DyRAD: Radar Novel View Synthesis for Dynamic Driving Scenes

Anonymous authors
Paper under double-blind review

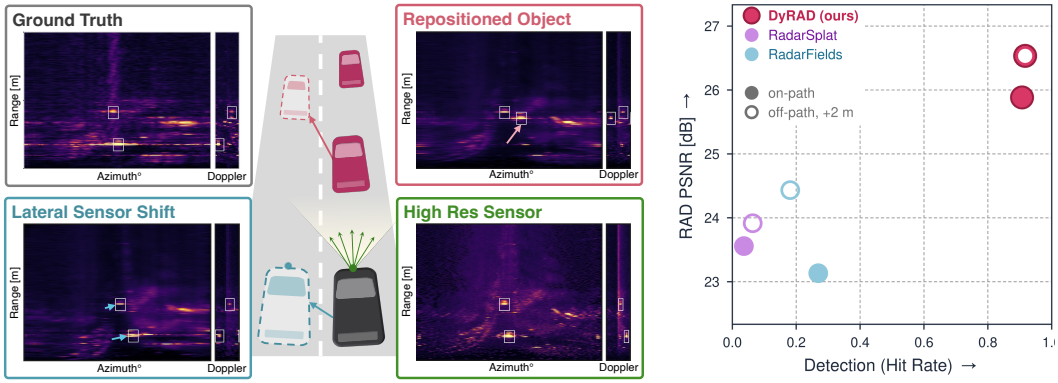


Figure 1: **Radar re-simulation with DyRAD.** From recorded radar measurements, poses, and object boxes, DyRAD reconstructs dynamic scenes supporting **novel viewpoints**, **sensor-configuration changes**, and **scene editing**. **Right:** Higher RAD PSNR and detection hit rate than both baselines on RADIal, on-path and after a $2\,\mathrm{m}$ lateral-shift round trip. See the supplementary HTML viewer for videos of all real-data off-path and synthetic experiments.

Abstract

Reconstructing dynamic driving scenes from recorded sensor data supports closed-loop evaluation of autonomous driving systems by synthesizing observations beyond the original trajectory. Unlike cameras and LiDAR, radar measures radial velocity directly through Doppler. Yet existing radar novel-view synthesis fails to exploit this capability: methods addressing dynamic scenes reconstruct only range–azimuth tensors, while methods that render Doppler assume static scenes. Moreover, because radar processing spreads each reflection across multiple bins, existing representations absorb this spread into scene geometry, causing it to render incorrectly when the viewpoint moves. We present DyRAD, which models dynamic driving scenes using static background reflectors and motion-tracked dynamic point reflectors to render complete range–azimuth–Doppler (RAD) tensors. Reflector velocities are derived from object tracks and projected onto the line of sight, making Doppler both a rendered output and supervision for those tracks. Crucially, we render reflectors through a fixed analytic point-spread function (PSF) derived from the radar's signal-processing chain, preventing sensor-induced spread from being baked into the scene representation. Beyond improving scene reconstruction, this separation also enables zero-shot sensor-configuration transfer, allowing the same reconstructed scene to be rendered under different radar specifications without refitting. We evaluate DyRAD on RADIal, Boreas, and a synthetic benchmark across both on-path poses and displaced viewpoints untested by prior work. On RADIal, DyRAD recovers radar detections in 90.7% of reference-detected objects, compared with 26.9% for the strongest baseline.

1 Introduction

Accurate sensor simulation is critical for training and evaluating automotive systems. Reconstructing dynamic driving scenes from recorded sensor data supports this goal by enabling realistic ob-

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

servations to be synthesized beyond the original trajectory. Capturing radar measurements is particularly important given radar's distinct capabilities within a vehicle's sensor suite: direct radial velocity measurements via Doppler and robust operation under adverse weather. Yet its sparse measurements and limited spatial resolution make reconstruction particularly challenging. Received power varies with range and viewing aspect, while Doppler shifts depend on relative sensor–object motion. Re-simulation must therefore recover the scene from measurements that couple sensor pose, object dynamics, and signal processing.

Neural scene representations such as NeRF (Mildenhall et al., 2021) and 3D Gaussian Splatting (Kerbl et al., 2023) now support novel-view synthesis (NVS) of dynamic driving scenes for camera and LiDAR (Yan et al., 2024; Chen et al., 2025; Turki et al., 2026), with related approaches emerging for radar (Borts et al., 2024; Kung et al., 2025; Zhang et al., 2026; Huang et al., 2024). Three capabilities, however, remain unexplored: 1) reconstructing dynamic scenes together with their Doppler signature, 2) recovering scene structure rather than the sensor's own response, and 3) evaluating synthesis beyond the recorded trajectory.

1) Dynamics and Doppler. Radar measures motion through Doppler, a property routinely exploited in radar perception (Haitman & Bialer, 2025; Ding et al., 2024; Musiat et al., 2024; Paek et al., 2022). Yet existing radar NVS methods discard this measurement where it matters most: methods addressing dynamic driving scenes reconstruct only range–azimuth (RA), omitting Doppler (Zhang et al., 2026). Methods that render Doppler assume static environments, where it arises solely from sensor motion (Huang et al., 2024; Chen et al., 2026; Kung et al., 2026). No existing radar NVS method reconstructs dynamic driving scenes from complete RAD measurements while using rendered Doppler to constrain object motion.

2) Separating the scene from the sensor. A radar measurement is not a direct picture of the scene: the radar's signal-processing chain spreads each reflector's return across range, azimuth, and Doppler bins, as described by its point-spread function (PSF). Although existing methods estimate occupancy or transmittance to recover scene structure (Kung et al., 2025; Borts et al., 2024), the spatial extent of that structure is learned directly from the measurements. A broad return can therefore be explained by a broad scene element rather than the sensor's PSF, leaving scene structure entangled with sensor-induced spread. Once absorbed into the scene representation, this spread behaves like static geometry rather than a measurement artifact, causing it to render incorrectly from displaced viewpoints. Recovering true scene structure therefore requires explicitly modeling the sensor's response from its signal-processing chain. Furthermore, isolating the sensor response in this way unlocks a capability prior methods cannot provide: resimulating the recovered scene under a different sensor configuration by simply swapping the analytic PSF.

3) Evaluating away from the recorded path. Existing radar NVS evaluates only at held-out frames along the driven trajectory (Borts et al., 2024; Kung et al., 2025; Zhang et al., 2026). Because these test frames sit right between nearly identical training poses, interpolating directly in measurement space can score just as well as a correct forward model. As a result, this protocol fails to test synthesis at displaced viewpoints, which are precisely what closed-loop simulation requires.

To close these gaps we propose DyRAD, a dynamic point-reflector scene representation coupled with a differentiable, physics-grounded RAD renderer. Radar returns originate from discrete scattering centers, so we represent the scene as zero-extent point reflectors whose positions, base reflected power, and aspect-dependent reflectivity are optimized from the measurements. Static reflectors represent the background, while dynamic reflectors follow learned rigid object tracks initialized from bounding boxes. We derive reflector velocities from these tracks and project relative sensor–reflector motion onto the line of sight to render Doppler. Jointly optimizing the reflectors and tracks against recorded RAD measurements makes Doppler both a rendered output and a constraint on those tracks.

We address the second gap by explicitly modeling the sensor's response. Each point reflector is rendered through a sensor-specific PSF across range, azimuth, and Doppler, derived from the radar's signal-processing chain and held fixed throughout optimization. Using zero-extent reflectors with a fixed sensor response constrains how measurement spread is explained, reducing ambiguity between scene structure and sensor-induced blur. This makes rendering at sensor poses away from the driven trajectory a forward model rather than interpolation in measurement space, and it makes the recovered scene portable across sensors (Fig. 1).

We evaluate DyRAD on RADIal, Boreas, and a synthetic benchmark, testing reconstruction at held-out poses along the recorded trajectory as well as at displaced viewpoints; the latter addressing the third gap. We assess displaced views directly against synthetic ground truth and through a render-and-refit cycle on real data. On RADIal, DyRAD increases full-RAD correlation from 0.068 to 0.272 and foreground hit rate from 26.9% to 90.7% over the strongest respective baselines. Ablations show that Doppler supervision reduces vehicle Doppler peak error by 60% relative to RA-only fitting, while fixing the analytic PSF more than doubles joint detection recall compared with learning Gaussian extents or the PSF. Finally, the same reconstructed scene supports coarse-to-fine sensor-configuration transfer without refitting, improving object-region RAD PSNR and nearly doubling detection F1 over linear upsampling.

Our main contributions are as follows:

1. **Physics-grounded RAD rendering of dynamic scenes.** To our knowledge, we introduce the first radar NVS method to render Doppler for dynamic driving scenes. We reconstruct scenes as static and dynamic point reflectors following rigid object tracks and render complete RAD tensors with Doppler derived from relative sensor–reflector motion.
2. **Decoupling scene structure from sensor response.** We render point reflectors through a fixed, sensor-specific PSF, separating scene structure from sensor-induced spread. This improves object-region reconstruction and detection while also enabling the same scene to be rendered under alternative radar configurations without refitting.
3. **Off-path evaluation for radar novel-view synthesis.** Prior radar NVS evaluates only along the recorded trajectory, testing interpolation rather than true spatial generalization. We introduce off-path evaluation via displaced ground-truth views in a synthetic benchmark and a cycle-consistency protocol adapted to real radar recordings.

2 Related Work

2.1 Novel-view synthesis for dynamic driving scenes

Reconstructing driving scenes from recorded data has become an established approach to sensor simulation for autonomous driving. Once reconstructed, a scene can be rendered along new ego trajectories, enabling training and closed-loop evaluation beyond the recorded drive (Tancik et al., 2022; Manivasagam et al., 2020; Yang et al., 2023; Tonderski et al., 2024). Early neural approaches focused on static camera scenes (Tancik et al., 2022; Rematas et al., 2022), later extending to dynamic driving environments (Turki et al., 2023; Yang et al., 2024; Yan et al., 2024; Chen et al., 2025). Reconstruction-based simulation also expanded to LiDAR, progressing from static (Manivasagam et al., 2020; Huang et al., 2023) to dynamic scenes (Wu et al., 2024; Zheng et al., 2024), and subsequently to joint camera–LiDAR rendering (Yang et al., 2023; Tonderski et al., 2024; Hess et al., 2025; Turki et al., 2026). For camera and LiDAR, dynamic driving scenes can now be reconstructed and replayed from viewpoints that were never driven. Our work makes this possible also for radar.

2.2 Radar reconstruction and Doppler

Radar novel-view synthesis. Radar reconstruction has developed along similar lines as its camera and LiDAR counterparts. Radar Fields (Borts et al., 2024) and RadarSplat (Kung et al., 2025) reconstruct RA measurements of *static* driving scenes, the former with a frequency-space neural field, the latter with Gaussian primitives combining modeling of radar noise and multipath effects. RF4D (Zhang et al., 2026) extends the Radar Fields framework to dynamic driving scenes through a spatiotemporal field and a scene-flow module that predicts motion offsets between adjacent frames. Yet it too reconstructs RA alone. NeuRadar (Rafidashti et al., 2025) extends joint camera–LiDAR rendering of dynamic driving scenes to radar point clouds, but renders radar as sparse detections without Doppler information rather than as raw, dense measurements.

Doppler as a motion cue. Doppler has been used as an input motion cue for visual scene reconstruction, as in 4DRadar-GS (Tang et al., 2025), which uses per-point radar velocities but does not render radar measurements. Other methods render Doppler but assume static environments (Huang et al., 2024; Chen et al., 2026; Kung et al., 2026), where it arises solely from sensor motion. We instead reconstruct dynamic driving scenes directly from full RAD measurements, where Doppler is both rendered and used to refine the object motion that produced it.

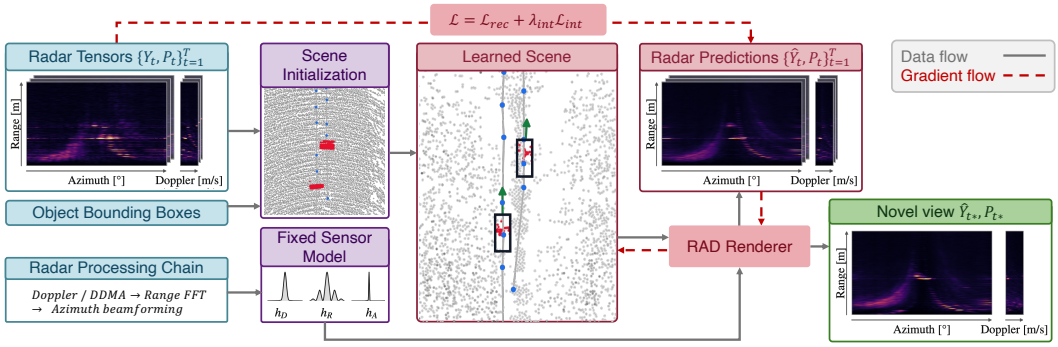


Figure 2: **DyRAD overview. Inputs:** recorded RAD tensors, sensor poses, object bounding-boxes, and radar processing information when available. **Initialization:** measurements and annotations initialize the scene; processing information defines the fixed sensor response. **Optimization:** the differentiable RAD renderer jointly fits reflectors and object tracks to recorded measurements. **Inference:** the optimized scene produces RAD predictions $\hat{\mathbf{Y}}(t^\star, \mathbf{P}^\star)$ at query times and sensor poses. Solid gray arrows show **data flow**; dashed red arrows show **gradient flow**.

