RadioDiff-FS: Physics-Informed Manifold Alignment in Few-Shot Diffusion Models for High-Fidelity Radio Map Construction

Xiucheng Wang, *Graduate Student Member, IEEE*, Zixuan Guo, Nan Cheng, *Senior Member, IEEE*, Zhisheng Yin, *Member, IEEE*, Ruijin Sun, *Member, IEEE*, and Xuemin Shen, *Fellow, IEEE*

***Abstract*—Radio maps (RMs) provide spatially continuous propagation characterizations essential for 6G network planning, but high-fidelity RM construction remains challenging. Rigorous electromagnetic solvers incur prohibitive computational latency, while data-driven models demand massive labeled datasets and generalize poorly from simplified simulations to complex multipath environments. This article proposes RadioDiff-FS, a few-shot diffusion framework that adapts a pretrained main-path (MP) generator to multipath-rich target domains with only a small number of high-fidelity samples. The adaptation is grounded in a power-domain decomposition of the multi-path RM into a dominant MP component and a structured residual. This decomposition indicates that the dominant path model (DPM)-to-IRT4 domain shift can be approximated as a bounded feature-space correction under matched geometry and transmitter semantics. A direction-consistency loss (DCL) is then introduced to constrain diffusion updates along the empirically estimated DPM-to-IRT4 feature-shift direction, thereby suppressing propagation-inconsistent artifacts in the low-data regime. Experiments show that RadioDiff-FS reduces normalized mean squared error by 59.5% on static RMs and by 74.0% on dynamic RMs relative to the vanilla diffusion baseline, achieving an structural similarity index measure of 0.9752 and a peak signal-to-noise ratio of 36.37 dB under severely limited supervision. Even in a one-shot setting with a single target-domain sample per scene, RadioDiff-FS maintains competitive structural fidelity, confirming that the directional constraint provides an effective inductive bias under extreme data scarcity. Code is available at https://github.com/UNIC-Lab/RadioDiff-FS**

***Index Terms*—Diffusion models, few-shot learning, physics-informed learning, radio map (RM).**

I. Introduction

RADIO maps (RMs) provide a spatially continuous abstraction of large-scale wireless propagation and are increasingly recognized as a core spatial infrastructure for 6G intelligence [1], [2]. By decoupling channel state information (CSI) acquisition from real-time pilot overhead, RMs serve as a unified [3], queryable substrate for latency-critical functionalities, including user-centric coverage optimization [4], proactive beamforming [5], and electromagnetic digital twin construction [6]. Unlike transient pointwise channel estimates, RMs explicitly encode global spatial layouts and capture mesoscopic propagation features that govern network behavior, from occlusion boundaries defined by building geometries to diffraction at sharp edges and high-transmission apertures [7], [8]. This structural awareness is particularly valuable at high frequencies. In dense millimeter-wave (mmWave) deployments [9], a precomputed RM enables anticipation of nonline-of-sight (NLoS) blockage [10], [11], facilitating proactive UAV-assisted relay planning and robust connectivity under dynamic environmental conditions [12].

Despite this potential, high-fidelity RM construction is constrained by two fundamental bottlenecks regarding computational efficiency [8], [13] and distributional generalization [7], [14], [15]. Rigorous electromagnetic solvers, including full-wave methods and high-fidelity ray-tracing engines, provide authoritative approximations of real-world propagation physics [8], [13]. However, when modeling carrier-frequency dependence, fine spatial resolution, and higher-order multipath interactions in material-heterogeneous environments, these solvers frequently incur minute-level latency even for modest scene scales [8]. Such overhead precludes rapid design-space exploration and the real-time scene updates required by dynamic network digital twins [16]. Data-driven generative models, including U-Net [17], Transformer architectures [18], and diffusion-based backbones, can synthesize kilometer-scale RMs within sub-second runtimes [7], [14], [19]. Yet their generalization is limited by the training distribution. Because scalable training often relies on simplified propagation surrogates, these models frequently fail to capture nonstationary phenomena that dominate real-world environments, including high-order reflections and edge diffraction [7], [20]. Diffusion models further exacerbate this data demand, as their iterative reverse process requires comprehensive coverage of the data manifold for stable trajectory learning, intensifying the dependence on massive multipath datasets that are prohibitively expensive to acquire.

The underlying obstacle is the unfavorable economics of high-fidelity data acquisition, as illustrated in Fig. 1.

Received 25 March 2026; revised 17 May 2026; accepted 16 June 2026. Date of publication 24 June 2026; date of current version 26 August 2026. This work was supported in part by the National Key Research and Development Program of China under Grant 2024YFB907500 and in part by the National Natural Science Foundation of China under Grant 62571402. *(Corresponding author: Nan Cheng.)*

Xiucheng Wang, Nan Cheng, Zhisheng Yin, and Ruijin Sun are with the State Key Laboratory of ISN and the School of Telecommunications Engineering, Xidian University, Xi'an 710071, China (e-mail: xcwang_1@stu.xidian.edu.cn; dr.nan.cheng@ieee.org; zsyin@xidian.edu.cn; sunruijin@xidian.edu.cn).

Zixuan Guo is with the School of Physics, Xidian University, Xi'an 710071, China (e-mail: 24179100067@stu.xidian.edu.cn).

Xuemin Shen is with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada (e-mail: sshen@uwaterloo.ca).

Digital Object Identifier 10.1109/JIOT.2026.3706895

2327-4662 © 2026 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission.
See https://www.ieee.org/publications/rights/index.html for more information.

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

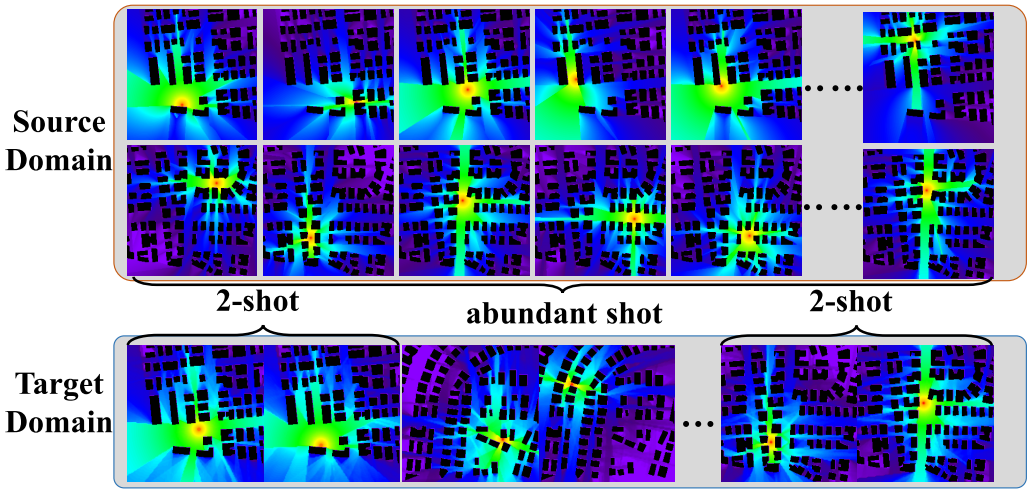


Fig. 1. Illustration of the few-shot learning for RM construction.

Field measurements incur prohibitive costs in skilled labor, specialized equipment, and logistical coordination, rendering large-scale collection across diverse cities, material classes, and architectural morphologies impractical. High-fidelity ray-tracing simulation faces comparable scalability barriers. Generating fifty thousand samples via rigorous multipath tracing would typically require over eight hundred compute-hours [21], and this burden increases substantially with higher-order reflections, realistic material catalogs, and vertical scene complexity. Exhaustively expanding datasets to cover all relevant multipath interactions is neither scalable nor economically viable [22]. These constraints motivate a paradigm shift from indiscriminate data accumulation to principled cross-domain adaptation. The central hypothesis is that a tractable solution should leverage abundant, computationally inexpensive main-path (MP) data to learn global spatial layouts, while relying on only a small set of high-fidelity samples to correct structural biases and inject domain-specific multipath details. This few-shot strategy reframes the objective as transfer learning that couples generative modeling with physical consistency.

To address these challenges, we propose RadioDiff-FS, a framework that reformulates high-fidelity RM construction as a physics-informed few-shot diffusion adaptation problem. Rather than treating RM generation as an opaque image-to-image translation task, our approach builds on a rigorous decomposition of the electromagnetic field. The multipath RM is formalized as a dominant MP component augmented by a structured residual. This decomposition provides a strong physical inductive bias, indicating that adaptation from a MP prior to a multipath target involves learning a bounded geometric correction rather than synthesizing a new distribution from scratch. RadioDiff-FS therefore transfers global spatial layout knowledge from the pretrained backbone and concentrates generative capacity on localized diffraction and reflection effects. A direction-consistency loss (DCL) is introduced to regularize the fine-tuning process. The DCL constrains score function updates to evolve along physically plausible propagation directions by penalizing gradient components orthogonal to the valid propagation manifold. This suppresses propagation-inconsistent artifacts and mode hallucinations in low-data regimes while maintaining geometric fidelity at diffraction edges and shadow boundaries. RadioDiff-FS targets pathloss or received-power RMs, rather than phase-resolved complex fields. Thus, the noncoherent residual model is used as a power-domain approximation consistent with RadioMapSeer labels, while coherent cancellation is left to future complex-valued channel-map studies. The main contributions are summarized as follows.

1) We propose RadioDiff-FS, which is, to the best of our knowledge, the first framework to formulate high-fidelity RM construction as a few-shot diffusion adaptation problem. By transferring knowledge from a pretrained backbone learned on computationally inexpensive MP priors to multipath-rich target domains, it bridges the distributional gap caused by complex propagation environments and circumvents prohibitive data acquisition costs.
2) We establish a power-domain decomposition of the multipath RM into a dominant MP component and a structured residual. This formulation provides a physical inductive bias and shows that the dominant path model (DPM)-to-IRT4 transition can be modeled as a bounded feature-space correction under matched environment and transmitter semantics. We further clarify the boundary of this approximation. It applies to pathloss or received-power maps, while phase-resolved coherent cancellation requires a complex-valued formulation.
3) We design a DCL that functions as a lightweight feature-space manifold projector during few-shot adaptation. The loss estimates an average DPM-to-IRT4 shift direction from paired samples and constrains generated features to move toward the multipath-aware domain. It suppresses orthogonal feature drifts and negative domain shifts, while still allowing sample-dependent projection magnitudes that reflect local multipath variations.
4) Extensive experiments on the RadioMapSeer benchmark demonstrate that RadioDiff-FS achieves state-of-the-art performance in both static and dynamic scenarios,

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

with superior edge preservation and significantly lower reconstruction error compared to existing baselines under extremely limited supervision. Ablation studies validate the critical role of DCL in stabilizing the adaptation trajectory.

II. Related Work and Preliminary

A. Rm Construction Methods

RMs offer a spatially continuous abstraction of wireless propagation environments and serve as a fundamental enabler for environment-aware 6G applications, such as coverage optimization, localization, and digital twin construction [1]. In recent years, learning-based approaches have shown great promise in improving the accuracy and efficiency of RM construction, particularly in complex urban and indoor environments where traditional statistical models based on radial symmetry fail to account for diffraction, shadowing, and occlusion effects. However, despite their effectiveness, these neural models often rely on large-scale training data that is expensive to obtain either through simulation or measurement.

