

RaLD: Generating High-Resolution 3D Radar Point Clouds with Latent Diffusion

Ruijie Zhang¹, Bixin Zeng¹, Shengpeng Wang¹, Fuhui Zhou², Wei Wang^{3*}

¹Huazhong University of Science and Technology,

²Nanjing University of Aeronautics and Astronautics,

³Wuhan University

{ruijiezhang, bixinzeng, wsp666}@hust.edu.cn, zhoufuhui@ieee.org, wangw@whu.edu.cn

Abstract

Millimeter-wave radar offers a promising sensing modality for autonomous systems thanks to its robustness in adverse conditions and low cost. However, its utility is significantly limited by the sparsity and low resolution of radar point clouds, which poses challenges for tasks requiring dense and accurate 3D perception. Despite that recent efforts have shown great potential by exploring generative approaches to address this issue, they often rely on dense voxel representations that are inefficient and struggle to preserve structural detail. To fill this gap, we make the key observation that latent diffusion models (LDMs), though successful in other modalities, have not been effectively leveraged for radar-based 3D generation due to a lack of compatible representations and conditioning strategies. We introduce RaLD, a framework that bridges this gap by integrating scene-level frustum-based LiDAR autoencoding, order-invariant latent representations, and direct radar spectrum conditioning. These insights lead to a more compact and expressive generation process. Experiments show that RaLD produces dense and accurate 3D point clouds from raw radar spectrums, offering a promising solution for robust perception in challenging environments.

Introduction

Millimeter-wave radar has attracted growing interest in a range of applications, including perception, localization, and mapping, owing to its affordability, strong penetration capability, and robustness under adverse weather and unilluminated conditions (Lu et al. 2020; Gao et al. 2022; Fan et al. 2024b). These advantages make radar an indispensable sensing modality for autonomous systems operating in complex real-world environments.

However, the utility of radar remains limited by the inherent sparsity and low resolution of its point cloud outputs, which stem from fundamental hardware constraints. This poses significant challenges for downstream tasks that require dense and accurate spatial information. To address this issue, there is a pressing need for effective methods that can enhance or reconstruct high-quality point clouds from raw radar measurements, enabling more reliable and precise environmental understanding.

*Corresponding author.

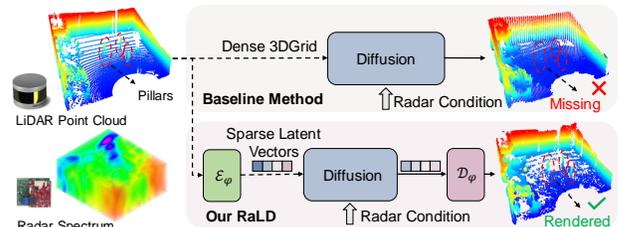


Figure 1: RaLD uses latent diffusion with sparse vectors, enabling high-fidelity 3D radar point cloud generation beyond dense-grid methods.

Diffusion models have recently shown remarkable success in cross-modal generation tasks, such as text-to-image synthesis (Rombach et al. 2022; Peebles and Xie 2023), inspiring their application in radar-based 3D scene understanding. In the context of 3D radar point cloud generation, they offer a promising path to generate high-resolution, LiDAR-like point clouds from sparse and noisy radar measurements.

Despite their success in image domains, conventional diffusion models struggle to scale effectively to 3D point clouds, where dense representations incur significant computational and resolution limitations. While prior work such as SDDiff (Wang et al. 2025) adopts a dense voxel-based representation for radar-conditioned diffusion, it suffers from high computational cost and memory consumption, which further compromises the achievable resolution. For example, SDDiff predicts occupancies over a dense $128 \times 128 \times 64$ voxel grid, yet the actual number of generated points remains on the order of 1%—leading to substantial inefficiency and loss of fine-grained structural details.

In contrast, we observe that latent diffusion models (LDMs) (Rombach et al. 2022) offer a more efficient and flexible alternative. By operating in a compact latent space learned via a dedicated autoencoder, LDMs alleviate the burden of modeling sparse and unordered 3D point clouds directly in data space. This two-stage framework not only enables higher-resolution generation with reduced resource demands but also better preserves structural details, making it well-suited for the challenges of radar-based 3D point cloud super-resolution.

While latent diffusion offers a promising direction, applying it to radar-conditioned 3D point cloud generation poses several technical challenges. First, a key requirement is a robust autoencoder that can compress high-resolution, scene-level LiDAR point clouds into informative latent codes. Most existing point cloud autoencoders are designed for object-level shapes (Zhang et al. 2023) and fall short in handling the sparse, large-scale, and geometrically complex nature of LiDAR data. Second, raw radar spectrums are extremely noisy and lack strong semantic cues, especially compared to LiDAR. These spectrums are affected by multipath effects, specularities, resulting in ambiguous or aliased observations. Conditioning generation on such data is non-trivial, as naive approaches (e.g., radar point-based inputs) often fail to provide sufficient guidance for high-quality reconstruction.

