

X-band Radar Non-Line-of-Sight Imaging

Dongyu Du^{1,2*} Mingkun Zhao^{1*} Yutong Yang³ Dominik Scheuble³ Xiaolong Huang¹
Zijian Shao¹ Mario Bijelic^{1,4} Kaushik Sengupta¹ Felix Heide^{1,4}

¹ Princeton University ² University of Toronto ³ Mercedes-Benz AG ⁴ Torc Robotics

Abstract

Conventional imaging systems capture objects visible in the direct line-of-sight (LOS). A decade of research on non-line-of-sight (NLOS) imaging approaches has made it possible to reconstruct hidden geometry outside the line of sight by analyzing indirect light transport. However, most existing methods operate in the optical visible or IR range. Relying on diffuse inter-reflections, every bounce incurs a quadratic intensity falloff. As such, with illumination power limited by eye-safety limitations, existing methods are fundamentally restricted to short ranges on the order of a few meters. We propose an X-band radar-based NLOS imaging method that leverages the long wavelength to convert diffuse reflections into predominantly specular ones, allowing for large-scale hidden-scene perception. We develop a neural reconstruction method that combines a learned dense prediction module and a geometry-aware NLOS reconstruction module, tackling the inherently low spatial resolution of long-wavelength radar. We assess our method with a prototype system and in simulation. Synthetic validation shows that, under the same transmit power, X-band radar achieves $10\times$ longer NLOS reconstruction range than optical systems, while experimental results further demonstrate accurate hidden-object reconstructions up to 40 m, establishing a practical pathway toward real-world long-range NLOS sensing.

1. Introduction

Human vision and conventional imaging systems are fundamentally constrained by their reliance on direct line-of-sight visibility, leaving occluded regions inaccessible. Extending perception beyond human capabilities via non-line-of-sight imaging enables us to “see around corners” [4, 5, 61]. In theory, this offers unprecedented sensing capabilities for remote sensing [5], medical imaging, and especially autonomous driving [61], where detecting approaching traffic and monitoring otherwise unreachable areas can provide

*indicates joint first authorship.




Figure 1. **Wavelength-dependent Energy Transport in NLOS Imaging.** Typical relay-wall roughness ($\sigma \approx 10^{-4}$ – 10^1 mm) yields mostly specular interactions at 10 GHz ($\lambda_1 \approx 30$ mm), directing stronger reflected energy into hidden regions, but predominantly diffuse scattering at 77 GHz ($\lambda_2 \approx 3.9$ mm), resulting in much weaker indirect illumination.

significant benefits.

Over the past decade, extensive NLOS imaging research [15, 23, 33, 53, 68, 73, 82] has demonstrated the ability to peek into such occluded areas. Thereby, the hidden geometry is inferred by reconstructing light transport through diffuse inter-reflections, as shown in Fig. 1. However, each diffuse bounce incurs roughly quadratic attenuation, shrinking the returning signal faster than the noise can be averaged. Limited to eye-safe power, visible/IR NLOS systems operate over only a few meters [22, 38, 49]. As a result, optical NLOS techniques are impractical for large-scale or outdoor scenarios that require long-range sensing.

To mitigate the attenuation of optical signals, existing methods have extended NLOS imaging into the millimeter-wave regime [9, 33, 61, 65, 73, 75], typically operating at 77 GHz ($\lambda = 3.9$ mm) [33, 61, 73, 75]. However, at 77 GHz, the wavelength is comparable to the surface rough-

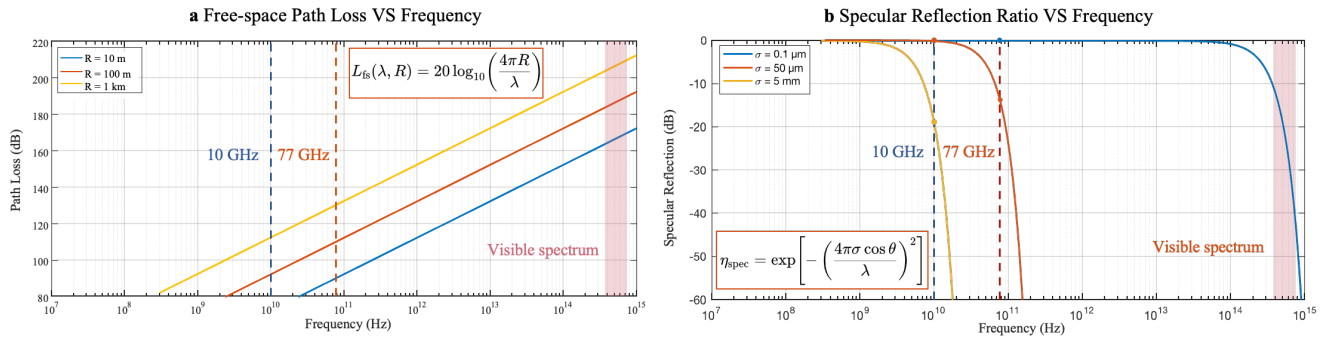


Figure 2. **Frequency-Dependent Propagation and Reflection Characteristics.** (a) Free-space path loss versus frequency at 10, 100, and 1000 m, showing stronger attenuation at 77 GHz and optical/THz bands. (b) Specular reflection ratio versus frequency for surface roughness $\sigma = 0.1 \mu\text{m}$, $50 \mu\text{m}$, and 5mm , where 10 GHz remains largely mirror-like while 77 GHz becomes diffuse.

ness of typical relay walls (10^{-4} – 10^1 mm), resulting in strong diffuse reflections, as illustrated in Fig. 1.

To extend the range further, we propose a lower-frequency X-band radar operating at 10 GHz ($\lambda = 30$ mm). As shown in Fig. 2(a), the longer wavelength substantially reduces free-space path loss by about 20 dB relative to typical 77 GHz mmWave systems and by 150 dB compared to optical sensors. In addition, longer wavelengths strongly suppress rough-surface scattering (smaller σ), yielding a higher specular-to-diffuse ratio [2, 29, 70], as illustrated in Fig. 2(b), allowing for improved illumination of occluded regions. The transmission component of X-band waves through typical outdoor relay walls can be safely ignored in our setting, as demonstrated in the Supp. Sec. 2.1.

However, the X-band, while reducing signal loss, degenerates spatial resolution due to its long wavelength, which in turn degrades the accuracy of scene reconstruction. Existing radar-based NLOS approaches [9, 33, 61, 73] typically rely on backpropagation methods of multi-bounce signals or assume known relay geometries and controlled tracking conditions, limiting their applicability in complex scenes. To overcome these limitations, we design a neural architecture that integrates a dense prediction module with a geometry-aware NLOS reconstruction module, achieving improved spatial resolution, reducing angular distortions, and disentangling direct LOS echoes from NLOS returns.

To this end, we present an X-band radar NLOS imaging approach that leverages the long wavelength to operate in a more specular scattering regime, enabling large-scale hidden-scene perception. We first derive an X-band NLOS image-formation model that captures propagation, reflection, and multipath effects specific to this band. Then, we introduce a neural reconstruction method to improve reconstruction fidelity and counteract the inherently low spatial resolution of long-wavelength radar. Building on the image formation model, we develop a physics-informed simulation method that integrates antenna design, material-dependent reflectance, and multi-bounce propagation for data generation and algorithm evaluation. We validate the method in simulation and with a large-aperture X-band pro-

tototype on real-world scenes, confirming accurate hidden-scene reconstructions. In summary, the contributions of this work are the following:

- We develop an X-band NLOS image formation and radar data synthesis method that models antenna radiation, multipath propagation, and reflection characteristics.
- We design a neural reconstruction method that combines learned dense prediction with geometry-aware refinement, substantially improving reconstruction quality from low-resolution radar data.
- We implement a large-aperture X-band radar NLOS imaging system, and validate that the method achieves accurate hidden-scene reconstruction up to 40 m in real-world environments.

2. Related Work

Radar Perception has become indispensable for object recognition and scene understanding in robotics [66], indoor settings [32, 43], and outdoor environments [3, 16, 24, 56]. Beyond robustness, radar supports rich environmental modeling, including depth [37], semantics [48, 62, 85], scene flow [13], object presence [3, 24, 27, 35]. In automotive applications, radar is a cornerstone for safety and efficiency [7, 12], with integration rates projected to grow rapidly [6]. Similar benefits extend to mobile robotics, both indoors and outdoors [16, 32, 43, 56, 66]. Thereby, approaches can be grouped by input representation into point-cloud approaches [62] and range–Doppler–azimuth (RAD) waveform methods [11, 50, 52, 77]. Point-cloud pipelines adapt 3D detection frameworks such as PointNets [54] or PointPillars [71], often with temporal aggregation [51], feature stacking [67], or graph modules [69] to capture spatiotemporal structure. However, radar point clouds are sparse and noisy, which can degrade feature completeness and reconstruction quality. Uncompressed waveform methods [11, 50, 84], in contrast, operate on dense spectral cubes across range, azimuth, and Doppler using 2D/3D CNNs [84]. Many employ multi-view fusion or attention to enhance feature extraction and detection [48]. Recent work

Table 1. Method comparison across modalities for NLOS imaging.

