

Differentiable Ray Tracing for THz Radio Channel Characterization with Point Clouds

Niklas Vaara*, Naveeth Lafir[†], Peize Zhang[‡], Tuomas Määttä*, Lam Huynh*,
Pekka Sangi*, Miguel Bordallo López*, Pekka Kyösti^{†§}, Janne Heikkilä*

*Center for Machine Vision and Signal Analysis, University of Oulu, Oulu, Finland, firstname.lastname@oulu.fi

[†]Centre for Wireless Communications, University of Oulu, Oulu, Finland

[‡]Centre for Wireless Innovation (CWI), Queen’s University Belfast, BT3 9DT Belfast, United Kingdom

[§]Keysight Technologies Finland Oy, 90590 Oulu, Finland

Abstract—Differentiable ray tracing-based radio propagation modeling is a powerful tool for learning material properties from channel measurements. Prior approaches rely on idealized, manually crafted triangle mesh models, which often lack smaller geometric details that become important at higher frequencies. In this paper, we perform differentiable ray tracing directly on a noisy point cloud created from RGB-D images, eliminating the need for handcrafted geometry. We optimize channel parameters using THz channel measurements collected at a center frequency of 318 GHz with a 4 GHz bandwidth. Our results demonstrate that the channel characteristics can be extracted well with available measurements and ray tracing. With the optimized parameters, we achieve a mean relative rms delay spread error of less than 5% and a mean absolute error for received power of about 0.36 dB at novel receiver locations.

Index Terms—Ray tracing, radio propagation, point clouds, machine learning

I. INTRODUCTION

Physically accurate simulation of radio wave propagation is essential for designing and optimizing wireless communication systems. Ray tracing (RT) is a powerful method to simulate the channel in such manner, as it can exploit GPU-acceleration in the path computation [1]. The electromagnetic computations in the RT-based radio propagation simulations utilize frequency-dependent material properties, which in public databases such as ITU-R-P.2040 [2] are limited in terms of available materials. In addition, they are only defined for specific frequency ranges.

RT produces accurate propagation paths that contains rich information about the propagation channel that is in many cases difficult to extract from measurements, such as angle of arrival and departure. Unlike radio channel measurements, it can be used to quickly estimate the radio channel characteristics at thousands of positions, which is why it is a great tool for creating synthetic datasets for machine learning applications [3]–[8].

Radio channel measurements are still an important reference for acquiring accurate material properties. Recent works have explored the utilization of differentiable RT to learn the material properties from radio channel measurements. In [9], a small multilayer perceptron (MLP) was trained to predict the material properties at each ray-surface intersection point based on conducted channel measurements. In [10], a digital twin framework was proposed that utilizes neural representation for

both electromagnetic properties and the interaction behavior. It segments the environment model into objects, enabling adaptation to environment changes.

However, a key limitation of these prior approaches is their reliance on manually crafted triangle mesh models to represent the environment. While effective in controlled settings, these meshes often lack the fine geometric details, such as small objects or irregular surfaces. Creating detailed meshes is time consuming and does not scale well to dynamic environments. Recent works have reconstructed triangle meshes from data acquired from depth sensors such as laser scanners or RGB-D cameras [11]–[14]. The reconstructed triangle meshes are often a simpler representation of the environment to enable utilization with traditional ray tracers that expect noiseless surfaces, while simultaneously introducing additional overhead due to the surface reconstruction.

Another research direction has been utilizing the reconstructed point clouds directly [15]–[19]. However, often these methods require heavy downsampling of the point cloud to achieve acceptable computation times, which may lead to the loss of environment details. Furthermore, approaches such as [16], [18], [19] use conical rays, which significantly increases the number of rays to be traced at each interaction level, and thus, reduces the applicability for real-time or large scale simulations.

To address these challenges, we propose a framework for optimizing material properties based on radio channel measurements directly on top of point clouds reconstructed from RGB-D images. Our method utilizes the newest version of our open source ray tracer NimbusRT¹, presented in [20]. It is capable of computing paths in just tens of milliseconds and it interfaces over Sionna [21] for differentiable electromagnetic computations. In this work, we perform the optimization by utilizing THz channel measurements at different locations. As a result, the method eliminates the needs for handcrafted geometry and enables fully automated and data-driven modeling of real-world environments using consumer-grade sensors. We validate our method by evaluating the measured and simulated power delay profiles (PDPs) at novel receiver (Rx) locations.