2.3 Sensor Modeling in Neural Rendering

Neural rendering has repeatedly improved by replacing an idealized sensor with a model of how the measurement is actually formed. Methods account for sampling footprints and beam divergence (Barron et al., 2021; Huang et al., 2023), processing and blur (Mildenhall et al., 2022; Ma et al., 2022), and appearance variation (Martin-Brualla et al., 2021). Radar reconstruction additionally accounts for propagation, antenna gain, and range–Doppler geometry (Borts et al., 2024; Huang et al., 2024; Chen et al., 2026; Armouti et al., 2026), as well as speckle and multipath (Kung et al., 2025). Measurement spreading has been addressed through PSF convolution for simulation (Bialer & Haitman, 2024), range smoothing and antenna-profile convolution (Kung et al., 2025), and Doppler soft binning (Kung et al., 2026). We instead represent the scene using static and dynamic point reflectors coupled with a fixed, sensor-specific RAD PSF, explicitly separating sensor-induced measurement spread from the learned scene representation.

3 Method

Method Overview. DyRAD reconstructs dynamic driving scenes from radar measurements to synthesize RAD tensors at novel sensor poses (Fig. 2). We represent the scene using static and dynamic point reflectors, with dynamic reflectors following rigid object tracks (Sec. 3.1). We couple this representation with a differentiable renderer that maps reflectors to radar measurements through a fixed sensor-specific response (Sec. 3.2). We initialize the reflectors and tracks from radar measurements and object bounding boxes, then jointly optimize them to reproduce the observations (Sec. 3.3).

Problem Setting. We are given a sequence of radar measurements $\{\mathbf{Y}_t\}_{t=1}^{T}$ and corresponding sensor poses $\{\mathbf{P}_t\}_{t=1}^{T}$. Each RAD tensor $\mathbf{Y}_t \in \mathbb{R}_{\geq 0}^{D\times R\times A}$ records return strength over D radial-velocity bins, R range bins, and A azimuth bins. We obtain RA and RD projections by averaging over Doppler and azimuth, respectively; for sensors without Doppler, we reconstruct RA maps. We additionally assume dynamic-object bounding-box annotations at a subset of frames, a common input in dynamic driving-scene reconstruction (Yan et al., 2024; Chen et al., 2025). These annotations initialize object tracks, providing spatial and motion information that is difficult to infer from sparse radar returns alone. Since the measurements do not resolve elevation, we model geometry and motion in the ground plane, with sensor poses defined by planar position and heading. Our goal is to synthesize measurements $\hat{\mathbf{Y}}(t^\star, \mathbf{P}^\star)$ at a query time $t^\star$ and sensor pose $\mathbf{P}^\star$.

3.1 Scene Representation

Scene Primitives. We represent the scene as N point reflectors with learned positions and reflectivities. Reflectors have no spatial extent; the renderer models sensor-induced spreading through a

fixed PSF. We partition the reflectors into a static background and M dynamic objects, with reflectors within each object sharing a rigid motion. The scene is then: $\mathcal{S} = \{(\mathbf{x}_i(\cdot), o_i, \boldsymbol{\eta}_i, q_i)\}_{i=1}^{N}$, where $\mathbf{x}_i(t) \in \mathbb{R}^2$ is reflector i's world-space position at time t. The learned logit $o_i \in \mathbb{R}$ defines reflector i's base reflected power $\alpha_i = \sigma(o_i) \in (0, 1)$. $\boldsymbol{\eta}_i \in \mathbb{R}^K$ contains spherical-harmonic coefficients for view-dependent reflectivity. The assignment $q_i \in \{0, \dots, M\}$ identifies the static background ($q_i = 0$) or a dynamic object ($q_i > 0$).

Motion Modeling. Following prior dynamic driving-scene reconstruction (Yan et al., 2024; Chen et al., 2025), we represent each moving object by a planar rigid track (Fig. 3).

Object j has one learned track point $\mathbf{c}_{j,k} \in \mathbb{R}^2$ per training timestamp t_k, with the first point defining its reference-frame origin. For $t \in [t_k, t_{k+1}]$, we linearly interpolate $\mathbf{c}_{j,k}$ and $\mathbf{c}_{j,k+1}$ to obtain its position $\mathbf{p}_j(t)$. We estimate its orientation from the direction of travel using neighboring track points and interpolate it between timestamps. The rotation $\mathbf{R}_j(t)$ describes the change in orientation relative to the object's initial orientation. Reflector i's world-space position is then:

$$\mathbf{x}_i(t) = \begin{cases} \mathbf{x}_i, & q_i = 0, \\ \mathbf{p}_j(t) + \mathbf{R}_j(t)\boldsymbol{\mu}_i, & q_i = j > 0. \end{cases} \tag{1}$$

Where $\mathbf{x}_i$ is a static reflector's learned world-space position and $\boldsymbol{\mu}_i$ is a dynamic reflector's learned position in its object's reference frame.

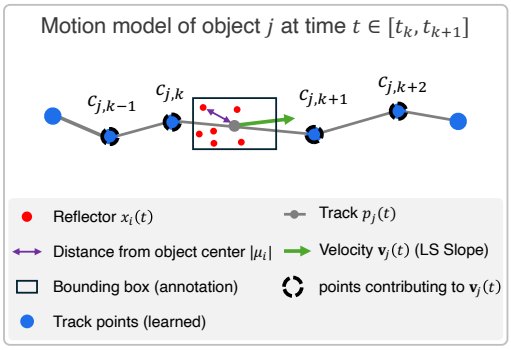


Figure 3: **Object motion along a learned track.** Four neighboring track points estimate translational velocity. Reflectors follow the track's translation and relative rotation, preserving their object-frame coordinates $\boldsymbol{\mu}_i$.

3.2 Physics-Grounded RAD Rendering

We render RAD tensors by projecting reflectors into sensor coordinates, computing their Doppler, applying the sensor-specific PSF, and accumulating their contributions. The pipeline is differentiable and implemented in CUDA for efficiency.

Reflector Projection. At time t, we transform each reflector's world-space position into the sensor frame using pose $\mathbf{P}$ and compute its range r_i, azimuth ϕ_i, and unit viewing direction $\mathbf{d}i$ from the sensor to the reflector. We model view-dependent reflected power as: $\rho_i(\mathbf{d}_i) = \alpha_i \exp\left(\sum_{k=0}^{K-1} \eta_{i,k} B_k(\mathbf{d}_i)\right)$, where B_k is the k-th real spherical-harmonic (SH) basis function. We fix the degree-zero coefficient to $\eta_{i,0} = 0$ to avoid redundancy between the constant harmonic and the base reflected power α_i.

Doppler. Doppler measures relative sensor–reflector velocity along the line of sight. Let $\mathbf{v}_i(t)$ denote reflector i's world-space velocity. Static reflectors have $\mathbf{v}_i(t) = \mathbf{0}$. For dynamic reflectors, we derive $\mathbf{v}_i(t)$ from the learned object track, accounting for changes in both position and orientation. We estimate the object's translational velocity as the slope of a linear fit to four neighboring object track points (Fig. 3) and additionally include the velocity induced by changes in object orientation. We obtain ego sensor velocity $\mathbf{v}_s(t)$ from the dataset when available and otherwise estimate it from timestamped sensor poses. A reflector's radial velocity is then $\nu_i(t) = s_D \cdot \hat{\mathbf{r}}_i(t)^\top \big(\mathbf{v}_i(t) - \mathbf{v}_s(t)\big)$, where $\hat{\mathbf{r}}_i(t)$ is the unit vector from the sensor to reflector i and $s_D \in \{-1, 1\}$ accounts for the sensor's Doppler sign convention.

Radar Doppler measurements are periodic: radial velocities outside the sensor's unambiguous interval wrap back into it, so different radial velocities can produce the same Doppler coordinate. We reproduce this behavior by wrapping $\nu_i(t)$ into this interval, yielding the rendered coordinate $\widetilde{\nu}_i(t)$.

Sensor Response. The sensor's PSF determines how each reflector's return spreads across range, azimuth, and Doppler axes. We model reflector i's contribution as:

$$\psi_i(r, \phi, \nu) = \rho_i(\mathbf{d}_i)\, h_R(r - r_i)\, h_A(\phi, \phi_i)\, h_D(\nu - \widetilde{\nu}_i(t)), \tag{2}$$

where $\rho_i(\mathbf{d}_i)$ is its view-dependent reflected power, and h_R, h_A, and h_D describe the sensor's spread kernel along each axis, where each captures range-FFT leakage, the azimuth beamformer response, and the Doppler-processing response, respectively. The Doppler kernel h_D is periodic, wrapping responses across the measured interval's boundaries. The azimuth kernel depends on both the evaluated angle ϕ and the reflector angle ϕ_i, allowing the PSF to vary across viewing directions.

All kernels are computed before scene optimization and remain fixed. We derive them from the signal-processing chain when available, and approximate otherwise (Appendix Sec. A.1).

Rasterization. Unlike depth-ordered alpha compositing in camera rendering, radar signal formation involves the superposition of echoes from multiple reflectors. We approximate it by summing reflected powers incoherently. We form the RAD tensor by summing the contributions of all reflectors' responses at the sensor's sampling grid: $\mathbf{\Psi}(t, \mathbf{P})[\ell, m, n] = \sum_{i=1}^{N} \psi_i(r_m, \phi_n, \nu_\ell)$, where r_m, ϕ_n, and ν_ℓ are the range, azimuth, and Doppler coordinates of cell $[\ell, m, n]$.

Finally, we convert this power tensor to the recorded intensity representation (e.g., amplitude, power, or log-compressed power), incorporating a sensor-specific measurement floor to account for the nonzero background. This floor is determined during preprocessing and held fixed throughout optimization. The resulting tensor is our prediction $\hat{\mathbf{Y}}(t, \mathbf{P})$.

3.3 Scene Initialization and Optimization

Initialization. We associate object bounding boxes across frames to initialize tracks and use the boxes to separate dynamic-object regions from the static background. Each track contains one point per training frame, initialized from annotated locations where available and linearly interpolated otherwise. We initialize background reflectors from peaks in Doppler-averaged RA maps. For each dynamic object, we extract peaks within its bounding box from the Doppler bin with the strongest response at its annotated location.

We project the peaks onto the ground plane and accumulate them across frames retaining background peaks in world coordinates and aligning dynamic peaks to their object's reference frame using the initialized tracks. We merge nearby peaks with the same object assignment to obtain the initial reflectors. Appendix Sec. B.1 provides details and an illustration.

Optimization. We jointly optimize the reflector parameters and track points using $\mathcal{L} = \mathcal{L}_{\text{rec}} + \lambda_{\text{int}}\mathcal{L}_{\text{int}}$, where $\mathcal{L}_{\text{rec}}$ is the mean squared error between predicted and recorded RAD tensors (or RA maps for sensors without Doppler). Inspired by RegNeRF (Niemeyer et al., 2022), the interpolation-consistency loss $\mathcal{L}_{\text{int}}$ regularizes predictions between training frames. We sample a time $\tilde{t}$ between adjacent training timestamps t_0 and t_1 and interpolate the sensor pose to obtain $\tilde{\mathbf{P}}$. We compare the predicted RA map with the neighboring measurements' average after spatial smoothing:

$$\mathcal{L}_{\text{int}} = \left\| \mathcal{G}\Big(\Pi_{\text{RA}}\big(\hat{\mathbf{Y}}(\tilde{t}, \tilde{\mathbf{P}})\big)\Big) - \mathcal{G}\left(\frac{\Pi_{\text{RA}}(\mathbf{Y}_{t_0}) + \Pi_{\text{RA}}(\mathbf{Y}_{t_1})}{2}\right) \right\|_1, \tag{3}$$

where Π_{RA} averages over Doppler and $\mathcal{G}$ applies a fixed spatial Gaussian blur. This provides a coarse appearance constraint at intermediate timesteps. We apply it only to RA, since the neighboring-frame average does not specify intermediate object positions or Doppler profiles. We evaluate its effect in the ablation studies (Sec. 4.4), with full results in Appendix Sec. D.

Densification. During optimization, we periodically extract peaks from positive reconstruction residuals and add static-background reflectors where no nearby reflector exists. Additional details and hyperparameters are provided in Sec. B.3 of the appendix.

4 Experiments

We encourage readers to explore the HTML viewer included in the supplementary material, which provides videos of all real-data off-path and synthetic experiments.

4.1 Experimental Setup

Datasets. We evaluate on two real-world datasets, *RADIal* (Rebut et al., 2022) and *Boreas* (Burnett et al., 2023), and one synthetic benchmark. Our primary benchmark uses ten 60-frame RADIal

Table 1: **On-path evaluation.** Mean ± sample standard deviation across ten RADIal sequences and three Boreas sequences. * indicates Doppler lifted RA-only baselines. Full results in Sec. C.2.

	RADIal: RAD reconstruction			RADIal: detection			Boreas: RA reconstruction		
Method	PSNR↑	$\rho\uparrow$	$\rho_{\text{obj}}\uparrow$	CD [m] ↓	F1 ↑	Hit rate ↑	PSNR↑	$\rho\uparrow$	$\rho_{\text{obj}}\uparrow$
RadarSplat*	$23.56_{\pm 0.85}$	$0.068_{\pm 0.028}$	$-0.010_{\pm 0.057}$	$6.53_{\pm 5.82}$	$0.073_{\pm 0.048}$	$0.036_{\pm 0.067}$	$23.92_{\pm 1.07}$	$0.666_{\pm 0.036}$	$0.178_{\pm 0.249}$
RadarFields*	$23.13_{\pm 1.23}$	$0.015_{\pm 0.009}$	$-0.011_{\pm 0.051}$	$6.97_{\pm 0.96}$	$0.016_{\pm 0.007}$	$0.269_{\pm 0.245}$	$19.96_{\pm 0.84}$	$0.279_{\pm 0.078}$	$0.016_{\pm 0.128}$
DyRAD-static	$25.67_{\pm 0.95}$	$0.213_{\pm 0.054}$	$-0.013_{\pm 0.072}$	$4.71_{\pm 1.29}$	$0.123_{\pm 0.041}$	$0.070_{\pm 0.083}$	$25.31_{\pm 0.90}$	$0.749_{\pm 0.018}$	$0.278_{\pm 0.217}$
DyRAD	$\mathbf{25.88}_{\pm 0.95}$	$\mathbf{0.272}_{\pm 0.062}$	$\mathbf{0.658}_{\pm 0.070}$	$\mathbf{4.01}_{\pm 1.02}$	$\mathbf{0.179}_{\pm 0.041}$	$\mathbf{0.907}_{\pm 0.086}$	$\mathbf{25.36}_{\pm 0.90}$	$\mathbf{0.752}_{\pm 0.017}$	$\mathbf{0.622}_{\pm 0.130}$

sequences, with raw radar measurements, vehicle annotations, and ego-motion measurements, supporting full-RAD reconstruction and evaluation of scene dynamics. For Boreas, we use three 50-frame sequences. Its 360° spinning Navtech radar provides RA measurements without Doppler, testing reconstruction in the baselines' native sensor setting and complementing RADIal's front-facing radar. Finally, we construct five synthetic scenes with an idealized RADIal-format sensor model to obtain ground-truth measurements at displaced poses for direct off-path evaluation. We will release the synthetic benchmark, including off-path ground-truth measurements, to support future research. Implementation details and the generation procedure are provided in Appendix Sec. C.4.

