

# mmDiff: A Noise-Robust Differentiable Ray-Tracing Framework for mmWave Scene Calibration and Channel Prediction

## Supplementary Material

### A. Proof of Theorem 1




Figure 10. Geometry setting for proof of Theorem 1.




Figure 11. Illustration of the effective area  $dS$ .

We prove the following equation holds:

$$\lim_{\alpha_r \rightarrow \infty} \lim_{S \rightarrow \infty} P_{\alpha_r}(S) = \left( \frac{\lambda \Gamma}{4\pi r_{\text{spec}}} \right)^2, \quad (10)$$

where  $P_{\alpha_r}(S)$  denotes the power received at RX that is carried by single-bounce diffuse-scattering paths whose scattering point lies in area  $S$ . We show that when the reflector surface is sufficiently large, the received scattering power approaches the specular reflection power asymptotically as the scattering pattern becomes more directional.

We start with an ideal, vertically polarized, isotropic antenna model and will generalize to an arbitrary antenna pattern towards the end.

*Proof.* Since the received power will always be a fraction of the transmitted power, for simplicity, we assume the transmitted power  $P_t = 1$ .

For an infinitesimal patch of the surface,  $dS'$ , the solid angle subtended by the patch at the TX is

$$d\Omega_i = \frac{\cos \theta_i dS'}{r_i^2}, \quad (11)$$

where  $\theta_i$  is the incident angle to the scattering point.

Hence, the differential incident power is

$$dP_{\text{inc}} = \frac{1}{4\pi} d\Omega_i = \frac{\cos \theta_i}{4\pi r_i^2} dS' \quad (12)$$

Since the scattered field amplitude is  $\Gamma \sqrt{f_s(\alpha_r; \theta_s, \phi_s)}$  times of the incident field (Eq. (8)), the differential scattered power is

$$dP_{\text{sca}} = \Gamma^2 f_s(\alpha_r; \theta_s, \phi_s) dP_{\text{inc}} \quad (13)$$

The scattering field propagates to RX in free space, hence its power decay follows the Friis equation of free-space path loss.

$$dP_r = \left(\frac{\lambda}{4\pi r_s}\right)^2 dP_{\text{sca}}, \quad (14)$$

Thus, the differential received power is:

$$dP_r = \left(\frac{\lambda\Gamma}{4\pi}\right)^2 f_s(\alpha_r; \theta_s, \phi_s) \frac{\cos \theta_i}{r_i^2 r_s^2} dS' \quad (15)$$

To get the total received power, we need to integrate over the whole reflective surface:

$$P_{\alpha_r}(S) = \iint_S \left(\frac{\lambda\Gamma}{4\pi}\right)^2 f_s(\alpha_r; \theta_s, \phi_s) \frac{\cos \theta_i}{r_i^2 r_s^2} dS' \quad (16)$$

It is essential to note that only a fraction of the total shined area is captured by the receiver's aperture. As illustrated in Fig. 10, the total illuminated area by a ray is denoted as  $dS'$ . The total power that reaches  $dS'$  is scattered (re-distributed) to different directions, only a fraction of which is captured by the effective aperture of the receiver, denoted as  $A_{\text{eff}}$ . According to the geometry, the effective area,  $dS$ , that contains the power that reaches the receiver is related to  $dS'$  by the following equation:

$$\frac{dS}{dS'} = \frac{A_{\text{eff}}}{A} = \frac{r_i^2}{(r_i + r_s)^2}, \quad (17)$$

where  $A$  is an auxiliary variable denoting the subtended area of the ray at distance  $r_i + r_s$ .

Plug  $dS$  into the power integration, we get:

$$P_{\alpha_r}(S) = \iint_S \left(\frac{\lambda\Gamma}{4\pi}\right)^2 f_s(\alpha_r; \theta_s, \phi_s) \frac{\cos \theta_i}{r_i^2 r_s^2} \frac{r_i^2}{(r_i + r_s)^2} dS \quad (18)$$

To calculate the integral, we perform a change of variable to spherical coordinates  $(\theta_s, \phi_s)$  at the scattering point. We use the geometry model in Fig. 10 to find the Jacobian matrix, where TX and RX are located at points  $\mathbf{T} = (-0, 0, h_T)$ ,  $\mathbf{R} = (d, 0, h_R)$ , respectively;  $\mathbf{P} = (x, y, 0)$  represents an arbitrary scattering point on the reflector plane. Since  $dS = dx dy$ , we want to change the basis from  $(x, y)$  to  $(\theta_s, \phi_s)$ . According to the geometry,

$$\mathbf{k}_s = \begin{bmatrix} \sin \theta_s \cos \phi_s \\ \sin \theta_s \sin \phi_s \\ \cos \theta_s \end{bmatrix} = \frac{1}{r_s} \begin{bmatrix} d - x \\ -y \\ h_R \end{bmatrix}, \quad (19)$$

where  $r_s = \sqrt{(d-x)^2 + y^2 + h_R^2}$ . And we can solve the determinant of the Jacobian:

$$|J| = \det(A) = \begin{vmatrix} \frac{\partial x}{\partial \theta_s} & \frac{\partial x}{\partial \phi_s} \\ \frac{\partial y}{\partial \theta_s} & \frac{\partial y}{\partial \phi_s} \end{vmatrix} = r_s^2 \tan \theta_s \quad (20)$$

