

# Active Sampling and Gaussian Reconstruction for Radio Frequency Radiance Field

Chi-Shiang Gau, Xingyu Chen, Wuqiong Zhao, *Graduate Student Member, IEEE*,  
Xinyu Zhang, *Senior Member, IEEE*, and Tara Javidi, *Fellow, IEEE*

**Abstract**—Reconstructing Radio-Frequency (RF) Radiance Fields is essential for modeling signal propagation in wireless systems but remains challenging due to complex wave phenomena such as reflection, scattering, and diffraction—particularly when the environment’s geometry and material properties are unknown. Recent neural approaches, such as NeRF2, employ large differentiable models for RF reconstruction but require dense spatial sampling and compute-intensive per-environment training, limiting their practicality in dynamic or quasi-static scenarios.

We propose a training-free, physics-assisted alternative based on a Gaussian Process (GP) with an adaptive spatio-temporal kernel, whose hyperparameters are estimated via maximum likelihood. Our model incorporates prior physical knowledge by representing the received signal as a linear combination of rays from different angles, allowing the GP to implicitly capture the spatial structure of the RF field in a physics-assisted manner. The kernel captures local structural and temporal patterns, while the GP framework naturally provides predictive uncertainty. Posterior variance guides active sampling, prioritizing regions of high information content to efficiently collect the most informative measurements.

Experiments show that our method substantially reduces the number of required RF observations while maintaining high reconstruction fidelity, enabling real-time adaptation in dynamic environments and offering a scalable, data-efficient solution for RF field estimation.

**Index Terms**—Active learning, Gaussian process, local kernel, fast adaptation, radio-frequency radiance field.

## I. INTRODUCTION

The modeling of electromagnetic fields in complex indoor and outdoor environments, particularly at radio frequency (RF) bands, has become increasingly critical for emerging applications in network planning and wireless communications [1]–[4]. RF signal propagation is fundamentally shaped by material properties and wave phenomena including reflection, refraction, and diffraction from object geometries in the scene. Similar to view synthesis for cameras, RF Radiance Field modeling aims to determine the total received power at any given spatial position.

A recent work, NeRF2 [5], explores adapting neural-based novel view synthesis (NVS) techniques—such as Neural Radiance Fields (NeRF) [6]—from the visual to the radio domain,

This work was supported in part by the sponsors of the Center for Wireless Communications at the University of California, San Diego, in part by the National Science Foundation under Grants CNS-2148313, CNS-2403124, and CNS-2312715, and in part by funds from federal agencies and industry partners through the Resilient and Intelligent NextG Systems (RINGS) program.

The authors are with the University of California, San Diego, CA 92093 USA (e-mail: {cgau, xic063, wqzhao, xyzhang, tjavidi}@ucsd.edu).

yielding promising initial results. However, two significant challenges limit the practical deployment of such approaches. First, due to the limited aperture size of radio antennas compared to optical sensors, these methods require an exceptionally high sampling density—on the order of 200 samples per square foot—compared to the dozens of images per scene typically sufficient for optical NVS. Second, RF systems often require real-time operation under stringent computational constraints, rendering neural-network-based reconstruction methods impractical for most real-time RF field reconstruction.

More recently, extensions of 3D Gaussian Splatting (GS) [7], originally developed to accelerate radiance-field training in vision, have been adapted to RF radiance-field reconstruction [8]. In these approaches, the RF field is represented using Gaussian primitives—deterministic Gaussian functions that encode local properties such as spatial location, extent, and signal strength—and the field is reconstructed by aggregating the contributions of all primitives. GSpaRC [9] further demonstrates a low-latency GS-based approach for RF signal reconstruction. While these methods were developed independently and appeared after the completion of our work, as shown later in our paper, they still require a large number of measurements to achieve accurate reconstruction. Moreover, neither NeRF2 nor RF variants of 3D Gaussian Splatting provide uncertainty estimates for the reconstructed fields. As a consequence, these methods do not support uncertainty-aware sequential decision-making or active sampling, nor do they offer a principled assessment of reconstruction reliability or guidance on which additional measurements would be most informative. These limitations significantly reduce their applicability in adaptive or resource-constrained settings.

To address these challenges, we propose a Gaussian reconstruction framework based on Gaussian processes for modeling RF Radiance Fields. Our key insight is to represent the RF field as a spatially continuous Gaussian random field, populated with virtual signal sources modeled as Gaussian random variables. This formulation enables accurate reconstruction from sparse RF measurements while explicitly quantifying uncertainty. Specifically, for any target location, our method provides not only a prediction of the RF Radiance Field—and thus received signal power—but also an associated predictive variance, enabling principled uncertainty-aware inference and adaptive measurement strategies.

The practical contribution of our work is to incorporate active sampling by taking measurements at the positions with the highest uncertainty, which significantly reduces the required sampling density. Moreover, with this approach, our method

can adapt to quasi-dynamic scene changes by collecting only a few measurements and comparing them with the past RF Radiance Field. When changes are detected, our method guides the agent in identifying which additional measurements are needed, allowing the RF Radiance Field to be updated accordingly. This capability distinguishes our method from existing approaches and makes it more suitable for real-world applications.

We evaluate our method in both synthetic and real environments, demonstrating superior performance in terms of computational and sampling efficiency, while maintaining high reconstruction accuracy. As a result, our proposed method is particularly advantageous when the number of available samples is limited.

Our contributions can be summarized as follows:

- **No pre-training required:** Our proposed method does not require any pre-training to make predictions at arbitrary target positions. Instead, the kernel hyperparameters of the Gaussian reconstruction model are estimated via maximum likelihood during model fitting. This procedure avoids environment-specific neural network training and eliminates the need for iterative optimization over large neural models.
- **An uncertainty model is provided:** We provide an uncertainty model for any given position in the RF Radiance Field, allowing users to determine if additional observations are needed to reduce uncertainty in specific areas—an ability that neither NeRF2 nor Gaussian Splatting methods offer.
- **Reduction in the number of samples:** With the help of our uncertainty model, highly informative samples are prioritized, reducing the number of measurements needed to achieve performance comparable to NeRF2. Our experiments show a 30% to 60% reduction in required samples on synthetic data. On real-world collected data, NeRF2 fails to achieve comparable accuracy—even when trained on the full dataset—so does GSpaRC which also struggles. In contrast, our method achieves lower error with significantly fewer samples across both small and larger-scale environments.
- **Fast adaptation to quasi-dynamic scene changes:** Our proposed method efficiently adapts to changes in the scene by focusing on learning the variations to reconstruct the new RF Radiance Field.