Initial studies, such as RadioUNet [17], demonstrated that convolutional encoder–decoder architectures can produce accurate 2-D pathloss maps from structured urban layouts, significantly reducing inference latency compared with ray tracing. Subsequent works like PMNet improved generalization and sample efficiency through advanced vision architectures. Although these discriminative models remain computationally efficient and interpretable, they typically require dense supervision and are difficult to adapt to out-of-distribution environments or scenarios with sparse sensor coverage. Moreover, most of these methods focus on planar pathloss estimation, ignoring vertical structure and multimetric coupling [14], [19]. To overcome the limitations of pointwise regression under sparse measurements, conditional generative models have been explored. RME-GAN [23] and DeepREM [24] propose using conditional generative adversarial networks and completion autoencoders to reconstruct dense RMs from limited measurements. These models integrate global propagation patterns and local texture variations to enhance estimation performance. While they significantly outperform interpolation baselines, they still rely heavily on simulation data and do not incorporate electromagnetic consistency into the learning process. Graph neural networks (GNNs) offer another route by encoding structural priors derived from building layouts and sensor topology [25]. Recent methods formulate sparse measurements as nodes and use edge structures informed by connectivity or spatial correlation to facilitate map completion. RadioGAT [26] extends this by capturing spatial–spectral dependencies across multiple frequency bands. These models offer improved robustness to sampling irregularities and can generalize under semisupervised settings. However, without explicit electromagnetic constraints, their behavior around complex propagation features, such as diffraction edges and shadow zones remains difficult to guarantee.

Diffusion models have recently emerged as a powerful alternative due to their generative structure and sample flexibility. RadioDiff [14] models RM generation as a conditional diffusion process and introduces attention mechanisms and fast Fourier modules to enhance feature extraction in dynamic environments. RadioDiff-3D builds upon this by incorporating 3-D convolutional networks and utilizing the UrbanRadio3D dataset, which contains multiheight, multiparameter RMs, including pathloss, direction-of-arrival (DoA), and time-of-arrival (ToA) [27]. These models achieve state-of-the-art performance in high-dimensional map estimation. However, without embedding physical priors, they may fail to reconstruct physically meaningful structures in unseen domains, especially in low-data regimes. To address this, physics-informed models have been introduced to embed electromagnetic theory into neural estimators. PEFNet integrates the volume integral equation into the learning objective, constraining the relationship between incident and total electric fields. More advanced methods, such as PhyRMDM [28] and RadioDiff-k^2 [7], incorporate Helmholtz residuals, source terms, and boundary conditions into diffusion model training. RadioDiff-k^2 in particular establishes a theoretical correspondence between wave propagation singularities and negative wavenumber domains, enabling the network to identify and reconstruct diffraction-sensitive structures via a two-stage generative process. These approaches significantly improve physical plausibility and generalization, but they often assume access to accurate simulations or well-calibrated measurements, which remain costly and difficult to obtain at scale. Task-specific strategies, such as RM inpainting based on propagation priority and radio depth maps have also been proposed to address areas where sensing is physically restricted [6]. In parallel, the development of large-scale datasets, such as UrbanRadio3D has greatly expanded the coverage of height, geometry, and multimetric labeling, providing an essential benchmark for high-fidelity 3-D $\times$ 3-D RM construction. Nevertheless, the cost and time required for data generation via ray tracing or controlled field measurements remain prohibitive in most deployment scenarios, and the resulting models often exhibit limited cross-region transferability.

In summary, although learning-based methods for RM construction have achieved impressive results across a variety of modeling paradigms, they remain heavily reliant on large volumes of high-quality data, either from simulation or measurement. However, due to the inherently high cost of generating such data, especially for diverse urban environments and 3-D geometries, reliance on large-scale supervision represents a critical bottleneck. This reality motivates a paradigm shift toward few-shot adaptation, wherein a pretrained model is efficiently fine-tuned to a new deployment scenario using only a handful of labeled samples.

B. EM Computing Methods for RM Construction

Recent advances in deep learning-based RM construction have underscored the importance of accurate and scalable ground truth generation. In practice, the DPM [29] has emerged as the standard solver for large-scale RM datasets due to its computational efficiency and fast rendering capabilities. By modeling only the strongest propagation path between transmitter and receiver while omitting multipath effects, DPM enables rapid data generation across dense spatial grids and diverse environmental layouts. As a result,

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

mainstream benchmarks, such as RadioMapSeer, UrbanRadio3D, and UrbanMIMOMap have all adopted DPM as their primary simulation engine. These datasets typically comprise tens of thousands of labeled samples, with UrbanRadio3D notably scaling up to over ten million images, making it the largest publicly available RM dataset to date. The widespread use of DPM reflects a practical tradeoff: while it sacrifices full physical fidelity, its low computational cost makes it ideal for supporting supervised training of large neural architectures. Shooting-and-bouncing ray (SBR) [30] methods, which adopt a ray-launching approach that allows for multiple specular reflections, provide a richer geometric representation than DPM. However, their sampling-based mechanism does not guarantee that rays will reach all regions of interest, particularly in dense, occluded, or NLOS environments. This coverage uncertainty introduces spatial sparsity that is undesirable in dense RM estimation tasks, particularly when complete spatial coverage is required for training or evaluation. As a result, although SBR is used in some site-specific simulations, it has seen limited adoption in large-scale RM dataset generation compared to DPM.

In contrast, intelligent ray tracing (IRT) [31] offers a more rigorous modeling paradigm that more faithfully adheres to electromagnetic propagation principles. By solving high-frequency approximations of Maxwell's equations, IRT explicitly accounts for reflection, diffraction, and scattering, thereby generating multipath-aware RMs that more accurately represent complex propagation effects. This makes IRT particularly suitable for reconstructing signal behavior near sharp edges, shadow boundaries, and other critical regions. However, the substantial computational overhead associated with full multipath tracing has constrained the scale at which IRT can be applied. As a result, only a small number of publicly available datasets incorporate IRT-simulated labels. For instance, RadioMapSeer includes an IRT subset of just over 1000 samples, which remains insufficient for training contemporary deep neural networks without relying on transfer learning or extensive data augmentation.

C. Denoising Diffusion Model

DMs have emerged as highly effective generative frameworks capable of learning complex data distributions and synthesizing high-fidelity samples with remarkable precision [32]. Among these, score-based DMs provide a principled approach grounded in stochastic differential equations (SDEs), wherein data is progressively perturbed by noise in the forward process and subsequently recovered by learning the score function, defined as the gradient of the log-probability density of the noisy data [33]. Unlike conventional DMs that rely on discrete-time Markov chains, score-based models employ continuous-time dynamics, making them particularly suitable for inverse problems, such as RM reconstruction under the Bayesian sampling paradigm. Formally, the forward SDE is defined as

$$d\boldsymbol{x} = f(\boldsymbol{x}, t)\, dt + g(t)\, d\boldsymbol{w} \tag{1}$$

where t is the diffusion time index, and $f(\boldsymbol{x}, t)$ denotes the drift term, $g(t)$ is a time-dependent diffusion coefficient, and $d\boldsymbol{w}$ represents a standard Wiener process. As t increases, the perturbed distribution $p_t(\boldsymbol{x})$ asymptotically approaches an isotropic Gaussian. The generative process involves solving the corresponding reverse-time SDE

$$d\boldsymbol{x} = \left[f(\boldsymbol{x}, t) - g^2(t)\, \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x})\right] dt + g(t)\, d\bar{\boldsymbol{w}} \tag{2}$$

where the score function $\nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x})$ is typically approximated by a neural network $s_\theta(\boldsymbol{x}, t)$, trained to satisfy

$$s_\theta(\boldsymbol{x}, t) \approx \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x}). \tag{3}$$

For deterministic sampling, the equivalent probability flow ordinary differential equation (ODE) is

$$d\boldsymbol{x} = \left[f(\boldsymbol{x}, t) - \frac{1}{2} g^2(t)\, \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x})\right] dt \tag{4}$$

which removes stochasticity while preserving sample quality, thereby drawing close connections to denoising diffusion probabilistic models (DDPMs). In the DDPM framework, the forward process is discretized as

$$q(\boldsymbol{x}_t | \boldsymbol{x}_{t-1}) = \mathcal{N}(\boldsymbol{x}_t; \alpha_t \boldsymbol{x}_{t-1}, \beta_t \boldsymbol{I}) \tag{5}$$

with the score function approximated via a learned denoiser ϵ_θ as

$$s_\theta(\boldsymbol{x}, t) = -\frac{\epsilon_\theta(\boldsymbol{x}, t)}{\sqrt{1 - \bar{\alpha}_t}}. \tag{6}$$

Thus, DDPM can be viewed as a discretized instance of a score-based continuous model under a variance-preserving schedule.

In the context of RM generation, a decoupled formulation has recently been proposed to enhance generative stability, notably in the RadioDiff framework [14]. Rather than injecting noise directly, the decoupled diffusion model (DDM) first attenuates the input signal to a baseline before introducing noise. The forward transition from clean input $\boldsymbol{x}_0$ to noisy state $\boldsymbol{x}_t$ is governed by

$$q(\boldsymbol{x}_t \mid \boldsymbol{x}_0) = \mathcal{N}\left(\gamma_t \boldsymbol{x}_0, \delta_t^2 \boldsymbol{I}\right) \tag{7}$$

where γ_t controls signal decay and δ_t determines noise variance. This process can be equivalently expressed through the following SDE

$$d\boldsymbol{x}_t = \boldsymbol{f}_t \boldsymbol{x}_t dt + g_t d\boldsymbol{\epsilon}_t \tag{8}$$

$$\boldsymbol{f}_t = \frac{d \log \gamma_t}{dt} \tag{9}$$

$$\int_0^t \boldsymbol{f}_t dt = -x_0 \tag{10}$$

$$g_t^2 = \frac{d\delta_t^2}{dt} - 2\boldsymbol{f}_t \delta_t^2 \tag{11}$$

where $\boldsymbol{f}_t$ defines the rate of attenuation and g_t^2 governs the rate of noise accumulation.

To generate samples, the reverse-time SDE is solved

$$d\boldsymbol{x}_t = \left[\boldsymbol{f}_t \boldsymbol{x}_t - g_t^2 \nabla_{\boldsymbol{x}} \log q(\boldsymbol{x}_t)\right] dt + g_t d\bar{\boldsymbol{\epsilon}}_t. \tag{12}$$

The structured separation between signal attenuation and noise addition enables more stable training and higher-quality

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

synthesis. Notably, the deterministic component of DDM simplifies the forward mapping as

$$q(\boldsymbol{x}_t|\boldsymbol{x}_0) = \mathcal{N}\left(\boldsymbol{x}_0 + \int_0^t \boldsymbol{f}_\tau d\tau, t\boldsymbol{I}\right) \tag{13}$$

and permits an efficient reverse-step formulation

$$q(z_{t-\Delta t} \mid z_t, z_0) = \mathcal{N}\left(z_t + \int_t^{t-\Delta t} \boldsymbol{f}_t \, \mathrm{d}t - \frac{\Delta t}{\sqrt{t}}\boldsymbol{\epsilon}, \frac{\Delta t\,(t-\Delta t)}{t}\boldsymbol{I}\right). \tag{14}$$

By decoupling signal dynamics and stochasticity, the DDM approach significantly improves generative quality, robustness, and computational efficiency. These properties render it highly advantageous for high-resolution and structure-aware RM reconstruction tasks under sparse supervision, where stability and controllability are essential.