Apply the change-of-variable:

$$\begin{aligned} \lim_{S \rightarrow \infty} P_{\alpha_r}(S) &= \int_0^{2\pi} \int_0^{\pi/2} \left(\frac{\lambda\Gamma}{4\pi}\right)^2 f_s(\alpha_r; \theta_s, \phi_s) \frac{\cos \theta_i}{r_i^2 r_s^2} \frac{r_i^2}{(r_i + r_s)^2} r_s^2 \tan \theta_s d\theta_s d\phi_s \\ &= \left(\frac{\lambda\Gamma}{4\pi}\right)^2 \int_0^{2\pi} \int_0^{\pi/2} f_s(\alpha_r; \theta_s, \phi_s) \sin \theta_s \frac{1}{(r_i + r_s)^2} \frac{\cos \theta_i}{\cos \theta_s} d\theta_s d\phi_s \end{aligned} \quad (21)$$

The integral range includes the entire upper hemisphere to adhere to the limit of a sufficiently large surface. According to the energy conservation (EC) property of the scattering pattern:


$$\int_0^{2\pi} \int_0^{\pi/2} f_s(\alpha_r; \theta_s, \phi_s) \sin \theta_s d\theta_s d\phi_s = 1. \quad (\text{EC})$$

As  $\alpha_r \rightarrow \infty$ , the scattering pattern approaches a Dirac Delta function in 3D, and  $\cos \theta_i = \cos \theta_s$  at the specular reflection point. Thus, according to the sifting property of the Delta function, we have:


$$\lim_{\alpha_r \rightarrow \infty} \lim_{S \rightarrow \infty} P_{\alpha_r}(S) = \left(\frac{\lambda\Gamma}{4\pi}\right)^2 \frac{1}{(r_i^* + r_s^*)^2} = \left(\frac{\lambda\Gamma}{4\pi r_{spec}}\right)^2, \quad (22)$$

where  $r_i^*$  and  $r_s^*$  are the incident and scattering segment lengths when the scattering point matches the specular reflection point.  $\square$


The above proof assumes a vertically-polarized isotropic antenna at both TX and RX. Since we model the antenna radiation pattern and polarization as a separated multiplicative factor (Sec. 3.1), the above proof can be easily generalized to arbitrary antenna patterns and polarization.



(a) Deviation analysis for different directivities and scales of the reflective surface.



(b) Deviation analysis for distances between TX/RX and reflection point.



(c) Deviation analysis for different incident angles.

Figure 12. Factors that impact the deviation of the directional scattering approximation from the specular reflection model.






Figure 13. Single planar reflector scene for deviation analysis.



(a) Normalization factor increases with large incident angles.



(b) Scattering pattern for different incident angles ( $\alpha_r = 120$ ).

Figure 14. Large incident angles cause a higher normalization factor, which amplifies the scattering lobe.

## B. Deviation Analysis

In the previous section, we have proved that under extreme conditions (i.e., sufficiently large surface and narrow beam), the directional scattering approximation converges to the specular reflection model. In this section, we analyze the deviation introduced by the scattering approximation under non-extreme scenarios and elucidate the source of deviations. Note that we call this "deviation" instead of "error" because the specular reflection model itself is an approximation of the real-world reflection behavior, and is derived under the infinitely large, perfectly flat surface assumption. Here, we treat the specular reflection model as a reference to understand the behavior of mmDiff's approximation under different conditions.

We found that the deviation is mainly caused by four factors: the scale of the reflective surface, the directivity of the scattering lobe, the distance between the surface and the TX/RX, and the incident angle. Fig. 12a shows the deviation versus the directivity of the scattering lobe for different surface scales. We conduct this experiment in the setup shown in Fig. 13. We fixed the distance between the TX/RX and the reflection point to 25 m and the incident angle to  $45^\circ$ . By sweeping across different directivities, we observe that the deviation decreases as the directivity increases. This is expected since a higher directivity means a narrower lobe, which better approximates the specular reflection. We also observe that for a given directivity, a larger surface scale results in a smaller deviation. This is because a larger surface can capture more energy, leading to a higher path gain that is closer to the specular reflection. For the scale of 5, the deviation is 2 dB even when the directivity has reached 120, indicating that the surface is not large enough to fully capture the energy.

Fig. 12b shows the deviation versus the distance between the TX/RX and the reflection point. In this experiment, we fixed the surface scale to 5 and directivity to 120, and changed the Tx-to-surface distance,  $d_1$ , while keeping the Rx-to-surface distance,  $d_2$ , fixed. The incident angle is also fixed to  $45^\circ$ . As shown, the deviation increases as the distance increases. This is because, for a fixed surface scale, as the TX/RX moves further away, the effective area of the surface that is illuminated by the TX and visible to the RX decreases, leading to a lower path gain.

The above two experiments reveal that the key factor that affects the deviation is the effective area of the surface that is illuminated by the TX and visible to the RX. When the effective area is large enough to capture all the energy, the directional scattering approximation matches the specular reflection path gains very well. Hence, the directional scattering model can be regarded as a softening of the specular reflection model for finite, non-ideal surface scenarios.