## II. RELATED WORKS

Neural representations for scene understanding have revolutionized computer vision and graphics [10]–[16], particularly since the seminal work of Neural Radiance Fields (NeRF) [6]. NeRF demonstrated that complex 3D scenes can be represented through implicit neural networks that encode both spatial geometry (via volumetric density) and view-dependent appearance (via directional radiance). This enables photorealistic novel view synthesis without explicit 3D modeling. Recent work such as NeRF2 [5] has extended neural scene representations beyond visible light to RF domains. Subsequent approaches have explored different representation strate-

gies. For example, WiNeRT [17] bridges neural and physics-based methods through learnable material-specific reflection parameters, while RFCanvas [18] introduces a fully explicit representation that can adapt to RF field dynamics. The most recently developed Gaussian Splatting methods, such as RF-3DGS [8] and GSpaRC [9], represent RF fields using deterministic Gaussian primitives, achieving improved reconstruction efficiency over neural networks. However, these methods are fundamentally limited by their reliance on optimization-based training from dense RF measurements and do not provide uncertainty estimates, making them impractical for active learning scenarios where sample efficiency is crucial.

In contrast, we model RF Radiance Fields using Gaussian processes, which provide an explicit measure of reconstruction uncertainty. This allows active learning, where measurements are prioritized in regions of high uncertainty, significantly reducing the number of samples required. As a result, our approach is particularly well-suited for real-world applications in which data collection is costly or time-constrained. Table I summarizes the key characteristics of different RF Radiance Field reconstruction paradigms, including neural-network-based methods, Gaussian Splatting-based methods, and our Gaussian Process approach.

## III. PRELIMINARY

In this section, we present an overview of the system model and the fundamental properties of Gaussian random processes.

### A. Notations

We use boldface letters to represent random variables and random vectors.  $(\cdot)$  denotes a vector, while capital letters represent matrices.  $(\cdot)^T$  denotes the transpose operation.  $\mathbb{E}$  is the expectation operator.

### B. System Model and Ray Structure

We assume that the scene remains static during a given time slot. We model each position in the room as an omnidirectional virtual signal source. Specifically, for time slot  $t$ , the virtual signal at position  $p$ , is denoted as  $x_t(p)$ . Furthermore, the received signal power at any position is modeled as a sum of  $R$  different rays from different angles. The received signal from each ray is modeled as a linear combination of the  $N$  virtual signal source sampled on the ray, as shown in Fig. 1.




Fig. 1. The Radiance Field in the RF domain, as measured by an omnidirectional antenna, exhibits a simple radial ray structure. Additionally, the RF signal experiences significant path loss, which simplifies the inverse problem by limiting long-range interference.

TABLE I  
COMPARISON OF REPRESENTATIVE RF RADIANCE FIELD RECONSTRUCTION PARADIGMS, INCLUDING NEURAL-NETWORK-BASED METHODS, GAUSSIAN SPLATTING-BASED METHODS, AND OUR GAUSSIAN PROCESS FRAMEWORK

Aspect	Neural-network-based (e.g., NeRF2 [5])	Gaussian Splatting-based (e.g., GSpaRC [9])	Gaussian Process (Ours)
Representation	Neural network	Deterministic Gaussian primitives	Gaussian Process
Training	Gradient descent	Gradient descent	ML estimation of kernel hyperparameters
Uncertainty	No	No	Yes (posterior variance)
Sample efficiency	Low	Low	High
Active sampling	No	No	Yes
Dynamic adaptation	Retraining required	Retraining required	Incremental update

For time slot  $t$  the received signal can be written as:

$$\mathbf{y}_t(p) \triangleq \sum_{r=1}^R \sum_{n=1}^N \alpha_{r,n} \cdot \mathbf{x}_t(p + s_{r,n}). \quad (1)$$

Here,  $\mathbf{y}_t(p)$  denotes the received signal power at position  $p$ . The terms  $\mathbf{x}_t(p + s_{r,n})$  represent virtual signal sources located at positions  $(p + s_{r,n})$ . The attenuation from each virtual source to the receiver is captured by  $\alpha_{r,n}$ . In this work, the receiver is assumed to use a single omnidirectional antenna, so the attenuation term reflects only propagation loss. To support different receiver configurations, one can incorporate the corresponding antenna pattern or array factor into the attenuation term  $\alpha_{r,n}$ .

The notation  $(p + s_{r,n})$  refers to the position of the  $n$ -th sample on the  $r$ -th ray for the receiver placed at  $p$ . Specifically, there are  $R \times N$  virtual signal sources that contribute to the received signal at  $p$  and their corresponding positions are given by:

$$(p + s_{r,n}) \triangleq p + n \cdot d \cdot \begin{bmatrix} \cos\left(2\pi \frac{r}{R}\right) \\ \sin\left(2\pi \frac{r}{R}\right) \end{bmatrix}, \quad (2)$$

for  $r = 1, 2, \dots, R$  and  $n = 1, 2, \dots, N$

where  $d$  is the sampling resolution on the ray. The propagation loss for the virtual signal source position  $(p + s_{r,n})$  toward receiver position  $p$  is:

$$\alpha_{r,n} \triangleq \frac{\beta}{\|s_{r,n}\|} = \frac{\beta}{n \cdot d} \quad (3)$$

where  $\beta$  is the parameter for the transmission medium. In this work, we adopt a simple path loss model as in [5].

A more accurate model could be used if additional information, such as the transmission frequency, is available.

For notation simplicity, we rewrite the received signal model for time slot  $t$  (1) into vector form:

$$\mathbf{y}_t(p) \triangleq \vec{\alpha}^T \cdot \vec{\mathbf{x}}_t(p). \quad (4)$$

Notice that, although the receiver is located within the room, some virtual signal sources may be located outside. This approach accounts for situations where the receiver is placed at the room's boundary or in a corner.

### C. Gaussian Random Field

For a given time slot  $t$ , we model the RF Radiance Field as a Gaussian random process, with the virtual signal sources—components of the RF Radiance Field—modeled as

Gaussian random variables. A Gaussian process is defined by its mean and covariance (also known as the kernel):

- **Mean:** For any position  $p$ , the virtual signal is modeled as a zero-mean Gaussian random variable.

$$\mathbb{E}[\mathbf{x}_t(p)] = 0 \quad (5)$$

- **Covariance:** The covariance between each pair of virtual signal sources is given by:

$$\text{Cov}(\mathbf{x}_t(p_i), \mathbf{x}_t(p_j)) = \sigma^2 \exp\left(-\frac{1}{2l^2} \|p_i - p_j\|^2\right) \quad (6)$$

The value of  $\sigma^2$  represents the variance of the virtual signal sources, which are selected based on the true transmission power of the device within the room. The length-scale parameter  $l$ , controls the rate of variation of the signal; smaller values of  $l$  indicate a sharper change in the function. In a later section, we will demonstrate how to estimate this parameter. Note that (6) represents the Radial Basis Function (RBF) kernel.

