SPHGAS: SPHERICAL GAUSSIAN SPLATTING FOR RF SPECTRUM SYNTHESIS UNDER VARYING TRANSMITTER–RECEIVER CONFIGURATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

This paper proposes SphGaS, a novel Gaussian splatting (GS)-based framework for rendering transmitter–receiver-conditioned (TX-RX) RF spectra directly on the spherical manifold. Existing GS-based methods typically synthesize planar spectra. Geometrically, this is not consistent with the practical spherical structure of the RF spectra, consequently leading to distorted Gaussian responses. Meanwhile, some methods model spectral variations at one transceiver endpoint only, neglecting the coupled variations induced by the TX and/or RX movement. To the best of our knowledge, SphGaS is the first framework that jointly addresses geometric misalignment and TX-RX mobility. It introduces spherical Gaussian primitives that deform with the TX-RX configuration and are rasterized directly in the spherical domain. This design preserves the geometry of Gaussian responses while efficiently capturing the coupled spectrum variations induced by TX–RX movement. With a customized CUDA rasterizer, SphGaS keeps fast rendering while substantially improving the modeling fidelity. Experiments show that SphGaS improves PSNR by 1.49 dB and achieves up to $40\times$ faster rendering than state-of-the-art GS-based methods on real-world RFID datasets.

1 INTRODUCTION

Radio frequency (RF) spectrum describes how RF signals propagate from a transmitter (TX) to a receiver (RX) in a given scenario. It captures directional propagation patterns shaped by reflection, blockage, and multipath, providing useful cues for tasks such as localization and floor plan reconstruction Levie et al. (2021); Amballa et al. (2025). However, densely collecting RF spectra across TX–RX configurations is costly, as it requires specialized hardware and labor-intensive measurements. This motivates synthesizing RF spectra for unseen configurations from limited observations.

Existing methods addressing this challenge learn the underlying RF propagation field from observations and query the learned model at novel configurations to synthesize unmeasured spectra. Early efforts employed neural radiance field (NeRF) formulations Mildenhall et al. (2020) to represent the field implicitly with neural networks Zhao et al. (2023); Lu et al. (2024). Despite their compact representations, these approaches incur substantial rendering costs due to the use of volumetric ray marching. More recently, 3D Gaussian Splatting (3DGS) Kerbl et al. (2023) emerged as an efficient alternative, explicitly characterizing the RF field with a collection of 3D Gaussian primitives Yang et al. (2025). Given a TX-RX configuration, these primitives are splatted onto a rendering plane and rasterized to synthesize the desired spectrum. Through replacing volumetric ray marching with CUDA-accelerated rasterization, 3DGS substantially reduces the rendering overhead, making it an attractive technique for RF spectrum synthesis.

However, current 3DGS-based RF spectrum synthesis still faces two key issues. The primary challenge lies in the geometric mismatch between the planar rendering space and the native spherical structure of RF spectra. In some work, the spectrum was projected from the spherical manifold onto a 2D plane, through e.g., equirectangular projection (ERP) Wen et al. (2026); Shen et al. (2026); Wen et al. (2025), as illustrated in Fig. 1**A**. This treatment may introduce two types of geometric artifacts in the splatted Gaussian responses: 1) *Gaussian response distortion*. The planar parameterization changes the local spherical geometry, causing the responses to be warped after being mapped

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

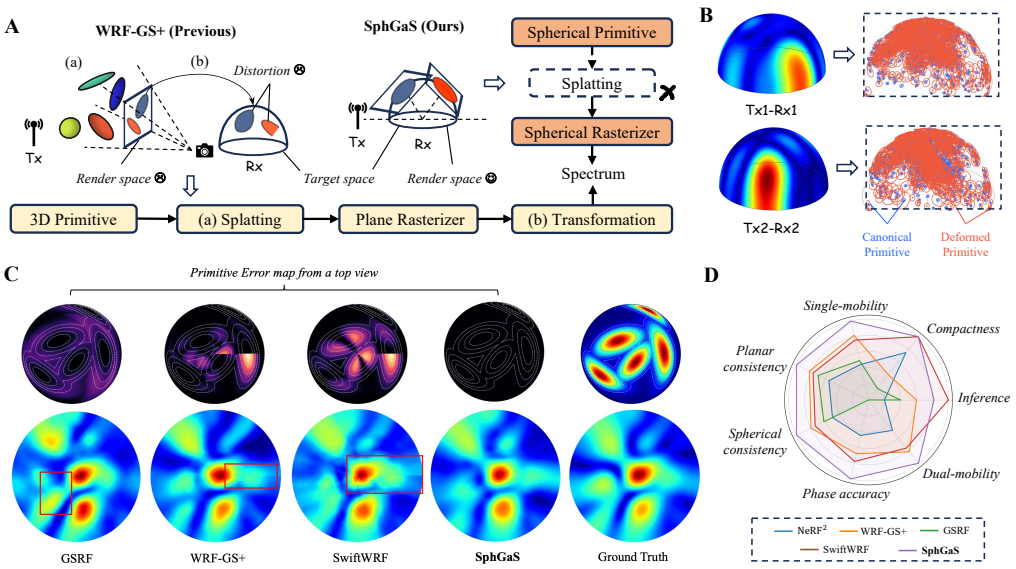

Figure 1: Overview and motivation of SphGaS. **A**: Comparison between conventional RF 3DGS rendering and our spherical formulation. **B**: TX-RX-conditioned deformation of spherical primitives under different TX–RX configurations. **C**: Qualitative comparison of primitive responses and reconstructed spectra, highlighting distortion and discontinuity in existing methods. **D**: Overall comparison across reconstruction quality, efficiency, compactness, and generalization.

back onto the sphere; 2) *artificial boundaries*. Directions that are adjacent on the sphere may be separated by the $0^\circ/360^\circ$ azimuth seam, causing the responses to be split or truncated. Together, these effects can disturb the Gaussian mixture structure assumed by the rasterizer, ultimately degrading the reconstruction fidelity (see Fig. 1**C** for visualization and Sec. C for a detailed error analysis). A second challenge comes from the TX-RX mobility. Most 3DGS-based approaches consider configuration changes at only one transceiver endpoint Liu et al. (2026); Yang et al. (2025), making it difficult to capture the coupled variations induced by joint TX–RX mobility. To our knowledge, simultaneously addressing the geometric mismatch and TX-RX mobility remains unexplored.

In this paper, we introduce **SphGaS**, a **Sph**erical **Ga**ussian **S**platting framework for RF spectrum synthesis under varying TX-RX configurations. SphGaS is based on the following insights: the localized virtual arrival response centered at a point on the sphere can be well modeled using a spherical Gaussian primitive (see Sec. B for a theoretical analysis). They are then combined using rasterization to reflect the combined contribution from multipath propagation. As the TX and RX move, these localized arrival responses vary accordingly, which are captured by the TX-RX-conditioned deformation of the primitives.

To build SphGaS, first, we propose spherical Gaussian primitives to eliminate the geometric misalignment caused by planar rendering. Instead of splatting primitives onto a flat plane, we define them directly on the sphere so that they can be rasterized in their native domain. To accommodate spherical Gaussian primitives, we develop a customized CUDA rasterizer that evaluates Gaussian responses under the Riemannian geometry and uses tile-based parallel processing for efficient rendering. Second, we design spherical primitive deformation to model spectrum variations under joint TX–RX mobility. Given a TX-RX configuration, a shared set of canonical primitives is deformed on the spherical manifold, which allows spectrum changes to be captured as principled variations of a common representation. Accordingly, the deformation is formulated under Riemannian geometry, preserving the spherical structure of the primitives throughout their evolution (see toy examples in Fig. 1**B** for an illustration). Finally, motivated by channel reciprocity, we construct bidirectional spectrum pairs for each link to provide reciprocal supervision for deformation learning. *In summary, SphGaS unifies RF spectrum representation, TX-RX-dependent variation, and rendering within a consistent spherical geometry, avoiding planar distortion across configurations.* It represents complex, high-dimensional RF field variations as structured changes on a low-dimensional spherical manifold, providing a compact and tractable representation of configuration-dependent propagation.

Compared with existing 3DGS-based methods, SphGaS achieves up to $40\times$ faster rendering and up to $15\times$ lower storage consumption.

Our contributions are summarized as follows: (1) We introduce **SphGaS**, which represents RF spectra directly with Gaussian primitives defined on the sphere; (2) We develop spherical CUDA rasterization and TX-RX-conditioned deformation under Riemannian geometry, keeping representation, deformation, and rendering consistent with the spherical structure of RF spectra; (3) Extensive experiments on real-world and simulated RF data demonstrate improved reconstruction quality over NeRF and Gaussian-based methods, while maintaining utility for downstream wireless applications such as received signal strength indicator (RSSI) prediction and line-of-sight (LOS) analysis.

2 Related Work

RF Field Reconstruction. RF field has traditionally been characterized using empirical, stochastic, and ray-tracing models, supporting the construction of channel statistics such as radio maps, spectrum cartography, and channel state information (CSI) Hata (1980); Saleh & Valenzuela (1987); Rappaport et al. (2017); Yun & Iskander (2015); 3GPP (2024); Bazerque et al. (2011); Romero & Kim (2022); Zeng & Xu (2021); Zeng et al. (2024); Li et al. (2022). More recently, learning-based approaches have been introduced to replace repeated physical simulation with efficient data-driven prediction, including convolutional neural networks, generative models, and diffusion-based methods Shrestha et al. (2022); Chen et al. (2025b); Levie et al. (2021); Zhang et al. (2023); Chaves-Villota & Viteri-Mera (2023); Bakirtzis et al. (2022); Lee & Molisch (2024); Jaensch et al. (2024). While effective for channel prediction, many of these methods operate on predefined spatial grids or target a limited set of channel quantities. Neural radiance field (NeRF)-based methods provide a continuous alternative by representing RF propagation as coordinate-conditioned neural fields Tancik et al. (2020); Müller et al. (2022); Zhao et al. (2023); Lu et al. (2024); Basu & Chakraborty (2024); Chen et al. (2024); Jiang et al. (2025); Chen et al. (2025c); Gao & Zheng (2025). Such continuous representations enable flexible spatial and directional queries but require repeated network evaluations, leading to higher rendering cost.

Gaussian Splatting for RF Fields. Gaussian splatting has recently been explored for synthesizing different forms of RF data. For radio maps and received-signal-strength fields, existing methods use 2D or 3D Gaussian representations to reconstruct continuous site-specific propagation maps from sparse measurements or simulated observations Zhang et al. (2026b); Wang et al. (2026b); Chen et al. (2025a). Gaussian representations have also been extended to RF domain and CSI, where compact 3D or planar primitives are used to capture spatial channel structure and enable efficient channel reconstruction Zhang et al. (2026b; 2025); Nukapotula et al. (2025). Our focus is RF spatial-spectrum synthesis, where each TX–RX configuration is associated with a directional response over the angular domain. Existing spectrum-oriented methods model such responses using complex-valued, deformable, or geometry-aware Gaussian representations, with subsequent work further considering bidirectional reconstruction, receiver generalization, and cross-frequency synthesis Yang et al. (2025); Wen et al. (2025; 2026); Shen et al. (2026); Zhang et al. (2026a); Yang & Srivastava (2026); Wang et al. (2026a). More recently, SwiftWRF adopts deformable 2D Gaussians to improve spectrum-rendering efficiency Liu et al. (2026). Despite these advances, existing GS-based RF spectrum representations remain largely Euclidean in their geometric formulation and rendering. In contrast, SphGaS places primitives directly on the spherical spectrum domain, aligning the representation and rasterization geometry with the practical structure of RF spectra.

3 Preliminaries

RF spectrum. An RF spectrum characterizes the directional responses for the RF signal from a TX by steering the RX beam over the spherical domain $\mathcal{M}$[1]. Let $\mathbf{s} = (\mathbf{p}_t, \mathbf{p}_r)$ denote a certain TX-RX configuration, where $\mathbf{p}_t, \mathbf{p}_r \in \mathbb{R}^3$ are the TX and RX positions. A beam-steering direction is parameterized by azimuth $\varphi \in [0, 2\pi)$ and elevation $\lambda \in [0, \pi/2]$. They specify a direction vector $\mathbf{q}(\varphi, \lambda) \in \mathcal{M}$ with unit length. For the m-th propagation path, the directional response is $g_m(\varphi, \lambda \mid \mathbf{s}) = \mathbf{w}(\mathbf{q}(\varphi, \lambda))^H \mathbf{h}_m(\mathbf{s})$, where $\mathbf{h}_m(\mathbf{s})$ is its complex response and $\mathbf{w}(\cdot)$ is the beam-

[1]In our setting, $\mathcal{M} = \{\mathbf{q} \in \mathbb{R}^3 : \|\mathbf{q}\|_2 = 1,\ q_z \geq 0\}$ is the receiving hemisphere. We refer to it as the spherical domain, or simply the sphere, throughout the paper.

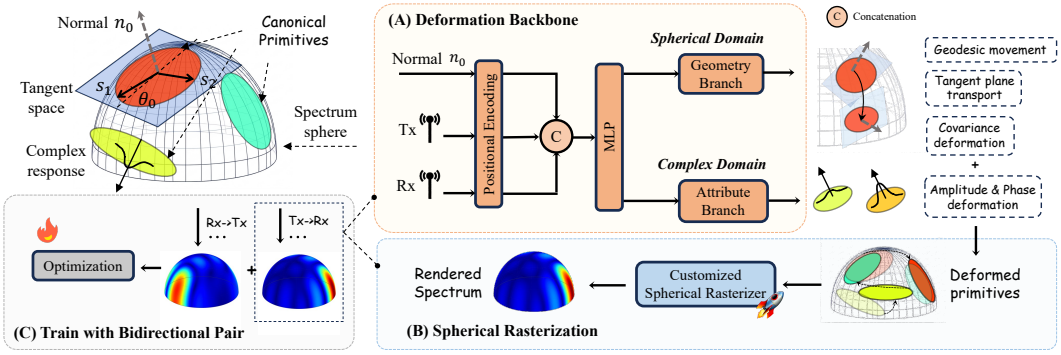

Figure 2: Framework of SphGaS.

steering vector. The RF spectrum is

$$S(\varphi, \lambda \mid \mathbf{s}) = \sum_{m=1}^{N_p} g_m(\varphi, \lambda \mid \mathbf{s}) \in \mathbb{C}, \tag{1}$$

with amplitude $A = |S|$ and phase $\psi = \arg S$. Our task is to synthesize this complex spectrum for varying TX-RX configurations.

Spherical geometry. Since the RF spectrum is defined on $\mathcal{M} \subset \mathbb{S}^2$, we shall apply the standard intrinsic operators on the unit (hemi)sphere. In particular, the tangent space at a given point on the sphere $\mathbf{p} \in \mathbb{S}^2$ is $T_{\mathbf{p}}\mathbb{S}^2 = \{\mathbf{v} \in \mathbb{R}^3 : \mathbf{p}^\top \mathbf{v} = 0\}$. It has an orthonormal basis $\mathbf{B}_{\mathbf{p}} = [\mathbf{b}_1, \mathbf{b}_2]$. With a displacement in the tangent space $\mathbf{v} \in T_{\mathbf{p}}\mathbb{S}^2$, and a target point $\mathbf{q} \in \mathbb{S}^2$ with $\mathbf{q} \neq -\mathbf{p}$, the exponential and logarithmic maps are

$$\mathrm{Exp}_{\mathbf{p}}(\mathbf{v}) = \cos r\, \mathbf{p} + \frac{\sin r}{r}\mathbf{v}, \qquad \mathrm{Log}_{\mathbf{p}}(\mathbf{q}) = \frac{\theta}{\sin\theta}\left(\mathbf{q} - \cos\theta\, \mathbf{p}\right), \tag{2}$$

where $r = \|\mathbf{v}\|_2$ and $\theta = \cos^{-1}(\mathbf{p}^\top \mathbf{q}) \in [0, \pi)$. By continuity, it is easy to verify that $\mathrm{Exp}_{\mathbf{p}}(\mathbf{0}) = \mathbf{p}$ and $\mathrm{Log}_{\mathbf{p}}(\mathbf{p}) = \mathbf{0}$.

The tangent vector at the target point $\mathbf{q}$ can be deduced using the parallel transport along the unique shortest geodesic. The amount of adjustment required by the transfer from $T_{\mathbf{p}}\mathbb{S}^2$ to $T_{\mathbf{q}}\mathbb{S}^2$ is

$$\mathrm{PT}_{\mathbf{p}\to\mathbf{q}}(\mathbf{v}) = \mathbf{v} - \frac{\mathbf{v}^\top \mathbf{q}}{1 + \mathbf{p}^\top \mathbf{q}}(\mathbf{p} + \mathbf{q}). \tag{3}$$

The above transport preserves norms and inner products; for a tangent basis or frame, the operator is applied column-wise to each basis vector.