Method	Mod.	λ	Range	Res.	Pen.	Obj.	Amb.
f-k [38]	optical	532nm	3-6m	0.5°	✗	(✓)	(✓)
P-F [40]		532nm	3-5m	0.5°	✗	(✓)	✗
SOCR [42]		532nm	3-5m	0.5°	✗	✗	✗
LCT [49]		675nm	3-6m	0.5°	✗	✗	(✓)
Pueyo [53]		675nm	2-3m	0.5°	✗	✗	✗
Xu [79]	mmW	3.9mm	9-11m	1°	✗	(✓)	✓
Li [33]		3.9mm	9-11m	3°	✗	(✓)	✓
Wei [73]		3.9mm	3-5m	3°	✗	(✓)	✓
Tang [68]	UWB	15.8cm	9-11m	10°	✓	✓	✓
Chen [9]		15.8cm	9-11m	14.5°	✓	✓	✓
Ours	XB	3cm	40m	4°	✓	✓	✓

Each criterion is fully ✓, partially (✓), or not ✗ met. Meth., λ , Res., Pen., Obj., and Amb. denote method, wavelength, angular resolution, penetration, object material independence, and robustness to ambient illumination, respectively.

also addresses azimuth resolution limits and reconstruction quality [36]. Progress has been driven by public datasets that provide scale and benchmarks [7, 44, 50, 55, 57]. Nonetheless, all of the approaches described above are built for direct line-of-sight imaging, while our work develops a radar perception model that allows us to reconstruct objects in occluded areas and provide an accompanying dataset.

Optical NLOS Imaging relies on time-resolved light transport to reconstruct hidden scenes from transient reflections on visible relay surfaces. Early works demonstrated geometry reconstruction using ultrafast streak cameras and pulsed illumination [17, 19, 28, 72], followed by SPAD-based transient imaging systems that improved robustness [38, 41, 47, 78]. Subsequent inverse methods formulated NLOS reconstruction as a geometric optics-based or wave-based problem, including backprojection [1, 19, 30], light transport deconvolution [49, 83], linear optimization [21, 26, 41, 81], phasor-field propagation [39, 40, 46, 60] and deep learning [10, 18, 25, 45]. As summarized in Tab. 1, benefiting from the short optical wavelength and directionality, optical NLOS methods can reach centimeter-level 3D accuracy [38, 39, 46, 49, 74]. However, severe multi-bounce attenuation, assumptions on relay-surface reflectance (often Lambertian), and sensitivity to ambient illumination restrict practical use outdoors [30, 40, 63, 82, 83]. Even though long-range demonstrations have been reported using large-aperture telescopes [76], the wall-to-target distance is still limited by several meters.

Radar NLOS Imaging leverages the longer wavelength and strong penetration of microwave signals to perceive hidden objects beyond obstacles under natural illumination [61], enabling long range NLOS applications. Existing research primarily focuses on two frequency regimes: ultra-wideband (UWB, 3–10 GHz) [9, 20, 34, 68] and millimeter-wave (mmWave, 60–80 GHz, typically 77 GHz)

[33, 61, 65, 73, 75, 79]. While UWB radars offer longer detection range, their limited angular resolution restricts reconstruction to coarse localization by backpropagating multi-bounce signals from known relay surfaces [9, 15, 80]. Combining 77 GHz radar with data-driven tracking supports NLOS detection [61] and 3D reconstruction [31] in real-world, dynamic settings. However, their reconstructions are easily degraded by multipath interference and diffuse reflection due to the relatively short wavelength. In this work, we employ 10 GHz radar and propose a learned neural reconstruction model to enhance both resolution and robustness, bridging the gap between physically grounded radar modeling and data-driven inference.

3. Image Formation

In this section, we introduce the NLOS radar forward model and derive X-band-specific transport approximations that describe wave propagation and surface interactions. Building on this foundation, we then introduce a coherent detection method and range-azimuth (RA) representation to extract spatial information from the measured waveforms.

3.1. Radar NLOS Image Formation Model

The NLOS Transport is illustrated in Fig. 3. We emit a radar wavefield g from transmitter \mathbf{l} , which reflects from the relay wall $\Pi \subset \mathbb{R}^3$ with normal \mathbf{n} at the point $\mathbf{w}_1 \in \Pi$ and reflectivity $\rho_\Pi(\mathbf{w}_1)$. The reflected wave propagates into the hidden region, interacts with an object O at the point \mathbf{o} with reflectivity $\rho_O(\mathbf{o})$, and is then captured by the receiver \mathbf{s} after a second reflection at the point $\mathbf{w}_2 \in \Pi$ with reflectivity $\rho_\Pi(\mathbf{w}_2)$. Integrating over all feasible paths $L(\mathbf{w}_1, \mathbf{o}, \mathbf{w}_2)$ yields the received signal ϕ for transmit and receive steering angles θ_t and θ_r at time t , where c denotes the wave propagation speed

$$\begin{aligned} \phi(\theta_t, \theta_r, t) = & \iiint_{\Pi, O} \underbrace{B_t(\theta_t) B_r(\theta_r)}_{\text{tx/rx beam patterns}} \\ & \times \underbrace{\rho_\Pi(\mathbf{w}_1) \rho_O(\mathbf{o}) \rho_\Pi(\mathbf{w}_2)}_{\text{surface reflectances}} \underbrace{A(\mathbf{w}_1, \mathbf{o}, \mathbf{w}_2)}_{\text{path attenuation}} \\ & \times \underbrace{g\left(t - \frac{L(\mathbf{w}_1, \mathbf{o}, \mathbf{w}_2)}{c}\right)}_{\text{path length constraint}} d\mathbf{w}_2 d\mathbf{o} d\mathbf{w}_1. \end{aligned} \quad (1)$$

Here, $B_t(\theta_t)$ and $B_r(\theta_r)$ denote the transmit and receive beam patterns evaluated at wall points \mathbf{w}_1 and \mathbf{w}_2 . The total path length is $L(\mathbf{w}_1, \mathbf{o}, \mathbf{w}_2) = r_{\mathbf{l}\mathbf{w}_1} + r_{\mathbf{w}_1\mathbf{o}} + r_{\mathbf{o}\mathbf{w}_2} + r_{\mathbf{w}_2\mathbf{s}}$ with $r_{\mathbf{x}\mathbf{y}} = \|\mathbf{y} - \mathbf{x}\|_2$ and $A(\mathbf{w}_1, \mathbf{o}, \mathbf{w}_2)$ denotes the corresponding path attenuation.

Specular-Dominant Reflections of X-band radar wavelengths allow us to simplify the reflectance model. In detail, the general model in Eq. (1) collapses to a single-path

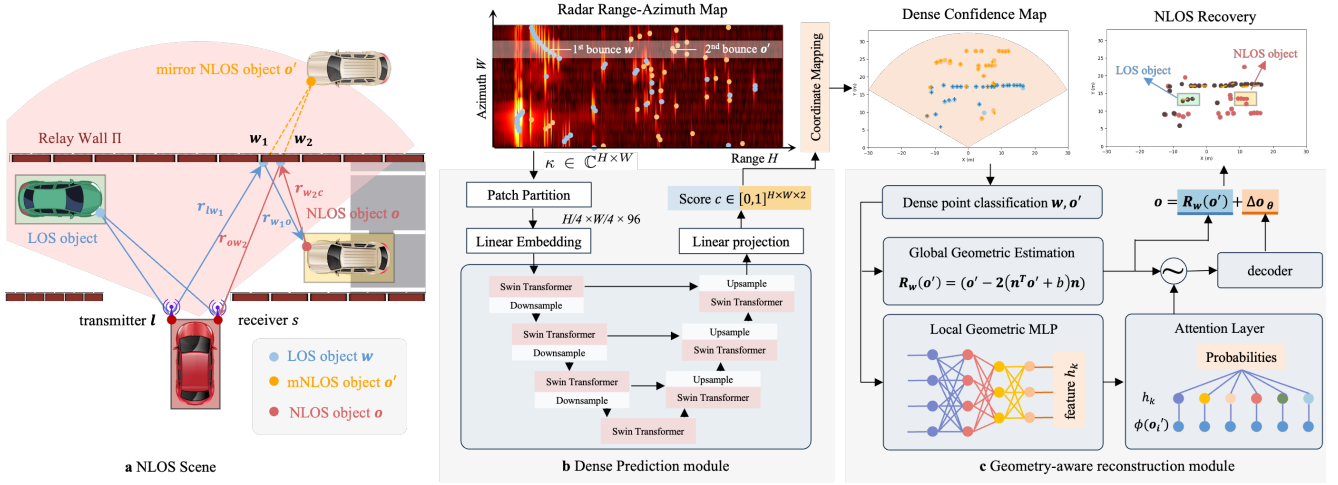


Figure 3. **X-band NLOS Imaging.** (a) An NLOS scenario is illustrated, where radar signals from hidden objects propagate via relay-wall reflections before reaching the receiver. The measured range–azimuth map is processed in (b) by a dense prediction module that predicts pixel-wise confidence and category maps for LOS and mirrored NLOS (mNLOS) responses. Finally, (c) uses a geometry-aware residual reconstruction network that leverages reflection geometry between visible and hidden regions to recover accurate hidden targets.

