

Towards Generative Ray Path Sampling for Faster Point-to-Point Ray Tracing

Jérôme Eertmans

ICTEAM

Université catholique de Louvain

Louvain-la-Neuve, Belgium

 0000-0002-5579-5360

Nicola Di Cicco

DEIB

Politecnico di Milano

Milano, Italy

 0000-0002-1524-8178

Claude Oestges

ICTEAM

Université catholique de Louvain

Louvain-la-Neuve, Belgium


 0000-0002-0902-4565

Laurent Jacques

ICTEAM

Université catholique de Louvain

Louvain-la-Neuve, Belgium


 0000-0002-6261-0328

Enrico Maria Vitucci

DEI

University of Bologna

Bologna, Italy

 0000-0003-4582-0953

Vittorio Degli-Esposti

DEI

University of Bologna

Bologna, Italy

 0000-0001-6589-8243

Abstract—Radio propagation modeling is essential in telecommunication research, as radio channels result from complex interactions with environmental objects. Recently, Machine Learning has been attracting attention as a potential alternative to computationally demanding tools, like Ray Tracing, which can model these interactions in detail. However, existing Machine Learning approaches often attempt to learn directly specific channel characteristics, such as the coverage map, making them highly specific to the frequency and material properties and unable to fully capture the underlying propagation mechanisms. Hence, Ray Tracing, particularly the Point-to-Point variant, remains popular to accurately identify all possible paths between transmitter and receiver nodes. Still, path identification is computationally intensive because the number of paths to be tested grows exponentially while only a small fraction is valid. In this paper, we propose a Machine Learning-aided Ray Tracing approach to efficiently sample potential ray paths, significantly reducing the computational load while maintaining high accuracy. Our model dynamically learns to prioritize potentially valid paths among all possible paths and scales linearly with scene complexity. Unlike recent alternatives, our approach is invariant with translation, scaling, or rotation of the geometry, and avoids dependency on specific environment characteristics.

Index Terms—Ray Tracing, Generative Sampling, Reinforcement Learning, Machine Learning, Radio Propagation.

I. INTRODUCTION

Over the past decade, Ray Tracing (RT) has emerged as a leading tool in the field of radio propagation modeling, offering a favorable balance between accuracy, ease of implementation, and computational efficiency [1], [2].

In Point-to-Point (P2P) RT, e.g., Image-based RT, the computational complexity is most impacted by the size of the scene or the number of interactions to be taken into account. Indeed, the number of path candidates, i.e., the *potential* ray paths to be considered, grows more than linearly with respect to both parameters.

In recent years, numerous papers have investigated the potential applications of Machine Learning (ML) to the modeling of radio propagation. Motivated by the need for reducing

the computation time, ML has been primarily utilized for two purposes: (1) the estimation of radio materials from measurements data or RT simulations, or scene reconstruction to later simulate radio propagation [3], [4]; and (2) the learning of wave propagation, based on measurements or simulations, in order to create a less expensive model that can make new predictions when the environment varies [5]–[12]. For the second purpose, which is closely related to our application goal, most models try to learn either a specific, usually fixed-size, input scene, or the significantly complex relationship between some input geometries and the electromagnetic (EM) field. Since EM fields usually vary with the frequency, the radio materials, and numerous other parameters, these ML models are constrained to operate within a specific frequency range or on a particular scene. Consequently, they are unable to generalize to other scenarios. Moreover, those models are usually trained to predict a specific channel characteristic, like the coverage map, and they cannot be reused to estimate other parameters. We hypothesize that ML models are not optimal for the complete replacement of channel modeling tools such as RT. This is particularly true given that the calculation of EM fields is usually straightforward once the ray paths are known. Instead, we believe that ML is best suited for the learning of the most complex parts of the RT pipeline, and should be used alongside already existing modeling tools, not as a complete replacement.

In this paper, we propose to address the issue of computational complexity by using an ML model that learns how to identify only the valid ray paths that connect two nodes in a scene, typically a transmitter (TX) and a receiver (RX), and that have a given number of interactions with the environment. A ray path is considered to be valid if (1) it fulfills Fermat’s principle and (2) it is not obstructed by any other object in the scene.

We summarize our major contributions as follows.

