

UCLA

UCLA Electronic Theses and Dissertations

Title

AI-Assisted Integrated Sensing and Communication for NextG Wireless Networks

Permalink

<https://escholarship.org/uc/item/7cc2m9rk>

ISBN

9798186171263

Author

Lu, Haofan

Publication Date

2026-07-30

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

AI-Assisted Integrated Sensing and Communication
for NextG Wireless Networks

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Computer Science

by

Haofan Lu

2026

© Copyright by

Haofan Lu

2026

ABSTRACT OF THE DISSERTATION

AI-Assisted Integrated Sensing and Communication for NextG Wireless Networks

by

Haofan Lu

Doctor of Philosophy in Computer Science

University of California, Los Angeles, 2026

Professor Omid Abari, Chair

Wireless technologies have been evolving significantly over the past few decades, forming the backbone of many modern applications. Two key paradigms have emerged for different purposes: communication and sensing. Wireless communication systems, such as WiFi and 5G, are optimized to encode as much information as possible into the limited time-frequency resources for high-throughput, low-latency data transmission. Wireless sensing systems, such as radar, focus on extracting information about the environment, such as the location and movement of objects, by analyzing the wireless signals reflected off them. These two paradigms have been evolving separately, developing different hardware and software stacks.

Next-generation wireless systems seek to integrate sensing and communication into a unified system to enhance performance and enable new applications. Many challenges exist in achieving this integration. One key challenge is designing unified radio software and hardware that can support both sensing and communication functionalities. Another challenge is to design synergistic mechanisms that create mutual benefits between sensing and communication. The advancement of machine learning and artificial intelligence has opened

up new opportunities for addressing these challenges.

This dissertation introduces three systems towards this objective. First, we present MilBack, a unified design of a millimeter-wave sensing and communication system based on the backscatter principle. MilBack enables high-accuracy target localization and high-throughput, two-way communication with ultra-low power consumption. Next, we present two systems, NeWRF and mmDiff, that materialize the notion of sensing the wireless environment for improving communication. NeWRF is a deep learning framework that reconstructs the wireless radiation field from sparse channel measurements and predicts the channel state information at unseen locations. It enables operators to identify coverage dead zones and optimize access point placement. mmDiff is a differentiable ray-tracing simulation framework that learns the radio material properties of the environment from sparse measurements. It enables robust mmWave beam prediction in complex 3D environments, which is critical for high-throughput, low-latency communication.

The dissertation of Haofan Lu is approved.

Sriram Narasimhan

Tony Nowatzki

Baharan Mirzasoleiman

Omid Abari, Committee Chair

University of California, Los Angeles

2026

*To my parents ...
who have been my strongest support.*

TABLE OF CONTENTS

1	Introduction	1
1.1	Unified Design of Sensing and Communication	2
1.2	Synergistic Mechanisms for Sensing and Communication	4
1.3	Organization of the Dissertation	6
2	MilBack: A Millimeter Wave Backscatter Network for Two-Way Commu- nication and Localization	7
2.1	Preliminaries	10
2.2	Related Work	12
2.3	MilBack Overview	13
2.4	MilBack’s Node Design	14
2.5	Localization and Orientation	17
2.5.1	Location Detection	17
2.5.2	Orientation Detection	18
2.6	Communication	19
2.6.1	Downlink Communication	19
2.6.2	OAQFM	21
2.6.3	Uplink Communication	22
2.7	Two-Way Communication and Localization Protocol	24
2.8	Implementation	26
2.9	Evaluation	27
2.9.1	Micro-benchmarks	28

2.9.2	Localization Performance	29
2.9.3	Orientation Performance	30
2.9.4	Downlink Performance	32
2.9.5	Uplink Performance	33
2.9.6	Power Consumption	35
2.10	Conclusion	36

3 NeWRF: A Deep Learning Framework for Wireless Radiation Field Reconstruction and Channel Prediction 37

3.1	Preliminaries	40
3.2	Related Work	43
3.3	NeWRF	45
3.3.1	Problem Description	45
3.3.2	Neural Wireless Channel Synthesis	46
3.3.3	Improving the Convergence of NeWRF	48
3.3.4	Optimizing NeWRF	49
3.3.5	NeWRF’s Scene Representation	50
3.3.6	Predicting Channels at Unvisited Locations	52
3.4	Evaluation	54
3.4.1	Channel Prediction Results	54
3.4.2	Measurement Density of Training Set	56
3.4.3	Impact of Materials Types and Signal Frequencies	57
3.4.4	Ray-searching Algorithm Performance	57
3.5	Limitations and Future Work	58

3.6	Conclusion	59
	Appendix	59
4	mmDiff: A Noise-Robust Differentiable Ray-Tracing Framework for mmWave Scene Calibration and Channel Prediction	64
4.1	Preliminaries	67
4.1.1	Angular Power Spectra	67
4.1.2	Physically-based rendering (PBR)	67
4.1.3	Physics of Wireless Ray-Tracing	68
4.2	Related Work	70
4.2.1	Wireless Channel Prediction	70
4.3	mmDiff Design	71
4.3.1	mmDiff’s Ray-Tracing Algorithm	72
4.3.2	Noise-Robust Ray-Surface Interaction	73
4.3.3	Optimization Framework	77
4.3.4	Implementation	77
4.4	Results	78
4.4.1	Real-world Experiments	78
4.4.2	Simulation Experiments	80
4.5	Discussion	84
4.6	Conclusion	85
	Appendix	85
5	Summary and Future Work	102
5.1	Summary	102

5.2	Future Work	103
	References	105

LIST OF FIGURES

2.1	Frequency Scanning Antenna (FSA).	10
2.2	FMCW Radar Signal.	11
2.3	Beam Pattern of MilBack’s Dual-Port FSA.	15
2.4	Block Diagram of MilBack’s Backscatter Node.	16
2.5	Orientation Detection at the Node.	20
2.6	OAQFM Modulation.	22
2.7	Block Diagram of MilBack’s AP.	24
2.8	MilBack Packet Structure.	24
2.9	MilBack’s Node Implementation.	27
2.10	Dual-port FSA Beam Pattern.	28
2.11	OAQFM Micro-benchmark.	29
2.12	Localization Performance.	30
2.13	Orientation Estimation Performance.	31
2.14	Downlink Performance.	33
2.15	Uplink Performance.	34
3.1	Wireless Field Tomography in a Simple Environment.	38
3.2	Wireless Communication Model.	41
3.3	Example Environment and Transmitter/Receiver Locations.	46
3.4	Neural Channel Synthesis Algorithm.	47
3.5	An Optimized Wireless Scene.	50
3.6	The Nature of Wireless Scenes.	51

3.7	3D Environment Models for Simulation.	54
3.8	Channel Prediction Results.	55
3.9	Ray-Searching Algorithm Performance.	58
3.10	NeWRF's Training Pipeline	60
3.11	NeWRF's MLP Architecture. Both Coarse and Fine model use the same architecture.	60
3.12	NeWRF's performance on adding noise to the channel	61
3.13	Subcarrier Channel Prediction.	61
4.1	AoA Power Spectra.	67
4.2	mmDiff's Learning Framework for Scene Calibration.	68
4.3	Issue with Pure Specular Reflection.	74
4.4	Single-reflector scene for the directional scattering path gain convergence test. .	75
4.5	Directional-scattering path gain converges to specular reflection path gain as the scattering footprint grows and directionality increases.	75
4.6	Real-World Testbed Setup.	79
4.7	Results of the Real-World Experiment.	79
4.8	Replica Scenes Used in Evaluation.	80
4.9	Robustness to Geometry Noise.	81
4.10	Geometry setting for proof of Theorem 1.	86
4.11	Illustration of the effective area dS	86
4.12	Factors that impact the deviation of the directional scattering approximation from the specular reflection model.	89
4.13	Single planar reflector scene for deviation analysis.	90

4.14	Large incident angles cause a higher normalization factor, which amplifies the scattering lobe.	90
4.15	Illustration of the sampling issue that causes the variation with respect to the incident angle (bird’s-eye view).	92
4.16	Maximum deviation for different distances between TX/RX and reflection point.	92
4.17	Additional Replica scenes used for evaluation: Office 1, Office 2, and Office 3. .	94
4.18	Sample poses in the TRAJ and RAND datasets (subsampled 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 that this is the 2D projection for visualization; the actual positions and orientations are in 3D.	95
4.19	mmDiff’s differentiability enables geometry calibration.	96
4.20	Illustration of the multipath interference	100
4.21	Multipath can cause inconsistent AoA spectra due to the loss of phase information	100
4.22	Material calibration with single dominant path.	101

LIST OF TABLES

2.1	Comparison with State-of-the-Art mmWave Backscatter Systems.	35
4.1	Robustness under Geometry Noise (TRAJ Dataset).	81
4.2	Prediction Accuracy on RAND Dataset.	83
4.3	Scene Statistics	96
4.4	Robustness under Geometry Noise for Additional Scenes (TRAJ Dataset). . . .	98
4.5	Prediction Accuracy Results for Additional Scenes (RAND Dataset).	99

ACKNOWLEDGMENTS

First and foremost, I would like to express my sincere gratitude to my advisor, Omid Abari, for his invaluable guidance and support. His keen insights and inspiring vision were instrumental in shaping my academic and personal growth. His support and openness to new ideas encouraged me to explore and innovate, and grow as a researcher and a person. I am fortunate to have had the opportunity to work with such a brilliant and supportive mentor, and I am deeply grateful for his guidance throughout my PhD journey.

I would also like to thank my committee members, Baharan Mirzasoleiman, Tony Nowatzki, and Sriram Narasimhan, for their valuable feedback and support. Their expertise and insights greatly enriched my work and helped me refine my ideas.

Many thanks to my fellow lab members: Mohammad Hossein Mazaheri, Reza Rezvani, Tianxiang Li, Ali Abedi, Haochen Zhao, Jian-Ting Ko, Christopher Vrattheuer, Shanmu Wang. I fondly remember the countless hours of brainstorming, discussion, and learning we shared. I learned so much from each of them, and I am grateful for the friendships we built along the way.

I would also like to thank my collaborators and mentors during my internships: Hao Chen, Vishnu Vardhan Ratnam, Lianjun Li, Pranav Madadi, Mohamed Ibrahim, Shivang Aggarwal, Puneet Sharma, Saadallah Kassir, Lola Awoniyi-Oteri, Miguel Griot, Peerapol Tinnakornsriruphap, Gavin Horn, Navid Abedini, Myungjin Lee, Humberto La Roche, and many others. Their invaluable insights, feedback, and encouragement contributed greatly to my professional and research development.

Special thanks to my friends, Yadi Cao, Wenxin Song, Huan Xu, Xiaojie Zhou, Keli Huang, Wanghao Yi, and many others, for sharing so many wonderful moments with me and for their encouragement during the difficult times. My life at UCLA could not have been as fulfilling without their support.

Finally, I would like to express my deepest gratitude to my family, especially my parents.

My mother, Jianmin Fan, you have been my greatest source of support and inspiration throughout my life. Your unwavering love, encouragement, and sacrifices have been the foundation of my success, and I am forever grateful for your belief in me. My father, Jing Lu, you have been a pillar of strength, providing the steady encouragement and guidance that helped me navigate the challenges of my academic journey. I am deeply thankful for your love, which has shaped who I am today and made this milestone possible.

VITA

- 2017–2021 B.Eng., Electrical Engineering and Automation, Zhejiang University, Hangzhou, China
- 2017–2021 B.S., Electrical Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA
- 2023 Summer Research Intern, Samsung Research America, Inc., Plano, TX, USA
- 2024 Summer Research Associate Intern, Hewlett Packard Enterprise, Inc., Milpitas, CA, USA
- 2025 Summer Interim Engineering Intern, Qualcomm Technology, Inc., San Diego, CA, USA

PUBLICATIONS

Haofan Lu, Yadi Cao, Wanghao Yi, Omid Abari. "mmDiff: A Noise-Robust Differentiable Ray-Tracing Framework for mmWave Scene Calibration and Channel Prediction", CVPR Findings 2026.

Haofan Lu, Christopher Vattheuer, Baharan Mirzasoleiman, Omid Abari "NeWRF: A Deep Learning Framework for Wireless Radiation Field Reconstruction and Channel Prediction", ICML 2024.

Haofan Lu, Mohammad Hossein Mazaheri, Reza Rezvani, Omid Abari, "A Millimeter Wave Backscatter Network for Joint Communication and Localization", SIGCOMM 2023.

Haofan Lu^{*}, Tianxiang Li^{*}, Reza Rezvani, Ali Abedi, Omid Abari, "Bringing WiFi Localization to Any WiFi Devices", HotNets 2022.

Tianxiang Li, Mohammad H. Mazaheri, Kalaivani Kamalakannan, **Haofan Lu**, Omid Abari "Can IoT Devices be Powered up by Future Indoor Wireless Networks?", HotMobile 2024.

Ali Abedi, **Haofan Lu**, Alex Chen, Charlie Liu, Omid Abari, "WiFi Physical Layer Stays Awake and Responds When Should Not", IEEE Internet of Things Journal.

Jian-Ting Ko, **Haofan Lu**, Myungjin Lee, Humberto La Roche, Omid Abari, "Bi-static mmWave Backscatter Communication and Localization". In submission.

Tianxiang Li, **Haofan Lu**, Reza Rezvani, Mohammad Hossein Mazaheri, Omid Abari. "Enhancing WiFi Protocols via Smarter Antenna Design". In submission.

CHAPTER 1

Introduction

The past few decades have witnessed rapid progress in wireless technology, giving rise to new applications such as the Internet of Things (IoT) [MRG23], unmanned aerial vehicles (UAVs) [LMA22, ZHM25], and virtual/augmented reality (VR/AR) [ABD17, EPB18]. These wireless technologies fall into two main categories: wireless communication and wireless sensing. Wireless communication serves as the pipe for data transmission, optimized for high data rate, low latency, and high reliability. Wireless sensing instead extracts information about surrounding humans and objects, enabling diverse applications such as indoor localization, occupancy detection, and device-free healthcare [XZY16, WLM25, ALC24, INO24]. Recent years have seen growing interest in integrating sensing into next-generation wireless communication systems, with two main goals:

- *Unified design*: building compact hardware and software systems that support both communication and sensing [LCM22, CLJ21]. The two functions share similar RF frontend components, but their backend signal processing algorithms have diverged substantially. Communication systems rely on complex modulation and demodulation to achieve high data rates, while sensing systems use range-Doppler analysis to extract target location and motion. A unified design consolidates these components, reducing hardware redundancy and improving resource efficiency, cost, size, and power consumption.
- *Synergistic mechanisms*: developing methods that let sensing and communication enhance each other's performance [CLJ21]. Sensing can inform communication, for ex-

ample through beamforming, channel prediction, and interference mitigation. Communication can in turn support sensing by coordinating multiple devices, sharing sensing data, and enabling collaborative sensing. Exploiting this complementarity between the two domains can improve overall system performance beyond what either could achieve alone.

Past research in integrated sensing and communication has taken two main directions: designing novel waveforms and signal processing algorithms that jointly perform communication and sensing [XYZ22, MTM24], or extending existing orthogonal frequency division multiplexing (OFDM)-based communication systems with sensing capabilities [AUC24]. Both directions focus mainly on signal processing and often overlook hardware design and physical layer considerations, which are crucial for a truly unified design. Few works have explored synergistic mechanisms that use sensing information to enhance communication performance. This dissertation tackles both aspects through a holistic system design, covering electromagnetic simulation, antenna and circuit design, signal processing algorithms, and machine learning techniques from the bottom up. The following sections outline the contributions of this dissertation.

1.1 Unified Design of Sensing and Communication

Sensing and communication are intrinsically different tasks from the perspective of the information source. A communication system’s information comes from an internal data source, so it needs a waveform as diverse as possible to carry more information. A sensing system’s information instead comes from the external environment, which acts as a modulator that imprints its information onto the transmitted signal. A sensing system therefore needs the opposite: a waveform as simple as possible, serving as a carrier that maximizes the information extracted from the environment. This fundamental difference in information source has driven communication and sensing waveforms apart, making a unified system that supports

both functions challenging to design.

To tackle this challenge, we introduce a new design approach that leverages the similarities between backscatter communication and radar sensing. Backscatter communication has a tag device that modulates and reflects an incident signal from a reader device to transmit information. Like radar sensing, its information source is external, so the reader waveform can stay simple while the tag modulates its information onto the incident signal. This similarity gives us common ground to design a unified system that supports both communication and sensing simultaneously.

Contribution: a unified design of mmWave backscatter communication and sensing system. Traditional backscatter communication systems, such as RFID and WiFi backscatter [BR14, ADM20], often suffer from low data rates that limit their applicability. mmWave backscatter communication has been proposed to address this limitation [MCA21, SPB21], since the much wider bandwidth available at mmWave frequencies enables significantly higher data rates. mmWave frequencies also support frequency-modulated continuous-wave (FMCW) radar, which is widely used for sensing [KHY25] and enables high-resolution localization and imaging. In this dissertation, we combine mmWave backscatter communication and FMCW radar sensing into a system named *MilBack* [LMR23], which enables orientation sensing, localization, and two-way high-data-rate communication.

This system addresses three main challenges. First, mmWave signals suffer from high path loss, which limits communication range. Beamforming is a common solution, but it requires sophisticated phased arrays and beam training algorithms that are too power-hungry for low-power IoT devices. Instead, we build on Frequency-Scanning Antenna (FSA) technology, which steers the beam direction passively by changing the transmitted signal’s frequency, and design a retro-reflective FSA-based structure that enables the tag to create a high-gain beam always pointing at the reader without any active components. Second, traditional backscatter tags typically avoid high-order modulation schemes due to their architectural complexity and power cost. We introduce a new second-order modulation

scheme, Orientation-Assisted Quadrature Frequency Modulation (OAQFM), which exploits the symmetry of FSA to double the data rate without adding significant complexity to the tag. Third, we integrate this tag design with a conventional FMCW radar waveform for orientation finding and localization, and design a communication protocol that coherently combines radar sensing and backscatter communication into a holistic solution. We built a full-fledged prototype with customized FSA antennas and off-the-shelf components. Empirical evaluation shows that MilBack achieves centimeter-level localization accuracy, with uplink and downlink data rates up to 40 Mbps and 36 Mbps at only 32 mW and 18 mW power consumption.

1.2 Synergistic Mechanisms for Sensing and Communication

Next-generation wireless communication systems are expected to operate in complex, dynamic environments, where signals can be blocked, reflected, or scattered by surrounding objects. In these scenarios, environmental awareness becomes crucial for optimizing communication performance, and sensing can provide the information needed to achieve it. While this high-level idea is intriguing, the concrete synergistic mechanisms remain an open question. What information should be sensed, and how? Which communication performance metrics can sensing improve, and how should the sensed information be used to improve them? These questions remain largely unexplored, and this dissertation aims to answer some of them.

We focus on the wireless environment itself as the target of sensing. Understanding how wireless signals propagate through the environment helps identify dead zones, optimize infrastructure deployment, and improve communication performance. To achieve this, we build on recent advances in computer vision and graphics, which have shown remarkable success in learning a scene’s 3D structure from sparse 2D images. We adapt these techniques to the wireless domain by grounding them in the physics of wireless signal propagation,

introducing two approaches for different scenarios. When the environment’s geometry is unavailable, we propose a data-driven approach that learns a neural representation of the wireless radiation field from sparse channel measurements. When the geometry is available, we propose a hybrid approach that combines physics-grounded wireless ray-tracing with differentiable simulation to learn the radio material properties in the environment from sparse angular power spectrum measurements.

Contribution: a neural representation of wireless radiation field. Wireless signals are invisible, and their behavior in the environment is complex and difficult to model, making it challenging to predict wireless coverage and channel quality. Traditional approaches rely on site surveys and drive tests, where engineers collect measurements at thousands or millions of locations to build a coverage map. This process is time-consuming, labor-intensive, and often impractical at scale: AT&T, one of the largest wireless carriers in the US, estimated that surveying just 25% of its nationwide network would cost \$45 million per year [Dan23]. To address this challenge, we propose a deep learning framework, NeWRF [LVM24], that learns to predict a wireless coverage map from sparse measurements, borrowing ideas from Neural Radiance Fields (NeRF) [MST21] in computer graphics to implicitly represent a wireless scene with a neural network. The network learns a continuous function mapping a 3D location to the radiated signal amplitude and phase at that point, and we introduce a physics-grounded neural channel synthesis algorithm that synthesizes the channel at any location from this representation. We implement the framework and demonstrate its effectiveness on a synthetic dataset, showing that it can accurately predict the wireless channel at any location with training sample density as low as 0.2 samples per ft^3 . This approach substantially reduces the need for site surveys and drive tests, enabling more efficient, cost-effective wireless network planning.

Contribution: a differentiable ray-tracing framework for mmWave scene calibration and channel prediction. Compared to sub-6GHz signals, mmWave signals are more sensitive to blockage and reflection from the surrounding environment. Their prop-

agation is highly directional, and channel quality can vary sharply with the geometry and material properties of the environment. Accurate mmWave channel prediction faces two new challenges. First, mmWave radios are highly directional, adding a spatial dimension to the prediction problem: one needs to predict both the signal’s arrival direction and its strength at any location. Second, mmWave signals are more sensitive to the material properties of surrounding objects, which can substantially affect reflection and absorption. To address these challenges, we propose mmDiff [LCY26], a differentiable ray-tracing framework that combines physics-grounded simulation with deep learning to learn the radio material properties in the environment from sparse angular power spectrum measurements. The framework embeds the physics of wireless signal propagation into an automatic-differentiation pipeline, enabling gradient-based optimization of the material properties. To make the framework robust to noise in the 3D-reconstructed geometry model, we introduce a directional scattering model to replace the traditional specular reflection model. We implement the framework and demonstrate its effectiveness on both synthetic and real-world datasets. Empirical evaluation shows that mmDiff is robust to noisy geometry, predicting the correct beam direction over 90% of the time with a power error within 1 dB, substantially outperforming baseline techniques.

