

Neural Reflectance Fields for Radio-Frequency Ray Tracing

Haifeng Jia¹, Xinyi Chen², Yichen Wei¹, Yifei Sun¹, and Yibo Pi¹

¹Shanghai Jiao Tong University, Shanghai, China

²East China Branch of State Grid Corporation of China, Shanghai, China

Email: yibo.pi@sjtu.edu.cn

Abstract—Ray tracing is widely employed to model the propagation of radio-frequency (RF) signal in complex environment. The modelling performance greatly depends on how accurately the target scene can be depicted, including the scene geometry and surface material properties. The advances in computer vision and LiDAR make scene geometry estimation increasingly accurate, but there still lacks scalable and efficient approaches to estimate the material reflectivity in real-world environment. In this work, we tackle this problem by learning the material reflectivity efficiently from the path loss of the RF signal from the transmitters to receivers. Specifically, we want the learned material reflection coefficients to minimize the gap between the predicted and measured powers of the receivers. We achieve this by translating the neural reflectance field from optics to RF domain by modelling both the amplitude and phase of RF signals to account for the multipath effects. We further propose a differentiable RF ray tracing framework that optimizes the neural reflectance field to match the signal strength measurements. We simulate a complex real-world environment for experiments and our simulation results show that the neural reflectance field can successfully learn the reflection coefficients for all incident angles. As a result, our approach achieves better accuracy in predicting the powers of receivers with significantly less training data compared to existing approaches.

I. INTRODUCTION

Ray tracing is a fundamental tool for the accurate modelling of electromagnetic (EM) wave propagation in complex environments, which enables a diverse range of applications, e.g., channel modelling [1], antenna design [2], 5G deployment [3], and indoor localization [4]. Employing ray tracing in real-world environment requires a high-quality scene model including the geometry of objects and their material properties. With the advances in computer vision and LiDAR, 3D scene geometry and material estimation has become increasingly accurate and efficient [5], [6]. Material estimation in computer vision focuses on identifying the material properties, e.g., surface color, roughness, and metallic, that affect visual effects [7]. In contrast, ray-frequency (RF) ray tracing cares about the material reflectivity, which determines the amplitude and phase changes of the RF signal before and after reflection. Material reflection coefficient is not a constant for the given material but a function of several factors including incident angle, thickness, etc [8]. It is challenging to accurately measure the material reflectivity in complex real-world environments in an efficient way.

The early work on material reflectivity measurement can be dated back to decades of years ago, where a pair of TX and RX nodes moves around the material sample to measure the attenuation of the reflected signal at different angles [9]. The material sample size and the antenna aperture have to be carefully selected, and electromagnetic absorbers are placed behind the material sample to combat the multi-path effects. Scaling this approach to complex environments requires measuring each material in a controlled experiment separately, which is apparently not scalable. Along the same line of work, the feasibility of material identification with a single mmWave radio is studied in [10]. While this approach greatly reduces the efforts for material reflectivity measurement, it is still labor-intensive for practical use in large-scale scenes. Further, compared to millimeter wave, it is more challenging to measure material reflectivity using low frequency waves due to rich paths between TX-RX pairs [11].

Scalable material reflectivity measurement is made possible with the material identification techniques in computer vision, where material types in real-world environments are first identified from the scene images and then used to infer material reflectivity. For outdoor environments, street images readily available from the Google map are used to construct a 3D material map for mmWave communications [12]. Together with ray tracing, the 3D material map can be used to determine the best direction for beamforming. However, as mentioned before, material reflectivity depends not only on material type, but also the incident angle and material thickness, which cannot be explicitly inferred from images. It is possible to compute the reflection coefficient using the Fresnel equations, but the estimated coefficients may deviate significantly from the actual ones [13].