¹<https://github.com/nvaara/NimbusRT>

II. METHOD

Our method computes PDPs from simulated channel impulse responses (CIRs), which are compared against measured PDPs to define a loss. This loss is then optimized through gradient-based refinement of the material labels assigned to the points visualized with distinct colors in Fig. 1.

A. Ray Tracing

For simulation, we employ NimbusRT [20], a RT-based tool for radio channel characterization that operates on point clouds. It supports high-performance modeling of electromagnetic wave propagation in complex environments and is robust to sensor noise, making it well-suited for real-world RGB-D camera data. This enables fully automated segmentation and 3D reconstruction pipelines, allowing the creation of environment models without manual intervention.

As an input the ray tracer expects an annotated point cloud, where the most relevant attribute for this work is the material label assigned to each point. The simulator allows tuning of parameters such as maximum interaction depth and the types of propagation mechanisms to be simulated. Based on these inputs, it computes exact propagation paths through gradient descent-based path length minimization based on [22] that adapts to noise and the underlying geometry. It is capable of simulating different propagation mechanisms, namely reflections, diffractions, and diffuse scattering. For electromagnetic computations, NimbusRT interfaces over Sionna [21], enabling differentiable computation with respect to parameters such as relative permittivity and conductivity. In our setup, CIRs are computed from the paths with

$$h_k = \sum_i^M \text{sinc}(k - \tau_i B) a_i, \quad (1)$$

where a_i and τ_i are the complex amplitude and time delay of i th path, B is the bandwidth, k is the delay bin index.

B. Training

We optimize material properties, namely relative permittivity $\epsilon_r \geq 1$ and conductivity $\sigma > 0$, assigned to each material label. We use activation functions to constrain them within their valid ranges. Our goal is to fit our ray traced paths into the measured PDPs through differentiable electromagnetic computations by optimizing the material properties. In addition, we learn a gain coefficient \mathcal{G} , which represents the generic total channel gain.

1) *Loss Functions*: In order to fit our simulated PDP to the measured one, we use a rms delay spread loss, inspired by [9]. In addition, we minimize the difference between the mean excess delays to further enforce a similar power distribution. The loss is formulated as

$$\mathcal{L}_{\text{ds}} = E_r(\tau_{\text{mean}}, \hat{\tau}_{\text{mean}}) + E_r(\tau_{\text{rms}}, \hat{\tau}_{\text{rms}}), \quad (2)$$

where τ_{mean} and $\hat{\tau}_{\text{mean}}$ are the simulated and measured mean excess delays, and similarly, τ_{rms} and $\hat{\tau}_{\text{rms}}$ are the simulated and measured rms delay spreads, respectively.

$$E_r(a, b) = \frac{|a - b|}{b} \quad (3)$$

computes the relative error between the inputs a and b ,

$$\tau_{\text{rms}} = \sqrt{\frac{\sum_k p_k (\tau_k - \tau_{\text{mean}})^2}{\sum_k p_k}}, \quad (4)$$

and

$$\tau_{\text{mean}} = \frac{\sum_k p_k \tau_k}{\sum_k p_k}. \quad (5)$$

In the aforementioned equations, p_i and τ_i are the power and delay of each delay bin, respectively. Lastly, as \mathcal{L}_{ds} does not account for the received power, we minimize it as follows

$$\mathcal{L}_{\text{pw}} = E_r(\mathcal{P}(\mathbf{h})\mathcal{G}, \mathcal{P}(\hat{\mathbf{h}})), \quad (6)$$

where \mathbf{h} and $\hat{\mathbf{h}}$ are the simulated and measured CIRs, and

$$\mathcal{P}(\mathbf{h}) = \sum_k |h_k|^2. \quad (7)$$

Finally, our total loss \mathcal{L} is formulated as

$$\mathcal{L} = \mathcal{L}_{\text{ds}} + \lambda_{\text{pw}} \mathcal{L}_{\text{pw}}, \quad (8)$$

where $\lambda_{\text{pw}} = 0.1$ is a regularization hyperparameter.