1.3 Organization of the Dissertation

This dissertation is organized into the following chapters. Chapter 2 details MilBack, providing hardware and algorithm designs to enable unified communication and sensing. Chapter 3 introduces NeWRF, focusing on the neural representation of a wireless scene and the learning technique for optimizing the scene from sparse measurements. In Chapter 4, mmDiff is presented, which incorporates differentiable ray-tracing to achieve noise-robust channel prediction. Finally, Chapter 5 concludes the dissertation and discusses future research directions in the field.

CHAPTER 2

MilBack: A Millimeter Wave Backscatter Network for Two-Way Communication and Localization

Millimeter wave (mmWave) technology refers to very high-frequency signals, typically between 24 GHz and 100 GHz. Over the past decade, there has been growing interest in using this technology in wireless networks, which offer three advantages over traditional wireless networks. First, they provide much higher throughput by exploiting the large bandwidth available in the high-frequency spectrum. The Federal Communications Commission (FCC) has released more than 14 GHz of bandwidth for licensed and unlicensed use in mmWave bands above 24 GHz, orders of magnitude more than the bandwidth allocated to today's WiFi networks [Coo16]. Second, mmWave networks can connect many nodes simultaneously on the same frequency band. The small wavelength of mmWave signals, on the millimeter scale, allows hundreds of antennas to be packed into a small area, creating an antenna array with a highly directional beam. These directional beams let mmWave networks perform space division multiplexing, supporting multiple links simultaneously on the same frequency band [NLJ15, HCX15, RFN13, VDP19]. Third, the short wavelength combined with large bandwidth and narrow beams enables mmWave networks to perform highly accurate range and direction localization of their devices, opening up new applications such as indoor navigation and asset tracking.

Despite these advantages, existing mmWave networks have a major limitation that makes them unsuitable for many IoT applications: their radios consume too much power for devices with limited energy sources. This problem is intrinsic to the technology, since RF circuit

power consumption scales with operating frequency [Lee03], so components like mixers and oscillators draw much more power at mmWave frequencies than at lower ones. mmWave radios also require steerable directional antennas, such as phased arrays, which further increase their power consumption compared to traditional radios [KSP18, YRG19].

To overcome the high power consumption problem of mmWave networks, researchers have designed mmWave backscatter devices [MCA21, SPB21]. The vision behind this approach is to let mmWave devices piggyback their data on a mmWave signal emitted by an access point (AP), instead of generating and transmitting their own signals. Eliminating the need for an active transmitter and power-hungry RF components allows mmWave backscatter devices to communicate on a very low energy budget. However, existing mmWave backscatter devices support either accurate localization or uplink-only communication, but not both [MCA21, SPB21, BAC22, NQZ21], so they cannot serve applications that need both uplink and downlink connectivity, such as Virtual Reality (VR) and Augmented Reality (AR). To the best of our knowledge, no current mmWave backscatter system enables localization, uplink, and downlink communication together.

In this chapter, we introduce MilBack, the first mmWave backscatter network that enables localization, uplink, and downlink communication. MilBack’s design consists of a novel backscatter node and an access point (AP), supporting both uplink and downlink communication between the two while also enabling accurate localization and orientation sensing of nodes. MilBack achieves this through several key innovations:

1) Two-Way Backscatter Communication and Localization: The first challenge is to develop a low-power backscatter node that supports both two-way communication and localization. To achieve this, MilBack introduces a novel backscatter node architecture built around a passive mmWave structure that creates two directional beams, each of which can operate in absorptive or reflective mode. In absorptive mode, the node receives the AP’s signal through its two beams and passes it to a simple baseband processor for demodulation, enabling downlink communication from the AP to the node. In reflective mode, each beam

reflects only a specific frequency component of the AP’s signal back to the AP, where these frequencies depend on the node’s orientation relative to the AP. This enables the AP to estimate the node’s location and orientation by transmitting a wideband signal and monitoring which frequencies reflect back to it. Finally, by switching the beams between reflective and absorptive mode, the node modulates the AP’s signal as it reflects it back, supporting uplink communication.

2) Orientation-Assisted Quadrature Frequency Modulation (OAQFM): The second challenge is to develop a modulation scheme that allows the backscatter node to modulate and demodulate signals using a simple, low-power baseband processor. Typical modulation schemes, such as Quadrature Amplitude Modulation (QAM), represent complex symbols using cosine and sine components at the same frequency, which requires a mixer and oscillator to separate and modulate or demodulate these components. These components are too complex and power-hungry for mmWave backscatter nodes. To address this, MilBack introduces Orientation-Assisted Quadrature Frequency Modulation (OAQFM), which represents complex symbols using two cosine tones at different frequencies instead of cosine and sine at the same frequency. MilBack’s node receives these tones from the AP through its beams and modulates or demodulates them using simple, low-power envelope detectors integrated at the output of its passive directional reflector/absorber.

In summary, we make the following contributions:

- We introduce the first mmWave backscatter system which enables localization, orientation sensing, uplink, and downlink communication.
- We present a new modulation scheme for mmWave backscatter nodes that exploits node orientation and two single-tone signals to enable two-way communication between the AP and nodes.
- We design a two-port mmWave structure that operates in reflective or absorptive mode. We integrate it with off-the-shelf components to build MilBack’s node and the AP. We




Figure 2.1: **Frequency Scanning Antenna (FSA)**. The passive FSA structure creates beams whose directions depend on the signal frequency. Here, we show beams for only seven frequencies; however, FSA enables continuous beam directions across its frequency band.

evaluate MilBack’s performance in real-world scenarios.

Our empirical evaluation shows that MilBack achieves accurate localization, downlink, and uplink communication while consuming only 18 mW and 32 mW, respectively, which is significantly lower than the power consumption of past mmWave backscatter networks that only support uplink [MCA21].

2.1 Preliminaries

FSA Technology: Frequency Scanning Antenna (FSA) is a passive structure that creates a beam whose direction depends on the frequency of the signal, as shown in Figure 2.1. Said differently, as the frequency of the received/transmitted signal changes, the direction of the beam changes too. Moreover, if the signal includes two or multiple different frequencies, the structure creates two or multiple different beams, where the direction of each beam corresponds to that specific frequency. The FSA structure consists of an array of emitting elements. As the signal propagates through the structure, it experiences phase shifts between the consecutive emitting elements. Therefore, the emitting signals will combine construc-




Figure 2.2: **FMCW Radar Signal.** FMCW radar transmits a chirp signal (blue). The signal gets reflected back by an object, and the radar receives the delayed version of it (orange). The time-of-flight (ToF) can be discovered from the frequency difference between the transmitted and received signal.

tively over the air, creating a beam. However, as the signal frequency changes, the amount of phase shift between two consecutive emitting elements changes too. Therefore, the direction in which signals combine constructively (i.e., the direction of the beam) will change with the frequency of the signal. FSA has been mostly used in imaging and weather radars to enable multi-dimensional environment sensing and imaging [ZKQ17, KZ13]. Recently, it has also been used to enable passive beamforming for relays [LMA22]. Another recent work [NST22] has also explored using FSA in localization and user identification; however, their design requires 10 GHz bandwidth to cover 48° azimuth angle. MilBack builds on past work on FSA, and designs a *dual-port* mmWave FSA to achieve passive and low-power beamforming for backscatter networks. Our FSA design covers over 60° azimuth angle with only 3 GHz bandwidth.

FMCW Technology: Frequency Modulated Continuous Wave (FMCW) radar is one of the most widely adopted wireless sensing technologies [Rao17, AKK14, AMK15, AHM15, AKK15]. This radar technology can measure the distance of objects from itself by trans-

mitting a signal and measuring the reflection from the objects. To do so, the FMCW radar transmits a chirp signal whose frequency changes linearly with time, as shown in Figure 2.2. This signal travels through the air and is reflected back by an object. The radar receives the reflected signal and compares its frequency to that of the transmitted signal. This comparison is done simply by multiplying the transmitted signal by the received signal and checking the frequency of the result. Note that, since the transmitted signal frequency is changing linearly in time, delays in the reflected signals translate into frequency shifts in comparison to the transmitted wave. Therefore, by comparing the frequency difference between the transmitted signal and the received signal, one can discover the time delay that the signal incurred, which corresponds to the time-of-flight (ToF) from the radar device to the object. In particular, ToF can be calculated as follows: $ToF = \Delta f / slope$, where Δf is the measured frequency difference between the transmitted and received signal, and slope is the slope of the frequency sweep. Finally, once the ToF is computed, the radar can estimate the round-trip distance of the object from itself by multiplying the ToF by the speed of light. Note that, although the above description is for a single object, it can be extended to an environment with multiple objects.

2.2 Related Work

mmWave and Backscatter Communication: The mmWave spectrum offers multi-GHz of licensed and unlicensed bandwidth that promises orders of magnitude higher network throughput in next-generation communication networks [NLJ15, WKK18, UM20, BHM18, MAA19]. Therefore, a wide range of applications has been proposed for this technology, such as Virtual Reality (VR) [ABD17, EPB18], data center networks [CXW17], and vehicular networks [VSB16, GZZ17, GKZ20]. Moreover, due to the availability of large bandwidth and short wavelength of mmWave signals, this technology can enable very accurate localization, enabling applications such as tracking and gesture sensing [XJM21, LMY16, LWZ20].

However, most existing mmWave work focuses on applications which have substantial energy and computing power. Therefore, today’s mmWave systems are not suitable for applications with limited energy sources.

Over the past decade, there has been a significant amount of work on backscatter technology. This technology is the most energy-efficient method for wireless communication. The most popular backscatter system is RFID, which operates at the 900 MHz band [XYZ18, NL19]. Recently, researchers have proposed backscatter systems operating on the WiFi band [ADM20, BJK15, KPG14, ZBJ16]. However, none of these systems enable accurate localization and high-data-rate uplink and downlink together. Moreover, these systems operate in sub-6GHz bands, which are already congested.

Researchers have recently proposed mmWave backscatter systems that bring mmWave technology to low-power devices that have limited energy and computational resources [MCA21, SPB21, BAC22, NQZ21, KGT17]. However, past work either enables accurate localization of the backscatter node or provides only uplink communication. In contrast, MilBack is the first mmWave backscatter system that provides accurate localization, uplink, and downlink communication. Table 2.1 shows the features of MilBack compared with the state-of-the-art mmWave backscatter systems.

2.3 MilBack Overview

In contrast to past mmWave backscatter systems, which enable either localization or communication (only uplink), we introduce MilBack, the first mmWave backscatter network that enables uplink, downlink, and accurate localization of its backscatter devices. MilBack’s network includes an AP and one or multiple backscatter nodes. The node operates in two modes: reflective or absorptive. In the reflective mode, it reflects the AP’s signal back to the AP. This enables the AP to accurately estimate the distance, direction, and orientation of the nodes. In particular, the AP uses directional antennas to create transmitting and receiving

beams. Then, it steers these beams together while transmitting its signal. When the beams are facing toward a node, the node reflects the signal back to the direction of arrival (i.e., the direction of the AP). The AP receives the backscatter signal and accurately estimates the distance, direction, and orientation of the node using FMCW. MilBack then exploits the node’s orientation information and enables both uplink and downlink communication between the node and the AP using its novel low-power design and modulation scheme.

Over the next few sections, we will explain how MilBack works in detail. We first present MilBack’s backscatter node design, and discuss how it creates a reflective or absorptive beam toward the AP without using active mmWave components such as phased arrays. We then explain how the AP can estimate the location and orientation of the backscatter nodes. Next, we describe how MilBack’s low-power, low-cost node enables both uplink and downlink communication. Finally, we explain MilBack’s protocol to enable two-way communication and localization.

2.4 MilBack’s Node Design

Due to the high-frequency nature of mmWave signals, these signals decay quickly with distance. Hence, a backscatter node operating at mmWave frequency needs to focus its reflected power into a narrow beam toward the AP to compensate for path loss. Otherwise, the AP does not hear the node’s reflected signal. A typical approach to create and steer beams in mmWave nodes is to use phased arrays. However, phased arrays are too complex and power-hungry for use in a low-power backscatter node. Hence, past work has proposed to use the Van Atta technique to create a beam in backscatter nodes [MCA21, SPB21, BAC22]. Van Atta is an array of antennas that are connected to each other through transmission line traces. This passive structure can create a beam and reflect signals in the same direction as their arrival direction [SD60]. However, we cannot use Van Atta in our design. In particular, Van Atta designs are targeted only for reflecting signals, and they cannot be used for receiv-




Figure 2.3: **Beam Pattern of MilBack's Dual-Port FSA.** Our dual-port FSA enables creating two sets of beams while their frequency assignments are mirrors of each other.

ing since they do not have any signal ports. Therefore, although they are great options for backscatter nodes that enable only localization or uplink, they are not suitable for MilBack, which targets both uplink and downlink.¹

To solve the above problem and create beams passively while supporting both uplink and downlink, our idea is to build a backscatter node using the FSA structure. As described in Section 2.1, FSA is a passive structure that creates a beam where the direction of the beam depends on the frequency of the signal. FSA does this without consuming any power, which makes it ideal for low-power designs such as backscatter nodes. Moreover, as shown in Figure 2.1, FSA has a signal port which can be used to receive and feed the signal to a receiver to enable downlink. In contrast to a typical FSA, described in Section 2.1, we design and build a *dual-port* FSA for MilBack's backscatter nodes.

Dual-port FSA: The FSA structure is symmetric, and hence by adding a second port to the other side of the FSA structure, we can enable two sets of beams, as shown in Figure 2.3.

¹In Van Atta, the length of traces between antennas should be carefully tuned. Thus, we cannot easily insert switches and/or ports in the middle of the traces to direct the signal to a local processor for downlink.




Figure 2.4: **Block Diagram of MilBack's Backscatter Node.** Our design is simple and low-power since it does not require power-hungry mmWave components such as amplifiers, mixers, and oscillators.

The red beams are created by Port A, and the blue beams are created by Port B. Therefore, for each direction, there will be two beams that correspond to different frequencies. For instance, in Figure 2.3, the rightmost beam corresponds to frequencies f_1 and f_7 for Port A and Port B, respectively. As we describe in the following sections, this dual-port FSA will enable MilBack to sense the orientation of nodes as well as support higher data-rate uplink and downlink communication.

The node's architecture: Figure 2.4 shows the block diagram of MilBack's backscatter node. Our design consists of a dual-port FSA antenna, where each port is connected to a switch. Each switch connects the FSA port to either the ground plane of the FSA (i.e., short circuit) or an envelope detector. When the FSA ports are connected to the ground, the FSA's beam reflects the signal back to the AP. On the other hand, when the FSA ports are connected to the envelope detectors, the FSA's beam absorbs the AP's signal and passes it to the envelope detectors. Hence, the node does not reflect any signals. This is because the envelope detector has a 50 ohm input impedance, which is matched with the impedance of the FSA's port. Finally, the outputs of the envelope detectors are connected to the ADC pin of a microcontroller unit (MCU), which also controls the switches. Note that the design

is simple, low-cost, and low-power. It does not use any costly and power-hungry mmWave components such as phased arrays, phase shifters, amplifiers, oscillators, or mixers. All it needs is two envelope detectors, two switches, and a low-power processor unit (such as a microcontroller).

2.5 Localization and Orientation

2.5.1 Location Detection

To find the location of a node with respect to the AP, MilBack uses FMCW as described in Section 2.1. In particular, the AP transmits an FMCW chirp signal while the node is in reflective mode. Then the AP measures the reflected signal from the node and estimates the signal’s time-of-flight and the node’s distance with respect to itself. However, to do this, MilBack needs to address a challenge; MilBack’s AP needs to extract the node’s reflected signal from the reflection of other objects in the environment. In particular, the node’s reflection is much weaker than the reflection of some other objects.

To solve the above challenge, we use a similar technique to the background subtraction technique used in radar systems [AKK14]. In our design, the node switches between reflective and absorptive mode at a 10 kHz rate (i.e., modulating the reflected signal) while the AP is transmitting FMCW chirps. The AP then takes the FFT of the received signal of five consecutive chirps, and subtracts every two pairs from each other. It then uses the results of these subtractions to detect the node’s reflected signal. Note that, since the chirp duration is much shorter than any changes in the environment, the subtraction will remove the reflected signals of other objects. On the other hand, since the node’s reflection is modulated, it will experience changes between consecutive chirps. Hence, the node’s reflected signal will not be removed by subtraction.

Using this technique, the AP can detect the node’s reflection from the signal reflected by other objects, and estimate the distance of the node with respect to itself.

2.5.2 Orientation Detection

The orientation sensing of a node can be crucial for applications such as VR and AR in determining the user’s gesture and direction. MilBack enables both the AP and the node to estimate the orientation of the node with respect to the AP.

(a) AP detecting node’s orientation: MilBack’s design enables the AP to estimate the node’s orientation. To do so, the AP transmits an FMCW signal while the node operates in the reflective mode. Recall that the node’s beam direction depends on the signal frequency. Therefore, the node creates a reflective beam toward the AP only for some specific frequencies of the FMCW signal. Said differently, the node reflects only some frequencies of the FMCW signal back to the AP, where those frequencies depend on the orientation of the node with respect to the AP. Hence, the AP can estimate the orientation of the node by monitoring which frequencies of the FMCW signal have resulted in the highest reflection power.

Similar to the localization process, we also need to distinguish between the node’s reflected signal and reflections from other objects. Therefore, we put one port of the node’s FSA in absorptive mode and switch the other port between absorptive and reflective mode to create different reflected signals over time. This enables the AP to take the FFT of the signal and perform background subtraction. After removing the background, the AP then takes an inverse Fourier Transform and measures the reflected signal power across MilBack’s mmWave FMCW band. This enables the AP to measure which signal frequency has resulted in the highest reflection power and use it to accurately estimate the orientation of the node.

(b) Node detecting its own orientation: MilBack’s design enables the node to also estimate its own orientation with respect to the AP. To do so, the node needs to find out which one of its beams is aligned toward the AP. Recall that the node uses FSA to passively create beams whose directions depend on the signal frequency. Moreover, for the beams that are aligned toward the AP, the node receives the highest power from the AP. Therefore, there is a one-to-one mapping between the node’s orientation and the signal frequency that

results in the highest received power at the node. Now the question is how the node can measure which signal frequency results in the highest received power.

The AP transmits an FMCW signal while the node is measuring the received signal power using its envelope detector. Note that although the envelope detector can measure the received power, it cannot identify the signal frequency. Hence, the node cannot detect which frequency has resulted in the highest received power. To solve this issue, our idea is to use a triangular chirp instead of the sawtooth chirp for the FMCW signal. In this case, although the node cannot measure the frequency of the signal that resulted in the highest received power, it can measure the delay between the two times when the node receives the highest power, as shown in Figure 2.5. In particular, the V-shape property of the triangular chirp creates different delays between the two peaks in the received power, where the amount of the delay depends on the orientation of the node. This enables the node to measure the delay and estimate its orientation. In the following section, we explain how this orientation information is used at the AP to choose correct carrier frequencies for both uplink and downlink communication.

2.6 Communication

Past backscatter mmWave systems enable either localization or communication (only uplink). In the previous section, we explained how MilBack enables localization and orientation sensing. In this section, we explain how MilBack also enables both downlink and uplink communication using its novel architecture.

2.6.1 Downlink Communication

Due to the significant path loss of mmWave signals, any mmWave transmitters and receivers must focus their energy into narrow beams to compensate for the loss. Therefore, communication is only possible once their transmitter and receiver beams are aligned. Recall




Figure 2.5: **Orientation Detection at the Node.** (a) a triangular FMCW waveform, and (b) received signal power at the output of the node’s power detector for three different node orientations.

that MilBack’s node uses our dual-port FSA to passively create two beams, as shown in Figure 2.3, where the direction of each beam depends on the signal frequency, and the signal of each beam is received by only one of the ports. Hence, for the AP to send data to the backscatter node, the AP first needs to know which two signal frequencies to use for communication. In particular, one frequency aligns the beam for Port A of the FSA toward the AP while the other frequency aligns the beam for Port B of the FSA toward the AP. Note that these frequencies can be selected based on the orientation of the node. In particular,

the AP estimates the nodes' orientation by measuring the reflected signal power across different frequencies, as described in Section 2.5.2. The frequencies that result in the highest reflected power are the ones that align the node beams toward the AP. Once the AP chooses the correct frequencies to align the node's beams toward the AP, it can then send signals at those frequencies, which are received at the two ports of the node's FSA. As shown in Figure 2.4, each port is connected to a dedicated envelope detector, which enables a simple microcontroller to demodulate and decode the signal in a simple and low-power approach, as we describe next.

2.6.2 OAQFM

MilBack takes advantage of the two-port FSA and introduces a new modulation technique, called Orientation-Assisted Quadrature Frequency Modulation (OAQFM). In this technique, the AP encodes its bits by modulating two carrier signals, where their frequencies are selected based on the orientation of the node. In particular, the presence or absence of each carrier signal can be used to represent four different symbols, as shown in an example in Figure 2.6. In this example, the AP has selected the f_A and f_B frequencies for communication since they have aligned the node's beams toward the AP. Here, if the AP wants to send bits "01" or "10", it transmits a single tone at f_B or f_A , respectively. Similarly, if it wants to send bits "11", it transmits two tones at f_A and f_B , simultaneously.

On the node side, each port of the FSA is connected to a switch, which can connect the FSA port to either ground or an envelope detector, as shown in Figure 2.4. In the downlink mode, the switch connects the FSA ports to envelope detectors whose outputs are connected to the microcontroller. This enables the node to measure the power of the signal received by each port. Since each port receives only one of the carrier frequencies (f_A or f_B), the node can seamlessly decode the bits by checking the presence or absence of a signal at each of the two FSA ports.

So far, we have explained how MilBack enables downlink communication by introducing




Figure 2.6: **OAQFM Modulation.** It uses the presence or absence of two tones (f_A and f_B) to represent different symbols.

dual-port FSA and OAQFM modulation. However, in cases where the node is normal to the AP (i.e., zero incidence angle), the FSA's beams for Port A and Port B will be operating at exactly the same frequency (i.e., $f_A = f_B$). Hence, the AP cannot use two carrier frequencies in these cases and needs to use a single carrier frequency. Note that we already explained how the AP and the node can estimate the orientation of the node. Therefore, when they sense that the node is normal to the AP, they both rely on a single carrier on-off-keying (OOK) modulation.