4 Proposed SphGaS

As illustrated in Fig. 2, SphGaS represents RF spectra with a set of learnable canonical spherical Gaussian primitives shared across TX-RX configurations. Given the TX and RX positions, a TX-RX-conditioned deformation network predicts primitive-wise geometric and (complex) attribute variations. The deformed primitives are then rendered directly on the sphere using a customized differentiable rasterizer. The canonical primitives and deformation network are jointly optimized, with bidirectional spectrum pairs used during training to regularize the learned variations.

4.1 Parameterization of Spherical Primitives

We represent RF spectra using a collection of canonical spherical Gaussian primitives shared across TX–RX configurations. The i-th primitive is parameterized by

$$\Theta_i^0 = \left\{\mathbf{n}_i^0,\ \mathbf{s}_i^0,\ \rho_i^0,\ a_i^0\right\}, \qquad \begin{cases} \mathbf{n}_i^0 \in \mathcal{M}, & \text{spherical center,} \\ \mathbf{s}_i^0 = (s_{1,i}^0, s_{2,i}^0) \in \mathbb{R}_+^2, & \text{anisotropic scale,} \\ \rho_i^0 \in \mathbb{R}, & \text{in-plane rotation,} \\ a_i^0 \in \mathbb{C}, & \text{complex spectrum attribute.} \end{cases} \tag{4}$$

The spherical center is parameterized by learnable azimuth φ_i^0 and elevation λ_i^0, whose valid ranges are enforced via

$$\mathbf{n}_i^0 = \mathbf{d}(\varphi_i^0, \lambda_i^0), \qquad \varphi_i^0 \in [0, 2\pi), \quad \lambda_i^0 \in [0, \pi/2]. \tag{5}$$

where $\mathbf{d}(\varphi, \lambda)$ denotes the spherical coordinate mapping from azimuth and elevation to $\mathcal{M}$:

$$\mathbf{d}(\varphi, \lambda) = [\cos\lambda\cos\varphi \quad \cos\lambda\sin\varphi \quad \sin\lambda]^\top. \tag{6}$$

The angular parameterization ensures $\|\mathbf{n}_i^0\|_2 = 1$ and $(n_i^0)_z \geq 0$. The anisotropic scales are parameterized using unconstrained variables $\boldsymbol{\ell}_i^0 = (\ell_{1,i}^0, \ell_{2,i}^0)$ as

$$s_{k,i}^0 = s_{\min} + \mathrm{softplus}(\ell_{k,i}^0), \qquad k \in \{1, 2\}. \tag{7}$$

This guarantees the non-negativity of Gaussian scales.
Given $\mathbf{n}_i^0$, we construct a reference orthonormal tangent basis $\mathbf{B}_i^0 = [\mathbf{b}_{1,i}^0, \mathbf{b}_{2,i}^0]$ of $T_{\mathbf{n}_i^0}\mathbb{S}^2$ by selecting one of the Cartesian axes that is not nearly aligned with $\mathbf{n}_i^0$ as the reference direction and forming the basis through normalized cross product. We then rotate this reference basis by an angle ρ_i^0 to obtain the Gaussian principal-axis frame:

$$\mathbf{T}_i^0 = \mathbf{B}_i^0\mathbf{R}(\rho_i^0) = [\mathbf{t}_{1,i}^0, \mathbf{t}_{2,i}^0], \qquad \mathbf{R}(\rho) = \begin{bmatrix} \cos\rho & -\sin\rho \\ \sin\rho & \cos\rho \end{bmatrix}. \tag{8}$$

The scales $(s_{1,i}^0, s_{2,i}^0)$ act along $\mathbf{t}_{1,i}^0$ and $\mathbf{t}_{2,i}^0$, respectively, determining the anisotropic extent of the primitive. ρ_i^0 denotes the in-plane orientation and a_i^0 is the complex spectrum attribute. These primitives form the canonical representation, which are subsequently deformed according to the TX–RX configuration and rasterized on $\mathcal{M}$ to synthesize the RF spectrum.

4.2 TX-RX-Conditioned Primitive Deformation

Deformation backbone. Given a TX-RX configuration $\mathbf{s} = (\mathbf{p}_t, \mathbf{p}_r)$, a shared deformation network f_θ takes the positional encodings of the canonical primitive center and the TX-RX positions as input and predicts primitive-wise attribute residuals:

$$(\Delta\boldsymbol{\xi}_i^{\mathbf{s}}, \Delta\boldsymbol{\ell}_i^{\mathbf{s}}, \Delta\rho_i^{\mathbf{s}}, \Delta a_i^{\mathbf{s}}) = f_\theta\big(\gamma(\mathbf{n}_i^0), \gamma(\mathbf{p}_t), \gamma(\mathbf{p}_r)\big), \tag{9}$$

where $\gamma(\cdot)$ denotes the positional encoding (PE), $\Delta\boldsymbol{\xi}_i^{\mathbf{s}} \in \mathbb{R}^2$ is the center displacement in local tangent coordinates, $\Delta\boldsymbol{\ell}_i^{\mathbf{s}} \in \mathbb{R}^2$ is the scale residual, $\Delta\rho_i^{\mathbf{s}} \in \mathbb{R}$ is the in-plane rotation residual, and $\Delta a_i^{\mathbf{s}} \in \mathbb{C}$ is the complex-attribute residual. The deformation network is shared across all primitives. The residuals are applied via manifold-consistent primitive deformation, as illustrated in Fig. 3.

Spherical deformation. The predicted center displacement is first lifted from its local 2D coordinates to a tangent vector at the canonical center and then mapped onto the sphere through the exponential map:

$$\begin{aligned} \mathbf{v}_i^{\mathbf{s}} &= \mathbf{B}_i^0\Delta\boldsymbol{\xi}_i^{\mathbf{s}} \in T_{\mathbf{n}_i^0}\mathbb{S}^2, \\ \mathbf{n}_i^{\mathbf{s}} &= \mathrm{Exp}_{\mathbf{n}_i^0}(\mathbf{v}_i^{\mathbf{s}}) \in \mathbb{S}^2. \end{aligned} \tag{10}$$

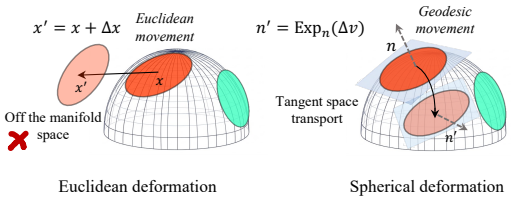

Figure 3: Illustration of our primitive deformation. Direct Euclidean offsets can move primitive centers away from the valid spherical domain $\mathcal{M}$, whereas tangent-space deformation followed by the exponential map preserves manifold-consistent evolution.

We do not restrict $\mathbf{n}_i^{\mathbf{s}}$ to remain inside $\mathcal{M}$, since a virtual arriving response may move outside the observable hemisphere as the TX-RX configuration changes, thereby avoiding boundary clipping. After the center moves from $\mathbf{n}_i^0$ to $\mathbf{n}_i^{\mathbf{s}}$, the canonical tangent basis is parallel-transported to the new location:

$$\mathbf{B}_i^{\mathbf{s}} = \mathrm{PT}_{\mathbf{n}_i^0 \to \mathbf{n}_i^{\mathbf{s}}}(\mathbf{B}_i^0), \qquad (\mathbf{B}_i^{\mathbf{s}})^\top\mathbf{n}_i^{\mathbf{s}} = \mathbf{0}. \tag{11}$$

The in-plane orientation and anisotropic scales are then updated in the transported tangent space:

$$\mathbf{T}_i^{\mathbf{s}} = \mathbf{B}_i^{\mathbf{s}}\mathbf{R}(\rho_i^0 + \Delta\rho_i^{\mathbf{s}}), \qquad s_{k,i}^{\mathbf{s}} = s_{\min} + \mathrm{softplus}(\ell_{k,i}^0 + \Delta\ell_{k,i}^{\mathbf{s}}). \tag{12}$$

Finally, the complex spectrum attribute is updated as

$$a_i^{\mathbf{s}} = a_i^0 + \Delta a_i^{\mathbf{s}} \in \mathbb{C}. \tag{13}$$

The resulting primitive $\Theta_i^{\mathbf{s}}$ therefore captures configuration-dependent variations in center, scale, orientation, and complex response while remaining stay on the sphere.

4.3 SPHERICAL RASTERIZATION

Given the deformed primitives $\{\Theta_i^{\mathbf{s}}\}_{i=1}^N$, each query direction $\mathbf{d}(\varphi,\lambda) \in \mathcal{M}$ is first mapped to the local tangent frame of primitive i through the logarithmic mapping $\mathbf{z}_i^{\mathbf{s}}(\varphi,\lambda) = (\mathbf{T}_i^{\mathbf{s}})^\top \operatorname{Log}_{\mathbf{n}_i^{\mathbf{s}}}(\mathbf{d}(\varphi,\lambda))$. This provides intrinsic local coordinates whose norm is the geodesic distance on the sphere. The complex spectrum is rendered by combining the contributions from all primitive responses using

$$S(\varphi,\lambda \mid \mathbf{s}) = \sum_{i=1}^{N} a_i^{\mathbf{s}} G_i^{\mathbf{s}}(\varphi,\lambda), \qquad G_i^{\mathbf{s}}(\varphi,\lambda) = \exp\left(-\frac{1}{2}\sum_{k=1}^{2}\frac{z_{k,i}^{\mathbf{s}}(\varphi,\lambda)^2}{(s_{k,i}^{\mathbf{s}})^2}\right). \tag{14}$$

Here, $G_i^{\mathbf{s}}$ denotes the intrinsic spherical response of primitive i, and S denotes the reconstructed complex RF spectrum. Since both response evaluation and accumulation are performed directly on $\mathcal{M}$, the Gaussian kernel is applied without intermediate geometric transformation and therefore preserves its intrinsic shape.

4.4 TRAINING STRATEGY

We train SphGaS end-to-end through the customized differentiable rasterizer. For each spectrum, the overall objective is

$$\mathcal{L} = \lambda_c \mathcal{L}_c + \lambda_a \mathcal{L}_a + \lambda_p \mathcal{L}_p + \lambda_s(1 - \text{S-SSIM}), \tag{15}$$

where $\mathcal{L}_c$ is the ℓ_1 loss on the normalized real and imaginary channels, $\mathcal{L}_a$ is the ℓ_1 loss on the normalized amplitude, and $\mathcal{L}_p$ is the ℓ_1 loss on the wrapped phase difference in radians. The S-SSIM term further encourages structural consistency over the spherical spectrum; its definition and implementation details are provided in Sec. F.4. For our customized datasets, both directions of each link are jointly sampled during training, providing reciprocal supervision for the shared deformation model. The canonical spherical primitives and deformation network are jointly optimized through our differentiable rendering pipeline.

Table 1: Comparison on the RFID real-world dataset. Best results are in bold and second-best results are underlined.

Method	Venue	Tx-RX Mobility	Planar Rendering	Model Space	PSNR (dB) ↑	SSIM ↑	S-SSIM ↑	MAE ↓	Size (MB) ↓	Infer. (ms) ↓
NeRF^2	MobiCom'23	✓	✗	$\mathbb{R}^3$	21.76	0.7732	0.7852	0.0710	8.00	230
WRF-GS+	TWC'25	✓	✓	$\mathbb{R}^3$	<u>25.23</u>	0.8885	<u>0.8896</u>	<u>0.0530</u>	11.30	8
GSRF	NeurIPS'25	✗	✗	$\mathbb{R}^3$	22.05	0.8497	0.8591	0.0697	69.00	23
SwiftWRF	TVCG'26	✗	✓	$\mathbb{R}^2$	24.75	0.8714	0.8536	0.0580	**3.50**	**0.02**
GAI-GS	CVPR'26	✗	✓	$\mathbb{R}^3$	–	<u>0.9100</u>	–	–	–	17.3
SphGaS	**Ours**	✓	✗	$\mathcal{M}$	**26.72**	**0.9320**	**0.9583**	**0.0467**	**3.50**	<u>0.2</u>

5 EXPERIMENT

5.1 EXPERIMENTAL SETTING

Implementation. Each SphGaS model uses 10K canonical primitives together with a shared deformation network implemented as an 8-layer multilayer perceptron (MLP) with hidden dimension 256, followed by lightweight prediction heads for primitive residuals. Positional encoding uses 10 frequency levels for primitive directions and 6 for TX–RX coordinates. The differentiable rasterizer follows the `gsplat` rasterization pipeline Ye et al. (2025). Training consists of 20K iterations of canonical primitive optimization followed by 400K iterations of joint optimization of the primitives and deformation network. Spectra are reconstructed at resolution of 90×360. We use an $80/20$ train/test split over individual samples when bidirectional spectra are unavailable, and over bidirectional link pairs otherwise.

Datasets. We evaluate SphGaS on three categories of datasets: ① *RFID spectrum dataset*, where we follow NeRF^2 Zhao et al. (2023) and use the S23 dataset containing real-world spectra collected in a laboratory environment with fixed RX and moving TX, corresponding to TX-only mobility; ② *Ray-traced spectrum dataset*, including indoor scenes (BEDROOM,

CONFERENCEROOM, and OFFICE) generated using the MATLAB Communications Toolbox ray tracerMathWorks (2024), and outdoor scenes (MUNICH, ETOILE, and FLORENCE) generated using NVIDIA Sionna RT Hoydis et al. (2023) from realistic urban layouts, where both TX and RX locations vary to evaluate joint TX-RX mobility; and ③ *RadioMapSeer dataset*, where we additionally use simulated RSSI measurements Yapar et al. (2022) to evaluate downstream applications. The RFID and outdoor city datasets provide power angular spectra (PAS), while the indoor ray-traced datasets provide complex-valued spectra.

Baselines. We compare SphGaS with representative neural-field and Gaussian-based RF reconstruction methods. **(1)** NeRF^{2}Zhao et al. (2023) uses an implicit neural field with volumetric rendering. **(2)** WRF-GS+Wen et al. (2025; 2026) adopts deformable 3DGS for RF field modeling. **(3)** GSRF Yang et al. (2025) uses complex-valued 3DGS for RF spectrum synthesis. **(4)** GAI-GS Shen et al. (2026) augments 3DGS with geometric-algebra-informed attention. **(5)** SwiftWRF Liu et al. (2026) uses deformable 2DGS for efficient spectrum rendering. **(6)** Kriging serves as a classical spatial-interpolation baseline. Methods without native support for joint TX–RX mobility are excluded from the corresponding benchmarks, except SwiftWRF, whose deformation model is extended to condition on both TX and RX embeddings as a Euclidean counterpart to SphGaS.

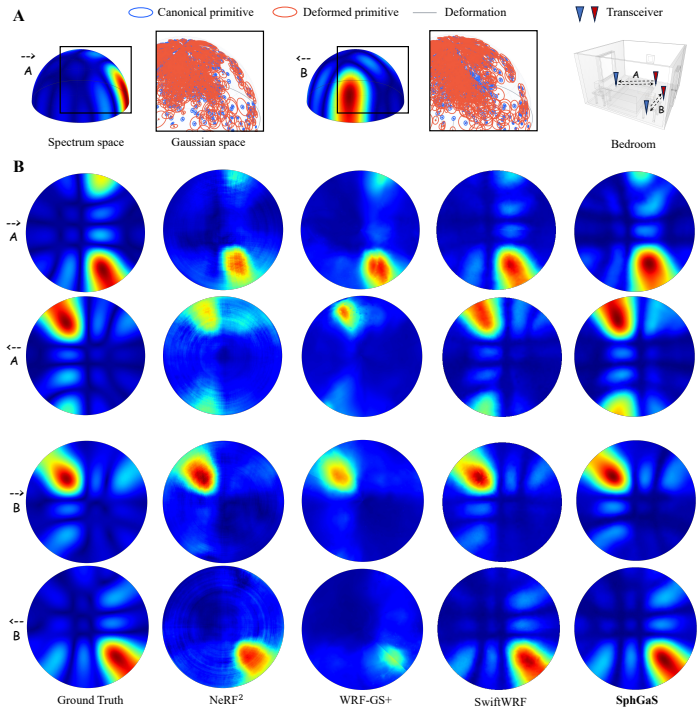

Figure 4: Visual comparison of reconstructed spectra under different TX-RX configurations. Zoom in for details.

Evaluation metrics. To maintain comparability with existing RF spectrum reconstruction benchmarks, we report peak signal-to-noise ratio (PSNR), mean absolute error (MAE), and structural similarity index (SSIM) Wang et al. (2004), where MAE corresponds to the ℓ_1 reconstruction error. Since the spectra are defined over a spherical angular domain, we additionally report S-SSIM Chen et al. (2018), which more appropriately accounts for the underlying spherical geometry. For complex-valued spectra, phase mean absolute error (Phase MAE) and root mean squared error (Phase RMSE) are used to evaluate phase reconstruction. Higher PSNR, SSIM, and S-SSIM indicate better performance, while lower MAE and phase errors are preferred.