III. System Model and Problem Formulation

We consider a wireless propagation scenario over a discretized planar region represented as an $N \times N$ uniform grid. Each cell in this grid corresponds to a spatial resolution element of width Δ. The cell size is assumed small enough that the pathloss within it can be treated as constant. A RM over this region is defined as a matrix $\boldsymbol{P} \in \mathbb{R}^{N\times N}$, where each entry $P(i, j)$ denotes the pathloss at the center of cell (i, j). The pathloss is measured at a fixed carrier frequency f_c with the corresponding wavelength $\lambda = c/f_c$, where c denotes the speed of light in free space. A single transmitter, modeled as a point radiator, is positioned at $\boldsymbol{r} = \langle d_x, d_y, d_z\rangle$, where (d_x, d_y) denotes the horizontal coordinates in the 2-D region and $d_z > 0$ indicates the height above the ground. The transmitter emits electromagnetic energy isotropically or according to a specified antenna gain pattern, and serves as the sole radiation source. Although we focus on the single-transmitter setting, the formulation can be extended to multiantenna or multisite deployments by considering multiple independent transmitter locations.

The surrounding environment is composed entirely of static structures, such as buildings. The geometric layout of these structures is captured by a binary occupancy matrix $\boldsymbol{H} \in \{0, 1\}^{N\times N}$. An entry $H(i, j) = 1$ indicates that the cell at location (i, j) is fully occupied by a nonpenetrable obstacle, while $H(i, j) = 0$ represents free space. These obstacles are assumed to fully block direct wave propagation. Given a transmitter position $\boldsymbol{r}$ and a static environment $\boldsymbol{H}$, the goal is to estimate the pathloss field $\boldsymbol{P}_\star(\boldsymbol{r}) = \mathcal{F}(\boldsymbol{H}; \boldsymbol{r})$, where $\mathcal{F}$ denotes the environment-to-field operator. In this work, we consider two distinct types of RMs. The first type is the static RM (SRM), which models only the dominant propagation path and captures large-scale attenuation effects, such as free-space spreading and shadowing due to occlusions. The second type is the multipath RM (MRM), which provides a more detailed characterization of the propagation field by incorporating multipath effects. These effects include reflections and diffractions induced by the surrounding static geometry. The MRM offers a more physically complete description of the electromagnetic field, particularly in NLoS regions.

To support the construction of these RMs, we assume the availability of two labeled datasets. The first dataset is denoted as $\mathcal{D}_{\mathrm{SRM}} = \{\boldsymbol{H}_k, \boldsymbol{r}_k, \boldsymbol{P}_k^{\mathrm{SRM}}\}_{k=1}^K$ and contains a large number of examples generated under the SRM setting. These samples are efficient to compute and cover a wide variety of layouts and transmitter placements. The second dataset is denoted as $\mathcal{D}_{\mathrm{MRM}} = \{\boldsymbol{H}_\ell, \boldsymbol{r}_\ell, \boldsymbol{P}_\ell^{\mathrm{MRM}}\}_{\ell=1}^L$ and contains a small number of high-fidelity samples that capture multipath propagation effects. Due to the computational overhead of accurate multipath modeling, the size of this dataset is limited, i.e., $L \ll K$. The objective of the above problem can be formulated as follows:

$$\min_{\boldsymbol{\theta}} \sum_{\ell=1}^{L} \frac{1}{L}\mathcal{L}\left(\boldsymbol{P}_\ell^{\mathrm{MRM}}, \hat{\boldsymbol{P}}_\ell^{\mathrm{MRM}}\right) \tag{15}$$

$$\text{s.t.} \quad \hat{\boldsymbol{P}}_\ell^{\mathrm{MRM}} = \boldsymbol{\mu}_{\boldsymbol{\theta}}(\boldsymbol{H}, \boldsymbol{r} \mid \mathcal{D}_{\mathrm{SRM}}, \mathcal{D}_{\mathrm{MRM}}) \tag{15a}$$

where $\boldsymbol{\theta}$ denotes the trainable parameters of the neural network $\boldsymbol{\mu}(\cdot)$.

IV. Few-Shot DM for RM Construction

A. MP Decomposition

Consider a bounded planar domain $\Omega \subset \mathbb{R}^2$ and a fixed transmitter–environment configuration. At a receiver location $\boldsymbol{x} \in \Omega$, let $\mathcal{P}(\boldsymbol{x})$ denote a finite set of propagation paths induced by the environment. Under the power-incoherent superposition model commonly adopted in RM construction, the received power field is expressed as the nonnegative sum

$$U(\boldsymbol{x}) = \sum_{p\in\mathcal{P}(\boldsymbol{x})} P_p(\boldsymbol{x}), \quad P_p(\boldsymbol{x}) \geq 0 \tag{16}$$

where $P_p(\boldsymbol{x})$ aggregates geometric spreading and interaction-dependent losses along path p. This is a power-domain RM model. It aggregates path powers after phase information has been marginalized, and should not be read as coherent complex-field superposition. Thus, the monotone residual relation is restricted to pathloss or received-power maps, while destructive interference in phase-resolved channels is outside this formulation. Let $p^\star(\boldsymbol{x}) \in \arg\max_{p\in\mathcal{P}(\boldsymbol{x})} P_p(\boldsymbol{x})$ denote a strongest path at $\boldsymbol{x}$. The MP-RM is then defined as

$$U_{\mathrm{MP}}(\boldsymbol{x}) = P_{p^\star(\boldsymbol{x})}(\boldsymbol{x}). \tag{17}$$

Let $\mathcal{P}_k(\boldsymbol{x}) \subseteq \mathcal{P}(\boldsymbol{x})$ denote the subset that collects all paths whose interaction order does not exceed k, and assume $p^\star(\boldsymbol{x}) \in \mathcal{P}_k(\boldsymbol{x})$. The MRM is defined by aggregating all such paths

$$U_{\mathrm{MU}}(\boldsymbol{x}) = \sum_{p\in\mathcal{P}_k(\boldsymbol{x})} P_p(\boldsymbol{x}). \tag{18}$$

Subtracting (17) from (18) yields the fundamental decomposition

$$U_{\mathrm{MU}}(\boldsymbol{x}) = U_{\mathrm{MP}}(\boldsymbol{x}) + \Delta U(\boldsymbol{x}) \tag{19}$$

$$\Delta U(\boldsymbol{x}) \triangleq \sum_{p\in\mathcal{P}_k(\boldsymbol{x})\setminus\{p^\star(\boldsymbol{x})\}} P_p(\boldsymbol{x}). \tag{20}$$

Since every summand in (20) is nonnegative, the transition from MP-RM to MU-RM is pointwise monotone

$$U_{\mathrm{MU}}(\boldsymbol{x}) \geq U_{\mathrm{MP}}(\boldsymbol{x}) \tag{21}$$

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

$$\Delta U(\boldsymbol{x}) \geq 0 \quad \forall\, \boldsymbol{x} \in \Omega \tag{22}$$

and $\Delta U(\boldsymbol{x}) = 0$ holds if and only if no additional path contributes power at $\boldsymbol{x}$ within $\mathcal{P}_k(\boldsymbol{x})$. The residual is structured rather than universally sparse. It is often concentrated near visibility changes, street corners, and shadow boundaries, but may spread more broadly in extreme NLoS scenes with dense higher-order interactions. DCL therefore uses this structure as a soft directional regularizer, not as a hard sparsity constraint.

To describe the spatial regularity of the residual, we introduce a visibility indicator $V_p(\boldsymbol{x}) \in \{0,1\}$ for each path p and factorize the per-path power as

$$P_p(\boldsymbol{x}) = V_p(\boldsymbol{x})\, \widetilde{P}_p(\boldsymbol{x}) \tag{23}$$

$$\widetilde{P}_p(\boldsymbol{x}) \geq 0 \tag{24}$$

where $\widetilde{P}_p(\boldsymbol{x})$ captures the interaction losses along path p. This factor varies smoothly with $\boldsymbol{x}$ whenever the interaction sequence defining p remains valid. We define the visibility state as $S(\boldsymbol{x}) = \{V_p(\boldsymbol{x})\}_{p\in\mathcal{P}_k(\boldsymbol{x})}$ and partition Ω into regions of constant visibility state

$$\mathcal{R}_s = \{\boldsymbol{x} \in \Omega :\ S(\boldsymbol{x}) = s\}. \tag{25}$$

On each region $\mathcal{R}_s$, every indicator $V_p(\boldsymbol{x})$ is constant. Therefore, both U_{MP} and U_{MU} reduce to finite sums of smooth functions $\widetilde{P}_p(\boldsymbol{x})$ and inherit the same smoothness within $\mathcal{R}_s$. Nonsmooth behavior of spatial derivatives can only occur on the union of visibility boundaries $\partial\mathcal{R}_s$, where at least one indicator $V_p(\boldsymbol{x})$ toggles. Under standard geometric optics rules, these boundaries form a finite union of curves and occupy zero area measure in Ω. The MU-RM therefore differs from the MP-RM through a nonnegative residual ΔU that is smooth almost everywhere and changes sharply only across visibility boundaries. The decomposition (19) implies an additive relation under any linear integral operator. Let $\mathcal{L}_\sigma$ be a low-pass operator with bandwidth parameter σ, represented by a nonnegative kernel κ_σ of unit mass

$$\mathcal{L}_\sigma[U](\boldsymbol{x}) = \int_\Omega \kappa_\sigma(\boldsymbol{x}-\boldsymbol{y})\, U(\boldsymbol{y})\, \mathrm{d}\boldsymbol{y} \tag{26}$$

$$\int_{\mathbb{R}^2} \kappa_\sigma(\boldsymbol{u})\, \mathrm{d}\boldsymbol{u} = 1, \quad \kappa_\sigma(\boldsymbol{u}) \geq 0. \tag{27}$$

By linearity and (19)

$$\mathcal{L}_\sigma[U_{\mathrm{MU}}] = \mathcal{L}_\sigma[U_{\mathrm{MP}}] + \mathcal{L}_\sigma[\Delta U]. \tag{28}$$

A quantitative bound on the low-frequency increment follows from standard convolution inequalities. We denote the L^p norm of a function f over Ω by $\|f\|_p$. Applying Young's convolution inequality yields

$$\|\mathcal{L}_\sigma[\Delta U]\|_2 \leq \|\kappa_\sigma\|_1\ \|\Delta U\|_2 = \|\Delta U\|_2 \tag{29}$$

and

$$\|\mathcal{L}_\sigma[\Delta U]\|_\infty \leq \|\kappa_\sigma\|_\infty\ \|\Delta U\|_1. \tag{30}$$

Suppose ΔU is effectively supported on a subset $\mathcal{S} \subseteq \Omega$ with small area and is bounded by $\Delta U(\boldsymbol{x}) \leq U_{\max}$. Then the L^1 norm of the residual satisfies

$$\|\Delta U\|_1 = \int_\Omega \Delta U(\boldsymbol{x})\, \mathrm{d}\boldsymbol{x} = \int_{\mathcal{S}} \Delta U(\boldsymbol{x})\, \mathrm{d}\boldsymbol{x} \leq U_{\max}|\mathcal{S}| \tag{31}$$

where $|\mathcal{S}|$ denotes the area of $\mathcal{S}$. Combining (30) with (31) yields

$$\|\mathcal{L}_\sigma[\Delta U]\|_\infty \leq \|\kappa_\sigma\|_\infty\, U_{\max}|\mathcal{S}|. \tag{32}$$

For typical low-pass kernels, $\|\kappa_\sigma\|_\infty$ decreases as σ increases. Thus the coarse-scale effect of a spatially localized residual is attenuated by smoothing. This result formalizes a central structural implication of (19): the MP-RM dominates the large-scale skeleton of the RM, while the MU-RM introduces structured positive increments concentrated in geometric neighborhoods where additional paths become visible.