Figure 15. Illustration of the sampling issue that causes the variation with respect to the incident angle (bird's-eye view).




Figure 16. Maximum deviation for different distances between TX/RX and reflection point.

Interestingly, we found that the incident angle also affects the deviation. In this experiment, we fixed the surface scale to 200 (sufficiently large), and the distance between the TX/RX and the reflection point to 25m, then swept the incident angle  $\theta_i$  from 0 to  $90^\circ$ . We move both TX and RX to keep the incident angle and reflection angle the same, such that we can use the specular reflection path gain as the reference. We choose three directivities: 10, 50, and 120, and plot the deviation versus the incident angle in Fig. 12c. We observe that for small incident angles (up to  $40^\circ$  for  $\alpha_r = 10$ , and  $70^\circ$  for  $\alpha_r = 50$  and 120), the deviation is relatively constant with respect to the incident angle. However, as the incident angle increases over  $40^\circ$ , there is an interesting variation pattern of the directional path gain: the deviation first decreases, then increases again after  $70^\circ$ , and finally decreases again after  $80^\circ$ .

Next, we perform a qualitative analysis to elucidate the reason behind this variation. First, we found that this variation is related to the normalization factor of the scattering lobe. As we discussed in Sec. A, the scattering pattern must satisfy the energy conservation equation (Eq. (EC)). In our case, we adopt the scattering pattern model in [7], which has a normalization factor,  $F_{\alpha_r}(\theta_i)^{-1}$ , (Eq. (7)) to ensure the energy conservation. We plot the normalization factor versus the incident angle in Fig. 13a. We observe that the normalization factors share the same pivot point at around  $40^\circ$  for  $\alpha_r = 10$  and around  $70^\circ$  for  $\alpha_r = 50$  and 120. After the pivot point, the normalization factor increases rapidly as the incident angle increases. This means that for large incident angles, the scattering lobe is amplified more (to compensate for the cutting by the scattering plane) to satisfy the energy conservation equation. Fig. 13b shows the scattering pattern for different incident angles when  $\alpha_r = 120$ . We can see that as the incident angle increases to  $85^\circ$ , the scattering lobe is cut more by the scattering plane, and the remaining lobe is amplified more to satisfy the energy conservation equation.

Now that we know that the variation is related to the amplification of the scattering lobe for large incident angles, we further investigate why this amplification causes the variation in deviation. According to Eq. (18), the path gain of a single scattering path is also related to the scattering path length ( $r_i + r_s$ ). If we plot the level set of the scattering path length on the scattering surface, the contours will be ellipses, as illustrated in Fig. 15. We also plot the curve of  $\theta_i = \theta_i^*$  on the figure, where  $\theta_i^*$  is the incident angle that results in the specular reflection (our reference). The resulting curve is a circle. Within the circle,  $\theta_i < \theta_i^*$ , and  $F_{\alpha_r}(\theta_i) \leq F_{\alpha_r}(\theta_i^*)$ . On the contrary,  $\theta_i > \theta_i^*$ , and  $F_{\alpha_r}(\theta_i) \geq F_{\alpha_r}(\theta_i^*)$  outside the circle. The intersection of the circle and the ellipses divides the graph into two parts. The part inside the circle has a smaller area, while the part outside the circle has a larger area. To this end, the sampling density comes into the scope, since the number of samples in each area and their strengths will both affect the final integration result of the path gain. In our implementation, we sample the TX radiation sphere uniformly, which results in a non-uniform sampling on the scattering surface. The closer to the TX's projection, the denser the sampling. Therefore, the part inside the circle has more samples compared to the part outside the circle. This results in two contradictory forces affecting the path gain integration: the part inside the circle has a smaller area, less amplification, but more samples, while the part outside the circle has a larger area, more amplification, but fewer samples. Jointly, these two forces lead to the interesting variation pattern of the deviation with respect to the incident angle.

In summary, the variation is an artifact of the non-uniform sampling on the scattering surface. We identify this as a future research direction for better sampling strategies to fix this artifact. Finally, we investigate the maximum deviation caused by this artifact. In this experiment, we sweep the TX/RX to reflection point distance from 1 m to 50 m, and for each combination, we sweep the incident angle from 0 to  $90^\circ$  to find the maximum deviation. We plot the heatmap of maximum deviation versus

the distances in Fig. 16. We note that for most areas of the plot, the maximum deviation is less than 1 dB, which is good enough for coverage estimation. The maximum deviation increases for larger distances, but it is still bounded by 3 dB even when both distances are 50 m. This indicates that the artifact is not significant in practical scenarios.

### C. Training and Model Details

Here, we provide additional details regarding the training procedure and model architecture used in our experiments. We train our mmDiff model using the Adam optimizer with an initial learning rate of  $5 \times 10^{-4}$ . In each scene, we train the model for 100 epochs, which takes approximately 4 hours on a single NVIDIA RTX 3090 GPU. We use a single receiver per batch due to the current architecture limitation. We expect that using multiple receivers per batch can further stabilize and accelerate the training process.

