

WiNeRF: Measurement Constrained Radiance Fields for Actionable Wireless Channel Modeling



Saif
Ur Rahman¹



Rafid Umayer
Murshed¹



Anton
Dmitriev¹



Cagri
Tanriover²



Rahul C.
Shah²



Elahé
Soltanaghai¹

¹ University of Illinois Urbana-Champaign ² Intel
EWSN 2026

Connectivity-aware robots




plan the

Coverage and AP planning



Wireless digital twins



3D model

All three need the channel at unseen locations
none of them gets it from a device's own sparse measurements

needs: link quality at locations it has not visited

today: statistical path-loss models [1]


needs: coverage at every point before deploying

today: interpolated site-survey RSSI [2]


needs: the channel at every point of the model

today: ray tracing from a CAD model [3]

What the receiver measures: the CSI matrix



OFDM
decoding



CSI (H): 4 antennas $\times$ 50 subcarriers

one measurement, at the receiver's own position

Wireless multipath: direct, walls, ceiling, floor, furniture

OFDM decoding: one complex gain per antenna per subcarrier


CSI (H) is the sum of all paths

$$H = \text{red} + \text{yellow} + \text{green} + \text{blue} + \text{purple}$$

CSI measurement is coarse: two hardware limits



Angle: 4 antennas cannot separate nearby arrivals




array output vs. look direction; dashed: true angle of arrivals (AoA)

limit: antenna array configuration puts limit on AoA resolution

WiNeRF: samples the rays the way the hardware sees them, as sets of directions it cannot tell apart

Phase: the absolute phase reference is unstable




absolute phase jumps every packet; the difference stays put

limit: absolute phase is unusable

WiNeRF: learns and is scored on relative phase only

left: array factor of a 4-element $\lambda/2$ linear array [8]; right: toy illustration

The problem: predicting the channel at unseen locations



Have

a few hundred coarse channel snapshots along a trajectory


Want

the channel at every unseen location

Underconstrained

many different arrangements of walls and furniture would produce the same few snapshots


Prior neural wireless fields trade extra information for output fidelity




Why cones? Sample only what the array can actually resolve




thin rays (NeRF², NeWRF)



samples a set of directions as if the array could tell them apart




(a) more antennas: narrower cones




(b) fewer antennas: wider cones

cone (WiNeRF)



one sample volume for the whole sector the array resolves



(c) linear array: front and back mirrored


two arrivals inside one sector look identical to the antenna array

thin rays sample them as distinct directions: aliasing

a cone integrates the sector; it is **not** a beam pattern

cone width and shape follow the antenna array

Restrict what the field may represent to what the antenna array can actually observe



array geometry → **cone shape**

a line of 4 antennas: no elevation resolution, front and back look alike, so the cone is full-height with a mirror twin

finite angular resolution → **cone width**

the cone spans the directions the array cannot resolve (previous slide)

unstable phase → **relative-phase loss**


only phase differences between antennas are trusted, in the loss and in the output

12 such cones around the receiver: the field is only asked to explain what these can observe

WiNeRF: from sparse CSI to a continuous complex-valued channel field



three **measurement-driven** modules turn sparse CSI into a **queryable channel field**



cone sampling


answers: **finite angular resolution**
one sample volume per resolvable sector

multi-scale scene

answers: **reflectors at all scales**
hash grid, room to furniture scale

CSI reconstruction

answers: **unstable absolute phase**
renders the 4×50 complex CSI; trained with a relative-phase loss



Input


sparse CSI at the measured positions

differentiable rendering


loss gradients flow back through all modules

Output

complex CSI at any location



Evaluation: Commodity-WiFi CSI collected from a robot in three rooms



the measurement platform

off-the-shelf parts: ASUS router on a custom 4-antenna mount; depth camera for pose (SLAM)




Meditation Room

4 × 2.25 m, small and cluttered




Office Space

5 × 3 m, metal cabinets



Conference Room

12 × 6 m, table blocks line of sight




measurement positions along random trajectories (illustrative)

~30k synchronized CSI + pose samples
≈ 4 per ft³ (NeRF² needs ~178)

LoS and NLoS conditions in every room
5-fold cross-validation at held-out receiver positions

WiNeRF predicts complex CSI $\sim 3\times$ more accurately than prior neural baselines





prediction SNR, all environments (higher is better)

5.3 dB

median prediction SNR

+4.9 dB

vs. baselines on average ($\approx 3\times$)



4 vs 178

samples / ft³: ours vs. NeRF²


Accurate CSI is not enough: it has to be directly usable




standard signal-processing algorithms consume WiNeRF's **complex CSI as-is**

predicted complex CSI

4 × 50, at any location




equal-gain combining




beamforming gain

MUSIC



line-of-sight AoA

$|H|^2$ averaged over antennas and subcarriers



RSSI coverage map

no downstream model is trained: the algorithms are unchanged, WiNeRF estimates the CSI input

Measurement constrained

hardware limits become inductive biases


Sparse → continuous

complex CSI at unseen locations from sparse real measurements
(5.3 dB, $\approx 3\times$ baselines)

Actionable

beamforming, AoA and coverage maps straight from the
predicted CSI

Thank you! Questions?



- [1] A. Muralidharan and Y. Mostofi, “Communication-aware robotics: exploiting motion for communication,” Annual Review of Control, Robotics, and Autonomous Systems, 2021.
- [2] D. Romero and S.-J. Kim, “Radio map estimation: a data-driven approach to spectrum cartography,” IEEE Signal Processing Magazine, 2022.
- [3] A. Alkhateeb, S. Jiang, and G. Charan, “Real-time digital twins: vision and research directions for 6G and beyond,” IEEE Communications Magazine, 2023.
- [4] X. Zhao, Z. An, Q. Pan, and L. Yang, “NeRF²: neural radio-frequency radiance fields,” ACM MobiCom, 2023.
- [5] T. Orekondy et al., “WiNeRT: towards neural ray tracing for wireless channel modelling and differentiable simulations,” ICLR, 2023.
- [6] X. Chen, Z. Feng, K. Sun, K. Qian, and X. Zhang, “RFCanvas: modeling RF channel by fusing visual priors and few-shot RF measurements,” ACM SenSys, 2024.
- [7] H. Lu, C. Vattheuer, B. Mirzasoleiman, and O. Abari, “A deep learning framework for wireless radiation field reconstruction and channel prediction” (NeWRF), arXiv:2403.03241, 2024.
- [8] H. L. Van Trees, Optimum Array Processing, Wiley, 2002.
- [9] A. Goldsmith, Wireless Communications, Cambridge University Press, 2005.