In recent years, the neural radiance field (NeRF) [14] has been shown a great success in the ray tracing of light. Given a few images of the scene, it can learn a continuous representation of the scene's volumetric properties to generate new images from different angles without the need for a scene model. NeRF² [15] extends the neural radiance field from optics to electromagnetism by considering the reflection, diffraction, and scattering of RF signal. NeRF² represents scenes as neural radiance fields and optimize the neural networks with RF signal measurements. Given a transmitter with known position, NeRF² can accurately predict the receive




Fig. 1. Our proposed RF ray tracing framework with the neural reflectance field. Our framework includes three components: (1) **Ray tracing module**, which employs current ray tracing algorithms to discover the strong rays for each TX-RX pair, maps the reflection points to their surface IDs, and steams the tuples, (incident angle, surface ID), to the neural reflectance fields. (2) **Neural reflectance field**, which takes as input the tuples and outputs the attenuation and phase change of the RF signal. (3) **Rendering module**, which composes point-level attenuation into ray-level attenuation, aggregates across multiple rays to compute the received signal for each TX-RX pair, and sums up the received signal from all transmitters to compute the receive power.

power at any position when the data density is sufficient. However, NeRF² has three key drawbacks. First, its performance depends on the data density and may experience high prediction errors under low data density. Second, it fails to achieve accurate prediction in regions where sudden jumps in measurements exist even when data density is high. Most importantly, although the predicted receive power is accurate, the learned volumetric scene function cannot reflect the real radio propagation environment. These issues are inherent to data-driven approaches [16].

To resolve the above issues, we propose RF ray tracing with accurate scene models. Specifically, to predict the receive power at a position, we first leverage ray tracing algorithms to discover all strong rays from the transmitter to the position, compute the receive power at each ray, and add up ray-level complex receive powers. As mentioned, an accurate scene model includes the 3D geometry of objects and their material properties. Considering the success of 3D scene reconstruction (e.g., OpenStreetMap [17] and Google Map), we simply assume that the 3D scene geometry is known and focus on learning the complex material reflectivity functions, which are challenging to measure at scale. To address this challenge, we leverage the strong descriptive power of neural networks and propose neural reflectance fields for materials in real-world environments, where the neural reflectance fields are capable of learning the material reflection coefficients at different incident angles. For efficient and scalable measurement, we can learn the neural reflectance fields in complex outdoor environments using the receive power of user equipments (UEs) in cellular networks and the 3D scene geometry provided by OpenStreetMap. Our approach can be easily extended to indoor scenarios with 3D scene geometry constructed by popular mapping tools, e.g., SLAM [18].

Specifically, we translate the neural reflectance fields from optics to RF domain with a complex-valued multi-layer perception (MLP), where both the amplitude and phase of RF signals are considered to model the constructive and destructive effects resulting from multipaths. The neural reflectance fields take as input the incident angle and surface ID to predict the amplitude and phase changes due to reflection of the

surface. In our neural network, we compute attenuation at the ray level using the surface-level attenuation predicted by the neural reflectance fields. This allows the neural network to decompose the ray-level attenuation from the combined receive power of a UE, which is a mixture of signals from both the serving and interfering base stations (BSs). Experimental results show that our neural network excels at predicting not only the combined receive power of a UE, but also the respective attenuation for each of the communication and interference channels, which have great potentials for resource allocation and scheduling tasks.

In summary, our contributions in this paper are as follows.

- We propose a RF ray tracing framework that combines traditional ray tracing with the neural reflectance field to learn the complex material reflectivity functions in real-world environments.
- We translate the neural reflectance field from optics to RF domain with a complex-valued MLP and optimize the neural reflectance field to match the receive power of RF signals.
- We demonstrate with simulations that our approach outperforms NeRF² in predicting the total receive power and that our approach can accurately predict the receive power at the ray level, making our approach valuable to a diverse range of applications.

II. RAY TRACING WITH NEURAL REFLECTANCE FIELDS

A. Framework

We combine the general ray tracing framework with the neural reflectance field that accounts for both scene geometry and material reflectance. We decompose the scene geometry into separate surfaces and assign each surface an ID as its reference. Each surface is assumed to adopt the same material such that material reflectance can be measured at the surface level. Unlike many existing works, we do not assume that the material types of surfaces are known in advance and thus measure material reflectance separately for each surface in the scene. As shown in Fig. 1, given the scene geometry and the locations of TX and RX nodes, we can employ traditional ray tracing algorithms to discover all strong rays from TX

nodes to RX nodes. For each reflection point along a ray, we obtain the incident angle of the ray and the surface ID at the reflection point. Then, the neural reflectance field regresses attenuation $\delta(\theta, s)$ from the incident angle and surface ID for all reflection points along the ray, where θ is the ray's incident angle at the reflection point and s is the surface ID. The renderer takes as input the attenuations at all reflection points along a ray, the free-space path loss (FSPL), and the transmit powers of transmitters to calculate the receive powers at different locations. This framework is differentiable, allowing us to fit the neural reflectance field to match the RF signal measurements.