For information on other types of kernel functions and guidance on selecting the appropriate kernel, we refer the reader to [19].

With the above definitions,  $\vec{\mathbf{x}}_t(p)$ , the collection of virtual signal sources contributing to  $\mathbf{y}_t(p)$  is a Gaussian random vector:

$$\vec{\mathbf{x}}_t(p) \sim N(\vec{0}, \Sigma_p). \quad (7)$$

Here,  $\Sigma_p$  is the covariance matrix where the  $(i, j)$ -th component represents the covariance between the  $i$ -th and  $j$ -th components of  $\vec{\mathbf{x}}_t(p)$ .

Also,  $\mathbf{y}_t(p)$  is a linear combination of Gaussian random variables. Therefore, from (4), it is itself a Gaussian random variable with the following distribution:

$$\mathbf{y}_t(p) \sim N(0, \vec{\alpha}^T \Sigma_p \vec{\alpha}). \quad (8)$$

### D. Conventional Gaussian Prediction

Next, we describe how to predict the received signal power  $\mathbf{y}_t(p_{\text{Target}})$  at a given target position  $p_{\text{Target}}$  using  $M$  observations  $\mathbf{y}_t(p_1), \mathbf{y}_t(p_2), \dots, \mathbf{y}_t(p_M)$ , collected at positions  $\mathcal{P} = \{p_1, \dots, p_M\}$ , within the same time slot  $t$ . Since each observation is a Gaussian random variable, the collection of  $M$  observations can be represented as a Gaussian random vector:

$$\vec{\mathbf{y}}_t(\mathcal{P}) \triangleq \begin{bmatrix} \mathbf{y}_t(p_1) \\ \mathbf{y}_t(p_2) \\ \vdots \\ \mathbf{y}_t(p_M) \end{bmatrix} = A \cdot \vec{\mathbf{x}}_t(\mathcal{P}) \quad (9)$$

where  $\vec{\mathbf{x}}_t(\mathcal{P})$  represents the collection of virtual signal sources contributing to the  $M$  observations for time slot  $t$ :

$$\vec{\mathbf{x}}_t(\mathcal{P}) \triangleq \begin{bmatrix} \vec{\mathbf{x}}_t(p_1) \\ \vec{\mathbf{x}}_t(p_2) \\ \vdots \\ \vec{\mathbf{x}}_t(p_M) \end{bmatrix} \quad (10)$$

The matrix  $A$  is the linear combination matrix, given by:

$$A \triangleq I_M \otimes \vec{\alpha}^T \quad (11)$$

where  $I_M$  is the  $M \times M$  identity matrix and  $\otimes$  represents the Kronecker product.

Next, by stacking the received signal power  $\mathbf{y}_t(p_{Target})$  for the target position  $p_{Target}$  into (9), we obtain a new Gaussian random vector:

$$\begin{bmatrix} \vec{\mathbf{y}}_t(\mathcal{P}) \\ \mathbf{y}_t(p_{Target}) \end{bmatrix} = \begin{bmatrix} A & \vec{0} \\ \vec{0} & \vec{\alpha}^T \end{bmatrix} \cdot \begin{bmatrix} \vec{\mathbf{x}}_t(\mathcal{P}) \\ \vec{\mathbf{x}}_t(p_{Target}) \end{bmatrix} \quad (12)$$

The mean and covariance matrix can be computed using (8). We denote this distribution as:

$$\begin{bmatrix} \vec{\mathbf{y}}_t(\mathcal{P}) \\ \mathbf{y}_t(p_{Target}) \end{bmatrix} \sim N \left( \begin{bmatrix} \vec{\mu}_{t,Past} \\ \mu_{t,Target} \end{bmatrix}, \begin{bmatrix} \Sigma_{t,Past} & \Sigma_{t,Past,Target} \\ \Sigma_{t,Target,Past} & \Sigma_{t,Target} \end{bmatrix} \right) \quad (13)$$

Note that from (8)) and (9),  $\vec{\mu}_{t,Past}$  is a zero vector and  $\mu_{t,Target}$  is a zero scalar, respectively.

With (13), the predicted value  $\mathbf{y}_t(p_{Target})$  for the target position  $p_{Target}$  can be calculated using the conditional probability formula:

$$\hat{\mathbf{y}}_t(p_{Target}) \triangleq \mathbf{y}_t(p_{Target}) | \{ \vec{\mathbf{y}}_t(\mathcal{P}) = \vec{\mathbf{y}}_t(\mathcal{P}) \} \quad (14)$$

- **Prediction Mean:**

$$\begin{aligned} & \text{mean}(\hat{\mathbf{y}}_t(p_{Target})) \\ & = \Sigma_{t,Target,Past} \Sigma_{t,Past}^{-1} (\vec{\mathbf{y}}_t(\mathcal{P}) - \vec{\mu}_{t,Past}) \end{aligned} \quad (15)$$

- **Prediction Variance:**

$$\begin{aligned} & \text{var}(\hat{\mathbf{y}}_t(p_{Target})) \\ & = \Sigma_{t,Target} - \Sigma_{t,Target,Past} \Sigma_{t,Past}^{-1} \Sigma_{t,Past,Target} \end{aligned} \quad (16)$$

Notice that in (16), the prediction variance at the selected target position depends only on the target position and the locations of all previously sampled observations.

#### IV. PROPOSED METHODS

Our approach consists of three main parts: Local Kernel Estimation, Active Sampling, and Quasi-Dynamic Reconstruction. First, we use the given observations to reconstruct the RF Radiance Field with a more precise uncertainty model compared to equation (16). Next, leveraging this uncertainty model, we show how observations can be collected more efficiently. Finally, we demonstrate that when the RF Radiance Field changes, only a few new samples are needed to reconstruct the updated field.

##### A. Local Kernel Estimation

The computational bottleneck for conventional Gaussian prediction arises from the matrix inversion of  $\Sigma_{t,Past}$  in (15) and (16). The size of  $\Sigma_{t,Past}$  for  $M$  observations is  $M \times M$ . As a result, the computational complexity is  $\mathcal{O}(M^3)$ . This complexity is particularly high in large-scale Radiance Fields which entail a large number of observations.