5.2 SPECTRUM RECONSTRUCTION

Reconstruction quality. We first evaluate spectrum reconstruction on the real-world RFID dataset, where SphGaS achieves the best overall quality with 26.7 dB PSNR and 0.96 S-SSIM using a compact representation (Tab. 1). On the ray-traced city scenes, it improves the average S-SSIM over SwiftWRF from 0.66 to 0.81, while on indoor scenes it reaches 0.90 S-SSIM and reduces phase MAE by 11.5% to 0.74 rad (Tab. 2). These consistent gains demonstrate the advantage of spherical modeling under varying TX-RX configurations, as also reflected in Fig. 4. Phase visualizations are provided in Sec. D.5.

Efficiency and storage. As shown in Tab. 1, SphGaS requires only 3.50 MB storage and 0.2 ms per spectrum, providing a favorable trade-off between reconstruction quality, model size, and rendering efficiency. Compared with SwiftWRF, spherical modeling introduces additional latency in exchange for more accurate reconstruction, while remaining efficient relative to other GS-based baselines.

Generalization to unseen TX configurations. We evaluate two train–test settings that differ in

Table 2: Spectrum reconstruction performance on outdoor and indoor scenes. Higher is better for PSNR, SSIM, and S-SSIM, while lower is better for phase errors. Phase errors are reported in radians. Best results are in bold and second-best results are underlined.

City Scenes					Indoor Scenes						
Scene	Method	PSNR (dB) ↑	SSIM ↑	S-SSIM ↑	Scene	Method	PSNR (dB) ↑	SSIM ↑	S-SSIM ↑	Phase MAE ↓	Phase RMSE ↓
Munich	Kriging	16.45	0.6048	0.5871	Conf.Room	Kriging	17.86	0.5666	0.5373	1.334	1.607
	NeRF^2	22.44	0.6761	0.6814		NeRF^2	19.57	0.6579	0.6548	1.227	1.493
	WRF-GS+	21.32	0.6881	0.6828		WRF-GS+	18.24	0.6314	0.6047	1.094	1.361
	SwiftWRF	22.91	0.7018	0.6937		SwiftWRF	24.31	0.8187	0.7646	0.838	1.132
	SphGaS	**24.23**	**0.8351**	**0.8792**		**SphGaS**	**25.33**	**0.8524**	**0.9104**	**0.752**	**1.073**
Etoile	Kriging	16.03	0.5921	0.5885	Bedroom	Kriging	18.85	0.6287	0.5703	1.311	1.579
	NeRF^2	21.17	0.6514	0.6643		NeRF^2	19.65	0.6608	0.6633	1.203	1.468
	WRF-GS+	22.12	0.6924	0.6901		WRF-GS+	19.25	0.6515	0.6306	1.079	1.342
	SwiftWRF	22.05	0.6885	0.6811		SwiftWRF	25.03	0.8386	0.7940	0.823	1.118
	SphGaS	**25.77**	**0.8293**	**0.8643**		**SphGaS**	**25.91**	**0.8561**	**0.9182**	**0.738**	**1.068**
Florence	Kriging	15.21	0.5584	0.5562	Office	Kriging	16.19	0.5009	0.5206	1.329	1.592
	NeRF^2	17.98	0.6133	0.6124		NeRF^2	17.39	0.5447	0.5237	1.221	1.481
	WRF-GS+	17.53	0.6177	0.6057		WRF-GS+	17.54	0.5992	0.5405	1.136	1.397
	SwiftWRF	18.42	0.6215	0.6148		SwiftWRF	21.91	0.7481	0.7121	0.847	1.137
	SphGaS	**18.75**	**0.6370**	**0.6774**		**SphGaS**	**22.84**	**0.7815**	**0.8613**	**0.731**	**1.066**
Average	Kriging	15.90	0.5851	0.5773	Average	Kriging	17.63	0.5654	0.5427	1.325	1.593
	NeRF^2	20.53	0.6469	0.6527		NeRF^2	18.87	0.6211	0.6139	1.217	1.481
	WRF-GS+	20.32	0.6661	0.6595		WRF-GS+	18.34	0.6274	0.5919	1.103	1.367
	SwiftWRF	21.13	0.6706	0.6632		SwiftWRF	23.75	0.8018	0.7569	0.836	1.129
	SphGaS	**22.92**	**0.7671**	**0.8070**		**SphGaS**	**24.69**	**0.8300**	**0.8966**	**0.740**	**1.069**

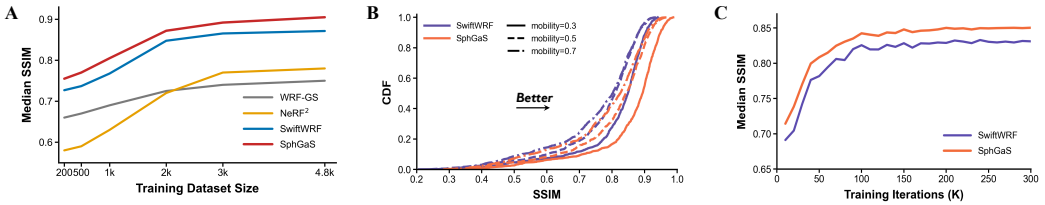

Figure 6: Analysis of SphGaS under different settings: **(A)** training dataset size, **(B)** TX mobility, and **(C)** training convergence.

whether the test TX has been observed during training. In the seen-TX setting, each TX is included in training, while 20% of its RX links are held out for testing; this evaluates spatial interpolation to unseen RX locations under a known TX. In the unseen-TX setting, entire TX locations together with all their associated RX links are excluded from training, so the model must predict spectra for a transmitter-induced RF field that has not been observed during training. As shown in Fig. 5, SphGaS performs best in both settings and shows only a 0.94 dB PSNR drop on unseen TXs, compared with 2.21 dB for SwiftWRF.

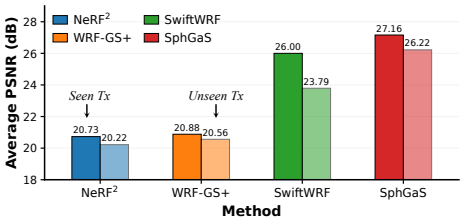

Figure 5: Performance under seen- and unseen-TX settings.

Robustness and optimization analysis. Fig. 6 further examines SphGaS under different training conditions. In **(A)**, SphGaS consistently achieves higher SSIM across different training-set sizes and reaches stable performance with fewer samples, indicating better data efficiency under limited measurements. In **(B)**, SphGaS maintains a better SSIM distribution than SwiftWRF across different TX mobility ranges, showing stronger robustness to increasing TX-RX variation. Here, the mobility ratio denotes the TX movement range relative to the room size, and a ratio of zero corresponds to a static RF field with only RX movement. In **(C)**, SphGaS improves faster during the early stage of training and converges to a higher SSIM, indicating more favorable and stable optimization behavior for spectrum reconstruction.

5.3 Ablation Study

We conduct ablation studies on the indoor scenes, focusing on the effects of deformation and training components. Additional analyses are provided in Sec. D.4.

Table 3: Ablation of TX-RX-conditioned deformation components. Phase errors are reported in radians. Higher is better for PSNR and SSIM, while lower is better for MAE and phase errors. For PSNR, values in parentheses denote absolute drops in dB from the full model; for the remaining metrics, they denote relative changes.

Deformation Components				Reconstruction Quality				
$\Delta\boldsymbol{\xi}$	$\Delta\boldsymbol{\ell}$	$\Delta\rho$	Δa	PSNR (dB) ↑	SSIM ↑	MAE ↓	Phase MAE (rad.) ↓	Phase RMSE (rad.) ↓
✓	✓	✓	✓	**24.69**	**0.8300**	**0.0531**	**0.740**	**1.069**
✗	✓	✓	✓	21.84 (-2.85 dB)	0.7618 (-8.2%)	0.0612 (+15.3%)	0.832 (+12.4%)	1.185 (+10.9%)
✓	✗	✓	✓	23.27 (-1.42 dB)	0.8017 (-3.4%)	0.0554 (+4.3%)	0.771 (+4.2%)	1.119 (+4.7%)
✓	✓	✗	✓	22.91 (-1.78 dB)	0.7936 (-4.4%)	0.0576 (+8.5%)	0.786 (+6.2%)	1.148 (+7.4%)
✓	✓	✓	✗	23.58 (-1.11 dB)	0.8089 (-2.5%)	0.0541 (+1.9%)	0.763 (+3.1%)	1.101 (+3.0%)

Deformation components. Tab. 3 evaluates the contribution of each component in the TX-RX-conditioned spherical deformation. Center displacement $\Delta\boldsymbol{\xi}$ is the most influential component: removing it reduces PSNR from 24.69 to 21.84 dB and increases MAE by 15.3%, with similarly clear degradation in phase reconstruction. Rotation and scale deformation are also important for adapting the orientation and angular extent of the primitives, while complex-attribute deformation further improves both reconstruction and phase accuracy. Overall, removing any component degrades all metrics, showing that center, shape, and complex-response variations provide complementary information for modeling TX-RX-dependent spectrum changes.

Training components. Tab. 4 shows that all three training components contribute positively but play different roles. Bidirectional training has the largest effect, causing a 5.47% PSNR drop when removed, highlighting the benefit of physical reciprocity, where reversing the TX–RX direction provides complementary supervision on spectrum variation. In contrast, complex-domain supervision stabilizes training and improves complex-response fidelity, while the S-SSIM loss further improves spherical structural consistency. Together, these components provide complementary supervision for accurate spectrum reconstruction.

5.4 Applications

Beyond spectrum synthesis, SphGaS can be readily adapted to support downstream RF applications. First, the synthesized spectrum can be aggregated over the reception hemisphere with solid-angle weighting to obtain an RSSI estimate at a given TX–RX configuration, enabling coverage-aware applications such as TX deployment and communication-aware navigation. Second, its directional structure preserves propagation cues that can be used for LOS analysis and coarse scene understanding. On the real-world BLE dataset, SphGaS achieves an RSSI reconstruction MAE of 2.85 dB and the lowest direct-path angular errors among the compared methods. Detailed evaluations are provided in Sec. E.1 and Sec. E.2.

Table 4: Ablation study of the training settings. Higher is better for all metrics. For PSNR, values in parentheses denote absolute drops in dB from the full setting; for SSIM and S-SSIM, they denote relative changes.

Training Setting	PSNR (dB) ↑	SSIM ↑	S-SSIM ↑
Full setting	**24.69**	**0.8300**	**0.8966**
w/o Complex loss	24.46 (-0.23 dB)	0.8288 (-0.15%)	0.8894 (-0.80%)
w/o Bidirectional training	23.34 (-1.35 dB)	0.7802 (-6.00%)	0.8594 (-4.15%)
w/o S-SSIM loss	24.58 (-0.11 dB)	0.8292 (-0.10%)	0.8881 (-0.95%)

6 Conclusion

We presented SphGaS, a spherical Gaussian splatting framework for RF spectrum synthesis under varying TX-RX configurations. SphGaS represents virtual arriving signal responses as Gaussian primitives on the reception sphere and models configuration-dependent variations through deformation on the spherical manifold. Together with our customized differentiable rasterizer, this formulation avoids geometric mismatch introduced by planar rendering pipelines while supporting joint TX–RX mobility. Experiments on real-world and simulated datasets demonstrate improved spectrum reconstruction while maintaining efficient rendering. Beyond spectrum synthesis, the learned representation adapts naturally to downstream RF tasks and shows potential for spherical visual modeling, such as 360° image and video representation (see Sec. E.3 for an example).

AI USE STATEMENT

Generative AI tools were used to refine the writing of the manuscript, assist with software code development, and support correctness checking of parts of the theoretical error analysis. The research ideas, methodology, and main theoretical arguments were developed by the authors. All AI-assisted code and theoretical analyses were independently reviewed and verified by the authors. We take full responsibility for the final content of this work, including all text, code, mathematical claims, and results.

REFERENCES

3GPP. Study on channel model for frequencies from 0.5 to 100 ghz. Technical Report TR 38.901, 3rd Generation Partnership Project, 2024.

Chaitanya Amballa, Yu-Lin Wei, Sattwik Basu, Zhijian Yang, Mehmet Ergezer, and Romit Roy Choudhury. Can nerfs "see" without cameras? In *Advances in Neural Information Processing Systems*, volume 38, 2025.

Stefanos Bakirtzis, Jun Chen, Kehai Qiu, Ian Wassell, and Jie Zhang. Em deepray: An expedient, generalizable, and realistic data-driven indoor propagation model. *IEEE Transactions on Antennas and Propagation*, 70(6):4140–4154, 2022.

Amartya Basu and Ayon Chakraborty. Specnerf: Neural radiance field driven wireless coverage mapping for 5g networks. In *Proceedings of the 25th International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (MobiHoc)*, pp. 440–445, 2024.

Juan A. Bazerque, Gonzalo Mateos, and Georgios B. Giannakis. Group-lasso on splines for spectrum cartography. *IEEE Transactions on Signal Processing*, 59(10):4648–4663, 2011.

Ana Chaves-Villota and Carlos A. Viteri-Mera. Deeprem: Deep-learning-based radio environment map estimation from sparse measurements. *IEEE Access*, 2023.

Anthony Chen, Shiwen Mao, Zhu Li, Hongliang Zhang, Dusit Niyato, and Zhu Han. A gaussian splatting approach to continuous radio map construction. In *IEEE INFOCOM 2025 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 1–6, 2025a. doi: 10.1109/INFOCOMWKSHPS65812.2025.11153005.

Sijia Chen, Yingxue Zhang, Yiming Li, Zhenzhong Chen, and Zhou Wang. Spherical structural similarity index for objective omnidirectional video quality assessment. In *2018 IEEE International Conference on Multimedia and Expo (ICME)*, pp. 1–6. IEEE, 2018. doi: 10.1109/ICME.2018.8486584.

Xiaonan Chen, Jun Wang, and Qingyang Huang. Dynamic spectrum cartography: Reconstructing spatial-spectral-temporal radio frequency map via tensor completion. *IEEE Transactions on Signal Processing*, 73:1184–1199, 2025b. doi: 10.1109/TSP.2025.3531872.

Xingyu Chen, Zihao Feng, Ke Sun, Kun Qian, and Xinyu 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 (SenSys)*, 2024. doi: 10.1145/3666025.3699351.

Xingyu Chen, Zihao Feng, Kun Qian, and Xinyu Zhang. Radio frequency ray tracing with neural object representation for enhanced rf modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 21339–21348, 2025c.

Zhiwen Fan, Kevin Wang, Kairun Wen, Zehao Zhu, Dejia Xu, and Zhangyang Wang. Lightgaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps. In *Advances in Neural Information Processing Systems*, volume 37, pp. 140138–140158, 2024.

Bo Gao and Tianyue Zheng. Poster: Gluert: Bridging the gap between ray-tracing and real-world wireless channel. In *Proceedings of the 23rd Annual International Conference on Mobile Systems, Applications and Services (MobiSys)*, pp. 585–586, 2025. doi: 10.1145/3711875.3734546.

Masaharu Hata. Empirical formula for propagation loss in land mobile radio services. *IEEE Transactions on Vehicular Technology*, 29(3):317–325, 1980.

Jakob Hoydis, Fayçal Ait Aoudia, Sebastian Cammerer, Merlin Nimier-David, Nikolaus Binder, Guillermo Marcus, and Alexander Keller. Sionna rt: Differentiable ray tracing for radio propagation modeling. In *IEEE Global Communications Conference Workshops (GLOBECOM Workshops)*, 2023.

Fabian Jaensch, Giuseppe Caire, and Begüm Demir. Radio map estimation—an open dataset with directive transmitter antennas and initial experiments, 2024. arXiv:2402.00878.

Shuaifeng Jiang, Qi Qu, Xiaqing Pan, Abhishek Agrawal, Richard Newcombe, and Ahmed Alkhateeb. Learnable wireless digital twins: Reconstructing electromagnetic field with neural representations. *IEEE Open Journal of the Communications Society*, 6:1568–1590, 2025. doi: 10.1109/OJCOMS.2025.3535959.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. In *ACM SIGGRAPH*, 2023.

Puja Kumari, Junil Choi, Nuria Gonzalez-Prelcic, and Robert W. Heath. Ieee 802.11ad-based radar: An approach to joint vehicular communication-radar system. *IEEE Journal of Selected Topics in Signal Processing*, 12(6):1471–1485, 2018. doi: 10.1109/JSTSP.2018.2872511.