Baselines. We compare with RadarSplat (Kung et al., 2025), and RadarFields (Borts et al., 2024), which reconstruct static scenes and predict RA measurements. We use their published configurations on Boreas and adapt them to RADIal's sensor geometry and resolution. For RAD reconstruction and detection evaluation, we assign Doppler assuming a static scene and using the sensor's ego-motion; these results are marked with an asterisk. We omit RF4D (Zhang et al., 2026) because its released code is "still under checking". We also evaluate DyRAD-static, which retains our point-reflector representation and renderer but treats all reflectors as static and omits bounding-box separation during initialization. Comparing this variant with the full model measures the combined benefit of object separation and motion modeling.

Metrics. *Reconstruction.* Following RadarSplat and RadarFields, we report PSNR and SSIM (Wang et al., 2004) over full measurements and object regions of RAD, RA, and RD, and full-measurement LPIPS (Zhang et al., 2018) over full RA and RAD measurements. However, agreement at the noise floor can dominate these global scores in sparse radar data (Sec. C.1). We therefore report Pearson correlation, a scale-invariant measure of structural agreement. ρ and ρ_{obj} denote correlation over full measurements and object regions, respectively.

Detection. To assess whether synthesized measurements preserve usable radar returns, we apply identical RD-CFAR detection to reference and synthesized RAD tensors and compare the resulting point clouds using metrics adapted from RadarGen (Borreda et al., 2026). We report spatial Chamfer distance (CD), joint F1 over position, Doppler, and power, and foreground hit rate: the fraction of reference-occupied annotated boxes also containing predicted detections. Detection evaluation covers RADIal and synthetic data; Boreas lacks the Doppler dimension required by RD-CFAR. Metric definitions and additional results appear in Appendix Sec. C.

4.2 Novel View Synthesis

On-Path Evaluation. Following prior work (Kung et al., 2025; Borts et al., 2024), we hold out every fifth frame and exclude its measurements and bounding boxes from initialization and optimization (Table 1). Our static variant already improves reconstruction over both baselines, including in their native RA output: on RADIal, DyRAD-static achieves correlation 0.533, compared with 0.291 for RadarSplat and 0.055 for RadarFields. However, this stronger aggregate reconstruction still misses most vehicle returns. Moving from DyRAD-static to the full model increases full-RAD correlation from 0.213 to 0.272, while object-region RAD correlation rises from -0.013 to 0.658 and foreground hit rate from 7.0% to 90.7%. On Boreas, the gap over RadarSplat is smaller in full-measurement correlation (0.752 versus 0.666) than within object regions (0.622 versus 0.178). These results support the benefit of explicit object modeling even without Doppler supervision. Figure 4 provides qualitative examples on both datasets, with enlarged regions highlighting vehicle returns that DyRAD preserves but the static variant and baselines weaken or miss.

Real-Data Off-Path Evaluation. Without ground-truth measurements at displaced viewpoints, we adopt the cycle-consistency protocol of Neural LiDAR Fields (Huang et al., 2023). We fit each

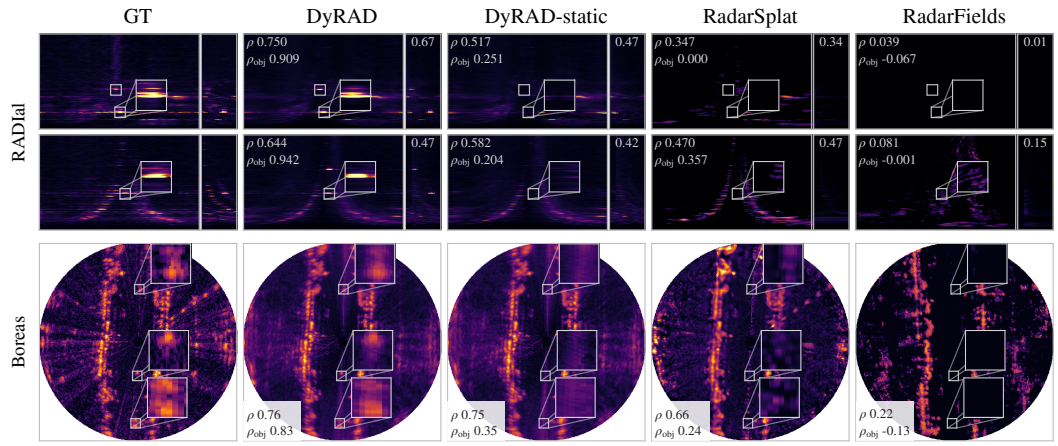


Figure 4: **On-path reconstruction on RADIal and Boreas.** The top two rows show held-out RADIal frames with RA (left) and RD (right) projections; the bottom row shows a held-out Boreas RA frame in Cartesian bird's-eye view. White boxes mark objects, with selected regions enlarged in the insets. DyRAD preserves object returns that the static variant and baselines weaken or miss, even when surrounding scene structure appears similar.

Table 2: **Real-data off-path evaluation.** Reconstruction at the original poses after a $2\mathrm{m}$ lateral-shift round trip. Mean $\pm$ sample standard deviation across ten RADIal sequences and three Boreas sequences. * indicates Doppler lifted RA-only baselines. See the supplementary HTML viewer for off-path videos and Sec. C.3 for full quantitative results.

	RADIal: RAD reconstruction			RADIal: detection			Boreas: RA reconstruction		
Method	PSNR↑	$\rho\uparrow$	$\rho_{\text{obj}}\uparrow$	CD [m] ↓	F1 ↑	Hit rate ↑	PSNR↑	$\rho\uparrow$	$\rho_{\text{obj}}\uparrow$
RadarSplat*	23.92 $\pm$0.85	0.076 $\pm$0.032	−0.004 $\pm$0.053	6.35 $\pm$4.75	0.077 $\pm$0.050	0.064 $\pm$0.065	24.29 $\pm$1.16	0.689 $\pm$0.042	0.211 $\pm$0.237
RadarFields*	24.43 $\pm$0.75	0.023 $\pm$0.018	−0.011 $\pm$0.048	9.19 $\pm$6.07	0.027 $\pm$0.020	0.181 $\pm$0.183	21.09 $\pm$1.28	0.538 $\pm$0.081	0.095 $\pm$0.123
DyRAD-static	26.41 $\pm$0.73	0.335 $\pm$0.054	0.056 $\pm$0.064	3.47 $\pm$0.61	0.210 $\pm$0.042	0.209 $\pm$0.144	24.45 $\pm$0.62	0.741 $\pm$0.007	0.261 $\pm$0.215
DyRAD	**26.53** $\pm$0.74	**0.363** $\pm$0.052	**0.536** $\pm$0.077	**3.25** $\pm$0.57	**0.236** $\pm$0.040	**0.917** $\pm$0.101	**24.51** $\pm$0.62	**0.744** $\pm$0.009	**0.641** $\pm$0.165

method to the recordings, render a trajectory shifted laterally by $2\mathrm{m}$, and train a new instance on those renders. Its predictions at the original poses and training timestamps are then compared with the real recordings; scores are therefore not directly comparable to the held-out on-path results. This tests whether scene information remains recoverable through displaced rendering and refitting; direct off-path accuracy is evaluated separately on synthetic data. DyRAD achieves the best means across the metrics in Table 2. On RADIal, it substantially outperforms the prior baselines in both reconstruction and detection, recovering detections in 91.7% of reference-detected objects, compared with 18.1% for the strongest prior baseline on this measure, and 20.9% for the static model. On Boreas, RadarSplat nearly matches our PSNR (24.29 versus 24.51 dB), yet its object-region RA correlation is substantially lower (0.211 versus 0.641). The object-reconstruction benefit observed on-path thus persists through displaced rendering and refitting, including in the baselines' native RA setting without Doppler supervision.

Synthetic Off-Path Evaluation. We directly assess off-path accuracy by fitting each method to five synthetic scenes and evaluating without refitting at lateral offsets of -1.75, $+1.75$, and $+3.5\mathrm{m}$ and yaw offsets of $+5^\circ$ and $+10^\circ$. Across the 25 scene–offset pairs, DyRAD achieves the highest mean correlation scores, globally and within objects. Object-region RAD correlation reaches 0.356, compared with 0.161 for DyRAD-static and 0.086 for RadarFields, the strongest prior baseline, demonstrating improved object reconstruction at displaced viewpoints. CD is nearly identical to RadarSplat's (5.13 versus $5.14\mathrm{m}$), but joint detection F1 is substantially higher (0.227 versus 0.106), indicating better agreement in location, Doppler, and power despite similar spatial distances. F1 is close to DyRAD-static (0.223), whereas foreground hit rate is much higher: 74.6%, compared with 22.4% for DyRAD-static and 36.0% for RadarSplat. However, this greater coverage comes with more reference-empty annotated boxes containing predicted detections: 4.9 on average, versus 0.7 and 3.4, respectively. Videos of all synthetic scenes and evaluated offsets appear in the supplementary HTML viewer; the generator and full results are described in Sec. C.4.

4.3 Additional Capabilities

Separating the scene from the sensor response enables sensor-configuration transfer without retraining. We demonstrate coarse-to-fine transfer using two configurations processed from the same raw ADC recordings. After fitting to coarse measurements, we replace the PSF and sampling grid with the finer configuration's, keeping scene parameters fixed. Against real fine-resolution measurements, direct rendering reduces CD from 7.27 to $3.20\mathrm{m}$ and increases F1 from 0.101 to 0.195 over linear upsampling of the same model's coarse renders. It also improves object-region RAD PSNR, although upsampling scores higher on most correlation and SSIM metrics. These detection gains without refitting suggest that the learned reflectors capture scene structure beyond the coarse measurement resolution. Details and full results appear in Appendix E. Additionally, the explicit scene and motion representation supports scene editing, demonstrated qualitatively in Fig. 1.

4.4 Ablation Studies

We ablate Doppler supervision, interpolation consistency, and PSF modeling under the RADIal on-path and off-path protocols. Complete results appear in Tables 12- 15 in Sec. D of the appendix.

Doppler supervision. Doppler supervision raises joint detection F1 from 0.144 to 0.179 on-path and from 0.120 to 0.236 off-path, while reducing object Doppler peak error from 1.843 to 0.732 bins and from 2.264 to 1.375 bins, respectively. Yet foreground hit rate remains nearly unchanged under both protocols: bounding-box initialization and RA-only fitting already recover vehicle coverage, while Doppler supervision refines the velocity structure of those returns. Object-region RAD correlation also increases from 0.504 to 0.658 on-path and from 0.389 to 0.536 off-path, at the cost of reductions of 0.055 and 0.033 in full-measurement RA correlation, respectively.

Interpolation consistency. On-path, the loss improves RA, RD, and RAD correlation by 0.219, 0.143, and 0.051 and precision from 0.162 to 0.233, with little change in recall or vehicle coverage. Off-path, it improves RA PSNR by 1.50 dB and SSIM by 0.109, but lowers RAD correlation (0.447 to 0.363) and F1 (0.296 to 0.236), both still above all baselines. The loss regularizes intermediate timestamps evaluated on-path, whereas off-path evaluation uses the training timestamps. We retain it to balance temporal interpolation and RA fidelity against off-path RAD fidelity.

PSF modeling. We compare the fixed sensor-derived PSF with Gaussian primitives with learned extent and point reflectors with learned PSF bandwidths. On-path, the fixed PSF more than doubles detection recall relative to either alternative, improving F1 at similar precision. Off-path, F1 reaches 0.236, versus 0.074 for learned extent and 0.124 for learned bandwidths. Across both protocols, the fixed PSF also improves object-region correlation in RA, RD, and RAD, despite lower full-RAD PSNR and SSIM. These results support fixing the sensor response rather than allowing learnable spread to compensate for errors in reflector placement and reflectivity.

5 Conclusion, Limitations, and Future Work

We presented DyRAD, a method for synthesizing complete RAD measurements of dynamic driving scenes from novel viewpoints. It combines static and motion-tracked point reflectors with a fixed sensor response, separating scene structure from measurement spreading. This formulation improves object reconstruction and detection in both on-path and off-path evaluations and enables sensor-configuration changes without refitting the scene. To our knowledge, it is the first radar NVS method to reconstruct dynamic driving scenes with Doppler measurements.

Several limitations remain. Reflector-to-object assignments rely on object annotations and remain fixed during optimization, even as tracks and reflector positions are refined. Jointly refining these assignments could help correct initialization errors, while modeling multipath and speckle noise could improve measurement realism. While our formulation supports 3D reconstruction, in this work we use a ground-plane implementation because the evaluated measurements lack elevation. Future work may evaluate elevation-resolving radar and explore transfer across physical sensors, beyond the processing configurations demonstrated here.