III. CHANNEL MEASUREMENTS AND ENVIRONMENT RECONSTRUCTION

The radio propagation measurement campaign was conducted in the Saalastinsali Hall at the University of Oulu, Finland. The floor layout of the measurement setup is shown in Fig. 2. The hall consists of a large open area at ground level and an elevated stage.

A. Channel Measurements

The transmitter (Tx) antenna was placed on the stage at a height of 1.85 m, while the Rx antenna was positioned at a height of 1.42 m and moved to multiple locations across the hall. The hall walls are composed of metal, concrete, and glass surfaces. To minimize unwanted reflections, all furniture was removed before the measurements. The measurements were conducted using a vector network analyzer (VNA)-based channel sounder operating at a center frequency of 318 GHz with a signal bandwidth of 4 GHz, with 1001 frequency points spanning the bandwidth. Both the Tx and Rx used horn antennas with half-power beamwidths (HPBWs) of 14.2° and 14.4° in the azimuth and elevation planes, respectively.

During the measurement campaign, both the Tx and Rx antennas were rotated to capture the angular characteristics of the radio channel. The Tx elevation angle was fixed at 90° , while its azimuth was varied from 38° to 143° in 15° increments. The Rx antenna was rotated at two elevation angles (i.e., 70° and 90°) while sweeping the azimuth from 0° to 360° in 10° steps. The channel frequency responses (CFRs) were recorded for all Tx-Rx angular combinations. The corresponding directional channel impulse responses (DCIRs) were then extracted by applying an inverse fast Fourier transform (IFFT) combined with a Blackman-Harris window.

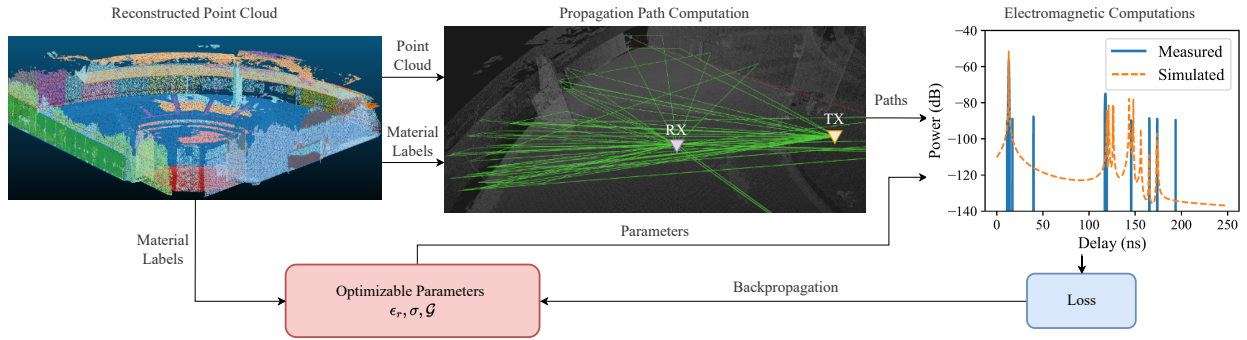


Fig. 1. Overview of the steps involved in our method. Our method takes a reconstructed labeled point cloud as the input, and creates a new material for each one of the labels. Each ray traced propagation path contains a link to the intersection material label. This information is passed to differentiable electromagnetic computations, which enables optimization of the material parameters.

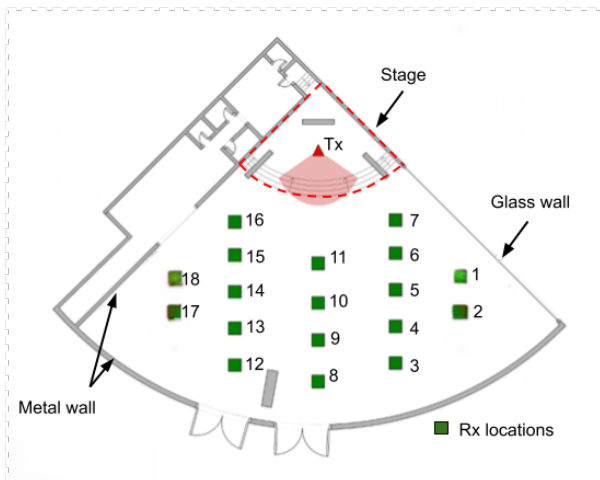


Fig. 2. The layout of Saalastinsali hall

1) *Measurement Pre-processing*: The constructed PDP from the DCIRs appears as white noise due to phase offsets between the DCIRs, as we show in Fig. 3a. However, if we look at the separate PDPs constructed from the DCIRs in Fig. 3b, we can see distinct peaks of the contributing multipath components. Thus, we threshold each complex amplitude a by