integral. A wall behaves as a quasi-mirror that redirects the transmitted field around the corner, so the measurements are equivalent to those produced by a virtual scene behind a transparent interface. Accordingly, the BRDF at the relay wall can be approximated by a specular kernel such that

$$\rho_{\Pi}(\mathbf{w}_1) \approx \alpha(\mathbf{w}_1) \delta(\mathbf{n}_{\mathbf{r}_{\mathbf{w}_1 \mathbf{o}}} - \mathbf{n}_{\mathbf{r}_{\mathbf{w}_1}}). \quad (2)$$

Through the Dirac-delta constraint, the integral in (1) collapses to the specular reflection points \mathbf{w}_1^* , \mathbf{w}_2^* , yielding

$$\begin{aligned} \phi(\boldsymbol{\theta}_t, \boldsymbol{\theta}_r, t) &\approx \int_{\mathcal{O}} \alpha(\mathbf{w}_1^*) \alpha(\mathbf{w}_2^*) B_t(\boldsymbol{\theta}_t) B_r(\boldsymbol{\theta}_r) \\ &\times A(\mathbf{w}_1^*, \mathbf{o}, \mathbf{w}_2^*) \rho_{\mathcal{O}}(\mathbf{o}) g\left(t - \frac{L(\mathbf{w}_1^*, \mathbf{o}, \mathbf{w}_2^*)}{c}\right) d\mathbf{o}. \end{aligned} \quad (3)$$

Further, the path loss A changes from the general form introduced above. Let r_t and r_r denote the propagation ranges from the transmitter to the object and from that point to the receiver, respectively. For specular or corner reflections, the wall/edge redirects the incident wavefront without re-radiating energy diffusely; the two bounces are equivalent to a single free-space hop of total length $r_t + r_r$. By spherical propagation, the received power therefore scales as the inverse square of this effective range, that is

$$A(\mathbf{w}_1, \mathbf{o}, \mathbf{w}_2) = \frac{1}{L(\mathbf{w}_1, \mathbf{o}, \mathbf{w}_2)^2}. \quad (4)$$

Scene Symmetry allows us to further rewrite the integral to simplify the search for \mathbf{w}_1^* , \mathbf{w}_2^* . Under specular reflection, the wall Π acts as a mirror and maps a scene object \mathcal{O} to its virtual counterpart \mathcal{O}' through the reflection operator R_{Π} . This yields the compact form

$$\begin{aligned} \phi(\boldsymbol{\theta}_t, \boldsymbol{\theta}_r, t) &\approx \int_{\mathcal{O}'} B_t(\boldsymbol{\theta}_t) B_r(\boldsymbol{\theta}_r) \\ &\cdot \tilde{\alpha}(\mathbf{o}') \tilde{A}(\mathbf{o}') \tilde{\rho}_{\mathcal{O}}(\mathbf{o}') g\left(t - \frac{L'(\mathbf{o}')}{c}\right) d\mathbf{o}', \end{aligned} \quad (5)$$

3.2. Detection and Range-Azimuth Representation

Our emitted signal g is a frequency modulated continuous wave (FMCW) modeled as $g(t) = \exp[j(2\pi f_c t + \pi(B/T)t^2)]$, where f_c , B and T denote the carrier frequency, sweep bandwidth and chirp duration, respectively. To extract the true object position, we multiply the received signal $\phi(\boldsymbol{\theta}_t, \boldsymbol{\theta}_r, t)$ by the conjugate of the transmitted waveform to perform a dechirping operation, and the subsequent low-pass filtering removes the high-frequency carrier term, yielding the baseband coherent signal

$$\tilde{m}(\boldsymbol{\theta}_t, \boldsymbol{\theta}_r, t) = \text{LPF}\{\phi(\boldsymbol{\theta}_t, \boldsymbol{\theta}_r, t) g^*(t)\}$$

$$\propto \int_{\mathcal{O}'} W(\mathbf{o}') \exp[j(2\pi f_b(\mathbf{o}') t + \varphi(\mathbf{o}'))] d\mathbf{o}', \quad (6)$$

where

$$W(\mathbf{o}') = B_t(\mathbf{o}'; \boldsymbol{\theta}_t) B_r(\mathbf{o}'; \boldsymbol{\theta}_r) \tilde{\alpha}(\mathbf{o}') \tilde{A}(\mathbf{o}') \tilde{\rho}_{\mathcal{O}}(\mathbf{o}'),$$

is the complex amplitude of the reflected path; $\text{LPF}\{\cdot\}$ denotes low-pass filtering; $\varphi(\mathbf{o}')$ is the residual phase term; $f_b(\mathbf{o}')$ is the beat frequency that encodes the target's round-trip delay $L'(\mathbf{o}')$. This operation converts the time delay into a measurable frequency shift, allowing the range information to be directly recovered from the spectrum of the baseband signal. By electronically steering the antenna array across transmit and receive angles, the system samples $\tilde{m}(\boldsymbol{\theta}_t, \boldsymbol{\theta}_r, t)$ over multiple directions, thereby forming the Range-Azimuth representation of the scene.

4. Neural Reconstruction

Due to the relatively long wavelength, X-band radar measurements suffer from low angular resolution, see Fig. 3.

This makes extracting precise RA information challenging. Our neural reconstruction method addresses this by first locating each target as a dense prediction. This enables our approach to effectively upsample the angular resolution while classifying each point as LOS or mNLOS reflection. Then, the NLOS targets are mapped to their correct Cartesian positions via a geometry-aware reconstruction module.

4.1. Feature Extraction and Geometry-aware Recovery

Our neural reconstruction approach processes the complex-valued RA measurements $\kappa \in \mathbb{C}^{H \times W}$ with a transformer-based encoder-decoder backbone that models long-range spatial dependencies. The backbone follows the Swin-UNet design [8] to extract both local geometric details and global contextual information necessary for precise peak localization. The final decoded feature map is projected through a linear layer to a two-channel tensor, followed by element-wise Sigmoid activation to produce a dense confidence map $\mathbf{c} \in [0, 1]^{H \times W \times 2}$. Each channel encodes the pixel-wise probability of a location corresponding to a LOS point \mathbf{w} or mNLOS reflection \mathbf{o}' , respectively. This generates precise range-angle estimates for detected peaks while also effectively classifying them.

The geometry-aware reconstruction head takes the detected LOS and mNLOS points and explicitly encodes the physical reflection geometry between visible and hidden regions. Each mNLOS point is first anchored to its most plausible LOS cluster by geometric angle consistency, then a local relay wall is formed via clustering. The normal of the relay wall is estimated by line fitting. Since LOS samples are often sparsely distributed, analytic normal estimation generates noisy normal estimates. To address this, the geometry-aware module predicts a residual correction that refines the analytic mirror reflection. For each detected mNLOS point \mathbf{o}' , the hidden point \mathbf{o} is reconstructed by combining the geometric reflection with a learned residual

$$\mathbf{o} = R_{\Pi}(\mathbf{o}') + \Delta\mathbf{o}_{\theta} = (\mathbf{o}' - 2(\mathbf{n}^{\top}\mathbf{o}' + b)\mathbf{n}) + \Delta\mathbf{o}_{\theta}, \quad (7)$$

where \mathbf{n} is the normal of the relay wall, b is the line offset, and $\Delta\mathbf{o}_{\theta}$ is predicted using attention-based aggregation of neighboring points on the relay wall. This allows the geometry-aware reconstruction head to correct systematic deviations caused by wall curvature and other non-idealities, achieving accurate and interpretable reconstruction of hidden objects even under sparse radar measurements. Additional details about the neural reconstruction approach in general are provided in the Supp. Sec. 3.1.