AI USE STATEMENT

We used generative AI tools to assist with writing and editing code, drafting and revising manuscript text, and reviewing related literature. The authors reviewed and revised the AI-assisted manuscript text and code. We take responsibility for the final content of this work, including all AI-assisted code, text, claims, and artifacts.

REPRODUCIBILITY STATEMENT

The appendix details sensor-specific PSF construction (Sec. A.1), scene initialization (Sec. B.1), densification (Sec. B.3), and hyperparameter selection, including search ranges and selected values (Sec. B.2). It also provides evaluation protocols, metric definitions, and additional experimental results (Sec. C), together with ablation settings (Sec. D) and configuration-transfer procedures (Sec. E). Upon acceptance, we will release the code with usage examples and experiment configurations, along with the synthetic dataset described in Sec. C.4 as a benchmark for off-path radar novel-view synthesis.

REFERENCES

Adnan Armouti, Yixuan Gao, and Rajalakshmi Nandakumar. mmir: Frequency-space inverse rendering for 3d millimeter-wave radar adc synthesis. *arXiv preprint arXiv:2608.28913*, 2026.

Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *ICCV*, 2021.

Oded Bialer and Yuval Haitman. Radsimreal: Bridging the gap between synthetic and real data in radar object detection with simulation. In *CVPR*, 2024.

Tomer Borreda, Fangqiang Ding, Sanja Fidler, Shengyu Huang, and Or Litany. Radargen: Automotive radar point cloud generation from cameras. In *ECCV*, 2026.

David Borts, Erich Liang, Tim Broedermann, Andrea Ramazzina, Stefanie Walz, Edoardo Palladin, Jipeng Sun, David Brueggemann, Christos Sakaridis, Luc Van Gool, et al. Radar fields: Frequency-space neural scene representations for fmcw radar. In *SIGGRAPH*, 2024.

Keenan Burnett, David J Yoon, Yuchen Wu, Andrew Z Li, Haowei Zhang, Shichen Lu, Jingxing Qian, Wei-Kang Tseng, Andrew Lambert, Keith YK Leung, et al. Boreas: A multi-season autonomous driving dataset. *The International Journal of Robotics Research*, 2023.

Chuhan Chen, Tianshu Huang, Akarsh Prabhakara, Chaithanya Kumar Mummadi, Zhongxiao Cong, Anthony Rowe, Matthew O'Toole, and Deva Ramanan. Radarsim: Simulating single-chip radar via multimodal neural fields. In *3DV*, 2026.

Ziyu Chen, Jiawei Yang, Jiahui Huang, Riccardo De Lutio, Janick Martinez Esturo, Boris Ivanovic, Or Litany, Zan Gojcic, Sanja Fidler, Marco Pavone, et al. Omnire: Omni urban scene reconstruction. In *ICLR*, 2025.

Fangqiang Ding, Xiangyu Wen, Yunzhou Zhu, Yiming Li, and Chris Xiaoxuan Lu. Radarocc: Robust 3d occupancy prediction with 4d imaging radar. *NeurIPS*, 2024.

Yuval Haitman and Oded Bialer. Doppdrive: Doppler-driven temporal aggregation for improved radar object detection. In *ICCV*, 2025.

Georg Hess, Carl Lindström, Maryam Fatemi, Christoffer Petersson, and Lennart Svensson. Splatad: Real-time lidar and camera rendering with 3d gaussian splatting for autonomous driving. In *CVPR*, 2025.

Shengyu Huang, Zan Gojcic, Zian Wang, Francis Williams, Yoni Kasten, Sanja Fidler, Konrad Schindler, and Or Litany. Neural lidar fields for novel view synthesis. In *ICCV*, 2023.

Tianshu Huang, John Miller, Akarsh Prabhakara, Tao Jin, Tarana Laroia, Zico Kolter, and Anthony Rowe. Dart: Implicit doppler tomography for radar novel view synthesis. In *CVPR*, 2024.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, George Drettakis, et al. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 2023.

Pou-Chun Kung, Skanda Harisha, Ram Vasudevan, Aline Eid, and Katherine A Skinner. Radarsplat: Radar gaussian splatting for high-fidelity data synthesis and 3d reconstruction of autonomous driving scenes. In *ICCV*, 2025.

Pou-Chun Kung, Yuan Tian, Zhengqin Li, Yue Liu, Eric Whitmire, Wolf Kienzle, and Hrvoje Benko. Radarsplat-rio: Indoor radar-inertial odometry with gaussian splatting-based radar bundle adjustment. *arXiv preprint arXiv:2604.13492*, 2026.

Li Ma, Xiaoyu Li, Jing Liao, Qi Zhang, Xuan Wang, Jue Wang, and Pedro V Sander. Deblur-nerf: Neural radiance fields from blurry images. In *CVPR*, 2022.

Sivabalan Manivasagam, Shenlong Wang, Kelvin Wong, Wenyuan Zeng, Mikita Sazanovich, Shuhan Tan, Bin Yang, Wei-Chiu Ma, and Raquel Urtasun. Lidarsim: Realistic lidar simulation by leveraging the real world. In *CVPR*, 2020.

Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *CVPR*, 2021.

Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 2021.

Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul P Srinivasan, and Jonathan T Barron. Nerf in the dark: High dynamic range view synthesis from noisy raw images. In *CVPR*, 2022.

Alexander Musiat, Laurenz Reichardt, Michael Schulze, and Oliver Wasenmüller. Radarpillars: Efficient object detection from 4d radar point clouds. In *2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC)*, 2024.

Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *CVPR*, 2022.

Dong-Hee Paek, Seung-Hyun Kong, and Kevin Tirta Wijaya. K-radar: 4d radar object detection for autonomous driving in various weather conditions. *Advances in Neural Information Processing Systems*, 2022.

Mahan Rafidashti, Ji Lan, Maryam Fatemi, Junsheng Fu, Lars Hammarstrand, and Lennart Svensson. Neuradar: Neural radiance fields for automotive radar point clouds. In *CVPR*, 2025.

Julien Rebut, Arthur Ouaknine, Waqas Malik, and Patrick Pérez. Raw high-definition radar for multi-task learning. In *CVPR*, 2022.

Konstantinos Rematas, Andrew Liu, Pratul P Srinivasan, Jonathan T Barron, Andrea Tagliasacchi, Thomas Funkhouser, and Vittorio Ferrari. Urban radiance fields. In *CVPR*, 2022.

Matthew Tancik, Vincent Casser, Xinchen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *CVPR*, 2022.

Xiao Tang, Guirong Zhuo, Cong Wang, Boyuan Zheng, Minqing Huang, Lianqing Zheng, Long Chen, and Shouyi Lu. 4dradar-gs: Self-supervised dynamic driving scene reconstruction with 4d radar. *arXiv preprint arXiv:2509.12931*, 2025.

Adam Tonderski, Carl Lindström, Georg Hess, William Ljungbergh, Lennart Svensson, and Christoffer Petersson. Neurad: Neural rendering for autonomous driving. In *CVPR*, 2024.

Haithem Turki, Jason Y Zhang, Francesco Ferroni, and Deva Ramanan. Suds: Scalable urban dynamic scenes. In *CVPR*, 2023.

Haithem Turki, Qi Wu, Xin Kang, Janick Martinez Esturo, Shengyu Huang, Ruilong Li, Zan Gojcic, and Riccardo de Lutio. Simuli: Real-time lidar and camera simulation with unscented transforms. In *ICLR*, 2026.

Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 2004.

Hanfeng Wu, Xingxing Zuo, Stefan Leutenegger, Or Litany, Konrad Schindler, and Shengyu Huang. Dynamic lidar re-simulation using compositional neural fields. In *CVPR*, 2024.

Yunzhi Yan, Haotong Lin, Chenxu Zhou, Weijie Wang, Haiyang Sun, Kun Zhan, Xianpeng Lang, Xiaowei Zhou, and Sida Peng. Street gaussians: Modeling dynamic urban scenes with gaussian splatting. In *ECCV*, 2024.

Jiawei Yang, Boris Ivanovic, Or Litany, Xinshuo Weng, Seung Wook Kim, Boyi Li, Tong Che, Danfei Xu, Sanja Fidler, Marco Pavone, et al. Emernerf: Emergent spatial-temporal scene decomposition via self-supervision. In *ICLR*, 2024.

Ze Yang, Yun Chen, Jingkang Wang, Sivabalan Manivasagam, Wei-Chiu Ma, Anqi Joyce Yang, and Raquel Urtasun. Unisim: A neural closed-loop sensor simulator. In *CVPR*, 2023.

Jiarui Zhang, Zhihao Li, Chong Wang, and Bihan Wen. Rf4d: Neural radar fields for novel view synthesis in outdoor dynamic scenes. In *CVPR*, 2026.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.

Zehan Zheng, Fan Lu, Weiyi Xue, Guang Chen, and Changjun Jiang. Lidar4d: Dynamic neural fields for novel space-time view lidar synthesis. In *CVPR*, 2024.

DyRAD Appendix

Contents

A Radar Signal Formation **13**

A.1 Sensor-Specific PSF Construction 13

B Scene Initialization and Optimization **14**

B.1 Scene Initialization Details 14

B.2 Hyperparameter Selection 14

B.3 Densification 15

C Additional Experimental Results **15**

C.1 Evaluation Protocol 15

C.2 On-Path Evaluation 16

C.3 Real-Data Off-Path Evaluation 16

C.4 Synthetic Off-Path Evaluation 18

C.5 Runtime 20

D Ablations **21**

E Sensor Configuration Transfer **21**

A Radar Signal Formation

A.1 Sensor-Specific PSF Construction

A localized reflector produces a response across multiple measurement bins because of the radar's acquisition and signal-processing chain. This spreading is a sensor property: changing the reconstructed scene should change the reflectors, not the response applied to them. We therefore construct the sensor response before scene optimization and hold it fixed. For RADIal, we use its known processing operations and calibration table. For Boreas, where these are unavailable, we use an approximate response with fixed parameters selected from sensor specifications and validation data. For the synthetic data we use the same formulation as in RADIal.

RADIal. We construct the kernels from RADIal's reference processing (Rebut et al., 2022). Range and Doppler use Hamming-windowed FFTs. We approximate their normalized power responses with the corresponding analytic Hamming-FFT kernel, evaluated over a five-bin neighborhood.

RADIal separates its transmitters through phase codes that shift their returns by multiples of 16 Doppler FFT bins. From the 256-bin spectrum, we gather the transmitter channels using RADIal's prescribed shifts and retain 16 reduced Doppler offsets. We do not disambiguate individual returns onto the original 256-bin grid. Gathering compensates for the transmitter shifts, leaving the Hamming-FFT response modeled by our Doppler kernel; Doppler indices wrap around the reduced grid.

For azimuth, we compute each source direction's response through RADIal's calibrated beamformer, including its antenna window, then square the magnitude and normalize the peak. This preserves the response's dependence on viewing direction. We use the same elevation slice as RADIal's reference RA processing (index 5, corresponding to $+1^\circ$) for both measurement preprocessing and kernel construction. All kernels remain fixed during scene optimization.

Boreas. The processing chain is unavailable, so we approximate the range and azimuth responses using fixed, tapered sinc-squared kernels:

$$h_u(x) = \mathrm{sinc}^2(x/w_u)\,\frac{1+\cos(\pi x/k_u)}{2}\,\mathbf{1}_{|x|\le k_u}, \qquad u \in \{R, A\}. \tag{4}$$

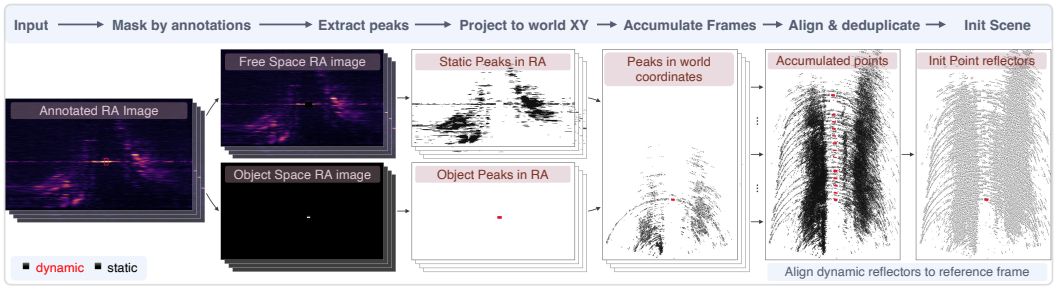


Figure 5: **Point-reflector initialization.** Object bounding boxes separate background and object regions. Selected returns are projected into world coordinates and accumulated across training frames. Dynamic returns are aligned by track translation and expressed in object coordinates. Deduplication yields the initial static and dynamic reflectors.

We use $(w_R, k_R) = (5.7, 16)$ and $(w_A, k_A) = (2.26, 4)$, in bins. The azimuth width is based on the nominal $1.8°$ beamwidth and $0.9°$ sampling interval; the taper yields an effective width of approximately $1.64°$. The range width is selected by validation reconstruction performance because the available specifications do not provide the range-response width. This is an empirical approximation, held fixed throughout scene optimization. Boreas provides no Doppler axis, so $h_D = 1$.

B Scene Initialization and Optimization

B.1 Scene Initialization Details

Figure 5 illustrates the initialization procedure. We initialize the scene using only training-frame measurements and bounding boxes.

Object association and peak extraction. We associate bounding boxes across annotated frames by extrapolating each bounding box from the previous annotated frame using a constant-velocity model and matching it to the current annotations. Each resulting track receives a unique object index.