Ju-Hyung Lee and Andreas F. Molisch. A scalable and generalizable pathloss map prediction. *IEEE Transactions on Wireless Communications*, 23(11):17793–17806, 2024. doi: 10.1109/TWC.2024.3457431.

Ron Levie, Çağkan Yapar, Gitta Kutyniok, and Giuseppe Caire. Radiounet: Fast radio map estimation with convolutional neural networks. *IEEE Transactions on Wireless Communications*, 20(6):4001–4015, 2021. doi: 10.1109/TWC.2021.3054977.

Kun Li, Peiming Li, Yong Zeng, and Jie Xu. Channel knowledge map for environment-aware communications: Em algorithm for map construction. In *IEEE Wireless Communications and Networking Conference (WCNC)*, 2022.

Mufan Liu, Cixiao Zhang, Qi Yang, Yujie Cao, Yiling Xu, Yin Xu, Shu Sun, Mingzeng Dai, and Yunfeng Guan. Deformable 2d gaussian splatting for efficient wireless radiance field rendering. *IEEE Transactions on Visualization and Computer Graphics*, 32(7):6363–6378, 2026. doi: 10.1109/TVCG.2026.3681115.

Haofan Lu, Christopher Vattheuer, Baharan Mirzasoleiman, and Omid Abari. Newrf: A deep learning framework for wireless radiation field reconstruction and channel prediction. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024.

MathWorks. Ray tracing for wireless communications. MATLAB Communications Toolbox Documentation, 2024.

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision (ECCV)*, 2020.

Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4), 2022.

Bhavya Sai Nukapotula, Rishabh Tripathi, Seth Pregler, Dileep Kalathil, Srinivas Shakkottai, and Theodore S. Rappaport. Gsparc: Gaussian splatting for real-time reconstruction of rf channels. *arXiv preprint arXiv:2511.22793*, 2025.

Theodore S. Rappaport, Yunchou Xing, George R. MacCartney, Andreas F. Molisch, Evangelos Mellios, and Jianhua Zhang. Overview of millimeter wave communications for fifth-generation (5g) wireless networks—with a focus on propagation models. *IEEE Transactions on Antennas and Propagation*, 65(12):6213–6230, 2017.

Daniel Romero and Seung-Jun Kim. Radio map estimation: A data-driven approach to spectrum cartography. *IEEE Signal Processing Magazine*, 39(6):53–72, 2022.

Adel A. M. Saleh and Reinaldo A. Valenzuela. A statistical model for indoor multipath propagation. *IEEE Journal on Selected Areas in Communications*, 5(2):128–137, 1987.

Jingzhou Shen, Tianya Zhao, and Xuyu Wang. A geometric algebra-informed 3dgs framework for wireless channel prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4973–4982, 2026.

Sagar Shrestha, Xiao Fu, and Mingyi Hong. Deep spectrum cartography: Completing radio map tensors using learned neural models. *IEEE Transactions on Signal Processing*, 70, 2022. doi: 10.1109/TSP.2022.3145190.

Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Sheng Wang, Hengtao He, Chaozheng Wen, Jingwen Tong, Xinyu Li, Xiao Li, Jun Zhang, and Shi Jin. Xfreq-gs: Cross-frequency wireless radiation field reconstruction with 3d gaussian splatting, 2026a. arXiv:2605.11432.

Yiheng Wang, Ye Xue, Shutao Zhang, Hongmiao Fan, and Tsung-Hui Chang. Radcloudsplat: Scatterer-driven 3d gaussian splatting with point-cloud priors for radiomap extrapolation. In *IEEE INFOCOM 2026 - IEEE Conference on Computer Communications*, 2026b.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. doi: 10.1109/TIP.2003.819861.

Chaozheng Wen, Jingwen Tong, Yingdong Hu, Zehong Lin, and Jun Zhang. Wrf-gs: Wireless radiation field reconstruction with 3d gaussian splatting. In *IEEE INFOCOM*, 2025.

Chaozheng Wen, Jingwen Tong, Yingdong Hu, Zehong Lin, and Jun Zhang. Neural representation for wireless radiation field reconstruction: A 3d gaussian splatting approach. *IEEE Transactions on Wireless Communications*, 25:7490–7504, 2026. doi: 10.1109/TWC.2025.3631663.

Kang Yang and Mani Srivastava. Rxgs: Receiver-generalizable 3d gaussian splatting for radio-frequency data synthesis, 2026. arXiv:2605.24290.

Kang Yang, Gaofeng Dong, Sijie Ji, Wan Du, and Mani Srivastava. Gsrf: Complex-valued 3d gaussian splatting for efficient radio-frequency data synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.

Çağkan Yapar, Ron Levie, Gitta Kutyniok, and Giuseppe Caire. Dataset of pathloss and toa radio maps with localization application. IEEE DataPort, 2022.

Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari, Jianbo Ye, Jeffrey Hu, Matthew Tancik, and Angjoo Kanazawa. gsplat: An open-source library for gaussian splatting. *Journal of Machine Learning Research*, 26(34):1–17, 2025.

Zhengqing Yun and Magdy F. Iskander. Ray tracing for radio propagation modeling: Principles and applications. *IEEE Access*, 3:1089–1100, 2015.

Yong Zeng and Xiaoli Xu. Toward environment-aware 6g communications via channel knowledge map. *IEEE Wireless Communications*, 28(3):84–91, 2021.

Yong Zeng, Junting Chen, Jie Xu, Di Wu, Xiaoli Xu, Shi Jin, Xiqi Gao, David Gesbert, Shuguang Cui, and Rui Zhang. A tutorial on environment-aware communications via channel knowledge map for 6g. *IEEE Communications Surveys & Tutorials*, 26(3):1478–1519, 2024.

Jinghan Zhang, Xitao Gong, Qi Wang, Richard A. Stirling-Gallacher, and Giuseppe Caire. Planar gaussian splatting with bilinear spatial transformer for wireless radiance field reconstruction, 2026a. arXiv:2604.25945.

Lihao Zhang, Zongtan Li, and Haijian Sun. Rf-pgs: Fully-structured spatial wireless channel representation with planar gaussian splatting. *arXiv preprint arXiv:2508.16849*, 2025.

Lihao Zhang, Haijian Sun, Samuel Berweger, Camillo Gentile, and Rose Qingyang Hu. Rf-3dgs: Wireless channel modeling with radio radiance field and 3d gaussian splatting. *IEEE Transactions on Wireless Communications*, 25:10419–10433, 2026b. doi: 10.1109/TWC.2026.3652154.

Songyang Zhang, Achintha Wijesinghe, and Zhi Ding. Rme-gan: A learning framework for radio map estimation based on conditional generative adversarial network. *IEEE Internet of Things Journal*, 10(20):18016–18027, 2023.

Xiaopeng Zhao, Zhenlin An, Qingrui Pan, and Lei Yang. Nerf2: Neural radio-frequency radiance fields. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2023. doi: 10.1145/3570361.3592527.

– Appendix –

Appendix Contents

A Nomenclature **14**

B Physical Interpretation and Modeling Assumptions **14**

C Error Analysis **16**

C.1 Geometric Analysis of Primitive Distortion 16

C.2 Analysis of Tile Assignment Accuracy 18

D Supplementary Experiments **19**

D.1 Spherical vs. Euclidean Gaussian Representations 19

D.2 Primitive Count and Pruning 19

D.3 Novel Spectrum vs. Novel Field Synthesis 19

D.4 Ablation Study 21

D.5 Visualization 22

E Applications **23**

E.1 RSSI Reconstruction 23

E.2 Directional Propagation Analysis 26

E.3 Extension to Visual Applications 26

F Implementation details **27**

F.1 Dataset Information 27

F.2 Baseline Implementation Details. 28

F.3 SphGaS Implementation 28

F.4 Evaluation Metrics Details 30

G Limitations **31**

A NOMENCLATURE

Symbol	Description
RF spectrum and configuration	
$\mathcal{M} \subset \mathbb{S}^2$	Reception hemisphere on which the RF spatial spectrum is observed.
$\mathbf{s} = (\mathbf{p}_t, \mathbf{p}_r)$	TX–RX configuration with transmitter and receiver positions $\mathbf{p}_t$ and $\mathbf{p}_r$.
(φ, λ)	Azimuth and elevation angles, with $\varphi \in [0, 2\pi)$ and $\lambda \in [0, \pi/2]$.
$\mathbf{d}(\varphi, \lambda) \in \mathcal{M}$	Spherical direction corresponding to azimuth–elevation coordinates (φ, λ).
$S(\varphi, \lambda \mid \mathbf{s})$	Complex-valued RF spatial spectrum under configuration $\mathbf{s}$.
$\mathbf{h}_m(\mathbf{s})$	Complex array-response vector of the m-th propagation component.
Beam-steering interpretation	
$\mathbf{w}(\mathbf{d})$	Array steering vector toward spherical direction $\mathbf{d}$.
$g_m(\mathbf{d} \mid \mathbf{s})$	Beam-steered directional response of the m-th propagation path.
$\alpha_m(\mathbf{s})$	Complex coefficient of the m-th propagation path.
$\mathbf{r}_n,\ k$	Position of the n-th antenna element and carrier wavenumber $k = 2\pi/\lambda_c$.
$\mathbf{B}_m(\mathbf{s})$	Orthonormal tangent basis at the path direction $\mathbf{d}_m(\mathbf{s})$.
$\boldsymbol{\xi}_m(\mathbf{d}; \mathbf{s})$	Intrinsic tangent coordinate of query direction $\mathbf{d}$ around $\mathbf{d}_m(\mathbf{s})$.
$\mathbf{H}_m(\mathbf{s})$	Local quadratic form of the beam-steered response, $\mathbf{H}_m = k^2 \mathbf{B}_m^\top \mathbf{C}_r \mathbf{B}_m$, where $\mathbf{C}_r$ is the antenna-position second-moment matrix.
Spherical Gaussian representation	
Θ_i^0	Primitive parameterized by $\{\mathbf{n}_i^0, \boldsymbol{\ell}_i^0, \rho_i^0, a_i^0\}$ and treated as the canonical representation for deformation.
$\Theta_i^{\mathbf{s}}$	Primitive after transceiver-conditioned deformation under configuration $\mathbf{s}$.
$\mathbf{n}_i^0, \mathbf{n}_i^{\mathbf{s}}$	Initial center $\mathbf{n}_i^0 \in \mathcal{M}$ and deformed center $\mathbf{n}_i^{\mathbf{s}} \in \mathbb{S}^2$.
$\boldsymbol{\ell}_i^0, s_{k,i}^{\mathbf{s}}$	Scale parameters and deformed anisotropic scales.
$\rho_i^0, \Delta\rho_i^{\mathbf{s}}$	Initial in-plane orientation and its configuration-dependent residual.
$a_i^0, a_i^{\mathbf{s}}$	Initial and deformed complex spectrum attributes.
$\mathbf{B}_i^0, \mathbf{B}_i^{\mathbf{s}}$	Reference tangent basis and its parallel-transported counterpart.
$\mathbf{T}_i^0, \mathbf{T}_i^{\mathbf{s}}$	Gaussian principal-axis frames before and after deformation.
Spherical deformation	
f_θ	Transceiver-conditioned deformation network.
$\Delta\boldsymbol{\xi}_i^{\mathbf{s}}$	Center displacement in local tangent coordinates.
$\Delta\boldsymbol{\ell}_i^{\mathbf{s}}$	Scale residual under configuration $\mathbf{s}$.
$\Delta\rho_i^{\mathbf{s}}$	In-plane rotation residual under configuration $\mathbf{s}$.
$\Delta a_i^{\mathbf{s}}$	Complex spectrum-attribute residual under configuration $\mathbf{s}$.
$\mathbf{v}_i^{\mathbf{s}}$	Tangent displacement lifted from $\Delta\boldsymbol{\xi}_i^{\mathbf{s}}$.
$\mathrm{Exp}_{\mathbf{n}}(\cdot)$	Exponential map from a tangent space onto the sphere.
$\mathrm{PT}_{\mathbf{n}\to\mathbf{n}'}(\cdot)$	Parallel transport between tangent spaces on $\mathbb{S}^2$.
Spherical rasterization	
$\mathrm{Log}_{\mathbf{n}}(\cdot)$	Logarithmic map from the sphere to the tangent space at $\mathbf{n}$.
$\mathbf{z}_i^{\mathbf{s}}(\mathbf{d})$	Local coordinates of spherical direction $\mathbf{d}$ in the Gaussian frame.
$G_i^{\mathbf{s}}(\mathbf{d})$	Spherical Gaussian response of primitive i along direction $\mathbf{d}$.
$S(\varphi, \lambda \mid \mathbf{s})$	Rendered complex RF spectrum at direction $\mathbf{d}(\varphi, \lambda)$ under configuration $\mathbf{s}$.

Table 5: Main notation used in SphGaS.

B PHYSICAL INTERPRETATION AND MODELING ASSUMPTIONS

This section provides a local physical interpretation of the spherical Gaussian primitives used in SphGaS. Under a standard narrowband far-field array model, we show that the directional response of an individual propagation path admits a local second-order approximation by an anisotropic Gaussian in the intrinsic tangent coordinates of the sphere.

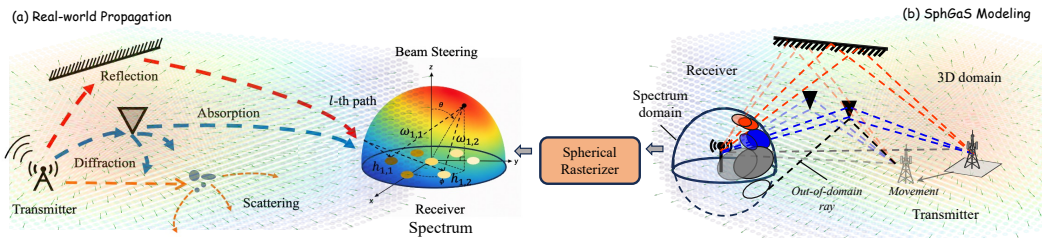

Figure 7: Physical interpretation of spherical primitives and their deformation. Localized responses vary in direction, angular profile, and complex contribution as the TX–RX configuration changes.

We assume that each arriving path is approximately planar across the RX array aperture, as is standard in the far-field regime. A commonly used criterion is $R \gg 2D^2/\lambda_c$, where R is the propagation distance, D is the array aperture, and λ_c is the carrier wavelength. This condition depends on the array aperture rather than the overall size of the TX–RX environment. We further assume a narrowband response and a fixed array orientation across TX–RX configurations. These assumptions are introduced only for the physical interpretation below and are not explicit constraints of the SphGaS representation or rasterizer.

Recall from the main text that the complex RF spatial spectrum is obtained by beam steering and superposing multiple propagation paths:

$$S(\mathbf{d} \mid \mathbf{s}) = \sum_{m=1}^{N_p} g_m(\mathbf{d} \mid \mathbf{s}), \qquad g_m(\mathbf{d} \mid \mathbf{s}) = \mathbf{w}(\mathbf{d})^H \mathbf{h}_m(\mathbf{s}), \tag{16}$$

where $\mathbf{d} = \mathbf{d}(\varphi, \lambda) \in \mathcal{M} \subset \mathbb{S}^2$ denotes the spherical query direction.