4.2. Training

We pretrain the dense model using Gaussian heatmaps and a focal loss [86]. For each reflection class $k \in \{1 = \text{LOS}, 2 = \text{mNLOS}\}$, we generate target heatmaps \mathbf{Y} \in

$[0, 1]^{H \times W \times 2}$ with Gaussian kernels centered at ground-truth peak locations. The focal loss emphasizes hard examples while suppressing background measurements

$$\mathcal{L}_{\text{heat}} = -\frac{1}{N} \sum_{x,y,k} \begin{cases} (1 - c_{xyk})^{\alpha} \log(c_{xyk}), & Y_{xyk} = 1, \\ (1 - Y_{xyk})^{\beta} (c_{xyk})^{\alpha} \\ \quad \times \log(1 - c_{xyk}), & \text{otherwise,} \end{cases} \quad (8)$$

where $\alpha = 2$ and $\beta = 4$ are focusing parameters, N is the total number of ground truth peaks, and c_{xyk} and Y_{xyk} denote the predicted confidence and target heatmap value at location (x, y) for class k , respectively.

The geometry-aware module is supervised using an ℓ_1 loss on the predicted residual corrections

$$\mathcal{L}_{\text{geo}} = \frac{1}{K} \sum_{j=1}^K \|\Delta\mathbf{o}_{\theta}^{(j)} - (\mathbf{o}^{(j)} - \hat{\mathbf{o}}_{\text{geom}}^{(j)})\|_1, \quad (9)$$

where $\mathbf{o}^{(j)}$ is the ground truth hidden point, $\hat{\mathbf{o}}_{\text{geom}}^{(j)}$ is the analytic mirror reflection, and $\Delta\mathbf{o}_{\theta}^{(j)}$ is the predicted residual. This approach enforces geometric consistency while allowing the network to learn corrections for non-ideal reflections. See Supp. Secs. 3.2–3.4 for details.

5. Assessment

We train our model on simulated data and evaluate it on both synthetic scenes with ground truth (Sec. 5.1) and real-world measurements with a X-band radar (Sec. 5.2).

5.1. Synthetic Validation

5.1.1. Simulation

Our simulation method is an end-to-end radar simulator that builds directly upon the image formation model established in Sec. 3. Unlike prior signal-level simulators [14, 64], our simulation integrates antenna array design, electromagnetic field modeling, scene rendering, and ray-based propagation to generate realistic RA-maps, as illustrated in Fig. 4. We begin by physically modeling the radar antenna array, incorporating 3D radiation patterns, element coupling effects, and precise spatial geometry to faithfully replicate the sensor’s electromagnetic behavior. Wave interactions with diverse materials are then simulated using a Finite-Difference Time-Domain (FDTD) solver, which determines reflection coefficients and propagation losses. Realistic materials and scene geometry are extracted through a 3D computer graphics engine (Unreal Engine). A dedicated ray tracer models multi-path propagation by integrating antenna simulation and material parameters. Finally, the simulated RF signals undergo coherent detection to produce RA data for training and validation. Specifically, we synthesize a dataset of 2,160 RA maps with ground truth spanning diverse urban,

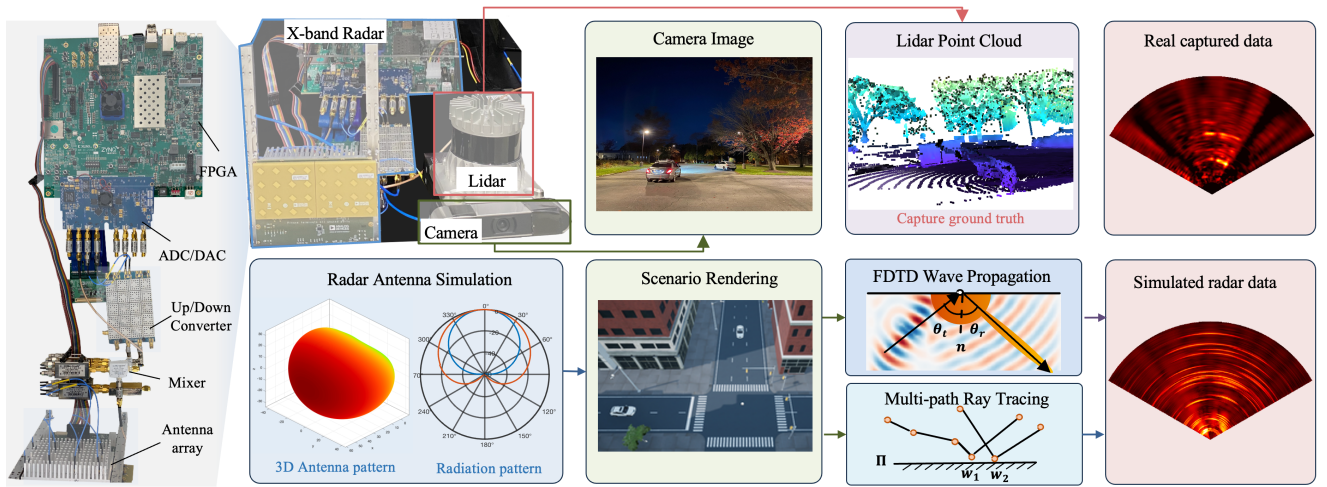


Figure 4. **Radar Prototype and Simulation Method.** The prototype system integrates an X-band radar (zoom-in view on the left), a camera, and a LiDAR sensor that provides ground-truth supervision. To accurately model the X-band radar sensing process, we simulate the complete RF signal chain using antenna radiation pattern synthesis, electromagnetic wave propagation, and multi-path ray tracing.

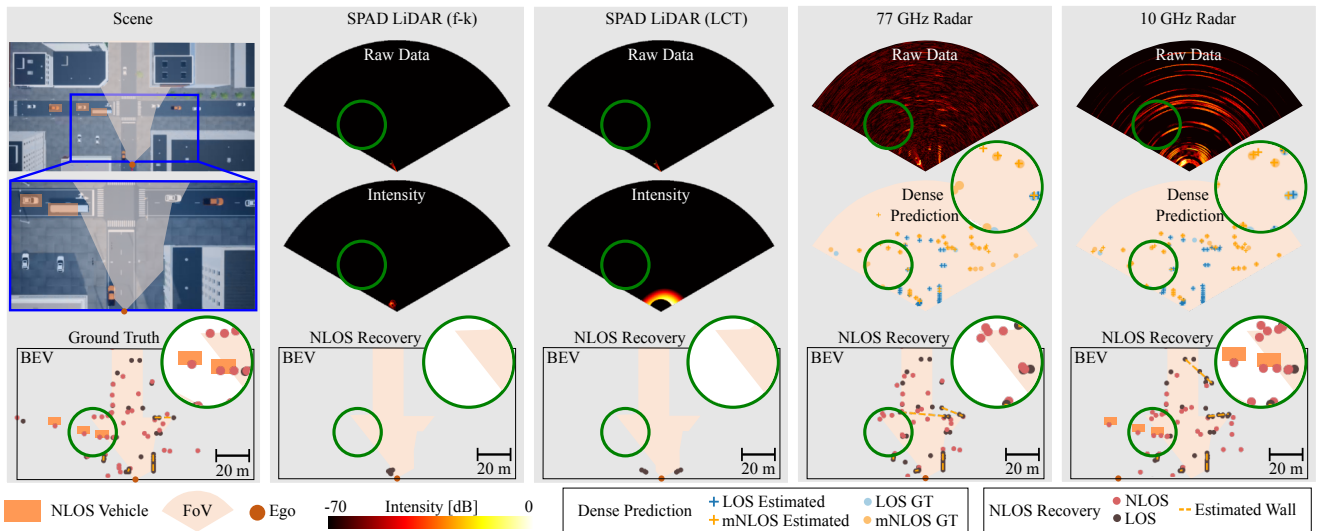


Figure 5. **Cross-Modality NLOS Comparison.** The first column shows the scene and a zoomed-in view, where the FOV and NLOS objects are marked in light orange and light yellow, respectively. The remaining columns present the raw measurements (top), intensity or dense prediction (middle), and NLOS reconstructions (bottom) of SPAD-LiDAR, 77 GHz and 10 GHz radar.

parking, and residential environments. Additional details on the simulator are provided in the Supp. Sec. 2.1.