2.6.3 Uplink Communication

MilBack enables low-power uplink communication between its nodes and AP using the backscatter technique. In particular, the AP transmits a query signal consisting of two tones at frequencies that align the node's beams toward the AP. Then the node uses the OAQFM modulation to piggyback its data on the AP's query signal by selectively reflecting or absorbing each tone. For example, to send '01' to the AP, the node reflects the tone at f_A while absorbing the tone at f_B . Similarly, to send '10' to the AP, the node reflects the tone at f_B while absorbing the tone at f_A . Finally, to send '00', the node absorbs both tones, and to send '11', it reflects both tones. Now, the question is how the node can change each

beam to absorptive or reflective mode independently.

Recall that each FSA port is connected to a switch, which connects the port to either the ground plane or an envelope detector. Moreover, each port receives only one of the tones in the AP's query signal. Hence, when an FSA port is connected to the ground, the FSA reflects back the tone which corresponds to that specific port to the AP. On the other hand, when the FSA port is connected to the envelope detector, the FSA absorbs that tone since the input impedance of the envelope detector matches the impedance of the FSA's port. Therefore, the node can switch each beam to reflective or absorptive mode independently by connecting each FSA port to either ground or the envelope detector.

So far we have explained how the node can modulate the AP's signal and encode its data using OAQFM. However, when the AP receives the node's signal, it first needs to extract it from interference signals. There are two types of interference in our system. The first one is self-interference, which is caused by the AP itself. This is because the AP receives the node's weak signal while transmitting a strong query signal. The second is caused by signals reflected by other objects, such as walls and desks in the environment, which also interfere with the node's signal. To mitigate this interference at the AP, our idea is to multiply the received signal by each tone of the AP's query and then filter out the interference, as we explain below.

Figure 2.7 shows the block diagram of MilBack's AP. The AP transmits a query signal consisting of two tones (i.e., $\cos(2\pi f_A t) + \cos(2\pi f_B t)$) using one antenna and receives the reflected signals using two antennas. The signal at each receiving antenna is first amplified using an LNA and then multiplied by one of the query signal's tones using a mixer. Note that, since the interference signals (both the self-interference and reflection from surrounding objects) are just a delayed version of the transmitted signal, when they are multiplied by the tone (e.g. $\cos(2\pi f_A t)$), the results will be a DC signal and some high-frequency signals (i.e. $\cos(2\pi 2f_A t)$, $\cos(2\pi f_{A+B} t)$ and $\cos(2\pi f_{A-B} t)$) which can be easily filtered out using a band-pass filter (BPF). On the other hand, the node's reflected signal is the modulated




Figure 2.7: **Block Diagram of MilBack's AP.**




Figure 2.8: **MilBack Packet Structure.** Preamble is used for localization, sensing the node's orientation, and informing the node whether the system operates in uplink or downlink mode.

version of the AP's transmitted signal. Hence, when the AP multiplies it by the tone, the results will be some high-frequency signals (which are filtered out using a BPF) and the node's response at baseband frequency. Therefore, by doing this, the AP can easily extract the node's response from interference signals.

2.7 Two-Way Communication and Localization Protocol

Previous sections presented different components of MilBack, enabling localization and two-way communication between its backscatter nodes and the AP. Here, we explain MilBack's joint communication and localization protocol.

MilBack packets consist of a preamble followed by a payload, as shown in Figure 2.8. The preamble includes two fields: Field 1 and Field 2. The purpose of the first field is to enable the node to estimate its orientation and also inform the node whether the AP will send or receive data during the payload. The purpose of the second field is to enable the AP to localize the node. Finally, the payload will carry the uplink or downlink data.

During the first field of the preamble, the AP sends multiple triangular FMCW chirps while the node has its ports (beams) in the absorptive mode. This enables the node to sense its orientation, and the AP to inform the node whether the system will operate in the uplink or downlink mode during the payload. In particular, if the AP sends three chirps during this field, it means that the system operates in the uplink mode. In this case, during the payload part of the packet, the AP sends a query signal (i.e., two tones) while the node modulates it and sends its data to the AP, as described in Section 2.6.3. On the other hand, if the AP sends two chirps (with a gap in the middle, as shown in Figure 2.8) during the first field, it means that the system operates in the downlink mode. In this case, during the payload, the AP sends data and the node receives and decodes it, as described in Section 2.6.1. During the second field of the preamble, as shown in Figure 2.8, the AP sends five FMCW sawtooth chirps while the node switches its ports between the reflective and absorptive modes. This enables the AP to receive the node's reflection, localize it, and sense its orientation, as described in Section 2.5. Following the preamble, there will be a payload, which is used either for uplink or downlink. The length of the payload is predefined for both the AP and the nodes; however, it can be adjusted based on the application and data-rate requirements. Finally, it is noteworthy that MilBack can potentially support multiple nodes by using a spatial division multiplexing (SDM) technique. In particular, the AP can create multiple beams towards different nodes and establish communication links with them concurrently.

2.8 Implementation

MilBack Node: Figure 2.9 shows the prototype of MilBack’s node. We implemented the node using off-the-shelf circuit components and our customized dual-port FSA antenna. For the FSA design, we reproduced an FSA designed in a previous work [LMA22] and extended it to a dual-port FSA. We evaluated the design using ANSYS HFSS software and fabricated it on a PCB using Rogers Substrate. For the two switches, we used ADRF5020 SPDT RF Switch [Ana23c], which enables the node to connect the FSA ports to either a short circuit or the envelope detectors. For the envelope detectors, we used ADL6010 [Ana23a], which has $50\ \Omega$ input impedance, matched to the FSA impedance. Finally, we used a TI MSP-EXP430FR6989 [Tex23] module as a low-power microcontroller on the node to control the switches and process the downlink and uplink data.

MilBack AP: Figure 2.7 shows the block diagram of MilBack’s AP. For waveform generation, we used a Keysight M9384B VXG, which can generate the FMCW signal for localization and the two-tone signal for communication. The signal is then amplified using an Analog Devices ADPA7005 power amplifier [Ana23b] and fed to a directional antenna. The signal power is 27dBm. Each FMCW chirp takes $45\ \mu\text{s}$ and $18\ \mu\text{s}$ in Field 1 and Field 2 of the preamble, respectively. We have chosen slower chirps for Field 1 since the sampling rate of the node’s microcontroller is lower than the AP’s sampling rate. Each chirp spans 3 GHz bandwidth from 26.5 GHz to 29.5 GHz². At the receiver, we used two directional antennas and ADL8142 Low Noise Amplifiers (LNAs) to amplify the signal. We used Mini-Circuits ZMDB-44H-K+ RF mixers to mix the received signal with the transmitting signal. ZFHP-0R50-S+ and ZFHP-0R23-S+ are used to implement band-pass filters (BPFs). The filter outputs are fed to a Keysight DSOX3102G oscilloscope controlled by a laptop to capture and process the baseband signal. The waveform generator and oscilloscope are synchronized

²The maximum bandwidth of our signal generator is 2 GHz. We transmitted two 2 GHz chirps centered at 27.25 GHz and 28.75 GHz, respectively, and patched the results together.




Figure 2.9: **MilBack’s Node Implementation.** Our prototype uses evaluation boards for switches and envelope detectors. However, one can design a compact prototype of MilBack’s node by integrating all components on the same PCB as the FSA.

externally. Two points are worth mentioning. First, we used Mi-Wave 261(34)-20/595 horn antennas with 20dB gain and mechanically steered the AP’s beams. However, one can instead use a phased array to electronically steer the AP’s beams. Second, we have used a Keysight waveform generator and oscilloscope in our setup since they provide more flexibility for debugging and evaluating the end-to-end performance of the system. However, one can simply design a complete MilBack AP using off-the-shelf components. In fact, all components required for our AP are common in most commercial mmWave radars, such as those used in automobiles [Tex22].

2.9 Evaluation

We first evaluate the performance of different components of MilBack. We then evaluate the performance of the complete MilBack in an indoor environment with objects such as tables, chairs, and shelves.




Figure 2.10: **Dual-port FSA Beam Pattern.**

2.9.1 Micro-benchmarks

Dual-port FSA We evaluate the beamforming performance of MilBack’s dual-port FSA using HFSS. Figure 2.10 shows the beam pattern of each FSA port. Note that the FSA steers beams continuously; however, we only plot the beams for seven sample frequencies. These results show that the FSA creates two sets of beams with more than 10dB gain, where the direction of each beam depends on the signal frequency. Moreover, the beam direction for each port can change by up to 60 degrees as we change the signal frequency over our FMCW bandwidth (i.e., 26.5 GHz to 29.5 GHz).

OAQFM As mentioned in Section 2.6.2, MilBack introduces the OAQFM modulation scheme for communication. Here, we provide the micro-benchmark for this modulation scheme. We place MilBack’s node 2 m away from the AP. The AP estimates the node’s orientation and identifies 27.5 GHz and 28.5 GHz as the correct signal frequencies to send data to the node. The AP then sends symbols "00", "01", "10" and "11", consecutively. In particular, the transmitter first sends nothing for a symbol duration (i.e., $1\mu s$ in this experiment). It then sends a single tone at 27.5 GHz for a symbol duration and another single tone at 28.5 GHz for another symbol duration. Finally, it sends two tones simultaneously.




Figure 2.11: **OAQFM Micro-benchmark.** Signals at the node’s FSA ports when symbols 00, 01, 10, 11 was sent by the AP.

Figure 2.11 shows the output voltage of the envelope detectors at the two ports of the FSA. This result shows that MilBack’s node has successfully separated the two tones at two different FSA ports. Moreover, the envelope detector at each FSA port measures the presence or absence of the mmWave signal, enabling the microcontroller to decode the bits.

2.9.2 Localization Performance

We evaluate the localization performance of MilBack. To localize a node in a 2-D plane, the AP needs to find the distance and angle of the node with respect to itself. In this experiment, we place our node at various angles and distances with respect to the AP while the AP transmits its signal and estimates the node’s location. In particular, to measure the distance, the AP uses the FMCW technique, as described in Section 2.5, and to measure the angle, the AP compares the phase of the node’s baseband signal at the AP’s two antennas and uses that to estimate the direction of the node. We perform this experiment for different node-to-AP distances. For each distance, we repeat the experiment 20 times. We measure




Figure 2.12: **Localization Performance.**


the ground truth using a BOSCH Laser meter and a protractor.

Figure 2.12a shows the mean and 90th percentile error in estimating the distance of the node with respect to the AP. As distance increases, the system’s accuracy decreases because the SNR degrades. However, the mean accuracy is less than 5 cm and 12 cm, even when the node is 5 m and 8 m away, respectively. Figure 2.12b shows the CDF of the error in estimating the angle of the node with respect to the AP. The median and 90th percentile errors are 1.1 and 2.5 degrees, respectively. Note that the angle estimation can also be further improved if the AP uses a phased array with a large number of elements.


2.9.3 Orientation Performance

Next, we evaluate MilBack’s orientation sensing performance. Recall that our orientation estimation is performed at both the AP and node sides; therefore, we evaluate them separately.

Orientation Sensing at Node In this experiment, the AP transmits the triangular FMCW pulses as described in Section 2.8. We place the node at 2 m from the AP. Both ports of the node are set to absorptive mode, where the microcontroller samples the output of the envelope detectors at 1 MHz and measures the time separation between the two peaks on



(a) **Orientation Estimation at the Node.**



(b) **Orientation Estimation at the AP.**

Figure 2.13: **Orientation Estimation Performance.**

each port to estimate the node orientation using the technique described in Section 2.5.2. The estimates from the two ports are averaged to produce the final estimate of the node orientation. We perform this experiment for different node orientations. For each orientation, we repeat the experiment 25 times. We measure the ground truth using a protractor.

Figure 2.13a shows the mean and variance of error in estimating the orientation of the node for different orientations. The results show that the mean error in orientation estimation is always less than 3 degrees, which is sufficient for many IoT applications. For comparison, smartphones can estimate the orientation of the device with an error between 0.5 and 3 degrees [KGR21].

Orientation Sensing at AP Next, we evaluate how accurately the AP can estimate the node’s orientation. To do so, we put the node 2 m away from the AP. The AP transmits an FMCW signal and estimates the node’s orientation, as described in Section 2.9.2. We perform this experiment for different node orientations. For each orientation, we repeat the experiment 25 times. We measure the ground truth using a protractor.

Figure 2.13b shows the mean and variance of the error in estimating the orientation of the node by the AP. In general, the mean error is less than 1.5 degrees for all orientations.

However, the error is slightly higher for orientations between -6° and -2° . This is because the mirror reflection of the FSA structure (created by its ground plane) is colliding with the modulated backscatter reflection. The mirror reflection increases the interference in our system. Ideally, the mirror reflection should be removed by the background subtraction technique; however, since it varies due to node switching, it will not be removed completely. Despite this issue, the AP can still estimate the orientation of the node with less than 3-degree error on average. As mentioned in Section 2.6, the AP uses the orientation estimation to detect which signal frequency to use for the OAQFM modulation scheme. However, as the beam width of the node is around 10 degrees, a 3–4 degree error in estimating the node’s orientation will not impact communication performance, as we evaluate next.

2.9.4 Downlink Performance

Next, we evaluate MilBack’s performance in establishing a communication link from the AP to the node. Recall that we use the OAQFM modulation scheme for downlink communication, where the AP sends data to the node using the presence or absence of two tones, as described in Section 2.6. The frequencies of these two tones depend on the node’s orientation, which is estimated in the orientation sensing phase. In this experiment, we place the node at a fixed location. Then, the AP first finds the correct frequencies for two tones and sends its bits to the node using the OAQFM scheme. The node receives the AP’s signal using its FSA antenna and passes it to the envelope detectors, which are connected to the microcontroller. We measure the Signal-to-Interference-plus-Noise Ratio (SINR) of the signal at the microcontroller input. Note that we measure SINR instead of SNR since the beam created by each port has sidelobes, which may be in the same direction as the main beam of the other port. This non-ideality results in a small interference between the signals of the two ports. Hence, we report our results in terms of SINR (instead of SNR) to account for both noise and this interference.

Figure 2.14 shows the result of this experiment across different distances of the node with





Figure 2.14: **Downlink Performance.**


respect to the AP for a downlink bandwidth of 1 GHz. The figure shows that as the distance between the node and the AP increases, the SINR is reduced. However, MilBack enables an SINR of more than 12 dB even when the node is 10 m away from the AP. This SINR is more than enough to enable very low BER (i.e., less than 10^{-8}). Finally, it is worth mentioning that the maximum downlink data rate of MilBack is 36 Mbps, which is mainly limited by the rise and fall time of the envelope detector. Hence, one can increase the data rate further by using a faster envelope detector. Another option is to define denser OAQFM modulation schemes, where each symbol represents more bits by considering different amplitudes for each OAQFM tone.

2.9.5 Uplink Performance

Finally, we evaluate the performance of MilBack in enabling uplink communication between the node and the AP. We place the node at a fixed location with respect to the AP. The AP first estimates the orientation of the node and uses it to find the correct frequencies for two tones. Then, the AP transmits a query signal consisting of these tones as described.



(a) Data-rate of 10 Mbps.



(b) Data-rate of 40 Mbps.

Figure 2.15: **Uplink Performance.**

The node piggybacks its data on the query signal by switching each port of its FSA between reflective and absorptive mode. The AP receives the node's reflection, and separates it from the interference by down-converting it to the baseband signal for decoding, as described in Section 2.6.

Figure 2.15 shows the SNR of the node's signal at the AP across different distances for uplink data rates of 10 Mbps and 40 Mbps. Note that as the data-rate increases, the operating bandwidth increases. Higher bandwidth results in a higher noise floor and lower SNR. The figure also shows the corresponding BER for each SNR. These results show that MilBack achieves a very low BER for uplink data rates of 10 Mbps and 40 Mbps even when the node is 8 m and 6 m away from the AP, respectively. Note that the maximum uplink data rate that the node can operate at is 160 Mbps. This rate is limited by the switching speed of the node's switches. One can increase the uplink data rate by using faster switches and improving the range (or BER) by using larger FSA's with higher gains. Finally, it is worth mentioning that MilBack achieves higher SNR in downlink compared to uplink. This is because, in uplink, the signal gets attenuated by the channel twice, as opposed to downlink, in which the signal goes through the channel once. However, for both uplink and downlink, MilBack's novel design enables very low-power communication.

System	Uplink Comm.	Localization	Downlink Comm.	Orientation Sensing
mmTag [MCA21]	Yes	No	No	No
Millimetro [SPB21]	No	Yes	No	No
Omniscatter [BAC22]	Yes	Yes	No	No
MilBack (This Work)	Yes	Yes	Yes	Yes

Table 2.1: **Comparison with State-of-the-Art mmWave Backscatter Systems.**

2.9.6 Power Consumption

A key promise of MilBack is that it enables low-power two-way mmWave communication and localization for devices with limited energy resources. We evaluate the power consumption of MilBack’s node in different modes. The design of MilBack’s node mostly uses passive components and does not require any power-hungry mmWave components such as amplifiers, mixers, and oscillators. The only active components used in the design are two switches and two envelope detectors, which are simple and low-power. The total power consumption of MilBack’s node is 18 mW during localization and downlink, and 32 mW during uplink.³ The higher power consumption of uplink is due to the fact that switches are operating at higher rates. Note that, considering the downlink and uplink data rates of 36 Mbps and 40 Mbps, the energy efficiency of MilBack is 0.5 nJ/bit and 0.8 nJ/bit, respectively. This is much lower than the energy efficiency of past mmWave backscatter networks, which consume 2.4 nJ/bit while they only support uplink [MCA21].

³Note that this power consumption does not include the power consumption of the microcontroller since it is already available in user devices such as AR headsets and most IoT devices. A typical microcontroller such as the one used in our prototype consumes 5.76 mW.

2.10 Conclusion

mmWave backscatter technology promises to bring mmWave connectivity to devices with limited energy resources such as IoT devices. However, existing mmWave backscatter systems enable either accurate localization or uplink communication. Hence, they are not suitable for most applications since they cannot enable downlink. This chapter introduces MilBack, the first mmWave backscatter system that enables orientation sensing, localization, uplink, and downlink. MilBack achieves this by proposing a novel architecture and modulation scheme for mmWave backscatter nodes. We believe that with the advancement of mmWave backscatter systems, future mmWave access points (such as 5G/6G access points) and radars (such as automobile radars) can directly communicate with low-power IoT devices using mmWave signals. Finally, it is worth mentioning that MilBack is very low-power (i.e., consumes only 18 mW) while achieving 8 m range and 40 Mbps data-rate. This range and data-rate are sufficient for most IoT applications. However, both range and data-rate can be further increased by designing a larger FSA and faster switches in our prototype, respectively.

CHAPTER 3

NeWRF: A Deep Learning Framework for Wireless Radiation Field Reconstruction and Channel Prediction

Wireless networks (such as WiFi and 5G) have become an essential part of our lives. Real-world WiFi and cellular deployments commonly encounter many issues such as dead spots, dropped signals, sudden outages, and slow throughput, potentially undermining the connectivity of wireless devices. To address these issues, site surveys [KD18] are often conducted to measure wireless network quality across various locations and, if required, optimize base station deployment. However, to conduct an effective site survey, an experienced engineer needs to perform measurements for a very dense grid of points. Today, this exhaustive measurement is rarely done due to its intractable time and cost. Instead, a sparse measurement approach is usually taken, where the engineer walks with a radio receiver on the site and measures the wireless channels at random locations. A wireless channel is a measure of the distortions imposed on wireless signals as they propagate from a transmitter to a receiver, which determines the quality of communication. The wireless channel is represented as a complex number that characterizes factors such as signal attenuation, phase rotation, and interference. Although the random sparse measurement approach is much more efficient than the exhaustive grid-based survey, it fails to uncover the signal quality at unvisited locations; therefore, it can potentially miss many dead spots. As a supplement, wireless ray-tracing simulations [Rem23] are also performed to analyze the distribution of wireless




Figure 3.1: **Wireless Field Tomography in a Simple Environment.** Like water ripples, wireless signals can add up constructively or destructively, causing fine-scale spatial variation. Our algorithm, NeWRF, reconstructs this field from a sparse set of channel measurements.

fields, which use Computer-Aided Design (CAD) models of the environment. However, ray-tracing simulations are not reliable since CAD models cannot fully replicate how real-world environments interact with wireless signals. Therefore, successive site surveys are always required for further validation and calibration.

In this chapter, we introduce a deep learning framework that reduces the time and cost of wireless site surveys. Our goal is to predict the wireless channel at every location in space using only a sparse set of channel measurements. One straightforward approach is to train a machine-learning model that maps a spatial location to the corresponding wireless channel. However, this approach fails to incorporate the underlying physics of wireless propagation, and thus still requires very dense measurements to achieve reasonable performance. The challenge stems from the fine-scale spatial variation of the channel, caused by the constructive/destructive interference of electromagnetic waves propagating via different paths (as shown in Figure 3.1).

To solve this problem, we are inspired by recent advances in computer vision and computer graphics, in particular, Neural Radiance Fields (NeRF) [MST21]. NeRF is a novel

deep learning framework that achieves impressive performance in 3D reconstruction and view synthesis tasks. Specifically, by learning from a collection of photos of a scene, NeRF can synthesize photorealistic images from any viewpoint, even those not captured in the original dataset. The key innovation of NeRF is to represent a scene’s radiance field as a continuous function, parameterized by a Multilayer Perceptron (MLP) network. This model learns to correlate 3D spatial coordinates with colored radiance and density from 2D images. NeRF uses ray-tracing and volume-rendering techniques to synthesize images, which incorporate the underlying physics of light propagation.

We build on NeRF to predict wireless channels at unvisited locations by learning the wireless radiation scenes from a sparse set of channel measurements. However, adapting NeRF to predict wireless channels requires addressing several important challenges due to fundamental differences between visible light and wireless signals. First, wireless environments, typically at the scale of a room, are much larger and more complicated than the scenes handled in NeRF. Although there have been recent attempts to learn room-scale scenes with NeRF [RBM22, AMG22], they all require additional depth information for regularization, which is hard to obtain for wireless measurements. Second, each wireless measurement results in a single complex number, which contains much less information than an image of thousands of pixels. As a consequence, training a model with merely wireless channel measurements suffers from restricted information. Third, unlike image sensors in a camera that capture light rays from a single direction, wireless antennas capture signals from all directions. Tracing rays for antennas requires sampling all directions, which significantly enlarges the search space. We find that training the model to trace all ray directions is both memory-hungry and prone to poor convergence. Last but not least, the propagation of wireless signals is much more complex than that of visible light. Effects such as attenuation, phase rotation, reflection, and interference need to be characterized carefully to enable the model to learn a realistic wireless radiation scene.

