments, offering high physical interpretability and accuracy. However, such methods are computationally expensive and often struggle to deliver the fine-grained spatial resolution and low latency required for real-time and large-scale applications, especially in complex urban or high-frequency scenarios [1, 5]. Probabilistic models, while providing a statistical framework that can better accommodate uncertainty and environmental diversity, still face limitations in capturing highly complex propagation dynamics under realistic conditions. In response to these challenges, emerging neural network-based approaches have demonstrated substantial potential by offering more flexible and adaptive mechanisms for environmental reconstruction and wireless signal behavior modeling [11, 38].

The Wireless Radiation Field (WRF) reconstruction technique WRF-GS [33] ranks among the most promising advancements in this field. Based on 3D Gaussian Splatting (3D-GS), this approach revolutionizes the modeling of complex scenes by using a set of 3D Gaussian functions to represent the environment, allowing for high-quality, real-time wireless environment rendering. Leveraging the smooth nature of Gaussian functions, WRF-GS efficiently models signal propagation in scenarios where traditional methods underperform. Compared to computationally expensive alternatives like NeRF^2 [38], WRF-GS offers a lighter, more scalable solution ideally suited for real-time prediction applications. Nevertheless, its efficacy heavily depends on high-quality initial point clouds for scene construction. Acquiring such point clouds typically requires resource-intensive methods (e.g., LiDAR scanning [20]), rendering them impractical for real-time scenarios and creating significant barriers to consumer-level deployment as well. This constitutes a fundamental limitation, particularly in real-world applications where point cloud data may be noisy, sparse, or even unavailable.

Despite recent advances in random initialization for visual 3D Gaussian splatting (e.g., RAIN-GS [8]), directly transferring such techniques to wireless signal modeling is non-trivial. Unlike visual rendering, wireless radiation fields are inherently complex-valued, phase-sensitive, and governed by multipath superposition effects, where signal contributions interact constructively and destructively across space. Moreover, wireless receivers perceive signals in the angular domain rather than image space, introducing additional challenges in representation and optimization.

To bridge this gap, we propose a novel framework named **WRF-RAIN**, which enables robust and accurate wireless radiation field reconstruction even when point cloud data is imperfect or incomplete. Rather than directly adopting prior methods, WRF-RAIN introduces a wireless-specific scene representation framework that integrates random initialization with signal-aware modeling and angular-domain optimization. This innovation allows 3D-GS models to be trained from sparse, randomly initialized point clouds without compromising performance. As illustrated in Fig. 1, WRF-RAIN leverages sparse initialization with large-variance Gaussians, paired with progressive Gaussian low-pass filtering, allowing the model to adaptively learn scene structures despite suboptimal initial conditions. Importantly, these mechanisms are co-designed with the wireless signal formation process, enabling stable optimization under complex multipath interactions.

Integrating random initialization into wireless signal modeling presents several fundamental challenges. First, the absence of reliable geometric priors makes the initialization inherently ill-posed, as the model must infer spatial structures purely from sparse signal observations. Second, the presence of sparse and noisy measurements leads to unstable optimization dynamics, where random initialization can easily result in poor convergence or degenerate solutions. Third, there exists an intrinsic tension between global structure modeling and local detail reconstruction, requiring the model to simultaneously capture large-scale propagation patterns and fine-grained multipath interactions. These challenges are further exacerbated by the unique characteristics of wireless signals. In particular, wireless radiation fields are complex-valued and phase-sensitive, and their superposition behavior introduces strong non-linearity in optimization. As a result, naive adaptations of existing random initialization strategies are prone to instability and suboptimal local minima.

Contributions. The main contributions of this paper are summarized as follows.

- We identify and formalize the challenges of applying random initialization to wireless radiation field reconstruction, highlighting fundamental differences between visual and RF modeling.
- We propose WRF-RAIN, a wireless-specific scene representation framework that enables robust reconstruction directly from randomly initialized Gaussians without geometric priors.
- We devise a co-optimized training strategy integrating variance annealing, angular-domain filtering, and adaptive Gaussian density control, ensuring stable optimization under intricate multipath interactions.
- Through extensive experimental evaluations, we demonstrate that WRF-RAIN achieves competitive performance under substantially relaxed assumptions, particularly in sparse and noisy settings.

Roadmap. The remainder of the paper is structured as follows. Section 2 presents related work. Section 3 introduces preliminaries. Section 4 details our framework. Section 5 presents experiments, followed by a case study in Section 6. Finally, Section 7 concludes the paper.

2 RELATED WORK

This section reviews existing literature on wireless channel modeling and RF-based localization/tracking that is related to our work.

2.1 Wireless Channel Modeling

Wireless channel modeling has long been a foundational problem in wireless communications and sensing, aiming to accurately characterize signal propagation behaviors under complex environmental interactions.

Early wireless channel models predominantly relied on probabilistic formulations [14, 23], where signal attenuation and fading effects are approximated using statistical distributions as functions of distance and environmental type. While computationally efficient and analytically tractable, such models inherently lack spatial specificity and struggle to capture fine-grained multipath structures. In contrast, deterministic ray-tracing approaches [5, 37] explicitly simulate wave propagation through geometric interactions such as reflection, diffraction, and penetration. Although these methods offer high physical interpretability and spatial accuracy, their computational complexity grows rapidly with scene scale and environmental complexity, limiting real-time applicability.

Recent advances in neural field representations have introduced a paradigm shift in wireless channel modeling. NeRF^2 [38] pioneered the adaptation of neural radiance fields [13] from optical rendering to electromagnetic signal modeling. By representing the environment as a continuous neural field parameterized by multi-layer perceptrons (MLPs) [28], NeRF^2 treats each spatial location as a potential radiation source with learnable complex-valued signal attributes, enabling implicit reconstruction of spatial spectrum that encodes multipath propagation phenomena. This formulation significantly improves modeling expressiveness over classical analytical methods, supporting tasks such as channel prediction and indoor localization; however, its reliance on scene-specific training and expensive MLP evaluations limits scalability and real-time performance.

To address these limitations and further enhance model generalization, recent works have explored hybrid and transformer-based wireless radiance representations. GWRF [34] introduces a generalizable wireless radiance field framework that incorporates geometry-aware transformer encoders to learn scene representations across different environments. Complementarily, Frequency-Aware Neural Radio-Frequency Radiance Fields [39] propose integrating frequency domain characteristics into the radiance field representation to better capture wideband effects and spectral variations in wireless propagation, addressing the multi-frequency generalization challenge.

Beyond purely MLP-based fields, explicit geometric representations have emerged as a promising direction for efficient and interpretable modeling. Building on explicit primitives, RFSPM [35] leverages 3D Gaussian splatting, a technique from real-time 3D scene synthesis, to develop a scalable RF signal spatial propagation modeling framework. In parallel, nGRF [30] applies explicit 3D Gaussian primitives to directly aggregate electromagnetic field contributions, acting as localized radio modulators.

Building upon this neural field foundation, WRF-GS [33] introduces a more explicit and physically grounded representation by leveraging 3D Gaussian primitives [9] to model wireless propagation environments. Instead of encoding the entire scene within a monolithic neural network, WRF-GS decomposes the environment into a set of spatially localized Gaussian components, each parameterized by position, covariance, and signal-related attributes. Environmental attenuation, reflection strength, and material-dependent signal characteristics are embedded directly into these learnable Gaussian parameters, enabling efficient reconstruction of multipath effects while preserving spatial interpretability. Crucially, the explicit Gaussian formulation allows the rendering process to

align with hardware-accelerated rasterization pipelines, resulting in substantial gains in computational efficiency compared to implicit neural representations. The modeling framework is further extended by WRF-GS+ [32], which enhances representational flexibility through deformable 3D Gaussians that dynamically adapt to spatial and temporal signal variations.

2.2 RF-Based Localization and Tracking

RF-based localization and tracking has been extensively studied over the past two decades, evolving along different trajectories depending on signal modality and application requirements. Early works exploited fine-grained physical-layer measurements, while recent studies emphasize robustness, scalability, and temporal consistency.

Phase-based localization offers high precision. Tagoram [36] achieves millimeter-level accuracy using coherent phase measurements from commodity RFID systems, but these methods are sensitive to multipath and phase ambiguity. Harmonic RFID systems [6] address this by leveraging higher-order harmonics to separate direct-path signals from multipath, enabling non-line-of-sight localization, though strict synchronization and controlled deployment remain necessary. Ultra-wideband (UWB) localization constitutes another important direction. By exploiting large signal bandwidth, systems such as [12] achieve centimeter-level accuracy via precise time-of-flight estimation and exhibit strong robustness to multipath. However, their dependence on specialized hardware and regulatory constraints limits large-scale adoption.

Fingerprinting-based methods provide a data-driven alternative by matching location-dependent RF signatures against offline databases. Recent work leverages fine-grained Channel State Information (CSI) for robust fingerprinting, using meta-learning to adapt to changing data [7] and hybrid CNN-LSTM architectures to model CSI fingerprints [29]. Despite their flexibility, these methods still face challenges such as database maintenance and environmental dynamics.

To address the limitations of single-modality methods, recent work increasingly explores hybrid RF localization and tracking systems. GLAC [4] integrates phase and Doppler measurements within a Kalman filtering framework to achieve temporally stable tracking of moving targets. Collectively, these advances reflect a clear shift from static, single-feature localization toward multi-feature, temporally coherent RF tracking systems.

3 PRELIMINARIES

In this section, we present the preliminaries of 3D-GS and RAIN-GS techniques.

3.1 3D Gaussian Splatting

3D Gaussian Splatting (3D-GS) provides an efficient approach for scene representation through a collection of 3D Gaussian primitives [9]. Each primitive is mathematically defined by its center position $\boldsymbol{\mu} \in \mathbb{R}^3$ and covariance matrix $\Sigma \in \mathbb{R}^{3\times3}$, which collectively characterize the spatial distribution and anisotropic properties of the Gaussian. The corresponding 3D Gaussian function is expressed as:

$$G(\mathbf{x}) = \exp\left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^\top \Sigma^{-1}(\mathbf{x}-\boldsymbol{\mu})\right), \tag{1}$$

where $\mathbf{x} = (x_0, x_1, x_2)$ represents coordinates in 3D space. The covariance matrix Σ can be factorized into scaling and rotation components:

$$\Sigma = \mathbf{R}\mathbf{S}\mathbf{S}^\top\mathbf{R}^\top, \tag{2}$$

with $\mathbf{S}$ being a diagonal scaling matrix and $\mathbf{R}$ an orthonormal rotation matrix.