To overcome these challenges, we propose a set of complementary designs that enable effective radar-to-LiDAR point cloud generation within the latent diffusion framework. We first introduce a point-based LiDAR autoencoder that compresses sparse scene-level point clouds into compact latent vectors and reconstructs them by predicting occupancy at arbitrary 3D query locations. Instead of using Cartesian voxel grids, we define occupancy labels based on volumetric frustums aligned with the LiDAR’s polar sampling geometry. This approach preserves the spatial structure of LiDAR returns, captures their non-uniform density across depth, and enables physically grounded interpolation during decoding. Building on this, we design an order-invariant latent encoding strategy that fuses local and global geometric features into a hybrid token representation, ensuring the latent codes are both expressive and order-invariant—crucial properties for stable and generalizable diffusion modeling. Finally, instead of conditioning on sparse radar point clouds, we leverage the richer, albeit noisy, radar spectrum. Our radar spectrum guidance module extracts spatial cues from the signal, providing strong priors for generation and improving the robustness and efficiency of decoding in sparse or ambiguous environments.

Through this integrated design, RaLD enables end-to-end generation of dense 3D structures from raw radar signals, offering a practical and scalable solution for robust 3D perception in adverse environments. In summary, this paper makes the following contributions:

- To our knowledge, this is the first study to explore a sparse, point-based latent diffusion model for radar spectrum-conditioned 3D point cloud generation, effectively enhancing the fidelity and structural quality of point clouds generated from raw radar spectrums.
- We develop a set of complementary designs tailored for radar-conditioned latent diffusion, including a frustum-based LiDAR autoencoder that preserves the polar geometry of LiDAR scans, an order-invariant latent encoding strategy that fuses local and global geometric features, and a radar spectrum guidance module that extracts semantic and spatial cues from raw radar signals to guide generation. Together, these components enable robust and high-resolution 3D point cloud synthesis from

sparse and ambiguous radar inputs.

- Compressive experiments on ColoRadar dataset demonstrate that our method significantly outperforms existing methods in terms of point cloud quality, achieving state-of-the-art results in radar 3D point cloud generation.

Related Work

Traditional Radar Super-Resolution. Extracting dense and reliable point clouds from low-resolution radar data has been a long-standing challenge in the field of radar signal processing. Traditional methods for radar super-resolution often rely on signal processing techniques, such as MUSIC (1986), and ESPRIT (1990), collaborating with a Constant False Alarm Rate (CFAR) detection strategy (Richards et al. 2005) to generate point clouds from radar spectrums. More recent approaches (Qian, He, and Zhang 2020; Lai et al. 2024; Geng et al. 2024) have focused on leveraging synthetic aperture radar (SAR) techniques (Soumekh 1999) to enhance the resolution of radar images. However, SAR-based methods typically rely on predefined moving trajectories and precise estimation of the platform’s motion, which limits their application scenarios.

Radar Super-Resolution with Generative Models. Early works (Guan et al. 2020; Sun et al. 2021; Cheng et al. 2022; Fan et al. 2024a) on radar super-resolution mainly used GANs (Goodfellow et al. 2020) with accumulated radar spectrums to generate high-resolution point clouds (Guan et al. 2020; Cheng et al. 2022; Prabhakara et al. 2022). However, GANs suffer from unstable training (Kurach et al. 2019), limiting their reliability. Recent success of diffusion models (Ho, Jain, and Abbeel 2020) has sparked interest in their use for this task. These methods typically use high-quality LiDAR point clouds as ground truth and train generative models that map radar data to LiDAR-like outputs. Zhang et al. (2024b) first introduced diffusion models for radar super-resolution, proposing an efficient framework to generate high-resolution point clouds in edge devices with limited computational resources. However, it is limited to generating point clouds in 2D, which leads to a significant loss of information in the 3D space. Luan et al. (2024) further extended to 3D point cloud generation model in bird’s-eye view (BEV) space, though BEV still poorly captures heights geometry. SDDiff (Wang et al. 2025) proposed a directional diffusion model, which starts from the radar prior distribution and diffuses toward the 3D LiDAR point clouds. Although both Luan et al. (2024) and SDDiff (2025) are capable of generating 3D point clouds, their resolution remains limited. This is largely due to their reliance on dense representations—such as BEV images or 3D voxel grids—as the generation target, which are computationally expensive and restrict output fidelity. In contrast, our method proposes a latent diffusion model that operates in a lower-dimensional latent space, significantly improving computational efficiency and enabling the generation of high-resolution 3D point clouds.

3D Latent Diffusion Models. Vision generation has been revolutionized by diffusion models, which have shown