5.1.2. Cross-Modality Validation

To assess the benefit of long wavelengths for NLOS recovery in simulation, we compare the imaging performance of three sensing modalities: a SPAD-based LiDAR [38, 41, 47, 78] operating at 850 nm wavelength, a 77 GHz radar ($\lambda = 3.9$ mm) [61], and an X-band radar ($\lambda = 3$ cm). For a fair comparison, all modalities are simulated with a total transmit power of 1.6 mW. Sensor details are provided in the Supp. Sec. 2.2. The distinct modalities require different reconstruction algorithms: for SPAD data, we use f - k migration [38] and LCT [49], whereas for radar inputs we use the proposed method.

Fig. 5 presents the reconstruction results for the three modalities, highlighting the critical role of wavelength. The SPAD LiDAR, operating at the shortest wavelength, can only reconstruct nearby objects, failing to capture distant NLOS regions. The 77 GHz radar, with a longer wavelength, shows improved penetration for longer-range LOS detection but remains limited in NLOS regions due to diffuse reflections from most surfaces. In contrast, the proposed X-band radar makes identical surfaces act more specular, enabling roughly a 10 \times longer NLOS reconstruction range than optical systems under the same transmit power.

5.1.3. Validation of Recovery Method

To validate the effectiveness of our neural reconstruction approach, we compare it against three radar reconstruction

Method	Dense Prediction						Reconstruction	
	F1 [-] \uparrow			CD [m^2] \downarrow			F1 [-] \uparrow	CD [m^2] \downarrow
	Macro	LOS	mNLOS	Macro	LOS	mNLOS		
<i>Simulation</i>								
NLOS-CFAR [61]	0.23	-	-	68.6	-	-	0.03	29.7
RTN[50]	0.57	0.68	0.47	104.0	140.6	67.2	0.20	14.0
Further Than CFAR[58]	0.41	0.52	0.31	65.9	47.6	84.2	0.16	13.5
Proposed + UNet[59]	0.67	0.73	0.61	44.0	43.0	44.9	0.27	11.9
Proposed w/o residual	0.85	0.88	0.82	17.3	11.6	23.0	0.40	9.2
Proposed	0.85	0.88	0.82	17.3	11.6	23.0	0.45	8.2
<i>Real</i>								
NLOS-CFAR [61]	0.21	-	-	42.8	-	-	0.01	36.4
RTN[50]	0.17	0.23	0.12	256.0	125.9	386.1	0.06	27.6
Further Than CFAR[58]	0.30	0.34	0.27	142.7	34.2	251.2	0.18	21.7
Proposed	0.47	0.42	0.52	37.4	20.2	54.6	0.20	24.3

Table 2. **Quantitative Reconstruction Results.** We evaluate dense prediction (separate LOS and mNLOS classes) and NLOS reconstruction on simulated and real data. ‘-’ indicates that CFAR cannot classify peaks. Our reconstruction method compares favorably to all baselines, and two ablations isolate the contributions of the Swin-UNet backbone and the residual formulation.

methods. The Constant False Alarm Rate (CFAR) detector, extended by Scheiner et al. [61] for NLOS recovery by assuming a known relay wall position, is employed as a recent radar NLOS baseline. Furthermore, it is benchmarked against two learned baselines: ‘‘Further Than CFAR’’ [58], which predicts heatmaps similar to our approach, and RTN [50], which we adapt from its original object detection purpose to perform LOS/NLOS recovery. Further details are available in the Supp. Sec. 4.1.

We evaluate performance separately for the dense prediction and geometric reconstruction tasks, using Chamfer Distance (CD) and F1 score. Dense prediction is assessed on the initial, un-mapped point clouds containing LOS or mNLOS points, evaluating the methods’ ability to localize peaks and classify them correctly. The final reconstruction quality is then measured on the complete point cloud after reflecting mNLOS points to their real positions using the geometric reconstruction described in Sec. 4.1.

As reported in Tab. 2, our method outperforms all baselines on simulated data. For dense prediction, it exceeds the next-best baseline (RTN) by 32.9% in Macro-F1 (mean of LOS and mNLOS-F1) score. The qualitative results in Fig. 6 show that the baselines struggle to accurately detect low-SNR peaks in the radar data (see zoom-ins). Our approach improves in the following two ways: first, the employed Swin-UNet’s transformer backbone captures scene-wide features essential for LOS and mNLOS distinction, unlike the baselines’ convolutional architecture. This is confirmed by the ablation in Tab. 2, where the performance reduces by 21.2% Macro-F1 score, when replacing the Swin-UNet with a regular UNet. Second, our Gaussian heatmap supervision (Eq. (8)) effectively captures the broad peaks characteristic of X-band radar. In contrast, Further

Than CFAR’s per-cell binary supervision struggles with peaks spanning multiple cells, and RTN’s object-detection fails to reconstruct areas with closely-spaced peaks.

The geometry-aware reconstruction is validated in Tab. 2, where it improves the F1 score by 55.6% and reduces the CD by 65.3% over the next best baseline. As shown in Fig. 6 (last column), this enables the correct reconstruction of NLOS points from vehicles in crossroad scenarios that are outside the FoV. This improved performance stems from our residual formulation of the geometry reconstruction (Eq. (7)). The incorporated geometrical prior facilitates learning as confirmed by the ablation experiments in Tab. 2.

5.1.4. Validation of Challenging Scenarios

We further evaluate challenging propagation effects at X-band, including beam splitting at corners and curved relay walls. Beam splitting introduces multiple propagation paths and mixed returns, while curvature causes deviations from ideal specular reflection. Our geometry-aware reconstruction jointly leverages local geometry and global context to handle these effects, enabling accurate reconstruction. Additional details and reconstruction results are provided in the Supp. Sec. 4.2.

5.2. Experimental Validation

Prototype We evaluate the proposed method experimentally with a custom X-band radar platform, as shown in Fig. 4. Our prototype is an FPGA-based RF System that integrates waveform generation, signal transmission, reception, and digital signal processing. The FMCW-modulated radar signal is sampled and up-converted before being sent to the antenna control board for beam steering. The transceiver module controls beam steering across the 4×8 X-band antenna array, while the received echo signals are down-converted and sampled by the FPGA. We collected 122 outdoor radar RA frames across 15 different scenes under varying environmental conditions. The radar wavefront is captured simultaneously with LiDAR point clouds and ego-view images. More prototype and data acquisition details are provided in the Supp. Sec. 1.

Evaluation As shown in Fig. 7, we evaluate the proposed method in multiple common NLOS scenarios during both daytime and nighttime. These experiments include different types of NLOS objects, such as cars and bicycles, as well as various relay-wall surfaces, including building facades, walls, and containers. The proposed method accurately reconstructs objects located in hidden regions; as shown in the first row, we detect NLOS cars at ranges of 40 m, corresponding to an 80 m round-trip path. Note that during nighttime, camera perception is challenging even in the line of sight. Quantitative experimental comparisons of NLOS recovery compared to baseline methods are reported in Table 2, where our method achieves the best F1 score, confirming its favorable accuracy and robustness. The Further-

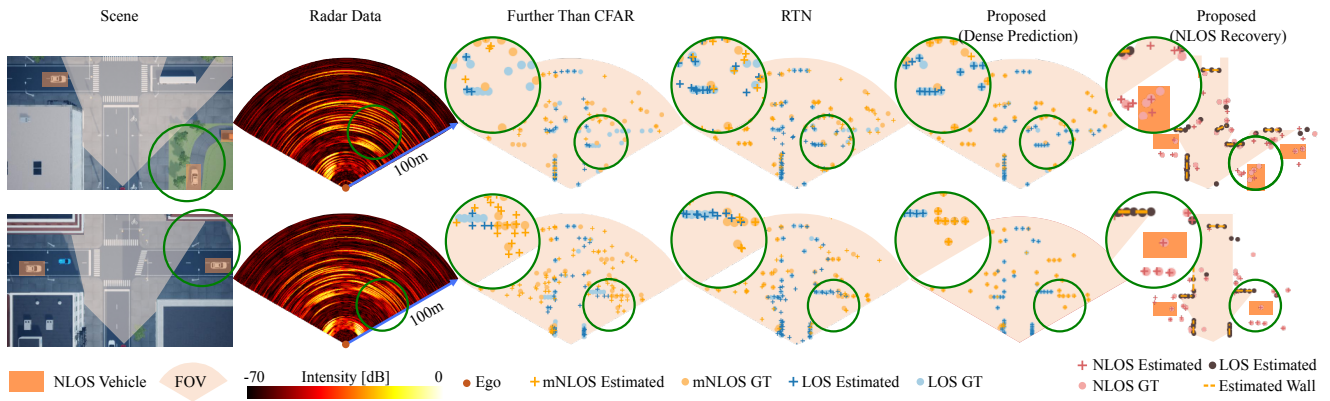


Figure 6. **Reconstruction Results on Synthetic Data.** The proposed method excels on the dense prediction task compared to the baseline methods, accurately distinguishing between mNLOS and LOS points. This enables NLOS recovery (last column), which, e.g., enables detection of vehicles outside the FoV.