- 1) We propose a **methodology** for sampling path candi-

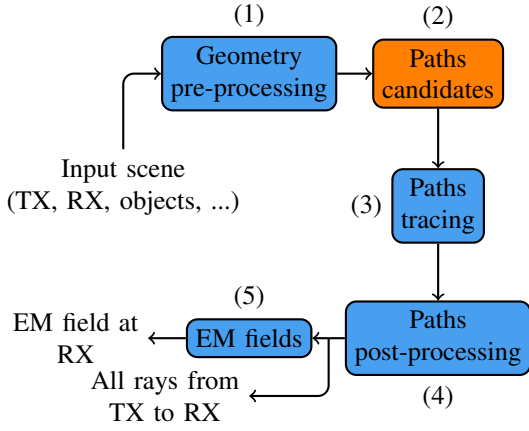


Fig. 1. The pipeline of a typical P2P RT software where relevant steps are numbered from 1 to 5. We propose to replace step 2 with an ML model.

dates based on generative ML, the first in the field to the best of our knowledge;

- 2) We develop a **radio frequency- and material-independent** model that can be trained using reinforcement learning, thus avoiding the need to generate a—potentially computationally expensive—ground truth dataset;
- 3) We implement an ML model that is invariant under the special Euclidean group $SE(3)$ in three dimensions (3D), composed of the rotation ($SO(3)$) and the translation groups, but also to arbitrary scaling, and that can support input scenes of **any size**, which enables generalization to novel scenes unseen at training time.

The remainder of this paper is organized as follows. Sec. II motivates the use of a sampling-based approach to reduce the overall simulation time, and introduces the necessary notations. Then, Sec. III discusses the recent applications of ML to radio propagation. Subsequently, Sec. IV presents our model in a general framework, elucidating its operating principle, and Sec. V demonstrates its efficiency through a 3D RT example and preliminary results. Finally, Sec. VI reiterates the contributions and preliminary outcomes of the model, while also outlining future avenues for research. The source code¹ and a comprehensive tutorial² are also made available to guide readers through the simulation procedure.

II. MOTIVATIONS AND NOTATIONS

In the majority of P2P RT frameworks, the process of estimating the EM fields in a given scene can be summarized as outlined in Fig. 1. It should be noted that the computational complexity of individual steps varies, with paths computation (steps 2 to 4 in Fig. 1) being particularly demanding in the context of large scenes or higher order paths.

Specifically, let the size of the scene, i.e., the number of objects, be N , and the number of interactions be the variable




Fig. 2. Illustration of the many path candidates vs. few valid paths issue in a 2D scene with blue obstacles. Here, only 10 % of all path candidates form valid ray paths (orange solid lines) between TX and RX.

K , then the number of possible path candidates (step 2) is at most N^K . A path candidate is simply an ordered sequence, of length K , indicating the objects with which the ray path interacts. Then, for each of the path candidates, an actual geometric ray path is constructed (step 3), e.g., using the image method or other methods, such as the minimization-based ones [1], [13], [14]. Constructing a ray path from a path candidate usually takes a time proportional to the number of interactions, i.e., $\mathcal{O}(K)$. After that, each ray path is checked (step 4) to see if it is an actual valid path or not. This check is usually bounded by $\mathcal{O}(NK)$ operations, assuming that checks are performed for each ray path segment. Finally, the number of ray paths used to compute EM fields (step 5) is usually orders of magnitude smaller than the number of possible path candidates (see Fig. 2).

In city-scale 3D scenes, N can easily exceed 10^6 (e.g., the number of triangular facets), and N^K thus becomes too large either to fit in memory or to be simulated in a reasonable time. There are three principal countermeasures to this fundamental problem: (i) resort to Ray Launching (RL), a method that launches rays with a given angular separation from a transmitter node, e.g., TX, and collects all rays passing in the vicinity of a receiver node, e.g., RX; (ii) use a heuristic approach to reduce the number of path candidates (e.g., the *Fibonacci* method used in Sionna [15]); or (iii) compute some kind of visibility tree—or visibility matrix—to avoid generating path candidates that eventually lead to invalid ray paths [2], [14]. Solution (iii) is potentially time-efficient; however, its implementation in practice is challenging, since computing the visibility in a 3D scene is itself computationally expensive. Solution (ii) is a hybrid of RL and P2P RT, and may depend on how good the heuristic is.