Our model architecture consists of a 4-layer MLP with 64 hidden units per layer. We use a 6-octave positional encoding to encode the 3D coordinates of the ray-surface interaction points before feeding them into the MLP. The output of the MLP is clipped and transformed using an exponential function to the expected range of material parameters, i.e., the relative permittivity  $\epsilon_r \in [1, 200]$  and the conductivity  $\sigma \in [10^{-3}, 10^6]$  S/m; the scattering coefficient and cross-polarization discrimination parameters are both in the range of  $[0, 1]$ . In total, our model has 29,700 learnable parameters. We found that this architecture strikes a good balance between model capacity and training efficiency for our task.

For the Sionna baseline models, we follow the training procedure in [11] and train the model for 10000 iterations. We observe that the learning curves have converged at this point. For the NeRF<sup>2</sup> baseline, we follow the training procedure in [36] and train the model for 30,000 iterations to converge.

Table 3. Scene Statistics

Scene Name	Dimensions (m)	# Objects	# Materials	# Test Samples	Training Sample Density (#samples/m <sup>3</sup> )
Office 0	4.4 × 5.0 × 3.0	68	10	489	42.4
Office 1	4.3 × 2.9 × 2.7	52	8	500	81.1
Office 2	4.0 × 7.2 × 2.8	94	8	468	35.3
Office 3	8.0 × 5.0 × 3.1	113	7	476	22.6
Office 4	6.5 × 6.5 × 2.8	71	7	500	23.3
Room 0	7.8 × 4.7 × 2.8	94	8	494	27.3
Room 1	6.1 × 3.5 × 2.8	57	8	474	47.4

## D. Dataset Details




Figure 17. Additional Replica scenes used for evaluation: Office 1, Office 2, and Office 3.

Here, we present three additional Replica scenes, making the total dataset cover seven realistic indoor environments. The additional scenes, Office 1, Office 2, and Office 3, are illustrated in Fig. 17. We follow the same dataset generation and experimental procedures as described in the main paper. The camera/receiver poses for the TRAJ and RAND datasets are visualized in Fig. 18. For the sake of clarity, we down-sample the poses by 10x in the visualization. In each scene, the TRAJ dataset contains 2,800 training samples, and the RAND dataset contains around 500 test samples (we manually removed some out-of-scene poses). Detailed scene information is summarized in Tab. 3. We note that this dataset is built on top of the open-source Replica dataset [26], with additional radio material assignments, trajectory sampling, and wireless channel simulations. We will make this dataset publicly available to facilitate future research.

## E. Geometric Differentiability

We demonstrate the differentiability of mmDiff by showing its capability to calibrate simple geometry offsets. We adopt the single-reflector scene shown in Fig. 4, where both TX and RX arrays are placed symmetrically on the same side of a  $1 \times 1 \text{ m}^2$  planar reflector. We first record the reference AoA power spectrum when the reflector is centered at the origin with  $0^\circ$  rotation. We then translate/rotate the reflector by a small distance/angle. And run mmDiff’s simulator again to obtain the AoA power spectrum in the distorted scene. Over the course of training, we minimize the log-MAE of AoA power spectrum using Adam optimizer with a learning rate of 0.02. For each experiment, we make the changed geometry variable (translation distance/rotation angle) the sole optimization parameter.

Fig. 19 shows the calibration of geometry offset over the course of training for 200 iterations. Fig. 19a shows the calibration of offsets along the X/Y/Z axis. As expected, the translation offsets gradually converge to 0 m, correctly recovering the original reflector position. Similarly, we test the calibration of rotational offset along the X/Y axis. As shown in Fig. 19b, in both cases, the calibrated angles converge back to the ground-truth  $-0^\circ$ . We note that rotation around the Z axis (surface normal) has a negligible effect on the AoA power spectrum, leading to vanishing gradients; thus, calibration along this axis is ineffective.

This experiment shows that mmDiff is differentiable with respect to scene geometry and can successfully correct geometric offsets in a simple scene. However, calibrating large-scale complex scene meshes remains highly challenging, and we deem this as an interesting future direction.

## F. Additional Results

Here, we provide the evaluation results on additional scenes to further validate the effectiveness of our proposed mmDiff model. Table 4 shows the robustness evaluation under geometry noise on three additional scenes. As is shown, our mmDiff model consistently outperforms the Sionna baseline across all metrics, preserving almost perfect prediction accuracy even in the presence of geometry noise. The Average Angular Error (AAE) and Average Power Error (APE) are also consistently lower for mmDiff, demonstrating its robustness to geometry noise.

Table 4. Robustness under Geometry Noise for Additional Scenes (TRAJ Dataset).

Scene	Methods	Top-1 Peak			Top-5 Peaks		
		F1↑	AAE↓	APE↓	F1↑	AAE↓	APE↓
Office 1	Sionna	0.9339	2.43	0.72	0.8576	2.23	1.14
	<b>mmDiff</b>	<b>0.9893</b>	<b>0.36</b>	<b>0.16</b>	<b>0.9897</b>	<b>0.33</b>	<b>0.19</b>
Office 2	Sionna	0.8407	2.79	0.75	0.8120	2.48	1.30
	<b>mmDiff</b>	<b>0.9989</b>	<b>0.18</b>	<b>0.02</b>	<b>0.9924</b>	<b>0.19</b>	<b>0.04</b>
Office 3	Sionna	0.6743	4.38	1.77	0.6570	3.77	2.15
	<b>mmDiff</b>	<b>0.9975</b>	<b>0.25</b>	<b>0.02</b>	<b>0.9874</b>	<b>0.29</b>	<b>0.03</b>

Table 5. Prediction Accuracy Results for Additional Scenes (RAND Dataset).