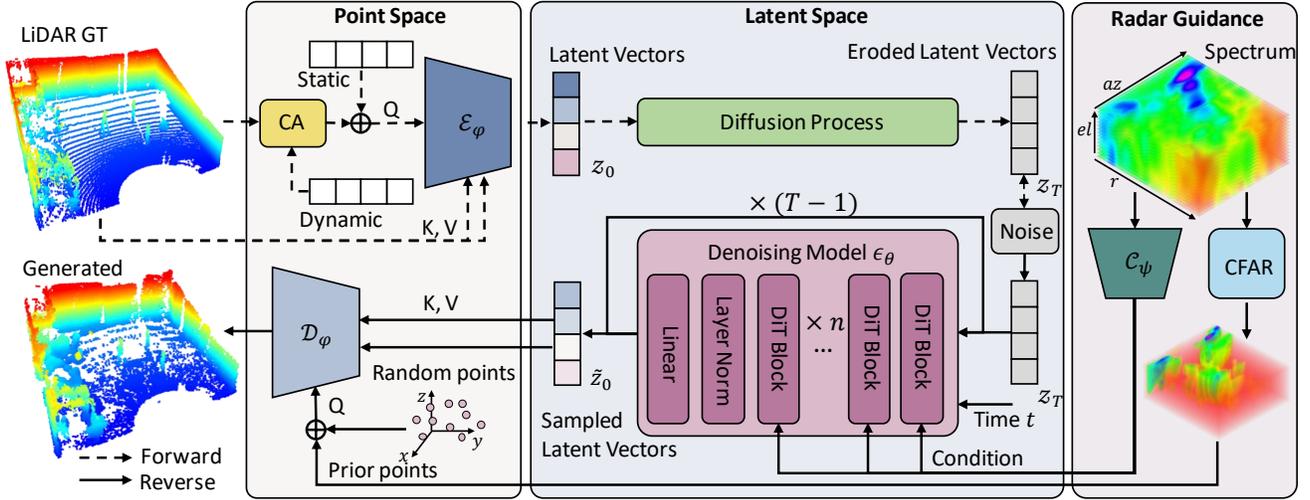


Figure 2: Framework of RaLD pipeline.

remarkable performance in generating high-quality images (Ho, Jain, and Abbeel 2020). Stable Diffusion (Rombach et al. 2022) has further advanced the field by introducing a latent diffusion model that operates in a lower-dimensional latent space, significantly improving computational efficiency and enabling the generation of high-resolution images. Following the success of 2D latent diffusion models, researchers have explored their application in 3D content generation. 3DShape2Vecset (Zhang et al. 2023) introduced an effective representation for object-level 3D shapes in the latent space of 1D vectors, which serves as a foundation 3D variational autoencoder (VAE) for the following 3D latent diffusion models (Zhao et al. 2023; Zhang et al. 2024a; Li et al. 2024). However, these methods primarily focus on object-level 3D generation, which limits their applicability in generating scene-level 3D content. Moreover, these models are designed to generate 3D shapes on condition of text or images, leaving an unexplored gap in generating 3D LiDAR-like point clouds conditioned on radar spectrums.

Methodology

In this section, we first provide a brief overview of latent diffusion models as a preliminary. Then, we present the overall system architecture of RaLD, followed by design details of the frustum-based LiDAR autoencoder, the order-invariant latent encoding, and the radar spectrum guidance.

Preliminary

Diffusion Models Diffusion models are a class of generative models that learn to generate data by reversing a diffusion process (Ho, Jain, and Abbeel 2020). The process begins with a simple distribution, such as Gaussian noise, and gradually transforms it into a complex data distribution through a series of steps.

The forward diffusion process is defined as a Markov

chain that adds Gaussian noise to the data at each step:

$$q(x_t | x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I}), t = 1, \dots, T \quad (1)$$

where x_0 is the original data, x_T is the final noisy data, and β_t controls the noise schedule. The reverse process is parameterized by a neural network and aims to recover the original data by removing noise step by step:

$$p_\theta(x_{t-1} | x_t, c) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t, c), \Sigma_\theta(x_t, t, c)), \quad (2)$$

where μ_θ and Σ_θ are the mean and variance predicted by the network, and c denotes optional conditioning information, such as labels, text, or radar signals. The model is trained to predict the added noise, typically using a mean squared error loss.

Latent Diffusion Models Latent Diffusion Models (LDMs) (Rombach et al. 2022) extend diffusion models by operating in a lower-dimensional latent space instead of directly in the data space. For example, a pre-trained autoencoder ($\mathcal{E}_\varphi, \mathcal{D}_\varphi$) is used to compress high-dimensional point cloud data $x \in \mathbb{R}^{N \times 3}$ into a compact latent representation $z \in \mathbb{R}^{M \times d}$, where N is the number of points, d is the dimension of channel, and M is the number of latent points (typically $M \ll N$):

$$z = \mathcal{E}_\varphi(x), \quad x \approx \mathcal{D}_\varphi(z), \quad (3)$$

where \mathcal{E}_φ and \mathcal{D}_φ denote the encoder and decoder, respectively.

The diffusion model is then trained in the latent space by perturbing z and learning to reverse the corruption:

$$z_t = \sqrt{\alpha_t}z_0 + \sqrt{1 - \alpha_t}\epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (4)$$

$$\mathcal{L}_{\text{LDM}} = \mathbb{E}_{z_t, \epsilon, t} \left[\|\epsilon - \epsilon_\theta(z_t, t)\|^2 \right]. \quad (5)$$

Once denoised, the final latent z_0 is decoded to reconstruct the data $x \approx \mathcal{D}(z_0)$.

System Overview

Given an input radar spectrum \mathbf{S} , RaLD aims to generate a dense and accurate 3D point clouds $\mathbf{P} \in \mathbb{R}^{N \times 3}$ that reconstruct the scene with LiDAR-like fidelity. Following prior works (Zhang et al. 2024b; Wang et al. 2025), we adopt a conditional diffusion framework that learns to synthesize point clouds conditioned on radar observations.

To achieve this, RaLD operates in a compact latent space, where a diffusion model is trained to generate point cloud embeddings guided by the radar spectrum. The overall pipeline, as illustrated in Figure 2, begins with an autoencoder that compresses LiDAR point clouds into structured latent codes. A radar-conditioned latent diffusion model then samples from this space, and a decoder reconstructs the final 3D point clouds guided by radar priors.

We next present the key designs that collectively realize RaLD’s capability: a frustum-based LiDAR autoencoder, an order-invariant latent representation, and a radar-aware generation strategy. Together, these components form a cohesive system that maps radar signals to detailed LiDAR-like point clouds.