The rendering pipeline involves three key steps:

(1) **Projection**: 3D Gaussians $G(\mathbf{x})$ are projected onto the image plane to form 2D Gaussians $G_i'(\mathbf{x}')$ [40];
(2) **Sorting**: Gaussians are depth-sorted to ensure proper occlusion handling;

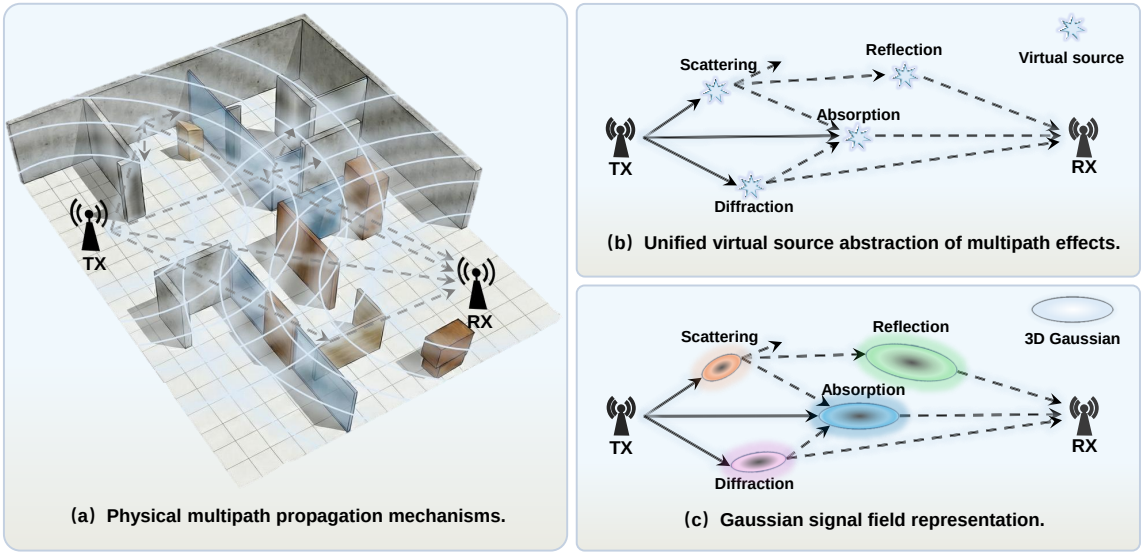


Fig. 2. Enriched multipath signal representation with multi-stage virtual sources and 3D Gaussians.

(3) **Blending**: Per-pixel colors are computed through α-blending [18]:

$$C = \sum_{i=1}^{N} c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j). \tag{3}$$

Here, C is the final color at screen pixel $\mathbf{p}_s$, N is the number of overlapping Gaussians at that pixel, and $c_i \in \mathbb{R}^3$ is the RGB color of the i-th Gaussian. The transparency contribution $\alpha_i = o_i G'_i(\Delta \mathbf{p}_i)$ combines the opacity $o_i \in [0, 1]$ with the 2D Gaussian evaluation $G'_i(\Delta \mathbf{p}_i)$, where $\Delta \mathbf{p}_i = \mathbf{p}_i - \mathbf{p}_s$ is the pixel offset between the projected Gaussian center $\mathbf{p}_i$ and the screen pixel $\mathbf{p}_s$. The complete image is rendered by processing tiles in parallel, followed by loss computation against ground truth for parameter optimization.

This representation delivers distinct advantages for wireless channel modeling: 1) Gaussian primitives naturally model signal propagation as virtual transmitters; 2) The explicit representation enables efficient Wireless Radiance Field (WRF) reconstruction; 3) Projection onto receiver antenna planes facilitates rapid power spectrum synthesis. The mathematical formulation in Eq. (1) and the rendering pipeline in Eq. (3) form the foundation for our wireless channel modeling framework.

As illustrated in Fig. 2, complex wireless multipath propagation in structured environments can be compactly represented through an enriched virtual source formulation based on 3D Gaussian primitives. Instead of explicitly modeling individual electromagnetic interactions with surrounding objects, diverse propagation effects are implicitly captured by a collection of Gaussian-weighted virtual sources distributed in space. This representation enables a unified and differentiable signal field description, naturally aligning physical multipath behavior with the Gaussian splatting framework. Consequently, 3D-GS provides an effective foundation for modeling wireless radiance fields with high efficiency and interpretability.

3.2 RAIN-GS

RAIN-GS emerges as a transformative approach for 3D scene representation by fundamentally rethinking the initialization requirements of conventional 3D Gaussian splatting. At its core, this technique redefines how

Gaussian primitives interact with the optimization landscape, enabling remarkable robustness against imperfect initial conditions.

Central to this approach is the *sparse-large-variance initialization strategy*, which deliberately begins with an extremely limited set of Gaussians possessing expansive coverage. This intentional sparsity forces the system to prioritize global structure before local details, mirroring the coarse-to-fine learning pattern observed in successful Structure-from-Motion (SfM) [24] initialized pipelines, where an approximate global geometry is first recovered from sparse keypoints before being refined through dense reconstruction. Mathematically, the initialization strategy can be described as follows:

$$\mathbf{G}_0 = \{\mathcal{N}(\mu_i, \Sigma_i) \mid \mu_i \in \mathbb{R}^3,\ \Sigma_i = \alpha \cdot I_3,\ i \in \{1, \ldots, N\}\}. \tag{4}$$

Here, $\mathbf{G}_0$ denotes the set of initial Gaussian primitives, each with a mean μ_i and covariance matrix Σ_i. The covariance matrix is initialized with a large variance factor α, which ensures broad coverage of the scene. N represents the number of Gaussian primitives used for the initial scene representation.

The large initial variances create overlapping influence regions that promote collaborative optimization across Gaussians, rather than isolated primitive refinement. This sparsity, combined with the large initial variances, compels the system to prioritize global structure before local details, ensuring that the overall scene context is learned first.

Complementing this initialization strategy, the *progressive Gaussian low-pass filtering mechanism* dynamically regulates the effective receptive field of each Gaussian primitive during training. By systematically adjusting the filter strength in relation to the evolving Gaussian density, the method maintains an optimal balance between global coherence and local precision throughout the optimization trajectory. The effect of the low-pass filter on the receptive field can be modeled as:

$$R(\mathbf{G}_i) = \mathcal{L}(\mathbf{G}_i, t) = \exp\left(-\frac{d(\mu_i, \mu_j)^2}{\sigma(t)^2}\right), \tag{5}$$

where $d(\mu_i, \mu_j)$ denotes the Euclidean distance between the means of two Gaussians μ_i and μ_j, and $\sigma(t)$ is the time-varying scale factor controlling the filter strength at iteration t. The progressive nature of $\sigma(t)$ ensures that the receptive field expands as the training progresses, initially focusing on large-scale features and later refining local details.

The *adaptive bound-expanding split algorithm* addresses one of the most persistent challenges in Gaussian-based representations—the tendency for Gaussians to become trapped in local regions. By actively transporting Gaussians beyond their initial bounds during the critical warm-up phase, this component guarantees comprehensive scene coverage regardless of starting conditions. Mathematically, the expansion of a Gaussian bound can be modeled as:

$$\mu_i(t+1) = \mu_i(t) + \delta(t) \cdot \mathbf{v}_i, \tag{6}$$

where $\mu_i(t)$ is the mean of the i-th Gaussian at iteration t, $\delta(t)$ is a scaling factor that controls the expansion, and $\mathbf{v}_i$ is a direction vector that guides the expansion of the Gaussian. The expansion mechanism is designed to adaptively scale the bound relative to the scene dimensions, thereby ensuring appropriate exploration without excessive dispersion of Gaussians.

What distinguishes RAIN-GS most profoundly is its ability to maintain the exceptional rendering quality and real-time performance of conventional 3D-GS while dramatically expanding the range of viable starting conditions. The technique demonstrates that accurate initialization, while helpful, is not fundamentally necessary for high-quality novel view synthesis when the optimization process itself incorporates appropriate inductive biases.

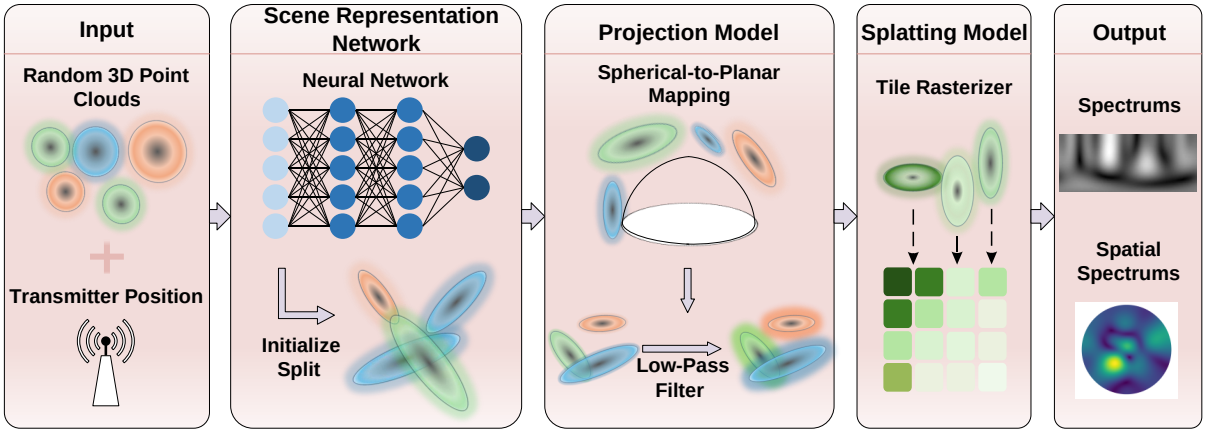


Fig. 3. An overview of the WRF-RAIN framework.

4 THE WRF-RAIN FRAMEWORK

In this section, we present WRF-RAIN, a novel wireless signal modeling framework designed to address the fundamental challenges of random initialization in wireless radiation field reconstruction.

As discussed in Section 1, three key challenges arise in this setting: (1) the absence of reliable geometric priors renders initialization inherently ill-posed; (2) sparse and noisy measurements lead to unstable optimization; and (3) the need to simultaneously model global propagation structures and local multipath effects introduces a coarse-to-fine modeling conflict.

To tackle these challenges, WRF-RAIN adopts a unified design that integrates three complementary strategies: (1) randomized Gaussian initialization with variance annealing for stable global exploration; (2) adaptive density control for data-driven refinement under sparse observations; and (3) progressive angular-domain filtering for robust coarse-to-fine signal reconstruction. The overall architecture of WRF-RAIN is illustrated in Fig. 3.

Problem Setting. We consider a typical indoor scenario where a transmitter (TX) is placed at arbitrary positions within a space, and a fixed receiver (RX) equipped with a spatial antenna array captures the incoming electromagnetic field. The environment contains various objects that induce multipath propagation, and the RX's task is to perceive the spatial spectrum—i.e., the angular power distribution—of the received signals. From a modeling perspective, each multipath component can be interpreted as a virtual transmitter embedded within the environment.

Motivation. In the original WRF-GS framework, virtual TXs are modeled using 3D Gaussians whose initial positions and shapes are derived from precomputed point clouds (e.g., those obtained via LiDAR). While effective in controlled settings, this reliance on dense geometric priors limits applicability to real-world scenarios where such data is unavailable or unreliable. To address this, WRF-RAIN adopts a fundamentally different initialization paradigm that removes reliance on geometric priors. Instead of initializing from explicit geometry, the model begins with randomly distributed large-variance Gaussian primitives, enabling it to explore the solution space in a data-driven manner.

Overview. The WRF-RAIN framework is built around three core components:

- **Scene Representation Network**: This module characterizes the wireless propagation environment using a set of randomly initialized 3D Gaussian primitives. It leverages variance annealing and gradient-driven splitting strategies to iteratively refine the spatial resolution throughout the training process.