Scene	Methods	Top-1 Peak			Top-5 Peaks		
		F1↑	AAE↓	APE↓	F1↑	AAE↓	APE↓
Office 1	Sionna	0.8756	2.79	0.92	0.8222	2.57	1.57
	NeRF <sup>2</sup>	0.1433	12.45	7.95	0.4662	12.62	6.45
	mmDiff w/o Calib.	0.7892	4.87	3.22	0.7498	3.95	3.45
	<b>mmDiff Full</b>	<b>0.9297</b>	<b>1.72</b>	<b>0.80</b>	<b>0.9145</b>	<b>1.70</b>	<b>1.02</b>
Office 2	Sionna	0.8219	2.93	0.91	0.7845	2.57	1.74
	NeRF <sup>2</sup>	0.1496	11.36	3.80	0.2338	9.59	4.64
	mmDiff w/o Calib.	0.7906	3.55	1.08	0.7829	3.11	1.99
	<b>mmDiff Full</b>	<b>0.9722</b>	<b>1.04</b>	<b>0.27</b>	<b>0.9534</b>	<b>0.99</b>	<b>0.47</b>
Office 3	Sionna	0.7882	4.92	4.80	0.7376	3.61	4.83
	NeRF <sup>2</sup>	0.0273	14.09	1.38	0.0450	13.01	3.40
	mmDiff w/o Calib.	0.8655	3.25	1.37	0.8210	2.85	1.93
	<b>mmDiff Full</b>	<b>0.9685</b>	<b>1.09</b>	<b>0.32</b>	<b>0.9437</b>	<b>0.96</b>	<b>0.50</b>

Table 5 presents the channel prediction accuracy results on additional scenes on the RAND poses dataset. Note that in this experiment, we train all the models on the TRAJ dataset and evaluate them on the RAND dataset to assess their generalization capabilities to unseen poses. We also include an ablation study of our mmDiff model without the calibration, in which we use a randomly initialized neural network for predicting the material parameters.

The results indicate that our full mmDiff model significantly outperforms both the Sionna and NeRF<sup>2</sup> baselines and the ablated version without calibration across all metrics. The failure of Sionna is mainly due to its inconsistency in the tracing paths, i.e., the paths traced in the ground truth scene are different from those traced in the reconstructed scene. This inconsistency causes difficulties for the model to learn the material parameters at the ray-surface interactions, leading to poor generalization to unseen poses. NeRF<sup>2</sup> performs poorly due to the sparsity of training samples. Basically, as presented in [17, 36], NeRF<sup>2</sup> requires extremely dense sampling (around 6,000 samples/m<sup>3</sup>) of the scene to effectively learn the radiation field representation. In our datasets, the sampling density is around 40 samples/m<sup>3</sup>, which is far below the training condition of NeRF<sup>2</sup>, leading to its failure in learning a good representation of the scene. The ablated version of mmDiff without calibration also shows inferior performance compared to the full model, highlighting the importance of the calibration step in accurately estimating material parameters for better generalization.

We also provide visual comparisons of the predicted AoA power spectrum at sample poses in Fig. 20 to qualitatively

demonstrate the superior performance of our mmDiff model in accurately predicting the mmWave channel characteristics compared to the baselines.

## G. Multipath Interference

mmDiff adopts a Monte-Carlo-based algorithm to integrate the power of signals arriving from different directions. However, it discards the phase of rays at the ray-surface interaction point and accumulates the real-valued power of the ray weighted by the relative phases across the antenna array.

This is not a problem when there is a single dominant path – the AoA and power estimation remain correct. This is usually the case in mmWave due to the narrow beam. Even when there is interference from multiple paths, for instance, from a sidelobe, it can be suppressed using the windowing method.

As illustrated in Fig. 21, we consider the single planar reflector case, where both TX and RX are equipped with 8x8 URA, both facing the reflector. There are two paths: the LoS and the reflection path, each with different amplitude and phase. Fig. 22 shows the AoA spectra produced by the coherent sum of two paths versus the incoherent sum produced by mmDiff. We apply a Hanning window to both TX and RX precoding vectors to reduce the side-lobe level. Note there is a minor difference between the two spectra, highlighted via yellow boxes. This is due to the inconsistent phase between the two ways of calculating the channel.

Despite the inconsistency, the calibration works fine. As shown in Fig. 23, since the reflection path (at  $0^\circ$ ) is much stronger than the LoS path (at  $45^\circ$ ), the loss is mostly contributed by the dominant path. By calibrating against the ground truth reference map, we can recover the path gain.