$$10 \log_{10}(|a|^2) \geq t_{\text{dB}}, \quad (9)$$

where t_{dB} is the threshold value the sample power. Heuristically, we chose t_{dB} as -90 dB. Next, we choose the sample with the largest power at each delay bin among all of the DCIRs, and lastly, with these filtered multipath components we compute a bandlimited CIR with Eq. (1). We showcase the resulting PDP in Fig. 3c.

B. Environment Reconstruction

The environment was reconstructed from RGB-D data and refined using the pre-processing, segmentation, diffraction-edge detection, and material annotation modules of the pipeline described in [23]. Chairs and other movable objects were manually removed during post-processing to match the

measurement setup, while glass surfaces, not captured by the depth sensor, were also manually added.

IV. EXPERIMENTS RESULTS AND DISCUSSION

We compute the ray traced propagation paths with the reconstructed point cloud for each Tx-Rx pair presented in Fig. 2. The Rx's are split into train and test datasets, where the Rx1 and Rx13 are used for testing and the rest are used for training. The optimization is performed for 1000 iterations using an Adam optimizer with a learning rate of $1e-2$. The experiments were ran on a computer equipped with an NVIDIA GeForce RTX 3080 GPU.

A. Training Results and Analysis

The qualitative results on the test data are illustrated in Fig. 4, while the quantitative metrics are summarized in Table I. We assess the relative delay spread error in Fig. 4a. Both test scenarios exhibit clear convergence, demonstrating a consistent reduction in delay spread error. Specifically, the relative delay spread error τ_{err} falls below 9% for Rx1 and under 1% for Rx13. Similarly, a consistent per iteration improvement can be observed for the gain coefficient \mathcal{G} in Fig. 4b. After the training iterations, \mathcal{G} applied to \mathcal{P} reaches 42.87 dB. The resulting mean absolute error of the received power is less than 0.36 dB. While the peak antenna gain is 25 dBi, it is important to note that not all propagation paths benefit from this maximum gain in the measurements. Due to the directional nature of the antenna patterns, many paths arrive at non-optimal angles, resulting in lower effective gain. The learned gain coefficient \mathcal{G} therefore reflects a composite contribution across all paths, not just those aligned with the peak beam direction.

The comparison of measured and simulated PDPs at Rx1 and Rx13 is shown in Fig. 4c and Fig. 4d, respectively. From the figures we can observe that our method is capable of finding the significant reflection peaks with similar amplitudes in both cases. In the case of Rx1, a diffuse-like peak can be observed, suggesting additional arriving paths. To confirm this, we performed the optimization with second order reflections as well. As one can observe from the qualitative results shown