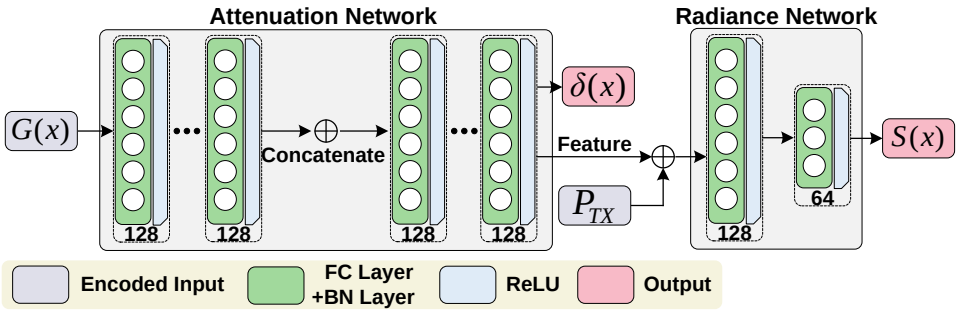


Fig. 4. Architecture of the neural model with two MLPs.

- **Projection Module**: This module maps each Gaussian onto a hemispherical perception plane that is aligned with the RX antenna array, employing Mercator projection to facilitate accurate angular-domain signal synthesis.
- **Electromagnetic Splatting Module**: This module aggregates signal contributions from all projected Gaussians, and incorporates distance-dependent attenuation to compute the spatial spectrum observed at the RX side.

Together, these modules form a fully differentiable pipeline that reconstructs the spatial signal field from randomly initialized volumetric elements, removing the need for fine initialization with geometric priors. The following subsections detail each component of our WRF-RAIN framework.

4.1 Scene Representation via Learnable 3D Gaussians

To effectively model the spatial distribution of wireless signals, WRF-RAIN represents the environment as a set of learnable 3D Gaussian primitives. Each Gaussian $\boldsymbol{G}(\mathbf{x})$ is parameterized by its position $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, following the standard formulation (see Eq. (1)). To capture wireless signal characteristics, this representation is further extended with two additional components: a complex-valued signal amplitude $\boldsymbol{S}(\mathbf{x})$ and an attenuation factor $\boldsymbol{\delta}(\mathbf{x})$, which together account for multipath effects and signal decay due to path loss. These parameters are dynamically predicted by the scene representation network.

4.1.1 Randomized Initialization with Large Variance. In contrast to the original WRF-GS framework, which initializes Gaussians based on external 3D point clouds, our approach introduces a robust randomized initialization strategy inspired by the RAIN-GS paradigm. Specifically, the centers $\boldsymbol{\mu}$ are uniformly sampled from a bounded spatial region $\mathcal{B} \subset \mathbb{R}^3$, independent of any prior 3D scans or structural information. Each initial Gaussian is assigned an isotropic covariance matrix:

$$\Sigma = \sigma_0^2 \mathbf{I}, \tag{7}$$

where $\sigma_0^2 \in [0.25, 1.0]$ allows for flexible initial coverage.

To stabilize training and improve convergence, an annealing schedule is applied to progressively reduce the variance of Gaussians over time:

$$\sigma^2(t) = \sigma_0^2 \cdot \exp(-\lambda t). \tag{8}$$

Here, t denotes the training step and λ is a decay coefficient. This mechanism mimics a coarse-to-fine optimization trajectory, where early training emphasizes capturing large-scale signal geometry and later stages allow finer structural details to emerge.

4.1.2 Adaptive Gaussian Density Control. Beyond initial coverage and attenuation modeling, WRF-RAIN introduces a progressive refinement mechanism based on adaptive Gaussian splitting. During the training phase, the contribution of each Gaussian is monitored via the gradient of the reconstruction loss with respect to its parameters. For a given Gaussian, the saliency metric is computed as:

$$\Delta_i = \left\| \nabla_{\boldsymbol{\mu}} L_{\text{WRF}}(G(x)) \right\|. \tag{9}$$

If Δ_i exceeds a predefined threshold τ, the corresponding Gaussian is considered under-expressive and is split along its principal eigen-axis of Σ_i, forming two new Gaussians with perturbed means and halved weights. This dynamic refinement allows the model to increase representational density only in regions exhibiting complex signal behavior, such as areas around edges, occlusions, or regions with high Signal-to-Noise Ratio (SNR) gradients.

Therefore the total number of Gaussians evolves adaptively over time. Early training proceeds with a small number of broad Gaussians for coarse modeling, while later iterations concentrate more Gaussians in informative regions, achieving high-fidelity representations without incurring excessive computational overhead. This strategy ensures efficient resource allocation, directing modeling capacity toward spatial regions that contribute most significantly to signal reconstruction accuracy.

4.1.3 Scene Representation Network. The learnable signal and attenuation properties of each 3D Gaussian are modeled using a two-branch Multi-Layer Perceptron (MLP) architecture, as illustrated in Fig. 4.

In the first stage, the coordinates of each Gaussian center are input into an 8-layer fully connected network, where each layer consists of 128 channels with ReLU activation. This branch is tasked with learning the spatial attenuation profile:

$$\delta(x) = \Delta A(x) e^{j\Delta\psi(x)}, \tag{10}$$

where $\Delta A(x)$ and $\Delta\psi(x)$ denote the attenuation amplitude and phase shift introduced by the medium at point x, respectively. The output includes a feature embedding that encodes local spatial context.

In the second stage, this feature vector is concatenated with the current transmitter position P_{TX} and input to a secondary MLP containing two fully connected layers with 128 and 64 channels. The output of this branch models the complex signal field:

$$S(x) = A(x) e^{j\psi(x)}, \tag{11}$$

yielding the amplitude $A(x)$ and phase $\psi(x)$ of the wireless signal emitted by a virtual TX at location x.

Together, the two branches define a mapping:

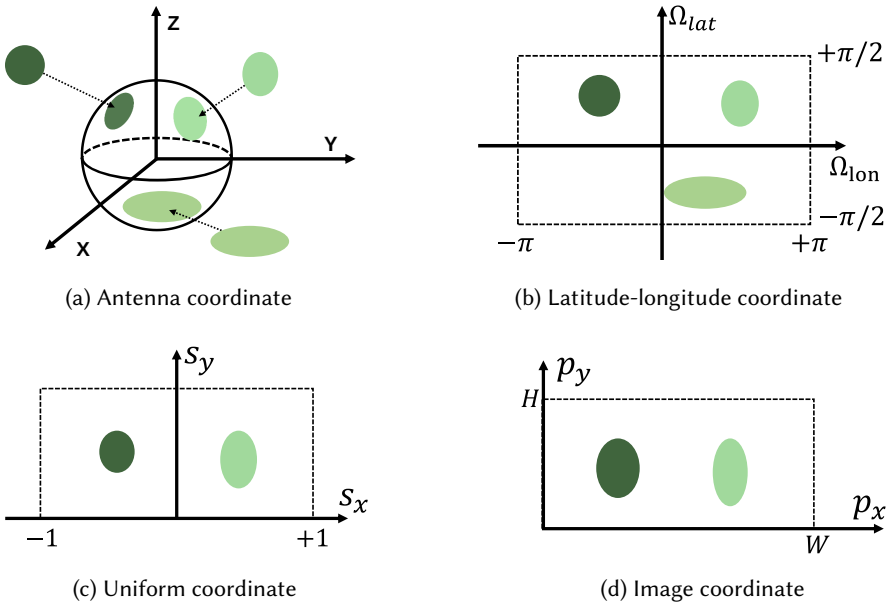
$$F_\Theta : (G(x), P_{\text{TX}}) \mapsto (\delta(x), S(x)), \tag{12}$$

where Θ denotes the shared learnable parameters of the network. While this design builds upon the DeepSDF framework [16], it incorporates critical extensions to support mobile TX positions and complex-valued outputs—critical capabilities for accurate wireless channel modeling.

The proposed scene representation network unifies the above components, namely, random initialization, variance annealing, adaptive splitting, and neural modeling of signal fields, into an end-to-end differentiable pipeline. The parameters of both the Gaussians and the neural network are jointly optimized using backpropagation, guided by the reconstruction loss between predicted and observed wireless field measurements.

4.2 Spherical-to-Planar Mapping for RX Signal Perception

In order to simulate the spatial reception of wireless signals at the RX antenna array, each 3D Gaussian primitive must be projected onto the RX's angular perception plane. Fig. 5 visualizes the spatial transformation process of projecting the 3D Gaussian center onto the RX perception plane. This projection model differs fundamentally from conventional camera projection models in optical domains, such as pinhole or fisheye cameras, owing to the hemispherical nature of radio signal reception. Unlike images that preserve photometric properties under



(a) Antenna coordinate

(b) Latitude-longitude coordinate

(c) Uniform coordinate

(d) Image coordinate

Fig. 5. Visualization of spatial transformation in projection.

perspective transformations, the spatial spectrum formed by RX antennas reflects the angular structure of incoming signals, where directionality takes precedence over radiometric fidelity.

Let the RX location be denoted by $\boldsymbol{x}_{\text{RX}} \in \mathbb{R}^3$, and consider a Gaussian primitive centered at $\boldsymbol{\mu}_i$. The first step is to compute its position relative to the RX:

$$\boldsymbol{d}_i = \boldsymbol{\mu}_i - \boldsymbol{x}_{\text{RX}} = [d_{ix}, d_{iy}, d_{iz}]^\top. \tag{13}$$

This 3D vector is then converted into spherical coordinates to describe the direction of arrival:

$$\Omega_i = (\Omega_i^{\text{lon}}, \Omega_i^{\text{lat}}) = \left(\arctan 2(d_{iy}, d_{ix}), \arcsin\left(\frac{d_{iz}}{\|\boldsymbol{d}_i\|_2}\right)\right). \tag{14}$$

Here, $\Omega_i^{\text{lon}} \in [-\pi, \pi]$ denotes the azimuth angle, and $\Omega_i^{\text{lat}} \in [0, \frac{\pi}{2}]$ represents the elevation, constrained to the upper hemisphere, as signals originating from below the RX are typically blocked or neglected in practical scenarios.

This angular representation offers a natural domain for projection onto a cylindrical 2D perception plane. Instead of the nonlinear Mercator projection widely used in geographic cartography, we adopt a linear cylindrical mapping similar to orthographic projection, which better preserves angular uniformity:

$$s_{ix} = \frac{\Omega_i^{\text{lon}}}{\pi}, \quad s_{iy} = \frac{2\Omega_i^{\text{lat}}}{\pi}. \tag{15}$$

This transformation maps the angular range to a normalized domain $s_{ix} \in [-1, 1]$, $s_{iy} \in [0, 1]$, forming a rectangular image-like grid with a one-to-one correspondence to directions in 3D space.

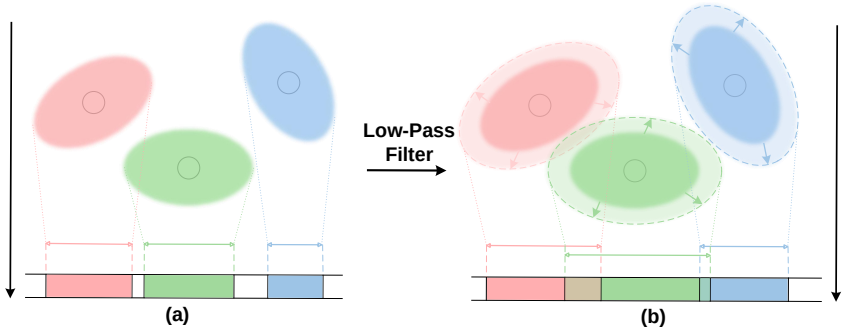


Fig. 6. Visualization of low-pass filter.