Let $\mathbf{d}_m(\mathbf{s}) \in \mathbb{S}^2$ and $\alpha_m(\mathbf{s}) \in \mathbb{C}$ denote the arrival direction and complex coefficient of the m-th path, respectively. For an array with element locations $\{\mathbf{r}_n\}_{n=1}^{N_a}$, we adopt the standard far-field steering vector Kumari et al. (2018)

$$\mathbf{a}(\mathbf{d}) = \begin{bmatrix} e^{jk\mathbf{r}_1^\top \mathbf{d}} & \dots & e^{jk\mathbf{r}_{N_a}^\top \mathbf{d}} \end{bmatrix}^\top, \tag{17}$$

where $k = 2\pi/\lambda_c$ is the carrier wavenumber. Under the single-path plane-wave model,

$$\mathbf{h}_m(\mathbf{s}) = \alpha_m(\mathbf{s}) \mathbf{a}(\mathbf{d}_m(\mathbf{s})). \tag{18}$$

Using the normalized matched beamformer $\mathbf{w}(\mathbf{d}) = \mathbf{a}(\mathbf{d})/N_a$ gives

$$g_m(\mathbf{d} \mid \mathbf{s}) = \alpha_m(\mathbf{s}) \frac{1}{N_a} \sum_{n=1}^{N_a} \exp\left(jk\mathbf{r}_n^\top [\mathbf{d}_m(\mathbf{s}) - \mathbf{d}]\right). \tag{19}$$

We express the antenna locations relative to the array centroid, such that

$$\frac{1}{N_a} \sum_{n=1}^{N_a} \mathbf{r}_n = \mathbf{0}. \tag{20}$$

Proposition 1 (Local second-order Gaussian approximation)**.** *Let $\mathbf{B}_m(\mathbf{s}) \in \mathbb{R}^{3\times 2}$ be an orthonormal basis of the tangent space $T_{\mathbf{d}_m(\mathbf{s})}\mathbb{S}^2$, and define*

$$\boldsymbol{\xi}_m(\mathbf{d}; \mathbf{s}) = \mathbf{B}_m(\mathbf{s})^\top \operatorname{Log}_{\mathbf{d}_m(\mathbf{s})}(\mathbf{d}). \tag{21}$$

Then, as $\mathbf{d} \to \mathbf{d}_m(\mathbf{s})$,

$$g_m(\mathbf{d} \mid \mathbf{s}) = \alpha_m(\mathbf{s}) \left[\exp\left(-\frac{1}{2} \boldsymbol{\xi}_m^\top \mathbf{H}_m(\mathbf{s}) \boldsymbol{\xi}_m \right) + \mathcal{O}(\|\boldsymbol{\xi}_m\|^3) \right], \tag{22}$$

where

$$\mathbf{H}_m(\mathbf{s}) = k^2 \mathbf{B}_m(\mathbf{s})^\top \mathbf{C}_r \mathbf{B}_m(\mathbf{s}), \qquad \mathbf{C}_r = \frac{1}{N_a} \sum_{n=1}^{N_a} \mathbf{r}_n \mathbf{r}_n^\top. \tag{23}$$

Proof. Let

$$\mathbf{v}_m = \mathrm{Log}_{\mathbf{d}_m(\mathbf{s})}(\mathbf{d}) = \mathbf{B}_m(\mathbf{s})\boldsymbol{\xi}_m. \tag{24}$$

Using the local expansion of the exponential map on $\mathbb{S}^2$,

$$\mathbf{d}_m(\mathbf{s}) - \mathbf{d} = -\mathbf{B}_m(\mathbf{s})\boldsymbol{\xi}_m + \frac{1}{2}\|\boldsymbol{\xi}_m\|^2\mathbf{d}_m(\mathbf{s}) + \mathcal{O}(\|\boldsymbol{\xi}_m\|^3). \tag{25}$$

Expanding equation 19 around $\boldsymbol{\xi}_m = \mathbf{0}$ gives

$$\begin{aligned}\frac{g_m(\mathbf{d} \mid \mathbf{s})}{\alpha_m(\mathbf{s})} = 1 + \frac{jk}{N_a}\sum_{n=1}^{N_a}\mathbf{r}_n^\top[\mathbf{d}_m(\mathbf{s}) - \mathbf{d}] \\ - \frac{k^2}{2N_a}\sum_{n=1}^{N_a}\left(\mathbf{r}_n^\top[\mathbf{d}_m(\mathbf{s}) - \mathbf{d}]\right)^2 + \mathcal{O}(\|\boldsymbol{\xi}_m\|^3).\end{aligned} \tag{26}$$

Since $\sum_n \mathbf{r}_n = \mathbf{0}$, the linear term vanishes. Retaining the second-order contribution gives

$$\frac{g_m(\mathbf{d} \mid \mathbf{s})}{\alpha_m(\mathbf{s})} = 1 - \frac{1}{2}\boldsymbol{\xi}_m^\top\mathbf{H}_m(\mathbf{s})\boldsymbol{\xi}_m + \mathcal{O}(\|\boldsymbol{\xi}_m\|^3). \tag{27}$$

Meanwhile,

$$\exp\left(-\frac{1}{2}\boldsymbol{\xi}_m^\top\mathbf{H}_m(\mathbf{s})\boldsymbol{\xi}_m\right) = 1 - \frac{1}{2}\boldsymbol{\xi}_m^\top\mathbf{H}_m(\mathbf{s})\boldsymbol{\xi}_m + \mathcal{O}(\|\boldsymbol{\xi}_m\|^4). \tag{28}$$

The two expressions therefore agree through second order. □

Applying Proposition 1 pathwise gives the local approximation

$$S(\mathbf{d} \mid \mathbf{s}) \approx \sum_{m=1}^{N_p}\alpha_m(\mathbf{s})\exp\left(-\frac{1}{2}\boldsymbol{\xi}_m(\mathbf{d};\mathbf{s})^\top\mathbf{H}_m(\mathbf{s})\boldsymbol{\xi}_m(\mathbf{d};\mathbf{s})\right). \tag{29}$$

Remark 1: Primitive interpretation. Eq. 29 provides a local physical interpretation of the spherical Gaussian representation. The Gaussian center corresponds to a dominant arrival direction, while the eigenstructure of $\mathbf{H}_m(\mathbf{s})$ determines the local anisotropic angular profile: its eigenvectors specify the principal directions and its eigenvalues determine the corresponding concentration. These quantities naturally correspond to the orientation and anisotropic scales of a spherical Gaussian primitive, while $\alpha_m(\mathbf{s})$ represents its complex magnitude and phase contribution. The learned primitives need not correspond one-to-one with physical propagation paths; their superposition instead provides a compact representation of the complete directional spectrum.

Remark 2: Deformation interpretation. The quantities $\mathbf{d}_m(\mathbf{s})$, $\mathbf{H}_m(\mathbf{s})$, and $\alpha_m(\mathbf{s})$ vary with the TX–RX configuration. Changes in $\mathbf{d}_m(\mathbf{s})$ shift the localized response over the sphere, changes in $\mathbf{H}_m(\mathbf{s})$ modify its anisotropic angular extent and orientation, and changes in $\alpha_m(\mathbf{s})$ modify its complex magnitude and phase. These variations correspond, respectively, to the center, scale/orientation, and complex-attribute deformation modeled by SphGaS.

C Error Analysis

C.1 Geometric Analysis of Primitive Distortion

We analyze how different Gaussian formulations induce different response profiles when represented on the spherical RF domain. The formulations considered in our comparison can be grouped into three representative cases: (1) *direct ERP parameterization*, where the Gaussian is defined directly on the azimuth–elevation plane, as in SwiftWRF Liu et al. (2026); (2) *panoramic projection*, where a 3D Gaussian is projected onto the ERP plane through a panoramic camera model, as in WRF-GS Wen et al. (2025), WRF-GS+ Wen et al. (2026), and GAI-GS Shen et al. (2026); and (3) *chord-distance evaluation*, where the primitive remains on the sphere but its response is evaluated using the ambient Euclidean distance, as in GSRF Yang et al. (2025). In contrast, SphGaS evaluates the

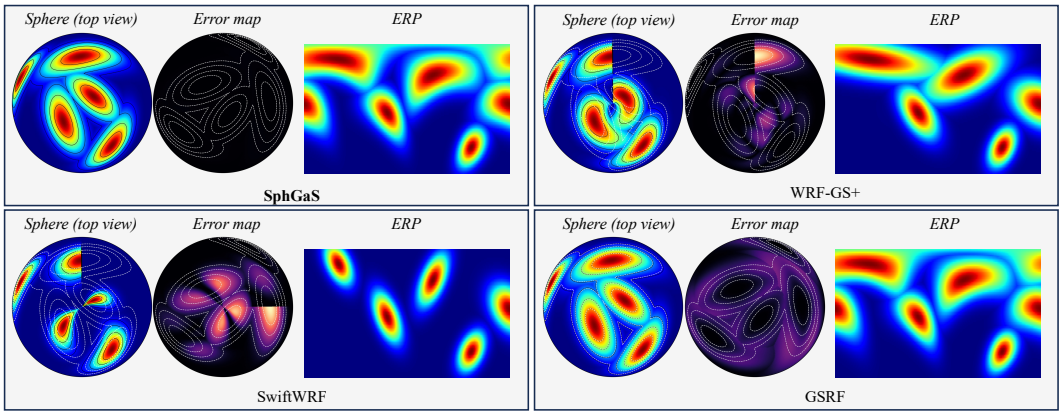

Figure 8: Visualization of geometric distortion induced by different RF spectrum representations. For each method, we show the rendered response on the sphere (top view), the corresponding error map with respect to the intrinsic SphGaS response, and its ERP visualization.

Table 6: Geometric comparison of representative RF Gaussian formulations.

Method	Formulation	Evaluation Domain	Planar Domain	Azimuth Seam	Intrinsic Spherical Geometry
SwiftWRF	Direct ERP	ERP plane	✓	✓	✗
WRF-GS / WRF-GS+ / GAI-GS	Panoramic projection	ERP plane	✓	✓	✗
GSRF	Chord-distance	Sphere	✗	✗	✗
SphGaS	**Intrinsic spherical**	**Sphere**	✗	✗	✓

Gaussian directly using the intrinsic geodesic geometry of the sphere. We derive below the response induced by each formulation under a common primitive configuration.

To isolate the geometric effect, we consider the same primitive center

$$\mathbf{n} = \mathbf{d}(\varphi_0, \lambda_0), \tag{30}$$

the same query direction $\mathbf{d}(\varphi, \lambda) \in \mathcal{M}$, and the same isotropic scale σ for all formulations. The anisotropic case follows analogously along the two principal directions.

For SphGaS, the response is evaluated intrinsically through the logarithmic map:

$$G_{\text{geo}}(\varphi, \lambda) = \exp\left(-\frac{\|\text{Log}_{\mathbf{n}}\left(\mathbf{d}(\varphi, \lambda)\right)\|_2^2}{2\sigma^2}\right). \tag{31}$$

Since the norm of the logarithmic map equals the geodesic distance on $\mathbb{S}^2$, this response follows the intrinsic spherical geometry.

For *1) direct ERP parameterization*, the Gaussian is instead evaluated using Euclidean distance in the azimuth–elevation coordinates:

$$G_{\text{ERP}}(\varphi, \lambda) = \exp\left(-\frac{(\varphi - \varphi_0)^2 + (\lambda - \lambda_0)^2}{2\sigma^2}\right). \tag{32}$$

This replaces the spherical metric with a uniform Euclidean metric in ERP coordinates and introduces the $0°/360°$ azimuth boundary.

For *2) panoramic projection*, we isolate the angular deformation introduced by locally projecting an isotropic Gaussian onto the ERP plane. Around elevation λ_0, the local spherical metric is

$$\mathrm{d}s^2 = \cos^2 \lambda_0 \, \mathrm{d}\varphi^2 + \mathrm{d}\lambda^2. \tag{33}$$

Accordingly, the corresponding local ERP covariance is

$$\mathbf{\Sigma}_{\text{proj}} = \sigma^2 \begin{bmatrix} 1/\cos^2 \lambda_0 & 0 \\ 0 & 1 \end{bmatrix}, \tag{34}$$

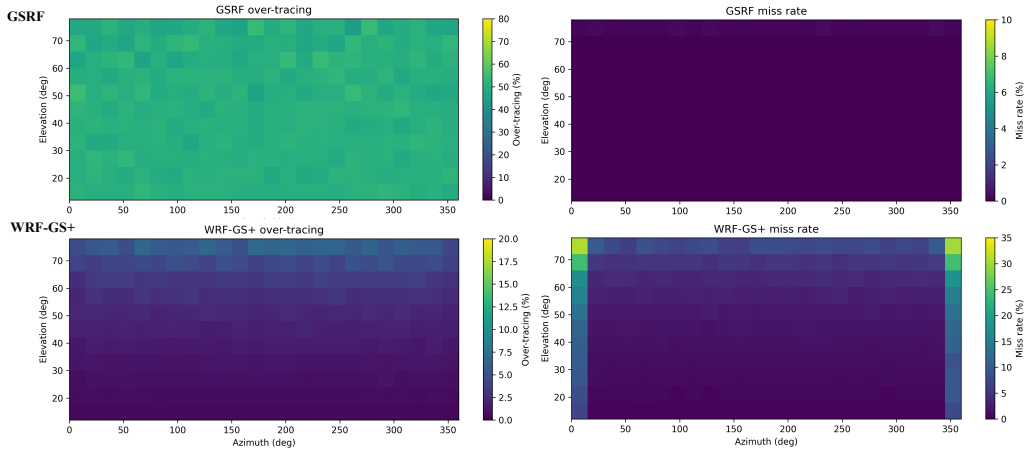

Figure 9: Analysis of tile assignment induced by different Gaussian support formulations. Left: support recall, measuring the fraction of the intrinsic spherical 3σ support covered by the selected region. Right: over-coverage, measuring the fraction of the selected region lying outside the intrinsic support. WRF-GS+ produces a relatively tight support with slight under-coverage, while GSRF improves recall through a more conservative expansion at the cost of substantially higher over-coverage.

which gives the locally projected response

$$G_{\text{proj}}(\varphi, \lambda) = \exp\left(-\frac{\cos^2 \lambda_0(\varphi - \varphi_0)^2 + (\lambda - \lambda_0)^2}{2\sigma^2}\right). \tag{35}$$

This expression isolates the local angular component of the panoramic projection; the complete projected response of a 3D Gaussian additionally depends on its 3D covariance and projection geometry.

For *3) chord-distance evaluation*, both the primitive center and the query remain on the sphere, but their response is determined by the ambient Euclidean distance in $\mathbb{R}^3$:

$$G_{\text{chord}}(\varphi, \lambda) = \exp\left(-\frac{\|\mathbf{d}(\varphi, \lambda) - \mathbf{n}\|_2^2}{2\sigma^2}\right). \tag{36}$$

For an angular separation $\theta = \arccos(\mathbf{n}^\top \mathbf{d})$, the chord distance satisfies

$$\|\mathbf{d} - \mathbf{n}\|_2 = 2\sin\frac{\theta}{2} < \theta = \|\text{Log}_{\mathbf{n}}(\mathbf{d})\|_2, \qquad \theta \in (0, \pi), \tag{37}$$

showing that chord-distance evaluation systematically underestimates the intrinsic spherical distance away from the primitive center.

The resulting response differences are visualized in Fig. 8.

C.2 Analysis of Tile Assignment Accuracy

During splatting, each Gaussian primitive is assigned to a set of image tiles according to its estimated support and is evaluated only within these tiles. Therefore, any mismatch between the estimated support and the true spherical support directly affects the tile list. If the support is underestimated, some tiles containing valid Gaussian responses are missed. If the support is overestimated, the primitive is unnecessarily assigned to additional tiles, increasing redundant evaluation.

To quantify this effect, we use the intrinsic spherical 3σ support of each primitive as the reference. The reference support is mapped to ERP space and compared with the support used for tile assignment by WRF-GS+ and GSRF. We report *support recall*, which measures how much of the reference support is covered by the selected region, and *over-coverage*, which measures how much of the selected region lies outside the reference support.

WRF-GS+ achieves 96.18% support recall with 2.11% over-coverage, indicating a relatively tight support estimate but with a small portion of valid regions missed. GSRF improves the recall to 99.96% through elevation-dependent compensation and periodic boundary handling, but its more conservative support also includes substantially more unnecessary regions, resulting in 51.92% over-coverage. In contrast, SphGaS directly preserves the intrinsic spherical support, yielding 100% recall and zero geometry-induced over-coverage.

D Supplementary Experiments

D.1 Spherical vs. Euclidean Gaussian Representations

To isolate the effect of representation geometry, we compare spherical Gaussian primitives with conventional 2D and 3D Euclidean Gaussian primitives in a controlled single-spectrum fitting experiment. Five RF spectra are independently selected for repeated evaluation. For each target spectrum,

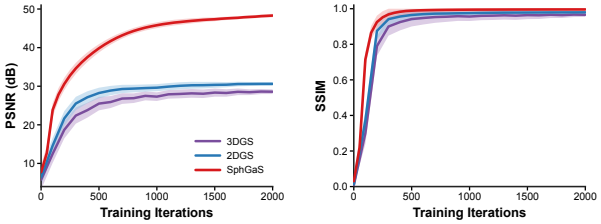

Figure 10: Spectrum fitting with 3D, 2D, and spherical Gaussian representations. The spherical representation converges faster and achieves higher fitting accuracy.

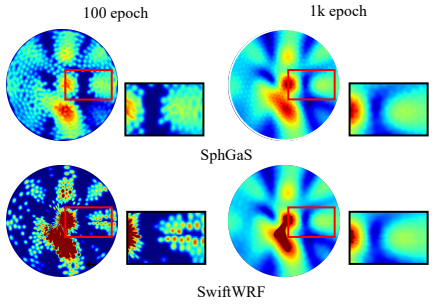

Figure 11: Comparison of Gaussian responses under planar and spherical geometry.

all representations use the same budget of 1,000 Gaussian primitives, and only the primitive parameters are optimized, without any transceiver-conditioned deformation. The spherical representation is rasterized directly on the spherical domain, while 2DGS operates on a planar parameterization and 3DGS follows the standard splatting-and-rasterization pipeline.