To address these challenges, we make the following contributions:

- We propose NeWRF, the first NeRF-based wireless channel prediction framework that integrates wireless propagation characteristics into NeRF, enabling the learning of wireless radiation scenes from sparse sets of channel measurements.
- We improve the convergence of our model through a novel ray-casting scheme that uses direction-of-arrival measurements from a wireless antenna array to guide the model to search the most critical directions.
- By inspecting the learned model of NeWRF, we discover, for the first time, the simple nature of wireless scenes that enables our model to learn, even in complex room-scale spaces using only sparse channel measurements.
- We propose a novel ray-searching algorithm to estimate the direction-of-arrival, enabling accurate channel predictions at any unvisited location.

In summary, we propose NeWRF, a novel deep learning framework for predicting wireless channels. We believe our work represents the first step towards using NeRF to predict wireless channels in complex indoor environments with sparse channel measurements, opening up many new opportunities to optimize the performance of future wireless networks.

3.1 Preliminaries

In this section, we provide background knowledge on both wireless communication and Neural Radiance Fields.

Wireless Channel A wireless communication system consists of a transmitter and a receiver, illustrated in Figure 3.2. The transmitter emits wireless signals, whose amplitude and phase are carefully adjusted to encode data. The transmitted signal can be represented as a complex number $x = Ae^{j\psi}$, where A is the amplitude, and ψ is the phase. As the signal propagates, its amplitude is attenuated by an attenuation factor, A_{att} , and the phase




Figure 3.2: **Wireless Communication Model.** Wireless signals emitted by the transmitter experience attenuation, reflection, absorption, multipath interference, etc. before reaching the receiver. Our framework models all of these factors.

is rotated by $\Delta\psi$. The signal captured by the receiver y can be written as

$$y = Ae^{j\psi} \times A_{att}e^{j\Delta\psi} = A \cdot A_{att}e^{j(\psi+\Delta\psi)}. \quad (3.1)$$

For free-space propagation, i.e., with no objects blocking the propagation path, the amplitude attenuation is proportional to the inverse of the propagation distance, and phase rotation is proportional to the propagation distance [Mol12]:

$$A_{att}(d) = \frac{c}{4\pi df}, \quad \Delta\psi(d) = -2\pi fd/c. \quad (3.2)$$

where d is the propagation distance, f is the signal frequency, and c is the speed of light. Interactions with objects, for example, reflection and penetration, introduce additional attenuation and phase rotations to the signal.

Additionally, transmitted signals propagate via multiple paths to reach the receiver. As a result, the received signal is the sum of multiple copies of the transmitted signal, each with different attenuation factors and phase rotations:

$$y = Ae^{j\psi} \times \sum_{l=0}^{L-1} A_{att}^l e^{j\Delta\psi^l} \quad (3.3)$$

where L is the total number of propagation paths, and l is the path index. The wireless channel, h , is defined as the ratio of the received signal and transmitted signal, and characterizes the distortions introduced by the environment:

$$h = \frac{y}{x} = \sum_{l=0}^{L-1} A_{att}^l e^{j\Delta\psi^l}. \quad (3.4)$$

Neural Radiance Fields (NeRF). NeRF learns a 3D volumetric representation of a scene's radiance field from 2D images. It represents a scene using an MLP model with 5-dimensional inputs: $f : (x, y, z, \theta, \phi) \rightarrow (r, g, b, \sigma)$, where x, y, z are the 3D spatial coordinates, and θ, ϕ are the 2D view direction of a ray traced from the camera. This function translates the 5D coordinates to the radiance of red (r), green (g), and blue (b) light, as well as the volume density, σ , which describes the degree to which the environment blocks or attenuates light at a location. To generate images from a particular viewpoint, NeRF performs ray tracing from each pixel of the image, takes discrete sample points $\{\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_N\}$ along each ray $\mathbf{r}$, and queries the MLP model with the coordinates of these samples for the radiance and volume density predictions. The volume density is translated to transmittance, T , and opacity, α , of the corresponding volume block. Transmittance refers to the amount of light that passes through volume blocks without being absorbed or scattered, while opacity describes the extent to which the light is absorbed or scattered by a particular volume block.

$$T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right), \quad \alpha_i = 1 - \exp(-\sigma_i \delta_i), \quad (3.5)$$

where i is the index of sample point, $\delta_i = \|\mathbf{p}_{i+1} - \mathbf{p}_i\|$ is the distance between consecutive samples.

A volume rendering algorithm is then employed to approximate the pixel value, $\hat{C}(\mathbf{r})$, as the weighted sum of the radiance at the sample points:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N T_i \alpha_i \mathbf{c}_i, \quad (3.6)$$

where N is the number of sample points along each ray, and $\mathbf{c}_i = [r, g, b]$ is the vector of radiance emitted at the sample point $\mathbf{p}_i$. The generated image is compared with the ground truth image to calculate the loss and update model parameters during training.

3.2 Related Work

NeRF has recently received huge attention from the computer vision and graphics communities. There has been a lot of work trying to improve and extend the original NeRF in various ways. For instance, [MRS21] extends NeRF to synthesize large scenes from unconstrained online photo collections; [PCP21] enables NeRF to learn dynamic scenes; [MHM22] handles raw images taken in dark environments. These works all use NeRF in the visible light spectrum, which is fundamentally different from wireless signals. In contrast, we address several challenges and extend the NeRF formulation to the radio frequency spectrum to learn the radiation scenes of wireless signals.

Predicting the wireless channel has been a long-standing problem in the wireless community [KM23, MM12, KHC21]. The existing studies are mainly divided into two categories: temporal prediction and spatial prediction. Temporal prediction work focuses on predicting the wireless channel at a future timestamp given the past observations [FSC23, VGS23]. Spatial prediction work attempts to predict the channel at an unseen location given the measurements at some other locations [KM23]. Here, we focus on solving the spatial prediction problem.

There have been some pioneering attempts in this field; however, the fine-scale spatial variation of channels makes this task extremely challenging. [KM23] proposes an analytical method to predict the wireless channel by extrapolating measurements taken at the boundaries of the area of interest. However, their method cannot deal with complex indoor environments, where multipath interference is more severe. It is also limited to 2D receiver layouts with carefully designed boundaries. In contrast, our method can predict wireless

channels in 3D indoor environments with sparse measurements at random locations. A recent work, NeRF² [ZAP23], proposed a NeRF-inspired technique to predict the direction-of-arrival heatmap for a fixed wireless receiver when the transmitter is placed at different locations in the space. We note that the varying location of the transmitter results in a varying wireless field at each measurement, which deviates from the nature of NeRF, which captures a static field representation¹. As a consequence, NeRF² needs to take the transmitter locations as a direct input to the MLP model and essentially learns a direct mapping from transmitter locations to the receiving heatmaps. This is evidenced by their requirement of extremely dense measurements ($178.1 \text{ measurements}/ft^3$ on average in their datasets) for training. We believe the reverse problem, i.e., learning a static radiation field sourced from a fixed transmitter, is more challenging. In this work, we bridge this gap and devise a novel algorithm to learn a wireless radiation field with significantly lower measurement density (i.e., $0.2 \text{ measurements}/ft^3$).

Another work [OKK23] proposes a NeRF-based neural surrogate for wireless ray-tracing simulation. Their method requires knowledge of the environment geometry (given as a CAD model) and focuses on learning the ray-surface interaction behavior for a wireless signal propagating in space. In contrast, our method directly learns the radiation field of the scene and requires no knowledge about the environment geometry (except for the boundary of the scene). This enables us to eliminate the need for CAD models as a replica of the real world, and thus our method can be easily extended to predict wireless channels in real-world scenarios.

¹This is analogous to learning a NeRF where the lighting condition of the scene changes from one image to another.

3.3 NeWRF

In this section, we present the design of NeWRF. We first provide a concrete formulation of the problem in Section 3.3.1. Then we present our customized neural channel synthesis algorithm in Section 3.3.2. We improve the convergence of our model with a novel DoA-guided ray-casting method introduced in Section 3.3.3. We present the optimization process of NeWRF in Section 3.3.4. In Section 3.3.5, we uncover the nature of wireless scenes and provide insights on how it enables our model to reconstruct wireless fields in complex environments. Finally, in Section 3.3.6, we present our novel ray-searching algorithm for predicting wireless channels at inference time.

3.3.1 Problem Description

We consider indoor environments (such as the conference room shown in Figure 3.3) with a single transmitter at a fixed location and multiple receivers placed randomly throughout the space. This setting represents a realistic scenario where a WiFi access point talks to multiple client devices, such as mobile phones, laptops, and various Internet-of-Things (IoT) devices. Each receiver (client device) measures the wireless channel and reports its measurements to a central database. Note that when multiple transmitters co-exist in the same space, communication systems naturally perform time/frequency multiplexing to avoid interference. Therefore, we assume the receivers can connect to each transmitter in a round-robin manner and construct a wireless scene for each transmitter. The propagation of wireless signals can be modeled as rays that are emitted from the transmitter, reflected by objects and walls, and finally captured by the receiver [Mit16]. The reflection of wireless signals follows the law of reflection, i.e., the angle of the reflected ray is equal to the angle of the incident ray, since the wavelengths of wireless signals (e.g., 12 cm for 2.4 GHz WiFi signals) are typically much larger than the roughness of object surfaces [Mit16]. Our goal is to predict the wireless channel at any location in the environment using the sparse measurements obtained from




Figure 3.3: **Example Environment and Transmitter/Receiver Locations.** We fix the transmitter location and measure the wireless channel at various receiver locations. The wireless channel is the sum of multiple propagation paths from the transmitter to the receiver.

the receivers.

3.3.2 Neural Wireless Channel Synthesis

The first step to synthesizing the wireless channel is learning a wireless radiation scene. A wireless radiation scene is a continuous function that maps a point in space to the wireless signal emitted from the point. This function is also direction-dependent, since a point may emit signals of different amplitudes and phases towards different directions. We approximate this wireless radiation scene with an MLP model that takes a 5D input of spatial (x, y, z) and direction (θ, ϕ) coordinates. The model outputs (A, ψ, σ) , which are radiated signal amplitude, signal phase, and volume density, respectively. Volume density indicates the




Figure 3.4: **Neural Channel Synthesis Algorithm.** We back-trace each propagation path from a receiver and query an MLP with the 5D coordinates (location and viewing direction) of points along the path, to predict the radiated signal (amplitude and phase) at that point. Then we integrate signals from all paths to synthesize a channel. We model reflections as virtual transmitters at the point of reflection.

probability that a point is a wireless radiator.

Figure 3.4 illustrates the synthesis process of a wireless channel for a particular receiver location. We back-trace the rays captured by the receiver, and along each ray $\mathbf{r}$, we take multiple samples to obtain a discrete set of points ². We query the MLP with the spatial and direction coordinates for each sample point $\mathbf{p}$ to get the radiated signal $A_{\mathbf{rp}}e^{j\psi_{\mathbf{rp}}}$ and the corresponding volume density σ . We translate the volume density into a weight $w_{\mathbf{rp}} = T_{\mathbf{rp}}\alpha_{\mathbf{rp}}$, which characterizes the contribution of the signal radiated at $\mathbf{p}$ to the channel at the receiver. We follow the definition of T and α in Eq. (3.5).

Note that two kinds of points contribute most to the channel at the receiver. One is the transmitter itself, which directly acts as a source of radiation, and the other is any point that reflects the transmitter signal. We consider the latter as virtual transmitters. The

²To avoid confusion, the term "samples" is used exclusively for the points taken along the rays. We refer to the channel measurement data samples in our datasets as "measurements".

radiated signal from a virtual transmitter includes the effects of absorption, attenuation, and phase rotation due to the reflection. The effects of free space propagation (Eq. (3.2)) are then explicitly included in our synthesis algorithm. Hence, the channel at a receiver can be presented as the following equation, where it is calculated by summing up signals radiated from all sample points along all rays:

$$H(f) = \sum_{\mathbf{r} \in \mathbb{L}} \sum_{\mathbf{p} \in \mathbb{P}} w_{\mathbf{rp}} A_{\mathbf{rp}} \frac{c}{4\pi d_{\mathbf{p}} f} e^{j(\psi_{\mathbf{rp}} - 2\pi f d_{\mathbf{p}}/c)}, \quad (3.7)$$

where $\mathbb{L}$ is the set of rays, and $\mathbb{P}$ is the set of points taken along a ray.

3.3.3 Improving the Convergence of NeWRF

One key difference between wireless measurements and images is that wireless antennas capture signals from all directions. Thus, in wireless, finding the ray direction to trace is a non-trivial problem. One potential solution is to search exhaustively along a discrete set of azimuth and elevation angle combinations that span all directions. However, this approach has two fatal drawbacks. First, the actual ray directions lie in a continuous space. Approximating them with a discrete set of ray directions would require an extremely dense angle grid, which significantly increases memory consumption for training the model. Second, for each receiver location, there are only 10~20 ray directions contributing significantly to the channel measurement, whereas a search grid of 1° resolution involves $360 \times 180 = 64,800$ ray angles. As a result, optimizing the model to find the actual ray directions from such a large ray space is very difficult, if not impossible. We find through experiments that the model fails to converge with such a large search space of ray directions. To solve this issue, we leverage the fact that the Direction-of-Arrival (DoA) of wireless signals can be measured using an antenna array at the receiver location [FSB22]. We use DoA measurements as prior knowledge and let the learning algorithm trace rays primarily toward directions identified by the DoA measurements. Through our experiments, we find that this DoA-guided ray-casting approach significantly reduces training time and memory consumption, and enables

the learning algorithm to converge quickly.

3.3.4 Optimizing NeWRF

As in NeRF, we simultaneously optimize two MLP models for each environment: one "coarse" and one "fine". We use mini-batch gradient descent to train our models; at each iteration, we randomly pick a batch of receiver locations from the training set and obtain the ground-truth channel and DoA measurements. We cast rays toward each DoA and use the stratified sampling strategy [MST21] to take coarse-grained samples uniformly distributed on each ray. We use these samples to query the "coarse" model and use the procedures described in Section 3.3.2 to synthesize coarse predictions of the channel for each receiver location. We use the predicted weights $w_{\mathbf{p}}$ of coarse samples to perform a round of hierarchical sampling which densely samples regions with high weight values. We query the "fine" model with the union of coarse and fine samples and synthesize the fine channel predictions. The coarse and fine channel predictions are each compared with the ground truth to calculate loss and update the parameters of each model.

In practice, we let the model predict the real and imaginary parts of the radiated signal, (I, Q) , instead of amplitude and phase, (A, ψ) . This is because phase is modulo 2π , which is not differentiable. A and ψ can be easily translated from I and Q using Euler's formula:

$$A = \sqrt{I^2 + Q^2}, \psi = \arctan(Q/I)$$

We use the Normalized Mean Square Error (NMSE) as the loss function:

$$L(\mathbf{h}, \hat{\mathbf{h}}) = \frac{\sum |\hat{\mathbf{h}} - \mathbf{h}|^2}{\sum |\mathbf{h}|^2} \quad (3.8)$$

where $\mathbf{h}$ is the ground truth channel, and $\hat{\mathbf{h}}$ is the predicted channel. Unlike NeRF, we do not use Mean Squared Error (MSE) as the loss function since wireless channels are typically at the scale of $10^{-1} \sim 10^{-5}$ and vary by orders of magnitude in the dataset. NMSE is more robust to scale changes.




Figure 3.5: **An Optimized Wireless Scene.** Spots with high volume density values indicate the locations of (virtual) transmitters. The gray cube shows the 3D environment model presented in Figure 3.3.

Finally, similar to other NeRF work, we noticed floaters, i.e., spurious particles that "float" in space, in the trained model. We find that casting a small number (5~10) of extra rays toward random directions during the training process can effectively guide the model to remove the floaters, resulting in a clean scene geometry. Further details about our training pipeline and model architecture can be found in Section 3.6.

3.3.5 NeWRF's Scene Representation

We train the model as described in the previous section and find it has good generalization performance. However, we find that the model learns a representation of the wireless radiation scene that was different from what we originally expected. Despite that, the scene still allows the model to predict wireless fields (channels) even in room-scale environments with merely a sparse set of channel measurements. In this section, we present our findings and




Figure 3.6: **The Nature of Wireless Scenes.** While back-tracing direct rays leads to the transmitter, back-tracing reflected rays consistently intersect at the virtual images of the transmitter.

insights.

As mentioned in Section 3.3.2, we model the reflection of wireless signals as a virtual transmitter emitting signals at the reflection point on the surface. Therefore, we expect our model to learn a representation where points with high volume densities appear on the reflector surfaces of the environment. Eventually, this would allow us to reconstruct a 3D model of the environment. However, we find that our model instead learns to assign high volume densities to the virtual images of the real transmitter against reflective surfaces, as shown in Figure 3.5. In this figure, we query the trained MLP model with a fine-resolution grid of 3D locations for the volume density predictions and plot the points with significant volume densities. The red dot indicates the volume blocks that match the real transmitter location. The green dots indicate the locations of virtual transmitters, which appear at virtual images of the real transmitter against each reflective surface, such as walls and tables.

To better understand why this is the case, we illustrate a simple wireless scene with a

single transmitter and two reflective surfaces in Figure 3.6. Three receivers are placed in random locations to measure the wireless channel. The signals emitted by the transmitter propagate via three major paths to reach each receiver: one direct line-of-sight path and two reflection paths ³. Since the receivers have no notion about the environment, they simply trace the rays back in the direction from which they arrived. Because of the law of reflection, the rays traced by all the receivers intersect at the locations of the virtual images of the transmitter and the transmitter itself. As a result, these points are sampled more frequently than other points in the space, giving the model an inductive bias to put higher volume densities at these locations.

We note that this representation is one of many possible solutions to this optimization problem that fit the training set. However, the model seems to have such a bias that it is more likely to converge to this representation of the scene. Although this representation is not what we had initially expected, it does reveal the simple nature of radiation scenes and allows for accurate channel synthesis from sparse measurements.

3.3.6 Predicting Channels at Unvisited Locations

While DoA information is known for our training data, it is not for the test data. Because of this, a new ray-casting technique was required for inference time predictions. Simply casting rays across a grid of all directions is problematic. Depending on the angle resolution used for sampling, we could under- or overestimate the contribution of a virtual transmitter on a channel. When the resolution is too low, virtual transmitters could be missed, and when it is too high, virtual transmitters could be sampled multiple times and contribute too much to the channel. For this reason, we developed a novel ray-searching algorithm that leverages our discovery of the nature of wireless scenes to extract DoA information and

³We omit the second-order reflection here for simplicity; however, it can also be characterized in the same manner.

thereby intelligently cast rays at new locations.

Our ray-searching algorithm consists of five steps:

1. **Intersection Point Identification:** For each receiver location, $\mathbf{m}$, in the training set, perform ray marching along each DoA to identify intersection points, forming a set of potential (virtual) transmitters, $\mathbb{O}$.
2. **Refine (Virtual) Transmitters:** Apply DBSCAN [EKS96] to cluster the estimated positions in $\mathbb{O}$. Add the centroid of each cluster to a reduced set, $\mathbb{V}$.
3. **(Virtual) Transmitter Assignment:** For each receiver location, $\mathbf{m}$, in the training set, find the closest (virtual) transmitter $\mathbf{v} \in \mathbb{V}$ along each DoA and add it to a set $\mathbb{V}_m$ of pertinent virtual transmitters for $\mathbf{m}$.
4. **(Virtual) Transmitter Count Estimation for Testing Locations** Train a fully connected neural network $g : \mathbf{m} \rightarrow |\mathbb{V}_m|$ to map receiver locations to the corresponding number of (virtual) transmitters. Then use this model to predict the size of $\mathbb{V}_t$, the set of pertinent (virtual) transmitters for test location $\mathbf{t}$, for all test locations.
5. **Voting-based (Virtual) Transmitter Selection:** For each $\mathbf{t}$, select its neighboring receiver locations from the training set (we use 6 nearest neighbors in practice). Populate $\mathbb{V}_t$ with the (virtual) transmitters most frequently assigned to the selected train locations. Finally, DoA can be calculated using the position of $\mathbf{t}$ and the estimated positions of (virtual) transmitters in $\mathbb{V}_t$.

We note that although this ray-searching algorithm can identify the locations of (virtual) transmitters, it cannot replace the role of training a NeWRF model for the scene, as we still need the prediction of signal amplitude and phase for each virtual transmitter.




Figure 3.7: **3D Environment Models for Simulation.**

3.4 Evaluation

To better understand the performance of NeWRF in representing a wireless scene versus its performance in predicting channels in unseen locations, we first evaluate how accurately NeWRF can predict channels, assuming DoA is known for each receiver location. We then evaluate the performance of NeWRF in channel prediction while the DoA is not known and we use our ray-searching algorithm (presented in Section 3.3.6) to discover DoA of unseen locations.

Datasets: We generate simulation datasets for evaluation using MATLAB. We use the Image Method [YI15] for ray tracing. We create datasets for three 3D environment models [MAT23a, MAT23b, Fre23] of different complexities. Figure 3.7 shows the geometry and dimensions of the three models. We place a transmitter in each environment and randomly distribute 443, 975, and 1907 receivers, respectively, to collect channel measurements. We use 80% of the measurements to train the model and the remaining 20% to test.

3.4.1 Channel Prediction Results

We first evaluate our algorithm using our conference room dataset simulated with perfect reflector materials at a signal frequency of 2.412 GHz. We will discuss the impact of material types and signal frequencies in Section 3.4.3.