To align with the image resolution of the spatial spectrum map, these normalized coordinates are further scaled to pixel coordinates:

$$p_{ix} = \frac{(s_{ix} + 1)}{2} \cdot W, \quad p_{iy} = s_{iy} \cdot H, \tag{16}$$

where W and H represent the width and height (in pixels) of the RX's angular resolution grid. In practice, we set $W = 360$ and $H = 90$ to achieve a 1-degree angular resolution in azimuth and elevation, respectively. This design guarantees that each pixel corresponds to a unique Angle-of-Arrival (AoA) bin.

4.3 Progressive Angular Splatting for Wireless Signal Accumulation

After projecting each 3D Gaussian onto the RX's angular perception plane, we aim to accumulate signal contributions to form the spatial spectrum observable at the receiver. This accumulation process involves a variance-aware splatting operation that distributes each Gaussian's signal energy over nearby angular directions, followed by physics-informed attenuation modeling. To further improve stability and fidelity, we introduce a progressive low-pass filtering mechanism inspired by RAIN-GS, which is specifically adapted to the angular domain.

4.3.1 Variance-Aware Angular Splatting with Progressive Low-Pass Filtering. Each 3D Gaussian primitive is first projected onto the angular perception plane, resulting in an angular coordinate $\boldsymbol{\theta}_i = \Omega(\boldsymbol{\mu}_i - \boldsymbol{x}_{\text{RX}}) \in \mathbb{R}^2$, and an associated angular covariance matrix $\Sigma_{\text{ang},i} \in \mathbb{R}^{2\times2}$, obtained via Jacobian-based transformation:

$$\Sigma_{\text{ang},i} = \mathbf{J}_i \Sigma_i \mathbf{J}_i^\top. \tag{17}$$

To avoid aliasing and facilitate coarse-to-fine optimization, we apply a progressive low-pass filter directly in the angular space. Specifically, we augment the angular covariance matrix by adding an isotropic regularization term:

$$\Sigma_{\text{eff},i}(t) = \Sigma_{\text{ang},i} + s(t) \cdot \mathbf{I}, \tag{18}$$

where $s(t)$ is a time-dependent filter width that gradually decreases over training iterations. Early in training, a larger $s(t)$, as shown in Fig. 6, encourages broader angular smoothing, enabling stable optimization from noisy or sparse initializations. As the optimization procedure progresses, $s(t) \to 0$, allowing the model to capture fine-grained angular structures of the wireless signal field.

The contribution of each Gaussian to a given angular location $\boldsymbol{\theta} \in \mathbb{R}^2$ is computed as:

$$\mathcal{S}_i(\boldsymbol{\theta}) = \exp\left(-\frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{\theta}_i)^\top \Sigma_{\text{eff},i}^{-1}(t)(\boldsymbol{\theta} - \boldsymbol{\theta}_i)\right). \tag{19}$$

This defines an anisotropic Gaussian kernel centered at $\boldsymbol{\theta}_i$, whose spatial spread adapts both to the intrinsic uncertainty of 3D Gaussian and the current filtering bandwidth, which ensures consistent angular modeling throughout the training process.

4.3.2 Attenuation Modeling and Signal Aggregation. Each projected Gaussian is associated with a complex-valued amplitude $\mathbf{a}_i \in \mathbb{C}$, which encodes the strength and phase of its radiated signal. The contribution of each Gaussian to angular location $\boldsymbol{\theta}$ is scaled by a distance-based attenuation factor:

$$T_i = \exp\left(-\beta \cdot \|\boldsymbol{\mu}_i - \boldsymbol{x}_{\mathrm{RX}}\|\right), \tag{20}$$

where β is a learnable or pre-defined attenuation coefficient that controls the rate of signal decay over propagation distance.

To model realistic occlusion and multipath effects in wireless environments, we adopt a depth-aware masking strategy. The angular perception plane is partitioned into discrete tiles (i.e., pixels), and Gaussians are sorted by their radial depth within each tile. The accumulated signal at a given angular location $\boldsymbol{\theta}_k$ is:

$$R_k = \sum_{i=1}^{N_k} \left(\prod_{j=1}^{i-1} \delta(x_j) \right) \cdot \mathbf{a}_i \cdot T_i \cdot \boldsymbol{\mathcal{S}}_i(\boldsymbol{\theta}_k), \tag{21}$$

where $\delta(x_j) \in [0, 1]$ denotes the masking weight of the j-th Gaussian in the tile, and N_k is the number of Gaussians contributing to angular location $\boldsymbol{\theta}_k$.

Finally, the observable spatial spectrum $P_k \in \mathbb{R}$ at pixel k is obtained by computing the power of the accumulated complex signal:

$$P_k = |R_k|^2, \tag{22}$$

which serves as the supervisory target during model training, grounded in real-world physical wireless signal measurements.

This splatting formulation offers several key advantages for wireless signal modeling. First, it ensures angular smoothness and coarse-to-fine convergence through progressive filtering. Second, it incorporates signal strength, distance-dependent attenuation, and occlusion effects in a fully differentiable manner, facilitating end-to-end optimization of the entire model. Third, it aligns closely with the physics of RF propagation, providing a faithful modeling of how virtual transmitters contribute to the RX's directional signal field. Importantly, the filtering bandwidth $s(t)$ can be annealed according to iteration count or validation loss, enabling dynamic adjustment of spatial resolution over time, thereby achieving an optimal balance between training stability and reconstruction fidelity.

4.4 Physics-Informed Frequency-Aware Wireless Field Modeling

While the proposed WRF-RAIN framework affords robust reconstruction of wireless radiation fields from sparse observations, the current formulation assumes a single carrier frequency, which limits applicability in practical wireless systems where signal propagation is inherently frequency-dependent. Specifically, both amplitude attenuation and phase evolution vary with frequency, giving rise to phenomena such as frequency-selective fading and phase inconsistency across channels. To address this, we extend WRF-RAIN by introducing a physics-informed frequency-aware modeling mechanism, building on recent advances in neural wireless field representations (e.g., NeRF^2), while ensuring consistency with electromagnetic propagation principles.

Wave Number Parameterization. Instead of directly using frequency f, we adopt the wave number formulation:

$$k = \frac{2\pi f}{c}, \tag{23}$$

where c is the speed of light. This parameterization is physically grounded, as the phase evolution of electromagnetic waves is linear with respect to the wave number.

Frequency-Conditioned Field Representation. We extend the original neural field formulation to incorporate wave number conditioning:

$$F_\Theta : (G(x), P_{\text{TX}}, k) \mapsto (\delta(x,k), S(x,k)), \tag{24}$$

where both the attenuation $\delta(x,k)$ and the complex signal $S(x,k)$ are functions of the wave number.

To capture high-frequency variations, the wave number is encoded using a sinusoidal embedding:

$$\gamma(k) = \left[\sin(2^l k), \cos(2^l k)\right]_{l=0}^{L-1}, \tag{25}$$

which is concatenated with spatial coordinates and transmitter position before being fed into the network.

Physically Consistent Phase Modeling. A key requirement for frequency-aware wireless modeling is to maintain phase consistency across frequencies. To this end, we explicitly decompose the phase into a deterministic propagation term and a learnable residual:

$$\psi(x,k) = k \cdot d(x) + \Delta\psi(x), \tag{26}$$

where $d(x) = \|x - P_{\text{TX}}\|$ denotes the propagation distance, and $\Delta\psi(x)$ captures environment-induced phase shifts due to reflections, scattering, and diffractions. This formulation ensures that phase evolution with respect to frequency follows physical laws, while allowing the network to model complex multipath effects through the residual term.

Frequency-Dependent Amplitude Modeling. Signal attenuation is also frequency-dependent due to material properties and propagation effects. We model this as:

$$A(x,k) = A_0(x) \cdot \exp(-\beta(k) \cdot d(x)), \tag{27}$$

where $A_0(x)$ is a learnable spatial amplitude and $\beta(k)$ is a frequency-dependent attenuation coefficient parameterized by a lightweight neural function.

Integration with Electromagnetic Splatting. The frequency-aware signal representation is seamlessly integrated into the existing angular splatting pipeline. Each Gaussian primitive contributes a wave-number-dependent complex amplitude:

$$\mathbf{a}_i(k) = A_i(k) e^{j\psi_i(k)}. \tag{28}$$

The accumulated signal at angular location $\boldsymbol{\theta}_k$ becomes:

$$R_k(k) = \sum_{i=1}^{N_k} \left(\prod_{j=1}^{i-1} \delta(x_j, k) \right) \cdot \mathbf{a}_i(k) \cdot T_i(k) \cdot \mathcal{S}_i(\boldsymbol{\theta}_k), \tag{29}$$

and the corresponding spatial spectrum is:

$$P_k(k) = |R_k(k)|^2. \tag{30}$$

Discussion. Unlike naive frequency-conditioned neural fields, our proposed formulation enforces physically consistent phase evolution with respect to frequency through wave number parameterization and structured phase decomposition. This design significantly improves generalization across frequencies and provides a principled foundation for modeling wideband wireless signals. During training, the model is jointly optimized over multiple frequencies, enabling it to learn a continuous frequency-dependent representation rather than fitting each frequency independently.

Moreover, this extension enables WRF-RAIN to readily accommodate fine-grained wireless features such as Channel State Information (CSI), which inherently captures frequency-dependent amplitude and phase variations across subcarriers, broadening its applicability in real-world communication systems.

4.5 Summary and Capabilities

Building on the aforementioned technologies, our WRF-RAIN framework eliminates the reliance on dense geometric priors, enabling robust training directly from randomly initialized 3D Gaussians. The seamless integration of core components, including variance annealing for stable coarse-to-fine optimization, adaptive Gaussian splitting for dynamic spatial resolution refinement, and physics-aware electromagnetic splatting for accurate directional signal aggregation, empowers faithful wireless signal field reconstruction even under challenging real-world conditions (e.g., sparse measurement data and noisy initializations). This makes WRF-RAIN a practical and scalable solution for real-time wireless channel modeling in next-generation communication systems, addressing key limitations of prior methods that depend on high-quality geometric priors.

Limitation. While WRF-RAIN implicitly captures complex propagation effects such as reflection and diffraction through learned Gaussian representations, it does not explicitly model these physical mechanisms in a deterministic manner. This may limit interpretability and could affect generalization under extreme propagation conditions or across drastically different environments. Future work could explore incorporating physics-based priors or hybrid modeling to further enhance robustness. Additionally, the current WRF-RAIN framework assumes a static propagation environment. Adapting to dynamic scenes would require temporal extensions, such as deformable 3D Gaussians or time-variant density control, which constitute an important direction for future work.

5 IMPLEMENTATION AND EVALUATION

We implement WRF-RAIN using PyTorch and evaluate it in a real-world indoor wireless environment. This section focuses on the details of implementation and evaluation, including model initialization, optimization settings, dataset preparation, and performance evaluation.

5.1 Training Setup

The model initializes $N = 1024$ 3D Gaussian primitives with centers $\mu_i \in \mathbb{R}^3$ sampled uniformly from a bounded scene volume of the indoor wireless environment. Covariance matrices are set as shown in Eq. (7), with $\sigma_0^2 = 1.0$ to achieve maximum initial coverage. Notably, no LiDAR scans or other geometric priors are used.