Moreover, in a given scene, certain areas may consist of empty space, where virtual signal sources are expected to be highly correlated. In contrast, other areas may be densely populated with objects such as tables, chairs, and other furniture. In these crowded regions, the correlation between virtual signal sources will differ from that in the empty spaces. As a result, the uncertainty model should be adapted based on the characteristics of the observations. Specifically, we assign low uncertainty to smoother areas and higher uncertainty to regions with sharper variations.

To address the two issues mentioned above, we propose a *local kernel estimation* strategy that tackles both the computational complexity and the correlation locality problems. For a selected target position  $p_{Target}$ , we use only the past observations sampled close to  $p_{Target}$  to perform the local kernel estimation and predict  $\mathbf{y}_t(p_{Target})$ . Specifically, observation  $\mathbf{y}_t(p)$  will be used for the local kernel estimation for position  $p_{Target}$ , if the position of the observation,  $p$ , is sufficiently close  $p_{Target}$ , that is:

$$\{p \mid \|p - p_{Target}\| \leq L\} \quad (17)$$

Here,  $L$  denotes the locality parameter, which can be adjusted according to the structural complexity of the scene. In practice,  $L$  may be selected based on available side information. For example:

- **Signal frequency:** Different operating frequencies lead to different path losses, which affect the spatial correlation between positions.
- **Room layout:** Knowledge of the environment can guide the choice of  $L$ ; for instance, in rooms densely filled with furniture, spatial correlation tends to be lower, suggesting a smaller  $L$  is appropriate.

By incorporating such side information, users can select  $L$  in a principled way to balance reconstruction accuracy and computational complexity. In this work, we introduce the concept of local kernel estimation, while leaving a systematic quantitative method for determining the optimal  $L$  to future research.

Once the observations for the selected target are obtained, we estimate the kernel's length-scale parameter  $l$ , as defined in (6). However, instead of using a global length-scale, we adopt a local kernel estimation approach, where the parameter  $l$  varies based on the position of the target. Therefore, we modify (6) as follows:

$$\text{Cov}(\mathbf{x}_t(p_i), \mathbf{x}_t(p_j)) = \sigma^2 \exp \left( -\frac{1}{2l_t(p_{Target})^2} \|p_i - p_j\|^2 \right) \quad (18)$$

Next, we perform maximum likelihood estimation to estimate the local length-scale parameter  $l_t(p_{Target})$  in (18):

$$\hat{l}_t(p_{Target}) = \arg \max_l \frac{1}{(2\pi)^{\frac{M_{local}}{2}} |\Sigma_{t,l}|^{\frac{1}{2}}} e^{-\frac{1}{2} \bar{\mathbf{y}}_t(\mathcal{P}_{local})^T \Sigma_{t,l}^{-1} \bar{\mathbf{y}}_t(\mathcal{P}_{local})} \quad (19)$$

where  $\bar{\mathbf{y}}_t(\mathcal{P}_{local})$  is a subset of  $\bar{\mathbf{y}}_t(\mathcal{P})$  that satisfies the condition in (17) and  $M_{local}$  denotes the number of elements of  $\bar{\mathbf{y}}_{t,local}$ . Note that the local length-scale parameter  $l_t(p_{Target})$  indicates which areas are smooth and which exhibit sharper spatial variations. After estimating the local kernel, we can use equations (15) and (16) to calculate the prediction mean and variance of  $\mathbf{y}_t(p_{Target})$ .

Since only  $M_{local}$  observations are used for the prediction, the computation complexity is reduced from  $\mathcal{O}(M^3)$  to  $\mathcal{O}(M_{local}^3)$ . If the observations are uniformly distributed within the room, then the ratio  $\frac{M_{local}}{M}$  is approximately  $\frac{\pi L^2}{\text{size of the scene}}$ .

To illustrate the concept of local kernel estimation, we collect 200 observations from synthetic data generated by the NeRF2 pre-trained model, which simulates a room measuring 10 m  $\times$  6 m, as shown in Fig. 3. The locality parameter is set to  $L = 1$  m. As shown in the red circle of Fig. 2,

only 7 observations are needed to predict the received signal at the selected target position.

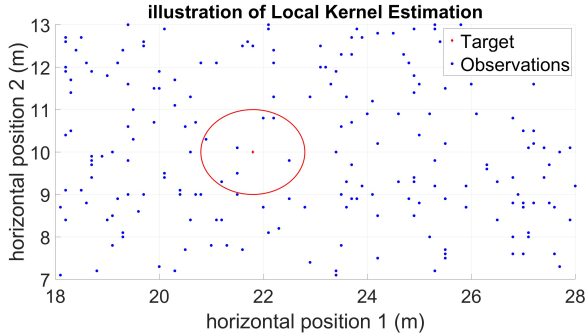


Fig. 2. Illustration of local kernel estimation for a selected target position.

Fig. 5 and Fig. 6 show the reconstruction results with and without our proposed local kernel estimation. Compared to the ground truth in Fig. 3, it is clear that the local kernel estimation helps capture the shape details of the RF Radiance Field, leading to a more accurate reconstruction.

### B. Active Sampling

As mentioned earlier, collecting new observations can be costly, and some may offer limited or negligible information. In such cases, adding these samples to the reconstruction may not significantly improve the model's accuracy. To address this, we propose an active sampling strategy to focus on collecting more informative observations.

Our approach starts with an initial prediction of the RF Radiance Field using a small number of random observations. We then choose the next observation position based on the highest predicted variance, as given by (16). Fig. 4 shows an example of the prediction variance of the received signal power




Fig. 3. Simulated RF Radiance Field.




Fig. 4. Prediction variance and selected next samples.




Fig. 5. Prediction mean with local kernel estimation.




Fig. 6. Prediction mean without local kernel estimation.

across different positions. After a few initial observations, only certain areas of the scene exhibit high uncertainty (i.e., large variance). Thus, further observations are focused on these high-variance regions to improve the model's accuracy most efficiently. Note that the high-variance regions in Fig. 4 match the areas of sharp change in Fig. 3, which are exactly the regions that require more samples to learn. The proposed method for static RF Radiance Field is summarized in Algorithm 1.

In this work, instead of only selecting the largest variance position each time, we divide the scene into smaller sections and select the position with the highest variance within each section as the next observation point. This approach helps avoid repeatedly computing the matrix inverse in (15) and (16) by adding only one new observation at a time. Note that some sections may be skipped if the predicted variances of all positions within a section are below a certain threshold.

In summary, Local Kernel Estimation helps us learn the shape and smoothness of different regions in the RF Radiance Field. Prediction variance shows which areas need more samples and which have little uncertainty. Using this variance, the active sampler efficiently avoids oversampling and undersampling.