Frustum-Based LiDAR Autoencoder

While latent diffusion enables efficient generation of 3D point clouds, designing an effective autoencoder remains challenging due to the extreme sparsity and irregular structure of scene-level LiDAR data. To address this, we propose a tailored autoencoder architecture that leverages the geometric characteristics of LiDAR point clouds.

Inspired by prior vector set-based representations for 3D shapes (Zhang et al.), we compress scene-level LiDAR point clouds into a set of compact latent vectors, and reconstruct them by predicting the occupancy of query points in space. By treating the decoder as a continuous interpolation function for occupancy prediction, this design allows high-fidelity reconstruction with a flexible number of output points.

A key challenge lies in defining the occupancy label for each query point, as required by the decoder. A common approach, widely adopted in perception tasks such as 3D object detection (Zhou and Tuzel 2018; Yin, Zhou, and Krahenbuhl 2021), defines occupancy and extracts features using voxel grids in Cartesian space, where a query point is considered occupied if its voxel contains at least one LiDAR return. However, it fails to capture the sensing characteristics of LiDAR. In particular, LiDAR sensors emit beams at fixed angular resolutions, resulting in a non-uniform point distribution—dense in the near field and sparse at greater distances. Cartesian voxelization ignores this property, imposing a uniform spatial grid that does not align with how the data is captured.

To better align with LiDAR sensing geometry, we define occupancy using frustums in polar coordinates as shown in Figure 3. These grids respect the angular sampling pattern of the sensor and preserve spatial regularity across depth. Formally, a frustum $\mathcal{F}_{i,j,k}$ is defined as a volumetric cell in polar coordinate space, bounded by range, azimuth, and elevation

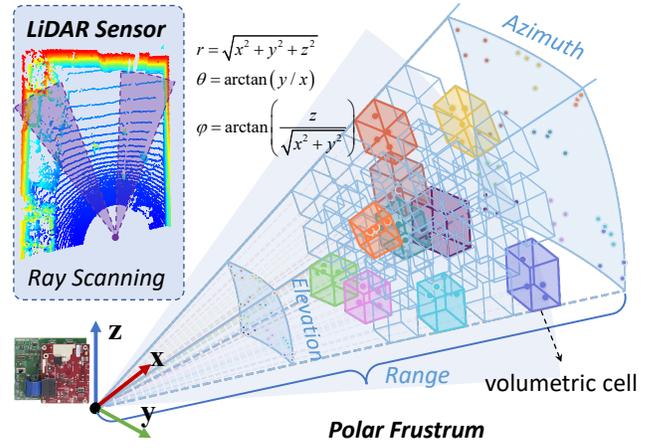


Figure 3: Frustum-based occupancy partitioning aligned with the capture geometry of LiDAR and radar. Volumetric cells are defined in polar coordinates—range, azimuth, and elevation—to match the sensors’ angular sampling patterns.

intervals. Specifically, the frustum is defined as:

$$\mathcal{F}_{i,j,k} = \left\{ (r, \theta, \phi) \left| \begin{array}{l} r \in [r_i, r_{i+1}] \subseteq [r_{\min}, r_{\max}] \\ \theta \in [\theta_j, \theta_{j+1}] \subseteq [\theta_{\min}, \theta_{\max}] \\ \phi \in [\phi_k, \phi_{k+1}] \subseteq [\phi_{\min}, \phi_{\max}] \end{array} \right. \right\}, \quad (6)$$

where r is the radial distance, θ is the azimuth angle, and ϕ is the elevation angle. The intervals $[r_{\min}, r_{\max}]$, $[\theta_{\min}, \theta_{\max}]$, and $[\phi_{\min}, \phi_{\max}]$ denote the LiDAR’s range, azimuth, and elevation field of view, respectively. Each frustum $\mathcal{F}_{i,j,k}$ thus captures a local volume along a LiDAR ray direction.

We define the occupancy O of a frustum $\mathcal{F}_{i,j,k}$ as a binary indicator:

$$O_{i,j,k} = \begin{cases} 1, & \text{if } \exists \mathbf{p} \in \mathcal{P} \text{ such that } \mathbf{p} \in \mathcal{F}_{i,j,k} \\ 0, & \text{otherwise} \end{cases}, \quad (7)$$

where \mathcal{P} denotes the set of all LiDAR points in the scene. The occupancy of a query point \mathbf{q} is then determined by the occupancy of the frustum $\mathcal{F}_{i,j,k}$ that contains it:

$$O(\mathbf{q}) = O_{i,j,k} \quad \text{if } \mathbf{q} \in \mathcal{F}_{i,j,k}. \quad (8)$$

Critically, frustum-based partitioning also facilitates learning occlusion relationships, since occupied frustums tend to be the closest along the same angular ray. This geometric alignment makes the decoder’s task of occupancy prediction, which can be viewed as an interpolation over 3D space, more structured and physically grounded. Besides, partitioning the space into frustums in polar coordinates provides a more consistent representation for radar spectrum, which makes it easier to utilize the radar spectrum as the conditioning information for the diffusion model.

Order-Invariant Latent Encoding

While the frustum-based autoencoder provides a compact latent representation aligned with LiDAR sensing geometry, it is also essential to ensure that the learned latent space respects the unordered nature of point clouds. To this end, we

introduce an enhanced order-invariant encoding mechanism that improves the generalization and stability of the latent diffusion process.