B. Feature Shift Geometry

The MP-RM and the MU-RM satisfy the residual decomposition

$$U_{\mathrm{MU}}(\boldsymbol{x}) = U_{\mathrm{MP}}(\boldsymbol{x}) + \Delta U(\boldsymbol{x}),\ \Delta U(\boldsymbol{x}) \geq 0 \quad \forall\, \boldsymbol{x} \in \Omega. \tag{33}$$

To connect this physical structure with learning, we introduce a fixed feature encoder

$$\Phi :\ L^2(\Omega) \to \mathbb{R}^d, \quad \boldsymbol{z} = \Phi(U) \tag{34}$$

which maps a RM field to a d-dimensional representation. The encoder Φ is kept fixed during adaptation and serves only to measure domain shift and structural consistency. We assume that the representation varies stably with the RM field, as stated below.

Assumption 1: There exists a constant $L_\Phi > 0$ such that

$$\|\Phi(U) - \Phi(\widetilde{U})\|_2 \leq L_\Phi \|U - \widetilde{U}\|_2 \quad \forall\, U, \widetilde{U} \in L^2(\Omega). \tag{35}$$

Applying (33) to (35) yields a deterministic bound on the feature increment induced by multipath contributions

$$\begin{aligned}\|\Phi(U_{\mathrm{MU}}) - \Phi(U_{\mathrm{MP}})\|_2 &= \|\Phi(U_{\mathrm{MP}} + \Delta U) - \Phi(U_{\mathrm{MP}})\|_2 \\ &\leq L_\Phi \|\Delta U\|_2.\end{aligned} \tag{36}$$

If ΔU is spatially localized or energy-limited as established in Section IV-A, the induced feature shift remains controlled. This result supports stable transfer from MP-RM to MU-RM in the learned representation.

Let U^{MP} and U^{MU} be random fields representing the MP-RM and MU-RM domains under matched transmitter and environment semantics. Define the corresponding feature vectors as

$$\boldsymbol{z}^{\mathrm{MP}} = \Phi\left(U^{\mathrm{MP}}\right), \quad \boldsymbol{z}^{\mathrm{MU}} = \Phi\left(U^{\mathrm{MU}}\right) \tag{37}$$

and denote the domain means by

$$\boldsymbol{\mu}_{\mathrm{MP}} = \mathbb{E}\left[\boldsymbol{z}^{\mathrm{MP}}\right], \quad \boldsymbol{\mu}_{\mathrm{MU}} = \mathbb{E}\left[\boldsymbol{z}^{\mathrm{MU}}\right]. \tag{38}$$

The average feature shift induced by multipath is defined as

$$\boldsymbol{w} \triangleq \boldsymbol{\mu}_{\mathrm{MU}} - \boldsymbol{\mu}_{\mathrm{MP}}. \tag{39}$$

Using (33) and (37), the vector $\boldsymbol{w}$ in (39) admits the equivalent representation

$$\boldsymbol{w} = \mathbb{E}\left[\Phi\left(U^{\mathrm{MP}} + \Delta U\right) - \Phi\left(U^{\mathrm{MP}}\right)\right] \tag{40}$$

which shows that $\boldsymbol{w}$ is the mean feature displacement caused by the multipath residual. The structural implication of Section IV-A is that the MP-RM skeleton is preserved while ΔU

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

injects localized increments. In a representation that captures global geometry and coarse power layout, this observation motivates a translation-dominant approximation in the feature space.

Consider paired samples $\{(U_i^{\mathrm{MP}}, U_i^{\mathrm{MU}})\}_{i=1}^N$ sharing the same underlying configuration, and define

$$z_i^{\mathrm{MP}} = \Phi\left(U_i^{\mathrm{MP}}\right), \quad z_i^{\mathrm{MU}} = \Phi\left(U_i^{\mathrm{MU}}\right). \tag{41}$$

We model the cross-domain relation as

$$z_i^{\mathrm{MU}} = z_i^{\mathrm{MP}} + \boldsymbol{w} + \boldsymbol{\delta}_i \tag{42}$$

$$\mathbb{E}\left[\boldsymbol{\delta}_i\right] = \mathbf{0}, \quad \|\boldsymbol{\delta}_i\|_2 \le \eta \tag{43}$$

where $\boldsymbol{\delta}_i$ captures sample-specific multipath variability not explained by the common direction $\boldsymbol{w}$.

A first-order justification of (42) follows from local linearization. Assume that Φ is Fréchet differentiable in a neighborhood of U_i^{MP}, and write

$$\Phi\left(U_i^{\mathrm{MP}} + \Delta U_i\right) = \Phi\left(U_i^{\mathrm{MP}}\right) + \mathcal{J}_\Phi\left(U_i^{\mathrm{MP}}\right)[\Delta U_i] + \boldsymbol{r}_i \tag{44}$$

where $\mathcal{J}_\Phi(U)$ denotes the Fréchet derivative as a bounded linear operator and $\boldsymbol{r}_i$ is the higher-order remainder. Taking expectations and using (40) yields

$$\boldsymbol{w} = \mathbb{E}\left[\mathcal{J}_\Phi\left(U^{\mathrm{MP}}\right)[\Delta U]\right] + \mathbb{E}\left[\boldsymbol{r}\right]. \tag{45}$$

The translation model (42) then follows by identifying $\boldsymbol{\delta}_i$ with the deviation of $\mathcal{J}_\Phi(U_i^{\mathrm{MP}})[\Delta U_i] + \boldsymbol{r}_i$ from its mean.

Under (42), pairwise relations are stable up to the fluctuation level η. Then, for any i and j, we have

$$z_i^{\mathrm{MU}} - z_j^{\mathrm{MU}} = z_i^{\mathrm{MP}} - z_j^{\mathrm{MP}} + \boldsymbol{\delta}_i - \boldsymbol{\delta}_j. \tag{46}$$

By the reverse triangle inequality

$$\left| \|z_i^{\mathrm{MU}} - z_j^{\mathrm{MU}}\|_2 - \|z_i^{\mathrm{MP}} - z_j^{\mathrm{MP}}\|_2 \right| \le \|\boldsymbol{\delta}_i - \boldsymbol{\delta}_j\|_2 \le \|\boldsymbol{\delta}_i\|_2 + \|\boldsymbol{\delta}_j\|_2 \le 2\eta. \tag{47}$$

Equation (47) formalizes that the relational geometry of MP-RM features is largely preserved in the MU-RM domain. The dominant interdomain difference is captured by the common shift direction $\boldsymbol{w}$.

Given the paired features in (41), define empirical means as

$$\widehat{\boldsymbol{\mu}}_{\mathrm{MP}} = \frac{1}{N}\sum_{i=1}^N z_i^{\mathrm{MP}}, \quad \widehat{\boldsymbol{\mu}}_{\mathrm{MU}} = \frac{1}{N}\sum_{i=1}^N z_i^{\mathrm{MU}} \tag{48}$$

and define the empirical shift direction as

$$\widehat{\boldsymbol{w}} = \widehat{\boldsymbol{\mu}}_{\mathrm{MU}} - \widehat{\boldsymbol{\mu}}_{\mathrm{MP}}. \tag{49}$$

From (42) and (48), it follows that:

$$\widehat{\boldsymbol{w}} = \boldsymbol{w} + \frac{1}{N}\sum_{i=1}^N \boldsymbol{\delta}_i \tag{50}$$

and therefore

$$\|\widehat{\boldsymbol{w}} - \boldsymbol{w}\|_2 \le \frac{1}{N}\sum_{i=1}^N \|\boldsymbol{\delta}_i\|_2 \le \eta. \tag{51}$$

Equation (51) shows that the shift direction can be reliably estimated from a small number of paired samples. This guarantee holds when the translation-dominant approximation is valid with bounded fluctuations. Equation (51) only requires a dominant mean direction with bounded sample-dependent fluctuations, so RadioDiff-FS uses the estimated DPM-to-IRT4 direction as a soft regularizer rather than a deterministic one-to-one translation.

C. Directional Consistency Principle

Let $\mathcal{G}_\theta$ denote a conditional generator that maps an MP-RM input U^{MP} and condition c to a predicted MU-RM output

$$\widehat{U}^{\mathrm{MU}} = \mathcal{G}_\theta\left(U^{\mathrm{MP}}, c\right). \tag{52}$$

Let Φ be the fixed feature encoder defined in (34). The feature displacement induced by the generator is defined as

$$\Delta z \triangleq \Phi\left(\widehat{U}^{\mathrm{MU}}\right) - \Phi\left(U^{\mathrm{MP}}\right). \tag{53}$$

Let $\boldsymbol{w} = \boldsymbol{\mu}_{\mathrm{MU}} - \boldsymbol{\mu}_{\mathrm{MP}}$ be the mean shift direction defined in (39). The central requirement is that the generator update should align with the dominant cross-domain direction. At the same time, the alignment should allow sample-dependent magnitudes that reflect spatially varying multipath increments.

Definition 1: The directionally consistent set associated with $\boldsymbol{w}$ is

$$\mathcal{C}(\boldsymbol{w}) = \left\{\boldsymbol{v} \in \mathbb{R}^d : \ \boldsymbol{v} = \alpha\boldsymbol{w} + \boldsymbol{e}, \ \alpha \ge 0, \ \boldsymbol{e}^\top\boldsymbol{w} = 0\right\} \tag{54}$$

where α determines the amount of multipath injection and $\boldsymbol{e}$ represents orthogonal distortion. Directional consistency requires suppressing the orthogonal component $\boldsymbol{e}$.

To enforce this principle, we decompose Δz into its projection onto $\boldsymbol{w}$ and its orthogonal residual. Define

$$\alpha(\Delta z) = \frac{\Delta z^\top \boldsymbol{w}}{\|\boldsymbol{w}\|_2^2} \tag{55}$$

$$\Delta z_\parallel = \alpha(\Delta z)\, \boldsymbol{w} \tag{56}$$

$$\Delta z_\perp = \Delta z - \Delta z_\parallel. \tag{57}$$

The directional consistency loss penalizes $\Delta z_\perp$ and discourages negative alignment with $\boldsymbol{w}$. For a paired dataset $\{(U_i^{\mathrm{MP}}, U_i^{\mathrm{MU}}, c_i)\}_{i=1}^N$, define

$$\widehat{U}_i^{\mathrm{MU}} = \mathcal{G}_\theta\left(U_i^{\mathrm{MP}}, c_i\right) \tag{58}$$

$$\Delta z_i = \Phi\left(\widehat{U}_i^{\mathrm{MU}}\right) - \Phi\left(U_i^{\mathrm{MP}}\right). \tag{59}$$

The proposed directional consistency loss is

$$\mathcal{L}_{\mathrm{dir}}(\theta) = \frac{1}{N}\sum_{i=1}^N \left\| \Delta z_i - \frac{\Delta z_i^\top \boldsymbol{w}}{\|\boldsymbol{w}\|_2^2}\, \boldsymbol{w} \right\|_2^2 + \frac{\beta}{N}\sum_{i=1}^N \left[\max\left(0, -\Delta z_i^\top \boldsymbol{w}\right)\right]^2 \tag{60}$$

where β controls the penalty for moving opposite to the DPM-to-IRT4 direction. Since $\alpha(\Delta z_i)$ remains sample-dependent, DCL preserves local multipath diversity while suppressing orthogonal and negative feature drifts. The second term enforces nonnegative alignment, which is consistent with the nonnegative multipath increment property established in Section IV-A.