As shown in Eq. (8), a variance annealing schedule is applied during training, with $\lambda = 5 \times 10^{-5}$, which gradually sharpens the Gaussians to focus on finer structure as training progresses.

The model is trained using the Adam optimizer with a learning rate of 10^{-3}. To ensure consistent optimization of all learnable components, we perform joint optimization of both the Gaussian parameters and the neural network weights for a total of 40,000 training iterations. The training objective minimizes a weighted combination of the ℓ_1 loss and Structural Similarity Index Measure (SSIM) [31] between the predicted and measured spatial spectra:

$$\mathcal{L} = (1 - \eta) \cdot \|I_{\text{gt}} - I_{\text{pred}}\|_1 + \eta \cdot (1 - \text{SSIM}(I_{\text{gt}}, I_{\text{pred}})), \tag{31}$$

with $\eta = 0.2$ to balance the two loss terms. The SSIM is computed over the angular domain as a perceptual metric.

Gaussian splitting is initiated after 10,000 optimization steps, based on gradient magnitude thresholds. When triggered, a Gaussian primitive is duplicated with perturbed centers and reduced variance. This adaptive splitting mechanism effectively increases spatial resolution in geometrically complex regions while maintaining computational efficiency by avoiding unnecessary Gaussian proliferation in areas with simple signal variations.

5.2 Experimental Environment

All experiments are conducted on a desktop equipped with an Intel Core i5 CPU, 32 GB RAM, and an NVIDIA RTX 4060Ti GPU with 16 GB memory. Custom CUDA kernels are used for tile-wise Gaussian splatting onto the

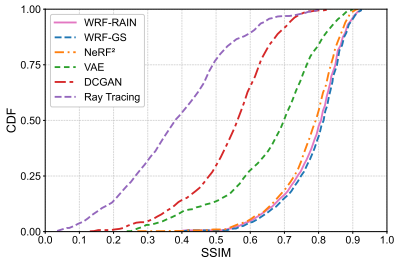


Fig. 7. CDF of SSIM scores.

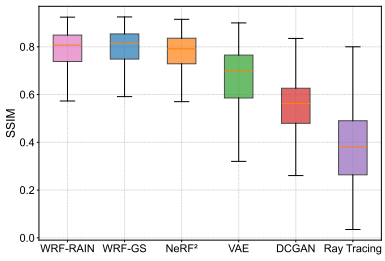


Fig. 8. Box plots of SSIM distributions.

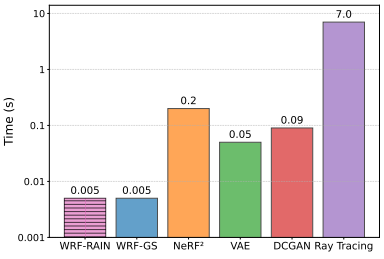


Fig. 9. Rendering time comparison.

angular domain, leveraging the inherent sparsity and locality of Gaussian support to accelerate computation. The angular resolution is set to 360×90, corresponding to 1-degree bins in azimuth and elevation.

The dataset is derived from the NeRF^2 benchmark [38], which captures TX-RX spectra in a cluttered indoor lab environment. The RX is configured as a fixed 4×4 antenna array operating at 915 MHz. The TX is an RFID tag transmitting RN16 messages, moving to various positions across the space. Each data sample comprises a 3D TX position and a corresponding measured spatial spectrum at the RX. The total dataset includes 4,800 samples, of which 4,320 are used for training and 480 for testing. Unlike prior baselines, our method requires no external point clouds or pre-built scene models.

5.3 Performance Evaluation

5.3.1 Overall Performance. To comprehensively evaluate the performance of our system, we compare it with the following baselines, covering different wireless signal modeling paradigms:

- **WRF-GS** [33]: Utilizes LiDAR-based initialization for 3D Gaussians, which serves as a baseline for Gaussian-based wireless modeling;
- **NeRF^2** [38]: Synthesizes wireless fields using volumetric neural rendering, which represents the neural rendering paradigm for wireless modeling;
- **VAE** [10]: Learns latent representations of spatial spectra, representing a typical unsupervised generative modeling approach;
- **DCGAN** [19]: Adversarially trained to generate angular signal maps, serving as another representative of generative models;
- **Ray Tracing** [37]: Simulates physical propagation using pre-built mesh geometry, which is a classic physics-driven modeling method.

SSIM is employed as the principal evaluation metric, given that the structural similarity it quantifies between predicted and measured spatial spectra serves as a key indicator for assessing wireless signal field reconstruction quality. Fig. 7 shows the Cumulative Distribution Function (CDF) [3] of SSIM scores for synthesized spatial spectra generated by different methods. WRF-RAIN demonstrates competitive performance, ranking second among all evaluated methods. Its SSIM curve lies significantly above those of NeRF^2, VAE, DCGAN, and Ray Tracing across nearly the entire distribution, indicating superior structural fidelity. The median SSIM value for WRF-RAIN is approximately 0.81, placing it just below the top-performing WRF-GS (0.82) and notably higher than NeRF^2 (0.78), VAE (0.70), DCGAN (0.56), and Ray Tracing (0.38). Such strong performance validates that WRF-RAIN can effectively preserve the structural similarity between synthesized and ground truth spectra, even without relying on dense geometric priors. Its proximity to WRF-GS suggests that both approaches benefit from a common foundation based on 3D Gaussian modeling, which robustly captures the complex interactions between wireless signals and the surrounding environment. The slight performance gap between WRF-RAIN and WRF-GS

may stem from differences in their specific parameterizations or computational approximations. It is worth noting that WRF-GS leverages LiDAR priors, while WRF-RAIN uses only random initialization.

To further characterize the stability and consistency of each method, we analyze the SSIM distributions using box plots, as illustrated in Fig. 8. This visualization clearly highlights the variability, median, and overall spread of SSIM scores for each model. WRF-RAIN maintains strong and stable performance, nearly matching that of the top-performing WRF-GS. Both methods exhibit tight interquartile ranges, indicating consistent performance across the entire test set. In contrast, VAE and DCGAN present broader SSIM distributions, reflecting their relatively limited ability to capture fine-grained spatial structures of wireless signals. Ray Tracing performs weakest overall, underscoring its inherent challenges in modeling complex wireless environments.

In addition to reconstruction accuracy, we also compare the inference efficiency of all methods by evaluating their per-scene rendering time, as shown in Fig. 9. Results are plotted on a logarithmic scale to reflect the wide range of computational costs across models. Both WRF-RAIN and WRF-GS demonstrate exceptional inference efficiency, requiring only approximately 0.005 seconds per scene—making them the fastest among all evaluated methods by a substantial margin. NeRF^2, despite its relatively strong SSIM performance, incurs a rendering time of 0.2 seconds per scene, highlighting its considerable computational burden in practical real-time deployments. VAE and DCGAN show moderate performance in both accuracy and speed, with rendering time of 0.05 and 0.09 seconds per scene, respectively. Ray Tracing remains the most computationally intensive method, taking approximately 7 seconds per scene, which severely limits its scalability for time-sensitive applications. These results emphasize the practicality of WRF-RAIN, which strikes a near-optimal balance between rendering accuracy and inference efficiency. Its rendering time is three orders of magnitude faster than Ray Tracing while delivering significantly higher structural fidelity. Moreover, its computational parity with WRF-GS further confirms that the key innovations in WRF-RAIN do not come at the cost of runtime performance, preserving its suitability for time-sensitive wireless sensing tasks.

In summary, WRF-RAIN enables high-precision wireless signal field modeling without reliance on geometric priors, while simultaneously maintaining high accuracy and low latency. This key advantage addresses the critical limitation of existing methods that depend on high-quality initial point clouds, making it uniquely suited for real-world scenarios where such data is often imperfect, costly to acquire, or entirely unavailable. By achieving an optimal balance between reconstruction precision and computational efficiency, our framework represents a practical advancement for dynamic, real-time wireless modeling in next-generation communication systems.

5.3.2 Efficiency Discussion. We provide additional results on training efficiency, convergence behavior, memory consumption, and representation complexity for the evaluated dataset.

Training Time. Under the same hardware configuration, WRF-GS requires approximately 5.3 hours to converge, while NeRF^2 requires around 9.8 hours due to its volumetric neural rendering process. Our WRF-RAIN introduces extra optimization overhead from random initialization and adaptive refinement, resulting in a moderately increased training time of about 6.6 hours. It is worth noting that training time is sensitive to scene geometry and signal complexity and may vary across different environments.

Convergence Behavior. We analyze convergence via both training loss and SSIM evolution. WRF-GS typically reaches stable convergence after around 35k iterations, whereas WRF-RAIN converges after approximately 38k iterations. The slightly slower convergence is mainly attributed to the delayed Gaussian splitting strategy (after 10k iterations), which prioritizes stable coarse reconstruction before progressive refinement. Despite this difference, both methods exhibit stable and monotonic convergence trends.

Memory Consumption. The peak GPU memory usage of WRF-GS is approximately 9.4 GB, while WRF-RAIN requires around 11.8 GB due to the additional buffers introduced by adaptive density control and progressive filtering. NeRF^2 exhibits higher memory consumption (approximately 15.6 GB), mainly due to its dense neural field evaluation and volumetric sampling.

Table 1. Robustness under different training data ratios.

Method	SSIM (100%)	SSIM (50%)	SSIM (25%)	SSIM (10%)
WRF-GS	0.82	0.78	0.72	0.68
NeRF^2	0.78	0.74	0.65	0.59
WRF-RAIN	**0.81**	**0.77**	**0.72**	**0.66**

Number of Primitives. Both WRF-GS and WRF-RAIN adopt 3D Gaussian representations initialized with 1024 primitives. After adaptive refinement, WRF-GS typically converges to about 88,000 Gaussian primitives, while WRF-RAIN produces a denser representation with about 100,575 Gaussians. This difference stems primarily from the absence of geometric priors, which necessitates additional primitives to capture complex wireless propagation effects. In contrast, NeRF^2 relies on an implicit neural field representation, making direct comparison in terms of primitive count not applicable.

In summary, we note that efficiency-related metrics such as training time, convergence behavior, memory usage, and representation complexity are inherently dependent on the specific dataset and scene characteristics. In this work, we primarily focus on reconstruction quality (SSIM) and inference efficiency, as they are more directly relevant to practical deployment scenarios.

5.3.3 Robustness to Data Size. To evaluate the robustness of WRF-RAIN under data-scarce conditions, we conduct an experiment by varying the amount of available training data. Specifically, we randomly subsample the training dataset at different ratios (100%, 50%, 25%, and 10%), while keeping the test set unchanged. All models are retrained under each setting with identical hyperparameters.

Table 1 reports the SSIM trends of different methods as the training data ratio decreases from 100% to 10%. As observed, all methods experience performance degradation as the amount of training data diminishes. However, WRF-RAIN consistently demonstrates superior robustness compared to NeRF^2, particularly under highly sparse conditions. When only 10% of the training data is available, WRF-RAIN maintains an SSIM of 0.66, significantly outperforming NeRF^2 (0.59) and remaining close to WRF-GS (0.68), which relies on LiDAR-based initialization.

These results indicate that our WRF-RAIN is more data-efficient and can more effectively leverage sparse observations. We attribute this advantage to the structured Gaussian representation and the coarse-to-fine optimization strategy, which offer stronger inductive bias for modeling wireless propagation compared to pure neural field approaches.