ML models have already shown very impressive results in generating discrete objects such as graphs [16] and protein

¹Repository: <https://github.com/jeertmans/DiffeRT>.

²Tutorial: <https://differt.rtdfd.io/icmlcn2025/notebooks/sampling-paths.html>.

structures [17]. Following this line of investigation, instead of relying on a heuristic such as solution (ii), our motivation is to construct an ML model that will learn how to sample valid path candidates in the set of all possible path candidates (step 2). In essence, our objective is to learn an ML model that constructs a probability-based visibility tree, similarly to solution (iii).

III. RELATED WORK

ML methods have already been applied to a variety of radio propagation applications, frequently in combination with RT, most of the time to train the model by generating data. For example, ML has been used to predict the propagation of electromagnetic waves in wireless networks [5]–[12], [18]. In many of these cases, the ML model will eventually replace the entire channel modeling tool, thereby eliminating the ability to obtain any intermediate, interpretable outputs, such as those provided by RT in the form of geometrical ray paths. Furthermore, we have identified the following limitations in the existing models: (i) they often necessitate the acquisition of measurements or many simulations for training, (ii) they depend on a specific frequency and the radio materials from the scene, and, most importantly, (iii) they are typically trained on a fixed scene and cannot generalize to arbitrarily-sized inputs.

As previously stated, we posit that ML is ill-suited to predict the final outputs of RT tools, such as coverage maps, as they are often chaotic and depend on an excessive number of parameters. Instead, we propose the creation of an ML-assisted RT pipeline that employs the ML model solely to reduce the overall computational complexity, rather than the deterministic and computationally simple procedures, such as EM coefficients, which can be calculated with certainty once a ray path is known. To the best of our knowledge, the WiNeRT model [12] is the only attempt to address the challenge of path tracing by learning the relationship between incident and outgoing rays for each interaction, e.g., reflections, as well as the EM propagation.

In comparison to existing literature, our model offers the following main novelties:

- 1) To the best of our knowledge, our model is the first to use ML to sample path candidates;
- 2) Because our model only samples path candidates, it is therefore radio-material- and frequency-independent;
- 3) It is capable of accommodating input scenes of any size due to the manner in which the input scene is transformed, as detailed in the subsequent section.
- 4) Furthermore, the model does not require any a priori ground truth training set and can be fully trained using any ray tracer, while completely ignoring the EM fields.

The subsequent section will present the ML model architecture that incorporates all the aforementioned properties.

IV. METHODOLOGY

As previously stated, the objective of our model is to serve as a surrogate for the conventional path candidates generation

process (step 2 in Fig. 1). The objective is thus to learn a generative random function f_w , depending on TX, RX and the list of objects OBJECTS and parametrized by a set of learnable weights w , and such that each call to f_w provides, with good probability, a valid path candidate. Moreover, the learning of w is expected to maximize the probability

$$\mathbb{P}[f_w(\text{TX}, \text{RX}, \text{OBJECTS}) = \text{VALID PATH}]. \quad (1)$$

To generate all paths candidates, we construct a *search tree* (see Fig. 3), in that it is strictly related to the concept of visibility tree in RT [2]. For a given number of interactions K , we can construct each possible path candidate by descending from the root node, also referred to as the initial state, and take one of the possible branches. Each branch corresponds to an object in the scene, identified here with a unique letter, and a path candidate is constructed by identifying a list of objects to visit. As a ray path cannot usually interact with the same object twice in a row, those branches should be **unreachable** (dotted lines). The question mark “?” is used as an indicator of a path candidate being incomplete.

Our approach leverages the principles of Generative Flow Networks (GFlowNets) [19] ML for learning to efficiently traverse the search tree. Instead of relying on exhaustive search, we model the process of generating a path candidate, \mathcal{P} , as *flowing* in a Directed Acyclic Graph (DAG), starting from an initial state, and ending at a terminal state, the latter representing a complete path candidate. Each child state s' has one unique parent state s , e.g., state “AB” is a child state of “A?”.

Each terminal state, i.e., a path candidate, is associated with a scalar reward R . Intuitively, greater reward values represent “better” paths. An example reward function could return 1 if the path candidate ends up generating a valid ray path, and 0 otherwise. We will discuss later how we can define a better reward function. The central property of GFlowNets is that, after successful training, the model is expected to sample terminal states, i.e., complete path candidates, with a probability p that is proportional to their corresponding reward, R . That is,

$$p(\mathcal{P}) \propto R(\mathcal{P}). \quad (2)$$

Given the example reward function, GFlowNets will learn to sample from the distribution of valid paths only.