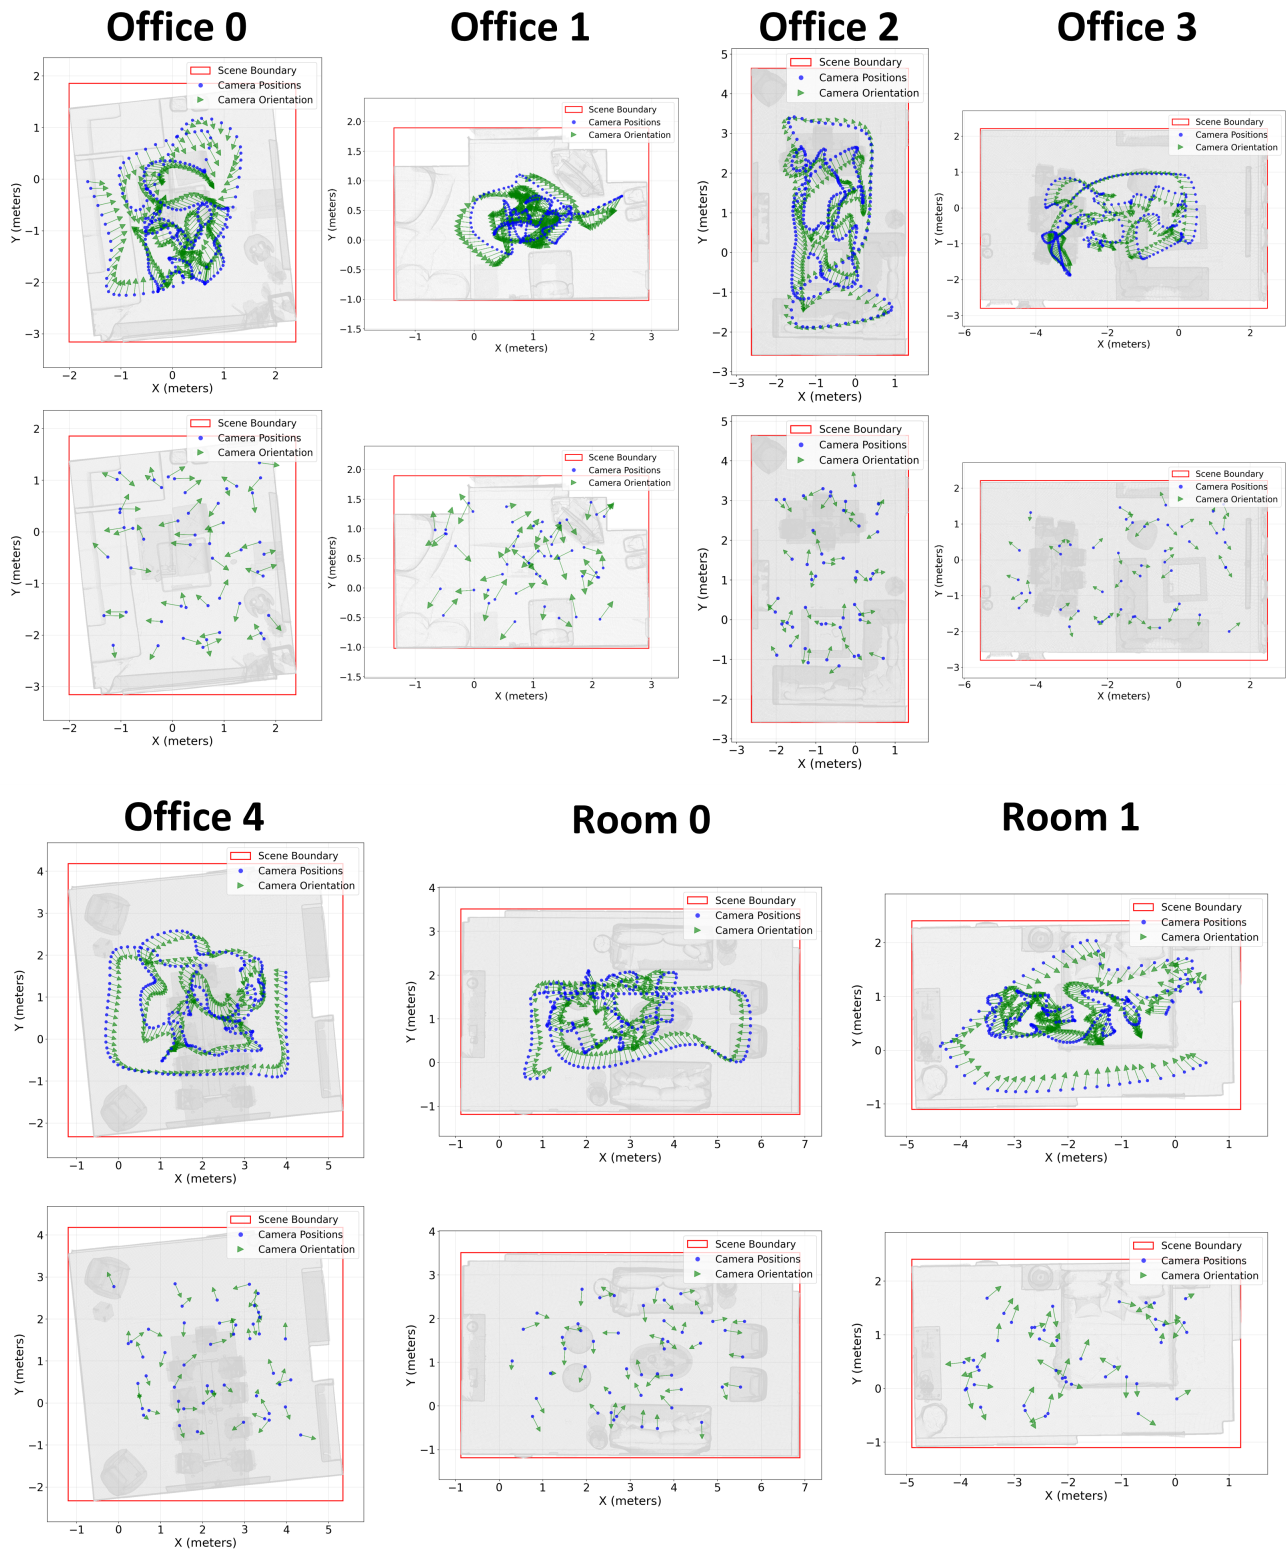
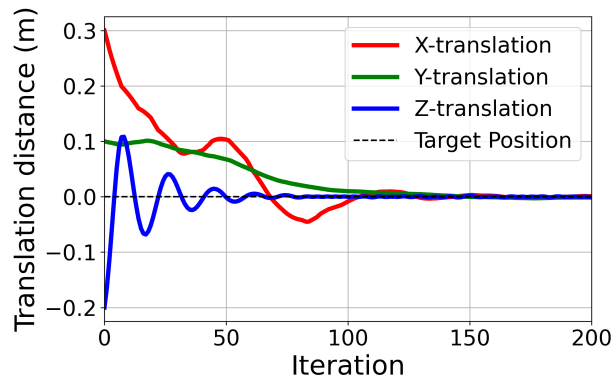
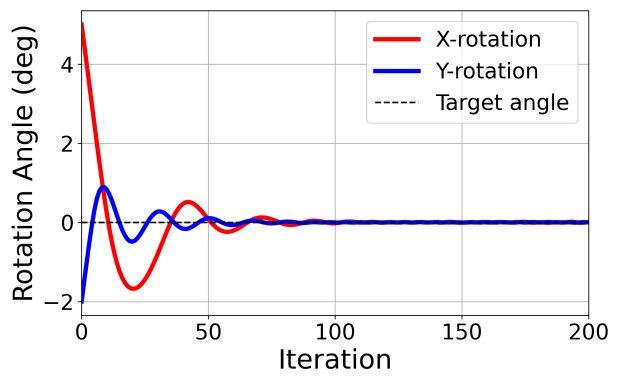