The annotated bounding box size defines an approximate object footprint, which we dilate by the sensor PSF to obtain its measurement region. The static-background region is the complement of these regions. On RADIal, for the background, we threshold the Doppler-averaged RA map. For each dynamic object, we select the Doppler bin with the strongest response at its annotated location and threshold that slice within its region. For Boreas, we select returns directly from the RA maps, using annotated object regions for dynamic initialization and their complement for the static background. We project the selected cells onto the ground plane in world coordinates and accumulate them across frames. Cells in overlapping regions contribute to each corresponding object.

Dynamic-peak alignment. Let $\mathbf{z}_i(t)$ be a selected return associated with object j, and let $\mathbf{p}_{0,j}$ be its first training-frame annotated position. Using the initial track position $\mathbf{p}_j(t)$, we align the return and express it relative to the object-frame origin:

$$\boldsymbol{\mu}_i = \mathbf{z}_i(t) - \mathbf{p}_j(t). \tag{5}$$

This initialization compensates for translation only, approximating the alignment of returns from turning objects.

Deduplication and parameter initialization. We deduplicate static returns on a $0.5\,\mathrm{m}$ voxel grid and aligned dynamic returns on a $0.15\,\mathrm{m}$ grid, separately for each object. Static reflectors retain their world-space positions and receive $q_i = 0$; dynamic reflectors use $\boldsymbol{\mu}_i$ and receive $q_i = j$. We initialize the base reflected power to $\alpha_i = 0.2$ and all SH coefficients to zero, keeping the DC coefficient fixed.

B.2 Hyperparameter Selection

We select the interpolation-consistency weight λ_{int} and learning rates for reflector positions, opacities, spherical-harmonic coefficients, and object tracks separately for each dataset using held-out validation data. We use degree-three SH ($K = 16$). All remaining settings are fixed or data-derived.

Table 3: Search ranges and selected hyperparameters.

Hyperparameter	Range	RADIal	Boreas
Interpolation weight λ_{int}	10^{-3}–1	0.1	0.01
Position LR (`means_lr`)	10^{-5}–10^{-2}	0.0002	0.00005
Opacity LR (`opacities_lr`)	10^{-3}–1	0.02	0.02
SH LR (`sh_coeffs_lr`)	10^{-4}–10^{-1}	0.01	0.02
Motion LR (`deform_lr`)	10^{-4}–10^{-1}	0.01	0.0005

We tune five optimization hyperparameters. Each grid contains ten 1-2-5 logarithmic values. We sweep each hyperparameter independently while holding the remaining four at their pre-study values; the selections are therefore conditional rather than jointly optimal.

RADIal uses four held-out recordings. Boreas uses four held-out segments of a single annotated recording, each separated from evaluation segments by at least $1{,}000$ frames. Both datasets use every fifth frame for validation.

We select using a composite of RA correlation, RAD correlation, and peak F1, standardized across candidates within each window and averaged over the four windows. Boreas uses only RA correlation and peak F1, as it has no Doppler dimension. For RADIal, candidates whose object recall falls more than 0.05 below the best arm are disqualified.

We also compare per-window rank agreement. It confirms the interpolation-weight selection on RADIal and all five selections on Boreas. For the four RADIal learning rates, rank agreement disagreed or tied; after reviewing these results, we adopted the composite-score winners. Chosen values can be seen in table 3.

B.3 Densification

The final model uses residual-guided densification of the static background. Starting at step 1,000, we perform densification every 200 steps until step 2,300. At each event, we extract positive-residual peaks within the ego-motion-compensated static-Doppler band, project them into world coordinates, and remove duplicates and candidates within 1m of an existing reflector. We then add the 1,000 strongest remaining candidates as static reflectors. We use a residual threshold of 0.05 and limit the complete representation to 800,000 reflectors.

We evaluated densification in a 36-configuration sensitivity sweep on the same validation recordings used for hyperparameter selection, disjoint from the evaluation sequences. Performance was stable across the sweep, and we retain its central configuration for every setting.

C Additional Experimental Results

C.1 Evaluation Protocol

Implementation Details. We optimize DyRAD and DyRAD-static for 10k steps. For the baselines, we use their published training schedules on Boreas; on RADIal, we train RadarSplat for 10k steps and retain the published RadarFields schedule. All experiments run on a single NVIDIA GeForce RTX 4090.

Reconstruction evaluation. We evaluate reconstruction using PSNR, SSIM, LPIPS, and Pearson correlation ρ. Both datasets are evaluated in their native domain: RADIal is evaluated in the linear amplitude domain and Boreas in the sensor's log-compressed domain. PSNR, SSIM, and correlation are reported over full measurements and object regions; LPIPS is computed over full RA and RAD measurements only.

Approximately 97% of RADIal tensor bins lie at the noise floor, so agreement with the background can dominate full-measurement scores. To examine this effect, we evaluate two structureless diagnostic controls using the same normalization, held-out frames, and scoring procedure as the reconstruction methods. Const predicts the median value of each reference RA map uniformly, while Zero predicts zero everywhere.

Table 4: **Structureless controls.** Full-measurement RA reconstruction on held-out RADIal frames. Mean ± sample standard deviation across ten sequences; best means, including controls, are bold. Constant predictions are assigned $\rho = 0$ by the scorer.

Method	PSNR [dB] ↑	SSIM ↑	LPIPS ↓	ρ ↑
DyRAD	**25.78** ±1.09	0.690 ±0.025	**0.234** ±0.017	**0.612** ±0.031
DyRAD-static	25.16 ±1.15	0.686 ±0.025	0.249 ±0.021	0.533 ±0.055
RadarSplat	20.94 ±0.66	0.050 ±0.011	0.533 ±0.022	0.291 ±0.100
RadarFields	20.38 ±0.71	0.059 ±0.009	0.569 ±0.021	0.055 ±0.024
CONST	23.71 ±1.04	**0.723** ±0.023	0.583 ±0.025	0.000 ±0.000
ZERO	20.44 ±0.65	0.032 ±0.004	0.597 ±0.024	0.000 ±0.000

Despite containing no scene structure, CONST achieves higher SSIM than every reconstruction method and higher PSNR than both RadarSplat and RadarFields (Table 4). Thus, these metrics alone do not establish that a reconstruction preserves radar structure.

LPIPS ranks all reconstruction methods above the controls on RA, although its separation is much larger for DyRAD than for the prior baselines. Correlation measures agreement in spatial variation independently of an affine intensity transformation; it therefore complements PSNR without measuring intensity calibration. For constant predictions, where Pearson correlation is undefined, the scorer assigns zero.

We consequently interpret PSNR and SSIM alongside correlation and LPIPS, and report object-region scores to expose errors that full-measurement averages can conceal. Detection evaluation further tests whether synthesized measurements preserve usable radar returns.

Detection evaluation. We apply the CFAR detector from RADIal's original implementation, with unchanged settings, to both measured and synthesized tensors. This evaluates compatibility with the dataset's native detection pipeline without detector-specific tuning. We then evaluate the resulting point clouds using metrics from RadarGen (Borreda et al., 2026), with return power in place of calibrated RCS. Entire-area metrics comprise CD for spatial localization, CD-Full for joint distance in normalized position, Doppler, and power, and IoU@1 m for spatial overlap. DA precision, recall, and F1 require agreement in position and both radar attributes. Maximum Mean Discrepancy (MMD) measures distributional agreement separately for location, Doppler, and power, using five multi-scale RBF kernels and fixed-seed subsampling to at most 2,000 points. Doppler comparisons account for periodicity, with the DA Doppler threshold set to 10% of RADIal's Doppler period.

Foreground metrics are averaged over annotated boxes. Hit rate measures the fraction of boxes containing reference detections that also contain predicted detections; density similarity is the ratio of the smaller detection count to the larger, with value one when both are zero. Both penalize missing returns in reference-occupied boxes. We additionally report the number of boxes containing predictions but no reference detections (FP boxes).

We omit conditional foreground CD, CD-Full, and MMD: the baselines miss most objects, leaving these scores undefined or evaluated on a small detected subset. They therefore cannot support a comparison across methods. We instead report foreground hit rate and density similarity, which penalize missed objects. Tables report means and sample standard deviations across sequences, or scene–offset pairs for synthetic evaluation.

C.2 ON-PATH EVALUATION

Tables 5 and 6 provide the full reconstruction and detection results for the on-path evaluation described in Sec. 4.2. Figures 6 and 7 show qualitative comparisons on RADIal and Boreas, respectively.

C.3 REAL-DATA OFF-PATH EVALUATION

Tables 7 and 8 provide the full reconstruction and detection results for the real-data off-path protocol described in Sec. 4.2. Figures 8 and 9 illustrate the displaced renders and subsequent reconstructions

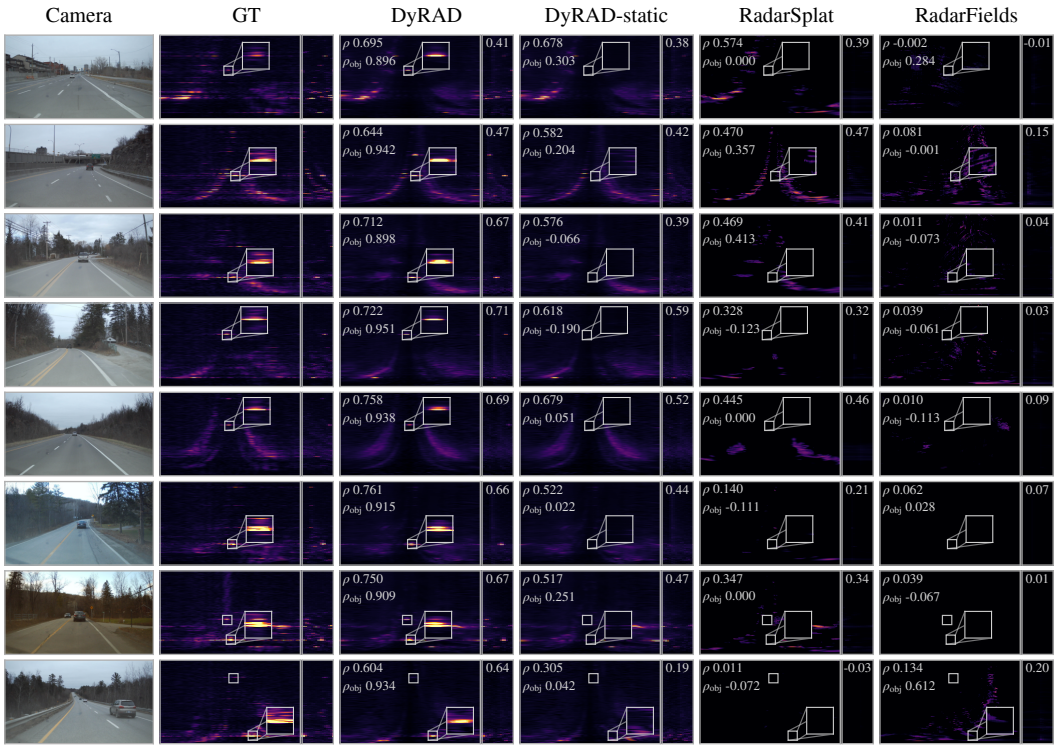


Figure 6: **Additional on-path examples on RADIal.** Each row compares a held-out measurement with the corresponding reconstructions, showing RA (left) and RD (right) projections within each panel. Displayed correlations are computed separately for each projection. White boxes mark objects; insets enlarge selected RA object regions. DyRAD preserves object returns that the static variant and baselines miss, even where the surrounding scene is reconstructed similarly.

Table 5: **On-path reconstruction: full results.** Mean ± sample standard deviation across ten RADIal sequences and three Boreas sequences. Best means within each projection are bold. * denotes the ego-Doppler lift. LPIPS is not reported for RD.

	Full				Object		
Method	$\rho \uparrow$	PSNR [dB] ↑	SSIM ↑	LPIPS ↓	$\rho \uparrow$	PSNR [dB] ↑	SSIM ↑
RADIal — RAD reconstruction							
RadarSplat*	0.068 ±0.028	23.56 ±0.85	0.215 ±0.006	0.634 ±0.008	−0.010 ±0.057	13.54 ±1.22	0.098 ±0.021
RadarFields*	0.015 ±0.009	23.13 ±1.23	0.203 ±0.012	0.637 ±0.007	−0.011 ±0.051	13.40 ±1.28	0.088 ±0.021
DyRAD-static	0.213 ±0.054	25.67 ±0.95	0.498 ±0.021	0.279 ±0.026	−0.013 ±0.072	13.76 ±1.24	0.199 ±0.035
DyRAD	**0.272** ±0.062	**25.88** ±0.95	**0.502** ±0.022	**0.275** ±0.025	**0.658** ±0.070	**16.97** ±1.06	**0.374** ±0.028
RADIal — RA reconstruction							
RadarSplat	0.291 ±0.100	20.94 ±0.66	0.050 ±0.011	0.533 ±0.022	−0.005 ±0.024	8.08 ±1.45	0.008 ±0.004
RadarFields	0.055 ±0.024	20.38 ±0.71	0.059 ±0.009	0.569 ±0.021	−0.004 ±0.063	8.18 ±1.43	0.019 ±0.007
DyRAD-static	0.533 ±0.055	25.16 ±1.15	0.686 ±0.025	0.249 ±0.021	0.021 ±0.088	8.70 ±1.56	0.232 ±0.058
DyRAD	**0.612** ±0.031	**25.78** ±1.09	**0.690** ±0.025	**0.234** ±0.017	**0.827** ±0.087	**14.93** ±1.89	**0.568** ±0.097
RADIal — RD reconstruction							
RadarSplat*	0.303 ±0.090	21.57 ±0.93	0.102 ±0.026	–	0.073 ±0.104	16.81 ±2.02	0.063 ±0.015
RadarFields*	0.131 ±0.062	21.44 ±0.93	0.129 ±0.036	–	0.019 ±0.063	16.87 ±2.06	0.098 ±0.043
DyRAD-static	0.375 ±0.064	24.78 ±1.68	0.611 ±0.038	–	0.073 ±0.084	18.56 ±2.39	0.438 ±0.111
DyRAD	**0.476** ±0.067	**25.33** ±1.70	**0.620** ±0.040	–	**0.564** ±0.100	**20.79** ±1.95	**0.558** ±0.086
Boreas — RA reconstruction							
RadarSplat	0.666 ±0.036	23.92 ±1.07	0.446 ±0.044	0.411 ±0.005	0.178 ±0.249	19.61 ±3.23	0.273 ±0.121
RadarFields	0.279 ±0.078	19.96 ±0.84	0.147 ±0.011	0.518 ±0.016	0.016 ±0.128	16.99 ±1.92	0.041 ±0.041
DyRAD-static	0.749 ±0.018	25.31 ±0.90	0.482 ±0.037	**0.304** ±0.011	0.278 ±0.217	21.13 ±2.89	0.372 ±0.087
DyRAD	**0.752** ±0.017	**25.36** ±0.90	**0.484** ±0.038	0.305 ±0.010	**0.622** ±0.130	**24.78** ±0.54	**0.565** ±0.129

Table 6: **On-path detection on RADIal: full results.** Mean ± sample standard deviation across ten sequences. Best means, including ties, are bold. * denotes the ego-Doppler lift.