B. Neural Reflectance Field

The neural reflectance field aims to learn the amplitude and phase changes of the ray at all reflection points. Let $\Delta A(\theta, s)$ and $\Delta\Theta(\theta, s)$ denote the amplitude and phase changes of the ray incident on the s -th surface at an angle of θ , respectively. We can represent the *reflection coefficient* as $\delta(\theta, s) = \Delta A(\theta, s)e^{j\Delta\Theta(\theta, s)}$. Fig. 2(a) shows an example of the relation between reflection attenuation (amplitude change) and the angle of incidence for several typical materials in the MATLAB toolbox. The neural reflectance field can be represented as:

$$F_w : (\theta, s) \mapsto (\Delta A(\theta, s), \Delta\Theta(\theta, s))$$

where w indicates the neural network weights. Unlike NeRF², the neural reflectance field takes as input the incident angle(s) of the ray, because the material reflectivity depends on the incident angle. Other impacting factors to material reflectivity, e.g., thickness, are not considered in the neural reflectance field because these factors are static. To consider the multipath effects in RF domain, we need to consider both the amplitude and phase of RF signal and thus fit the reflectance field using a complex-valued neural network. We use a classic MLP architecture composed of eight fully connected layers with the ReLU activations and 256 channels, where a skip connection concatenates the input with the fifth layer's activation [19]. As a data-driven approach, the neural reflectance field experiences the data density issue as shown in Fig. 2(b), where the density of incident rays varies across different angles. It is expected that the neural reflectance field can achieve better accuracy in estimating reflection attenuation in dense regions than in sparse ones. We want to ensure data density by measuring a sufficient number of RX nodes.

C. RF Domain Rendering

In the optics domain, rendering is the process to synthesize 2D images from a representation of the 3D scene. In the RF domain, we want the renderer to compute the desired features of the RF signals at given positions. In this work, we focus on the receive power of the RF signal. We consider a network with multiple transmitters, where the i -th transmitter transmits at the power of P_i^{tx} . Let $\mathcal{B}$ be the set of TX nodes and $H_{i,d}$ be the channel between TX i and RX d . The received signal




Fig. 2. Learning the relation between reflection attenuation and the angle of incidence with the neural reflectance field. Incident rays may differ in density at different angles, resulting in sparse regions to have larger learning errors for the reflection attenuation.

at RX d is a combination of the received signal from all TX nodes, which can be written as

$$S_k^{rx} = \sum_{i \in \mathcal{B}} \sqrt{P_i^{tx}} H_{i,d}. \quad (1)$$

Since the received signal is complex-valued, we compute the receive power as

$$P_k^{rx} = |S_k^{rx}|^2, \quad (2)$$

where the omnidirectional antenna is assumed with uniform gain in each direction. The channel between the TX and the RX can be further decomposed into multiple rays. Let $\mathcal{P}_i$ be the set of reflection points along the i -th ray, $r_{i,k}$ be the k -th reflection point on ray i , $s(r_{i,k})$ be the surface ID of the reflection point $r_{i,k}$, and $\theta(r_{i,k})$ be the incident angle of ray i at the reflection point $r_{i,k}$. We can obtain the reflection points along a ray after applying the ray tracing algorithms on the 3D scene geometry. Let d_i be the length of ray i . According to the Friis equation, we can compute the amplitude and phase changes of ray i as

$$H_i^{FS} = \frac{\lambda_c}{4\pi d_i} e^{j\frac{2\pi d_i}{\lambda_c}}, \quad (3)$$

where λ_c is the wavelength of the carrier wave and $2\pi d_i/\lambda_c$ is the phase rotation due to signal traveling from the TX to RX node. Let δ_k be the total reflection attenuation along ray k . We can express δ_k as

$$\delta_k = \prod_{r_{k,l} \in \mathcal{P}_k} \Delta A(\theta(r_{k,l}), s(r_{k,l})) e^{j\Delta\Theta(\theta(r_{k,l}))}. \quad (4)$$

Let $\mathcal{R}_{i,d}$ be the set of rays from TX i to RX d . Multiplying the FSPL and reflection attenuation, we can express the channel between TX i and RX d as

$$H_{i,d} = \sum_{k \in \mathcal{R}_{i,d}} H_k^{FS} \delta_k. \quad (5)$$

Combining all the above equations, we can rewrite the receive power of RX node d as in Eq. (6).