Point clouds are unordered sets, meaning any permutation of points in \mathbf{P} should represent the same geometry. However, neural encoders may inadvertently produce order-sensitive representations, which become problematic in diffusion training. During training, diffusion models are optimized to predict the noise added to latent variables at each timestep. If the encoded latent tokens change under different point orderings, then the target noise also varies inconsistently—even when the underlying geometry remains the same. This mismatch alters the optimization trajectory for the same training sample, impairing the model’s ability to learn a stable and generalizable denoising function. We visualize this issue in Figure 4, which highlights how different input orderings lead to inconsistent noise prediction targets.

To address this, we design a token encoding scheme that ensures consistent latent representations regardless of input point order. The key idea is to control the tokens fed into the cross-attention-based encoder. Rather than adopting either randomly sampled or fixed learned queries, as done in prior work (Zhang et al. 2023), we employ a *hybrid strategy* that integrates both static and dynamic queries.

Concretely, a fixed set of learned tokens serves as stable anchors, ensuring consistent token ordering across different samples. In parallel, dynamic queries are derived from \mathbf{P} via a learnable projection, allowing the encoder to capture geometry-specific features through cross-attention. The fusion of static and dynamic components forms the final set of query tokens for encoding.

Formally, we denote $\mathbf{Q}_s \in \mathbb{R}^{M \times d}$ as a fixed set of learnable query tokens that provide consistent ordering across samples. In parallel, dynamic queries \mathbf{Q}_d are computed from \mathbf{P} to capture input-specific geometric features. The final encoder queries \mathbf{Q}_{enc} are obtained by applying cross-attention between the dynamic queries and the input, followed by combining the result with the static queries:

$$\mathbf{Q}_{enc} = \text{Proj}(\mathbf{Q}_s + \text{CrossAttn}(\mathbf{Q}_d, \mathbf{P})). \quad (9)$$

This hybrid design maintains order invariance through fixed query structure, while enriching the latent space with geometry-aware features. It thus enables more effective diffusion modeling on unordered point clouds.

Radar Spectrum Guidance

Diffusion Conditioned on Radar Spectrum To incorporate radar observations into the generation process, we propose a radar-guided conditioning strategy that injects semantic and geometric cues from the radar spectrum into the latent diffusion model. Radar measurements are represented as a 3D tensor in polar coordinates (range, azimuth, elevation), where high-intensity values correspond to strong surface reflections. However, the radar signal is inherently noisy and exhibits anisotropic uncertainty—range measurements tend to be more reliable than angular ones due to beam characteristics (Kramer et al. 2022).

To handle this, we upsample the radar spectrum along the azimuth and elevation axes and apply a convolutional en-

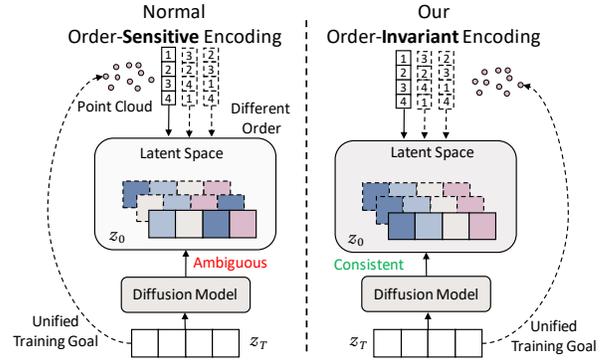


Figure 4: Order-sensitive (left) vs. order-invariant (right) latent encodings. Our method ensures consistent training targets despite the order of input point cloud.

coder \mathcal{C}_ψ to extract robust high-level features. This design enlarges the receptive field, suppressing angular noise while emphasizing more reliable range information. The resulting feature map serves as a compact and noise-aware representation of the radar input.

Since the latent space is formulated as 1D vectors, we adopt the transformer-based DiT architecture (Peebles and Xie 2023) as our diffusion model. The radar features are enriched with 3D positional embeddings to preserve spatial structure, then flattened to match the latent vector shape. This enables the diffusion model to capture spatially-aware relationships between the latent tokens and the radar signal.

Decoding with Radar-Guided Query Initialization After sampling from the latent diffusion model, we decode the resulting latent representations into point clouds. Due to the extreme sparsity of scene-level point clouds, this decoding process is computationally challenging. Although our autoencoder is trained with high-resolution frustums (centimeter-level in range and sub-degree in angle), it is infeasible to densely query the entire 3D space during decoding, as this would require evaluating tens of millions of points.

To improve efficiency, we introduce a radar-guided query initialization strategy. Specifically, we apply a low-threshold CFAR algorithm to the radar spectrum to identify candidate object regions. These CFAR detections guide the selection of query points for decoding. While imperfect, they provide a strong prior for object locations and reduce unnecessary queries in empty space. To maintain completeness, we additionally include a set of randomly sampled query points across the 3D space, allowing the decoder to capture undetected or low-reflectivity structures.

Experiments

Experiment Settings

Dataset We evaluate our method on ColoRadar (Kramer et al. 2022) dataset, which provides radar spectrum data paired with LiDAR point clouds. The dataset is collected in various scenarios, including laboratories, hallways, and

other environments. Each environment contains multiple sequences. We use the earlier sequences for training and reserve the last two sequences in each scene for validation and testing. To align radar and LiDAR data, we remove non-overlapping frames, transform LiDAR point clouds using the provided calibration parameters, and crop them to match the radar’s field of view and finally convert them to polar coordinates.