The necessity of a directional constraint follows from the nonidentifiability of distance-based alignment in Euclidean

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

feature spaces. Consider any objective that matches only pairwise distances between generated features and MP-RM features. Such an objective admits solutions related by orthogonal transformations, since for any orthogonal matrix $\boldsymbol{O}$

$$\|\boldsymbol{O}\boldsymbol{a} - \boldsymbol{O}\boldsymbol{b}\|_2 = \|\boldsymbol{a} - \boldsymbol{b}\|_2 \quad \forall\, \boldsymbol{a}, \boldsymbol{b} \in \mathbb{R}^d. \tag{61}$$

A generator can therefore satisfy distance constraints while rotating the feature configuration. Such rotation changes the semantic axes along which the MP-RM structure is represented and can distort the inherited skeleton. In contrast, minimizing (60) suppresses the orthogonal component for every sample. This fixes the cross-domain displacement direction and eliminates the rotational ambiguity.

The directional constraint also yields a stability property consistent with Section IV-B. Under the translation-dominant model, a directionally consistent generator satisfies

$$\Delta \boldsymbol{z}_i = \alpha_i \boldsymbol{w} + \boldsymbol{e}_i, \quad \boldsymbol{e}_i^{\top} \boldsymbol{w} = 0, \quad \|\boldsymbol{e}_i\|_2 \le \varepsilon. \tag{62}$$

Then, for any i and j, we have

$$\boldsymbol{z}_i^{\text{gen}} - \boldsymbol{z}_j^{\text{gen}} = \boldsymbol{z}_i^{\text{MP}} - \boldsymbol{z}_j^{\text{MP}} + \left(\alpha_i - \alpha_j\right) \boldsymbol{w} + \left(\boldsymbol{e}_i - \boldsymbol{e}_j\right) \tag{63}$$

where $\boldsymbol{z}_i^{\text{gen}} = \Phi(\widehat{U}_i^{\text{MU}})$. When $\boldsymbol{e}_i$ is controlled by the first term in (60), the deformation of the relational geometry is confined to a 1-D modulation along $\boldsymbol{w}$. This preserves the dominant MP-RM skeleton encoded by directions orthogonal to $\boldsymbol{w}$ and allocates the cross-domain change to the physically meaningful axis of multipath injection.

The overall fine-tuning objective combines a task fidelity term with the directional consistency regularizer. Let $\mathcal{L}_{\text{fit}}(\theta)$ denote a supervised or self-consistent reconstruction loss between $\widehat{U}_i^{\text{MU}}$ and U_i^{MU} in the MU-RM domain. The final training objective is

$$\min_{\theta} \; \mathcal{L}_{\text{fit}}(\theta) + \lambda_{\text{dir}} \mathcal{L}_{\text{dir}}(\theta) \tag{64}$$

where $\lambda_{\text{dir}} > 0$ controls the regularization strength. This regularizer steers adaptation along the mean shift direction from MP-RM to MU-RM while suppressing orthogonal distortions. It thereby aligns cross-domain statistics without sacrificing the geometric integrity of the RM skeleton. Unless otherwise stated, we set $\lambda_{\text{dir}} = 0.1$, estimate w with 8 mini-batches, and use the stochastic DCL scale $\alpha_s = 20^u$, $u \sim \mathcal{U}(0,1)$, as summarized in Table I.

D. Few-Shot Fine-Tuning

The objective is to adapt a pretrained generative model from the MP-RM domain to the MU-RM domain using only a small number of labeled MU-RM samples. The MP-RM domain provides abundant supervision that captures stable large-scale spatial priors and global propagation skeletons. In contrast, the MU-RM domain exhibits additional spatially localized power increments induced by multipath interactions. The adaptation should preserve the pretrained global structure and inject MU-RM-specific details in a controlled and physically consistent manner. Let c denote the conditioning input that encodes scene geometry and transmitter attributes. The fine-tuning set is composed of paired samples

$$\mathcal{D}_{\text{ft}} = \left\{\left(x_{0,i}^{\text{MP}}, x_{0,i}^{\text{MU}}, c_i\right)\right\}_{i=1}^{N} \tag{65}$$

where N is small and each pair shares the same configuration specified by c_i.

As shown in Fig. 2, we employ a variance-preserving diffusion model to parameterize the conditional generator. Let x_0 denote a clean MU-RM sample and let $q(x_t \mid x_0)$ be the forward noising process indexed by $t \in [0,1]$. The perturbed sample is written as

$$x_t = \alpha(t)\, x_0 + \sigma(t)\, \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, I) \tag{66}$$

where $\alpha(t)$ is a monotonically decreasing signal factor and $\sigma(t)$ is a monotonically increasing noise factor. The denoiser predicts the injected noise ε from the tuple (x_t, t, c). To allow different propagation components to be modeled with separate heads, we write the network output as $\varepsilon_{\theta,r}(x_t, t, c)$ for $r \in \mathcal{R}$ and define ε_r as the corresponding ground-truth noise in (66). The base objective is the denoising score matching loss

$$\mathcal{L}_{\text{base}} = \mathbb{E}_{x_0, t, \{\varepsilon_r\}_{r\in\mathcal{R}}} \left[\sum_{r\in\mathcal{R}} \|\varepsilon_r - \varepsilon_{\theta,r}(x_t, t, c)\|_2^2 \right]. \tag{67}$$

To apply the directional consistency principle to diffusion fine-tuning, we compute L_{dir} in (60) using the predicted clean sample $\hat{x}_{0,\theta}(x_t, t, c)$ as the generated MU-RM and the paired x_0^{MP} as the source-domain reference. Directional regularization should be weak at early denoising stages, where the predicted clean map is still unreliable. We therefore use a time-dependent weight that increases with the signal-to-noise ratio. The SNR at diffusion time t is

$$\text{SNR}(t) = \frac{\alpha(t)^2}{\sigma(t)^2} \tag{68}$$

and the scheduled DCL weight is

$$\lambda_{\text{dir}}(t) = \lambda_{\max} \frac{\text{SNR}(t)}{\text{SNR}(t) + 1}. \tag{69}$$

The final fine-tuning objective is

$$L_{\text{total}} = \mathbb{E}[L_{\text{base}}] + \mathbb{E}[\lambda_{\text{dir}}(t)\, L_{\text{dir}}]. \tag{70}$$

This objective preserves the pretrained MP-RM prior through the diffusion reconstruction loss, while steering the generated feature displacement toward the IRT4 domain through the DCL term.

V. Experimental Results

A. Datasets

We evaluate on RadioMapSeer [34], a large-scale urban benchmark for RM prediction. Each scene is represented by a 256×256 building occupancy grid with 1 m spatial resolution and a unified building height of 25 m. We consider the device-to-device (D2D) setting with transmitter and receiver heights of 1.5 m, a carrier frequency of 5.9 GHz, and a transmitter power of 23 dBm [34]. RadioMapSeer provides subsets with different propagation complexities. We use the multipath-rich subset with up to four interactions, denoted as IRT4 in the dataset release, as the target MU-RM domain. A key characteristic of IRT4 is its limited supervision density: each scene contains only two transmitter placements and thus only two MU-RM samples per scene. This results in 1402 MU-RM

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

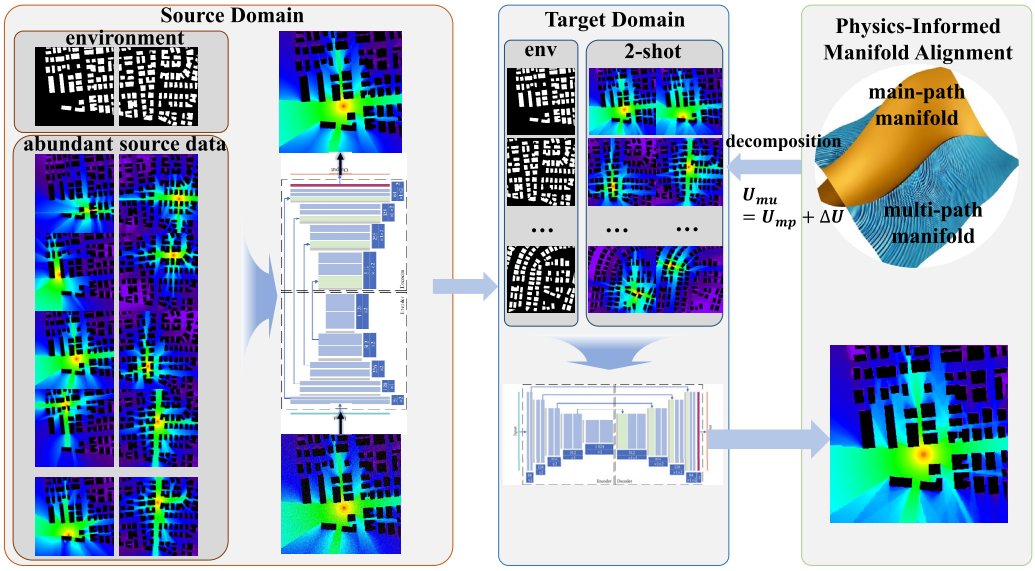


Fig. 2. RadioDiff-FS framework. The model pretrains on abundant MP data (left) and adapts to the multipath target domain using few-shot supervision (middle), whose transfer is enabled by physics-informed manifold alignment.

samples over 701 scenes, which naturally forms a few-shot adaptation setting when fine-tuning a pretrained generator. We evaluate both SRMs that account for building-induced attenuation and dynamic RMs (DRMs) that additionally include vehicle-induced blockage [34]. We adopt a scene-disjoint split with 500 scenes for training and validation and 200 unseen scenes for testing, ensuring generalization to new urban layouts. Unless otherwise stated, few-shot fine-tuning uses only a small subset of the IRT4 training samples, while the remaining samples are used for validation and ablation. The source generator is pretrained on RadioMapSeer DPM with 80 transmitter locations per scene, with DCL disabled. Few-shot adaptation uses IRT4 with two transmitter locations per scene. The frozen DCL encoder is an ImageNet-pretrained InceptionV3, where the 2048-D globally pooled feature is ℓ_2-normalized before shift estimation. The encoder is used only for feature measurement and introduces no trainable parameters. Table I summarizes the implementation.

B. Metrics and Baselines

We evaluate RM reconstruction using pixel-level fidelity and structural consistency metrics, following common practice in RM prediction [14], [17]. Let $Y, \widehat{Y} \in \mathbb{R}^{H\times W}$ denote the ground-truth and predicted maps. We report normalized mean squared error (NMSE) and root mean squared error (RMSE), where smaller values indicate higher accuracy. To quantify structure preservation, we compute the structural similarity index measure (SSIM) [35], and peak signal-to-noise ratio (PSNR). Higher SSIM and PSNR indicate better structural consistency and visual fidelity. All metrics are averaged over the test set.