	Entire area					
Method	CD [m] ↓	CD-Full ↓	IoU@1 m ↑	Precision ↑	Recall ↑	F1 ↑
RadarSplat*	6.53 ±5.82	0.1668 ±0.0758	0.141 ±0.066	0.084 ±0.036	0.087 ±0.056	0.073 ±0.048
RadarFields*	6.97 ±0.96	0.1612 ±0.0521	0.037 ±0.012	0.015 ±0.006	0.019 ±0.010	0.016 ±0.007
DyRAD-static	4.71 ±1.29	0.0783 ±0.0117	0.188 ±0.043	0.169 ±0.055	0.099 ±0.034	0.123 ±0.041
DyRAD	**4.01** ±1.02	**0.0689** ±0.0092	**0.238** ±0.040	**0.233** ±0.051	**0.149** ±0.037	**0.179** ±0.041

	Entire area			Foreground		
Method	MMD location ↓	MMD Doppler ↓	MMD power ↓	Hit rate ↑	Density sim. ↑	FP boxes ↓
RadarSplat*	0.9263 ±0.4340	1.1820 ±0.8409	5.6300 ±0.6544	0.036 ±0.067	0.048 ±0.048	**0.0** ±0.0
RadarFields*	0.5385 ±0.1566	1.3328 ±0.6162	4.7832 ±0.8515	0.269 ±0.245	0.153 ±0.121	**0.0** ±0.0
DyRAD-static	0.5842 ±0.2191	0.2386 ±0.1191	0.7471 ±0.2970	0.070 ±0.083	0.047 ±0.057	**0.0** ±0.0
DyRAD	**0.4208** ±0.1425	**0.1657** ±0.0741	**0.5933** ±0.2082	**0.907** ±0.086	**0.583** ±0.099	0.3 ±0.5

at the original poses. On RADIal, RadarFields produces a near-constant reconstruction on one sequence and an all-zero reconstruction on another; both are included in the reported averages.

C.4 Synthetic Off-Path Evaluation

Scene generation. We construct five scenes with crossing traffic, a curved road, sparse roadside structure, a dual carriageway, and an urban canyon. Static structures and vehicles comprise world-frame reflectors with positions, velocities, reflectivities, and surface normals. The scene evolves independently of the sensor trajectory, allowing the same world state to be rendered from different poses at each timestamp. Each sequence contains 40 frames at $0.1\,\mathrm{s}$ intervals.

The generator produces RADIal-format tensors with 512 range bins, 751 azimuth bins, and 16 Doppler bins, matching the real data; both are cropped identically to 447 range bins before training and scoring. These are generated together with sensor poses, ego velocities, and view-specific vehicle annotations. Measurements use analytic sensor responses in the amplitude domain and log-normal clutter. Doppler is wrapped to RADIal's $1.7968\,\mathrm{m/s}$ period; real DDMA demultiplexing

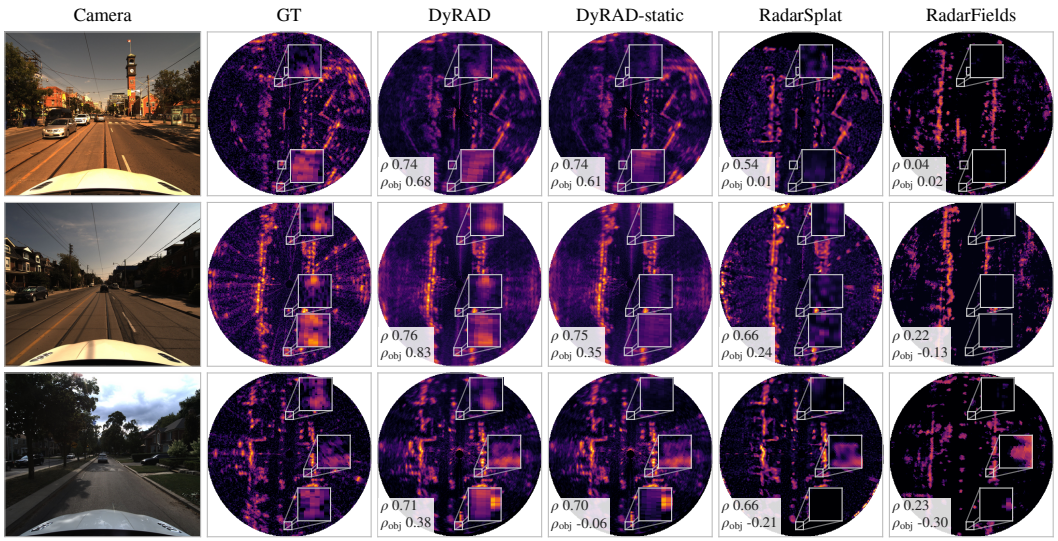


Figure 7: **On-path examples on Boreas.** Held-out RA measurements and reconstructions displayed in Cartesian bird's-eye view. White boxes mark annotated objects; insets enlarge selected object regions. DyRAD preserves object returns that the static variant and baselines miss, even where the surrounding scene is reconstructed similarly.

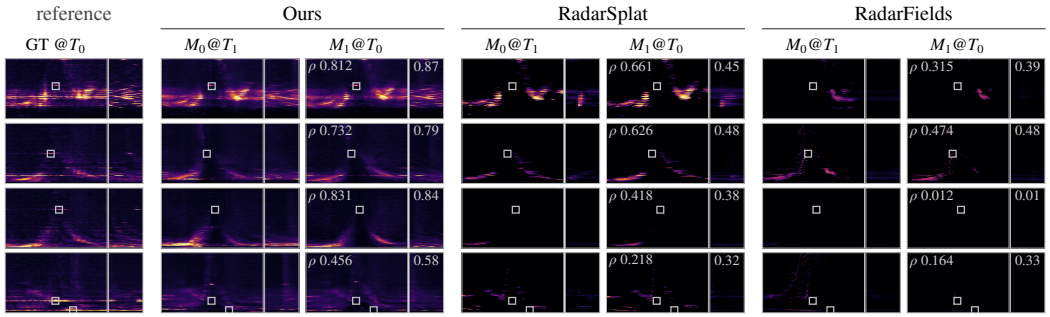


Figure 8: **Off-path rendering and round-trip reconstruction on RADIal.** Each panel shows RA (left) and RD (right) projections. For each method, $M_0@T_1$ shows measurements synthesized along the trajectory shifted laterally by 2 m. A new model is fitted to these renders and evaluated at the original poses ($M_1@T_0$), against the real measurements (GT@T_0). White boxes mark annotated vehicles. The values in the return panels report RA and RD correlations with the reference, respectively.

artifacts are not reproduced. The benchmark therefore provides ground truth under an idealized sensor model. We plan to publicly release the synthetic scenes and their off-path ground-truth measurements to support reuse of the benchmark.

Evaluation protocol. Each method is fitted to the base trajectory and evaluated without refitting at five displaced trajectories: lateral offsets of -1.75, $+1.75$, and $+3.5$ m, and yaw offsets of $+5°$ and $+10°$. Results are pooled over the 25 scene–offset pairs; the reported standard deviation describes variation across scenes and views, not training seeds.

Full results. Tables 9 and 10 provide the full reconstruction and detection results. For object-region RD metrics, the evaluation mask includes all Doppler bins at ranges occupied by annotated vehicles; collapsing azimuth therefore also includes background returns at those ranges. The supplementary HTML viewer provides videos of synthetic off-path sequences for all methods.

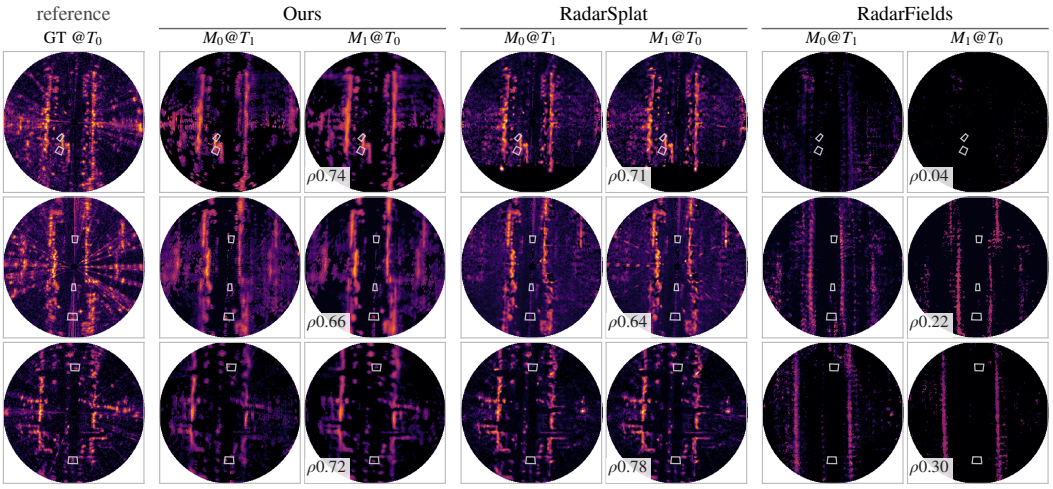


Figure 9: **Off-path rendering and round-trip reconstruction on Boreas.** RA measurements shown in Cartesian bird's-eye view. For each method, $M_0@T_1$ shows the initial model rendered along the trajectory shifted laterally by 2 m. A new model is fitted to these renders and evaluated at the original poses ($M_1@T_0$), against the real measurements (GT@T_0). White boxes mark annotated vehicles; ρ denotes correlation of the round-trip reconstruction with the reference.

Table 7: **Real-data off-path reconstruction: full results.** Results at the original poses after a 2 m lateral-shift round trip. Mean ± sample standard deviation across ten RADIal sequences and three Boreas sequences; best means are bold. * denotes the ego-Doppler lift. LPIPS is not reported for RD.

	Full				Object		
Method	ρ ↑	PSNR [dB] ↑	SSIM ↑	LPIPS ↓	ρ ↑	PSNR [dB] ↑	SSIM ↑
RADIal — RAD reconstruction							
RadarSplat*	0.076 ±0.032	23.92 ±0.85	0.217 ±0.006	0.634 ±0.008	−0.004 ±0.053	13.57 ±1.32	0.099 ±0.022
RadarFields*	0.023 ±0.018	24.43 ±0.75	0.217 ±0.006	0.638 ±0.008	−0.011 ±0.048	13.48 ±1.20	0.099 ±0.023
DyRAD-static	0.335 ±0.054	26.41 ±0.73	0.492 ±0.012	0.271 ±0.014	0.056 ±0.064	13.86 ±1.35	0.194 ±0.035
DyRAD	**0.363** ±0.052	**26.53** ±0.74	**0.498** ±0.012	**0.267** ±0.015	**0.536** ±0.077	**15.98** ±1.36	**0.344** ±0.041
RADIal — RA reconstruction							
RadarSplat	0.357 ±0.125	21.15 ±0.73	0.051 ±0.012	0.536 ±0.023	0.012 ±0.017	8.12 ±1.63	0.010 ±0.006
RadarFields	0.104 ±0.067	20.57 ±0.65	0.038 ±0.006	0.572 ±0.027	0.015 ±0.029	8.14 ±1.61	0.010 ±0.005
DyRAD-static	0.595 ±0.056	25.56 ±0.99	0.652 ±0.019	0.218 ±0.015	0.114 ±0.066	8.76 ±1.76	0.215 ±0.068
DyRAD	**0.661** ±0.035	**26.15** ±0.95	**0.665** ±0.023	**0.212** ±0.015	**0.807** ±0.057	**14.14** ±1.66	**0.517** ±0.089
RADIal — RD reconstruction							
RadarSplat*	0.301 ±0.115	21.50 ±0.94	0.091 ±0.023	–	0.063 ±0.098	16.93 ±2.05	0.058 ±0.014
RadarFields*	0.086 ±0.066	21.12 ±0.96	0.066 ±0.022	–	0.021 ±0.042	16.85 ±2.05	0.043 ±0.020
DyRAD-static	0.633 ±0.059	26.35 ±1.24	0.653 ±0.021	–	0.512 ±0.152	19.74 ±2.01	0.505 ±0.057
DyRAD	**0.653** ±0.048	**26.62** ±1.26	**0.661** ±0.022	–	**0.584** ±0.094	**20.56** ±1.97	**0.551** ±0.053
Boreas — RA reconstruction							
RadarSplat	0.689 ±0.042	24.29 ±1.16	**0.490** ±0.040	**0.415** ±0.004	0.211 ±0.237	19.25 ±3.51	0.289 ±0.123
RadarFields	0.538 ±0.081	21.09 ±1.28	0.466 ±0.044	0.550 ±0.007	0.095 ±0.123	16.54 ±2.30	0.296 ±0.069
DyRAD-static	0.741 ±0.007	24.45 ±0.62	0.359 ±0.023	0.462 ±0.024	0.261 ±0.215	20.02 ±2.54	0.296 ±0.065
DyRAD	**0.744** ±0.009	**24.51** ±0.62	0.360 ±0.020	0.461 ±0.023	**0.641** ±0.165	**24.19** ±1.39	**0.519** ±0.180

C.5 RUNTIME

We measure computational cost on RADIal using an otherwise idle NVIDIA RTX 4090 (Tab. 11). DyRAD requires a median of 7.2 minutes for 10k optimization steps and 5.65 ms to render a full RAD tensor, with 1055 MiB peak allocated GPU memory. Rendering measurements cover the complete forward pass after warm-up, excluding loading and file output. The baselines render RA maps, so their latency measures a different output. Training times follow each method's stated budget.