### C. Quasi-Dynamic Reconstruction

Now, we explore the extension of Algorithm 1 in a quasi-dynamic environment. We focus on the scenario where the RF Radiance Field of the scene has already been fully measured at time slot  $t$ , and at time slot  $t + 1$ , certain regions of the RF Radiance Field change due to environmental dynamics, such as object movement, partial blockage, or material variations. Instead of reconstructing the entire RF field from scratch, we model the temporal evolution as an incremental update. Specifically, the received signal power at a target position  $P_{Target}$  can be written as:

$$\mathbf{y}_{t+1}(p_{Target}) = \mathbf{y}_t(p_{Target}) + \mathbf{e}_t(p_{Target}) \quad (20)$$

**Algorithm 1** Proposed method for static RF Radiance Field**Input:**

- $M$ : Initial number of observations
- targets: Set of target positions

**Output:** Predicted mean and variance for each position**Initialize:** Collect  $M$  initial observations**For each target position** in targets:Step 1: Find local samples close to  $P_{Target}$  from the initial observations

Step 2: Estimate the local kernel based on the selected local samples

Step 3: Calculate the prediction mean and variance using the estimated local kernel and the local samples

**if** Active Sampling **then**

Collect new samples at the position with the largest variance (within each section or for the entire scene), then return to Step 1.

**end if**

where  $e_t(p_{Target})$  denotes the difference of the value between time slot  $t$  and  $t + 1$  for position  $p_{Target}$ . This formulation allows us to treat the previously reconstructed field at time  $t$  as a prior reference. By conditioning on this prior, we reconstruct the RF Radiance Field at time  $t + 1$  by focusing on learning the difference field induced by environmental changes.

Operationally, at time slot  $t + 1$ , sparse new measurements, are collected. We compute the difference between the newly observed measurements and the corresponding field values measured at time  $t$ . Next, the same Local Kernel Estimation framework used in the static case is applied to estimate the difference field. The updated RF Radiance Field for time slot  $t + 1$  is then obtained by adding the estimated difference field to the RF Radiance Field for time slot  $t$ . In this way, we avoid making large number of observation at those regions that remain unchanged, instead allocate sensing resources to regions that exhibit significant temporal variations.

A key observation supporting this strategy is that environmental changes in RF propagation are spatially correlated and sparse. For example, the received signal power typically varies over a localized region caused by a small number of reflectors/objects' movements. Based on this property, our Active Sampling strategy prioritizes measurements in the spatial neighborhood of significant variations, efficiently targeting dynamically evolving regions of the RF Radiance Field, which is exactly what is needed to adapt to the evolving field. Examples will be provided in the next section.

## V. EXPERIMENTS

In this section, we present the experimental validation of our proposed method and compare it with representative approaches from two distinct methodological categories: NeRF2, a neural network-based reconstruction method, and GSpaRC, a Gaussian Splatting-based reconstruction method.

**Experimental Setup and Data Collection:** We use a set of scenarios consistent with the Bluetooth Low Energy (BLE) localization experiments reported in NeRF2: 1) omnidirectional

antenna measurements in a single-carrier system, with only the received signal power being recorded; and 2) high operating SNR (negligible noise).

**Evaluation Metrics:** We evaluate the performance of our algorithm, NeRF2, and GSpaRC in terms of the number of measurements, using the mean and median absolute error between each method's predicted signal strength and the ground truth.

### A. Static RF Radiance Field

1) *Static RF Radiance Field Scenarios:* In our static RF Radiance Field experiments, we consider four datasets. Scenarios 1–3 are based on the datasets provided in the original NeRF2 paper [5, Fig. 12], corresponding to indoor environments of sizes 10, m  $\times$  6, m, 8, m  $\times$  4, m, and 8, m  $\times$  8, m, respectively. These environments are referred to as Scenario 1, Scenario 2, and Scenario 3 throughout this section. Scenario 4 corresponds to a real-world measurement dataset collected in a 10, m  $\times$  6, m meeting room at a local office, following a similar data collection procedure to provide independent validation. In this experiment, we use a TurtleBot4 robot as the mobile measurement platform, which carries two wireless systems: a 60,GHz mmWave system using MikroTik wAP 60G $\times$ 3 routers, and a 5,GHz WiFi system consisting of an IEEE 802.11ac access point and an LG Nexus 5 smartphone running CSIKit to extract received signal strength information (RSSI). Each scenario is illustrated in Fig. 7((a))–((d)), where the red dots indicate the real-world RF measurements collected along a trajectory within indoor environments containing blockers and scatterers.

In addition, for Scenarios 1–3, we use a pre-trained NeRF2 model to generate synthetic received power maps across the entire environment at a spatial resolution of 0.1, m, illustrated as mesh surfaces in Fig. 7((a))–((c)). These synthetic maps provide dense reference fields for evaluating reconstruction accuracy. For Scenario 4, however, the number of available measurements is insufficient to train a reliable NeRF2 model; therefore, no synthetic reference field is generated.

2) *Static RF Radiance Field Results:* Here, we present the results of the experiment for static and RF Radiance Fields. Fig. 8 and Fig. 9 compare the performance of our proposed method, NeRF2, and GSpaRC, using synthetic data and real measurement data, respectively. The results on synthetic data show that our method outperforms NeRF2, reducing the number of observations by 30%–60%, and also achieves better performance than GSpaRC. For real-world measurement data, our method demonstrates significant improvements in both mean absolute error (MAE) and median absolute error (Median AE) compared to NeRF2 and GSpaRC. In both synthetic and real scenarios, our method exhibits much more stable performance, whereas the performance of NeRF2 and GSpaRC can vary considerably with small changes in the number of samples. This variability arises from the randomness in the initial weights of the neural network and the initialization of Gaussian primitives in the Gaussian Splatting framework. With a limited number of samples, both NeRF2 and GSpaRC struggle to converge to accurate solutions. Furthermore, the results

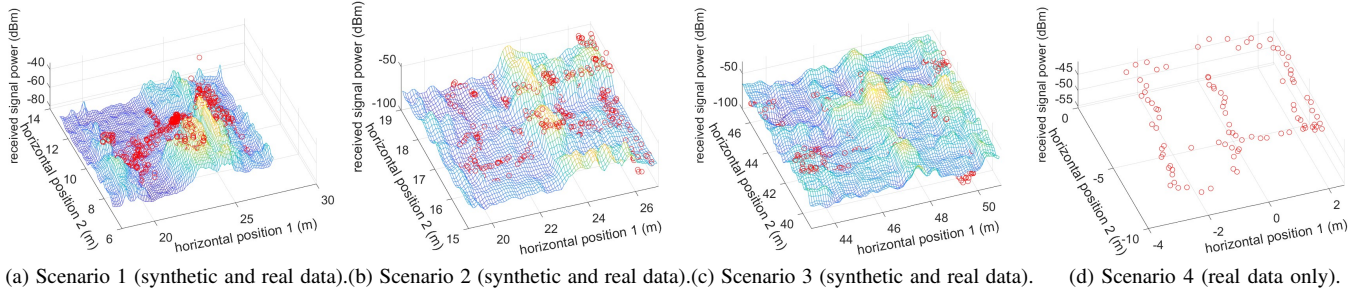


Fig. 7. Synthetic and real data for various scenarios. The synthetic field is shown as a mesh surface; measured samples are indicated by red dots.

show that active sampling outperforms non-active sampling, as it prioritizes collecting measurements from positions with high uncertainty, improving overall reconstruction accuracy and efficiency.