We compare RadioDiff-FS with representative baselines covering supervised regression, adversarial generation, and diffusion-based generation. To isolate the effect of the proposed few-shot adaptation, all methods are trained from scratch on the same few-shot supervision from the IRT4 subset. Specifically, all methods use 500 training scenes with 2 transmitter placements per scene, and are evaluated on the identical scene-disjoint test split. We additionally report RadioDiff-FS in a one-shot setting, where only a single MU-RM sample per scene is used for fine-tuning, to examine performance under extreme data scarcity. The baselines are summarized as follows.

1) *RadioUNet [17]*: A sampling-free convolutional regressor based on a U-Net backbone that predicts RMs directly from environment and transmitter inputs.
2) *RME-GAN [23]*: An adversarial generator adapted to the sampling-free protocol by using only environmental features and transmitter metadata for direct RM synthesis.
3) *PhyRMDM [28]*: A conditional diffusion framework that generates high-resolution RMs in the image domain, augmented with a physics-alignment loss based on a Helmholtz-equation residual.
4) *RadioDiff [14]*: A diffusion-based sampling-free RM generator that serves as a strong baseline under the same conditioning inputs but without the proposed directional regularization.

RadioDiff is particularly relevant because it belongs to the same backbone family as our approach. It directly maps the environment and transmitter inputs to the RM without the proposed MP to multipath guidance mechanism. It therefore serves as a controlled ablation baseline that isolates the contribution of the directional consistency constraint while keeping the generative capacity unchanged.

C. Performance on SRM

Quantitative results on the SRM are reported in Table II. RadioDiff-FS achieves the best performance across all metrics, attaining an NMSE of 0.0049 and a PSNR of 36.37 dB. Compared with the vanilla RadioDiff baseline at NMSE 0.0121 and PhyRMDM at NMSE 0.0100, RadioDiff-FS reduces the reconstruction error by 59.50%. This result indicates that few-shot training alone is insufficient for high-fidelity multipath reconstruction without additional adaptation guidance.

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

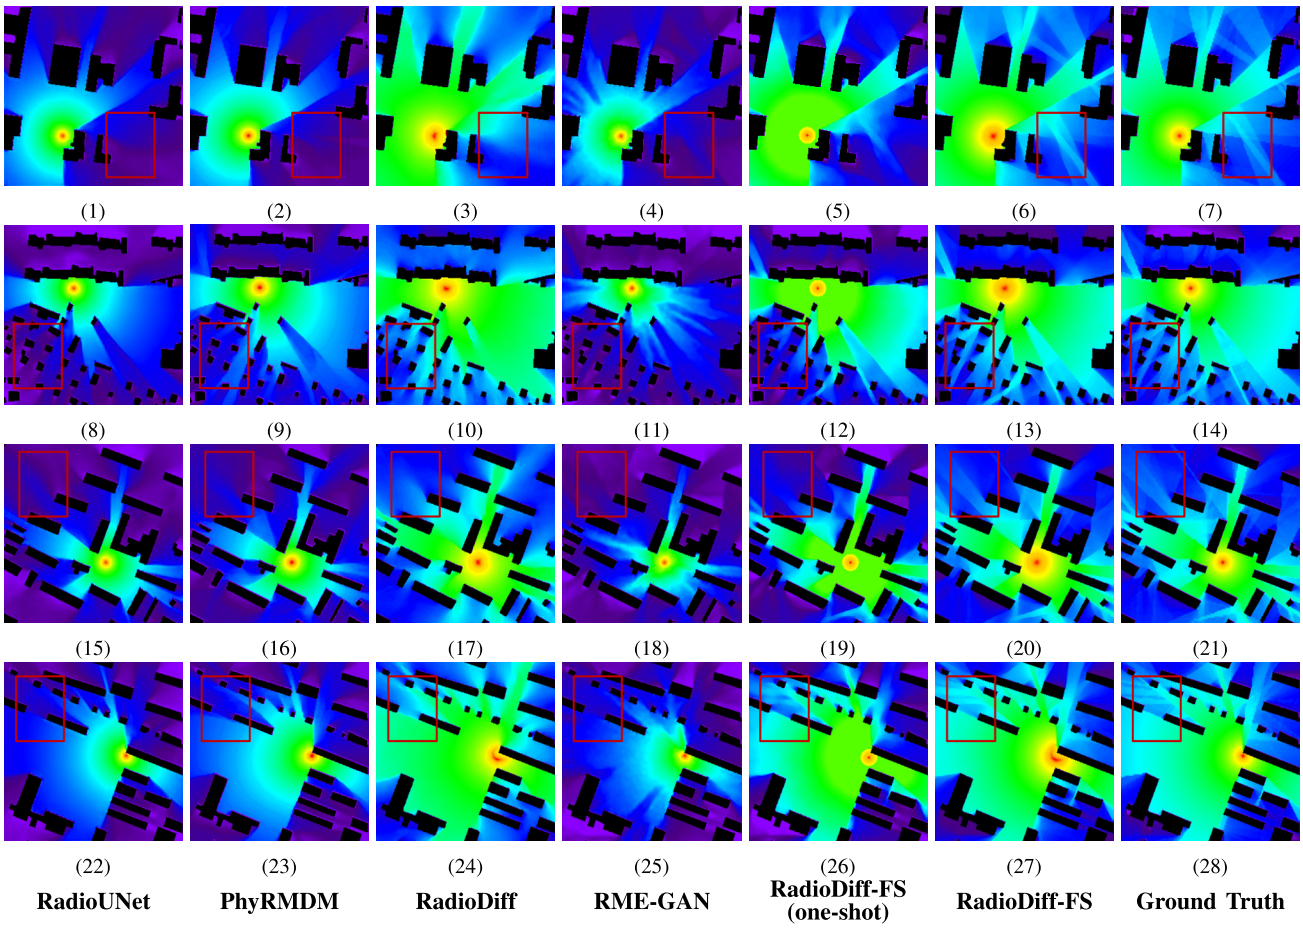


Fig. 3. Visual comparison of constructed SRM across different methods.

TABLE I
IMPLEMENTATION DETAILS OF RADIODIFF-FS

Item	Setting
Source pretraining data	RadioMapSeer DPM, `numTx`=80
Target adaptation data	RadioMapSeer IRT4, `numTx`=2
DCL in source pretraining	Disabled
Fixed feature encoder	ImageNet-pretrained InceptionV3
Feature dimension	2048
Feature normalization	ℓ_2 normalization
Direction estimation	8 mini-batches in the main setting
λ_{dir}	0.1 in the main IRT4 adaptation setting
Progressive training	Enabled
Stage 1 endpoint	2000 steps
Stage 2 endpoint	6000 steps in the main setting
Stochastic DCL scale	$\alpha_s = 20^u,\ u \sim \mathcal{U}(0,1)$

RadioDiff-FS also yields the highest SSIM of 0.9752, which substantially exceeds all other methods and suggests improved preservation of spatial structures, such as diffraction edges and shadow boundaries. The one-shot variant of RadioDiff-FS also demonstrates competitive performance. With only a single MU-RM sample per scene for fine-tuning, RadioDiff-FS achieves the second-best SSIM of 0.9285 and the second-best PSNR of 31.77 dB, achieving competitive structural fidelity among the compared methods, although its pixel-level error remains higher than the full two-shot adaptation setting. This result validates that the directional consistency constraint provides a strong inductive bias even under extreme data scarcity, as it leverages the pretrained MP-RM prior to guide adaptation with minimal supervision.

These quantitative trends are consistent with the qualitative results in Fig. 3. RadioUNet produces smooth attenuation fields that capture coarse pathloss gradients but blur sharp transitions near building silhouettes. These transitions are critical in multipath-dominant settings. RME-GAN generates sharper textures, yet it exhibits intermittent artifacts in low-signal regions where adversarial objectives may amplify visually plausible but physically inconsistent patterns. Diffusion-based baselines improve perceptual realism. However, without an explicit cross-domain constraint, they still under-represent boundary sharpness and local contrast in deep-shadow areas. In contrast, RadioDiff-FS better matches the ground truth both globally and locally. It preserves the large-scale attenuation skeleton while recovering fine-grained transition bands around occluded corridors, aligning with the intended MP to multipath generalization behavior.

D. Performance on DRM

The DRM further introduces time-varying occlusions from vehicles, which produce localized attenuation spots and small-scale shadowing patterns along roads. As shown in Table III, most baselines degrade noticeably when moving from SRM to DRM. PhyRMDM suffers a sharp deterioration with NMSE

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

TABLE II

QUANTITATIVE COMPARISON ON IRT4. RESULTS IN BOLD RED AND UNDERLINED BLUE HIGHLIGHT THE BEST AND SECOND BEST, RESPECTIVELY. THE LAST COLUMN DENOTES THE RELATIVE IMPROVEMENT OF RADIODIFF-FS OVER THE VANILLA RADIODIFF BASELINE

Methods	RadioDiff-FS	RadioDiff-FS (one-shot)	PhyRMDM	RadioDiff	RME-GAN	RadioUNet	Rate vs. RadioDiff (%)
NMSE	**0.0049**	0.0204	0.0100	0.0121	0.0155	0.0159	↓ 59.50%
RMSE	**0.0196**	0.0278	0.0278	0.0309	0.0340	0.0344	↓ 36.57%
SSIM ↑	**0.9752**	0.9285	0.9180	0.9268	0.9123	0.9102	↑ 5.22%
PSNR ↑	**36.37**	31.77	31.40	30.44	29.74	29.64	↑ 19.48%

TABLE III

QUANTITATIVE COMPARISON ON IRT4 WITH CAR. RESULTS IN BOLD RED AND UNDERLINED BLUE HIGHLIGHT THE BEST AND SECOND BEST, RESPECTIVELY. THE LAST COLUMN DENOTES THE PERCENTAGE IMPROVEMENT OF RADIODIFF-FS OVER THE BEST BASELINE WITHOUT FINE-TUNING

Methods	RadioDiff-FS	RadioDiff-FS (one-shot)	PhyRMDM	RadioDiff	RME-GAN	RadioUNet	Rate vs. RadioDiff (%)
NMSE	**0.0121**	0.0214	0.1113	0.0465	0.0183	0.0381	↓ 73.98%
RMSE	**0.0214**	0.0296	0.2178	0.1319	0.0358	0.0307	↓ 83.78%
SSIM ↑	**0.9510**	0.8969	0.7182	0.7968	0.8823	0.8487	↑ 19.35%
PSNR ↑	**32.68**	30.61	25.23	27.86	30.74	29.57	↑ 17.30%

increasing to 0.1113, indicating that physics-alignment alone does not guarantee robustness when the target domain contains sparse and configuration-dependent perturbations under limited supervision. The vanilla RadioDiff baseline also degrades significantly, reflecting that the pretrained generative prior may not correctly allocate capacity to vehicle-induced details when trained with few DRM samples. RadioDiff-FS remains substantially more stable in this regime, achieving an NMSE of 0.0121 and an SSIM of 0.9510, which are markedly better than all baselines. The one-shot variant again delivers strong results, attaining the second-best RMSE of 0.0296 and the second-best SSIM of 0.8969. This confirms that the directional consistency constraint is effective even with a single fine-tuning sample per scene.

The qualitative comparison in Fig. 4 explains this stability. Vehicle occlusions manifest as compact attenuation patches whose locations are coupled to road geometry and line-of-sight corridors. Regression and diffusion baselines tend to smooth these patches into the background, leading to systematic underestimation of localized blockage. RadioDiff-FS more consistently recovers these attenuation signatures and their spatial extent, producing DRM reconstructions that are closer to the ground truth in both contrast and location. This behavior is particularly important for downstream tasks that depend on local reliability margins, where missing a small occlusion-induced dip can lead to incorrect link-quality assessment.