Table 8: **Real-data off-path detection on RADIal: full results.** Results at the original poses after a 2 m lateral-shift round trip. Mean ± sample standard deviation across ten sequences; best means are bold. * denotes the ego-Doppler lift.

	Entire area					
Method	CD [m] ↓	CD-Full ↓	IoU@1 m ↑	Precision ↑	Recall ↑	F1 ↑
RadarSplat*	6.35 ±4.75	0.1699 ±0.0702	0.155 ±0.076	0.085 ±0.042	0.093 ±0.061	0.077 ±0.050
RadarFields*	9.19 ±6.07	0.2211 ±0.0962	0.072 ±0.045	0.198 ±0.353	0.033 ±0.026	0.027 ±0.020
DyRAD-static	3.47 ±0.61	0.0681 ±0.0101	0.239 ±0.036	0.221 ±0.051	0.205 ±0.036	0.210 ±0.042
DyRAD	**3.25** ±0.57	**0.0629** ±0.0085	**0.268** ±0.034	**0.251** ±0.049	**0.227** ±0.034	**0.236** ±0.040

	Entire area			Foreground			
Method	MMD location ↓	MMD Doppler ↓	MMD power ↓	Hit rate ↑	Miss rate ↓	Density sim. ↑	FP boxes ↓
RadarSplat*	1.0231 ±0.5973	1.2247 ±0.7808	5.7598 ±0.6720	0.064 ±0.065	0.936 ±0.065	0.058 ±0.038	**0.0** ±0.0
RadarFields*	1.0864 ±0.4843	1.9832 ±1.1824	6.1715 ±0.7454	0.181 ±0.183	0.819 ±0.183	0.108 ±0.097	0.3 ±0.5
DyRAD-static	0.3965 ±0.1717	0.1867 ±0.0857	1.3521 ±0.2469	0.209 ±0.144	0.791 ±0.144	0.085 ±0.073	0.2 ±0.6
DyRAD	**0.3534** ±0.1606	**0.1749** ±0.0822	**1.2026** ±0.2107	**0.917** ±0.101	**0.083** ±0.101	**0.521** ±0.111	1.1 ±1.2

Table 9: **Synthetic off-path reconstruction.** Mean ± sample standard deviation across 25 scene–offset pairs; best displayed means are bold. * denotes the wrap-aware static ego-Doppler lift of RA predictions. LPIPS is reported only for full RA and RAD.

	Full				Object		
Method	ρ ↑	PSNR [dB] ↑	SSIM ↑	LPIPS ↓	ρ ↑	PSNR [dB] ↑	SSIM ↑
RAD reconstruction							
RadarSplat*	0.171 ±0.085	22.17 ±1.15	0.080 ±0.005	0.678 ±0.013	0.085 ±0.179	10.24 ±1.20	0.027 ±0.022
RadarFields*	0.035 ±0.024	22.39 ±1.18	0.075 ±0.001	0.703 ±0.013	0.086 ±0.171	10.44 ±0.99	0.022 ±0.005
DyRAD-static	0.297 ±0.158	23.94 ±0.58	0.290 ±0.014	0.522 ±0.047	0.161 ±0.270	11.14 ±1.32	0.121 ±0.045
DyRAD	**0.324** ±0.153	**24.06** ±0.55	**0.291** ±0.014	**0.516** ±0.045	**0.356** ±0.151	**11.89** ±1.49	**0.215** ±0.073
RA reconstruction							
RadarSplat	0.467 ±0.195	26.67 ±5.19	0.333 ±0.224	0.300 ±0.269	0.070 ±0.060	13.69 ±5.90	0.092 ±0.078
RadarFields	0.051 ±0.053	25.30 ±4.09	0.313 ±0.213	0.360 ±0.270	−0.005 ±0.018	13.22 ±5.49	0.028 ±0.026
DyRAD-static	0.514 ±0.239	28.38 ±4.45	0.635 ±0.215	0.255 ±0.216	0.106 ±0.143	13.80 ±5.72	0.103 ±0.071
DyRAD	**0.574** ±0.227	**28.85** ±4.54	**0.640** ±0.215	**0.247** ±0.217	**0.593** ±0.237	**17.39** ±3.38	**0.542** ±0.126
RD reconstruction							
RadarSplat*	0.299 ±0.154	21.78 ±4.38	0.142 ±0.094	–	0.209 ±0.140	20.42 ±5.11	0.091 ±0.065
RadarFields*	0.057 ±0.046	21.38 ±4.20	0.124 ±0.088	–	0.034 ±0.034	20.19 ±4.94	0.066 ±0.046
DyRAD-static	0.492 ±0.185	23.90 ±4.04	0.530 ±0.189	–	0.333 ±0.127	22.00 ±4.93	0.389 ±0.132
DyRAD	**0.512** ±0.177	**24.19** ±4.15	**0.539** ±0.192	–	**0.365** ±0.097	**22.80** ±4.36	**0.457** ±0.139

The two splatting methods have similar rendering latencies, with DyRAD's modest overhead consistent with motion evaluation and full-RAD synthesis. RadarFields has higher rendering latency, consistent with its neural-field rendering approach.

D Ablations

We evaluate the ablations using the RADIal on-path and off-path protocols, applying each ablation in both fitting stages of the off-path cycle. We report mean ± sample standard deviation across ten sequences. Tables 12 and 13 report all on-path ablations; Tables 14 and 15 report the corresponding off-path results.

E Sensor Configuration Transfer

We evaluate coarse-to-fine sensor-configuration transfer against linear upsampling of the same model's coarse renders. Both measurement configurations are derived from the same RADIal raw ADC recordings; the scene and acquisition platform remain unchanged.

Table 10: **Synthetic off-path detection.** Mean ± sample standard deviation across 25 scene–offset pairs; best displayed means are bold. * denotes the wrap-aware static ego-Doppler lift.

	Entire area					
Method	CD [m] ↓	CD-Full ↓	IoU@1 m ↑	Precision ↑	Recall ↑	F1 ↑
RadarSplat*	5.14 ±0.69	0.1773 ±0.0351	0.185 ±0.093	0.090 ±0.051	0.132 ±0.091	0.106 ±0.065
RadarFields*	8.06 ±0.86	0.2104 ±0.0438	0.099 ±0.064	0.022 ±0.019	0.028 ±0.026	0.024 ±0.022
DyRAD-static	**5.08** ±1.03	0.0885 ±0.0237	0.230 ±0.095	0.330 ±0.131	0.172 ±0.093	0.223 ±0.109
DyRAD	5.13 ±1.01	**0.0884** ±0.0231	**0.236** ±0.098	**0.348** ±0.129	**0.174** ±0.094	**0.227** ±0.109

	Entire area			Foreground			
Method	MMD location ↓	MMD Doppler ↓	MMD power ↓	Hit rate ↑	Miss rate ↓	Density sim. ↑	FP boxes ↓
RadarSplat*	**0.2278** ±0.0597	0.6362 ±0.3841	5.1296 ±1.0148	0.360 ±0.238	0.640 ±0.238	0.244 ±0.153	3.4 ±5.0
RadarFields*	0.3196 ±0.0521	0.7134 ±0.5546	5.9267 ±1.0330	0.305 ±0.199	0.695 ±0.199	0.220 ±0.111	5.0 ±7.4
DyRAD-static	0.5109 ±0.2132	**0.6353** ±0.3296	1.8323 ±1.0790	0.224 ±0.132	0.776 ±0.132	0.240 ±0.214	**0.7** ±0.8
DyRAD	0.5345 ±0.2215	0.6398 ±0.2961	**1.7327** ±1.1115	**0.746** ±0.222	**0.254** ±0.222	**0.341** ±0.155	4.9 ±4.5

Table 11: **Computational cost on RADIal.** Median training time, rendering latency, and peak allocated GPU memory on an RTX 4090.

Method	Training budget	Train [min]	Render [ms/frame]	Peak [MiB]	Rendered output
RadarSplat	10k steps	4.3	4.68	2138	RA map
RadarFields	800 iterations/image	1.4	15.10	4669	RA map
DyRAD	10k steps	7.2	5.65	1055	RAD tensor

Configurations. The native (fine) configuration uses 512 ADC samples per chirp and 256 chirps per coherent processing interval. The coarse configuration retains 128 samples per chirp and 128 consecutive chirps, yielding fourfold coarser range resolution and twofold coarser Doppler resolution. The native and coarse RAD tensors contain $16 \times 447 \times 751$ and $8 \times 112 \times 751$ bins after range cropping, respectively. The azimuth grid and Doppler wrap period of $1.7968\,\mathrm{m/s}$ remain unchanged. Both sets of measurements are processed directly from raw ADC; the coarse measurements are not obtained by resizing native tensors.

Training and rendering. We fit a model to the coarse measurements for each of the ten sequences in the RADIal on-path split, using 10k optimization steps. Initialization uses only the coarse training measurements, with configuration-specific normalization. For direct rendering, we replace the coarse PSF and sampling grid with those of the native configuration, keeping all reflector and motion parameters fixed. For linear upsampling, we render the same model under the coarse configuration and interpolate along range and Doppler, using circular interpolation over the shared Doppler wrap period.

Evaluation. Both outputs are evaluated against the same real held-out native measurements using the same scorer. Table 16 reports reconstruction and detection results as means and sample standard deviations across the ten sequences. Figure 10 compares the two rendering routes with the native measurements.

Table 12: **Ablations on-path on RADIal: reconstruction.** Mean ± sample standard deviation across ten sequences; best displayed means, including ties, are bold.

	Full				Object		
Method	$\rho \uparrow$	PSNR [dB] ↑	SSIM ↑	LPIPS ↓	$\rho \uparrow$	PSNR [dB] ↑	SSIM ↑
RAD reconstruction							
w/o Doppler	0.226 ±0.030	24.889 ±0.832	0.448 ±0.017	0.283 ±0.015	0.504 ±0.125	15.374 ±1.165	0.292 ±0.046
w/o $\mathcal{L}_{\text{int}}$	0.221 ±0.089	23.988 ±1.510	0.511 ±0.024	0.277 ±0.019	**0.667** ±0.087	17.038 ±1.082	**0.383** ±0.039
Learned extent	0.263 ±0.020	26.114 ±0.977	0.535 ±0.018	0.567 ±0.012	0.610 ±0.090	16.801 ±1.322	0.360 ±0.049
Learned PSF	**0.305** ±0.027	**26.607** ±0.940	**0.547** ±0.012	0.392 ±0.017	0.636 ±0.089	**17.071** ±1.074	0.336 ±0.047
DyRAD	0.272 ±0.062	25.884 ±0.951	0.502 ±0.022	**0.275** ±0.025	0.658 ±0.070	16.966 ±1.056	0.374 ±0.028
RA reconstruction							
w/o Doppler	**0.667** ±0.032	**26.439** ±1.016	0.682 ±0.024	**0.217** ±0.016	**0.834** ±0.056	**15.818** ±0.885	**0.616** ±0.040
w/o $\mathcal{L}_{\text{int}}$	0.393 ±0.106	21.395 ±1.806	0.566 ±0.023	0.243 ±0.012	0.790 ±0.114	14.036 ±1.939	0.465 ±0.120
Learned extent	0.600 ±0.046	25.707 ±1.109	**0.721** ±0.018	0.477 ±0.017	0.783 ±0.075	13.415 ±0.987	0.494 ±0.066
Learned PSF	0.645 ±0.039	26.267 ±1.059	0.712 ±0.020	0.346 ±0.018	0.788 ±0.081	12.554 ±1.184	0.438 ±0.055
DyRAD	0.612 ±0.031	25.780 ±1.091	0.690 ±0.025	0.234 ±0.017	0.827 ±0.087	14.927 ±1.889	0.568 ±0.097
RD reconstruction							
w/o Doppler	0.418 ±0.041	24.759 ±1.701	0.547 ±0.033	–	0.440 ±0.135	19.854 ±2.175	0.489 ±0.059
w/o $\mathcal{L}_{\text{int}}$	0.333 ±0.135	21.824 ±1.721	0.554 ±0.041	–	0.510 ±0.177	19.669 ±2.453	0.487 ±0.094
Learned extent	0.460 ±0.043	25.422 ±1.786	**0.663** ±0.035	–	0.500 ±0.111	19.677 ±2.273	0.539 ±0.104
Learned PSF	**0.493** ±0.033	**25.717** ±1.770	0.649 ±0.032	–	0.502 ±0.104	19.863 ±2.148	0.530 ±0.093
DyRAD	0.476 ±0.067	25.331 ±1.700	0.620 ±0.040	–	**0.564** ±0.100	**20.786** ±1.948	**0.558** ±0.086

Table 13: **Ablations on-path on RADIal: detection and Doppler.** Mean ± sample standard deviation across ten sequences; best displayed means, including ties, are bold. Doppler MAE is peak-location error at annotated cars, in bins; it is reported here only for the Doppler-supervision comparison.