Computation time comparison: NeRF2 (running on a A6000 server) – 6 hours; GSPaRC (running on a A6000 server) – 1 minute; our method (running on a 4 core CPU laptop) – 1 minute total (non-adaptive); 10 seconds per iteration (adaptive)

## B. Quasi-Dynamic RF Radiance Field

1) *Static RF Radiance Field Scenarios*: In the quasi-dynamic experiments, we consider two representative cases to evaluate the effectiveness of the proposed conditional update strategy.

**Case 1 (Synthetic Perturbation)**: We apply a controlled perturbation to the pre-trained NeRF2 model to generate a modified RF Radiance Field at time slot  $t + 1$ . This setup allows us to isolate and analyze the behavior of the incremental update mechanism under well-defined changes.


**Case 2 (Realistic Indoor Environment)**: We import the actual layout of our office building into the simulator to construct a more complex RF Radiance Field within a  $8\text{ m} \times 14\text{ m}$  area. The environment includes typical indoor furniture such as bookshelves, tables, and desks. Between two consecutive time slots, selected furniture items are repositioned to emulate realistic environmental dynamics, producing spatially correlated changes in the RF propagation field.

It is important to emphasize that the proposed method does not require prior knowledge of the object locations or the transmitter position. The description of the scenarios studied is only used as the ground truth for evaluation purposes. The Algorithm, however, operates solely based on observed RF measurements and their temporal variations.

2) *Quasi-Dynamic RF Radiance Field Results*: We first use the pre-trained NeRF2 model to generate the synthetic RF Radiance Field at time slot  $t + 1$ , shown as a mesh surface in Fig. 10(a). The difference between this field and the previous field at time  $t$  is also visualized as a mesh surface in Fig. 10(b).

Directly learning the RF Radiance Field at  $t + 1$  without conditioning on the field at  $t$  can lead to excessive sampling in regions that remain unchanged or structurally sharp. In contrast, by conditioning on the previous field, our method focuses on learning only the differences between consecutive

time steps, thereby prioritizing regions where meaningful changes occur. The locations of the observations used to learn this difference are indicated by the green dots in Fig. 10(b). The high-density region of green dots (highlighted by the red square) aligns with the area experiencing significant changes in the RF Radiance Field, demonstrating that the proposed sampling strategy successfully concentrates measurements in informative regions.



(a) RF Radiance Field at time slot  $t + 1$ . (b) Difference of the RF Radiance Field between time slot  $t$  and  $t + 1$  with the active learning samples.

Fig. 10. Comparison of RF Radiance Fields: the field at time  $t + 1$  and its temporal difference relative to time  $t$ .

Next, we import the actual layout of our office building into the NVIDIA Sionna simulator, creating a more realistic indoor environment that includes furniture such as bookshelves, tables, and desks. The layouts at two consecutive time slots are shown in Fig. 11(a) and Fig. 11(b). A signal source is placed in the middle of the scene (indicated by the red triangle). Using these layouts, we generate the RF Radiance Field at time slot  $t + 1$ , shown as a mesh surface in Fig. 12(a). The difference between the RF Radiance Fields at time slots  $t$  and  $t + 1$  is also visualized as a mesh surface in Fig. 12(b). The green dots indicate the locations of the observations selected by our active sampling strategy to learn this difference. As highlighted by the red square, the high-density sampling region corresponds to the area where the RF Radiance Field changes significantly, demonstrating that the proposed uncertainty-driven sampling strategy effectively focuses measurements on informative regions.

We show the MAE of the two quasi-dynamic scene change examples in Fig. 13(a) and Fig. 13(b), respectively. The results show that, by conditioning on the RF Radiance Field at time slot  $t$ , our approach updates the field at time slot  $t + 1$  by learning the differences between the two time slots, achieving better performance than re-estimating the entire field