E. Feature-Space Analysis of DPM-to-IRT4 Adaptation

We examine the DPM-to-IRT4 transition using exactly paired samples. A frozen InceptionV3 encoder maps each DPM or IRT4 RM to a 2048-D ℓ_2-normalized feature. We compute the feature displacement for each pair and average all displacements to obtain the global DPM-to-IRT4 shift direction. For visualization, features are projected onto the plane spanned by this mean direction and the leading orthogonal residual direction.

As shown in Fig. 5, the shifted DPM cluster centers move toward their paired IRT4 centers, while the local cluster layout is largely preserved. Across all 398 matched pairs, 79.9% of feature shifts have positive projection on the global mean direction. This supports a dominant but nonexhaustive feature translation. It also justifies using DCL as a soft directional regularizer that suppresses orthogonal and negative drifts while retaining sample-specific multipath deviations.

TABLE IV

PERFORMANCE UNDER DIFFERENT BUILDING DENSITIES

Group	Scenes	Samples	Building Ratio	NMSE	RMSE	SSIM	PSNR
Overall	199	398	0.0017–0.5158	0.0049	0.0196	0.9752	36.37
Sparse	67	134	0.0017–0.1891	0.0032	0.0187	0.9820	37.15
Medium	66	132	0.1893–0.2918	0.0055	0.0210	0.9705	35.85
Dense	66	132	0.2937–0.5158	0.0060	0.0200	0.9710	36.15

F. Performance Under Different Blockage Densities

To examine robustness under more complex blockage conditions, we group test scenes by the building-pixel ratio into sparse, medium, and dense subsets using test-set tertiles. This grouping is a proxy for blockage complexity rather than a complete NLoS characterization. As shown in Table IV, RadioDiff-FS remains stable across all groups. The dense group achieves an NMSE of 0.0060 and an SSIM of 0.9710, which remain close to the overall results. This suggests that the proposed adaptation remains effective under increased blockage density, while more severe coherent NLoS effects are left to future complex-valued channel-map studies.

G. Ablation Study

We ablate the proposed DCL to quantify its contribution under both full fine-tuning and low-rank adaptation

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

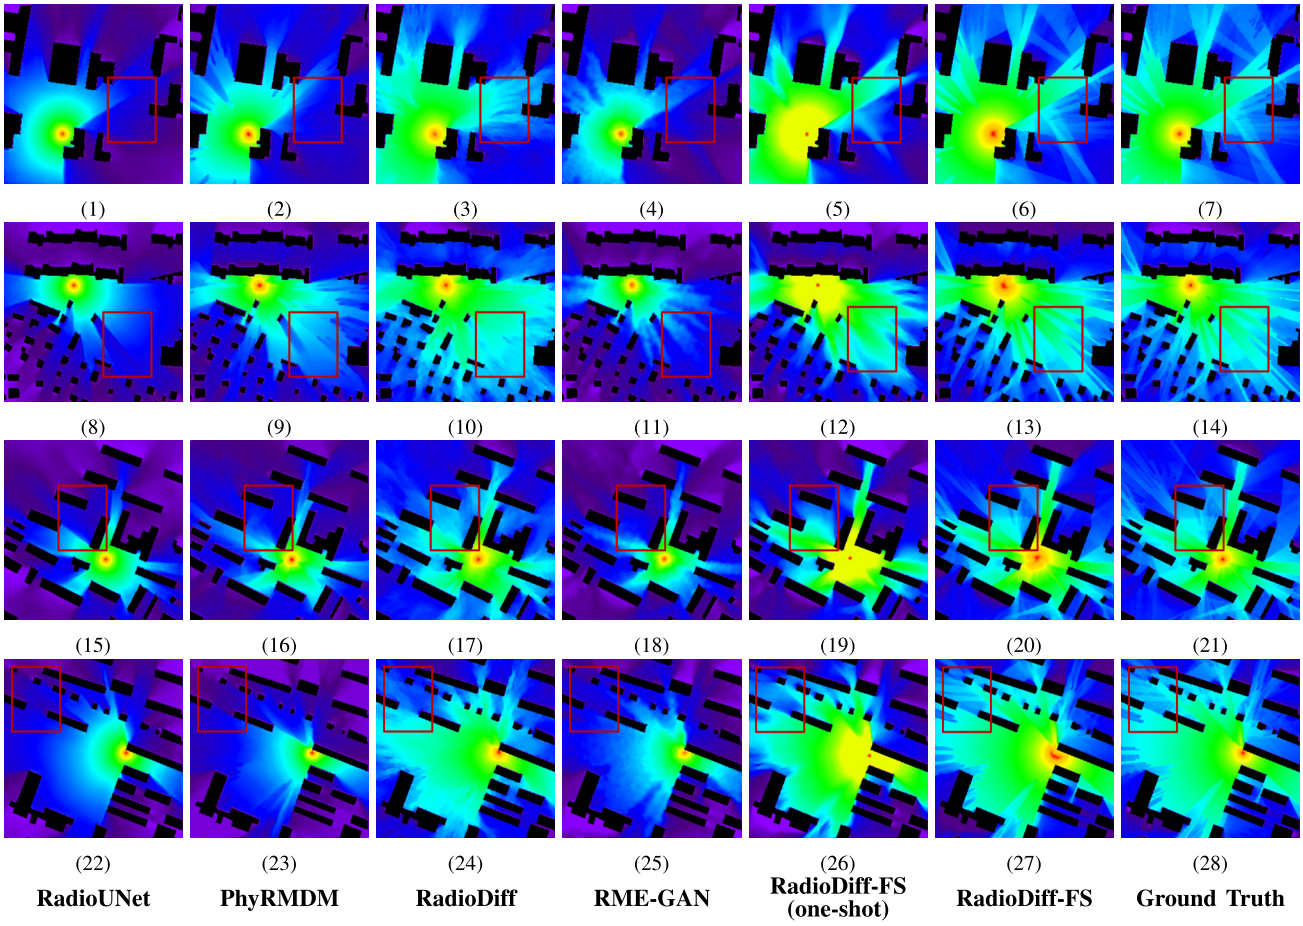


Fig. 4. Visual comparison of constructed DRM across different methods.

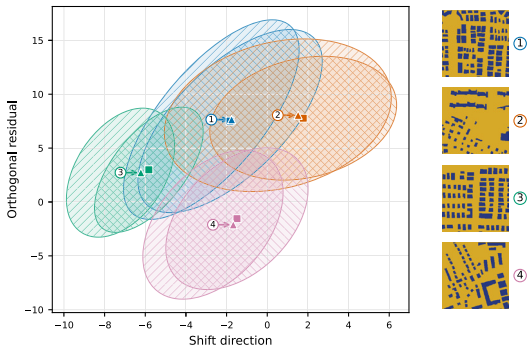


Fig. 5. Feature-shift geometry from DPM to IRT4. The horizontal axis is the global mean shift direction, and the vertical axis is the leading orthogonal residual direction. Circles, triangles, and squares denote DPM centers, shifted DPM centers, and paired IRT4 centers, respectively.

(LoRA). The DCL constrains the feature displacement of generated maps to follow the dominant MP to multipath shift while suppressing orthogonal distortions. This empirically trans lates to sharper boundaries and fewer spurious artifacts in low-signal regions. Table V presents results on the more challenging DRM setting. Removing DCL degrades performance consistently. Under full fine-tuning, adding DCL reduces NMSE from 0.0163 to 0.0121 and improves SSIM from 0.9105 to 0.9510. The SSIM gain indicates that DCL primarily improves structural consistency rather than merely rescaling pixel intensities, which matches the goal of preserving interaction-induced patterns. The improvement remains visible under LoRA, where DCL reduces NMSE from 0.0259

TABLE V

ABLATION STUDY ON IRT4 WITH CAR. COMPARISON OF FULL FINE-TUNING AND LORA STRATEGIES WITH AND WITHOUT DCL

Configuration		NMSE	RMSE	PSNR (dB)	SSIM
Full FT	w/o DCL	0.0163	0.0264	31.79	0.9105
	w/ DCL	**0.0121**	**0.0214**	**32.68**	**0.9510**
LoRA	w/o DCL	0.0259	0.0302	28.88	0.8750
	w/ DCL	0.0195	0.0270	30.65	0.8982

TABLE VI

ABLATION STUDY ON IRT4. RESULTS HIGHLIGHT THE IMPACT OF DCL ACROSS DIFFERENT FINE-TUNING STRATEGIES

Configuration		NMSE	RMSE	PSNR (dB)	SSIM
Full FT	w/o DCL	0.0053	0.0204	35.79	0.9715
	w/ DCL	**0.0049**	**0.0196**	**36.37**	**0.9752**
LoRA	w/o DCL	0.0069	0.0233	31.88	0.9550
	w/ DCL	0.0064	0.0225	32.65	0.9582

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

TABLE VII
INFERENCE LATENCY OF RADIODIFF-FS

Metric	Mean	Std
Compute time for 398 RMs	85.09 s	2.14 s
Time / RM	0.2138 s	0.0054 s
Time / configured diffusion step	0.0428 s	0.0011 s
Time / effective reverse step	0.0356 s	0.0009 s
Throughput	4.62 samples/s	0.18 samples/s

to 0.0195 and improves SSIM from 0.8750 to 0.8982. This suggests that DCL is complementary to parameter-efficient adaptation and helps concentrate limited update capacity on multipath-relevant corrections.

Similar trends are observed on SRM in Table VI. Although SRM is less challenging than DRM, DCL still yields consistent gains for both adaptation strategies. Under full fine-tuning, DCL improves PSNR from 35.79 to 36.37 dB and SSIM from 0.9715 to 0.9752. Under LoRA, the gains follow the same direction with PSNR improving from 31.88 to 32.65 dB. Taken together, these ablations confirm that the advantage of RadioDiff-FS does not stem solely from additional training steps. The proposed directional constraint provides a principled and effective inductive bias for few-shot generalization from MP structure to multipath detail.

H. Computational Cost

Timing is measured on a single GPU with batch size one and includes reverse denoising and VAE decoding only. As shown in Table VII, RadioDiff-FS takes 0.2138 s per RM and 0.0428 s per configured diffusion step. DCL is used only during fine-tuning, so it does not increase online inference latency. The fine-tuning setting uses four GPUs with an effective global batch size of 32 and takes 1.16 s per optimization step.

VI. CONCLUSION

We have presented RadioDiff-FS, a few-shot adaptation framework for multipath RM construction that transfers a pretrained MP prior to multipath-rich environments via a direction-consistency constraint in a frozen feature space. Experiments have demonstrated that the proposed adaptation improves both pixel-level accuracy and structural fidelity on static and dynamic RMs, which can translate into more reliable coverage assessment and blockage-aware resource planning in practical communication systems under scarce high-fidelity supervision. Future work will extend the framework to multifrequency and multiantenna RM targets and validate the adaptation with real-world measurement-driven datasets.

REFERENCES

[1] X. Wang et al., "A tutorial on learning-based radio map construction: Data, paradigms, and physics-awareness," 2026, *arXiv:2603.17499*.

[2] X. Shen et al., "Toward immersive communications in 6G," *Frontiers Comput. Sci.*, vol. 4, Jan. 2023, Art. no. 1068478.