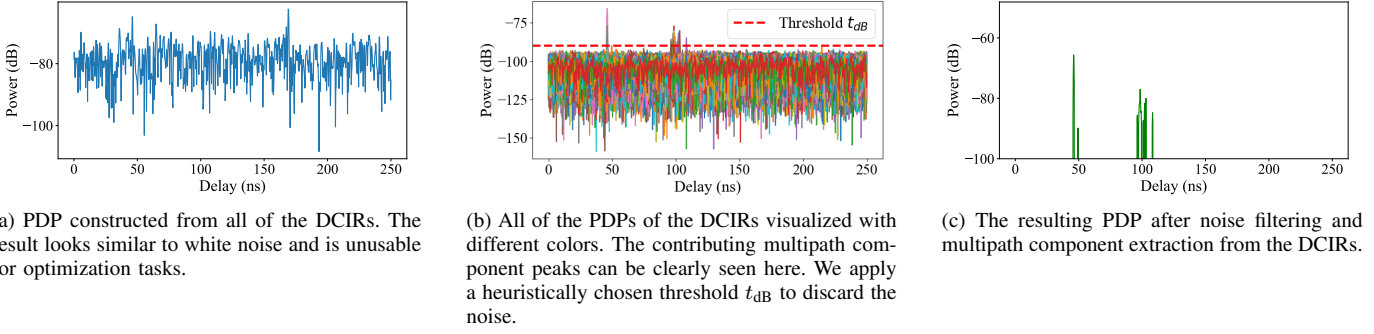


Fig. 3. Measurement processing steps to determine contributing multipath components with minimal noise for training.

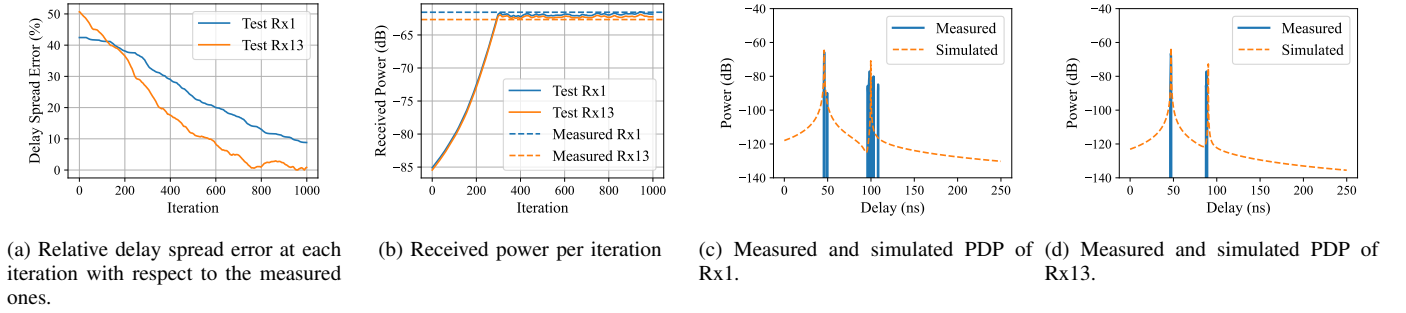


Fig. 4. Training and testing results for one reflection case with 1000 iterations.

TABLE I

QUANTITATIVE RESULTS FOR THE TEST RXS. τ_{rms} AND \mathcal{P} ARE THE SIMULATED RMS DELAY SPREAD AND RECEIVED POWER. THE MEASURED COUNTERPARTS ARE $\hat{\tau}_{rms}$ AND $\hat{\mathcal{P}}$. THE SYMBOLS τ_{err} AND \mathcal{P}_{err} ARE THE RELATIVE DELAY SPREAD ERROR AND ABSOLUTE RECEIVED POWER ERROR, RESPECTIVELY.

Case	τ_{rms} (ns)	$\hat{\tau}_{rms}$ (ns)	τ_{err} (%)	\mathcal{P} (dB)	$\hat{\mathcal{P}}$ (dB)	\mathcal{P}_{err} (dB)
Rx1	18.11	19.86	8.82	-61.84	-61.56	0.28
Rx13	12.41	12.31	0.84	-62.24	-62.67	0.43

in Fig. 5, there is a new contributing multipath component near the diffuse peak in the case of Rx1 and multiple second order reflection peaks with Rx13. To further study these paths, we generated a new CIR from the measurements with (1) and (9) using a t_{dB} of -100 dB to see if these peaks can be seen in the data. The qualitative results are shown in Fig. 6, where similar peaks near the noise level can be observed, which we suspect to be from second order reflections. However, they are hard to distinguish from the noise, which is why we focused on first order reflections for the purpose of optimization.

B. Discussion

The presented results demonstrate that our method is capable of efficiently modeling the channel characteristics found in the measurements at novel receiver locations with a mean τ_{err} of less than 5%. This level of precision suggests that the ray traced paths agree with the measurements. Notably, the high accuracy is mainly due to the presence of only first order reflections.