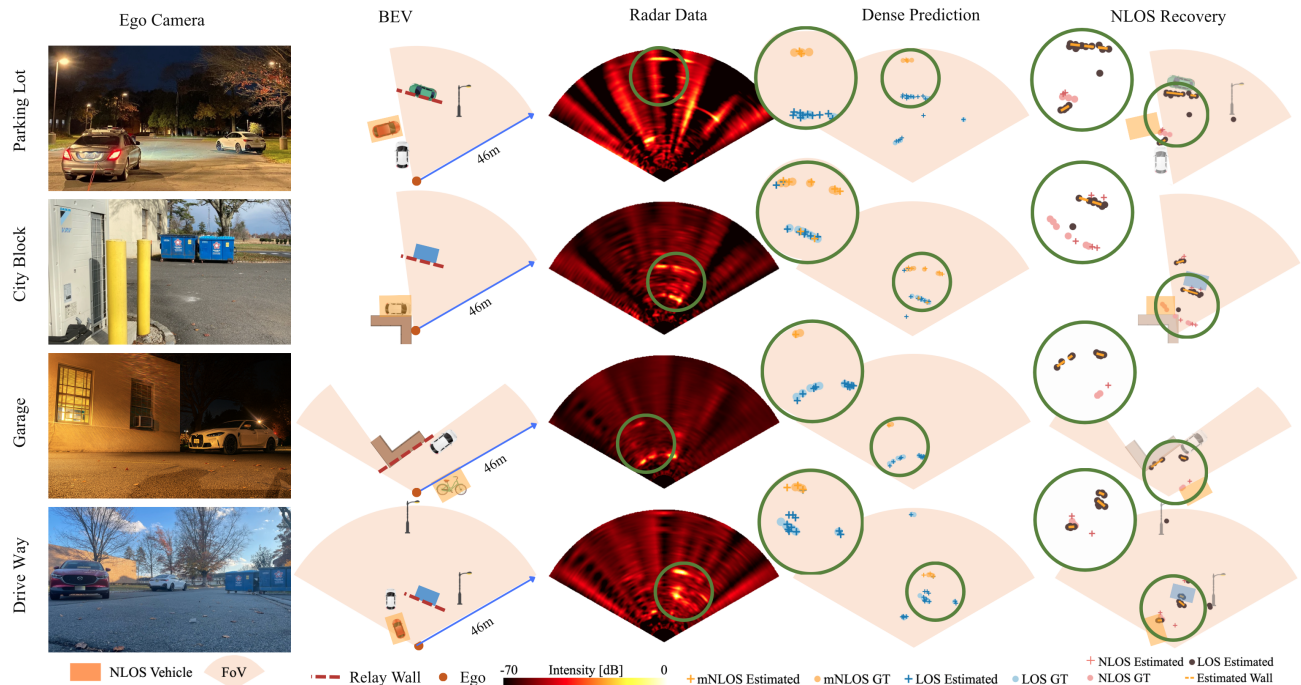


Figure 7. **Experimental Results.** The proposed method effectively reconstructs NLOS objects in common urban driving scenarios. From the raw radar measurements, our method accurately recovers hidden objects that are not visible in the ego-camera view. The recovered NLOS map agrees well with the bird's-eye-view ground-truth map.

Than-CFAR method yields a slightly lower CD as it predicts many false positive points near the wall (as shown by its inflated dense-prediction CD), which artificially simplifies the wall-fitting step of the NLOS recovery. Supporting visualizations are available in the Supp. Sec. 4.3.

6. Conclusion

We presented an X-band radar NLOS imaging method for large-scale hidden-scene perception by exploiting the stronger specular scattering of longer wavelengths. To address the limited angular resolution of X-band radar, we in-

troduced a neural reconstruction framework that combines dense prediction with geometry-aware reconstruction. We developed an X-band radar prototype to validate the proposed model across a variety of scenarios. Synthetic experiments demonstrate that X-band radar improves the NLOS reconstruction range by $10\times$ compared with optical systems. Tested on experimental data, our method achieves accurate hidden-scene reconstructions up to 40 m. Future work will incorporate Doppler cues, which are not yet used in the current prototype as discussed in Supp. Sec. 1, together with penetration effects to improve robustness in fog, foliage, and other challenging conditions.

Acknowledgements Felix Heide was supported by an NSF CAREER Award (2047359), a Packard Foundation Fellowship, a Sloan Research Fellowship, a Sony Young Faculty Award, a Project X Innovation Award, an Amazon Science Research Award, and a Bosch Research Award. Felix Heide is a co-founder of Algolux (now Torc Robotics), Head of AI at Torc Robotics, and a co-founder of Cephia AI. Dongyu Du was supported during this work by Tsinghua University and Tsinghua Innovation Center in Zhuhai (the NSF of Guangdong Province with Grant No.2023A1515012716). Mingkun Zhao would like to thank Zixi Li and Atsutse Kludze for their insightful discussions. Yutong Yang is also with the Institute of Signal Processing and System Theory (ISS), University of Stuttgart. The authors also would like to acknowledge Army Research Office (W911NF2110314, KS), Air Force Office of Scientific Research (FA9550-23-1-0176), and National Science Foundation (CNS-2211617, CNS-2148271) for funding.

References

- [1] Victor Arellano, Diego Gutierrez, and Adrian Jarabo. Fast back-projection for non-line of sight reconstruction. In *ACM SIGGRAPH 2017 Posters*, pages 1–2. 2017. 3
- [2] Petr Beckmann and Andre Spizzichino. The scattering of electromagnetic waves from rough surfaces. *Norwood*, 1987. 2
- [3] Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, and Felix Heide. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [4] Paulo V. K. Borges, Ash Tews, and Dave Haddon. Pedestrian detection in industrial environments: Seeing around corners. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4231–4232, 2012. 1
- [5] Katherine L. Bouman, Vickie Ye, Adam B. Yedidia, Frédo Durand, Gregory W. Wornell, Antonio Torralba, and William T. Freeman. Turning corners into cameras: Principles and methods. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2289–2297, 2017. 1
- [6] Ondrej Burkacky, Johannes Deichmann, Michael Guggenheimer, and Martin Keller. The automotive software and electronics market is poised for strong growth in the next decade. our latest market projections offer a glimpse into the industry’s future., 2022. 2
- [7] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 3
- [8] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang. Swin-UNet: UNet-like pure transformer for medical image segmentation. In *European Conference on Computer Vision (ECCV)*, pages 205–218, 2022. 5
- [9] Jiahui Chen, Shisheng Guo, Haolan Luo, Nian Li, and Guolong Cui. Non-line-of-sight multi-target localization algorithm for driver-assistance radar system. *IEEE Transactions on Vehicular Technology*, 72(4):5332–5337, 2023. 1, 2, 3
- [10] Wenzheng Chen, Simon Daneau, Fahim Mannan, and Felix Heide. Steady-state non-line-of-sight imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6790–6799, 2019. 3
- [11] Yuwei Cheng, Jingran Su, Mengxin Jiang, and Yimin Liu. A novel radar point cloud generation method for robot environment perception. *IEEE Transactions on Robotics*, 38(6): 3754–3773, 2022. 2
- [12] Juergen Dickmann, Nils Appenrodt, Hans-Ludwig Bloecher, C. Brenk, Thomas Hackbarth, Markus Hahn, Jens Klappstein, Marc Muntzinger, and Alfons Sailer. Radar contribution to highly automated driving. In *2014 44th European Microwave Conference*, pages 1715–1718, 2014. 2
- [13] Fangqiang Ding, Andras Palffy, Dariu M. Gavrilă, and Chris Xiaoxuan Lu. Hidden gems: 4d radar scene flow learning using cross-modal supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9340–9349, 2023. 2
- [14] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017. 5
- [15] Huagui Du, Chongyi Fan, Zhen Chen, Chun Cao, and Xiaotao Huang. Nlos target localization with an l-band uwb radar via grid matching. *Progress In Electromagnetics Research M*, 97:45–56, 2020. 1, 3
- [16] Ghina El Natour, Omar Ait Aider, Raphael Rouveure, François Berry, and Patrice Faure. Radar and vision sensors calibration for outdoor 3d reconstruction. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2084–2089, 2015. 2
- [17] Genevieve Garipey, Francesco Tonolini, Robert Henderson, Jonathan Leach, and Daniele Faccio. Detection and tracking of moving objects hidden from view. *Nature Photonics*, 10(1):23–26, 2016. 3
- [18] Javier Grau Chopite, Matthias B Hullin, Michael Wand, and Julian Iseringhausen. Deep non-line-of-sight reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 960–969, 2020. 3
- [19] Otkrist Gupta, Thomas Willwacher, Andreas Velten, Ashok Veeraraghavan, and Ramesh Raskar. Reconstruction of hidden 3d shapes using diffuse reflections. *Optics express*, 20(17):19096–19108, 2012. 3
- [20] Magnus Gustafsson, Åsa Andersson, Tommy Johansson, Stefan Nilsson, Ain Sume, and Anders Örbom. Extraction of human micro-doppler signature in an urban environment using a “sensing-behind-the-corner” radar. *IEEE Geoscience and Remote Sensing Letters*, 13(2):187–191, 2015. 3
- [21] Felix Heide, Matthias B. Hullin, James Gregson, and Wolfgang Heidrich. Diffuse mirrors: 3d reconstruction from diffuse indirect illumination using inexpensive time-of-flight