[3] Z. Wang et al., "A tutorial on extremely large-scale MIMO for 6G: Fundamentals, signal processing, and applications," *IEEE Commun. Surveys Tuts.*, vol. 26, no. 3, pp. 1560–1605, 3rd Quart., 2024.

[4] Y. Yang et al., "6G network AI architecture for everyone-centric customized services," *IEEE Netw.*, vol. 37, no. 5, pp. 71–80, Sep. 2023.

[5] N. Cheng et al., "Channel knowledge map-enabled 6D movable antenna systems with kinematic constraints: A manifold optimization approach," *IEEE Trans. Wireless Commun.*, vol. 25, pp. 8968–8981, 2026.

[6] X. Wang et al., "RadioDiff-inverse: Diffusion enhanced Bayesian inverse estimation for ISAC radio map construction," *IEEE Trans. Wireless Commun.*, vol. 25, pp. 14611–14626, 2026.

[7] X. Wang et al., "RadioDiff-k^2: Helmholtz equation informed generative diffusion model for multi-path aware radio map construction," *IEEE J. Sel. Areas Commun.*, vol. 44, pp. 2318–2333, 2025.

[8] G. A. Deschamps, "Ray techniques in electromagnetics," *Proc. IEEE*, vol. 60, no. 9, pp. 1022–1035, Sep. 1972.

[9] X. Wang et al., "Millimeter wave communication: A comprehensive survey," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 3, pp. 1616–1653, 3rd Quart., 2018.

[10] B. Liu, X. Liu, S. Gao, X. Cheng, and L. Yang, "Foundation model for intelligent wireless communications," 2025, *arXiv:2511.22222*.

[11] T. Yang et al., "WirelessGPT: A generative pre-trained multi-task learning framework for wireless communication," *IEEE Netw.*, vol. 39, no. 5, pp. 58–65, Sep. 2025.

[12] X. Wang et al., "Joint flying relay location and routing optimization for 6G UAV–IoT networks: A graph neural network-based approach," *Remote Sens.*, vol. 14, no. 17, p. 4377, Sep. 2022.

[13] D. S. Jones, *The Theory of Electromagnetism*. Amsterdam, The Netherlands: Elsevier, 2013.

[14] X. Wang, "RadioDiff: An effective generative diffusion model for sampling-free dynamic radio map construction," *IEEE Trans. Cognit. Commun. Netw.*, vol. 11, no. 2, pp. 738–750, Apr. 2024.

[15] V. N. Vapnik et al., *Statistical Learning Theory*. New York, NY, USA: Wiley, 1998.

[16] X. Wang et al., "RadioDiff-turbo: Lightweight generative large electromagnetic model for wireless digital twin construction," in *Proc. IEEE INFOCOM - IEEE Conf. Comput. Commun. Workshops (INFOCOM WKSHPS)*, May 2025, pp. 1–6.

[17] R. Levie, Ç. Yapar, G. Kutyniok, and G. Caire, "RadioUNet: Fast radio map estimation with convolutional neural networks," *IEEE Trans. Wireless Commun.*, vol. 20, no. 6, pp. 4001–4015, Jun. 2021.

[18] Y. Tian, S. Yuan, W. Chen, and N. Liu, "Transformer based radio map prediction model for dense urban environments," in *Proc. 13th Int. Symp. Antennas, Propag. EM Theory (ISAPE)*, vol. Volume1, Dec. 2021, pp. 1–3.

[19] X. Wang et al., "RadioDiff-flux: Efficient radio map construction via generative denoise diffusion model trajectory midpoint reuse," *IEEE Trans. Cognit. Commun. Netw.*, vol. 12, pp. 4882–4895, 2026.

[20] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.

[21] G. Wolfle, B. E. Gschwendtner, and F. M. Landstorfer, "Intelligent ray tracing—A new approach for field strength prediction in microcells," in *Proc. IEEE 47th Veh. Technol. Conf. Technol. Motion*, vol. 2, Apr. 1997, pp. 790–794.

[22] B. Zhang et al., "MetaDiff: Meta-learning with conditional diffusion for few-shot learning," in *Proc. AAAI Conf. Artif. Intell.*, 2024, vol. 38, no. 15, pp. 16687–16695.

[23] S. Zhang, A. Wijesinghe, and Z. Ding, "RME-GAN: A learning framework for radio map estimation based on conditional generative adversarial network," *IEEE Internet Things J.*, vol. 10, no. 20, pp. 18016–18027, Oct. 2023.

[24] A. Chaves-Villota and C. A. Viteri-Mera, "DeepREM: Deep-learning-based radio environment map estimation from sparse measurements," *IEEE Access*, vol. 11, pp. 48697–48714, 2023.

[25] G. Chen, Y. Liu, T. Zhang, J. Zhang, X. Guo, and J. Yang, "A graph neural network based radio map construction method for urban environment," *IEEE Commun. Lett.*, vol. 27, no. 5, pp. 1327–1331, May 2023.

[26] X. Li, H. Li, X. Li, G. Zhu, N. Qi, and M. Xiao, "RadioGAT: A model-based learning framework for radio map reconstruction via graph attention networks," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Apr. 2025, pp. 1–6.

[27] X. Wang et al., "RadioDiff-3D: A 3D× 3D radio map dataset and generative diffusion based benchmark for 6G environment-aware communication," *IEEE Trans. Netw. Sci. Eng.*, vol. 13, pp. 3773–3789, 2026.

[28] H. Jia et al., "Physics-informed representation alignment for sparse radio-map reconstruction," in *Proc. 33rd ACM Int. Conf. Multimedia*, Oct. 2025, pp. 12352–12360.

[29] R. Wahl, G. Wölfle, P. Wertz, P. Wildbolz, and F. Landstorfer, "Dominant path prediction model for urban scenarios," in *Proc. 14th IST Mobile Wireless Commun. Summit*, 2005, pp. 1–5.

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.

[30] R. Bhalla and H. Ling, "Three-dimensional scattering center extraction using the shooting and bouncing ray technique," *IEEE Trans. Antennas Propag.*, vol. 44, no. 11, pp. 1445–1453, Nov. 1996.
[31] T. Rautiainen, G. Wolfle, and R. Hoppe, "Verifying path loss and delay spread predictions of a 3D ray tracing propagation model in urban environment," in *Proc. IEEE 56th Veh. Technol. Conf.*, vol. 4, Apr. 2002, pp. 2470–2474.
[32] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 6840–6851.
[33] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," 2020, *arXiv:2011.13456*.
[34] Ç. Yapar, F. Jaensch, R. Levie, G. Kutyniok, and G. Caire, "The first pathloss radio map prediction challenge," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Aug. 2023, pp. 1–2.
[35] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.



Xiucheng Wang (Graduate Student Member, IEEE) is currently pursuing the Ph.D. degree with Xidian University, Xi'an, China.

His research area of interest is machine learning of the wireless network.



Zixuan Guo is currently pursuing the bachelor's degree with Xidian University, Xi'an, China.

Her research interests include diffusion models and radio maps.



Nan Cheng (Senior Member, IEEE) received the B.E. and M.S. degrees from the Department of Electronics and Information Engineering, Tongji University, Shanghai, China, in 2009 and 2012, respectively, and the Ph.D. degree from the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, Canada, in 2016.

He was a Post-Doctoral Fellow with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, Canada, from 2017 to 2019. He is currently a Professor with the State Key Laboratory of ISN and the School of Telecommunications Engineering, Xidian University, Xi'an, Shaanxi, China. He has published over 90 journal articles in IEEE Transactions and other top journals. His current research focuses on B5G/6G, AI-driven future networks, and space-air-ground integrated networks.

Prof. Cheng serves as an Associate Editor for IEEE TRANSACTIONS ON VEHICLE TECHNOLOGY, IEEE OPEN JOURNAL OF THE COMMUNICATION SOCIETY, and *Peer-to-Peer Networking and Applications*, and serves/served as a guest editor for several journals.



Zhisheng Yin (Member, IEEE) received the B.E. degree from Wuhan Institute of Technology, Wuhan, China, in 2012, the B.B.A. degree from the Zhongnan University of Economics and Law, Wuhan, in 2012, the M.Sc. degree from the Civil Aviation University of China, Tianjin, China, in 2016, and the Ph.D. degree from the School of Electronics and Information Engineering, Harbin Institute of Technology, Harbin, China, in 2020.

From September 2018 to September 2019, he visited the BBCR Group, Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, Canada. He is currently an Associate Professor with the School of Telecommunication Engineering, Xidian University, Xi'an, China. His research interests include space-air-ground integrated networks, wireless communications, cyber twin, and physical layer security.

Dr. Yin is also an Associate Editor of IEEE INTERNET OF THINGS JOURNAL.



Ruijin Sun (Member, IEEE) received the Ph.D. degree from Beijing University of Posts and Telecommunications, Beijing, China, in 2019.

From September 2017 to September 2018, she was a visiting student with the University of Waterloo. Waterloo, ON, Canada. From 2019 to 2021, she was a Joint Post-Doctoral Fellow with the Peng Cheng Laboratory, Shenzhen, China, and Tsinghua University, Beijing. She is currently a Lecturer with the School of Telecommunications Engineering and the State Key Laboratory of Integrated Services Networks, Xidian University, Xi'an, China. Her research interests include knowledge-driven wireless resource allocation.



Xuemin (Sherman) Shen (Fellow, IEEE) received the Ph.D. degree in electrical engineering from Rutgers University, New Brunswick, NJ, USA, in 1990.

He is a University Professor with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, Canada. His research focuses on network resource management, wireless network security, Internet of Things, 5G and beyond, and vehicular ad hoc and sensor networks.

Dr. Shen is a registered Professional Engineer of Ontario, Canada; an Engineering Institute of Canada Fellow; a Canadian Academy of Engineering Fellow; a Royal Society of Canada Fellow; a Chinese Academy of Engineering Foreign Member; and a Distinguished Lecturer of the IEEE Vehicular Technology Society and Communications Society. He received the R.A. Fessenden Award from IEEE Canada in 2019, the Award of Merit from the Federation of Chinese Canadian Professionals (Ontario) in 2019, the James Evans Avant Garde Award from the IEEE Vehicular Technology Society in 2018, the 2015 Joseph LoCicero Award and the 2017 Education Award from the IEEE Communications Society, the Technical Recognition Award from Wireless Communications Technical Committee in 2019, and the AHSN Technical Committee in 2013. He has also received the Excellent Graduate Supervision Award in 2006 from the University of Waterloo and the Premier's Research Excellence Award (PREA) from the Province of Ontario, Canada, in 2003. He served as the Technical Program Committee Chair/Co-Chair for IEEE Globecom'07, IEEE VTC'10 Fall, IEEE Infocom'14, and IEEE Globecom'16, and the Chair for the IEEE Communications Society Technical Committee on Wireless Communications. He is the Elected President of the IEEE Communications Society. He was the Vice President for Technical and Educational Activities, the Vice President for Publications, a Member-at-Large on the Board of Governors, the Chair of the Distinguished Lecturer Selection Committee, and a member of Selection Committee of ComSoc. He served as the Editor-in-Chief for IEEE INTERNET OF THINGS JOURNAL, *IEEE Network*, and *IET Communications*.

Authorized licensed use limited to: Nanjing University. Downloaded on October 30,2026 at 06:11:51 UTC from IEEE Xplore. Restrictions apply.