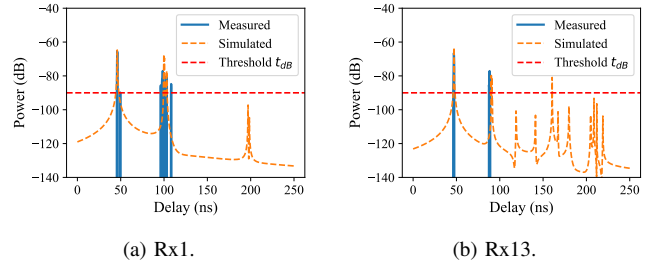


Fig. 5. Paths with up to second order reflections evaluated with the test Rx's after training.

A notable limitation are the holes in the point cloud, which can be seen in the model shown in Fig. 1. These holes can cause some paths to be missed that were present in the measurements. Incomplete scanning combined with noise may result in inaccurate normal vector estimation, which may lead to physically incorrect propagation paths. Similarly, this inaccurate normal estimation may happen near edges. Thus, normal estimation play a critical role in the simulation quality, motivating for further research in this domain.

As a result of the limitations in the segmentation method, the material labels visualized by different colors in the point cloud shown in Fig. 1 demonstrate multiple clusters across surfaces. This issue is particularly critical for NimbusRT, where these labels are used to prune duplicate propagation paths. Even minor noise can cause paths to converge near




Fig. 6. The PDP of Rx13 computed with a threshold t_{dB} of -100 dB, and the evaluation results of simulated PDP trained with measurement data generated with t_{dB} of -90 dB. Suspected second order reflection peaks near the simulation peaks can be observed.

segmentation boundaries, resulting in multiple nearly identical paths. Integrating semantic segmentation methods could enhance path quality and reduce parameter complexity, and thus, improving both material labeling and overall optimization.

Due to the high frequency measurements, the second order reflections are either indistinguishable from the noise, or show as weak peaks with respect to the noise floor. This makes the optimization more difficult for second order reflection paths.

V. CONCLUSION

This work demonstrated that differentiable RT applied to 3D point clouds reconstructed from RGB-D frames can effectively extract channel characteristics from THz measurements conducted at a center frequency of 318 GHz with a 4 GHz bandwidth. In this environment, our method achieves a mean relative rms delay spread error below 5% and a mean absolute received power error of approximately 0.36 dB at novel receiver locations.

Future work will explore advanced algorithms for normal vector estimation, applying semantic segmentation methods for generating material labels, and further validation using channel measurements across wide range of frequencies and in more complex environments.

ACKNOWLEDGMENT

This work was partly supported by the Research Council of Finland (former Academy of Finland) project (Grant no. 348980), 6G Flagship Programme (Grant no. 346208), and the Horizon Europe CONVERGE project (Grant 101094831). Keysight Technologies, Inc. has supported the research with measurement equipment donation.