D. Loss Function

Let P_i^{rx} be the measured receive power for RX i and $\hat{P}_i^{rx}$ be the predicted RX power by the neural network for RX i . Suppose that we have measurements for N RX nodes at

$$P_d^{rx} = \left| \sum_{i \in \mathcal{B}} \sqrt{P_i^{tx}} \sum_{k \in \mathcal{R}_{i,d}} \frac{\lambda_c}{4\pi d_k} e^{j \frac{2\pi d_k}{\lambda_c}} \prod_{p \in \mathcal{P}_k} \Delta A(\theta(r_{k,l}), s(r_{k,l})) e^{j \Delta \Theta(\theta(r_{k,l}))} \right|^2 \quad (6)$$

different locations. We can use the following loss function to train our neural network:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N |P_i^{rx} - \hat{P}_i^{rx}|^2. \quad (7)$$

The weights of the neural reflectance network will be optimized to minimize the loss function. We measure the receive power in the unit of decibels, such that the weak receive power is treated equally as the strong receive power. The reason for this is that the weak receive power may consist of weak rays that are the only ones reflecting on a remote surface. We want to use as many rays as possible for measurements to avoid prediction errors due to low data density. In contrast, using milliwatts for the receive power will make the neural network inclined to optimize for the strong receive power, which is a collection of strong rays. This equates to underplay the importance of weak rays, making the surfaces solely reflected by weak rays prone to large measurement errors.

III. IMPLEMENTATION DETAILS

Data prerequisites. Training the neural reflectance field requires the geometry of the 3D scene, the locations of the TX and RX nodes, and the transmit/receive powers of the TX/RX nodes at those locations. Since scenes differ in their geometries, we have to run ray tracing algorithms for each scene separately to obtain the rays for each TX-RX pair, where the number of rays can be specified as needed. We train a separate neural reflectance field for each scene.

One-hot encoding. Due to the unordered and categorical characteristics of the surface IDs, we apply one-hot encoding to represent surface IDs. For a reflection point on the s -th surface, the input is encoded into a high-dimension vector:

$$E : s \mapsto \underbrace{(0, 0, \dots, \overbrace{1}^{s\text{-th}}, \dots, 0, 0)}_{\text{Number of all reflection surfaces}}, \quad (8)$$

where the dimension is determined by the total number of reflection surfaces. In our experiments, the encoder is applied to the input reflection surfaces respectively.

Ray tracing. Our framework includes a ray tracing module that discovers the rays between a TX-RX pair. We adopt the MATLAB ray tracing module to find rays and meanwhile obtain the coordinates of reflection points, angles of incidence at reflection points as well as the propagation distance of each ray. Shooting and bouncing rays (SBR) method [20] works as the ray tracing algorithm, with the maximum reflection times limited to 3 and no diffraction allowed. Thereafter, we want to map reflection points to their corresponding surfaces. We export the geometric information of all the surfaces by the Blender Python scripts, and then use ray casting algorithm to determine which surface the reflection point is on.

Virtual reflection points. After running the ray tracing algorithms, we can find the rays between each pair of TX and RX nodes. The number of rays may vary across different TX-RX pairs and, similarly, the number of reflection points may vary across rays between a TX-RX pair. In order to parallelize tensor operations, we insert virtual reflection points to enforce rays to have equal number of reflection points and TX-RX pairs to have equal number of rays. We intentionally set the attenuation of virtual reflection points to zero and their incident angles to 90 degrees. Besides, the transmit power for the rays with all virtual reflection points should be set to zero. Since the attenuations of virtual reflection points are constant, their derivatives are zero in backpropagation, thereby causing no influence on the real reflection points during the training. When using the neural network for inference, we have to ensure that the virtual reflection points use the same values as in training.