Implementation Details For the autoencoder, the occupancy granularity is set to a frustum with 0.05 m, 0.25°, and 0.5° resolution for the range, azimuth, and elevation, respectively. Input point cloud coordinates are normalized to [-1,1]. During training, point clouds are downsampled to 10,000 points, with an equal number of decoder query points. To address sparsity, 6.25% of queries are positives; the rest are randomly sampled negatives. Points are encoded and compressed into 512 latent tokens, each with 32 dimensions. The autoencoder is trained for 150 epochs with a batch size of 28.

The diffusion model is trained for 100 epochs with a batch size of 16 using the EDM sampler (Karras et al. 2022). At inference, 500k query points are sampled from free space and 700k from CFAR regions to guide generation. All models are implemented in PyTorch and trained on two NVIDIA RTX 4090 GPUs. Training takes 28 hours for the autoencoder and 60 hours for the diffusion model.

Metrics and Baselines We evaluate Chamfer distance (CD) and Earth Mover’s Distance (EMD) to measure the similarity between the generated point cloud and the ground truth, both for autoencoder and diffusion model outputs. CD measures the average distance between points in the generated and ground truth point clouds, while EMD quantifies the minimum cost of transforming one point cloud into another, considering the distribution of points. Lower values indicate better performance. We compare our method with the traditional method OS-CFAR (Richards et al. 2005), as well as the GAN-based method RPDNet (Cheng et al. 2022) and diffusion-based method SDDiff (Wang et al. 2025).

Results

Main Results We first present the main results of the proposed autoencoder in Table 1. We compare the performance of different autoencoders across three scenes: Aspen Lab, Hallways, and ARPG Lab. To validate the effectiveness of

Scene	Encoder Query	Hybrid	Sample	Static	Hybrid
	Occupancy				
Aspen Lab	CD↓	0.133	0.082	0.090	0.088
	EMD↓	0.132	0.083	0.089	0.087
Hallways	CD↓	0.166	0.094	0.118	0.112
	EMD↓	0.162	0.095	0.118	0.112
ARPG Lab	CD↓	0.160	0.082	0.104	0.081
	EMD↓	0.155	0.083	0.104	0.080

Table 1: Performance comparison of different autoencoders.

Model	Aspen Lab		Hallways		ARPG Lab	
	CD↓	EMD↓	CD↓	EMD↓	CD↓	EMD↓
OS-CFAR	1.175	1.342	1.098	1.387	1.076	1.163
RPDNet	0.874	0.587	0.793	0.664	0.823	0.512
SDDiff	0.385	0.386	0.581	0.603	0.497	0.505
RaLD	0.339	0.356	0.576	0.515	0.488	0.450

Table 2: End-to-end radar point cloud generation results.

the method of occupancy partitioning, we compare the performance of the hybrid query with the voxel-based and the frustum-based partitioning. And we also compare the performance of the hybrid query with the static query and downsampled point query in the frustum-based partitioning. All the results are obtained using 500k randomly sampled points from the free space, with the voxel size in the voxel-based query set to [0.05, 0.05, 0.05] meters along the x, y, and z axes for fair comparison.

The results show that the frustum-based occupancy partitioning strategy outperforms the voxel-based counterpart with a significant margin, which provides a maximum decrease of 49.6% in CD and 48.3% in EMD. Within the frustum-based setting, the hybrid query consistently outperforms the static query across all scenes and achieves performance comparable to the downsampled point query. The results indicate the hybrid query in the frustum-based partitioning can provide a higher upper bound of the downstream latent diffusion task, while preserving the characteristics of order-invariant encoding.

Variant	Radar CFAR		Aspen Lab		Hallways		ARPG Lab	
	Enc.	Init	CD↓	EMD↓	CD↓	EMD↓	CD↓	EMD↓
(a)	w/o	w/	0.596	0.638	0.723	0.647	0.659	0.628
(b)	w/	w/o	0.348	0.381	0.586	0.545	0.535	0.547
(c)	w/	w/	0.339	0.356	0.576	0.515	0.488	0.450

Table 3: Ablation study on radar encoder conditioning and decoder query initialization. “w/” and “w/o” indicate the presence and absence of each component, respectively. Variants (a)–(c) correspond to different combinations of these modules.

Next, we report the end-to-end generation results of the proposed RaLD model in Table 2, with qualitative examples shown in Figure 5. RaLD consistently outperforms baseline methods across all scenes, achieving up to 11.9% and 14.6% improvements in CD and EMD, respectively, compared to the second-best method, SDDiff, in the Aspen Lab and Hallways scenes.

As shown in Figure 5, RaLD produces sharper and more detailed 3D radar point clouds, better preserving high-frequency structures, while baseline method SDDiff tends to generate overly smooth outputs. For example, in the Aspen Lab and Hallways scenes, RaLD successfully captures fine-grained structures such as pillars and wall edges, whereas