As shown in Fig. 10, the spherical representation converges faster and achieves substantially higher fitting accuracy. Across the evaluated spectra, 3DGS and 2DGS saturate at approximately 28.5 dB and 30.6 dB PSNR, respectively, whereas the spherical representation reaches nearly 48 dB. Fig. 11 further illustrates the geometric difference between the representations: planar Gaussian responses become distorted when mapped to the spherical observation domain, whereas spherical primitives retain their localized structure by operating directly on the sphere. These results demonstrate that spherical parameterization provides a more suitable representation for directional RF spectra under the same primitive budget.

D.2 Primitive Count and Pruning

We analyze the effect of the number of spherical Gaussian primitives and the robustness of SphGaS to primitive pruning. Pruning is performed using the criterion introduced in Fan et al. (2024). As shown in Fig. 12, increasing the number of primitives consistently improves reconstruction quality, with the gains gradually saturating beyond a sufficient primitive budget. SphGaS also retains most of its reconstruction quality after pruning, indicating that redundant primitives can be effectively removed to obtain a more compact representation.

D.3 Novel Spectrum vs. Novel Field Synthesis

TX and RX movements affect spectrum synthesis in different ways. For a fixed TX, changing the RX location queries different spectra within an observed propagation field, whereas moving the TX changes the propagation field itself. We therefore separately evaluate generalization to unseen

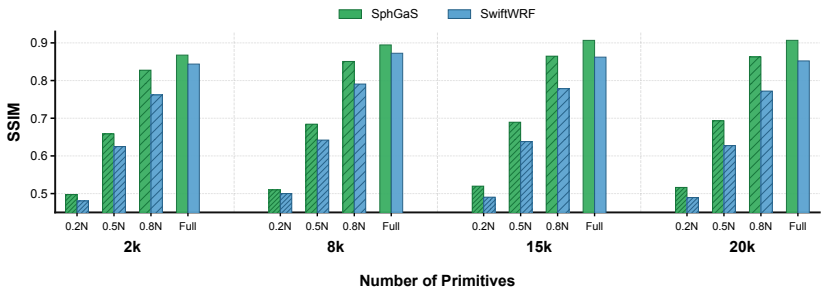

Figure 12: Effect of primitive count and pruning on reconstruction quality.

Table 7: Generalization performance for novel spectrum synthesis under seen TXs and novel field synthesis under unseen TXs. Higher is better for PSNR and SSIM, while lower is better for MAE. For PSNR, values in parentheses denote the absolute change in dB from seen to unseen TX; for MAE and SSIM, they denote relative changes. Best results are in bold and second-best results are underlined.

Scene	Method	Seen TX			Unseen TX		
		MAE ↓	PSNR (dB) ↑	SSIM ↑	MAE ↓	PSNR (dB) ↑	SSIM ↑
Bedroom	NeRF^2	0.0676	21.19	0.6228	0.0699 (+3.4%)	20.84 (-0.35 dB)	0.6125 (-1.7%)
	WRF-GS+	0.0676	20.81	0.6749	0.0685 (+1.3%)	20.63 (-0.18 dB)	0.6683 (-1.0%)
	SwiftWRF	0.0417	26.58	0.8691	0.0587 (+40.8%)	24.11 (-2.47 dB)	0.8415 (-3.2%)
	SphGaS	**0.0360**	**27.84**	**0.8863**	**0.0414** (+15.0%)	**26.93** (-0.91 dB)	**0.8729** (-1.5%)
Conf.Room	NeRF^2	0.0688	21.27	0.6097	0.0710 (+3.2%)	21.04 (-0.23 dB)	0.6077 (-0.3%)
	WRF-GS+	0.0648	21.37	0.7032	0.0655 (+1.1%)	21.32 (-0.05 dB)	0.7002 (-0.4%)
	SwiftWRF	0.0484	26.07	0.8623	0.0625 (+29.1%)	23.76 (-2.31 dB)	0.8293 (-3.8%)
	SphGaS	**0.0393**	**27.10**	**0.8800**	**0.0425** (+8.1%)	**26.75** (-0.35 dB)	**0.8745** (-0.6%)
Office	NeRF^2	0.0800	19.74	0.5563	0.0882 (+10.1%)	18.79 (-0.95 dB)	0.5308 (-4.6%)
	WRF-GS+	0.0705	20.47	0.6670	0.0765 (+8.5%)	19.72 (-0.75 dB)	0.6440 (-3.5%)
	SwiftWRF	0.0499	25.36	0.8511	0.0673 (+34.9%)	23.51 (-1.85 dB)	0.8264 (-2.9%)
	SphGaS	**0.0434**	**26.53**	**0.8684**	**0.0563** (+29.7%)	**24.99** (-1.54 dB)	**0.8529** (-1.8%)
Average	NeRF^2	0.0721	20.73	0.5963	0.0764 (+5.9%)	20.22 (-0.51 dB)	0.5837 (-2.1%)
	WRF-GS+	0.0676	20.88	0.6817	0.0702 (+3.7%)	20.56 (-0.32 dB)	0.6708 (-1.6%)
	SwiftWRF	0.0467	26.00	0.8608	0.0628 (+34.6%)	23.79 (-2.21 dB)	0.8324 (-3.3%)
	SphGaS	**0.0396**	**27.16**	**0.8782**	**0.0467** (+18.1%)	**26.22** (-0.94 dB)	**0.8668** (-1.3%)

RX configurations under seen TXs and to entirely unseen TXs. The former corresponds to *novel spectrum synthesis* within an observed field, while the latter provides a stronger test of *novel field synthesis*.

For each indoor scene, the first 90 TX locations are included in training. For each of these TXs, 80% of the associated RX links are used for optimization and the remaining 20% form the seen-TX test split. The final 10 TX locations, together with all of their associated RX links, are completely excluded from training and form the unseen-TX test split.

As shown in Tab.7, SphGaS achieves the best reconstruction performance in both settings. The transition from seen to unseen TXs also reveals the increased difficulty of novel field synthesis. Averaged across the three indoor scenes, SphGaS decreases from 27.16 to 26.22 dB PSNR, corresponding to a 0.94 dB drop, while SwiftWRF drops by 2.21 dB. Its SSIM decreases by only 1.3%, compared with 3.3% for SwiftWRF. Although NeRF^2 and WRF-GS+ exhibit smaller absolute changes between the two splits, their reconstruction accuracy remains substantially lower in both settings. Overall, SphGaS maintains strong reconstruction fidelity when generalizing from novel spectra within observed fields to spectra associated with previously unseen propagation fields.

Table 8: Effect of bidirectional training under the standard data construction. PSNR gains are shown as absolute dB differences, while SSIM gains are shown as relative percentages.

Scene	Variant	PSNR (dB) ↑	SSIM ↑
Conf.Room	NeRF^2 w/o bidi.	19.45	0.6459
	NeRF^2	19.57 (+0.12 dB)	0.6579 (+1.86%)
	SwiftWRF w/o bidi.	22.97	0.7689
	SwiftWRF	24.31 (+1.34 dB)	0.8187 (+6.48%)
	SphGaS w/o bidi.	24.02	0.8021
	SphGaS	**25.33** (+1.31 dB)	**0.8524** (+6.27%)
Bedroom	NeRF^2 w/o bidi.	19.32	0.6578
	NeRF^2	19.65 (+0.33 dB)	0.6608 (+0.46%)
	SwiftWRF w/o bidi.	23.63	0.7853
	SwiftWRF	25.03 (+1.40 dB)	0.8386 (+6.79%)
	SphGaS w/o bidi.	24.52	0.8068
	SphGaS	**25.91** (+1.39 dB)	**0.8561** (+6.11%)
Office	NeRF^2 w/o bidi.	16.83	0.4951
	NeRF^2	17.39 (+0.56 dB)	0.5447 (+10.02%)
	SwiftWRF w/o bidi.	20.60	0.6996
	SwiftWRF	21.91 (+1.31 dB)	0.7481 (+6.93%)
	SphGaS w/o bidi.	21.48	0.7317
	SphGaS	**22.84** (+1.36 dB)	**0.7815** (+6.81%)

Table 9: Sample-budget-controlled comparison of random and bidirectional training for SphGaS. Both settings use 4.8K directed training spectra and are evaluated on the same 600 bidirectional test pairs.

Scene	Variant	PSNR (dB) ↑	SSIM ↑
Conf.Room	SphGaS (Random)	22.74	0.7883
	SphGaS (Bidi.)	**23.58** (+0.84 dB)	**0.8061** (+2.26%)
Bedroom	SphGaS (Random)	22.45	0.7920
	SphGaS (Bidi.)	**23.76** (+1.31 dB)	**0.8102** (+2.30%)
Office	SphGaS (Random)	19.64	0.7038
	SphGaS (Bidi.)	**20.47** (+0.83 dB)	**0.7366** (+4.66%)

D.4 Ablation Study

Effect of Bidirectional Training. Electromagnetic reciprocity implies a structural relationship between the two directions of the same physical TX–RX link. We exploit this structure at the data level by organizing reversed links as bidirectional pairs during training, and study whether such pairing benefits spectrum reconstruction. We evaluate bidirectional training under two settings. The first follows our standard data construction and measures the overall gain from incorporating the reversed link. The second controls the number of training spectra to isolate the effect of reciprocal pairing from the effect of additional training data.

Setting 1: Standard bidirectional training. We first evaluate the overall effect of incorporating reversed-link spectra under the standard data construction. As shown in Tab. 8, bidirectional training consistently improves reconstruction across all three methods and scenes. The gains are particularly clear for the Gaussian-based methods, with SwiftWRF and SphGaS improving by approximately 1.3–1.4 dB in PSNR and 6–7% in SSIM. These results indicate that reversed-link spectra provide useful complementary supervision for modeling transceiver-dependent spectrum variation.

Setting 2: Sample-budget-controlled training. We further control the number of training spectra to distinguish the effect of reciprocal pairing from simply using more observations. Both variants use 4.8K directed spectra: the *random* variant samples them independently, whereas the *bidirectional* variant organizes them into reciprocal pairs; both are evaluated on the same 600 held-out bidirectional test pairs. As shown in Tab. 9, bidirectional pairing consistently improves PSNR by 0.83–1.31 dB and SSIM by 2.26–4.66% across the three scenes. These results confirm that the benefit arises from reciprocal pairing rather than merely from an increased number of training observations.

Ablation of Primitive Modeling Space. We further study how the geometric modeling space affects RF spectrum reconstruction. The comparison is organized at two levels. First, we compare SphGaS against deformable Gaussian representations defined in Euclidean spaces: WRF-GS+ models prim-

Table 10: Modeling choices of the compared Gaussian representations.

Variant	Primitive Space	Response Evaluation	Center Deformation
WRF-GS+	3D Euclidean	Euclidean	3D Euclidean
SwiftWRF	2D azimuth–elevation	Planar Euclidean	2D coordinate offsets
Spherical GS w/o Log map & geodesic deform.	Sphere	Local Euclidean approximation	Azimuth–elevation offsets
Spherical GS w/o Log map	Sphere	Local Euclidean approximation	Geodesic
Spherical GS w/o geodesic deform.	Sphere	Log-map based	Azimuth–elevation offsets
SphGaS	**Sphere**	**Log-map based**	**Geodesic**

Table 11: Comparison of Euclidean and spherical Gaussian modeling. PSNR drops are shown as absolute dB differences from full SphGaS, while SSIM drops are shown as relative percentages.

Dataset	Modeling Variant	PSNR (dB)↑	SSIM↑
RFID	3D Euclidean deformable GS (WRF-GS+)	25.23 (-1.49 dB)	0.8885 (-4.67%)
	2D Euclidean deformable GS (SwiftWRF)	24.75 (-1.97 dB)	0.8714 (-6.50%)
	Spherical GS w/o Log map & geodesic deform.	25.75 (-0.97 dB)	0.9138 (-1.95%)
	Spherical GS w/o Log map	26.15 (-0.57 dB)	0.9218 (-1.09%)
	Spherical GS w/o geodesic deform.	26.36 (-0.36 dB)	0.9261 (-0.63%)
	SphGaS (full manifold modeling)	**26.72**	**0.9320**
City	3D Euclidean deformable GS (WRF-GS+)	20.32 (-2.60 dB)	0.6661 (-13.17%)
	2D Euclidean deformable GS (SwiftWRF)	21.13 (-1.79 dB)	0.6706 (-12.58%)
	Spherical GS w/o Log map & geodesic deform.	22.12 (-0.80 dB)	0.7534 (-1.79%)
	Spherical GS w/o Log map	22.45 (-0.47 dB)	0.7591 (-1.04%)
	Spherical GS w/o geodesic deform.	22.64 (-0.28 dB)	0.7627 (-0.58%)
	SphGaS (full manifold modeling)	**22.92**	**0.7671**
Indoor	3D Euclidean deformable GS (WRF-GS+)	18.34 (-6.35 dB)	0.6274 (-24.41%)
	2D Euclidean deformable GS (SwiftWRF)	23.75 (-0.94 dB)	0.8018 (-3.40%)
	Spherical GS w/o Log map & geodesic deform.	23.82 (-0.87 dB)	0.8148 (-1.83%)
	Spherical GS w/o Log map	24.18 (-0.51 dB)	0.8213 (-1.05%)
	Spherical GS w/o geodesic deform.	24.38 (-0.31 dB)	0.8251 (-0.59%)
	SphGaS (full manifold modeling)	**24.69**	**0.8300**

itives in 3D Euclidean space, while SwiftWRF operates on a 2D azimuth–elevation plane. Second, within the spherical representation, we separately ablate the two intrinsic operations used by SphGaS: spherical response evaluation through the logarithmic map and geodesic center deformation through tangent-space displacement followed by the exponential map.

Tab. 10 summarizes the modeling choices of all compared variants. For *w/o Log map*, primitive centers remain on $\mathbb{S}^2$, but Gaussian responses are evaluated using a local Euclidean approximation rather than the Riemannian logarithmic map. For *w/o geodesic deform.*, spherical primitive centers are directly updated through azimuth–elevation offsets instead of tangent-space displacement followed by the exponential map. The *w/o Log map & geodesic deform.* variant removes both intrinsic operations while retaining the spherical primitive parameterization.

As shown in Tab. 11, moving from fully Euclidean representations to spherical primitive modeling already produces a substantial improvement. The complete Euclidean baselines exhibit considerably larger gaps from SphGaS, particularly on the City and Indoor datasets, where geometric mismatch becomes more pronounced. The controlled spherical ablations further show that both intrinsic response evaluation and geodesic deformation contribute consistently. Replacing the Log-map evaluation with a local Euclidean approximation reduces PSNR by 0.47–0.57 dB, while replacing geodesic deformation with direct azimuth–elevation offsets causes a 0.28–0.36 dB drop. Removing both operations further enlarges the degradation to 0.80–0.97 dB. The same trend is observed in SSIM across all three datasets. These results indicate that spherical primitive parameterization provides the main geometric advantage, while intrinsic response evaluation and geodesic deformation offer additional and complementary improvements.

D.5 Visualization

We provide additional visualizations of spherical We provide additional visualizations of spherical primitive deformation (Fig. 13), the effects of individual deformation components (Fig. 14), phase reconstruction (Fig. 15), and spectrum reconstruction on the real-world RFID dataset (Fig. 16).

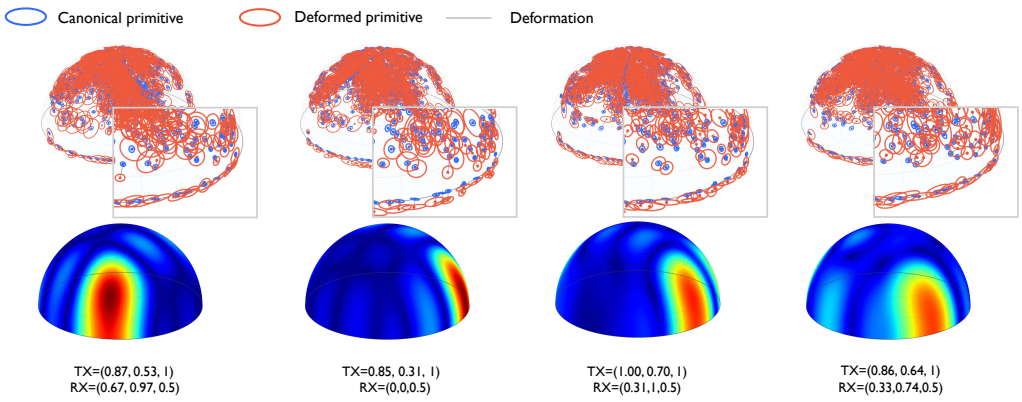

Figure 13: Visualization of primitive deformation on the spectrum space.