Figure 3.8a provides a visualization of the predicted channel for the test set measure-




Figure 3.8: **Channel Prediction Results.** a) Prediction results for conference room environment. Each dot represents the channel at one test location. b) Predicted channel components (paths) for a single location. Each dot represents a single propagation path. The arrow exhibits the sum of all paths; c) The channel prediction accuracy in different environments compared with baseline methods. d), e), and f) channel prediction accuracy for different training set sample densities, environment materials, and signal frequencies.

ments. Since the wireless channels are represented as complex numbers, we plot the real and imaginary parts of the channels on the x-axis and y-axis, respectively. Each dot in the plot represents the channel for a particular receiver location. As is shown, our predictions align very well with the ground truth. Moreover, although we supervise the model with merely the sum of all multipath components, NeWRF can accurately predict individual path components for each particular receiver location as shown in Figure 3.8b. In this plot, each dot represents one propagation path of the wireless signal; the arrow indicates the sum of all the components. Note that this capability of predicting individual path components is very important for optimizing the configurations of base stations and relays to enhance the

network coverage in practice [KM23].

Next, we compare NeWRF with three baseline methods.

1) K-Nearest Neighbors (KNN): We predict the wireless channel at a test location as the average channel of the closest K locations in the training set. We pick K=3 since that gives us the lowest error.

2) Multi-Layer Perceptrons (MLP): We train an MLP model to directly map from the 3D location to the channel. We use an MLP with 7 linear layers, each of width 128. We train the model with the same conditions as NeWRF, as described in Section 3.6.

3) NeRF²: We use the open-sourced code of NeRF² [ZAP23] and train it on our datasets. Since their code assumes a dynamic transmitter scenario, we generate the corresponding simulation datasets by keeping each device’s location the same as in our datasets while swapping the roles of transmitters and receivers.

We use Signal-to-Noise Ratio (SNR) as our evaluation metric, which is defined as $\text{SNR} = -10 \log_{10}(\text{NMSE})$, where NMSE is Normalized Mean Square Error and defined in Eq. (3.8). Higher SNR indicates a higher prediction accuracy. For example, a 20 dB SNR is equivalent to a 1% relative squared error in the prediction. Figure 3.8c shows the results for three environments. NeWRF outperforms all of the baseline methods in all environments. Note that NeRF² does not perform well on our datasets since the measurement density ($0.4 \text{ measurements}/ft^3$) of our datasets is much lower than what they require ($178 \text{ measurements}/ft^3$). Finally, our results show that although the performance of NeWRF degrades as the environment becomes larger and more complex (such as a large office space with lots of furniture), it still achieves more than 5 dB SNR, which is sufficient for many wireless applications.

3.4.2 Measurement Density of Training Set

We further investigate the effect of training set sample density on the performance of NeWRF. To demonstrate this, we use the conference room dataset. In this experiment,

we use 25%, 50%, 75%, and 100% of the training set measurements to train our model, which corresponds to 0.1, 0.2, 0.3, and 0.4 *measurements/ft³*, respectively. We evaluate the trained model on the same test set for a fair comparison. As shown in Figure 3.8d, NeWRF can effectively learn the wireless radiation scene and predict the channel accurately with sample density as low as 0.2 *measurements/ft³*, which is around 800× lower than that of NeRF²[ZAP23]. This would significantly reduce the cost of wireless site surveys and improve dead-spot detection capability.

3.4.3 Impact of Materials Types and Signal Frequencies

The attenuation and phase rotation of wireless signals during reflection are largely determined by the material that constitutes the reflector. In particular, the material’s electrical properties, such as permittivity and conductivity, play significant roles. As such, the wireless properties of an environment can change greatly as the materials it is made of change. Furthermore, even the same materials behave differently when they interact with signals of different frequencies. NeWRF is able to characterize both of these effects, as shown in Figure 3.8e and ??, where we train a separate model for each material type and signal frequency. NeWRF maintains high prediction accuracy in all these scenarios.

3.4.4 Ray-searching Algorithm Performance

So far, we evaluated the performance of NeWRF while the DoA is known. Here, we evaluate the performance of NeWRF when we use our ray-searching algorithm to discover the DoA for each location. Figure 3.9 shows the results. Our results show that due to the error in DoA estimation, the performance of NeWRF in channel estimation degrades by a small margin (3~5 dB) when we use the ray-searching algorithm. This error is mainly due to a slight discrepancy in DoA between the test location and its nearest neighbors. Although this SNR loss is negligible, one can improve it by using slightly higher measurement density.




Figure 3.9: **Ray-Searching Algorithm Performance.** Performance of NeWRF in channel prediction when DoA is known using ground truth or discovered using our Ray-Searching algorithm. GT stands for ground truth, and RS stands for the Ray-Searching algorithm presented in Section 3.3.6.

3.5 Limitations and Future Work

In this chapter, we bring NeRF to radio frequency and demonstrate the promise of reconstructing wireless radiation scenes to facilitate site surveys. However, there remains a vast range of directions to investigate to unleash the power of NeRF for wireless applications. Here we highlight a few of them to help stimulate research in these directions:

a) Relaxing phase requirement: accurate phase measurement is challenging in real communication systems due to the synchronization error between the transmitter and receiver. The presented NeWRF algorithm, which relies on narrowband channels for training, does not perform well if the phase is inaccurate or not provided. One potential solution is to use wideband channel measurements for training, since wideband channels translate phase into amplitude variation across different frequencies, which can be measured accurately.

b) Relaxing DoA requirement: accurate DoA measurement is required for ray-casting in the current NeWRF algorithm; the error tolerance is within one degree based on our experiments. In practice, achieving such a measurement resolution would require a very large antenna array. Further studies are needed to relax this constraint.

c) Extending to larger environments: although NeWRF outperforms all the baselines for the three environments studied, a performance reduction with increasing environment complexity is observed. Scaling to more complex environments requires further research. One potential solution would be fusing visual and wireless radiation fields, where visual images provide more information about the environment to facilitate the learning of wireless fields.

3.6 Conclusion

We propose NeWRF, the first NeRF-based learning framework for wireless channel prediction with sparse measurements. We solve a series of challenges to enable the neural synthesis of wireless channels and, for the first time, reveal the nature of wireless scenes. We believe NeWRF represents the first step towards the practical use of NeRF in future wireless applications.

Appendix

Implementation Details

We implement NeWRF in PyTorch. Figure 3.10 shows the training pipeline. We use a batch size of 32 for each training iteration. As in NeRF, we use positional encoding to map the input coordinates to higher-dimensional space to help the model learn high-frequency variation of the wireless scene [MST21]. We use 10 frequencies to encode the spatial coordinates (x, y, z) , and 4 frequencies for the direction coordinates.




Figure 3.10: NeWRF's Training Pipeline




Figure 3.11: NeWRF's MLP Architecture.

Both Coarse and Fine model use the same architecture.

Figure 3.11 shows the MLP model architecture of NeWRF. It consists of 7 linear layers of width 128 and one output layer of width 64. All layers, as well as the σ output, are activated by ReLU functions. The I, Q outputs are activated with a $\tanh$ function to keep the output value within the range $[-1, 1]$. A single skip connection concatenates the input with the fourth linear layer. $\gamma(\cdot)$ represents the positional encoding function.

We use the same architecture for both "coarse" and "fine" models. A weighted sum of "coarse" and "fine" model loss is used for supervision:

$$loss = 0.1 \times coarse\ loss + 0.9 \times fine\ loss$$

We use the Adam optimizer [KB14] with an initial learning rate of 5×10^{-4} , and a REDUCELRONPLATEAU learning rate scheduler with patience 10 and factor 0.9. The other parameters are left as default. We train our model on a single NVIDIA A100 GPU, and it typically takes $\sim 100k$ iterations to converge.

The sampling ranges and resolutions along each ray are adjusted for each environment due to their different dimensions. The sampling range is set to 9 m, 15 m, and 24 m for the

conference room, bedroom, and office, respectively. For the conference room and bedroom, we take 128 samples per ray for coarse stratified sampling, and another 128 samples per ray in hierarchical sampling to enhance granularity. For the office environment, due to its complexity, we double the number of samples for both sampling stages.




Figure 3.12: NeWRF’s performance on adding noise to the channel




Figure 3.13: **Subcarrier Channel Prediction**

Subcarrier Channel Prediction

Modern wireless communication systems, such as WiFi and 5G, use Orthogonal Frequency Division Multiplexing (OFDM) for modulation. OFDM divides the spectrum bandwidth into multiple narrow frequency bands, termed subcarriers. Each subcarrier modulates one data symbol. To demodulate data, the receiver has to perform channel estimation for all subcarriers. For instance, a 20 MHz WiFi channel is typically divided into 64 subcarriers, with each subcarrier 312.5 kHz apart from its neighbors. The channels of 52 (out of the 64) subcarriers are typically measured and reported by commercial off-the-shelf WiFi chips [MZW19], known as channel state information (CSI).

NeWRF is able to predict the subcarrier channels accurately using a model trained only on a single subcarrier without further fine-tuning. To demonstrate this capability, we use a model trained on the conference room dataset with a carrier frequency of 2.412 GHz to

infer the channels for its nearby subcarriers. To synthesize the channel at a specific subcarrier frequency, we simply change the signal frequency f in the channel synthesis algorithm (Eq. (3.7)) to the subcarrier frequency $f_{sc} = f_c + k \times \Delta f$ Hz, where k is the index of the subcarrier, ranging from -26 to 26; Δf is the subcarrier spacing, which is equal to 3.125 kHz for 802.11 a/g/n; and f_c is the carrier frequency of the channel. Figure 3.12a shows the predicted amplitude and phase versus the subcarrier indices at five different receiver locations (Pt.0~ 4). As shown, the predictions align very well with the ground truth. Figure 3.12b shows a different view of the same data as Figure 3.12a, where we plot the real and imaginary parts of the subcarriers on the x-axis and y-axis, respectively. In particular, we note that the prediction accuracy remains high even for highly frequency-selective channels. For instance, the channel amplitude of Pt.0 varies by orders of magnitude across different subcarriers. And NeWRF is still able to predict the CSI accurately.

The reason behind this capability of NeWRF is that the model learns to decouple the attenuation factor and phase rotation due to the reflection from that of free-space propagation. The model predicts only the radiated signal, which is nearly unchanged across the subcarrier frequencies. The dominant cause of the varying patterns across subcarriers is the phase shift during free-space propagation and constructive/destructive multipath interference, which are explicitly modeled by the channel synthesis algorithm.

Impact of Noise

Here, we are interested in seeing the robustness of NeWRF against the influence of noise in wireless measurements. In this experiment, we add artificial Gaussian white noise to the simulated channels. The ratio between the channel power and added noise power is denoted as channel SNR. Figure 3.12 shows the results. The blue line shows the results of adding fixed noise, i.e., bias, to the ground-truth channel in the training set. As expected, prediction accuracy decreases as channel SNR decreases. However, typical wireless systems today achieve more than 20 dB channel SNR in practice for reliable communication, where

NeWRF performs well. Also, we note that collecting multiple measurements with different biases can improve robustness. The orange line in Figure 3.12 demonstrates this effect, where we use measurements with different additive noise in each training iteration. The model is able to maintain good performance even at very low SNRs. This is because the model implicitly averages these measurements to reduce the bias.

CHAPTER 4

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

The rapid proliferation of consumer electronics, including smartphones, mixed reality headsets, and autonomous vehicles, has driven increasing demands for high-throughput, low-latency wireless connectivity. Millimeter wave (mmWave) technology [RXM17], operating across the 30–300 GHz frequency bands, offers extensive bandwidth that enables multi-gigabit data rates and sub-millisecond latency.

However, mmWave signals attenuate rapidly with distance and are easily blocked by everyday objects, limiting reliable communication range. To compensate, mmWave systems concentrate energy into narrow directional beams. This directionality demands precise alignment between the transmitter and receiver. In practice, the optimal beam directions are highly dependent on the physical environment: the layout of the space and the material composition. Conventional beam alignment algorithms assume no prior knowledge of the environment and must scan every possible direction, incurring significant overhead. Although many fast beam alignment methods have been proposed [HAR18, LZZ12, TPA11, YAD15, WZC21], they rarely exploit environmental priors to accelerate beam training.

Recent advances in 3D computer vision and neural representation have inspired researchers to model wireless radiation fields in the environment and use them for predicting wireless channels. There are two main lines of research:

Data-driven methods, such as [ZAP23, YDJ25, WTH25b, WTH25a, WXZ25, SZW25], adapt NeRFs and 3D Gaussian Splatting to the radio-frequency domain, implicitly encoding the radiation field in neural network weights or Gaussian parameters. Without prior knowledge of the scene geometry, these methods require dense RF measurements that are hard to obtain in practice. Recent works [CFS24, ZSB25] incorporate visual priors from cameras and LiDARs to relax this requirement, but the resulting representations remain implicit, coupling antenna pattern, polarization, path loss, and material properties in ways that limit explainability and adaptability to new transmitter/receiver configurations.

Simulation-driven methods [OKK23, HAC24, JCW25, ASK25], by contrast, decompose the scene into interpretable components: geometry, material properties, and transmitter/receiver configurations, and employ ray tracing to simulate radio propagation. This decomposition enables network operators to test diverse hypothetical configurations, making it well-suited for deployment planning. We follow the simulation-driven approach.

To bridge the gap between simulation and reality, simulation-driven methods require calibrating scene parameters, particularly the radio material properties. Early works manually assign material parameters and tune them using closed-form or simulated annealing optimization [KR23, CLT20, PJK12], which does not scale to large, complex scenes. The emerging paradigm is to make the ray-tracing pipeline differentiable to enable gradient-based optimization [HCA22], which is more efficient and scalable with GPU acceleration. This approach has been applied and verified at the sub-6GHz band [HAC24]. However, mmWave systems differ significantly from sub-6GHz systems, as mmWave employs large antenna arrays and exhibits vastly different signal scattering behavior. Existing differentiable calibration frameworks are not designed for these characteristics, leaving differentiable calibration for mmWave an open problem.

In pursuing this direction, we encountered a key practical challenge: geometry noise in 3D reconstructed models. Prior differentiable calibration works assume a high-quality CAD model that faithfully replicates the real-world scene, but such models are rarely available in

practice. While advances in 3D reconstruction technology have made scene modeling more accessible, the resulting models often exhibit flaws such as holes and noisy surfaces. These flaws cause critical propagation paths to be dropped or spurious paths to appear, degrading calibration performance. Manually fixing such flaws requires professional artists and days of effort, which is not scalable. We therefore ask: *Can we design the simulator itself to be robust to geometry noise?* Our key insight is that this is particularly tractable at mmWave frequencies. At sub-6GHz, specular reflection dominates, and geometry noise is difficult to handle – a small surface normal error directly translates to a large error in the reflected ray direction. At mmWave, however, scattering is stronger, creating an opportunity: we approximate the noise-sensitive specular reflection with a differentiable directional scattering model that smoothly distributes power among nearby ray directions, naturally absorbing the effect of geometry noise.

In this chapter, we present mmDiff, the first end-to-end differentiable ray-tracing framework for mmWave channel calibration, enabling gradient-based optimization of radio material properties at scale. Our main contributions are:

- We propose an end-to-end optimization framework for mmWave calibration using sparse mmWave channel measurements, bridging the gap in differentiable mmWave calibration.
- We introduce a differentiable approximation of the specular reflection. We prove that the approximation preserves the asymptotic accuracy of path-gain estimation.
- We evaluate mmDiff on real-world datasets and observe a 10.5 dB improvement in average received power prediction accuracy over the prior differentiable solver using specular reflection.




Figure 4.1: **AoA Power Spectra.** a) The 3D view of an AoA power spectrum, overlying a 4x4 antenna array. b) The 2D polar plot of the AoA power spectrum. The color indicates the received signal power from each direction.

4.1 Preliminaries

4.1.1 Angular Power Spectra

mmWave radios create highly directional beams to concentrate energy toward intended directions and overcome high path loss. Understanding how signal power arrives from different directions is therefore essential for characterizing channel behavior, optimizing coverage, and designing robust beam alignment strategies.

The angle-of-arrival (AoA) power spectrum provides this directional view by describing how much received signal power arrives from each direction in space, revealing the spatial structure of the wireless channel. Crucially, the AoA power spectrum is practically measurable: it can be obtained with a phased array or a mechanically rotated horn antenna that scans across directions, making it well-suited for mmWave scene calibration.

4.1.2 Physically-based rendering (PBR)

Physically-based rendering (PBR) is a computer graphics technique that attempts to simulate how light behaves in the real world for rendering a view. It simulates light rays propagating and bouncing in the scene to produce realistic images. Recent efforts have made




Figure 4.2: **mmDiff’s Learning Framework for Scene Calibration.** The inputs are the initial scene configuration and the AoA power spectrum measurements. The outputs are the calibrated scene configuration and the predicted AoA power spectrum.

the whole ray-tracing pipeline end-to-end differentiable, enabling the optimization of scene parameters from images, known as inverse rendering [JSR22b]. This capability has been widely used in various tasks, including 3D reconstruction [NJJ21] and material or illumination estimation [CLG21]. Recently, this concept has been extended to the wireless domain, represented by NVIDIA’s Sionna [HCA22], a differentiable ray-tracing engine for wireless communication. Sionna builds the wireless PHY and MAC layers on top of the differentiable rendering engine [JSR22a], providing GPU-accelerated simulation for wireless coverage estimation. The latest version of Sionna implements two solvers for the radio map and channel impulse response (CIR), respectively. In this chapter, we propose a new physics-grounded solver for mmWave channel simulation, enabling end-to-end calibration from the angular spectrum.

4.1.3 Physics of Wireless Ray-Tracing

Wireless ray tracing applies the ray approximation of electromagnetic (EM) propagation to predict how radio waves propagate through complex environments, incorporating wireless-specific physics such as antenna radiation patterns, array beamforming, polarization, phase evolution, and frequency-dependent material responses. Its goal is to compute the elec-

tromagnetic field received at a specific RX by modeling how rays depart the transmitter, interact with surfaces via reflection, diffraction, and scattering, and arrive at each receiver antenna element with the correct amplitude, phase, and polarization.

Following Sionna [HCA22], we adopt a physically grounded formulation for modeling these components. In particular, the total field arriving at the receiver depends on four factors: a) the radiation patterns of antennas; b) the beamforming weights of antenna arrays; c) the path gain of free-space propagation; and d) the change of signal amplitude, phase, and polarization due to the ray-surface interactions. We provide a brief formulation for some of these components below:

The radiation pattern of antennas is modeled as complex spherical functions:

$$\mathbf{C}(\theta, \varphi) = C_\theta(\theta, \varphi)\hat{\boldsymbol{\theta}} + C_\varphi(\theta, \varphi)\hat{\boldsymbol{\varphi}}, \quad (4.1)$$

where $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\varphi}}$ are unit basis vectors of the spherical coordinate system; $C_\theta(\theta, \varphi)$ and $C_\varphi(\theta, \varphi)$ are complex values, modeling the elevation and azimuth pattern of the antenna, respectively. The pattern can be any physically valid spherical function satisfying energy conservation [Bal16].

Under the far-field plane-wave assumption [HCA22], an antenna array can be treated as a single effective antenna with an aggregated radiation pattern. Rays are emitted from the array center, and beamforming gain is applied through the synthetic array weight:

$$\mathbf{W}_{\text{TX}}(\mathbf{k}) = \sum_{n=1}^N w_n \cdot e^{j\frac{2\pi}{\lambda}\mathbf{d}_n \cdot \mathbf{k}}, \quad (4.2)$$

where N is the number of TX antenna elements; $\mathbf{k}$ is the direction of departure vector; w_n is the precoding weight of the n -th antenna element to steer the TX beam direction; $\mathbf{d}_n$ is the relative position of the n -th antenna element with respect to the center of the array; λ is the signal wavelength.

The free-space propagation and field transformation due to ray-surface interaction can

generally be represented by a transformation matrix $\mathbf{T}$ that accounts for changes in amplitude, phase, and polarization during propagation. Hence, the E-field of a single ray i , reaching an RX antenna element, can be expressed as:

$$\mathbf{E}_i = \frac{\lambda}{4\pi} \mathbf{C}_T \mathbf{W}_{TX}(\mathbf{k}_i^{\text{AoD}}) \mathbf{T}_i \mathbf{C}_R^H, \quad (4.3)$$

where $\frac{\lambda}{4\pi}$ is a scaling factor accounting for the antenna aperture; $\mathbf{C}_T$ and $\mathbf{C}_R$ are short for $\mathbf{C}_T(\theta_T, \varphi_T)$ and $\mathbf{C}_R(\theta_R, \varphi_R)$ modeling the radiation patterns of TX and RX antennas. The final wireless channel is the accumulation of E-fields along all paths to the receiver antenna.

4.2 Related Work

4.2.1 Wireless Channel Prediction

Predicting the spatial distribution of the wireless channel is challenging. Recent studies employ deep learning techniques inspired by advances in computer vision to tackle this problem. There are two general approaches to this problem:

Data-Driven: [ZAP23, YDJ25, WTH25b, WTH25a, WXZ25, SZW25] adapt NeRFs and 3D Gaussian Splatting to the radio frequency domain by exploiting the shared propagation principles of visible light and wireless signals. These methods face three practical limitations. First, they require extremely dense datasets (6,216 samples/ m^3 [ZAP23]), where training and test samples are nearly homogeneous, which inflates apparent generalization performance. Second, the learned fields are encoded implicitly in neural network weights and cannot inform wireless network deployment decisions. While 3DGS-based methods use ellipsoidal pods to represent the scene, these works do not learn an explicit representation for the scene geometry, but use GS parameters to hold RF radiation fields. Third, the learned fields are fixed after training and cannot be adapted to new transmitter or receiver configurations, limiting their utility for evaluating alternative deployment plans. Recent

works [CFS24, ZSB25] solve the problem by incorporating visual priors from cameras and LiDARs to assist the reconstruction of radio radiation fields, achieving significantly better performance. However, they still face the challenge of producing physically rigorous channels and generalizing to new TX/RX setups.

Simulation-Driven: [OKK23, HAC24, JCW25, ASK25] build on the traditional wireless ray-tracing simulation and physics-based rendering pipelines, and integrate neural-based modules to bridge the gap between simulation and reality. Given the scene geometry provided as input, these methods are promising for complex, large-scale scenes. The key advantage is that they can easily generalize to new scene setups thanks to the decoupling of scene geometry, radio device setups, and materials. However, to adopt these methods in practice, one needs to prepare a detailed reconstruction of the 3D environment. Although 3D reconstruction technologies have advanced significantly in recent years and are becoming more accessible, reconstruction errors are often inevitable due to imperfections in hardware and algorithms. We found through our experiments that this noise is harmful to traditional ray-tracing pipelines, as it causes incorrect path predictions and thus poor performance in scene representation learning.