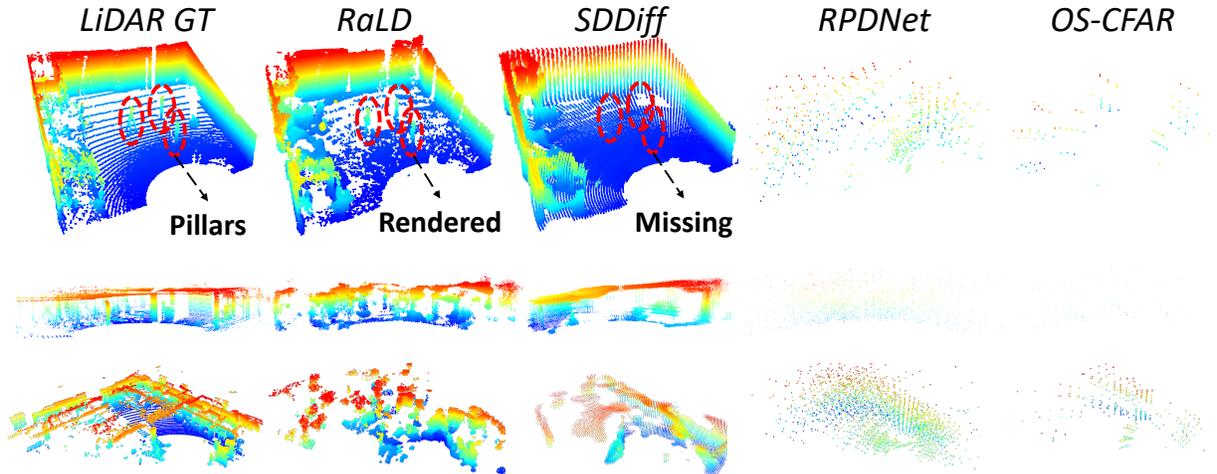


Figure 5: Visualization of 3D radar point cloud generation result.

SDDiff produces smoother outputs that omit these details. While SDDiff generates point clouds with a more continuous surface, it often lacks the key structural features present in the space, which can be crucial for downstream tasks like localization and mapping. In contrast, RaLD effectively captures these features, demonstrating its ability to leverage the learned latent space of the autoencoder to generate high-fidelity 3D radar point clouds.

Ablation Studies and Additional Results We conduct ablation studies to evaluate the impact of two key components in the RaLD model: radar encoder conditioning and decoder query initialization. The results are summarized in Table 3. We compare three model variants: (a) without radar encoder conditioning only, (b) with radar encoder conditioning only, and (c) with both components enabled. In variant (a), the raw radar spectrum is flattened and embedded with positional encodings to serve as the conditioning input to the diffusion model. For variants that do not use CFAR points for decoder query initialization, we randomly initialize the same number of points as used in the CFAR-based setup.

The results demonstrate that variant (c), which combines both radar encoder conditioning and decoder query initialization consistently achieves the best performance across all scenes. The comparison between variants (a) and (c) shows that the radar encoder conditioning significantly improves performance, indicating that the learned representation effectively captures radar signal characteristics and mitigates noise in the raw radar spectrum. Meanwhile, the comparison between variants (b) and (c) reveals that decoder query initialization further enhances performance, suggesting that initializing decoder queries with CFAR points provides a strong spatial prior for generating high-quality point clouds.

We further investigate how autoencoder design impacts RaLD performance. Specifically, we compare four types of encoder query strategies under different coordinate systems, as summarized in Table 4. Across all scenes, the diffusion model using the hybrid query in the frustum-based occupancy partitioning consistently achieves the best per-

formance in terms of EMD. For the CD metric, this configuration achieves the best result in the Aspen Lab scene and performs comparably to the downsampled point query in the other scenes. These results suggest that the hybrid query strategy within a frustum-aligned latent space is more effective at capturing the underlying 3D structure, leading to higher-quality point cloud generation. This also validates our design of order-invariant latent encoding, where combining static and dynamic queries enables stable supervision and better generalization for diffusion-based generation. While our method does not consistently produce the lowest CD—often sensitive to density and outliers (Wu et al. 2021)—it consistently improves EMD, reflecting better overall structure.

Scene	Encoder Query	Hybrid	Sample	Static	Hybrid
	Occupancy	Voxel	Frustum	Frustum	Frustum
Aspen Lab	CD↓	0.397	0.366	0.390	0.339
	EMD↓	0.519	0.422	0.412	0.356
Hallways	CD↓	0.695	0.562	0.633	0.576
	EMD↓	0.770	0.566	0.540	0.515
ARPG Lab	CD↓	0.609	0.475	0.564	0.488
	EMD↓	0.766	0.511	0.513	0.450

Table 4: Performance comparison of different autoencoder queries under different coordinate systems.

Conclusion

We propose a latent diffusion framework that generates high-resolution 3D point clouds from noisy radar spectrum. By combining a frustum-based autoencoder, order-invariant encoding, and radar spectrum guidance, our method effectively reconstructs detailed scene geometry, advancing radar-based 3D perception in challenging environments.

Acknowledgements

This work was supported in part by National Natural Science Foundation of China with Grant 62522214, 62471194, the National Key Research and Development Project under Grant 2023YFB2904500, and the Fundamental Research Funds for the Central Universities.