Network training configurations. Following the practice in NeRF², we employ the Adam optimizer with an initial learning rate of 1×10^{-3} and a weight decay of 5×10^{-5} . A cosine annealing scheduler is used to adjust the learning rate during the training with the maximum number of iterations $T_{max} = 10000$ and the minimum learning rate $\eta_{min} = 1 \times 10^{-6}$. Other hyper-parameters are set as default. We train our neural network on a single NVIDIA RTX 3090 GPU. It takes $90k \sim 110k$ iterations (about 2 hours) for our network training to converge.

IV. PERFORMANCE EVALUATION

A. Datasets

We synthesize our dataset using a campus-level 3D scene obtained from the OpenStreetMap, covering an area of roughly 1 km^2 and including 923 surfaces and 4 types of materials: water, vegetation, concrete, and brick. Each material has a simulated attenuation that is linear to the angle of incidence. As shown in Fig. 3, we deploy 3 TX nodes on the roof tops as micro BSs with a transmit power of 10 watts. Meanwhile, we randomly pick about 2000 RX nodes at the height of one meter from the ground and allow at most 21 rays for each RX and at most 3 reflections for each ray, where the number of rays is capped at 7 for each TX-RX pair. Each of the TX and RX nodes is equipped with a single omnidirectional antenna. We use 80% of the dataset for training and the rest for testing.

B. Accuracy of Neural Reflectance Field

Estimation accuracy of reflection coefficients. To evaluate the accuracy of the neural reflectance field, we want to compare its estimated reflection coefficient with the ground truth. Although our training set only includes rays with a limited number of incident angles for each surface, the neural



Fig. 3. Campus-level 3D scene geometry from OpenStreetMap





Fig. 4. Accuracy of the neural reflectance field

reflectance field is capable of predicting the reflection coefficient at any angle. For each surface, we generate the estimated reflection coefficients with randomly selected incident angles and compare them with the ground-truth coefficients computed from the simulated attenuation function. We define the estimation error as $\epsilon_\delta = |10 \log_{10}(|\delta_{\text{Estimate}}|/|\delta_{\text{GroundTruth}}|)|$, where δ_{Estimate} is the estimated reflection coefficient and $\delta_{\text{GroundTruth}}$ is the ground truth. When there is no error, ϵ_δ is equal to 0. Fig. 4(a) shows the mean estimation error and the 90-th, 95-th, and 99-th percentile errors under different angles of incidence, where the angles of incidence are evenly divided into 1000 bins and the percentile is computed for errors in each bin. It can be seen that the 99-th percentile error is below 2.5 dB for all incident angles and the mean error is below 0.6 dB. We can also see that the error decreases as the angle of incidence increases and rebounds when the angle comes near 90°. Based on our simulated attenuation function, reflection loss is large for small incident angles, resulting in a ray with strength weaker than others. Such weak rays have limited influence to the overall receive power and thus are likely to experience larger estimation errors. When the angle of incidence comes near 90°, the reflection coefficient is very small, which tends to suffer from large estimation errors.

Estimation accuracy by material type. We further look at the estimation error by material type. Fig. 4(b) shows the 99-th percentile estimation error by material type. We can see that the estimation error increases from water to brick, where water has the smallest attenuation and brick has the largest.



(a) Prediction error vs. data density (b) Estimation error vs. data density

Fig. 5. Comparison with existing approaches

From water to brick, the increased attenuation weakens the ray strength, thus greatly reducing its influence to the overall receive power. As a result, the reflection coefficients of brick surfaces are the most difficult to estimate.