To achieve this, the model must be trained to respect the following fundamental properties:

- 1) Each edge in the search graph must be assigned a positive *flow*, $F(s, s') > 0$, where s is the parent state and s' is the child state;
- 2) Flow conservation between ingoing and outgoing edges must be ensured:

$$\forall s', F(s, s') = R(s') + \sum_{s''} F(s', s''), \quad (3)$$

that is, the sum of output flows, $F(s', s'')$, must be equal to the input flow, $F(s, s')$, minus the reward;

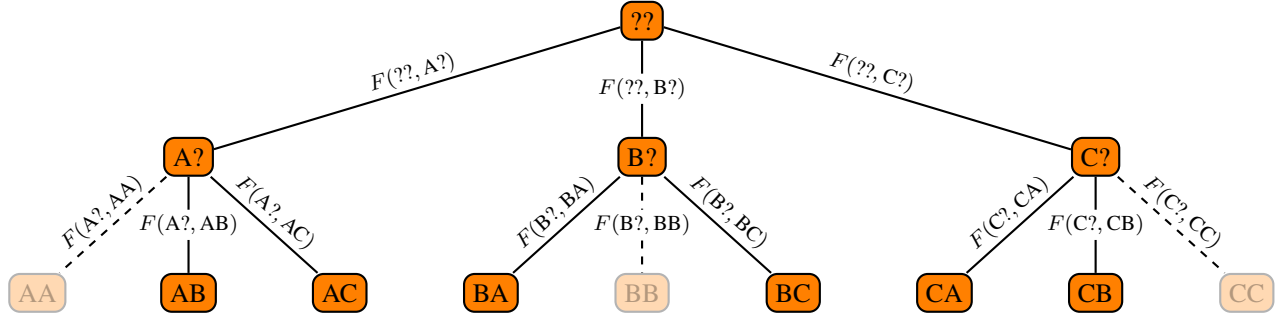


Fig. 3. Flowchart of all possible states for $K = 2$ path candidates in a scene with $N = 3$ objects: A, B, and C. Each state s can flow to three possible child states s' , to select interacting objects one at a time. As interacting with the same object twice in a row is **physically unsound**, this state is marked as **unreachable** (faded nodes and dotted lines). To account for that in the model, the flow is stopped, i.e., set to zero, to prevent reaching those states.




Fig. 4. Simplified representation of the proposed ML model replacing the second step in Fig. 1, with trainable weights as colored trapezoids. Each state, one-hot-encoded ($\mathbb{1}$) from the current path candidate (size 0 to K), records visited objects, starting from a zero matrix. A recurrent cell maps it to the feature space for positional encoding. Dashed lines indicate state updates, and randomness comes from using a different key per path candidate.

- 3) The probability of choosing state s' given state s must be defined as

$$p(s'|s) = \frac{F(s, s')}{\sum_{s''} F(s, s'')}, \quad (4)$$

that is, the probability of traversing an edge in the search graph is equal to its flow value normalized over all outgoing edges.

It should be noted that all non-terminal states are assigned a zero reward value. However, this is not a limitation of the model itself, and the potential benefits of integrating partial rewards in the learning process will be discussed in greater details later on.

As illustrated in Fig. 4, the proposed model is divided into two parts: the left part, depicted in green, is responsible for transforming the input scene into feature vectors, and the second part, represented in orange, generates the flow from a specified state and scene to all potential objects. The subsequent section will provide a more detailed analysis of the former part.

A. Extracting Feature Vectors From a Scene

The initial stage of the model is designed to extract feature vectors from the scene, including the positions of the TX and RX elements, as well as an arbitrary list of object coordinates. In Fig. 4, we assume 3D coordinates; however, our model is also applicable to any other dimension, such as 2D coordinates. The sole restriction on the input scene is that each object must be represented by a fixed number of vertices, denoted by V . To accommodate scenes of arbitrary size and order,

our model is based on the Deep Sets architecture [20]. To ensure invariance with respect to both translation and scaling, a standard normalization of the input scene coordinates is performed. Finally, the model itself is built on the $E3x$ ML library [21], which in turn utilizes two Multi Layer Perceptrons (MLPs) to convert XYZ coordinates into vectors of d features, thereby guaranteeing invariance with respect to rotation.