- sensors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3222–3229, 2014. 3
- [22] Felix Heide, Matthew O’Toole, Kai Zang, David B. Lindell, Steven Diamond, and Gordon Wetzstein. Non-line-of-sight imaging with partial occluders and surface normals. *ACM Trans. Graph.*, 2019. 1
- [23] Xin Huang, Ruilin Ye, Wenwen Li, Jian-Wei Zeng, Yi-Chen Lu, Huiqin Hu, Yijun Zhou, Lei Hou, Zheng-Ping Li, Hai-Feng Jiang, et al. Non-line-of-sight imaging and vibrometry using a comb-calibrated coherent sensor. *Physical Review Letters*, 132(23):233802, 2024. 1
- [24] Jyh-Jing Hwang, Henrik Kretzschmar, Joshua Manela, Sean Rafferty, Nicholas Armstrong-Crews, Tiffany Chen, and Dragomir Anguelov. Cramnet: Camera-radar fusion with ray-constrained cross-attention for robust 3d object detection, 2022. 2
- [25] Julian Iseringhausen and Matthias B Hullin. Non-line-of-sight reconstruction using efficient transient rendering. *ACM Transactions on Graphics (ToG)*, 39(1):1–14, 2020. 3
- [26] Achuta Kadambi, Hang Zhao, Boxin Shi, and Ramesh Raskar. Occluded imaging with time-of-flight sensors. *ACM Transactions on Graphics (ToG)*, 35(2):1–12, 2016. 3
- [27] Youngseok Kim, Juyeb Shin, Sanmin Kim, In-Jae Lee, Jun Won Choi, and Dongsuk Kum. Crn: Camera radar net for accurate, robust, efficient 3d perception. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 17615–17626, 2023. 2
- [28] Ahmed Kirmani, Tyler Hutchison, James Davis, and Ramesh Raskar. Looking around the corner using transient imaging. In *2009 IEEE 12th International Conference on Computer Vision*, pages 159–166. IEEE, 2009. 3
- [29] Eugene F Knott. *Radar cross section measurements*. Springer Science & Business Media, 2012. 2
- [30] Marco La Manna, Fiona Kine, Eric Breitbart, Jonathan Jackson, Talha Sultan, and Andreas Velten. Error back-projection algorithms for non-line-of-sight imaging. *IEEE transactions on pattern analysis and machine intelligence*, 41(7):1615–1626, 2018. 3
- [31] Haowen Lai, Zitong Lan, and Mingmin Zhao. Non-line-of-sight 3d reconstruction with radar. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 3
- [32] Seongwook Lee, Song-Yi Kwon, Bong-Jun Kim, Hae-Seung Lim, and Jae-Eun Lee. Dual-mode radar sensor for indoor environment mapping. *Sensors*, 21(7):2469, 2021. 2
- [33] Gen Li, Yun Ge, Yiyu Wang, Qingwu Chen, and Gang Wang. Detection of human breathing in non-line-of-sight region by using mmwave fmcw radar. *IEEE Transactions on Instrumentation and Measurement*, 71:1–11, 2022. 1, 2, 3
- [34] Songlin Li, Shisheng Guo, Jiahui Chen, Xiaqing Yang, Shihao Fan, Chao Jia, Guolong Cui, and Haining Yang. Multiple targets localization behind l-shaped corner via uwb radar. *IEEE Transactions on Vehicular Technology*, 70(4):3087–3100, 2021. 3
- [35] Yu-Jhe Li, Jinhyung Park, Matthew O’Toole, and Kris Kitani. Modality-agnostic learning for radar-lidar fusion in vehicle detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 918–927, 2022. 2
- [36] Yu-Jhe Li, Shawn Hunt, Jinhyung Park, Matthew O’Toole, and Kris Kitani. Azimuth super-resolution for fmcw radar in autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17504–17513, 2023. 3
- [37] Juan-Ting Lin, Dengxin Dai, and Luc Van Gool. Depth estimation from monocular images and sparse radar data. In *International Conference on Intelligent Robots and Systems (IROS)*, 2020. 2
- [38] David B. Lindell, Matthew O’Toole, and Gordon Wetzstein. Wave-based non-line-of-sight imaging using fast f-k migration. *ACM Transactions on Graphics*, 38(4):116:1–116:13, 2019. 1, 3, 6
- [39] Xiaochun Liu, Ibón Guillén, Marco La Manna, Ji Hyun Nam, Syed Azer Reza, Toan Huu Le, Adrian Jarabo, Diego Gutierrez, and Andreas Velten. Non-line-of-sight imaging using phasor-field virtual wave optics. *Nature*, 572(7771):620–623, 2019. 3
- [40] Xiaochun Liu, Sebastian Bauer, and Andreas Velten. Phasor field diffraction based reconstruction for fast non-line-of-sight imaging systems. *Nature communications*, 11(1):1645, 2020. 3
- [41] Xintong Liu, Jianyu Wang, Zhupeng Li, Zuoqiang Shi, Xing Fu, and Lingyun Qiu. Non-line-of-sight reconstruction with signal-object collaborative regularization. *Light: Science & Applications*, 10(1):198, 2021. 3, 6
- [42] Xintong Liu, Jianyu Wang, Leping Xiao, Zuoqiang Shi, Xing Fu, and Lingyun Qiu. Non-line-of-sight imaging with arbitrary illumination and detection pattern. *Nature Communications*, 14(1):3230, 2023. 3
- [43] Chris Xiaoxuan Lu, Stefano Rosa, Peijun Zhao, Bing Wang, Changhao Chen, John A. Stankovic, Niki Trigoni, and Andrew Markham. See through smoke: Robust indoor mapping with low-cost mmwave radar. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, pages 14–27, New York, NY, USA, 2020. Association for Computing Machinery. 2
- [44] Michael Meyer and Georg Kuschik. Automotive radar dataset for deep learning based 3d object detection. In *2019 16th European Radar Conference (EuRAD)*, pages 129–132, 2019. 3
- [45] Fangzhou Mu, Sicheng Mo, Jiayong Peng, Xiaochun Liu, Ji Hyun Nam, Siddeshwar Raghavan, Andreas Velten, and Yin Li. Physics to the rescue: Deep non-line-of-sight reconstruction for high-speed imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 3
- [46] Ji Hyun Nam, Eric Brandt, Sebastian Bauer, Xiaochun Liu, Marco Renna, Alberto Tosi, Eftychios Sifakis, and Andreas Velten. Low-latency time-of-flight non-line-of-sight imaging at 5 frames per second. *Nature Communications*, 12(1):6526, 2021. 3
- [47] Matthew O’Toole, David B. Lindell, and Gordon Wetzstein. Real-time non-line-of-sight imaging. In *ACM SIGGRAPH 2018 Emerging Technologies*, pages 1–2. 2018. 3, 6
- [48] Arthur Ouaknine, Alasdair Newson, Patrick Pérez, Florence Tupin, and Julien Rebut. Multi-view radar semantic segmen-