5.4 Ablation Study

To systematically evaluate the contribution of each core component within WRF-RAIN, we conduct a comprehensive ablation study. The study compares five distinct model configurations:

- **Baseline**: Utilizes precise LiDAR point clouds for Gaussian initialization (equivalent to WRF-GS [33]), serving as an upper-bound performance benchmark when perfect geometric priors are available;
- **w/o Annealing**: Employs random initialization but disables the proposed variance annealing schedule;
- **w/o Low-Pass**: Uses random initialization without the progressive low-pass filtering during early training;
- **w/o Density**: Relies on random initialization but removes the adaptive density control mechanism;
- **Full WRF-RAIN**: Our complete proposed framework integrating random initialization, variance annealing, progressive filtering, and adaptive density control.

The quantitative results are summarized in Table 2, which reports three key metrics: the SSIM scores for synthesized spatial spectra (higher is better), the pixel-wise L1 error (lower is better), and the convergence iteration count Conv. (k) required to reach a stable solution (lower is better). Each core component demonstrates a distinct

Table 2. Ablation study of WRF-RAIN components.

Variant	SSIM ↑	L1 Error ↓	Conv. (k) ↓
Baseline (with LiDAR)	0.82	0.042	35
w/o Annealing	0.74	0.068	–
w/o Low-Pass	0.77	0.055	48
w/o Density	0.79	0.048	42
Full WRF-RAIN	**0.81**	**0.045**	**38**

and statistically significant impact on the model's ability to learn an accurate signal field representation from scratch. The Baseline variant achieves the highest SSIM (0.82), confirming the inherent advantage of high-fidelity geometric initialization. However, this heavy reliance on LiDAR data severely limits its applicability in many practical scenarios.

Removing the variance annealing schedule (**w/o Annealing**) leads to the most severe performance degradation (SSIM 0.74) and even fails to converge to a stable solution (denoted by "–" in Conv. (k)). Without this mechanism, the Gaussians retain their large initial variances throughout the entire training process, preventing them from converging to sharp, detail-preserving representations of the scene's electromagnetic structure. This results in overly smoothed predictions and a critical failure to model fine-grained signal variations. These behaviors clearly suggest that variance annealing is crucial for enabling effective coarse-to-fine optimization from random initialization.

Disabling the progressive low-pass filter (**w/o Low-Pass**) causes a notable drop in SSIM to 0.77 and requires more convergence iterations (48k), indicating reduced training stability and reconstruction fidelity. In the absence of this filter, the model is exposed to high-frequency spectral components from the very beginning of training, which induces noisy and unstable gradient signals. This not only impedes stable convergence but also ultimately compromises the fidelity of the synthesized spatial spectra. These observations collectively indicate that progressive low-pass filtering plays a critical role in mitigating training artifacts and stabilizing the optimization process.

The variant without adaptive density control (**w/o Density**) achieves an SSIM of 0.79 and requires 42k iterations to converge. While it benefits from annealing and filtering, the fixed number of Gaussians cannot adapt dynamically to the scene's complexity. It lacks the capacity to allocate more primitives to regions with intricate geometry or complex multipath interactions, leading to a loss of fine-grained detail in these challenging areas and thus lower overall reconstruction accuracy.

The **Full WRF-RAIN** framework, integrating all three proposed core components, attains an SSIM of 0.81, only 0.01 lower than the LiDAR-initialized Baseline. This near-optimal performance is particularly remarkably, as it demonstrates that our method effectively compensates for the lack of geometric priors through a carefully designed training strategy. The synergy between the components is clear: variance annealing enables coarse-to-fine optimization, progressive filtering mitigates training artifacts and stabilizes gradients, and adaptive density control provides the necessary representational flexibility to capture complex signal structures. Collectively, these results quantitatively validate the necessity and effectiveness of each module in achieving robust, prior-free wireless field modeling.

Discussion. The three components of our framework form a coherent coarse-to-fine optimization strategy specifically designed for the wireless domain. Variance annealing addresses the exploration problem by enabling broad spatial sensing without geometric priors. Progressive filtering tackles the stability problem by smoothing the non-convex optimization landscape in the critical early phase. Adaptive density control handles the capacity problem by dynamically increasing representational fidelity where needed. Crucially, these components are

Table 3. Sensitivity analysis of initial Gaussians N_0.

N_0	SSIM@10k ↑	Final SSIM ↑	Iter to 0.75 (k) ↓
256	0.58	0.78	28
512	0.61	0.79	24
1024	0.64	0.81	20
2048	0.67	0.81	18
4096	0.69	0.81	16

sequentially dependent. Variance annealing creates the conditions for stable exploration, yet without filtering, the noisy angular gradients would destabilize the broad Gaussians. Filtering stabilizes gradients, but without annealing, the fixed large variances would prevent fine-detail convergence. Without density control, the model would lack the capacity to represent fine details even after variances have shrunk. This sequential interdependence explains why naive combinations or the removal of any individual component lead to failure or significant degradation.

5.5 Parameter Sensitivity Analysis

To provide a comprehensive understanding of the proposed framework and offer practical guidance for deployment, we conduct a detailed sensitivity analysis on three critical hyperparameters: the initial number of Gaussians N_0, the initial variance σ_0^2, and the variance annealing coefficient λ. These experiments not only illuminate the trade-offs between reconstruction accuracy, training stability, and computational efficiency, but also justify the hyperparameter choices adopted in our main experiments.

Evaluation Metrics. To quantify both reconstruction quality and optimization behavior, we adopt different evaluation metrics tailored to the role of each hyperparameter. For the initial number of Gaussians N_0, which primarily affects early-stage exploration and convergence behavior, we report: (i) **SSIM@10k**, the structural similarity index measured after 10k training iterations to reflect warmup-stage reconstruction quality; (ii) **Final SSIM**, evaluated at the end of training; and (iii) **Iter to 0.75**, defined as the number of training iterations (in thousands) required to reach an SSIM of 0.75, serving as a proxy for convergence speed. For the initial variance σ_0^2 and the variance annealing coefficient λ, which directly influence optimization stability and final reconstruction fidelity, we report: (i) **Final SSIM**, measured at convergence; (ii) **L1 reconstruction error**; and (iii) **Conv. (k)**, the total number of training iterations (in thousands) required to reach convergence.

Initial Number of Gaussians (N_0). The parameter N_0 determines the model's initial capacity before the adaptive density control mechanism engages. We vary N_0 from 256 to 4096 and evaluate its impact on early training performance, final reconstruction accuracy, and convergence speed. As shown in Table 3, a larger N_0 provides more comprehensive coverage of the scene volume during the warmup phase (i.e., first 10k iterations), leading to higher SSIM@10k scores and faster convergence to a reasonable quality threshold (iterations to reach SSIM 0.75). This is because a denser initial Gaussian sampling increases the probability of placing primitives near true signal sources and reflectors, laying a more reliable foundation for subsequent adaptive refinement. However, the final SSIM after full training (including adaptive splitting) shows diminishing returns beyond $N_0 = 1024$: increasing N_0 further does not improve final accuracy but introduces unnecessary memory overhead. This indicates that the adaptive density control mechanism is effective at dynamically growing the Gaussian set when needed, making the final performance less sensitive to the initial count. We select $N_0 = 1024$ as the optimal trade-off, ensuring robust early exploration without excessive computational burden during the initial training stage.

Initial Variance (σ_0^2). This parameter controls the spatial extent of each Gaussian primitive at the start of training, governing the model's ability to explore the scene broadly. We test $\sigma_0^2 \in \{0.25, 0.5, 1.0, 2.0, 4.0\}$ to

Table 4. Sensitivity analysis of initial variance σ_0^2.

σ_0^2	SSIM ↑	L1 Error ↓	Conv. (k) ↓
0.25	0.74	0.065	52
0.50	0.77	0.056	45
1.00	0.81	0.045	38
2.00	0.80	0.048	36
4.00	0.77	0.055	35

cover a wide range of initial spatial spreads. Results in Table 4 reveal a clear trend: moderate initial variances ($\sigma_0^2 = 1.0$ or 2.0) yield the best final SSIM, the lowest L1 error, while maintaining competitive convergence speed. A variance that is too small ($\sigma_0^2 \leq 0.5$) results in Gaussians with overly localized support, hindering their ability to capture global scene structure and leading to slower convergence, higher L1 error, and suboptimal local minima. Conversely, an excessively large variance ($\sigma_0^2 = 4.0$) causes severe over-smoothing in the early stages, blurring initial signal estimates and making it difficult for subsequent variance annealing to recover fine details. The choice of $\sigma_0^2 = 1.0$ provides an effective balance, offering sufficient coverage for initial coarse scene understanding while retaining the capacity for later refinement.

Variance Annealing Coefficient (λ). The coefficient λ in Eq. (8) dictates the speed at which Gaussian variances shrink, governing the smoothness of the coarse-to-fine optimization transition that is critical to WRF-RAIN's stability. We explore λ values on a logarithmic scale from 1×10^{-6} to 1×10^{-4}. As presented in Table 5, the optimal performance is achieved with $\lambda = 5 \times 10^{-5}$. An overly small λ (1×10^{-6}) decays the variance too slowly, causing the model to waste training iterations on coarse optimization and fail to capture high-frequency spatial spectrum details within the allotted iteration budget, even failing to converge to a stable solution (denoted by "–" in Conv. (k)). An overly large λ (1×10^{-4}), in contrast, shrinks the variances too rapidly: this prematurely confines the Gaussians to narrow spatial regions before a reliable coarse scene layout has been established, often trapping the optimization in poor local minima and reducing final SSIM. The selected value ensures a smooth and controlled progression, allowing the model to first lock in the major wireless propagation paths before dedicating capacity to modeling finer reflections, diffractions, and signal variations.

5.6 Initialization Robustness Analysis

To further evaluate the robustness of WRF-RAIN against different initialization strategies, we conduct experiments under various Gaussian initialization schemes. This analysis aims to investigate whether the proposed framework consistently converges to stable and accurate solutions when the initial spatial distribution of Gaussians varies.

Specifically, we consider three representative initialization strategies:

- **Uniform Random**: Gaussian centers are uniformly sampled within the bounded scene volume (default setting);
- **Grid Initialization**: Gaussian centers are initialized on a regular 3D grid covering the scene;

Table 5. Sensitivity analysis of annealing coefficient λ.

λ	SSIM ↑	L1 Error ↓	Conv. (k) ↓
1×10^{-6}	0.75	0.060	–
1×10^{-5}	0.78	0.052	48
5×10^{-5}	0.81	0.045	38
1×10^{-4}	0.76	0.057	32

Table 6. Initialization robustness under different Gaussian distributions.

Initialization	SSIM ↑	L1 Error ↓	Conv. (k) ↓
Uniform Random	0.81	0.045	38
Grid	0.80	0.047	40
Biased Random	0.78	0.051	44

- **Biased Random**: Gaussian centers are randomly sampled but with a spatial bias toward a subset of the scene (to simulate uneven coverage).

All configurations use the same number of initial Gaussians ($N_0 = 1024$) and identical training settings. The results are summarized in Table 6.