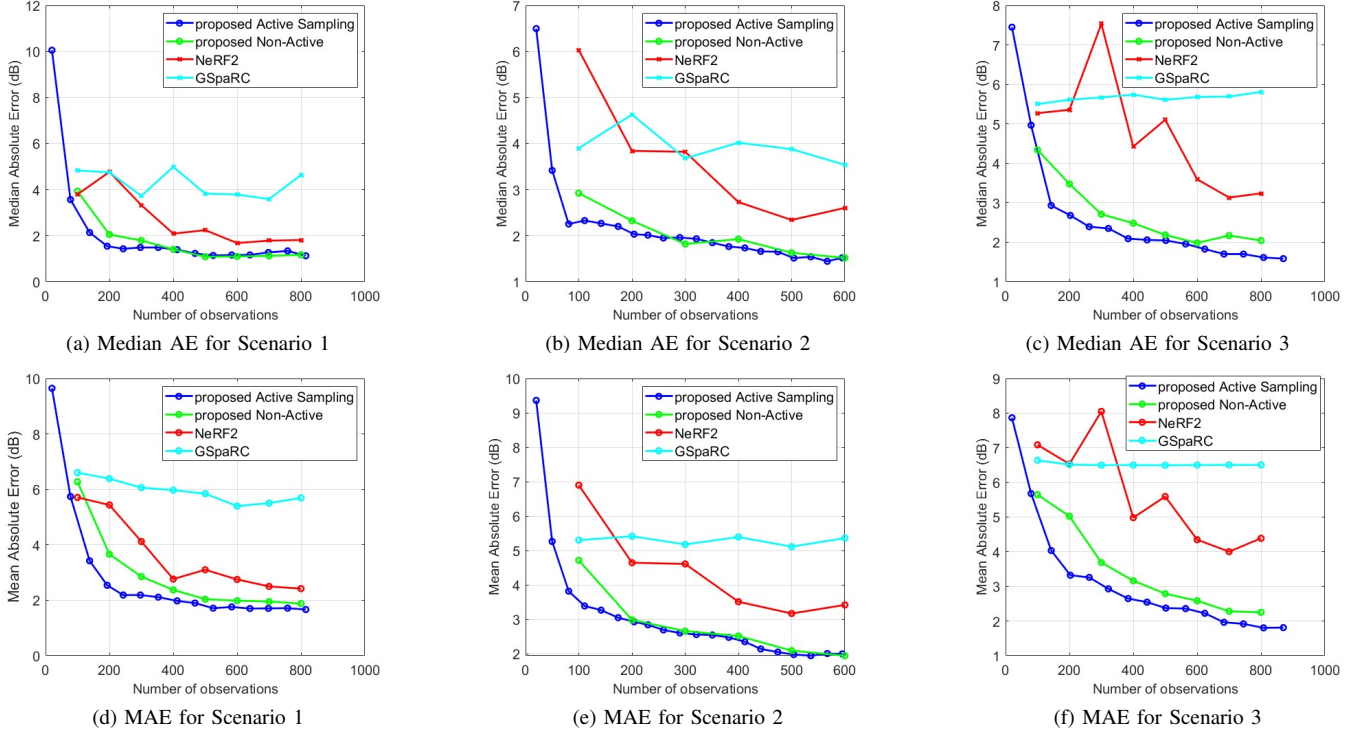


Fig. 8. Performance comparison for various scenarios using synthetic data.

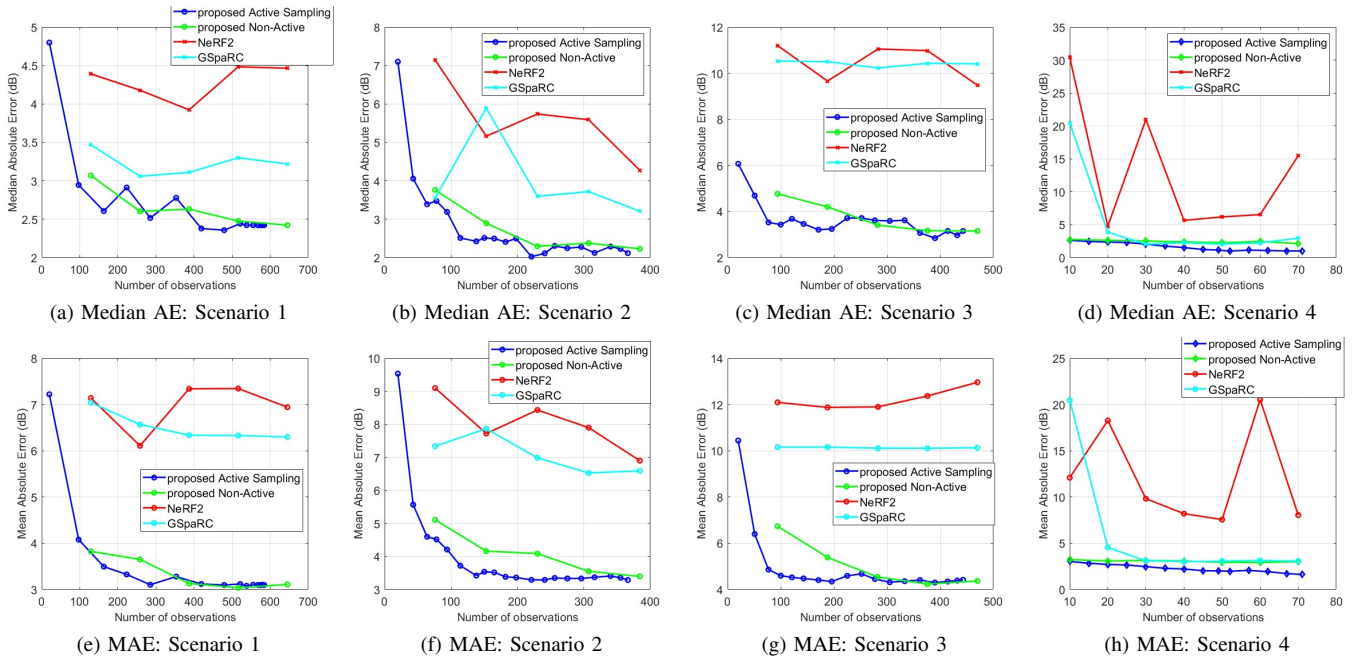


Fig. 9. Performance comparison for various scenarios using real measurement data.

from scratch. This strategy significantly reduces the number of required samples, making the method more efficient and practical for real-world wireless communication tasks such as channel estimation, signal tracking, and radio environment mapping.

In our experiments, instead of using the perfect RF Radiance Field at time  $t$ , we also use an estimated field with  $\sim 2$  dB

MAE. Even with this imperfect prior, learning the differences still outperforms constructing the  $t+1$  RF Radiance Field from scratch, demonstrating the robustness and practical advantage of the conditional update strategy.






Fig. 11. Comparison of room layouts at consecutive time slots  $t$  and  $t+1$ .



(a) RF Radiance Field at time slot  $t+1$ . (b) Difference of the RF Radiance Field between time slot  $t$  and  $t+1$  with the active learning samples.

Fig. 12. Comparison of RF Radiance Fields at time  $t+1$  and their changes relative to time  $t$ .



(a) MAE comparison from NeRF2 synthetic data. (b) MAE comparison from Nvidia Sionna synthetic data with real-world layout.

Fig. 13. Comparison of MAE results for different simulation setups.


## VI. CONCLUSIONS

In this work, we propose a training-free method for predicting received signal power in RF Radiance Fields. Our approach generates predictions quickly at any location within the scene and includes an uncertainty model. By utilizing local kernel estimation, it offers improved computational efficiency over traditional Gaussian prediction methods. Simulation results show that fewer samples are needed to achieve performance comparable to NeRF2 and GSpaRC. Additionally, we further reduce the number of required observations by incorporating active sampling, which selects the most informative data points. Importantly, our method is well-suited for dynamic wireless environments, enabling near-instantaneous reconstruction as soon as new observations are available. This real-time adaptability is essential for wireless communication systems operating in rapidly changing indoor environments, such as smart homes, offices, or industrial IoT settings.

## APPENDIX A QUASI-DYNAMIC RF RADIANCE FIELD (SIMPLE GEOMETRIC DISPLACEMENT):


Using the NVIDIA Sionna simulator [20], we construct a  $4\text{ m} \times 4\text{ m}$  indoor environment containing a  $2\text{ m} \times 2\text{ m}$  object. The object is then displaced by  $0.2\text{ m}$  to generate the RF Radiance Field at the subsequent time slot, enabling evaluation under localized structural variations. The scene layouts at the two time slots are shown in Fig. 14((a)) and Fig. 14((b)), where a signal source is placed at the right corner of the environment. Using the NVIDIA Sionna model, we generate the corresponding RF Radiance Fields for these layouts. The RF Radiance Field at time slot  $t+1$  is illustrated as a mesh surface in Fig. 15((a)). The mesh surface of the difference between the fields at time slots  $t$  and  $t+1$  is shown in Fig. 15((b)).