B. Flow and Loss Function

The second part of our model generates, for a given state s , an array of flows to each object. Then, the next state s' is chosen based on those flows. The diversity in the generated path candidates is provided by the random sampling performed between each state, where each possible output state s' is weighted according to its flow $F(s, s')$, see (4).

In practice, the flow function as detailed previously also depends on two other parameters, the scene feature vector and the object features, but were omitted for readability.

For training our model, we minimize the GFlowNets loss function, which rewrites (3) as a mean squared error:

$$L(s') = \left(F(s, s') - R(s') - \sum_{s''} F(s', s'') \right)^2. \quad (5)$$

We summarize the training procedure as follows:

- 1) Generate, for a given random scene, a batch of path candidates;
- 2) Evaluate the GFlowNets loss (5) for a batch of path candidates, and accumulate the results;




Fig. 5. Street canyon scene from Sionna [22].

- 3) Perform a gradient-based update of the ML model, minimizing the GFlowNets loss;

and repeat those steps for a large number of scenes.

Upon successful training, as a consequence of the reward function, the model is expected to sample all valid path candidates with the same probability, with a computational cost scaling linearly with the number of objects, reflections, and model size: $\mathcal{O}(cPNK)$, where c is a constant for model calls and P the number of sampled path candidates. For large scenes, this is significantly cheaper than the exhaustive approach, having a cost of $\mathcal{O}(N^K)$.

C. Accuracy and Hit Rate

As our loss function may not directly reflect the model performance, we introduce two metrics to better evaluate its effectiveness: accuracy and hit rate. The **accuracy** is defined as the ratio of valid paths to the number of path candidates generated, approaching the probability of sampling a valid path defined in (1). An accuracy of 100% indicates that the model exclusively samples path candidates that ultimately result in valid ray paths. The **hit rate** is defined as the number of valid generated ray paths divided by the total number of valid ray paths that could be generated. A hit rate of 100% means that the model samples all potential solutions.

V. APPLICATION TO A 3D SCENARIO

In this section, we present an application of our model to sampling path candidates in the context of 3D P2P RT in an urban street canyon, with first and second order specular-reflection paths only. The complete procedure is available in the aforementioned tutorial notebook.

A. Training Set

We generate our training set by randomly sampling modified versions of the “*simple street canyon*” scene from Sionna [15] (see Fig. 5):

- TX and RX are randomly positioned in the main street canyon;
- A random number of triangle facets is removed.

In each scene configuration, a number of path candidates are generated, and the corresponding ray paths are constructed. The validity of each generated path is determined using a standard ray-obstruction check, and a gradient-based update is performed to minimize the loss function (5).

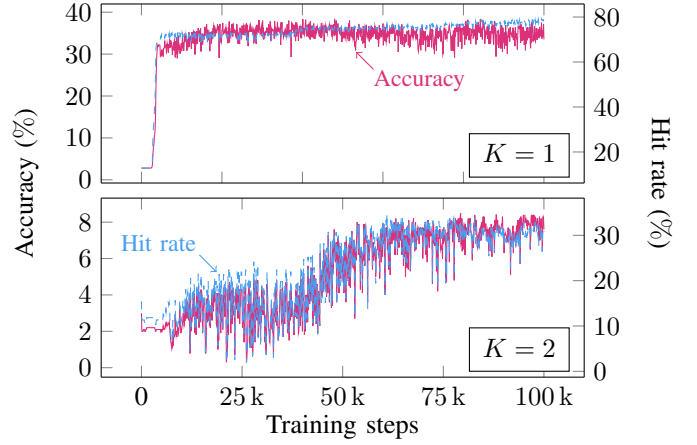


Fig. 6. Evolution of the accuracy and the hit rate throughout training steps.

For validation, we generate 100 new scenes that we will use to measure the evolution of the accuracy and the hit rate of the model throughout the learning steps.

B. Simulation Parameters

The model was trained by performing 100×10^3 steps with the Adam optimizer [23] (learning rate of 3×10^{-5}). For each step, the loss was accumulated by generating 100 path candidates. The number of features was set to $d = 100$. The first two MLPs have 3 linear layers and a hidden size matching the number of features. The last MLP has 3 linear layers with a hidden size of 500. Finally, the model’s accuracy and hit rate were evaluated by sampling 10 path candidates on each scene of the validation set.