As shown in the table, WRF-RAIN demonstrates strong robustness across different initialization strategies. The default uniform random initialization attains the best performance, while grid-based initialization yields comparable results with slightly slower convergence due to its less flexible spatial coverage in early stages. Even under biased random initialization, where certain regions of the scene are underrepresented, the model is still able to converge to a reasonable solution, albeit with a modest degradation in reconstruction accuracy. This indicates that the proposed optimization framework can effectively recover from imperfect initial coverage. We attribute this robustness to the proposed coarse-to-fine optimization design. In particular, variance annealing enables broad initial exploration of the scene, while progressive low-pass filtering stabilizes early-stage training and prevents the model from being trapped in poor local minima. Furthermore, the adaptive density control mechanism allows the model to dynamically refine underrepresented regions during training.

Overall, these results confirm that WRF-RAIN is not sensitive to the specific choice of Gaussian initialization, and can reliably learn meaningful wireless field representations even under imperfect or biased initial conditions.

5.7 Real-World Multi-Frequency Localization Experiment

To validate the frequency-aware module, we conduct evaluations using a local dataset collected from a multi-frequency RSSI localization system. The experiment is designed to demonstrate the framework's ability to predict RSSI values from transmitter positions at varying carrier frequencies, a capability essential for practical frequency-aware wireless channel modeling.

5.7.1 Experimental Setup. As shown in Fig. 10, the system employs an Impinj Speedway Revolution R420 reader, which is compliant with the EPC Gen2 standard. The reader is connected to a host machine via Ethernet, enabling real-time data acquisition at high throughput. A circularly polarized antenna of dimensions 22.5 cm×22.5 cm×4 cm with a gain of 8 dBi is used to ensure robust signal reception across the test environment. The RFID tag used in the system is an H47 model, which features an Impinj Monza 4 chip. This tag has a physical size of 50 mm × 50 mm and employs dual-linear polarization, providing polarization diversity with respect to the circularly polarized receiving antenna.

RSSI measurements are recorded at four different carrier frequencies: 921.375 MHz, 922.375 MHz, 923.375 MHz, and 924.375 MHz. These frequencies are selected to emulate realistic scenarios where multiple channels coexist in dense wireless deployments. Each measurement consists of the 3D position of the transmitter (TX) and the corresponding RSSI values received by the fixed receiver (RX). The dataset contains measurements from 200 distinct TX positions, and for each position the RSSI values at all four frequencies are recorded. We adopt a leave-one-frequency-out split: three frequencies are used for training, while the remaining frequency is held out for testing. This procedure is repeated for each of the four frequencies, yielding four independent test configurations.

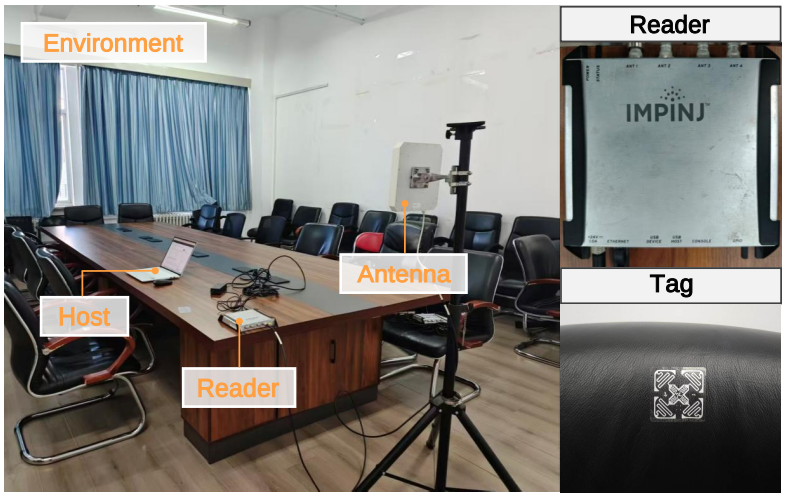


Fig. 10. Real-world localization experiment setup.

5.7.2 Results and Analysis. The primary goal of this experiment is to evaluate WRF-RAIN's generalization capability across different frequencies in real-world scenarios. Performance is measured by the RSSI prediction error, defined as the absolute difference between the predicted and actual RSSI values (in dB). Existing wireless signal modeling methods, such as NeRF^2 and WRF-GS, lack frequency-aware modules and are inherently limited to single-frequency modeling. As such, we are unable to compare our method directly with these approaches. Instead, the focus of this experiment is to validate the effectiveness of our method in real-world cross-frequency scenarios.

We report the results in Table 7, which lists the median RSSI prediction error together with the convergence speed for each test frequency. Convergence speed is quantified by the number of training iterations (in thousands, denoted as Conv. (k)) required to reach a stable solution. The results show that WRF-RAIN achieves consistently low prediction errors across all test frequencies, with the best performance at 923.375 MHz (median error of 2.3 dB). These results confirm that the frequency-aware representation enables effective adaptation to different carrier frequencies, while adaptive density control and progressive low-pass filtering jointly contribute to training stability and accurate predictions under varying signal conditions.

In summary, the above experiment validates the effectiveness of the frequency-aware extension of WRF-RAIN, demonstrating that the framework can reliably process multi-frequency data for wireless radiation field reconstruction. Such generalization is crucial for next-generation wireless systems, where operating frequency may vary due to environmental factors, spectrum regulations, or system design choices.

6 FIELD STUDY: BLE LOCALIZATION

In this section, we demonstrate the practical utility of WRF-RAIN through a field study on Bluetooth Low Energy (BLE) localization in resource-constrained single-antenna systems. Unlike MIMO systems which exploit spatial diversity through multiple antennas, BLE devices typically employ a single omnidirectional antenna and rely solely on Received Signal Strength Indicator (RSSI) measurements for localization. Traditional approaches, such as exhaustive fingerprinting or simplified analytical propagation models, face significant challenges in achieving both scalability and accuracy in complex indoor environments due to multipath effects and non-line-of-sight conditions.

Table 7. Performance of WRF-RAIN on the multi-frequency RSSI localization dataset.

Test Frequency (MHz)	RSSI Prediction Error (dB) ↓	Conv. (k) ↓
921.375	2.5	42
922.375	2.4	40
923.375	2.3	39
924.375	2.6	41

To address these limitations, we leverage the fundamental premise that RSSI measurements are inherently a function of the physical environment–shaped by signal attenuation, reflections, and diffraction along the propagation path. Consequently, there exists a deterministic mapping between a BLE node's spatial location and the RSSI vector observed across distributed gateways. We adapt WRF-RAIN to learn this complex mapping by retraining the model to synthesize the spatial distribution of RSSI values, effectively simulating BLE signal propagation in cluttered indoor spaces. The trained model can then generate large-scale, high-fidelity synthetic RSSI datasets, significantly enhancing the performance of downstream fingerprinting-based localization systems without requiring additional costly field measurements.

6.1 Experimental Setup

We conduct BLE localization experiments using the publicly available NeRF^2 dataset [38], captured in a representative indoor office environment. The dataset comprises RSSI measurements from mobile BLE beacons collected by 21 fixed-location gateways, totaling 6,000 signal-location pairs (4,800 for training, 1,200 for testing). Each data sample includes a 2D transmitter coordinate and the corresponding 21-dimensional RSSI vector $R(x)$, where each element represents the received signal strength at a specific gateway. Missing or undetected signals are marked as -100 dBm.

6.2 RSSI Prediction Accuracy

We first evaluate WRF-RAIN's capability to accurately predict RSSI vectors given arbitrary transmitter locations–a critical precursor to reliable localization. During inference, the model takes a 2D position x as input and renders the expected RSSI vector $R(x) \in \mathbb{R}^M$, where $M = 21$ denotes the number of gateways. This task is particularly demanding due to the intricate interplay of signal propagation phenomena in indoor settings.

We benchmark WRF-RAIN against several representative baselines, including methods that rely on geometric priors:

- **WRF-GS** [33]: A state-of-the-art method that initializes 3D Gaussians using precise LiDAR point clouds;
- **NeRF2** [38]: A volumetric neural renderer for wireless fields;
- **MRI** [26]: A classical radio propagation model based on log-distance path loss interpolation;
- **CGAN** [17]: A conditional generative adversarial network trained to predict RSSI vectors without explicit physical modeling.

The prediction error is quantified as the mean absolute deviation between predicted and ground-truth RSSI values across all gateways. Fig. 11 presents the CDF of RSSI prediction errors for all methods. WRF-RAIN achieves a median error of 2.3 dB, which is comparable to the LiDAR-dependent WRF-GS (2.2 dB) and outperforms NeRF^2 (2.6 dB), CGAN (4.3 dB), and MRI (5.9 dB). This performance is particularly noteworthy given that WRF-RAIN operates without any geometric prior, while WRF-GS benefits from high-fidelity LiDAR initialization. The error distributions, further visualized through box plots in Fig. 12, demonstrate that WRF-RAIN maintains consistently low prediction errors with a tight interquartile range, confirming its robustness in capturing complex signal-environment interactions.

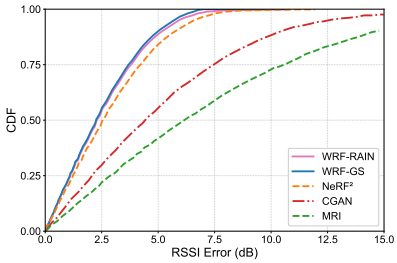


Fig. 11. CDF of RSSI prediction errors.

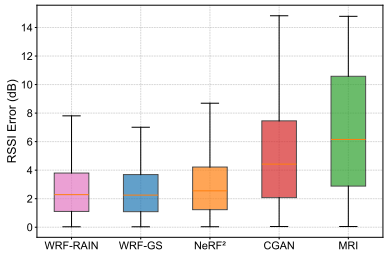


Fig. 12. Box plots of RSSI errors.

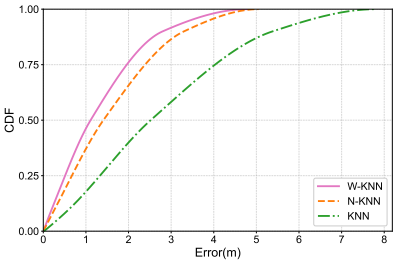


Fig. 13. CDF of localization errors.

6.3 Localization Results

We evaluate the downstream localization performance using a K-Nearest Neighbors (KNN) fingerprinting approach under three distinct training data configurations:

- **KNN (Baseline)**: Uses only the original 4.8k real-world RSSI samples from the dataset for nearest-neighbor matching;
- **NeRF^2-Augmented KNN (N-KNN)**: Augments the fingerprint database with 20k synthetic RSSI samples generated by a trained NeRF^2 model at randomly sampled 2D positions;
- **WRF-RAIN-Augmented KNN (W-KNN)**: Similarly augments the database with 20k synthetic samples generated by our WRF-RAIN model.

Given a query RSSI vector from the test set, the KNN algorithm computes the Euclidean distance to all entries in the fingerprint database, retrieves the K nearest neighbors (with $K = 5$ in our experiments), and estimates the transmitter's position as the distance-weighted average of their coordinates [15].

Fig. 13 illustrates the CDF of localization errors across all 1.2k test samples. Both data augmentation strategies significantly improve localization accuracy over the baseline KNN. Crucially, WRF-RAIN-augmented fingerprinting achieves a median localization error of 1.11 m, surpassing both NeRF^2-augmentation (1.40 m) and the baseline (2.55 m). This superior performance directly stems from WRF-RAIN's ability to generate more accurate and physically consistent synthetic RSSI samples, which enrich the fingerprint database and improve the discriminative power of the KNN matcher.