	Entire area						Doppler
Method	CD [m] ↓	CD-Full ↓	IoU@1 m ↑	Precision ↑	Recall ↑	F1 ↑	MAE ↓
w/o Doppler	**3.663** ±0.660	**0.069** ±0.006	**0.240** ±0.049	0.142 ±0.025	**0.153** ±0.031	0.144 ±0.021	1.843 ±0.585
w/o $\mathcal{L}_{\text{int}}$	3.974 ±0.877	0.070 ±0.010	0.193 ±0.057	0.162 ±0.065	0.152 ±0.058	0.155 ±0.061	–
Learned extent	7.468 ±2.520	0.100 ±0.015	0.126 ±0.039	0.224 ±0.052	0.067 ±0.024	0.101 ±0.032	–
Learned PSF	6.302 ±1.762	0.114 ±0.025	0.112 ±0.032	**0.243** ±0.031	0.058 ±0.022	0.089 ±0.028	–
DyRAD	4.008 ±1.021	**0.069** ±0.009	0.238 ±0.040	0.233 ±0.051	0.149 ±0.037	**0.179** ±0.041	**0.732** ±0.290

	Entire area			Foreground			
Method	MMD location ↓	MMD Doppler ↓	MMD power ↓	Hit rate ↑	Miss rate ↓	Density sim. ↑	FP boxes ↓
w/o Doppler	0.541 ±0.254	0.471 ±0.166	0.834 ±0.301	0.909 ±0.099	0.091 ±0.099	0.530 ±0.107	**0.200** ±0.422
w/o $\mathcal{L}_{\text{int}}$	0.496 ±0.254	**0.095** ±0.050	0.582 ±0.197	**0.910** ±0.119	**0.090** ±0.119	0.547 ±0.113	**0.200** ±0.422
Learned extent	0.545 ±0.223	0.133 ±0.049	**0.405** ±0.132	0.842 ±0.130	0.158 ±0.130	0.518 ±0.130	**0.200** ±0.422
Learned PSF	0.694 ±0.361	0.529 ±0.215	1.001 ±0.351	0.891 ±0.118	0.109 ±0.118	0.396 ±0.102	**0.200** ±0.422
DyRAD	**0.421** ±0.142	0.166 ±0.074	0.593 ±0.208	0.907 ±0.086	0.093 ±0.086	**0.583** ±0.099	0.300 ±0.483

Table 14: **Ablations off-path on RADIal: reconstruction.** Results after the $2\,\mathrm{m}$ lateral-shift cycle. Mean $\pm$ sample standard deviation across ten sequences; best displayed means, including ties, are bold.

	Full				Object		
Method	$\rho\uparrow$	PSNR [dB] $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	$\rho\uparrow$	PSNR [dB] $\uparrow$	SSIM $\uparrow$
RAD reconstruction							
w/o Doppler	$0.208\,\pm0.020$	$24.676\,\pm0.745$	$0.419\,\pm0.015$	$0.295\,\pm0.012$	$0.389\,\pm0.074$	$14.610\,\pm1.271$	$0.264\,\pm0.034$
w/o $\mathcal{L}_{\text{int}}$	$\mathbf{0.447}\,\pm0.062$	$\mathbf{27.153}\,\pm0.913$	$0.548\,\pm0.016$	$\mathbf{0.252}\,\pm0.017$	$\mathbf{0.549}\,\pm0.046$	$15.890\,\pm1.401$	$\mathbf{0.356}\,\pm0.053$
Learned extent	$0.313\,\pm0.031$	$26.712\,\pm0.880$	$\mathbf{0.551}\,\pm0.020$	$0.609\,\pm0.013$	$0.424\,\pm0.080$	$15.574\,\pm1.204$	$0.307\,\pm0.045$
Learned PSF	$0.349\,\pm0.026$	$26.938\,\pm0.875$	$0.539\,\pm0.011$	$0.381\,\pm0.010$	$0.488\,\pm0.116$	$15.849\,\pm1.213$	$0.318\,\pm0.049$
DyRAD	$0.363\,\pm0.052$	$26.534\,\pm0.736$	$0.498\,\pm0.012$	$0.267\,\pm0.015$	$0.536\,\pm0.077$	$\mathbf{15.977}\,\pm1.357$	$0.344\,\pm0.041$
RA reconstruction							
w/o Doppler	$\mathbf{0.694}\,\pm0.041$	$\mathbf{26.762}\,\pm1.210$	$0.675\,\pm0.021$	$0.232\,\pm0.018$	$\mathbf{0.835}\,\pm0.040$	$\mathbf{15.567}\,\pm0.778$	$\mathbf{0.595}\,\pm0.042$
w/o $\mathcal{L}_{\text{int}}$	$0.655\,\pm0.054$	$24.655\,\pm0.928$	$0.556\,\pm0.023$	$\mathbf{0.203}\,\pm0.014$	$0.818\,\pm0.049$	$13.590\,\pm1.775$	$0.438\,\pm0.083$
Learned extent	$0.640\,\pm0.024$	$26.190\,\pm1.056$	$\mathbf{0.729}\,\pm0.017$	$0.532\,\pm0.022$	$0.732\,\pm0.083$	$12.712\,\pm1.390$	$0.413\,\pm0.082$
Learned PSF	$0.665\,\pm0.022$	$26.419\,\pm1.043$	$0.698\,\pm0.020$	$0.344\,\pm0.016$	$0.761\,\pm0.075$	$12.581\,\pm1.506$	$0.409\,\pm0.079$
DyRAD	$0.661\,\pm0.035$	$26.150\,\pm0.951$	$0.665\,\pm0.023$	$0.212\,\pm0.015$	$0.807\,\pm0.057$	$14.138\,\pm1.662$	$0.517\,\pm0.089$
RD reconstruction							
w/o Doppler	$0.388\,\pm0.037$	$24.577\,\pm1.485$	$0.519\,\pm0.030$	–	$0.305\,\pm0.082$	$18.988\,\pm2.233$	$0.415\,\pm0.069$
w/o $\mathcal{L}_{\text{int}}$	$\mathbf{0.746}\,\pm0.029$	$25.604\,\pm1.008$	$0.596\,\pm0.032$	–	$\mathbf{0.679}\,\pm0.073$	$20.264\,\pm1.756$	$0.526\,\pm0.051$
Learned extent	$0.550\,\pm0.034$	$26.010\,\pm1.690$	$\mathbf{0.689}\,\pm0.032$	–	$0.523\,\pm0.114$	$19.682\,\pm2.194$	$\mathbf{0.551}\,\pm0.092$
Learned PSF	$0.585\,\pm0.024$	$26.203\,\pm1.705$	$0.671\,\pm0.029$	–	$0.545\,\pm0.099$	$19.936\,\pm2.174$	$0.537\,\pm0.087$
DyRAD	$0.653\,\pm0.048$	$\mathbf{26.616}\,\pm1.256$	$0.661\,\pm0.022$	–	$0.584\,\pm0.094$	$\mathbf{20.561}\,\pm1.973$	$\mathbf{0.551}\,\pm0.053$

Table 15: **Ablations off-path on RADIal: detection and Doppler.** Results after the $2\,\mathrm{m}$ lateral-shift cycle. Mean $\pm$ sample standard deviation across ten sequences; best displayed means, including ties, are bold. Doppler MAE is peak-location error at annotated cars, in bins; it is reported here only for the Doppler-supervision comparison.

	Entire area						Doppler
Method	CD [m] $\downarrow$	CD-Full $\downarrow$	IoU@1 m $\uparrow$	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	MAE $\downarrow$
w/o Doppler	$3.586\,\pm0.447$	$0.068\,\pm0.005$	$0.210\,\pm0.045$	$0.109\,\pm0.021$	$0.139\,\pm0.023$	$0.120\,\pm0.018$	$2.264\,\pm0.511$
w/o $\mathcal{L}_{\text{int}}$	$\mathbf{2.676}\,\pm0.565$	$\mathbf{0.063}\,\pm0.010$	$\mathbf{0.335}\,\pm0.057$	$\mathbf{0.292}\,\pm0.055$	$\mathbf{0.305}\,\pm0.062$	$\mathbf{0.296}\,\pm0.057$	–
Learned extent	$9.419\,\pm3.427$	$0.113\,\pm0.020$	$0.091\,\pm0.042$	$0.192\,\pm0.077$	$0.047\,\pm0.023$	$0.074\,\pm0.034$	–
Learned PSF	$5.220\,\pm1.232$	$0.094\,\pm0.016$	$0.142\,\pm0.031$	$0.241\,\pm0.051$	$0.087\,\pm0.023$	$0.124\,\pm0.031$	–
DyRAD	$3.246\,\pm0.569$	$\mathbf{0.063}\,\pm0.008$	$0.268\,\pm0.034$	$0.251\,\pm0.049$	$0.227\,\pm0.034$	$0.236\,\pm0.040$	$\mathbf{1.375}\,\pm0.431$

	Entire area			Foreground			
Method	MMD location $\downarrow$	MMD Doppler $\downarrow$	MMD power $\downarrow$	Hit rate $\uparrow$	Miss rate $\downarrow$	Density sim. $\uparrow$	FP boxes $\downarrow$
w/o Doppler	$0.496\,\pm0.266$	$0.436\,\pm0.170$	$1.027\,\pm0.299$	$0.927\,\pm0.071$	$0.073\,\pm0.071$	$\mathbf{0.532}\,\pm0.105$	$1.000\,\pm1.247$
w/o $\mathcal{L}_{\text{int}}$	$\mathbf{0.295}\,\pm0.222$	$\mathbf{0.080}\,\pm0.035$	$1.588\,\pm0.376$	$\mathbf{0.944}\,\pm0.083$	$\mathbf{0.056}\,\pm0.083$	$0.450\,\pm0.118$	$1.700\,\pm1.889$
Learned extent	$0.794\,\pm0.349$	$0.117\,\pm0.021$	$\mathbf{0.370}\,\pm0.082$	$0.587\,\pm0.255$	$0.413\,\pm0.255$	$0.328\,\pm0.143$	$\mathbf{0.200}\,\pm0.422$
Learned PSF	$0.561\,\pm0.247$	$0.341\,\pm0.150$	$1.287\,\pm0.341$	$0.776\,\pm0.234$	$0.224\,\pm0.234$	$0.368\,\pm0.115$	$0.800\,\pm0.789$
DyRAD	$0.353\,\pm0.161$	$0.175\,\pm0.082$	$1.203\,\pm0.211$	$0.917\,\pm0.101$	$0.083\,\pm0.101$	$0.521\,\pm0.111$	$1.100\,\pm1.197$

Table 16: **Coarse-to-fine configuration transfer versus linear upsampling.** Both outputs use the same coarse-trained model and are evaluated against held-out native measurements. Mean ± sample standard deviation across ten RADIal sequences. Object metrics use annotated vehicle regions. Best displayed means are bold. LPIPS is not reported for RD.

	Full			Object		
Method	ρ ↑	SSIM ↑	LPIPS ↓	ρ ↑	PSNR [dB] ↑	SSIM ↑
RAD reconstruction						
Linear upsampling	**0.329** ±0.021	**0.595** ±0.014	0.540 ±0.008	**0.556** ±0.044	12.87 ±2.70	**0.315** ±0.058
Direct rendering	0.201 ±0.037	0.461 ±0.012	**0.263** ±0.014	0.522 ±0.103	**15.32** ±1.10	0.301 ±0.040
RA reconstruction						
Linear upsampling	**0.624** ±0.033	**0.726** ±0.019	0.472 ±0.012	**0.629** ±0.056	10.30 ±2.08	**0.297** ±0.084
Direct rendering	0.501 ±0.035	0.549 ±0.015	**0.166** ±0.011	0.528 ±0.144	**11.15** ±1.03	0.285 ±0.070
RD reconstruction						
Linear upsampling	**0.516** ±0.046	**0.666** ±0.040	–	**0.513** ±0.058	**19.82** ±2.35	**0.536** ±0.097
Direct rendering	0.361 ±0.055	0.514 ±0.032	–	0.380 ±0.092	18.87 ±2.01	0.431 ±0.075

Detection					
Method	CD [m] ↓	IoU@1 m ↑	F1 ↑	Hit rate ↑	Density sim. ↑
Linear upsampling	7.27 ±1.71	0.103 ±0.035	0.101 ±0.038	0.906 ±0.124	0.339 ±0.077
Direct rendering	**3.20** ±0.54	**0.281** ±0.035	**0.195** ±0.046	**0.960** ±0.074	**0.593** ±0.076

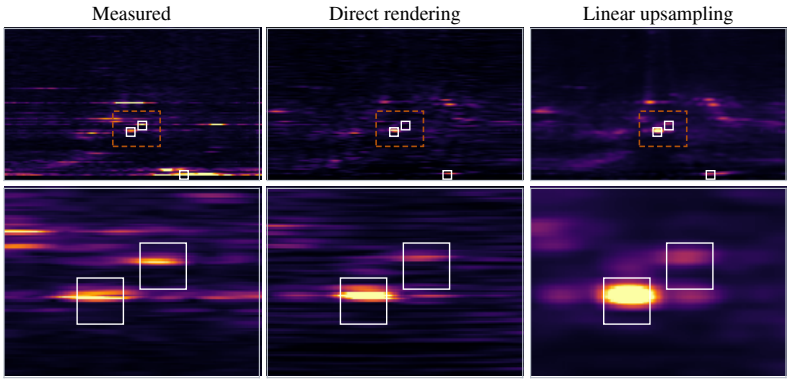


Figure 10: **Changing the measurement configuration of a reconstructed scene.** Left: measured native RA projection. Middle: the coarse-trained scene rendered with the native PSF and sampling grid. Right: linear upsampling of the same scene's coarse render. The bottom row enlarges the dashed region; white boxes mark vehicles.