C. Results

Fig. 6 compares the evolution of the accuracy and the hit rate throughout the training steps, for single-bounce reflection ($K = 1$) and double-bounce reflection ($K = 2$). First, we observe a short plateau in the learning curves for the first learning steps of the single-bounce model, where it starts with the same accuracy as a random sampler (3%). The double-bounce model is initialized from the single-bounce model, after training, which explains a higher starting accuracy than a random sampler (0.03%). Finally, a clear relationship can be observed between the hit rate and the accuracy, which shows that the model learns to generate different valid paths, with an increasing proportion of sampled paths being valid.

D. Discussion

The results are encouraging, particularly given the compact model size, when compared with random sampling performance. For $K = 1$, both accuracy and hit rate steadily improve, indicating that the model effectively learns to prioritize valid paths. With an average hit rate of nearly 80%, the model identifies more than three-quarters of all possible valid paths by sampling only 10 candidates out of an average of 73. Although promising, this performance remains insufficient to

fully replace exhaustive search, requiring further refinements for real-world applications.

For $K = 2$, the model initially surpasses random sampling but shows only marginal improvements during training before stabilizing. Training from scratch without initializing from $K = 1$ leads to collapse, with accuracy and hit rate dropping to zero, suggesting misalignment between the loss function and model performance. Higher K values further exacerbate this issue due to sparse rewards from the low proportion of valid paths. Alternative reward functions, such as those inversely proportional to path length, have accelerated convergence but typically yield lower final performance.

The current ML model consists of three small MLPs, enabling rapid sampling but potentially limiting scalability to complex scenes. Increasing layer depth could improve performance in larger environments, while expanding feature vector size has already shown benefits, albeit at the cost of slower training.

VI. CONCLUSION

This work presents a novel ML model that learns how to sample paths. When properly trained, it could potentially remove the exponential complexity behind the usual P2P RT toolboxes. The model learns how to prioritize some branches in the complex search tree that represents the set of all possible path candidates. Our results demonstrate that, particularly in the single-bounce case ($K = 1$), the model effectively learns to generate valid paths across diverse scenes. However, for more complex multi-bounce scenarios ($K \geq 2$), our preliminary results indicate that additional improvements are necessary.

Future work will explore several key areas to enhance the model's performance. First, addressing sparse reward functions, potentially through reward smoothing or continuous reward formulations [24]–[26], may aid in learning across more complex scenes. Additionally, scaling the model architecture by increasing layer size and depth could improve handling of intricate environments. Investigating alternative training approaches, such as pre-training on simpler scenes before scaling up, may also prove beneficial. Ultimately, our goal is to deploy this model in larger, more complex 3D urban scenes to substantively reduce computational demands, particularly within RT tools like Sionna or our own differentiable Ray Tracer, DiffeRT³, where complexity reduction could yield substantial computational savings.