The results underscore a key practical advantage of WRF-RAIN: by generating high-fidelity synthetic data, it dramatically reduces the burden of manual site surveying while simultaneously improving localization accuracy. This makes it particularly suitable for large-scale environments where exhaustive fingerprinting is infeasible. Importantly, the minimal performance gap between WRF-RAIN and the LiDAR-dependent WRF-GS highlights the efficacy of our random initialization strategy combined with progressive low-pass filtering and adaptive density control. By eliminating the dependency on expensive geometric priors, WRF-RAIN offers a more flexible and scalable solution for real-world deployment, particularly in scenarios where LiDAR or detailed architectural plans are unavailable.

The BLE localization experiment serves as a proof-of-concept validation of the utility of WRF-RAIN-generated synthetic data for downstream tasks. While we evaluate using KNN fingerprinting to isolate the effect of data augmentation, the synthetic RSSI vectors produced by our method are algorithm-agnostic and can be directly used to augment training sets for more advanced localization methods, such as deep neural networks. This seamless extensibility highlights the broader potential of our generative framework beyond classical fingerprinting, providing a versatile data-augmentation backbone for a wide range of wireless sensing applications where labeled real-world data remains a bottleneck.

7 CONCLUSION

This work addresses a fundamental limitation of neural network-based wireless radiance field reconstruction: the heavy reliance on a high-quality initial point cloud, which is often impractical in real-world scenarios due to environmental complexity, hardware constraints, and inefficient data acquisition. The proposed WRF-RAIN framework, with flexible random initialization, progressive low-pass filtering, and adaptive density control, synergistically compensates for data imperfections, enabling robust reconstruction from sparse, noisy, or low-fidelity input. We evaluate our system through extensive experiments across diverse benchmark datasets and practical communication scenarios, demonstrating that WRF-RAIN achieves comparable accuracy to state-of-the-art methods without high-precision point clouds, thus slashing initialization costs and breaking prior geometric dependency. By providing a scalable, real-time solution, this work advances practical wireless channel modeling and has significant implications for the development of next-generation communication systems. While WRF-RAIN exhibits strong robustness under varying data densities and initialization strategies, our current evaluation focuses primarily on static environments. Extending validation to more diverse and dynamic settings (e.g., different layouts, materials, and scales) remains an important direction for future work to establish broad generalizability.

Acknowledgments

This work was supported in part by the Beijing Natural Science Foundation under Grant No. 4262031 and the Beijing Institute of Technology Research Fund Program for Young Scholars. The work of Fan Li was supported in part by the National Natural Science Foundation of China under Grant No. 62372045.

References

[1] Jeffrey G. Andrews, Stefano Buzzi, Wan Choi, Stephen V. Hanly, Angel Lozano, Anthony C. K. Soong, and Jianzhong Charlie Zhang. 2014. What Will 5G Be? *IEEE Journal on Selected Areas in Communications* 32, 6 (2014), 1065–1082.

[2] Muhammad Zeeshan Asghar, Saima Arosha Memon, and Jari Hämäläinen. 2022. Evolution of Wireless Communication to 6G: Potential Applications and Research Directions. *Sustainability* 14, 10 (2022), 6356.

[3] Patrick Billingsley. 1995. *Probability and Measure* (3rd ed.). John Wiley & Sons, New York. Chapter 1: "Probability".

[4] Wei Gong, Haoyu Wang, Siyi Li, and Si Chen. 2024. GLAC: High-Precision Tracking of Mobile Objects With COTS RFID Systems. *IEEE/ACM Transactions on Networking* 32, 3 (2024), 2331–2343.

[5] Danping He, Bo Ai, Ke Guan, Ling Wang, Zhangdui Zhong, and Thomas Kürner. 2018. The Design and Applications of High-Performance Ray-Tracing Simulation Platform for 5G and Beyond Wireless Communications: A Tutorial. *IEEE Communications Surveys & Tutorials* 21, 1 (2018), 10–27. 1st Quarter.

[6] Xiaonan Hui and Edwin C Kan. 2019. Radio ranging with ultrahigh resolution using a harmonic radio-frequency identification system. *Nature Electronics* 2, 3 (2019), 125–131.

[7] Jiyu Jiao, Xiaojun Wang, Chenpei Han, Yuhua Huang, and Yizhuo Zhang. 2024. Dynamic Indoor Fingerprinting Localization based on Few-Shot Meta-Learning with CSI Images. arXiv:2401.05711 [cs.NI] https://arxiv.org/abs/2401.05711

[8] Jaewoo Jung, Jisang Han, Honggyu An, Jiwon Kang, Seonghoon Park, and Seungryong Kim. 2024. Relaxing Accurate Initialization Constraint for 3D Gaussian Splatting. arXiv:2403.09413 [cs.CV] https://arxiv.org/abs/2403.09413

[9] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics* 42, 4, Article 139 (2023), 14 pages.

[10] Diederik P Kingma and Max Welling. 2022. Auto-Encoding Variational Bayes. arXiv:1312.6114 [stat.ML] https://arxiv.org/abs/1312.6114

[11] Haoliang Lu, Christian Vattheuer, Baharan Mirzasoleiman, and Omid Abari. 2024. NeWRF: A Deep Learning Framework for Wireless Radiation Field Reconstruction and Channel Prediction. In *Proceedings of the International Conference on Machine Learning (ICML)*. Article 1345, 13 pages.

[12] Yunfei Ma, Nicholas Selby, and Fadel Adib. 2017. Minding the Billions: Ultra-wideband Localization for Deployed RFID Tags. In *Proceedings of the Annual International Conference on Mobile Computing and Networking (MobiCom)*. 248–260.

[13] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2021. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. *Commun. ACM* 65, 1 (2021), 99–106.

[14] Andreas F. Molisch. 2011. *Wireless Communications*. John Wiley & Sons, Chichester, UK.

[15] Lionel M. Ni, Yunhao Liu, Yiu Cho Lau, and Abhishek P. Patil. 2003. LANDMARC: Indoor Location Sensing Using Active RFID. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom)*. 407–415.
[16] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. 2019. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 165–174.
[17] Fernando Parralejo, Francisco J. Aranda, José A. Paredes, Fernando J. Álvarez, and Jesús Morera. 2021. Comparative Study of Different BLE Fingerprint Reconstruction Techniques. In *Proceedings of the IEEE International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. 1–8.
[18] Thomas Porter and Tom Duff. 1984. Compositing Digital Images. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*. 253–259.
[19] Alec Radford, Luke Metz, and Soumith Chintala. 2015. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. arXiv:1511.06434 [cs.LG] https://arxiv.org/abs/1511.06434
[20] Thinal Raj, Fazida Hanim Hashim, Aqilah Baseri Huddin, Mohd Faisal Ibrahim, and Aini Hussain. 2020. A Survey on LiDAR Scanning Mechanisms. *Electronics* 9, 5 (2020), 741.
[21] Theodore S. Rappaport. 2002. *Wireless Communications: Principles and Practice* (2nd ed.). Prentice Hall, Upper Saddle River, NJ, USA.
[22] Walid Saad, Mehdi Bennis, and Mingzhe Chen. 2020. A Vision of 6G Wireless Systems: Applications, Trends, Technologies, and Open Research Problems. *IEEE Network* 34, 3 (2020), 134–142.
[23] Tapan K Sarkar, Zhong Ji, Kyungjung Kim, Ahmed Medouri, and Magdalena Salazar-Palma. 2003. A Survey of Various Propagation Models for Mobile Communication. *IEEE Antennas and Propagation Magazine* 45, 3 (2003), 51–82.
[24] Johannes L Schönberger and Jan-Michael Frahm. 2016. Structure-from-Motion Revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 4104–4113.
[25] Wasswa Shafik. 2024. An Overview of Computational Modeling and Simulations in Wireless Communication Systems. In *Computational Modeling and Simulation of Advanced Wireless Communication Systems*. Springer, Cham, Switzerland, 8–40.
[26] Hojung Shin, Yohan Chon, Yonghun Kim, and Hojung Cha. 2014. MRI: Model-based Radio Interpolation for Indoor War-Walking. *IEEE Transactions on Mobile Computing* 14, 6 (2014), 1231–1244.
[27] Harsh Tataria, Mansoor Shafi, Andreas F. Molisch, Mischa Dohler, Henrik Sjöland, and Fredrik Tufvesson. 2021. 6G Wireless Systems: Vision, Requirements, Challenges, Insights, and Opportunities. *Proc. IEEE* 109, 7 (2021), 1166–1199.
[28] Hind Taud and Jean-François Mas. 2017. Multilayer Perceptron (MLP). In *Geomatic Approaches for Modeling Land Change Scenarios*. Springer International Publishing, Cham, 451–455.
[29] Muhammad Emad ud din. 2024. Hybrid CNN-LSTM based Indoor Pedestrian Localization with CSI Fingerprint Maps. arXiv:2412.13601 [cs.NI] https://arxiv.org/abs/2412.13601
[30] Muhammad Umer, Muhammad Ahmed Mohsin, Ahsan Bilal, and John M. Cioffi. 2026. Neural Gaussian Radio Fields for Channel Estimation. arXiv:2508.11668 [eess.SP] https://arxiv.org/abs/2508.11668
[31] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. 2004. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612.
[32] Chaozheng Wen, Jingwen Tong, Yingdong Hu, Zehong Lin, and Jun Zhang. 2025. Neural Representation for Wireless Radiation Field Reconstruction: A 3D Gaussian Splatting Approach. arXiv:2412.04832 [cs.NI] https://arxiv.org/abs/2412.04832
[33] Chaozheng Wen, Jingwen Tong, Yingdong Hu, Zehong Lin, and Jun Zhang. 2025. WRF-GS: Wireless Radiation Field Reconstruction with 3D Gaussian Splatting. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*. 1–10.
[34] Kang Yang, Yuning Chen, and Wan Du. 2025. GWRF: A Generalizable Wireless Radiance Field for Wireless Signal Propagation Modeling. arXiv:2502.05708 [cs.NI] https://arxiv.org/abs/2502.05708
[35] Kang Yang, Wan Du, and Mani Srivastava. 2025. Scalable 3D Gaussian Splatting-Based RF Signal Spatial Propagation Modeling. In *Proceedings of the ACM Conference on Embedded Networked Sensor Systems (SenSys)*. 680–681.
[36] Lei Yang, Yekui Chen, Xiang-Yang Li, Chaowei Xiao, Mo Li, and Yunhao Liu. 2014. Tagoram: Real-Time Tracking of Mobile RFID Tags to High Precision Using COTS Devices. In *Proceedings of the Annual International Conference on Mobile Computing and Networking (MobiCom)*. 237–248.
[37] Zhengqing Yun and Magdy F Iskander. 2015. Ray Tracing for Radio Propagation Modeling: Principles and Applications. *IEEE Access* 3 (2015), 1089–1100.
[38] Xiaopeng Zhao, Zhenlin An, Qingrui Pan, and Lei Yang. 2023. NeRF2: Neural Radio-Frequency Radiance Fields. In *Proceedings of the Annual International Conference on Mobile Computing and Networking (MobiCom)*. 1–15.
[39] Xiaopeng Zhao, Zhenlin An, Qingrui Pan, and Lei Yang. 2025. Frequency-Aware Neural Radio-Frequency Radiance Fields. *IEEE Transactions on Mobile Computing* 24, 11 (2025), 11401–11415.
[40] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. 2002. EWA Splatting. *IEEE Transactions on Visualization and Computer Graphics* 8, 3 (2002), 223–238.