4.3 mmDiff Design

We consider a practical application scenario, shown in Figure 4.2, where a network operator surveys a deployment environment using an RGB-D camera to scan the scene and 3D reconstruction techniques to build a mesh model. Simultaneously, the operator measures received signal power at different angles and beam configurations using an antenna array to calibrate material parameters. These parameters are represented implicitly by a neural network that maps 3D ray-surface intersection coordinates to material properties such as conductivity and permittivity. Predicted AoA power spectra are compared with physical measurements to compute the loss. Since mmDiff’s ray-tracing pipeline is fully differentiable, the loss gradi-

ents can be back-propagated to update neural network parameters. This process is repeated until convergence, yielding a calibrated scene representation that closely matches the real environment. The calibrated model can then predict channels for arbitrary TX-RX locations and antenna patterns without retraining, as predictions are computed dynamically.

Next, we detail mmDiff’s ray-tracing algorithm and noise-robust ray-surface interaction model, the core components enabling accurate and differentiable channel prediction.

4.3.1 mmDiff’s Ray-Tracing Algorithm

We introduce a pure Monte Carlo-based ray-tracing algorithm to compute the AoA power spectrum. It synthesizes received signal power at each RX antenna element by integrating contributions from many rays reaching the antenna. The algorithm comprises two tightly integrated components: path tracing, which finds signal propagation paths connecting transmitter and receiver, and power integration, which integrates EM power to produce the AoA power spectrum. We discuss each component in turn.

Path Tracing We adopt a backward ray tracing approach, where rays are launched from RX to TX to ensure signal power is captured from all receiving directions. For each ray emitted from the RX, we find its intersections with the environment mesh. At each intersection, we (i) perform a visibility test to the TX and, when unobstructed, spawn a deterministic ray toward the TX; and (ii) sample a next-hop direction according to the event type (reflection, diffusion, or refraction) and the surface scattering pattern. For highly directional scattering patterns, we use importance sampling to concentrate samples in high-gain directions, reducing simulation variance. We trace rays recursively until reaching a preset maximum depth.

Power Integration We integrate the received signal power at each receiver antenna element. We apply an additional phase shift based on the receiving direction of the ray, $\mathbf{k}$, and the relative position of the antenna element in the array, $\mathbf{d}_m$, where m is the index of the antenna

element:

$$\mathbf{W}_{\text{RX}}^m(\mathbf{k}) = e^{-j\frac{2\pi}{\lambda}\mathbf{d}_m \cdot \mathbf{k}}. \quad (4.4)$$

The Monte Carlo integration equation for the received signal power at RX antenna element m is:

$$\begin{aligned} P_m &= \int_{\Omega} \left| \frac{\lambda}{4\pi} \mathbf{C}_T \mathbf{W}_{\text{TX}}(\mathbf{k}^{\text{AoD}}) \mathbf{T} \mathbf{C}_R \right|^2 \mathbf{W}_{\text{RX}}^m(\mathbf{k}^{\text{AoA}}) d\Omega \\ &\approx \sum_{i=1}^L \left| \frac{\lambda}{4\pi} \mathbf{C}_T \mathbf{W}_{\text{TX}}(\mathbf{k}_i^{\text{AoD}}) \mathbf{T}_i \mathbf{C}_R^H \right|^2 \mathbf{W}_{\text{RX}}^m(\mathbf{k}_i^{\text{AoA}}), \end{aligned} \quad (4.5)$$

where Ω is the set of sampled angles; L is the number of ray directions we sampled uniformly on the receiver's sphere.

Once we have the received signal power at each RX antenna element, we can convert it to the AoA power spectrum by weighted summation with the steering vector of the RX antenna array:

$$P(\theta, \varphi) = \frac{1}{\sqrt{M}} \sum_{m=1}^M P_m \cdot e^{-j\frac{2\pi}{\lambda}\mathbf{d}_m \cdot \mathbf{k}(\theta, \varphi)}, \quad (4.6)$$

where M is the number of RX antenna elements. We compute $P(\theta, \varphi)$ for a mesh grid of angles θ and φ covering the upper hemisphere to generate the AoA power spectrum shown in Figure 4.1.

4.3.2 Noise-Robust Ray-Surface Interaction

So far, we have explained mmDiff's ray-tracing algorithm for generating the AoA power spectrum. Next, we discuss how we make the ray-tracing pipeline noise-robust and fully differentiable via a novel ray-surface interaction model.

Four major types of interaction events are considered in common wireless simulators [HCA22]: specular reflection, diffuse scattering, refraction, and diffraction. Diffraction is often ignored




Figure 4.3: **Issue with Pure Specular Reflection.** Small discrepancies in the reconstruction model cause significant differences in the simulated paths.

at mmWave frequencies due to its negligible contribution to overall path gain [SSS22]. Specular reflection is the dominant propagation mechanism, contributing over 90% of received power. However, specular reflection follows the ideal law of reflection, meaning incident and reflection angles are identical, so the outgoing direction is represented by a delta distribution. This makes ray tracing highly sensitive to geometric noise and non-differentiable with respect to surface normals. Even small perturbations in surface orientation can cause large deviations in reflection direction, leading to significant path discrepancies.

The sensitivity arises because the reflection lobe width depends on the ratio between the signal wavelength and the surface roughness [TPP11]. For lower-frequency radios (sub-6 GHz), the wavelength is much larger than typical material roughness, so most surfaces appear mirror-like, with energy concentrated in the specular direction. At mmWave frequencies, however, the wavelength is comparable to the roughness of common building materials such as concrete, marble, and brick. Recent measurements confirm that reflections at mmWave exhibit narrow but nonzero angular spread rather than an ideal delta peak [GZS25]. This motivates our directional scattering approximation, which replaces ideal specular reflection with a narrow-lobed scattering distribution. The resulting pattern preserves physical realism while maintaining differentiability of the ray-tracing pipeline with respect to geometry.

We adopt the directional scattering model from [DFV07], though any narrow-lobed formulation satisfying energy conservation applies. The scattering pattern is defined as:




Figure 4.4: Single-reflector scene for Figure 4.5: Directional-scattering the directional scattering path gain path gain converges to specular re-convergence test.

reflection path gain as the scattering footprint grows and directionality increases.

$$f(\psi) = F_{\alpha_r}(\theta_i)^{-1} \left(\frac{1 + \cos \psi}{2} \right)^{\alpha_r}, \quad (4.7)$$

where ψ is the deviation angle from the specular direction; $F_{\alpha_r}(\theta_i)^{-1}$ is a normalization factor that enforces energy conservation, depending on incident angle θ_i ; and α_r controls the lobe beamwidth, with larger values producing narrower, more directional patterns. Rather than relying on a single reflection point, directional scattering aggregates contributions from many nearby points to form a “heat spot”, averaging reflected power and angles and making the model robust to geometric noise. Our implementation uses 64-bit arithmetic to support ultra-narrow beam computation. The half-power beamwidth can be as narrow as 1.40° .

Since we approximate specular reflection using directional scattering, we next establish a theorem showing that this approximation preserves the path gain:

Theorem 1. *Directional-Scattering Approximation:* For a planar reflector with scattering pattern $f_s(\psi)$, and reflection coefficient Γ , when the ratio of scattered field amplitude

to incident field amplitude is given by the following equation: ¹

$$\frac{|E_{scat}|}{|E_{inc}|} = \Gamma \sqrt{f_s(\psi)}, \quad (4.8)$$

The directional-scattering path gain asymptotically converges to the specular reflection path gain as the scattering footprint grows sufficiently large and the scattering lobe becomes sufficiently narrow.

We refer readers to Supplemental Material Section 4.6 for the full proof of this theorem. The intuition is that sufficiently large reflectors capture almost all power radiated from the transmitter’s lower hemisphere. The captured power is then re-radiated towards the receiver, weighted by the scattering pattern. The highest weights are contributed by the scattering points near the specular reflection point. In the extreme case, when the scattering pattern collapses to a Dirac delta function, it matches the case of specular reflection. It is also worth noting that the classic Fresnel equations [Bal16], which are widely used to compute specular reflection gain in wireless simulators, are derived under the assumption of an infinitely large planar interface. This assumption rarely holds in real environments with finite-sized reflectors. Our approximation smoothly bridges the transition from idealized infinite surfaces to realistic finite-scale reflectors, resulting in behavior that more closely matches real-world scenarios.

In practice, the scattering footprint does not have to grow infinitely large to reach a good approximation. The convergence rate depends on the directionality of the scattering pattern. We set up a simple experiment to demonstrate this behavior. As shown in Figure 4.4, we put a pair of TX and RX in a scene with a single planar reflector. We put the TX and RX symmetrically on the same side of the reflector. We use Sionna’s path solver to find the specular reflection path gain as a reference. We then use mmDiff to simulate the gain of the directional scattering path for different reflector sizes and scattering-pattern lobewidths.

¹This equation is simplified here by eliminating the cross-polar discrimination coefficient, for the sake of clarity. The actual implementation includes an additional transform matrix to account for the polarization.

Figure 4.5 shows the result. As shown, all curves converge to the specular path gain as the reflector scale increases. The more directional the scattering pattern, the faster the convergence. More analysis about the approximation behavior can be found in Supplemental Material Section 4.6.

4.3.3 Optimization Framework

We use a fully-connected neural network with four hidden layers of 64 neurons each to parameterize the material properties in the reconstructed mesh. The neural network is evaluated at the 3D coordinates of each ray-surface intersection point (with 6-octave positional encoding [MST21]) to predict the permittivity, conductivity, scattering coefficients, and cross-polar discrimination coefficients of the surface. The predicted material parameters are fed into physics-based equations to compute the scattered fields. The scattered fields are eventually integrated to compute the predicted AoA power spectra. We adopt the log mean absolute error (log-MAE) as the loss function for comparing the predicted and measured AoA power spectra, defined as:

$$L_{\log\text{-MAE}} = \frac{1}{K} \sum_{i=1}^K |\log P_i^{\text{pred}} - \log P_i^{\text{meas}}|, \quad (4.9)$$

where the AoA power spectrum is represented as a real-valued matrix of size $K = N_{el} \times N_{az}$, with N_{el} and N_{az} denoting the numbers of elevation and azimuth angles, respectively. We select log-MAE because it effectively captures errors across a wide dynamic range of power values, which is common in mmWave channels.

4.3.4 Implementation

We implement mmDiff as a fork of Sionna v1.2.1 [HCA22], which is built on Mitsuba v3.7.1 [JSR22b] and Dr.Jit v1.2.0 [JSR22a]. We build a new stand-alone ray-tracing solver for synthesizing the AoA power spectrum, and a customized radio material implementation

that integrates the directional scattering approximation and neural network material parameterization. We keep the original diffuse scattering unchanged but only replace the specular reflection with the directional scattering model. The scattering coefficient is now used to account for the power distribution between diffuse scattering and directional scattering. We test the simulator on an NVIDIA RTX 3090 GPU and observe 45 fps for training and 120 fps for inference with one million rays and up to two bounces, satisfying the real-time inference requirement.

4.4 Results

We conduct both real-world and simulation experiments to evaluate mmDiff.

4.4.1 Real-world Experiments

We constructed a real-world mmWave testbed in an indoor laboratory environment. A Keysight M9384B vector signal generator serves as the transmitter, producing a continuous-wave 28 GHz sinusoidal signal that is radiated via a Mi-Wave 261(34)-20/595 horn antenna. On the receiver side, an N9021B spectrum analyzer captures the incoming signal through an identical horn antenna, which is mechanically rotated in the azimuth plane to measure the angular power spectrum across eight equally spaced azimuth angles, as illustrated in Figure 4.6(a).

The 3D environment is captured using the PolyCam app on an iPad, from which RGB-D frames and camera trajectories are exported and processed via Neural RGB-D Reconstruction to generate a 3D mesh model of the scene (Figure 4.6(b)). Received signal power measurements are collected at 29 sparsely distributed locations, whose spatial distribution is shown in Figure 4.6(c), of which 23 locations (80%) are reserved for training and the remaining 6 (20%) for testing. The mmDiff simulator is then configured to match the physical testbed parameters, including antenna radiation patterns, transmit power, and system losses.




Figure 4.6: **Real-World Testbed Setup.**




Figure 4.7: **Results of the Real-World Experiment.**

We train mmDiff for 20 epochs. The directional scattering lobe parameter is set to $\alpha_r = 300$, corresponding to a half-power beamwidth of 4.8° . We use Sionna v0.19.1’s path solver with line-of-sight (LoS) and specular reflection enabled as a baseline, and follow the procedure in [HAC24] to perform calibration.

Figure 4.7(a) presents the predicted AoA power spectrum at a representative test position, where mmDiff demonstrates substantially closer agreement with the ground truth compared to Sionna. Figure 4.7(b) shows the cumulative distribution function (CDF) of the absolute prediction error over the evaluation set ($6 \text{ test positions} \times 8 \text{ angles}$), where mmDiff achieves a mean absolute error of 4.60 dB, compared to 15.11 dB for Sionna — a reduction of more than 10 dB. The primary source of Sionna’s degraded performance is noisy geometry reconstruction, which introduces errors in both path identification and power calibration. In contrast, mmDiff’s directional scattering model is inherently robust to geometric noise, enabling reliable recovery of dominant propagation paths and yielding significantly more accurate power predictions. These results validate the practical effectiveness of mmDiff in real-world deployment scenarios, where geometrically perfect reconstructions are rarely attainable.



Figure 4.8: **Replica [SWM19] Scenes Used in Evaluation.**

4.4.2 Simulation Experiments

Datasets We evaluate mmDiff on four realistic 3D indoor scenes from the Replica dataset [SWM19] (Figure 4.8), using the original meshes as ground truth (GT) with ITU materials [Ser15] assigned based on semantic labels. For each scene, we create a camera trajectory in Blender and render RGB-D images using the Replica Renderer; each trajectory contains 2800 poses covering the majority of the scene area. These poses also serve as receiver locations in a wireless simulator to generate the corresponding mmWave channel datasets (TRAJ). We apply Neural RGB-D Reconstruction [AMG22] to recover scene geometry from the RGB-D images, producing reconstructed meshes (REC), and manually remove floaters introduced by the algorithm. This process introduces reconstruction noise similar to real-world captures. We further sample ~ 500 random receiver poses per scene as a held-out test set (RAND) to probe diverse viewpoints. For all datasets, the TX is fixed at a single pose. Both transmitter and receiver use a 4×4 antenna array; simulations run at 28 GHz with one million rays traced to a maximum depth of one. We empirically choose $\alpha_r = 100$ for all simulation experiments. Further dataset details appear in Supplemental Material Section 4.6.

Baselines We use Sionna v0.19.1’s path solver with line-of-sight (LoS) and specular reflection enabled as a baseline. Sionna traces diffuse scattering paths; however, these paths are used to predict the Channel Impulse Response (CIR), rather than to approximate specular path power. Hence, we disabled diffuse scattering for Sionna in our experiments. Sionna does not directly produce AoA spectra. We apply the same beamforming procedure described in Equation (4.5) and Equation (4.6) to obtain comparable AoA spectra from Sionna’s traced

Table 4.1: Robustness under Geometry Noise (TRAJ Dataset).

Scene	Methods	Top-1 Peak			Top-5 Peaks		
		F1 $\uparrow$	AAE $\downarrow$	APE $\downarrow$	F1 $\uparrow$	AAE $\downarrow$	APE $\downarrow$
Office 0	Sionna	0.8046	3.15	1.27	0.7431	3.30	1.93
	mmDiff	0.9954	0.29	0.14	0.9804	0.38	0.18
Office 4	Sionna	0.9546	2.00	0.79	0.8886	1.74	0.98
	mmDiff	0.9950	0.52	0.11	0.9774	0.62	0.16
Room 0	Sionna	0.9707	1.71	0.97	0.9133	1.64	1.02
	mmDiff	0.9936	1.33	0.55	0.9751	1.07	0.62
Room 1	Sionna	0.8086	3.89	1.75	0.7710	3.14	2.16
	mmDiff	0.9925	0.26	0.07	0.9879	0.28	0.10

paths.

Experiment 1: Robustness to Geometry Noise




Figure 4.9: **Robustness to Geometry Noise.** Sionna exhibits significant deviation in the REC environment while mmDiff performs consistently despite geometry noise.

We first evaluate the robustness of mmDiff to geometry noise that is commonly observed in 3D reconstruction.

In this experiment, we run both simulators on the ground-truth (GT) and reconstructed (REC) meshes across all four Replica scenes. To exclude the effect of material estimation, we assign the same uniform material to all scene objects. We generate AoA power spectra on REC and GT meshes along the same TRAJ datasets. Figure 4.9 illustrates the results for one of the receiver poses in Office 0. As shown, on the GT geometry, mmDiff and Sionna produce similar spectra; however, on the REC geometry, mmDiff closely matches the GT results while Sionna exhibits substantial deviations.

For quantitative evaluation, we detect dominant peaks in the AoA power spectra using a maximum-filter-based peak detection algorithm. We extract the top-1 and top-5 peaks and match them between the REC and GT spectra based on angular proximity (within 20°). We report the F1-score for peak detection accuracy and compute the Average Angle Error (AAE) in degrees and Average Power Error (APE) in dB for matched peaks. We choose these metrics as they directly reflect the accuracy of angle and power estimation for practical applications such as beam prediction. Table 4.1 summarizes the results: across all scenes and metrics, mmDiff exhibits almost perfect peak detection accuracy and only minor angle and power error, consistently outperforming the Sionna baseline. This demonstrates the robustness of mmDiff against geometry noise, a critical feature for real-world applications.

Experiment 2: Scene Calibration and Channel Prediction We run end-to-end experiments to validate that mmDiff’s robustness and differentiability benefit the material calibration and channel prediction. For Sionna and mmDiff, we separately generate synthetic datasets on the GT mesh as ground truth channels for supervision. During training, we use the ray-surface interaction points traced in the REC meshes to query the neural material model. We include another data-driven baseline approach, NeRF² [ZAP23], for this channel prediction task, which adopts NeRF-like ray tracing to learn the radiation fields of a wireless environment. The synthetic AoA power spectra generated by Sionna are used to

Table 4.2: Prediction Accuracy on RAND Dataset.

Scene	Methods	Top-1 Peak			Top-5 Peaks		
		F1 $\uparrow$	AAE $\downarrow$	APE $\downarrow$	F1 $\uparrow$	AAE $\downarrow$	APE $\downarrow$
Office 0	NeRF ²	0.0329	12.86	8.96	0.2082	13.50	8.51
	Sionna	0.6708	5.03	3.94	0.6794	3.88	4.52
	mmDiff w/o Calib.	0.6249	4.56	3.98	0.6845	4.02	4.97
	mmDiff Full	0.9168	1.84	1.08	0.9121	1.75	1.22
Office 4	NeRF ²	0.0460	14.28	3.15	0.0816	12.74	4.06
	Sionna	0.8154	4.16	2.03	0.7554	3.47	2.42
	mmDiff w/o Calib.	0.5643	8.16	7.15	0.6197	5.15	7.44
	mmDiff Full	0.9257	2.08	0.97	0.8994	1.58	1.11
Room 0	NeRF ²	0.2611	12.85	0.91	0.3822	9.38	3.32
	Sionna	0.9087	3.16	1.94	0.8345	2.71	2.11
	mmDiff w/o Calib.	0.6505	6.80	2.92	0.6948	4.71	3.65
	mmDiff Full	0.9534	1.36	0.56	0.9441	1.26	0.64
Room 1	NeRF ²	0.1868	12.29	2.43	0.3186	9.41	4.14
	Sionna	0.7859	5.39	4.34	0.7422	4.15	4.34
	mmDiff w/o Calib.	0.7473	4.59	2.10	0.7537	3.86	2.85
	mmDiff Full	0.9342	1.58	0.92	0.9287	1.43	0.95

train and test NeRF² models. We train all models using poses along a designated trajectory (TRAJ) and test on poses randomly sampled in the scene (RAND). More details about the training procedure can be found in Supplemental Material Section 4.6.

The results of this experiment are shown in Table 4.2: NeRF² fails to generalize under such sparse sampling density due to the lack of explicit geometric priors; mmDiff consistently outperforms both baselines across all scenes and metrics, achieving the highest accuracy in mmWave channel prediction. As an ablation study, we also report results for mmDiff without calibration, which perform poorly and underscore both the necessity of material calibration and the effectiveness of mmDiff. We also note that the overall performance on the RAND dataset is slightly lower than in Table 4.1, as it includes more diverse receiver locations where geometric reconstruction is imperfect due to occlusions. Nevertheless, mmDiff maintains high

prediction accuracy, sufficient for practical deployment scenarios.

Additional Results: We include more experiments and results in the Supplemental Material Section 4.6.

4.5 Discussion

- **Geometry optimization:** mmDiff is differentiable with respect to the scene geometry, which has the potential application of calibrating scene geometry. We demonstrate this capability in Supplemental Material Section 4.6, where we calibrate the translation and rotation of a planar reflector. However, calibrating the full geometry mesh is substantially harder and beyond this work’s scope; we believe this is an interesting direction for future exploration.
- **Dynamic scenarios:** mmDiff extends naturally to dynamic environments. In practice, large-scale reflectors, such as walls, ceilings, and building facades, are static and require no retraining as the environment evolves. For dynamic objects such as humans or moving furniture, visual tracking (e.g., from RGB-D cameras or LiDAR) can continuously update their positions and geometry within the scene representation. Leveraging mmDiff’s fast inference, these position updates translate directly into real-time channel predictions without re-running the full simulation pipeline, making mmDiff well-suited for dynamic scenarios.
- **Phase Errors:** mmDiff integrates received signal power over directions via Monte Carlo sampling, discarding per-ray phase to ensure accurate power integration. This may produce inconsistent interference patterns when multiple paths carry comparable strength, as we show in Supplemental Material Section 4.6. However, in most real-world scenarios, mmWave transceivers use very narrow beams; therefore, only a single path dominates the power compared to the other paths. Hence, the impact of phase

error is negligible, and the calibration is still functional.