C. Comparison with Existing Approaches

Prediction error. We compare our approach with the state-of-the-art approach, NeRF², in terms of the accuracy in predicting the receive power of RX nodes. We compute the prediction error as $\epsilon_P = |10 \log_{10}(P_{\text{Predict}}^{rx}/P_{\text{GroundTruth}}^{rx})|$, where P_{Predict}^{rx} and $P_{\text{GroundTruth}}^{rx}$ are the predicted and ground-truth receive powers, respectively. As shown in Fig. 5(a), our approach outperforms NeRF² in the mean prediction error under all data densities on the testing set. When data density is only 250 samples per km², our approach achieves a mean prediction error 14 dB less than NeRF². The prediction accuracy of NeRF² can be significantly improved by using higher data density, but its prediction error is still 2 dB higher than our approach at high data density. This is because NeRF² predicts the receive power of RX nodes at new locations with interpolation. Looking at the spatial distribution of receive powers, we find that there exist sudden jumps in receive power between nearby RX nodes, because the rays from a TX node to two close RX nodes may reflect on different reflection surfaces. This issue can be mitigated when data density is high, but is hard to entirely eliminate.

Robustness to low data density. Fig. 5(a) shows that our approach constantly achieves a low prediction error under different data densities. This is because ray tracing allows us to group reflections on the same surface together. In other words, reflections on the same surface can be used together to train the neural reflectance field for the surface. To achieve high prediction accuracy, our approach requires sufficient measurements at the surface level, while NeRF² requires high data density for all possible locations. When the space is large as in our experiments, NeRF² needs large volume of measurements to achieve high prediction accuracy. We further evaluate our approach's capability to accurately estimate material reflection coefficients, a task that NeRF² cannot perform. We follow the same way to compute the estimation error of reflection coefficients. From Fig. 5(b), we can see that the neural reflectance field can achieve low estimation errors with a mean of 0.8 dB when there are only 250 samples per km² and can also benefit from higher

data density. The reason for this is that more measurements at different angles improve the interpolation accuracy of the neural reflectance field.

D. SINR and Interference Channel Prediction

We evaluate the prediction error for the signal-to-interference-plus-noise ratio (SINR) and the interference channel gains. We define the *interference channel error* as the difference between the predicted and ground-truth interference channel gains between TX-RX pairs, where each RX node is served by the nearest TX node and interfered by the rest. We calculate the SINR error for each RX node in the testing set and show the distribution of SINR errors in Fig. 6(a), where almost all RX nodes have an accurate predicted SINR deviating from the actual one less than 0.05 dB. We further decompose the superimposed interference signals at the level of TX-RX pairs. Fig. 6(b) shows that our approach can accurately predict the channel gain for each interference channel, with only less than 1% of TX-RX pairs having a prediction error greater than 0.05 dB. Simulation results show that our ray tracing framework is powerful in predicting both communication and interference channel gains, which has great potentials to be used for resource allocation and scheduling tasks.

V. CONCLUSION

This paper focuses on solving the problem of scalable and efficient measurement of material reflectivity for ray tracing in complex real-world environments. To address this problem, we translated the neural reflectance field that have been proved a success in optics to the RF domain to learn the reflection coefficients at any angle from the receive power of RF signals. We designed a differential RF ray tracing framework to train the neural reflectance field such that the predicted receive powers match the measured ones. In our simulations, the neural reflectance field exhibits remarkable performance in learning the material reflection coefficients at different angles for surfaces in a large scene. Furthermore, our framework demonstrated the ability to decompose the superimposed interference signals and accurately predict the gain of each interference channel, which has great potentials for resource allocation and scheduling tasks.

VI. ACKNOWLEDGEMENT

We appreciate the constructive feedback from the anonymous reviewers. This work was supported by the East China Branch of State Grid Corporation of China under the Grant 529924240006.