- tation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15671–15680, 2021. 2
- [49] Matthew O’Toole, David B Lindell, and Gordon Wetzstein. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature*, 555(7696):338–341, 2018. 1, 3, 6
- [50] Dong-Hee Paek, Seung-Hyun Kong, and Kevin Tirta Wijaya. K-radar: 4d radar object detection for autonomous driving in various weather conditions. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. 2, 3, 7
- [51] Andras Palffy, Ewoud Pool, Srimannarayana Baratam, Julian FP Kooij, and Dariu M Gavrilă. Multi-class road user detection with 3+1d radar in the view-of-delft dataset. *IEEE Robotics and Automation Letters*, 7(2):4961–4968, 2022. 2
- [52] Akarsh Prabhakara, Tao Jin, Arnav Das, Gantavya Bhatt, Lilly Kumari, Elahe Soltanaghahi, Jeff Bilmes, Swarun Kumar, and Anthony Rowe. High resolution point clouds from mmwave radar. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4135–4142, 2023. 2
- [53] Oscar Pueyo-Ciudad, Julio Marco, Stephane Schertzer, Frank Christnacher, Martin Laurenzis, Diego Gutierrez, and Albert Redo-Sanchez. Time-gated polarization for active non-line-of-sight imaging. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–11, 2024. 1, 3
- [54] Charles R. Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 2
- [55] Julien Rebut, Arthur Ouaknine, Waqas Malik, and Patrick Pérez. Raw high-definition radar for multi-task learning. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17000–17009, 2022. 3
- [56] Giulio Reina, James Underwood, Graham Brooker, and Hugh Durrant-Whyte. Radar-based perception for autonomous outdoor vehicles. *Journal of Field Robotics*, 28(6):894–913, 2011. 2
- [57] Ignacio Roldan, Andras Palffy, Julian F. P. Kooij, Dariu M. Gavrilă, Francesco Fioranelli, and Alexander Yarovoy. A deep automotive radar detector using the radelft dataset. *IEEE Transactions on Radar Systems*, 2:1062–1075, 2024. 3
- [58] Ignacio Roldan, Andras Palffy, Julian F. P. Kooij, Dariu M. Gavrilă, Francesco Fioranelli, and Alexander Yarovoy. See further than CFAR: A data-driven radar detector trained by lidar. In *Proceedings of the 2024 IEEE Radar Conference (RadarConf24)*, 2024. 7
- [59] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 7
- [60] Diego Royo, Talha Sultan, Adolfo Muñoz, Khadijeh Masumnia-Bisheh, Eric Brandt, Diego Gutierrez, Andreas Velten, and Julio Marco. Virtual mirrors: Non-line-of-sight imaging beyond the third bounce. *ACM Transactions on Graphics (TOG)*, 42(4):1–15, 2023. 3
- [61] Nicolas Scheiner, Florian Kraus, Fangyin Wei, Buu Phan, Fahim Mannan, Nils Appenrodt, Werner Ritter, Jurgen Dickmann, Klaus Dietmayer, Bernhard Sick, and Felix Heide. Seeing around street corners: Non-line-of-sight detection and tracking in the wild using doppler radar. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2065–2074. IEEE, 2020. 1, 2, 3, 6, 7
- [62] Ole Schumann, Markus Hahn, Jürgen Dickmann, and Christian Wöhler. Semantic segmentation on radar point clouds. In *2018 21st International Conference on Information Fusion (FUSION)*, pages 2179–2186, 2018. 2
- [63] Sheila W. Seidel, Hoover Rueda-Chacón, Iris Cusini, Federica Villa, Franco Zappa, Christopher Yu, and Vivek K. Goyal. Non-line-of-sight snapshots and background mapping with an active corner camera. *Nature Communications*, 14(1):3677, 2023. 3
- [64] Shital Shah, Debadepta Dey, Chris Lovett, and Ashish Kapoor. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and service robotics: Results of the 11th international conference*, pages 621–635. Springer, 2017. 5
- [65] Dmitrii Solomitskii, Mikko Heino, Sreehari Buddappagari, Matthias A Hein, and Mikko Valkama. Radar scheme with raised reflector for nlos vehicle detection. *IEEE Transactions on Intelligent Transportation Systems*, 23(7):9037–9045, 2021. 1, 3
- [66] Christian Stetco, Barnaba Ubezio, Stephan Mühlbacher-Karrer, and Hubert Zangl. Radar sensors in collaborative robotics: Fast simulation and experimental validation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10452–10458, 2020. 2
- [67] Bin Tan, Zhixiong Ma, Xichan Zhu, Sen Li, Lianqing Zheng, Sihan Chen, Libo Huang, and Jie Bai. 3-d object detection for multiframe 4-d automotive millimeter-wave radar point cloud. *IEEE Sensors Journal*, 23(11):11125–11138, 2022. 2
- [68] Quan Tang, Jun Li, Lingyu Wang, Yong Jia, and Guolong Cui. Multipath imaging for nlos targets behind an l-shaped corner with single-channel uwb radar. *IEEE Sensors Journal*, 22(2):1531–1540, 2022. 1, 3
- [69] Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J. Guibas. KPConv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6411–6420, 2019. 2
- [70] Fawwaz T Ulaby, Richard K Moore, and Adrian K Fung. Microwave remote sensing: Active and passive. volume 2-radar remote sensing and surface scattering and emission theory. 1982. 2
- [71] Michael Ulrich, Sascha Braun, Daniel Köhler, Daniel Niederlöhner, Florian Faion, Claudius Gläser, and Holger Blume. Improved orientation estimation and detection with hybrid object detection networks for automotive radar. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pages 111–117. IEEE, 2022. 2

- [72] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Mounqi G. Bawendi, and Ramesh Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications*, 3: 745, 2012. 3
- [73] Shunjun Wei, Jinshan Wei, Xinyuan Liu, Mou Wang, Shan Liu, Fan Fan, Xiaoling Zhang, Jun Shi, and Guolong Cui. Non-line-of-sight 3-d imaging using millimeter-wave radar. *IEEE Transactions on Geoscience and Remote Sensing*, 60: 1–18, 2021. 1, 2, 3
- [74] Florian Willomitzer, Prasanna V. Rangarajan, Fengqiang Li, Muralidhar M. Balaji, Marc P. Christensen, and Oliver Coscairt. Fast non-line-of-sight imaging with high-resolution and wide field of view using synthetic wavelength holography. *Nature Communications*, 12(1):6647, 2021. 3
- [75] Timothy Woodford, Kun Qian, and Xinyu Zhang. Metasight: High-resolution nlos radar with efficient metasurface encoding. In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems*, pages 308–321, 2023. 1, 3
- [76] Cheng Wu, Jianjiang Liu, Xin Huang, Zheng-Ping Li, Chao Yu, Jun-Tian Ye, Jun Zhang, Qiang Zhang, Xiankang Dou, Vivek K Goyal, et al. Non-line-of-sight imaging over 1.43 km. *Proceedings of the National Academy of Sciences*, 118(10):e2024468118, 2021. 3
- [77] Jincheng Wu, Ruixu Geng, Yadong Li, Dongheng Zhang, Zhi Lu, Yang Hu, and Yan Chen. Diffradar: High-quality mmwave radar perception with diffusion probabilistic model. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8291–8295, 2024. 2
- [78] Shumian Xin, Sotiris Nousias, Kiriakos N. Kutulakos, Aswin C. Sankaranarayanan, Srinivasa G. Narasimhan, and Ioannis Gkioulekas. A theory of fermat paths for non-line-of-sight shape reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6793–6802, 2019. 3, 6
- [79] You Xu, Guanghua Liu, Xiaotong Lu, Chao Xie, Lei Xue, and Tao Jiang. Double sparse structure-enhanced mmwave nlos imaging under multiangle relay surface. *IEEE Transactions on Signal Processing*, 2024. 3
- [80] Xiqing Yang, Shihao Fan, Shisheng Guo, Songlin Li, Guolong Cui, and Wei Zhang. Nlos target localization behind an l-shaped corner with an l-band uwb radar. *IEEE Access*, 8: 31270–31286, 2020. 3
- [81] Juntian Ye, Yu Hong, Xiongfei Su, Xin Yuan, and Feihu Xu. Plug-and-play algorithms for dynamic non-line-of-sight imaging. *ACM Transactions on Graphics*, 43(5):1–12, 2024. 3
- [82] Jun-Tian Ye, Yi Sun, Wenwen Li, Jian-Wei Zeng, Yu Hong, Zheng-Ping Li, Xin Huang, Xianghui Xue, Xin Yuan, Feihu Xu, et al. Real-time non-line-of-sight computational imaging using spectrum filtering and motion compensation. *Nature Computational Science*, 4(12):920–927, 2024. 1, 3
- [83] Sean I Young, David B Lindell, Bernd Girod, David Taubman, and Gordon Wetzstein. Non-line-of-sight surface reconstruction using the directional light-cone transform. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1407–1416, 2020. 3
- [84] Ao Zhang, Farzan Erlik Nowruzi, and Robert Laganière. RADDet: Range-azimuth-doppler based radar object detection for dynamic road users. In *2021 18th Conference on Robots and Vision (CRV)*, pages 95–102, 2021. 2
- [85] Liwen Zhang, Xinyan Zhang, Youcheng Zhang, Yufei Guo, Yuanpei Chen, Xuhui Huang, and Zhe Ma. Peakconv: Learning peak receptive field for radar semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17577–17586, 2023. 2
- [86] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 5