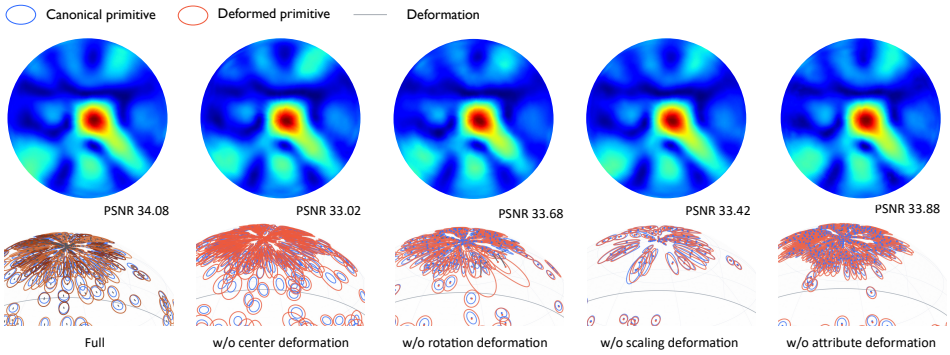

Figure 14: Visualization of different deformation components. Blue and red contours denote canonical and deformed primitives, respectively.

E Applications

E.1 RSSI Reconstruction

RSSI summarizes the received power of a wireless link into a scalar quantity by aggregating contributions from different arrival directions. This aggregation is particularly natural for SphGaS, since the reconstructed RF spectrum is defined directly over the spherical directional domain. In contrast, methods such as SwiftWRF and WRF-GS+ represent or render directional responses in projected planar spaces, where directly averaging uniformly sampled pixels does not preserve their corresponding spherical areas and can introduce projection-dependent bias during aggregation.

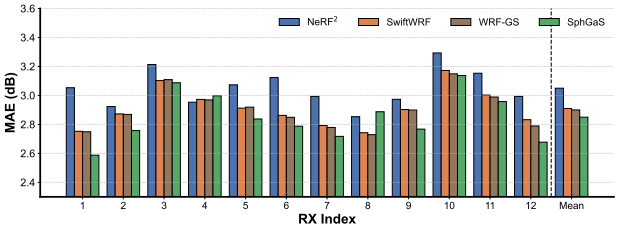

Figure 17: RSSI reconstruction error across receiver locations.

We therefore aggregate the reconstructed directional power directly on the reception sphere. Given the rendered complex spectrum $\hat{S}(\omega)$, we compute

$$\hat{P}_{\mathrm{rx}} = \frac{\sum_{j=1}^{M} |\hat{S}(\omega_j)|^2 \Delta\Omega_j}{\sum_{j=1}^{M} \Delta\Omega_j},$$

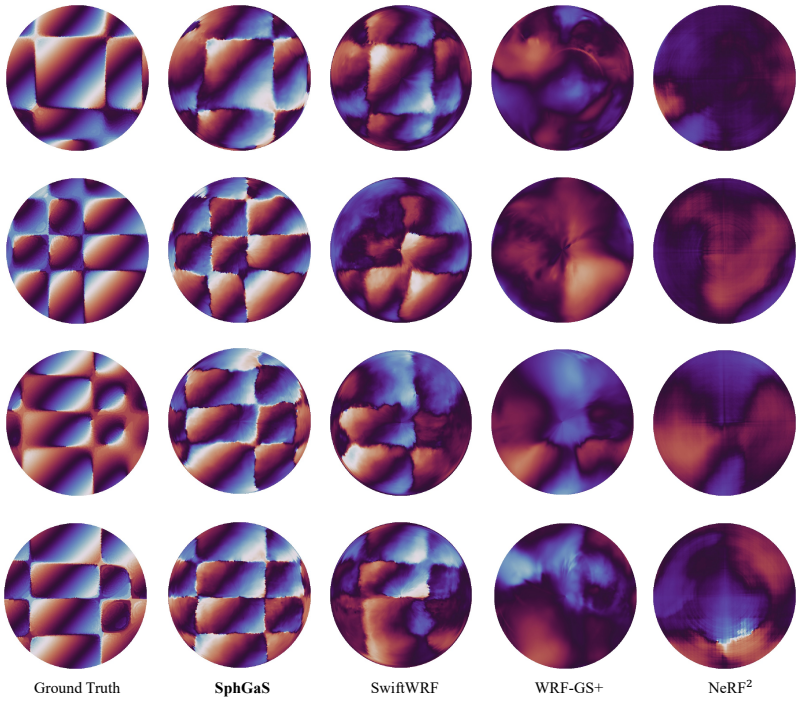

Figure 15: Qualitative comparison of RF spectrum phase reconstruction.

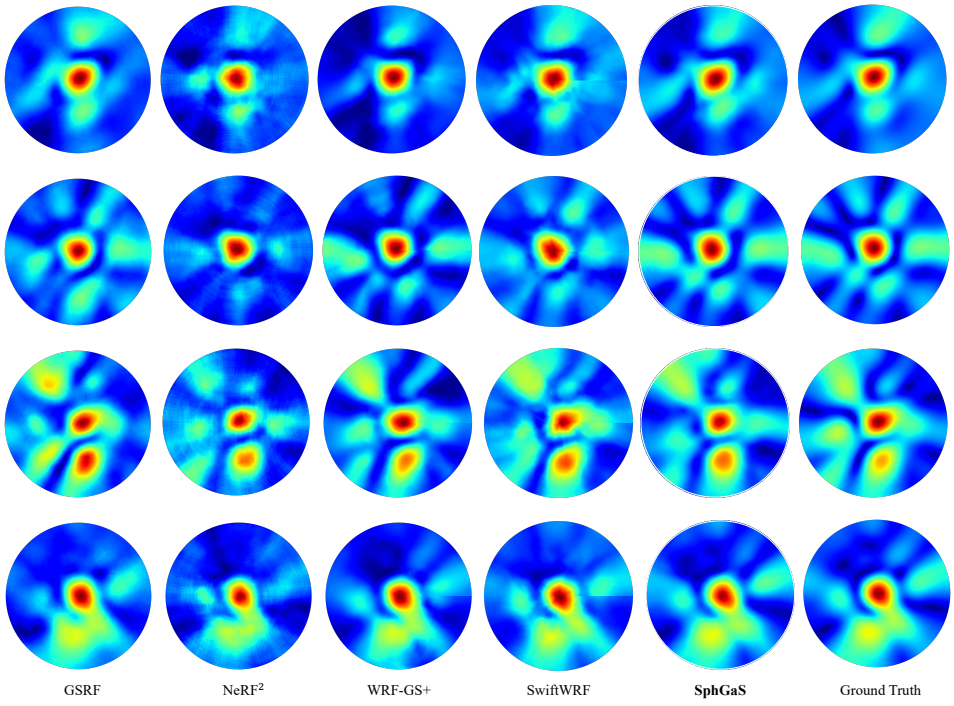

Figure 16: Visualization of spectrum reconstruction on the real-world RFID dataset.

where ω_j denotes an arrival direction and $\Delta\Omega_j$ is the solid angle represented by that sample. For uniformly sampled azimuth and elevation, the solid-angle weight is proportional to the cosine of elevation, accounting for the unequal spherical area represented by different angular bins. Since the

Table 12: Quantitative evaluation of RSSI reconstruction and downstream TX deployment. Lower is better for reconstruction errors, while higher is better for Top-8 accuracy. Best results are in bold and second-best results are underlined.

(a) Real-world RSSI reconstruction

Method	MAE in dB ↓
MRI	6.74 ± 0.27
NeRF^2	3.05 ± 0.18
SwiftWRF	2.91 ± 0.15
WRF-GS+	$\underline{2.90 \pm 0.16}$
SphGaS	$\mathbf{2.85 \pm 0.15}$

(b) RadioMapSeer RSSI reconstruction and TX deployment accuracy. Strength is normalized.

Method	RMSE ↓	Top-8 Acc. (%) ↑
NeRF^2	0.0808 ± 0.0071	71.03
WRF-GS+	0.0736 ± 0.0063	<u>74.87</u>
SwiftWRF	$\underline{0.0621 \pm 0.0038}$	<u>74.87</u>
SphGaS	$\mathbf{0.0558 \pm 0.0036}$	**83.33**

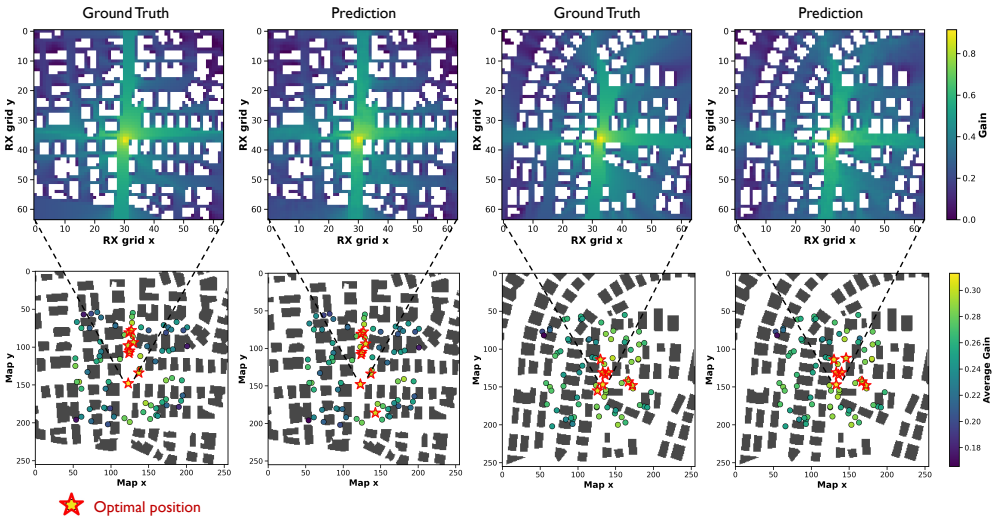

Figure 18: TX deployment selection using reconstructed RSSI fields. Predicted coverage maps are compared with the ground truth for representative TX candidates.

output is a single scalar rather than a directional spectrum, we use a coarse angular resolution of $36^\circ \times 9^\circ$ for efficient aggregation.

We evaluate RSSI reconstruction on both real-world RF measurements and the public RadioMapSeer benchmark. For BLE, we follow the original evaluation setting. For RadioMapSeer, we randomly select ten layouts, each with 80 candidate TX locations and a uniform 64×64 RX grid. We use 80% of the TX locations, together with all associated RX samples, for training, and reserve the remaining 20% of TX locations and their RX samples as unseen configurations.

RSSI reconstruction. As shown in Tab. 12, SphGaS achieves the lowest RSSI reconstruction error on both benchmarks, reaching an MAE of 2.85 dB on the real-world BLE measurements and an RMSE of 0.0558 on RadioMapSeer. Fig. 17 further shows consistently low reconstruction error across receiver locations. These results indicate that the reconstructed directional spectra preserve the received-power structure required for reliable spatial RSSI estimation. The reconstructed RSSI fields further enable communication-aware decisions under varying transceiver configurations. We consider two representative applications: transmitter deployment and communication-aware robot planning.

TX deployment. We evaluate whether the reconstructed RF field can support TX placement over the complete candidate set, where both observed and previously unobserved TX locations are valid deployment choices. For each candidate TX, we compute the average predicted received power over the accessible RX region as its coverage score and rank all candidates jointly. Top-8 accuracy measures the overlap between the predicted and ground-truth top-8 locations. Thus, this

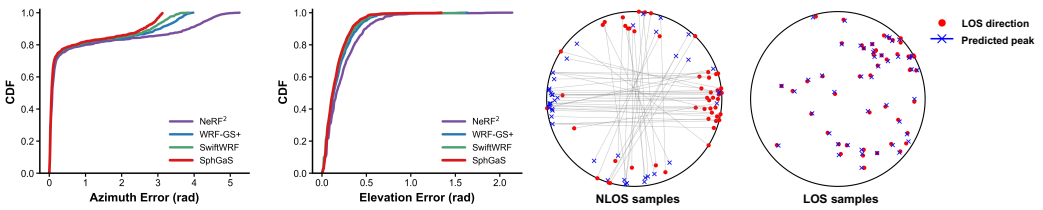

Figure 21: Directional consistency between reconstructed spectral responses and the geometric direct-path direction, with representative LOS and NLOS examples.

task assesses deployment over the full search space rather than unseen-TX generalization alone. As shown in Tab. 12, SphGaS achieves the highest Top-8 accuracy of 83.33%. Fig. 18 further shows that the reconstructed coverage patterns support consistent ranking of candidate TX locations.

Communication-aware robot planning. The reconstructed RSSI field can also serve as a communication prior for robot navigation. Combining predicted signal strength with geometric path cost discourages trajectories through weak-coverage regions. As illustrated in Fig. 19, the RSSI-aware trajectory takes a small detour from the geometry-only shortest path to maintain stronger connectivity, demonstrating the potential of SphGaS for communication-aware navigation.

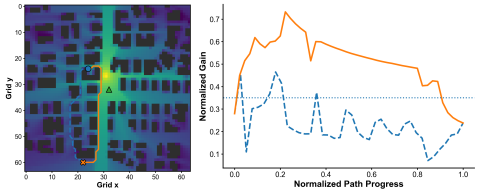

Figure 19: An illustrative example of communication-aware robot planning using the reconstructed RSSI field.

E.2 Directional Propagation Analysis

Beyond received signal strength, the reconstructed spectra also preserve directional propagation information. We evaluate this property by examining whether the direct-path direction can be recovered accurately and whether LOS and NLOS links can be distinguished from the reconstructed directional responses.

Table 13: Directional propagation evaluation on indoor scenes. Higher is better for LOS accuracy; lower is better for angular error.

Method	LOS Acc. ↑	Azimuth Err. ↓	Elevation Err. ↓
NeRF^2	0.86	0.7848	0.2277
SwiftWRF	<u>0.91</u>	<u>0.5760</u>	<u>0.1682</u>
WRF-GS+	0.88	0.6277	0.1822
SphGaS	**0.93**	**0.5237**	**0.1522**

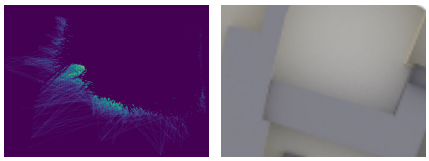

Figure 20: A simple example of layout reconstruction from aggregated directional propagation cues.

As shown in Tab. 13, SphGaS achieves the highest LOS prediction accuracy and the lowest azimuth and elevation errors, indicating that its reconstructed spectra better preserve directional propagation cues.

Fig. 21 further shows that LOS links typically retain a strong response near the direct-path direction, while NLOS links exhibit weaker or displaced responses due to indirect propagation. As a simple example, directional cues from multiple transceiver configurations can be aggregated to infer coarse scene structure, as illustrated in Fig. 20. This result highlights that the reconstructed spectra preserve propagation information beyond signal strength alone.

E.3 Extension to Visual Applications

The hemispherical observation domain used in our RF experiments reflects the field of view of the considered RX setup and does not constitute a limitation of the SphGaS representation itself. Since



Figure 22: Full-sphere visual reconstruction with SphGaS. Using 20K spherical Gaussian primitives, SphGaS reconstructs 832×416 360° images with an S-SSIM of 0.99, demonstrating that the representation naturally extends from hemispherical RF observations to complete spherical signals.

the primitives and rasterization are defined intrinsically on $\mathbb{S}^2$, the same formulation can be extended directly to full-sphere signals. As a proof of concept, we apply SphGaS to a 832×416 equirectangular 360° image and represent the complete sphere using 20K spherical Gaussian primitives. The reconstruction reaches an S-SSIM of 0.99, as illustrated in Fig. 22. This result shows that SphGaS can serve as a geometry-consistent spherical representation beyond RF sensing, providing a potential substrate for future 360° image and video modeling.

F Implementation Details

F.1 Dataset Information

NeRF^2 dataset. We use the public NeRF^2 dataset for real-world RF evaluation. The RFID subset contains measured PAS collected in a laboratory environment, where the RX is fixed and the TX is moved across different locations. We use this subset for real-world spectrum reconstruction. The BLE subset provides received-power measurements under different transceiver configurations and is used for RSSI-related evaluation. Following the original dataset convention, each RFID PAS is represented on a 90×360 (elevation $\times$ azimuth) angular grid.