4.6 Conclusion

In conclusion, we propose the directional scattering approximation of specular reflection to handle the inevitable geometry imperfections in 3D reconstructed models, enabling robust and differentiable simulation of mmWave channels. We build mmDiff, the first noise-robust, fully differentiable framework for mmWave scene calibration and channel prediction. Leveraging end-to-end differentiability, mmDiff learns scene-specific material parameters in geometrically complex environments through gradient-based optimization, eliminating the need for manual parameter tuning. Overall, these properties position mmDiff as a practical and deployable solution for real-world mmWave channel modeling, with applications spanning base station coverage prediction, network deployment planning, and beam training overhead reduction in next-generation wireless networks. Looking ahead, promising directions for future work include scaling mmDiff to larger outdoor environments, extending to dynamic scenes with continuous online learning, and integrating mmWave sensing capabilities to enable joint communication and sensing.

Appendix

Proof of Theorem 1

We prove the following equation holds:

$$\lim_{\alpha_r \rightarrow \infty} \lim_{S \rightarrow \infty} P_{\alpha_r}(S) = \left(\frac{\lambda \Gamma}{4\pi r_{\text{spec}}} \right)^2, \quad (4.10)$$

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




Figure 4.10: Geometry setting for proof of Theorem 1.




Figure 4.11: Illustration of the effective area dS .

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

$$d\Omega_i = \frac{\cos \theta_i dS'}{r_i^2}, \quad (4.11)$$

where θ_i is the incident angle to the scattering point.

Hence, the differential incident power is

$$dP_{\text{inc}} = \frac{1}{4\pi} d\Omega_i = \frac{\cos \theta_i}{4\pi r_i^2} dS' \quad (4.12)$$

Since the scattered field amplitude is $\Gamma \sqrt{f_s(\alpha_r; \theta_s, \phi_s)}$ times the incident field (Equation (4.8)), the differential scattered power is

$$dP_{\text{sca}} = \Gamma^2 f_s(\alpha_r; \theta_s, \phi_s) dP_{\text{inc}} \quad (4.13)$$

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

$$dP_r = \left(\frac{\lambda}{4\pi r_s} \right)^2 dP_{\text{sca}} , \quad (4.14)$$

Thus, the differential received power is:

$$dP_r = \left(\frac{\lambda\Gamma}{4\pi} \right)^2 f_s(\alpha_r; \theta_s, \phi_s) \frac{\cos \theta_i}{r_i^2 r_s^2} dS' \quad (4.15)$$

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

$$P_{\alpha_r}(S) = \iint_S \left(\frac{\lambda\Gamma}{4\pi} \right)^2 f_s(\alpha_r; \theta_s, \phi_s) \frac{\cos \theta_i}{r_i^2 r_s^2} dS' \quad (4.16)$$

It is essential to note that only a fraction of the total shined area is captured by the receiver's aperture. As illustrated in Figure 4.10, the total area illuminated 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_{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:

$$\frac{dS}{dS'} = \frac{A_{\text{eff}}}{A} = \frac{r_i^2}{(r_i + r_s)^2}, \quad (4.17)$$

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

Plugging dS into the power integration, we get:

$$P_{\alpha_r}(S) = \iint_S \left(\frac{\lambda\Gamma}{4\pi} \right)^2 f_s(\alpha_r; \theta_s, \phi_s) \frac{\cos \theta_i}{r_i^2 r_s^2} \frac{r_i^2}{(r_i + r_s)^2} dS \quad (4.18)$$

To calculate the integral, we perform a change of variable to spherical coordinates (θ_s, ϕ_s) at the scattering point. We use the geometry model in Figure 4.10 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 (θ_s, ϕ_s) . According to the geometry,

$$\mathbf{k}_s = \begin{bmatrix} \sin \theta_s \cos \phi_s \\ \sin \theta_s \sin \phi_s \\ \cos \theta_s \end{bmatrix} = \frac{1}{r_s} \begin{bmatrix} d - x \\ -y \\ h_R \end{bmatrix}, \quad (4.19)$$

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

$$|J| = \det(A) = \begin{vmatrix} \frac{\partial x}{\partial \theta_s} & \frac{\partial x}{\partial \phi_s} \\ \frac{\partial y}{\partial \theta_s} & \frac{\partial y}{\partial \phi_s} \end{vmatrix} = r_s^2 \tan \theta_s \quad (4.20)$$

Apply the change of variables:


$$\begin{aligned} \lim_{S \rightarrow \infty} P_{\alpha_r}(S) &= \int_0^{2\pi} \int_0^{\pi/2} \left(\frac{\lambda \Gamma}{4\pi} \right)^2 f_s(\alpha_r; \theta_s, \phi_s) \frac{\cos \theta_i}{r_i^2 r_s^2} \frac{r_i^2}{(r_i + r_s)^2} r_s^2 \tan \theta_s d\theta_s d\phi_s \\ &= \left(\frac{\lambda \Gamma}{4\pi} \right)^2 \int_0^{2\pi} \int_0^{\pi/2} f_s(\alpha_r; \theta_s, \phi_s) \sin \theta_s \frac{1}{(r_i + r_s)^2} \frac{\cos \theta_i}{\cos \theta_s} d\theta_s d\phi_s \end{aligned} \quad (4.21)$$

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:

$$\lim_{\alpha_r \rightarrow \infty} \lim_{S \rightarrow \infty} P_{\alpha_r}(S) = \left(\frac{\lambda \Gamma}{4\pi} \right)^2 \frac{1}{(r_i^* + r_s^*)^2} = \left(\frac{\lambda \Gamma}{4\pi r_{spec}} \right)^2, \quad (4.22)$$



(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 4.12: Factors that impact the deviation of the directional scattering approximation from the specular reflection model.

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 separate multiplicative factor (Section 4.3.1), the above proof can be easily generalized to arbitrary antenna patterns and polarization.

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.






Figure 4.13: Single planar reflector scene for deviation analysis.



(a) Normalization factor in-creases with large incident an-gles.



(b) Scattering pattern for differ-ent incident angles ($\alpha_r = 120$).

Figure 4.14: Large incident angles cause a higher normal-ization factor, which amplifies the scattering lobe.

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. Figure 4.12a shows the deviation versus the directivity of the scattering lobe for different surface scales. We conduct this experiment in the setup shown in Figure 4.13. We fixed the distance between the TX/RX and the reflection point to 25 m and the incident angle to 45° . 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.

Figure 4.12b 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° . As shown, the deviation increases with distance. This is because, for a fixed surface scale, as the TX/RX moves farther away, the effective area of

the surface illuminated by the TX and visible to the RX decreases, leading to 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.

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 25 m, then swept the incident angle θ_i from 0 to 90°. We moved both TX and RX to keep the incident angle and reflection angle the same, so we could 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 Figure 4.12c. We observe that for small incident angles (up to 40° for $\alpha_r = 10$, and 70° 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°, there is an interesting variation pattern of the directional path gain: the deviation first decreases, then increases again after 70°, and finally decreases again after 80°.

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 Section 4.6, the scattering pattern must satisfy the energy conservation equation (Equation (EC)). In our case, we adopt the scattering pattern model in [DFV07], which has a normalization factor, $F_{\alpha_r}(\theta_i)^{-1}$, (Equation (4.7)) to ensure the energy conservation. We plot the normalization factor versus the incident angle in Figure 4.13a. We observe that the normalization factors share the same pivot point at around 40° for $\alpha_r = 10$ and around 70° for $\alpha_r = 50$ and 120. After the pivot point, the normalization




Figure 4.15: Illustration of the sampling issue that causes the variation with respect to the incident angle (bird's-eye view).




Figure 4.16: Maximum deviation for different distances between TX/RX and reflection point.

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. Figure 4.13b shows the scattering pattern for different incident angles when $\alpha_r = 120$. We can see that as the incident angle increases to 85° , 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 Equation (4.18), 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 Figure 4.15. We also plot the curve of $\theta_i = \theta_i^*$ on the figure, where θ_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, sampling density comes into play, 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 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° to find the maximum deviation. We plot the heatmap of maximum deviation versus the distances in Figure 4.16. 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.

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 5×10^{-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 [HAC24] and train the model for 10,000 iterations. We observe that the learning curves have converged at this point. For the NeRF² baseline, we follow the training procedure in [ZAP23] and train the model for 30,000 iterations to converge.

Dataset Details



Figure 4.17: 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 Figure 4.17. We follow the same dataset generation and experimental procedures as described in the main chapter. The camera/receiver poses for the TRAJ and




Figure 4.18: 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 that this is the 2D projection for visualization; the actual positions and orientations are in 3D.

RAND datasets are visualized in Figure 4.18. For 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-

Table 4.3: Scene Statistics

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

of-scene poses). Detailed scene information is summarized in Table 4.3. We note that this dataset is built on top of the open-source Replica dataset [SWM19], with additional radio material assignments, trajectory sampling, and wireless channel simulations. We will make this dataset publicly available to facilitate future research.

Geometric Differentiability




Figure 4.19: mmDiff’s differentiability enables geometry calibration.

We demonstrate the differentiability of mmDiff by showing its capability to calibrate

simple geometry offsets. We adopt the single-reflector scene shown in Figure 4.4, 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° 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 the AoA power spectrum using the 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.

Figure 4.19 shows the calibration of geometry offset over the course of training for 200 iterations. Figure 4.19a shows the calibration of offsets along the X/Y/Z axes. 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 axes. As shown in Figure 4.19b, in both cases, the calibrated angles converge back to the ground truth – 0° . 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 an interesting future direction.

Additional Results

Here, we provide the evaluation results on additional scenes to further validate the effectiveness of our proposed mmDiff model. Table 4.4 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 Angle Error (AAE) and Average

Power Error (APE) are also consistently lower for mmDiff, demonstrating its robustness to geometry noise.

Table 4.4: Robustness under Geometry Noise for Additional Scenes (TRAJ Dataset).

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

Table 4.5 presents the channel prediction accuracy results on additional scenes using 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 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 NeRF² baselines and the ablated version without calibration across all metrics. The failure of Sionna is mainly due to inconsistency in the traced paths, i.e., the paths traced in the ground-truth scene differ from those traced in the reconstructed scene. This inconsistency makes it difficult for the model to learn material parameters at ray-surface interactions, leading to poor generalization to unseen poses. NeRF² performs poorly due to the sparsity of training samples. Basically, as presented in [ZAP23, LVM24], NeRF² requires extremely dense sampling (around 6,000 samples/m³) of the scene to effectively learn the radiation field representation. In our datasets, the sampling density is around 40 samples/m³, which

Table 4.5: Prediction Accuracy Results for Additional Scenes (RAND Dataset).

Scene	Methods	Top-1 Peak			Top-5 Peaks		
		F1 $\uparrow$	AAE $\downarrow$	APE $\downarrow$	F1 $\uparrow$	AAE $\downarrow$	APE $\downarrow$
Office 1	Sionna	0.8756	2.79	0.92	0.8222	2.57	1.57
	NeRF ²	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
	mmDiff Full	0.9297	1.72	0.80	0.9145	1.70	1.02
Office 2	Sionna	0.8219	2.93	0.91	0.7845	2.57	1.74
	NeRF ²	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
	mmDiff Full	0.9722	1.04	0.27	0.9534	0.99	0.47
Office 3	Sionna	0.7882	4.92	4.80	0.7376	3.61	4.83
	NeRF ²	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
	mmDiff Full	0.9685	1.09	0.32	0.9437	0.96	0.50

is far below the training condition of NeRF², 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.

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 side-lobe, it can be suppressed using the windowing method.

As illustrated in Figure 4.20, 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. Figure 4.21 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.




Figure 4.20: Illustration of the multipath interference




Figure 4.21: Multipath can cause inconsistent AoA spectra due to the loss of phase information

Despite the inconsistency, the calibration works fine. As shown in Figure 4.22, since the reflection path (at 0°) is much stronger than the LoS path (at 45°), the loss is mostly contributed by the dominant path. By calibrating against the ground truth reference map, we can recover the path gain.




Figure 4.22: Material calibration with single dominant path.

CHAPTER 5

Summary and Future Work

This dissertation explores two aspects of integrated sensing and communication in wireless networks: unified design, and synergistic mechanisms. Unified design requires a compact hardware-software system that supports both sensing and communication, while synergistic mechanisms calls for concrete schemes for incorporating sensing to enhance communication performance. This dissertation dives into both aspects and identifies respective challenges. We found that the difference of information source for sensing and communication leads contradictory design considerations in waveform and system architecture, and reconcile it with backscatter technology. We recognize the challenges in survey wireless network coverage, identifying dead zones, and efficiently aligning mmWave beams. We bring latest discoveries in the deep learning to learn the wireless environment for predicting channels and beam directions. Here I summarize the contributions of this dissertation and outline future research directions in this exciting area.

5.1 Summary

This dissertation introduces the following systems:

MilBack: A Millimeter Wave Backscatter Network for Two-Way Communication and Localization We design and implement a mmWave backscatter system that integrates localization, orientation sensing and two-way communication. The core contributions are a novel node architecture that leverages the dual-port frequency-scanning antenna

(FSA) for retro-reflective backscatter communication and localization; a new second-order modulation scheme, OAQFM, that encodes data by the amplitudes of two tones, enabling high-data-rate two-way communication; and a holistic system design including the communication protocol, localization algorithm, and orientation sensing algorithm, enabling a complete backscatter system for practical applications.

NeWRF: A Deep Learning Framework for Wireless Radiation Field Reconstruction and Channel Prediction We introduce a NeRF-inspired deep learning framework for learning the wireless radiation field from sparse channel measurements. The key findings are the simple nature of wireless scenes that enables our model to learn, even in complex room-scale spaces using only sparse channel measurements, and a novel ray-searching algorithm to estimate the direction-of-arrival, enabling accurate channel predictions at any unvisited location.

mmDiff: A Noise-Robust Differentiable Ray-Tracing Framework for mmWave Scene Calibration and Channel Prediction We propose a novel differentiable ray-tracing framework for mmWave networks. The key contributions are a directional scattering approximation model that enables robust path gain prediction in noisy 3D-reconstructed environments, and an end-to-end learning framework that calibrates the 3D environment from mmWave AoA power measurements, enabling accurate mmWave beam prediction.

5.2 Future Work

Learning and understanding the wireless environment is a promising direction to improve wireless communication. It connects to the broad field of wireless digital twins, which is an ambitious objective to create a high-fidelity digital replica of the physical wireless environment to enable monitoring, analysis, and optimization of wireless networks. This dissertation represent an initial step towards this vision. Future work includes exploring the following directions:

Large-scale Wireless Environment Modeling The work so far focuses on modeling the wireless environment for small-scale indoor spaces. In commercial wireless networks, such as 5G/6G networks, the wireless environment is much larger and more complex, which requires more scalable and efficient modeling techniques. Potential research directions include developing multi-scale modeling techniques that can capture the wireless environment at different levels of granularity.

Digital-Twin-In-the-Loop Wireless Network Optimization Current work focuses on modeling the wireless environment, but it does not directly address the optimization of wireless networks. How to use the digital twin of the wireless environment to optimize wireless network performance is an intriguing research question. This could involve developing algorithms that can leverage the digital twin to optimize network parameters such as beamforming, power control, and resource allocation.

Bridging the Sim-to-Real Gap for Wireless Digital Twins We have demonstrated the capability of differentiable ray-tracing for minimizing the sim-to-real gap in mmWave beam prediction. However, there are still many challenges to be addressed in this direction, such as handling dynamic environments, incorporating more complex propagation effects, and improving the efficiency of the learning process. Future research could focus on developing more robust and efficient algorithms for sim-to-real transfer in wireless digital twins.

REFERENCES

- [ABD17] O. Abari, D. Bharadia, A. Duffield, and D. Katabi. “Enabling High-Quality untethered virtual reality.” In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, pp. 531–544, Boston, MA, March 2017. USENIX Association.
- [ADM20] A. Abedi, F. Dehbashi, M. H. Mazaheri, O. Abari, and T. Brecht. “WiTAG: Seamless WiFi backscatter communication.” In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication*, SIGCOMM ’20, p. 240–252, New York, NY, USA, 2020. Association for Computing Machinery.
- [AHM15] F. Adib, C.-Y. Hsu, H. Mao, D. Katabi, and F. Durand. “Capturing the human figure through a wall.” *ACM Transactions on Graphics (TOG)*, **34**(6):1–13, 2015.
- [AKK14] F. Adib, Z. Kabelac, D. Katabi, and R. C. Miller. “3D tracking via body radio reflections.” In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*, pp. 317–329, 2014.
- [AKK15] F. Adib, Z. Kabelac, and D. Katabi. “Multi-person localization via {RF} body reflections.” In *12th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 15)*, pp. 279–292, 2015.
- [ALC24] A. Abedi, H. Lu, A. Chen, C. Liu, and O. Abari. “WiFi physical layer stays awake and responds when it should not.” *IEEE Internet of Things Journal*, **11**(3):4483–4496, 2024.
- [AMG22] D. Azinović, R. Martin-Brualla, D. B. Goldman, M. Nießner, and J. Thies. “Neural rgb-d surface reconstruction.” In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6290–6301, 2022.
- [AMK15] F. Adib, H. Mao, Z. Kabelac, D. Katabi, and R. C. Miller. “Smart homes that monitor breathing and heart rate.” In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, pp. 837–846, 2015.
- [Ana23a] Analog Devices, Inc. *Datasheet Envelope Detector ADL6010*, 2023. Accessed: 2023-07-08.
- [Ana23b] Analog Devices, Inc. *Datasheet Power Amplifier ADPA7005*, 2023. Accessed: 2023-07-08.
- [Ana23c] Analog Devices, Inc. *Datasheet SPDT ADRF5020*, 2023. Accessed: 2023-07-08.

- [ASK25] Z. An, L. Shangguan, J. Kaewell, P. Pietraski, and K. Jamieson. “RadioTwin: A Digital Building Material Twin for Wideband, Cross-link, Cross-band Wireless Channel Prediction.” In *2025 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, pp. 1–10, May 2025.
- [AUC24] I. Ahmad, A. Ullah, and W. Choi. “WiFi-based human sensing with deep learning: Recent advances, challenges, and opportunities.” *IEEE Open Journal of the Communications Society*, **5**:3595–3623, 2024.
- [BAC22] K. M. Bae, N. Ahn, Y. Chae, P. Pathak, S.-M. Sohn, and S. M. Kim. “OmniScatter: Extreme sensitivity MmWave backscattering using commodity FMCW radar.” In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services, MobiSys ’22*, p. 316–329, New York, NY, USA, 2022. Association for Computing Machinery.
- [Bal16] C. A. Balanis. *Antenna theory: analysis and design*. John Wiley & sons, 2016.
- [BHM18] S. A. Busari, K. M. S. Huq, S. Mumtaz, L. Dai, and J. Rodriguez. “Millimeter-wave massive MIMO communication for future wireless systems: A survey.” *IEEE Communications Surveys & Tutorials*, **20**(2):836–869, 2018.
- [BJK15] D. Bharadia, K. R. Joshi, M. Kotaru, and S. Katti. “BackFi: High throughput WiFi backscatter.” *SIGCOMM Comput. Commun. Rev.*, **45**(4):283–296, aug 2015.
- [BR14] C. Boyer and S. Roy. “— invited paper — backscatter communication and RFID: Coding, energy, and MIMO analysis.” *IEEE Transactions on Communications*, **62**(3):770–785, 2014.
- [CFS24] X. Chen, Z. Feng, K. Sun, K. Qian, and X. Zhang. “Rfcanvas: Modeling rf channel by fusing visual priors and few-shot rf measurements.” In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*, pp. 464–477, 2024.
- [CLG21] W. Chen, J. Litalien, J. Gao, Z. Wang, C. Fuji Tsang, S. Khamis, O. Litany, and S. Fidler. “DIB-R++: learning to predict lighting and material with a hybrid differentiable renderer.” *Advances in Neural Information Processing Systems*, **34**:22834–22848, 2021.
- [CLJ21] Y. Cui, F. Liu, X. Jing, and J. Mu. “Integrating sensing and communications for ubiquitous IoT: Applications, trends, and challenges.” *IEEE Network*, **35**(5):158–167, 2021.
- [CLT20] R. Charbonnier, C. Lai, T. Tenoux, D. Caudill, G. Gougeon, J. Senic, C. Gentile, Y. Corre, J. Chuang, and N. Golmie. “Calibration of Ray-Tracing With Diffuse

- Scattering Against 28-GHz Directional Urban Channel Measurements.” *IEEE Transactions on Vehicular Technology*, **69**(12):14264–14276, December 2020.
- [Coo16] Cooley LLP. *FCC opens millimeter wave spectrum for 5G*, July 2016. Accessed: 2023-07-08.
- [CXW17] Y. Cui, S. Xiao, X. Wang, Z. Yang, S. Yan, C. Zhu, X.-Y. Li, and N. Ge. “Diamond: Nesting the data center network with wireless rings in 3-d space.” *IEEE/ACM Transactions On Networking*, **26**(1):145–160, 2017.
- [Dan23] M. Dano. “Verizon, T-Mobile, AT&T balk at drive testing their networks.” <https://www.lightreading.com/operations/verizon-t-mobile-at-t-balk-at-drive-testing-their-networks>, 2023. Accessed: 2026-07-22.
- [DFV07] V. Degli-Esposti, F. Fuschini, E. M. Vitucci, and G. Falciaeseca. “Measurement and modelling of scattering from buildings.” *IEEE transactions on antennas and propagation*, **55**(1):143–153, 2007.
- [EKS96] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, et al. “A density-based algorithm for discovering clusters in large spatial databases with noise.” In *kdd*, volume 96, pp. 226–231, 1996.
- [EPB18] M. S. Elbamby, C. Perfecto, M. Bennis, and K. Doppler. “Toward low-latency and ultra-reliable virtual reality.” *IEEE Network*, **32**(2):78–84, 2018.
- [Fre23] Free3D. “Bedroom 3D model.” <https://free3d.com/3d-model/bedroom-86855.html>, 2023.
- [FSB22] J. Foutz, A. Spanias, and M. Banavar. *Narrowband direction of arrival estimation for antenna arrays*. Synthesis Lectures on Antennas. Springer International Publishing, 2022.
- [FSC23] G. Formis, S. Scanzio, G. Cena, and A. Valenzano. “Predicting wireless channel quality by means of moving averages and regression models.” In *2023 IEEE 19th International Conference on Factory Communication Systems (WFCS)*. IEEE, April 2023.
- [GKZ20] K. Z. Ghafoor, L. Kong, S. Zeadally, A. S. Sadiq, G. Epiphaniou, M. Ham-moudeh, A. K. Bashir, and S. Mumtaz. “Millimeter-wave communication for internet of vehicles: Status, challenges, and perspectives.” *IEEE Internet of Things Journal*, **7**(9):8525–8546, 2020.
- [GZS25] Y. Guo, T. Zhang, S. Sun, M. Tao, and R. Gao. “Measurement and Analysis of Scattering From Building Surfaces at Millimeter-Wave Frequency.”, February 2025.