REFERENCES

- [1] Z. Yun and M. F. Iskander, "Ray tracing for radio propagation modeling: Principles and applications," *IEEE access*, vol. 3, pp. 1089–1100, 2015.
- [2] "ITU-R P.2040," <https://www.itu.int/rec/R-REC-P.2040/en>, accessed: 11.14.2025.
- [3] T. Orekondy, P. Kumar, S. Kadambi, H. Ye, J. Soriaga, and A. Behboodi, "WiNeRT: Towards neural ray tracing for wireless channel modelling and differentiable simulations," in *The 11th International Conference on Learning Representations*, 2023.
- [4] S. Bakirtzis, K. Qiu, J. Zhang, and I. Wassell, "DeepRay: Deep learning meets ray-tracing," in *16th European Conference on Antennas and Propagation (EuCAP)*, 2022, pp. 1–5.
- [5] T. Hehn, M. Peschl, T. Orekondy, A. Behboodi, and J. Brehmer, "Differentiable and learnable wireless simulation with geometric transformers," in *The 13th International Conference on Learning Representations (ICLR)*, 2025.
- [6] G. Cao and Z. Peng, "Raypronet: A neural point field framework for radio propagation modeling in 3d environments," *IEEE Journal on Multiscale and Multiphysics Computational Techniques*, 2024.
- [7] A. Alkhateeb, "DeepMIMO: A generic deep learning dataset for millimeter wave and massive MIMO applications," *arXiv preprint arXiv:1902.06435*, 2019.
- [8] L. Zhang, H. Sun, J. Sun, and R. Q. Hu, "WiSegRT: Dataset for site-specific indoor radio propagation modeling with 3d segmentation and differentiable ray-tracing," in *International Conference on Computing, Networking and Communications (ICNC)*, 2024, pp. 744–748.
- [9] J. Hoydis, F. Ait Aoudia, S. Cammerer, F. Euchner, M. Nimier-David, S. Ten Brink, and A. Keller, "Learning radio environments by differentiable ray tracing," *IEEE Transactions on Machine Learning in Communications and Networking*, 2024.
- [10] S. Jiang, Q. Qu, X. Pan, A. Agrawal, R. Newcombe, and A. Alkhateeb, "Learnable wireless digital twins: Reconstructing electromagnetic field with neural representations," *IEEE Open Journal of the Communications Society*, 2025.
- [11] M. Pang, H. Wang, K. Lin, and H. Lin, "A gpu-based radio wave propagation prediction with progressive processing on point cloud," *IEEE Antennas and Wireless Propagation Letters*, vol. 20, no. 6, pp. 1078–1082, 2021.
- [12] N. Suga, Y. Maeda, and K. Sato, "Indoor radio map construction via ray tracing with rgb-d sensor-based 3d reconstruction: Concept and experiments in wlan systems," *IEEE Access*, vol. 11, pp. 24 863–24 874, 2023.
- [13] G. Xia, C. Zhou, F. Zhang, Z. Cui, C. Liu, H. Ji, X. Zhang, Z. Zhao, and Y. Xiao, "Path loss prediction in urban environments with sionna-rt based on accurate propagation scene models at 2.8 ghz," *IEEE Transactions on Antennas and Propagation*, 2024.
- [14] N. Suga, N. Yoshida, R. Gozono, Y. Maeda, and K. Sato, "Rgb-d sensor-aided radio map estimation using materials classification," *IEEE Wireless Communications Letters*, 2025.
- [15] U. T. Virk, J.-F. Wagen, and K. Haneda, "Simulating specular reflections for point cloud geometrical database of the environment," in *2015 Loughborough Antennas & Propagation Conference (LAPC)*. IEEE, 2015, pp. 1–5.
- [16] N. Vaara, P. Sangi, M. B. López, and J. Heikkilä, "Ray launching-based computation of exact paths with noisy dense point clouds," *IEEE Transactions on Antennas and Propagation*, 2025.
- [17] J. Järveläinen, K. Haneda, and A. Karttunen, "Indoor propagation channel simulations at 60 ghz using point cloud data," *IEEE Transactions on Antennas and Propagation*, vol. 64, no. 10, pp. 4457–4467, 2016.
- [18] P. Koivumäki and K. Haneda, "Point cloud ray-launching simulations of indoor multipath channels at 60 ghz," in *2022 IEEE 33rd Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*. IEEE, 2022, pp. 01–07.
- [19] M. F. De Guzman and K. Haneda, "Analysis of wave-interacting objects in indoor and outdoor environments at 142 ghz," *IEEE Transactions on Antennas and Propagation*, vol. 71, no. 12, pp. 9838–9848, 2023.
- [20] N. Vaara, P. Sangi, M. B. López, and J. Heikkilä, "Differentiable high-performance ray tracing-based simulation of radio propagation with point clouds," *arXiv preprint arXiv:2507.04021*, 2025.
- [21] J. Hoydis, F. Ait Aoudia, S. Cammerer, M. Nimier-David, N. Binder, G. Marcus, and A. Keller, "Sionna rt: Differentiable ray tracing for radio propagation modeling," in *2023 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 2023, pp. 317–321.
- [22] N. Vaara, P. Sangi, J. Pyhtilä, M. Juntti, and J. Heikkilä, "A refined path generation pipeline for radio channel propagation modeling," in *2023 17th European Conference on Antennas and Propagation (EuCAP)*. IEEE, 2023, pp. 1–5.
- [23] T. Määttä, "From rgb-d to diffraction edges: A 3d point cloud pipeline for digital twin and radio propagation modeling," Master's thesis, University of Oulu, 2025. (to be published).