Ray-traced spectrum datasets. We further construct six ray-traced spectrum datasets cover-

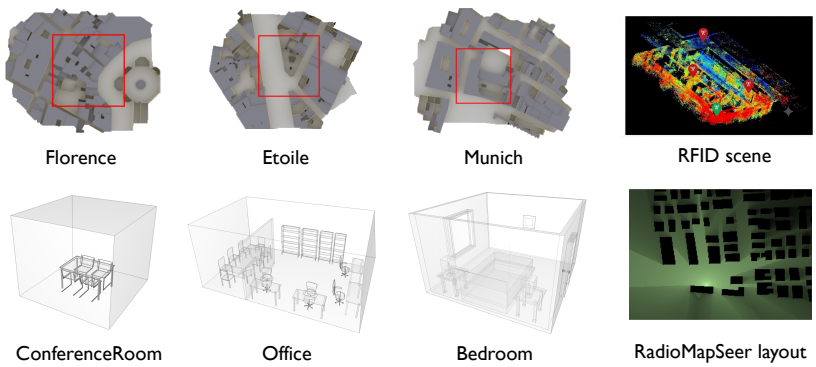

Figure 23: Overview of the evaluated scenes and datasets.

ing both outdoor and indoor environments, with all simulations conducted at a carrier frequency

of 2.4,GHz. For outdoor evaluation, we use realistic urban layouts from *Munich*, *Etoile*, and *Florence*, and simulate RF propagation with NVIDIA Sionna RT. The resulting spectra are represented as power angular spectra (PAS) and do not include phase information. For indoor evaluation, we construct synthetic *Conf.Room*, *Bedroom*, and *Office* scenes and simulate propagation using the MATLAB ray-tracing toolbox. The resulting spectra are complex-valued, with both amplitude and phase responses retained. On average, each spectrum contains approximately 183 traced propagation paths. Each scene contains 6,000 spatial spectra, with 80represented on a 90×360 (elevation $\times$ azimuth) angular grid. To further evaluate unseen-TX generalization, we additionally construct a denser 100-TX $\times$ 100-RX dataset for each indoor scene. For the first 90 TX locations, 80% of the corresponding RX links are used for training and the remaining 20% form the *Seen TX / Held-out RX* test split. The final 10 TX locations are completely excluded from training, and all of their RX links form the *Unseen TX / All RX* test split.

RSSI Datasets. We evaluate RSSI reconstruction on two datasets. The real-world BLE subset of NeRF^2 S23 provides measured received-power values under varying TX–RX configurations. RadioMapSeer provides simulated radio maps over diverse building layouts, where each TX location is associated with a dense 64×64 RX grid and corresponding pathloss/radio-strength values. For RadioMapSeer, we randomly select 10 layouts and use 80 TX locations per layout. The resulting dense maps are used to evaluate RSSI-field reconstruction at unseen configurations and downstream TX deployment.

Data Preprocessing. Before training and evaluation, all RF spectrum amplitudes and RSSI values are normalized to $[0, 1]$. The spatial coordinates of the transceiver locations are also normalized according to the corresponding scene extent. The same preprocessing is applied consistently to the training and test sets.

F.2 Baseline Implementation Details.

① *SwiftWRF.* We use $10{,}000$ primitives and an 8-layer MLP with hidden dimension 256, with the same positional encoding frequencies as SphGaS. Since the original SwiftWRF does not condition on both transceiver positions, we augment its deformation-network input with positional encodings of both TX and RX coordinates, providing a Euclidean planar counterpart under the same transceiver-conditioning setting. ② *WRF-GS+.* We follow the original implementation and settings. Although the original paper evaluates WRF-GS+ only under single-endpoint mobility, its formulation can in principle support joint TX–RX mobility through TX-dependent deformation and RX-dependent camera motion under the Mercator projection. ③ *Other baselines.* All remaining baselines follow their original implementations and settings. Since the code for GAI-GS is not publicly available, we report results from the original paper. GSRF takes only the RX location as positional input and is therefore evaluated in its original setting.

F.3 SphGaS Implementation

Primitive initialization and deformation constraints. Canonical primitive centers are initialized approximately uniformly over the upper hemisphere. For each canonical center $\mathbf{n}_i^0 \in \mathbb{S}^2$, we construct a deterministic orthonormal basis $\mathbf{B}_i^0 = [\mathbf{b}_{1,i}^0, \mathbf{b}_{2,i}^0]$ of its tangent space. Specifically, we use the global z-axis as the reference direction and switch to the x-axis when $\mathbf{n}_i^0$ is nearly aligned with the z-axis:

$$\mathbf{e}_i^{\mathrm{ref}} = \begin{cases} \mathbf{e}_z, & |(\mathbf{n}_i^0)_z| \le 0.95, \\ \mathbf{e}_x, & |(\mathbf{n}_i^0)_z| > 0.95. \end{cases} \tag{38}$$

The tangent basis is then constructed through normalized cross products:

$$\mathbf{b}_{1,i}^0 = \frac{\mathbf{e}_i^{\mathrm{ref}} \times \mathbf{n}_i^0}{\left\| \mathbf{e}_i^{\mathrm{ref}} \times \mathbf{n}_i^0 \right\|_2}, \qquad \mathbf{b}_{2,i}^0 = \mathbf{n}_i^0 \times \mathbf{b}_{1,i}^0. \tag{39}$$

This reference-axis switching avoids numerical degeneracy when the primitive center is nearly aligned with the z-axis. These canonical frames are constructed only once; during deformation, they are transferred to the deformed primitive centers by parallel transport rather than recomputed from spherical coordinates. The two tangent-space scales are initialized to 0.018, the in-plane orientations are randomly sampled from $[0, 2\pi)$, and the minimum Gaussian scale is set to 10^{-4}. Primitive centers evolve on the underlying unit sphere $\mathbb{S}^2$ and are not clipped to the reception hemisphere $\mathcal{M}$,

allowing responses to move smoothly outside the observable angular domain. To stabilize optimization, the geodesic displacement from each canonical center is bounded to 20°, and the one-sided angular extent of each primitive is limited to at most 90°. All the experiments are conducted on a single NVIDIA RTX 4090 GPU.

Network and optimization details. The shared deformation network takes the canonical primitive direction and the two transceiver positions as input. Positional encoding uses 10 frequency levels for primitive directions and 6 levels for each transceiver position. The network uses an 8-layer MLP with hidden dimension 256 and a middle skip connection, followed by lightweight prediction heads for the geometric residuals $(\Delta\boldsymbol{\xi}, \Delta\boldsymbol{\ell}, \Delta\rho)$ and the complex-attribute residual $\Delta\mathbf{a}$. The prediction heads are initialized close to zero so that training starts near the canonical representation. In the canonical stage, the primitive centers and geometric parameters use a learning rate of 2×10^{-3}, while the primitive attributes use 5×10^{-3}. During the deformation stage, the deformation-network learning rate is exponentially decayed from 8×10^{-4} to 1.6×10^{-5}. The canonical-center learning rate decays from 2×10^{-4} to 2×10^{-5} and the centers are frozen after 100K deformation-stage iterations. The scale and rotation learning rates decay from 5×10^{-4} to 5×10^{-5}, while the attribute learning rate decays from 2.5×10^{-3} to 2.5×10^{-4}. We use Adam with a batch size of one and gradient clipping at 5, without additional deformation regularization. For datasets with phase annotations, we supervise the predicted complex spectrum using the real–imaginary loss, amplitude loss, phase loss, and S-SSIM. For datasets without phase information, only the amplitude and S-SSIM losses are applied. The amplitude and phase are derived from the predicted real and imaginary channels, with phase represented in the range $[-\pi, \pi]$. The loss weights are set to $\lambda_a = 0.7$, $\lambda_s = 0.3$, $\lambda_c = 0.1$, and $\lambda_p = \frac{0.1}{2\pi}$. The real–imaginary and amplitude terms use ℓ_1 losses on normalized values, while the phase term uses an ℓ_1 error measured in radians. Training and evaluation are performed in FP32. The reported model size is measured from the complete FP32 inference representation, including both the canonical spherical Gaussian parameters and the deformation-network parameters.

CUDA Rasterizer Implementation. We implement spherical Gaussian rasterization using custom CUDA kernels with a tile-based parallel design. Importantly, tiling is used only to accelerate candidate selection and does not alter the intrinsic definition of the spherical Gaussian response. For each deformed primitive, we define a truncated geodesic support

$$r_i = \tau \max(s^s_{1,i}, s^s_{2,i}), \qquad \mathcal{C}_i = \{\mathbf{q} \in \mathcal{M} : d_{\mathbb{S}^2}(\mathbf{n}^s_i, \mathbf{q}) \le r_i\}, \tag{40}$$

and assign the primitive only to raster tiles intersecting this spherical support. We use 16×16 tiles with a two-pixel padding during tile assignment and set $\tau = 3$. As described above, the geodesic displacement from each canonical center is bounded to 20°, while the one-sided effective support of each primitive is limited to at most 90°. Together, these constraints keep all Gaussian evaluations away from the antipodal configuration, where the Riemannian logarithmic map and shortest-geodesic parallel transport become non-unique.

For each candidate primitive–sample pair, we first compute

$$c_i(\mathbf{q}) = \mathrm{clip}\big((\mathbf{n}^s_i)^\top \mathbf{q}, -1, 1\big), \qquad \theta_i(\mathbf{q}) = \arccos c_i(\mathbf{q}). \tag{41}$$

The clamping prevents floating-point errors from producing invalid arguments to $\arccos$. Samples satisfying $\theta_i(\mathbf{q}) > r_i$ are rejected before evaluating the logarithmic map. For the remaining samples, we compute

$$\mathrm{Log}_{\mathbf{n}^s_i}(\mathbf{q}) = \kappa_i \left(\mathbf{q} - c_i \mathbf{n}^s_i\right), \qquad \kappa_i = \frac{\theta_i}{\sin\theta_i}. \tag{42}$$

The removable singularity at $\theta_i \to 0$ is handled using the continuous small-angle limit, with a local series expansion in the numerically sensitive regime. The resulting tangent coordinate is

$$\mathbf{z}^s_i(\mathbf{q}) = (\mathbf{T}^s_i)^\top \mathrm{Log}_{\mathbf{n}^s_i}(\mathbf{q}), \tag{43}$$

and the anisotropic Gaussian response is evaluated as

$$G^s_i(\mathbf{q}) = \exp\left[-\frac{1}{2}\sum_{k=1}^{2} \frac{z^s_{k,i}(\mathbf{q})^2}{(s^s_{k,i})^2}\right]. \tag{44}$$

The resulting complex primitive responses are accumulated additively to form the RF spectrum.

We apply analogous numerical safeguards to the spherical deformation operators. For the exponential map, the factor $\sin r/r$ is evaluated using its small-r series expansion when the tangent

Algorithm 1 Intrinsic Spherical Gaussian Rasterization

Require: Deformed primitives $\Theta^s = \{\mathbf{n}_i^s, \mathbf{T}_i^s, \mathbf{s}_i^s, a_i^s\}_{i=1}^N$; angular samples $\mathcal{A}$; cutoff factor τ; raster tiles $\{\mathcal{T}_m\}$; numerical threshold ε

Ensure: Rendered complex RF spectrum $\widehat{S}(\varphi, \lambda \mid \mathbf{s})$

Initialize $\widehat{S}(\varphi, \lambda \mid \mathbf{s}) \leftarrow 0$ for all $(\varphi, \lambda) \in \mathcal{A}$

Spherical support binning

for $i = 1, \ldots, N$ **do**
  $r_i \leftarrow \tau \max(s_{1,i}^s, s_{2,i}^s)$
  $\mathcal{C}_i \leftarrow \{\mathbf{q} \in \mathcal{M} : d_{\mathbb{S}^2}(\mathbf{n}_i^s, \mathbf{q}) \le r_i\}$
  Assign primitive i to raster tiles intersecting $\mathcal{C}_i$
end for

Intrinsic Gaussian evaluation

for all $\mathcal{T}_m$ **do**
  for all $(\varphi, \lambda) \in \mathcal{A} \cap \mathcal{T}_m$ **do**
    $\mathbf{q} \leftarrow \mathbf{q}(\varphi, \lambda)$
    for all i assigned to $\mathcal{T}_m$ **do**
      $c_i \leftarrow \mathrm{clip}((\mathbf{n}_i^s)^\top \mathbf{q}, -1, 1)$
      $\theta_i \leftarrow \arccos(c_i)$
      if $\theta_i > r_i$ **then**
        continue
      end if
      $\mathbf{d}_i \leftarrow \mathbf{q} - c_i \mathbf{n}_i^s$
      if $\theta_i < \varepsilon$ **then**
        $\kappa_i \leftarrow 1 + \theta_i^2/6 + 7\theta_i^4/360$
      else
        $\kappa_i \leftarrow \theta_i / \sin\theta_i$
      end if
      $\mathbf{u}_i \leftarrow \kappa_i \mathbf{d}_i$
      $\mathbf{z}_i^s \leftarrow (\mathbf{T}_i^s)^\top \mathbf{u}_i$
      $G_i^s \leftarrow \exp\left[-\frac{1}{2}\sum_{k=1}^2 \frac{(z_{k,i}^s)^2}{(s_{k,i}^s)^2}\right]$
      $\widehat{S}(\varphi, \lambda \mid \mathbf{s}) \leftarrow \widehat{S}(\varphi, \lambda \mid \mathbf{s}) + a_i^s G_i^s$
    end for
  end for
end for
return $\widehat{S}(\cdot, \cdot \mid \mathbf{s})$

displacement approaches zero. For parallel transport, the 20° bound on center displacement keeps $1 + (\mathbf{n}_i^0)^\top \mathbf{n}_i^s$ safely away from zero, avoiding the antipodal singularity of the closed-form transport operator. Thus, the implementation explicitly handles the removable small-angle singularities while excluding the geometrically non-unique antipodal regime by construction. Both primitive support and response evaluation remain defined directly by spherical geometry, while the tile-based CUDA implementation serves purely as a computational acceleration of the intrinsic rasterization in Alg. 1.

F.4 Evaluation Metrics Details

Spherical SSIM. S-SSIM is evaluated on the azimuthal-equidistant hemispherical projection, rather than directly on the native equirectangular spectrum grid. The original S-SSIM Chen et al. (2018) constructs local structural comparisons while accounting for the nonuniform area represented by planar samples on the sphere. In our setting, both prediction and ground truth are first mapped to the same azimuthal-equidistant hemispherical domain; we therefore compute local SSIM directly on this projected domain and retain the spherical area-weighting principle of S-SSIM when aggregating the local scores. For the azimuthal-equidistant projection, where $r \propto \theta$, the spherical and projected area

elements are proportional to $\sin\theta\, d\theta\, d\phi$ and $\theta\, d\theta\, d\phi$, respectively. The corresponding area weight is

$$w(\theta) = \frac{\sin\theta}{\theta}, \tag{45}$$

with $w(0) = 1$, where θ denotes the angular distance from the projection center. We then compute

$$\text{S-SSIM} = \frac{\sum_{\mathbf{p}\in\Omega} w(\theta_{\mathbf{p}})\, \text{SSIM}_{\mathbf{p}}}{\sum_{\mathbf{p}\in\Omega} w(\theta_{\mathbf{p}})}, \tag{46}$$

where Ω denotes the valid samples on the projected hemispherical disk. S-SSIM is evaluated on the reconstructed and ground-truth amplitude spectra.

Phase Error. Since RF phase is periodic, we compute the phase error using the wrapped angular difference. Given the predicted phase $\widehat{\psi}(\varphi, \lambda)$ and the ground-truth phase $\psi(\varphi, \lambda)$, the wrapped phase difference is

$$\Delta\psi(\varphi,\lambda) = \text{atan2}\left(\sin\Big(\widehat{\psi}(\varphi,\lambda) - \psi(\varphi,\lambda)\Big), \cos\Big(\widehat{\psi}(\varphi,\lambda) - \psi(\varphi,\lambda)\Big)\right), \tag{47}$$

which lies in $[-\pi, \pi]$. We then compute

$$\text{Phase MAE} = \frac{1}{|\Omega_{\text{valid}}|} \sum_{(\varphi,\lambda)\in\Omega_{\text{valid}}} |\Delta\psi(\varphi,\lambda)|, \tag{48}$$

and

$$\text{Phase RMSE} = \sqrt{\frac{1}{|\Omega_{\text{valid}}|} \sum_{(\varphi,\lambda)\in\Omega_{\text{valid}}} \Delta\psi(\varphi,\lambda)^2}. \tag{49}$$

Here, Ω_{valid} denotes the angular samples whose ground-truth amplitude exceeds 1% of the peak amplitude. Both phase errors are reported in radians.

G LIMITATIONS

The current SphGaS formulation is primarily designed for scene-specific RF spectrum reconstruction, assuming a fixed environment while the TX–RX configuration varies. Its generalization across environments or changing scene geometry remains unexplored. A second limitation is the lack of real-world evaluation under joint TX–RX mobility. The real-world RFID benchmark used in our experiments mainly covers single-endpoint mobility, while our joint TX–RX evaluation is conducted on ray-traced synthetic data, leaving a potential sim-to-real gap. Nevertheless, the results on the real-world RFID benchmark support the effectiveness of the proposed spherical representation and rendering formulation under measured RF data. Future work will extend the evaluation to real-world settings with jointly moving TX and RX.