- [GZZ17] M. Giordani, A. Zanella, and M. Zorzi. “Millimeter wave communication in vehicular networks: Challenges and opportunities.” In *2017 6th International Conference on Modern Circuits and Systems Technologies (MOCASST)*, pp. 1–6. IEEE, 2017.
- [HAC24] J. Hoydis, F. A. Aoudia, S. Cammerer, F. Euchner, M. Nimier-David, S. T. Brink, and A. Keller. “Learning Radio Environments by Differentiable Ray Tracing.” *IEEE Transactions on Machine Learning in Communications and Networking*, **2**:1527–1539, 2024.
- [HAR18] H. Hassanieh, O. Abari, M. Rodriguez, M. Abdelghany, D. Katabi, and P. Indyk. “Fast millimeter wave beam alignment.” In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, SIGCOMM ’18*, p. 432–445, New York, NY, USA, 2018. Association for Computing Machinery.
- [HCA22] J. Hoydis, S. Cammerer, F. Ait Aoudia, M. Nimier-David, L. Maggi, G. Marcus, A. Vem, and A. Keller. “Sionna.”, 2022. <https://nvlabs.github.io/sionna/>.
- [HCX15] S. Han, I. Chih-Lin, Z. Xu, and C. Rowell. “Large-scale antenna systems with hybrid analog and digital beamforming for millimeter wave 5G.” *IEEE Communications Magazine*, **53**(1):186–194, 2015.
- [INO24] R. Iwatani, K. Negishi, T. Ohhira, and H. Hashimoto. “Breathing rate estimation for multiple people using Wi-Fi channel state information.” In *IECON 2024 - 50th Annual Conference of the IEEE Industrial Electronics Society*, pp. 1–6, 2024.
- [JCW25] H. Jia, X. Chen, Y. Wei, Y. Sun, and Y. Pi. “Neural Reflectance Fields for Radio-Frequency Ray Tracing.”, January 2025.
- [JSR22a] W. Jakob, S. Speierer, N. Roussel, and D. Vicini. “Dr.Jit: A just-in-time compiler for differentiable rendering.” *ACM Transactions on Graphics (TOG)*, **41**(4):1–16, 2022.
- [JSR22b] W. Jakob, S. Speierer, N. Roussel, M. Nimier-David, D. Vicini, T. Zeltner, B. Nicolet, M. Crespo, V. Leroy, and Z. Zhang. “Mitsuba 3 renderer.”, 2022. <https://mitsuba-renderer.org>.
- [KB14] D. P. Kingma and J. Ba. “Adam: A method for stochastic optimization.” *arXiv preprint arXiv:1412.6980*, 2014.
- [KD18] P. Kar and B. Dappuri. “Site survey and radio frequency planning for the deployment of next generation WLAN.” In *2018 International Conference on Wireless Communications, Signal Processing and Networking (WiSPNET)*, pp. 1–4, 2018.

- [KGR21] T. Kuhlmann, P. Garaizar, and U.-D. Reips. “Smartphone sensor accuracy varies from device to device in mobile research: The case of spatial orientation.” *Behavior research methods*, **53**:22–33, 2021.
- [KGT17] J. Kimionis, A. Georgiadis, and M. M. Tentzeris. “Millimeter-wave backscatter: A quantum leap for gigabit communication, RF sensing, and wearables.” In *2017 IEEE MTT-S International Microwave Symposium (IMS)*, pp. 812–815. IEEE, 2017.
- [KHC21] E. Krijestorac, S. Hanna, and D. Cabric. “Spatial signal strength prediction using 3D maps and deep learning.” In *ICC 2021-IEEE international conference on communications*, pp. 1–6. IEEE, 2021.
- [KHY25] H. Kong, C. Huang, J. Yu, and X. Shen. “A survey of mmWave radar-based sensing in autonomous vehicles, smart homes and industry.” *IEEE Communications Surveys & Tutorials*, **27**(1):463–508, 2025.
- [KM23] C. R. Karanam and Y. Mostofi. “A foundation for wireless channel prediction and full ray makeup estimation using an unmanned vehicle.” *IEEE Sensors Journal*, **23**(18):21452–21462, 2023.
- [KPG14] B. Kellogg, A. Parks, S. Gollakota, J. R. Smith, and D. Wetherall. “Wi-Fi backscatter: Internet connectivity for RF-powered devices.” In *Proceedings of the 2014 ACM conference on SIGCOMM*, pp. 607–618, 2014.
- [KR23] O. Kanhere and T. S. Rappaport. “Calibration of NYURay, a 3D mmWave and sub-THz Ray Tracer using Indoor, Outdoor, and Factory Channel Measurements.”, February 2023.
- [KSP18] K. Kibaroglu, M. Sayginer, T. Phelps, and G. M. Rebeiz. “A 64-element 28-GHz phased-array transceiver with 52-dBm EIRP and 8–12-Gb/s 5G link at 300 meters without any calibration.” *IEEE Transactions on Microwave Theory and Techniques*, **66**(12):5796–5811, 2018.
- [KZ13] S. Karimkashi and G. Zhang. “A dual-polarized series-fed microstrip antenna array with very high polarization purity for weather measurements.” *IEEE transactions on antennas and propagation*, **61**(10):5315–5319, 2013.
- [LCM22] F. Liu, Y. Cui, C. Masouros, J. Xu, T. X. Han, Y. C. Eldar, and S. Buzzi. “Integrated sensing and communications: Towards dual-functional wireless networks for 6G and beyond.” *IEEE journal on selected areas in communications*, 2022.
- [LCY26] H. Lu, Y. Cao, W. Yi, and O. Abari. “mmDiff: A noise-robust differentiable ray-tracing framework for mmWave scene calibration and channel prediction.” In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Findings*, pp. 6435–6444, June 2026.

- [Lee03] T. H. Lee. *The design of CMOS radio-frequency integrated circuits*. Cambridge university press, 2003.
- [LMA22] T. Li, M. H. Mazaheri, and O. Abari. “5G in the sky: the future of high-speed internet via unmanned aerial vehicles.” In *Proceedings of the 23rd Annual International Workshop on Mobile Computing Systems and Applications*, HotMobile ’22, p. 116–122, New York, NY, USA, 2022. Association for Computing Machinery.
- [LMR23] H. Lu, M. Mazaheri, R. Rezvani, and O. Abari. “A millimeter wave backscatter network for two-way communication and localization.” In *Proceedings of the ACM SIGCOMM 2023 Conference*, ACM SIGCOMM ’23, p. 49–61, New York, NY, USA, 2023. Association for Computing Machinery.
- [LMY16] F. Lemic, J. Martin, C. Yarp, D. Chan, V. Handziski, R. Brodersen, G. Fettweis, A. Wolisz, and J. Wawrzyn. “Localization as a feature of mmWave communication.” In *2016 International Wireless Communications and Mobile Computing Conference (IWCMC)*, pp. 1033–1038, 2016.
- [LVM24] H. Lu, C. Vattheuer, B. Mirzasoleiman, and O. Abari. “NeWRF: A deep learning framework for wireless radiation field reconstruction and channel prediction.” In R. Salakhutdinov, Z. Kolter, K. Heller, A. Weller, N. Oliver, J. Scarlett, and F. Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 33147–33159. PMLR, 21–27 Jul 2024.
- [LWZ20] H. Liu, Y. Wang, A. Zhou, H. He, W. Wang, K. Wang, P. Pan, Y. Lu, L. Liu, and H. Ma. “Real-time arm gesture recognition in smart home scenarios via millimeter wave sensing.” *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, **4**(4), dec 2020.
- [LZZ12] B. Li, Z. Zhou, W. Zou, X. Sun, and G. Du. “On the efficient beam-forming training for 60GHz wireless personal area networks.” *IEEE Transactions on Wireless Communications*, **12**(2):504–515, 2012.
- [MAA19] M. H. Mazaheri, S. Ameli, A. Abedi, and O. Abari. “A millimeter wave network for billions of things.” In *Proceedings of the ACM Special Interest Group on Data Communication*, SIGCOMM ’19, p. 174–186, New York, NY, USA, 2019. Association for Computing Machinery.
- [MAT23a] MATLAB. “Indoor MIMO-OFDM communication link using ray tracing.” <https://www.mathworks.com/help/comm/ug/indoor-mimo-ofdm-communication-link-using-ray-tracing.html>, 2023.

- [MAT23b] MATLAB. “Three-dimensional indoor positioning with 802.11az fingerprinting and deep learning.” <https://www.mathworks.com/help/wlan/ug/three-dimensional-indoor-positioning-with-802-11az-fingerprinting-and-deep-learning.html>, 2023.
- [MCA21] M. H. Mazaheri, A. Chen, and O. Abari. “MmTag: A millimeter wave backscatter network.” In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, SIGCOMM ’21, p. 463–474, New York, NY, USA, 2021. Association for Computing Machinery.
- [MHM22] B. Mildenhall, P. Hedman, R. Martin-Brualla, P. P. Srinivasan, and J. T. Barron. “Nerf in the dark: High dynamic range view synthesis from noisy raw images.” In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16190–16199, 2022.
- [Mit16] R. Mittra. *Computational electromagnetics*. Springer, 2016.
- [MM12] M. Malmirchegini and Y. Mostofi. “On the spatial predictability of communication channels.” *IEEE Transactions on Wireless Communications*, **11**(3):964–978, 2012.
- [Mol12] A. F. Molisch. *Wireless communications*. John Wiley & Sons, 2012.
- [MRG23] M. Mazaheri, R. Ruiz, D. Giustiniano, J. Widmer, and O. Abari. “Bringing millimeter wave technology to any IoT device.” In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, ACM MobiCom ’23, New York, NY, USA, 2023. Association for Computing Machinery.
- [MRS21] R. Martin-Brualla, N. Radwan, M. S. Sajjadi, J. T. Barron, A. Dosovitskiy, and D. Duckworth. “Nerf in the wild: Neural radiance fields for unconstrained photo collections.” In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7210–7219, 2021.
- [MST21] 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.” *Communications of the ACM*, **65**(1):99–106, 2021.
- [MTM24] S. Mura, D. Tagliaferri, M. Mizmizi, U. Spagnolini, and A. Petropulu. “Waveform design for OFDM-based ISAC systems under resource occupancy constraint.” In *2024 IEEE Radar Conference (RadarConf24)*, pp. 1–6, 2024.
- [MZW19] Y. Ma, G. Zhou, and S. Wang. “WiFi sensing with channel state information: A survey.” *ACM Computing Surveys (CSUR)*, **52**(3):1–36, 2019.

- [NJJ21] B. Nicolet, A. Jacobson, and W. Jakob. “Large steps in inverse rendering of geometry.” *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, **40**(6), December 2021.
- [NL19] J.-P. Niu and G. Y. Li. “An overview on backscatter communications.” *Journal of Communications and Information Networks*, **4**(2):1–14, 2019.
- [NLJ15] Y. Niu, Y. Li, D. Jin, L. Su, and A. V. Vasilakos. “A survey of millimeter wave communications (MmWave) for 5G: Opportunities and challenges.” *Wirel. Netw.*, **21**(8):2657–2676, nov 2015.
- [NQZ21] J. Nolan, K. Qian, and X. Zhang. “RoS: Passive smart surface for roadside-to-vehicle communication.” In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, SIGCOMM ’21, p. 165–178, New York, NY, USA, 2021. Association for Computing Machinery.
- [NST22] K. Neophytou, M. Steeg, J. Tebart, A. Stöhr, S. Iezekiel, and M. A. Antoniadis. “Simultaneous user localization and identification using leaky-wave antennas and backscattering communications.” *IEEE Access*, **10**:37097–37108, 2022.
- [OKK23] T. Orekondy, P. Kumar, S. Kadambi, H. Ye, J. Soriaga, and A. Behboodi. “WiN-eRT: Towards neural ray tracing for wireless channel modelling and differentiable simulations.” In *The Eleventh International Conference on Learning Representations*, 2023.
- [PCP21] A. Pumarola, E. Corona, G. Pons-Moll, and F. Moreno-Noguer. “D-nerf: Neural radiance fields for dynamic scenes.” In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10318–10327, 2021.
- [PJK12] S. Priebe, M. Jacob, and T. Kürner. “Calibrated broadband ray tracing for the simulation of wave propagation in mm and sub-mm wave indoor communication channels.” In *European Wireless 2012; 18th European Wireless Conference 2012*, pp. 1–10, April 2012.
- [Rao17] S. Rao. “Introduction to mmWave sensing: FMCW radars.” *Texas Instruments (TI) mmWave Training Series*, pp. 1–11, 2017.
- [RBM22] B. Roessle, J. T. Barron, B. Mildenhall, P. P. Srinivasan, and M. Nießner. “Dense depth priors for neural radiance fields from sparse input views.” In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12892–12901, 2022.
- [Rem23] Remcom. “Wireless InSite.” <https://www.remcom.com/wireless-insite-em-propagation-software>, 2023. Version 3.4.4.

- [RFN13] D. J. Richardson, J. M. Fini, and L. E. Nelson. “Space-division multiplexing in optical fibres.” *Nature photonics*, **7**(5):354–362, 2013.
- [RXM17] T. S. Rappaport, Y. Xing, G. R. MacCartney, A. F. Molisch, E. Mellios, and J. Zhang. “Overview of Millimeter Wave Communications for Fifth-Generation (5G) Wireless Networks-with a focus on Propagation Models.” *IEEE Transactions on Antennas and Propagation*, **65**(12):6213–6230, December 2017.
- [SD60] E. Sharp and M. Diab. “Van Atta reflector array.” *IRE Transactions on Antennas and Propagation*, **8**(4):436–438, 1960.
- [Ser15] P. Series. “Effects of building materials and structures on radiowave propagation above about 100 MHz.” *recommendation itu-r*, pp. 2040–1, 2015.
- [SPB21] E. Soltanaghaei, A. Prabhakara, A. Balanuta, M. Anderson, J. M. Rabaey, S. Kumar, and A. Rowe. “Millimetro: MmWave retro-reflective tags for accurate, long range localization.” In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking, MobiCom ’21*, p. 69–82, New York, NY, USA, 2021. Association for Computing Machinery.
- [SSS22] L. Souza, C. H. Souza Lopes, R. Santos, A. Cerqueira Sodré Junior, and L. L. Mendes. “A study on propagation models for 60 GHz signals in indoor environments.” *Frontiers in Communications and Networks*, **2**:757842, 2022.
- [SWM19] J. Straub, T. Whelan, L. Ma, Y. Chen, E. Wijmans, S. Green, J. J. Engel, R. Mur-Artal, C. Ren, S. Verma, A. Clarkson, M. Yan, B. Budge, Y. Yan, X. Pan, J. Yon, Y. Zou, K. Leon, N. Carter, J. Briaies, T. Gillingham, E. Mueggler, L. Pesqueira, M. Savva, D. Batra, H. M. Strasdat, R. D. Nardi, M. Goesele, S. Lovegrove, and R. Newcombe. “The Replica dataset: A digital replica of indoor spaces.” *arXiv preprint arXiv:1906.05797*, 2019.
- [SZW25] J. Shen, T. Zhao, Y. Wu, and X. Wang. “NeRF-APT: A New NeRF Framework for Wireless Channel Prediction.”, April 2025.
- [Tex22] Texas Instrument. *AWRL6432 single-chip 57- to 64-GHz automotive radar sensor*, 2022.
- [Tex23] Texas Instruments. *Datasheet Microcontroller MSP-EXP430FR6989*, 2023. Accessed: 2023-07-08.
- [TPA11] Y. M. Tsang, A. S. Poon, and S. Addepalli. “Coding the beams: Improving beamforming training in mmwave communication system.” In *2011 IEEE Global Telecommunications Conference-GLOBECOM 2011*, pp. 1–6. IEEE, 2011.
- [TPP11] F. Ticconi, L. Pulvirenti, N. Pierdicca, and V. Zhurbenko. “Models for scattering from rough surfaces.” *Electromagnetic waves*, **10**:203–226, 2011.

- [UM20] A. N. Uwaechia and N. M. Mahyuddin. “A comprehensive survey on millimeter wave communications for fifth-generation wireless networks: Feasibility and challenges.” *IEEE Access*, **8**:62367–62414, 2020.
- [VDP19] M. Vaezi, Z. Ding, and H. V. Poor. *Multiple access techniques for 5G wireless networks and beyond*, volume 159. Springer, 2019.
- [VGS23] R. Varshney, C. Gangal, M. Sharique, and M. S. Ansari. “Deep learning based wireless channel prediction: 5G scenario.” *Procedia Computer Science*, **218**:2626–2635, 2023. International Conference on Machine Learning and Data Engineering.
- [VSB16] V. Va, T. Shimizu, G. Bansal, R. W. Heath Jr, et al. “Millimeter wave vehicular communications: A survey.” *Foundations and Trends® in Networking*, **10**(1):1–113, 2016.
- [WKK18] X. Wang, L. Kong, F. Kong, F. Qiu, M. Xia, S. Arnon, and G. Chen. “Millimeter wave communication: A comprehensive survey.” *IEEE Communications Surveys & Tutorials*, **20**(3):1616–1653, 2018.
- [WLM25] H. Wang, X. Liang, D. Mishra, A. Seneviratne, and E. Ambikairajah. “Experimental demonstration of integrated WiFi communication and occupancy monitoring.” In *ICC 2025 - IEEE International Conference on Communications*, pp. 5591–5596, 2025.
- [WTH25a] C. Wen, J. Tong, Y. Hu, Z. Lin, and J. Zhang. “Neural representation for wireless radiation field reconstruction: A 3D gaussian splatting approach.”, 2025.
- [WTH25b] C. Wen, J. Tong, Y. Hu, Z. Lin, and J. Zhang. “Wrf-gs: Wireless radiation field reconstruction with 3d gaussian splatting.” In *IEEE INFOCOM 2025-IEEE Conference on Computer Communications*, pp. 1–10. IEEE, 2025.
- [WXZ25] Y. Wang, Y. Xue, S. Zhang, and T.-H. Chang. “RadSplat: Extending 3D Gaussian Splatting to Radio Frequencies for Wireless Radiomap Extrapolation.”, February 2025.
- [WZC21] T. Woodford, X. Zhang, E. Chai, K. Sundaresan, and A. Khojastepour. “Space-Beam: LiDAR-Driven One-shot MmWave beam management.” In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys ’21, pp. 389–401, New York, NY, USA, 2021. Association for Computing Machinery.
- [XJM21] H. Xue, Y. Ju, C. Miao, Y. Wang, S. Wang, A. Zhang, and L. Su. “MmMesh: Towards 3D real-time dynamic human mesh construction using millimeter-wave.” In *Proceedings of the 19th Annual International Conference on Mobile Systems*,

- Applications, and Services*, MobiSys '21, p. 269–282, New York, NY, USA, 2021. Association for Computing Machinery.
- [XYZ18] C. Xu, L. Yang, and P. Zhang. “Practical backscatter communication systems for battery-free Internet of Things: A tutorial and survey of recent research.” *IEEE Signal Processing Magazine*, **35**(5):16–27, 2018.
 - [XYZ22] J. Xiong, H. Yin, J. Zhu, and Y. Tang. “An overview of waveform design for integrated sensing and communication.” In *2022 IEEE/CIC International Conference on Communications in China (ICCC)*, pp. 991–996, 2022.
 - [XZY16] J. Xiao, Z. Zhou, Y. Yi, and L. M. Ni. “A survey on wireless indoor localization from the device perspective.” *ACM Comput. Surv.*, **49**(2), June 2016.
 - [YAD15] W. Yuan, S. M. Armour, and A. Doufexi. “An efficient and low-complexity beam training technique for mmWave communication.” In *2015 IEEE 26th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, pp. 303–308. IEEE, 2015.
 - [YDJ25] K. Yang, G. Dong, S. Ji, W. Du, and M. Srivastava. “GSRF: Complex-Valued 3D Gaussian Splatting for Efficient Radio-Frequency Data Synthesis.” In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.
 - [YI15] Z. Yun and M. F. Iskander. “Ray tracing for radio propagation modeling: Principles and applications.” *IEEE access*, **3**:1089–1100, 2015.
 - [YRG19] H. Yan, S. Ramesh, T. Gallagher, C. Ling, and D. Cabric. “Performance, power, and area design trade-offs in millimeter-wave transmitter beamforming architectures.” *IEEE Circuits and Systems Magazine*, **19**(2):33–58, 2019.
 - [ZAP23] X. Zhao, Z. An, Q. Pan, and L. Yang. “NeRF2: Neural radio-frequency radiance fields.” In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, New York, NY, USA, 2023. Association for Computing Machinery.
 - [ZBJ16] P. Zhang, D. Bharadia, K. Joshi, and S. Katti. “Hitchhike: Practical backscatter using commodity wifi.” In *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM*, pp. 259–271, 2016.
 - [ZHM25] W. Zhuang, X. He, Y. Mao, and J. Liu. “UAV-enabled wireless networks for integrated sensing and learning-oriented communication.” *IEEE Wireless Communications Letters*, **14**(2):340–344, 2025.
 - [ZKQ17] S. Zhan-shan, R. Ke, C. Qiang, B. Jia-jun, and F. Yun-qi. “3D radar imaging based on frequency-scanned antenna.” *IEICE Electronics Express*, **14**(12):20170503–20170503, 2017.

- [ZSB25] L. Zhang, H. Sun, S. Berweger, C. Gentile, and R. Q. Hu. “RF-3DGS: Wireless Channel Modeling with Radio Radiance Field and 3D Gaussian Splatting.”, July 2025.