REFERENCES

- [1] J.-H. Jung, J. Lee, J.-H. Lee, Y.-H. Kim, and S.-C. Kim, "Ray-tracing-aided modeling of user-shadowing effects in indoor wireless channels," *IEEE Trans. Antennas Propag.*, vol. 62, no. 6, pp. 3412–3416, 2014.
- [2] D. He, B. Ai, K. Guan, L. Wang, Z. Zhong, and T. Kürner, "The design and applications of high-performance ray-tracing simulation platform for 5G and beyond wireless communications: A tutorial," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 1, pp. 10–27, 2018.
- [3] A. W. Mbugua, Y. Chen, L. Raschkowski, L. Thiele, S. Jaeckel, and W. Fan, "Review on ray tracing channel simulation accuracy in sub-6 GHz outdoor deployment scenarios," *IEEE Open J. Antennas Propag.*, vol. 2, pp. 22–37, 2020.
- [4] P.-H. Tseng, Y.-C. Chan, Y.-J. Lin, D.-B. Lin, N. Wu, and T.-M. Wang, "Ray-tracing-assisted fingerprinting based on channel impulse response measurement for indoor positioning," *IEEE Trans. Instrum. Meas.*, vol. 66, no. 5, pp. 1032–1045, 2017.
- [5] A. Handa, T. Whelan, J. McDonald, and A. J. Davison, "A benchmark for RGB-D visual odometry, 3D reconstruction and SLAM," in *Proc. IEEE ICRA*, 2014, pp. 1524–1531.
- [6] T. Kühner and J. Kümmerle, "Large-scale volumetric scene reconstruction using LiDAR," in *Proc. IEEE ICRA*, 2020, pp. 6261–6267.
- [7] J. Zhang, Y. Yao, S. Li, J. Liu, T. Fang, D. McKinnon, Y. Tsin, and L. Quan, "NeILF++: Inter-reflectable light fields for geometry and material estimation," in *Proc. IEEE/CVF CVPR*, 2023, pp. 3601–3610.
- [8] J. Ahmadi-Shokouh, S. Noghmanian, and H. Keshavarz, "Reflection coefficient measurement for north american house flooring at 57–64 ghz," *IEEE Antennas Wireless Propag. Lett.*, vol. 10, pp. 1321–1324, 2011.
- [9] B. Langen, G. Lober, and W. Herzig, "Reflection and transmission behaviour of building materials at 60 GHz," in *Proc. PIMRC*, 1994, pp. 505–509.
- [10] C. Wu, F. Zhang, B. Wang, and K. R. Liu, "mSense: Towards mobile material sensing with a single millimeter-wave radio," in *Proc. ACM IMWUT*, vol. 4, no. 3, 2020, pp. 1–20.
- [11] W. Jiang, Q. Zhou, J. He, M. Habibi, S. Melnyk, M. El-Absi, B. Han, M. Di Renzo, H. Schotten, F. Luo, and T. El-Bawab, "Terahertz communications and sensing for 6G and beyond: A comprehensive review," *IEEE Commun. Surveys Tuts.*, 2024.
- [12] A. Kamari, Y. Chae, and P. Pathak, "mmSV: mmWave vehicular networking using street view imagery in urban environments," in *Proc. ACM MOBICOM*, 2023, pp. 1–16.
- [13] K. Sato, T. Manabe, T. Ihara, H. Saito, S. Ito, T. Tanaka, K. Sugai, N. Ohmi, Y. Murakami, M. Shibayama, and Y. Konishi, "Measurements of reflection and transmission characteristics of interior structures of office building in the 60-GHz band," *IEEE Trans. Antennas Propag.*, vol. 45, no. 12, pp. 1783–1792, 1997.
- [14] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *Commun. ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [15] X. Zhao, Z. An, Q. Pan, and L. Yang, "NeRF²: Neural radio-frequency radiance fields," in *Proc. ACM MOBICOM*, 2023, pp. 1–15.
- [16] G. Chi, Z. Yang, C. Wu, J. Xu, Y. Gao, Y. Liu, and T. X. Han, "RF-Diffusion: Radio signal generation via time-frequency diffusion," in *Proc. MOBICOM*, 2024, pp. 77–92.
- [17] OpenStreetMap contributors, "Planet dump retrieved from <https://planet.osm.org>," <https://www.openstreetmap.org>, 2017.
- [18] M. R. U. Saputra, A. Markham, and N. Trigoni, "Visual SLAM and structure from motion in dynamic environments: A survey," *ACM Comput. Surveys*, vol. 51, no. 2, pp. 1–36, 2018.
- [19] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "Deepsdf: Learning continuous signed distance functions for shape representation," in *Proc. IEEE/CVF CVPR*, 2019, pp. 165–174.
- [20] K. R. Schaubach, N. J. Davis, and T. S. Rappaport, "A ray tracing method for predicting path loss and delay spread in microcellular environments," in *Proc. IEEE VTS 42nd VTC-FOT*, 1992, pp. 932–935.




Fig. 6. Interference channel gain and SINR prediction