Figure 18. Sample poses in the TRAJ and RAND datasets (subsamped by 10x for visualization). Top: TRAJ poses follow a camera trajectory around the scene, used for training. Bottom: RAND poses are randomly distributed around the scene, used for testing. Note this is the 2D projection for visualization, the actual positions and orientations are in 3D.



(a) Calibration of translational offsets.



(b) Calibration of rotational offsets.

Figure 19. mmDiff's differentiability enables geometry calibration.

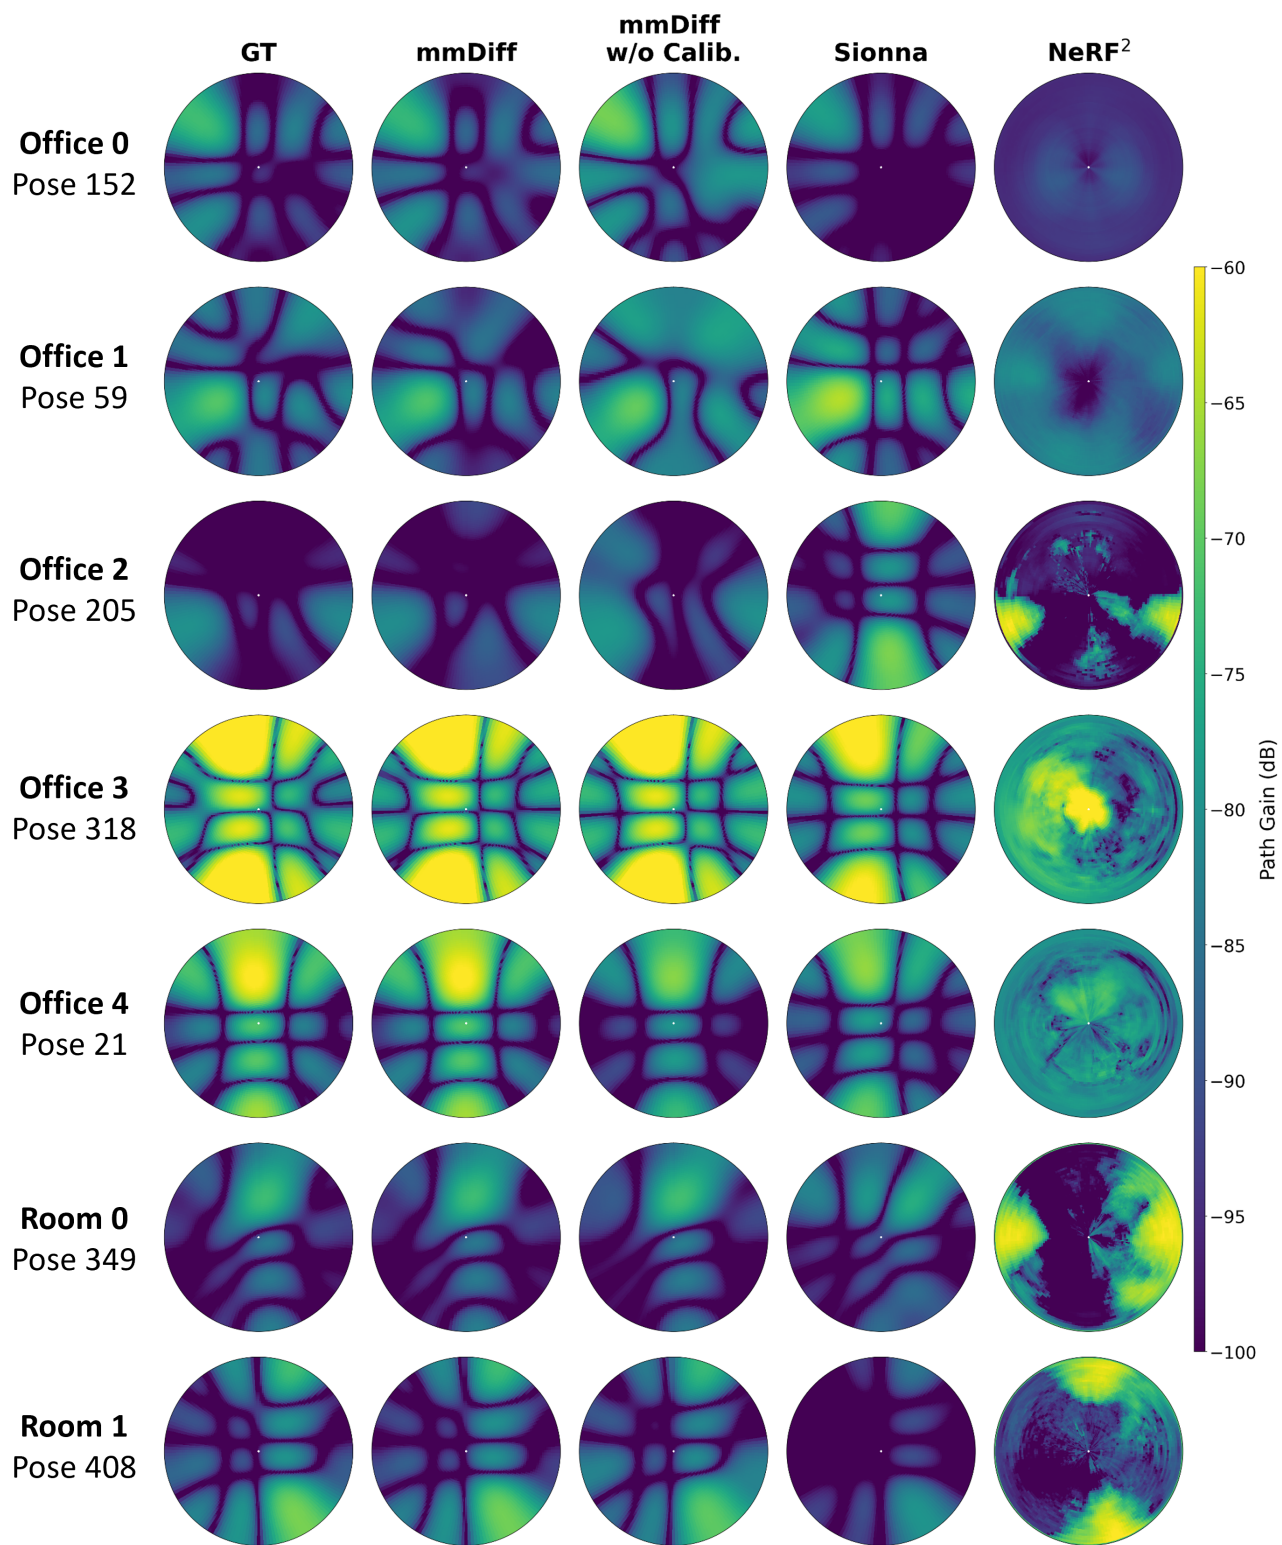


Figure 20. Prediction of AoA power spectrum at sample poses for visual comparison.




Figure 21. Illustration of the multipath interference

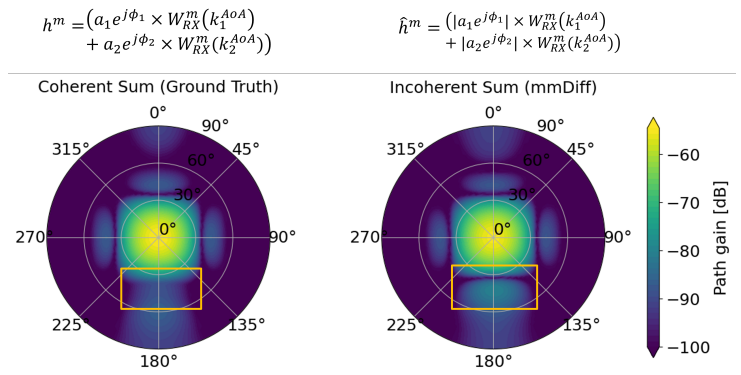


Figure 22. Multipath can cause inconsistent AoA spectra due to the loss of phase information




Figure 23. Material calibration with single dominant path.