As the object is shifted by  $0.2\text{ m}$ , part of the region experiences a signal power drop, while areas that were previously blocked receive stronger signals. In Fig. 15((b)), the green dots indicate the observation locations selected by our active sampling method. The high-density regions of the collected samples (highlighted by red squares) align with the areas where the RF Radiance Field changes significantly.



(a) Layout at time slot  $t$ . (b) Comparison of room layouts at consecutive time slots  $t$  and  $t+1$ .

Fig. 14. Comparison of room layouts at consecutive time slots  $t$  and  $t+1$ .



(a) RF Radiance Field at time slot  $t+1$ . (b) Difference of the RF Radiance Field between time slot  $t$  and  $t+1$  with the active learning samples.

Fig. 15. Comparison of RF Radiance Fields at time  $t+1$  and their changes relative to time  $t$ .

The MAE comparison is shown in Fig. 16. Consistent with the results in the previous section, conditioning on the RF Radiance Field at time slot  $t$  allows our method to update the field at time slot  $t+1$  by learning the difference between the

two time slots, which outperforms reconstructing the entire field from scratch. Furthermore, even when the prior field at time  $t$  is imperfect ( $\sim 2$  dB MAE), learning the difference still achieves better performance than reconstructing the RF Radiance Field at time  $t + 1$  from scratch.

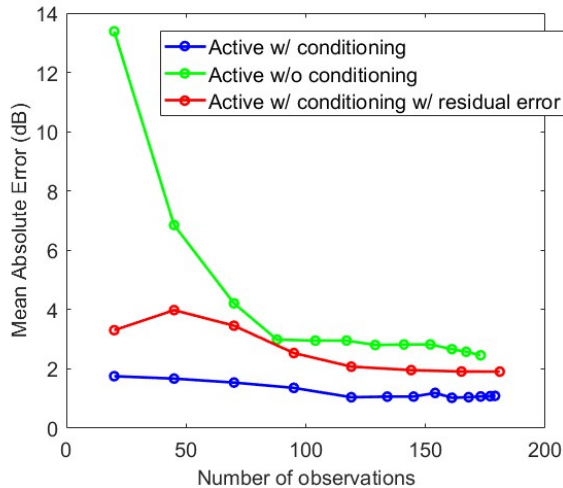


Fig. 16. MAE Comparison (from Nvidia Sionna).

#### REFERENCES

- [1] J. Hu, T. Zheng, Z. Chen, H. Wang, and J. Luo, "Muse-fi: Contactless multi-person sensing exploiting near-field wi-fi channel variation," in *Proceedings of the ACM MobiCom*, 2023.
- [2] H. Zhao, M. Huang, and Y. Shen, "High-accuracy localization in multipath environments via spatio-temporal feature tensorization," *IEEE Transactions on Wireless Communications*, vol. 21, no. 12, pp. 10 576–10 591, 2022.
- [3] S. Vakalis, L. Gong, and J. A. Nanzer, "Imaging with wifi," *IEEE Access*, vol. 7, pp. 28 616–28 624, 2019.
- [4] T. Woodford, X. Zhang, E. Chai, and K. Sundaresan, "Mosaic: Leveraging diverse reflector geometries for omnidirectional around-corner automotive radar," in *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services (MobiSys)*, 2022.
- [5] X. Zhao, Z. An, Q. Pan, and L. Yang, "NeRF2: Neural radio-frequency radiance fields," in *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, 2023, pp. 1–15.
- [6] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [7] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3D Gaussian splatting for real-time radiance field rendering," *ACM Trans. Graph.*, vol. 42, no. 4, pp. 139–1, 2023.
- [8] L. Zhang, H. Sun, S. Berweger, C. Gentile, and R. Q. Hu, "RF-3DGS: Wireless channel modeling with radio radiance field and 3D gaussian splatting," *IEEE Transactions on Wireless Communications*, vol. 25, pp. 10 419–10 433, 2026.
- [9] B. S. Nukapotula, R. Tripathi, S. Pregler, D. Kalathil, S. Shakkottai, and T. S. Rappaport, "GSpaRC: Gaussian splatting for real-time reconstruction of RF channels," *arXiv preprint arXiv:2511.22793*, 2025.
- [10] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," in *European Conference on Computer Vision*, 2020.
- [11] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "DeepSDF: Learning continuous signed distance functions for shape representation," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [12] V. Sitzmann, J. N. Martel, A. W. Bergman, D. B. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," in *Advances in Neural Information Processing Systems*, 2020.
- [13] A. Tewari, O. Fried, J. Thies, V. Sitzmann, S. Lombardi, K. Sunkavalli, R. Martin-Brualla, T. Simon, J. Saragih, M. Nießner *et al.*, "State of the art on neural rendering," *Computer Graphics Forum*, vol. 39, no. 2, 2020.
- [14] L. Yariv, Y. Kasten, D. Moran, M. Galun, M. Atzmon, B. Ronen, and Y. Lipman, "Multiview neural surface reconstruction by disentangling geometry and appearance," in *Advances in Neural Information Processing Systems*, 2020.
- [15] S. Chan, S. Saito, L. Li, H. Zhao, S. Xiao, K. Sugimoto, S. Sakashita, and T. Komura, "PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization," *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [16] Y. Xie, T. Takikawa, S. Saito, O. Litany, S. Yan, N. Khan, F. Tombari, J. Tompkin, V. Sitzmann, and S. Sridhar, "Neural fields in visual computing and beyond," in *Computer Graphics Forum*, 2021.
- [17] T. Orekondy, P. Kumar, S. Kadambi, H. Ye, J. Soriaga, and A. Behboodi, "WiNeRT: Towards neural ray tracing for wireless channel modelling and differentiable simulations," in *The Eleventh International Conference on Learning Representations*, 2023.
- [18] X. Chen, Z. Feng, K. Sun, K. Qian, and X. Zhang, "RFCanvas: Modeling RF channel by fusing visual priors and few-shot RF measurements," in *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*, 2024, pp. 464–477.
- [19] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.
- [20] J. Hoydis, S. Cammerer, F. A. Aoudia, A. Vem, N. Binder, G. Marcus, and A. Keller, "Sionna: An open-source library for next-generation physical layer research," *arXiv preprint arXiv:2203.11854*, 2022.