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] F. Fuschini, E. M. Vitucci, M. Barbiroli, G. Falciasecca, and V. Degli-Esposti, "Ray tracing propagation modeling for future small-cell and indoor applications: A review of current techniques," *Radio Science*, vol. 50, no. 6, pp. 469–485, 2015.
- [3] F. Zhang, C. Zhou, C. Brennan, R. Wang, Y. Li, G. Xia, Z. Zhao, and Y. Xiao, "A radio wave propagation modeling method based on high-precision 3-d mapping in urban scenarios," *IEEE Transactions on Antennas and Propagation*, vol. 72, no. 3, pp. 2712–2722, 2024.
- [4] J. Hoydis, F. A. Aoudia, S. Cammerer, F. Euchner, M. Nimier-David, S. ten Brink, and A. Keller, "Learning radio environments by differentiable ray tracing," 2023.
- [5] K. Ahmad and S. Hussain, "Machine learning approaches for radio propagation modeling in urban vehicular channels," *IEEE Access*, vol. 10, pp. 113 690–113 698, 2022.
- [6] S. K. Vankayala, S. Kumar, I. Roy, D. Thirumulanathan, S. Yoon, and I. S. Kanakaraj, "Radio map estimation using a generative adversarial network and related business aspects," in *2021 24th International Symposium on Wireless Personal Multimedia Communications (WPMC)*, 2021, pp. 1–6.
- [7] C. Ding and I. W.-H. Ho, "Digital-twin-enabled city-model-aware deep learning for dynamic channel estimation in urban vehicular environments," *IEEE Transactions on Green Communications and Networking*, vol. 6, no. 3, pp. 1604–1612, 2022.
- [8] S. Bakirtzis, K. Qiu, J. Zhang, and I. Wassell, "Deepray: Deep learning meets ray-tracing," in *2022 16th European Conference on Antennas and Propagation (EuCAP)*, 2022, pp. 1–5.
- [9] M. Yin, Y. Hu, T. Azzino, S. Kang, M. Mezzavilla, and S. Rangan, "Wireless channel prediction in partially observed environments," 2022.
- [10] K. Mao, Q. Zhu, M. Song, H. Li, B. Ning, G. F. Pedersen, and W. Fan, "Machine-learning-based 3-d channel modeling for u2v mmwave communications," *IEEE Internet of Things Journal*, vol. 9, no. 18, pp. 17 592–17 607, 2022.
- [11] Q. Zhu, F. Bai, M. Pang, J. Li, W. Zhong, X. Chen, and K. Mao, "Geometry-based stochastic line-of-sight probability model for a2g channels under urban scenarios," *IEEE Transactions on Antennas and Propagation*, vol. 70, no. 7, pp. 5784–5794, 2022.
- [12] 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 Eleventh International Conference on Learning Representations*, 2023.
- [13] F. Puggelli, G. Carluccio, and M. Albani, "A novel ray tracing algorithm for scenarios comprising pre-ordered multiple planar reflectors, straight wedges, and vertexes," *IEEE Transactions on Antennas and Propagation*, vol. 62, no. 8, pp. 4336–4341, 2014.
- [14] J. Eertmans, C. Oestges, and L. Jacques, "Min-path-tracing: A diffraction aware alternative to image method in ray tracing," in *2023 17th European Conference on Antennas and Propagation (EuCAP)*, 2023, pp. 1–5.
- [15] J. Hoydis, F. A. Aoudia, S. Cammerer, M. Nimier-David, N. Binder, G. Marcus, and A. Keller, "Sionna rt: Differentiable ray tracing for radio propagation modeling," 2023.
- [16] Y. Zhu, Y. Du, Y. Wang, Y. Xu, J. Zhang, Q. Liu, and S. Wu, "A survey on deep graph generation: Methods and applications," in *The First Learning on Graphs Conference*, 2022.
- [17] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko *et al.*, "Highly accurate protein structure prediction with alphafold," *Nature*, vol. 596, no. 7873, pp. 583–589, 2021.
- [18] T. Hehn, M. Peschl, T. Orekondy, A. Behboodi, and J. Brehmer, "Differentiable and learnable wireless simulation with geometric transformers," 2024.
- [19] Y. Bengio, S. Lahlou, T. Deleu, E. J. Hu, M. Tiwari, and E. Bengio, "Gflownet foundations," 2023.
- [20] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Poczos, R. Salakhutdinov, and A. Smola, "Deep sets," 2018.
- [21] O. T. Unke and H. Maennel, "E3x: E(3)-equivariant deep learning made easy," *arXiv preprint arXiv:2401.07595*, 2024.
- [22] J. Hoydis, S. Cammerer, F. Ait Aoudia, A. Vem, N. Binder, G. Marcus, and A. Keller, "Sionna: An open-source library for next-generation physical layer research," *arXiv preprint*, Mar. 2022.
- [23] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2017.
- [24] J. Eertmans, L. Jacques, and C. Oestges, "Fully differentiable ray tracing via discontinuity smoothing for radio network optimization," in *2024 18th European Conference on Antennas and Propagation (EuCAP)*, 2024, pp. 1–5.
- [25] L. Pan, D. Zhang, A. Courville, L. Huang, and Y. Bengio, "Generative augmented flow networks," 2022.
- [26] X. Ji, X. Zhang, W. Xi, H. Wang, O. Gadyatskaya, and Y. Li, "Meta generative flow networks with personalization for task-specific adaptation," 2023.

³Repository: <https://github.com/jeertmans/DiffeRT>.