References

- Cheng, Y.; Su, J.; Jiang, M.; and Liu, Y. 2022. A novel radar point cloud generation method for robot environment perception. *IEEE Transactions on Robotics*, 38(6): 3754–3773.
- Fan, C.; Zhang, S.; Liu, K.; Wang, S.; Yang, Z.; and Wang, W. 2024a. Enhancing mmwave radar point cloud via visual-inertial supervision. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 9010–9017. IEEE.
- Fan, J.; Yang, J.; Xu, Y.; and Xie, L. 2024b. Diffusion model is a good pose estimator from 3d rf-vision. In *European Conference on Computer Vision*, 1–18. Springer.
- Gao, P.; Zhang, S.; Wang, W.; and Lu, C. X. 2022. Dc-loc: Accurate automotive radar based metric localization with explicit doppler compensation. In *2022 International Conference on Robotics and Automation (ICRA)*, 4128–4134. IEEE.
- Geng, R.; Li, Y.; Zhang, D.; Wu, J.; Gao, Y.; Hu, Y.; and Chen, Y. 2024. Dream-pcd: Deep reconstruction and enhancement of mmwave radar pointcloud. *IEEE Transactions on Image Processing*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Guan, J.; Madani, S.; Jog, S.; Gupta, S.; and Hassanieh, H. 2020. Through fog high-resolution imaging using millimeter wave radar. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11464–11473.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Karras, T.; Aittala, M.; Aila, T.; and Laine, S. 2022. Elucidating the Design Space of Diffusion-Based Generative Models. In *Proc. NeurIPS*.
- Kramer, A.; Harlow, K.; Williams, C.; and Heckman, C. 2022. Coloradar: The direct 3d millimeter wave radar dataset. *The International Journal of Robotics Research*, 41(4): 351–360.
- Kurach, K.; Lučić, M.; Zhai, X.; Michalski, M.; and Gelly, S. 2019. A large-scale study on regularization and normalization in GANs. In *International conference on machine learning*, 3581–3590. PMLR.
- Lai, H.; Luo, G.; Liu, Y.; and Zhao, M. 2024. Enabling visual recognition at radio frequency. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, 388–403.
- Li, W.; Liu, J.; Yan, H.; Chen, R.; Liang, Y.; Chen, X.; Tan, P.; and Long, X. 2024. Craftsman3d: High-fidelity mesh generation with 3d native generation and interactive geometry refiner. *arXiv preprint arXiv:2405.14979*.
- Lu, C. X.; Rosa, S.; Zhao, P.; Wang, B.; Chen, C.; Stankovic, J. A.; Trigoni, N.; and Markham, A. 2020. See through smoke: robust indoor mapping with low-cost mmwave radar. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, 14–27.
- Luan, K.; Shi, C.; Wang, N.; Cheng, Y.; Lu, H.; and Chen, X. 2024. Diffusion-based point cloud super-resolution for mmwave radar data. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 11171–11177. IEEE.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205.
- Prabhakara, A.; Jin, T.; Das, A.; Bhatt, G.; Kumari, L.; Soltanaghaei, E.; Bilmes, J.; Kumar, S.; and Rowe, A. 2022. High resolution point clouds from mmwave radar. *arXiv preprint arXiv:2206.09273*.
- Qian, K.; He, Z.; and Zhang, X. 2020. 3D point cloud generation with millimeter-wave radar. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(4): 1–23.
- Richards, M. A.; et al. 2005. *Fundamentals of radar signal processing*, volume 1. Mcgraw-hill New York.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Roy III, R. H.; and Kailath, T. 1990. ESPRIT—Estimation of signal parameters via rotational invariance techniques. *Optical Engineering*, 29(4): 296–313.
- Schmidt, R. 1986. Multiple emitter location and signal parameter estimation. *IEEE transactions on antennas and propagation*, 34(3): 276–280.
- Soumekh, M. 1999. *Synthetic aperture radar signal processing*, volume 7. New York: Wiley.
- Sun, Y.; Huang, Z.; Zhang, H.; Cao, Z.; and Xu, D. 2021. 3DRIMR: 3D reconstruction and imaging via mmWave radar based on deep learning. In *2021 IEEE International Performance, Computing, and Communications Conference (IPCCC)*, 1–8. IEEE.
- Wang, S.; Luo, X.; Xie, Y.; and Wang, W. 2025. SD-Diff: Boost Radar Perception via Spatial-Doppler Diffusion. *arXiv preprint arXiv:2506.16936*.
- Wu, T.; Pan, L.; Zhang, J.; Wang, T.; Liu, Z.; and Lin, D. 2021. Density-aware chamfer distance as a comprehensive metric for point cloud completion. *arXiv preprint arXiv:2111.12702*.
- Yin, T.; Zhou, X.; and Krahenbuhl, P. 2021. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11784–11793.
- Zhang, B.; Tang, J.; Niessner, M.; and Wonka, P. 2023. 3dshape2vecset: A 3d shape representation for neural fields

and generative diffusion models. *ACM Transactions On Graphics (TOG)*, 42(4): 1–16.

Zhang, L.; Wang, Z.; Zhang, Q.; Qiu, Q.; Pang, A.; Jiang, H.; Yang, W.; Xu, L.; and Yu, J. 2024a. Clay: A controllable large-scale generative model for creating high-quality 3d assets. *ACM Transactions on Graphics (TOG)*, 43(4): 1–20.

Zhang, R.; Xue, D.; Wang, Y.; Geng, R.; and Gao, F. 2024b. Towards dense and accurate radar perception via efficient cross-modal diffusion model. *IEEE Robotics and Automation Letters*.

Zhao, Z.; Liu, W.; Chen, X.; Zeng, X.; Wang, R.; Cheng, P.; Fu, B.; Chen, T.; Yu, G.; and Gao, S. 2023. Michelangelo: Conditional 3d shape generation based on shape-image-text aligned latent representation. *Advances in neural information processing systems*, 36: 73969–73982.

Zhou, Y.; and Tuzel, O. 2018. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4490–4499.