Supplementary for Physics-informed line-of-sight learning for scalable deterministic channel modeling

Xiucheng Wang[1], Junxi Huang[1], Conghao Zhou[1*], Xuemin Shen[2], Nan Cheng[1*]

[1*]School of Telecommunications Engineering, Xidian University, No. 2 South Taibai Road, Xi'an, 710071, Shaanxi, China.
[2]Department of Electrical and Computer Engineering, University of Waterloo, 200 University Avenue West, Waterloo, N2L 3G1, Ontario, Canada.

*Corresponding author(s). E-mail(s): conghao.zhou@ieee.org; dr.nan.cheng@ieee.org;
Contributing authors: xcwang_1@stu.xidian.edu.cn; 24012100067@stu.xidian.edu.cn; sshen@uwaterloo.ca;

Supplementary Theorem S1: Complexity Analysis of D^{2}LoS

Theorem 1 (Computational complexity reduction) *Given m building vertices and n evaluation points, D^{2}LoS reduces the per-transmitter LoS preprocessing complexity from $O(m \log m + n)$ to $O(M \log M + n)$, where M is the number of predicted boundary vertices. In typical urban scenarios, $M \ll m$ and $M \log M \ll n$, so the effective per-transmitter cost is $O(n)$. Under GPU parallelisation with sufficient cores, the wall-clock time is bounded by a constant independent of both m and n.*

Proof Classical baseline. The rotational sweep algorithm computes the complete LoS polygon for a single transmitter in $O(m \log m)$: all m edges are sorted angularly and swept. Evaluation of n receiver points against this polygon costs $O(n)$ via angular binary search. The total per-transmitter cost is $O(m \log m + n)$.

D^{2}LoS cost decomposition. D^{2}LoS replaces the $O(m \log m)$ polygon construction with three operations.

Step 1: CNN inference. The network processes a fixed $H \times W$ input tensor with $H = W = 257$. For a U-Net with L layers each of bounded filter size, the computation cost is $O(L \cdot H \cdot W)$. Since L, H, and W are all fixed constants independent of m and n, the total inference cost is a constant C_{CNN}.

Step 2: Geometric post-processing. For each of the M predicted boundary vertices, we identify the nearest building edge within a search radius R_{search}. We assume a uniform grid-based spatial index is precomputed for each scenario in $O(|E|)$ time, where $|E|$ is the total number of building edges. Given this index, each query retrieves at most k candidate edges, where $k \leq \pi R_{\mathrm{search}}^2 \cdot \rho_{\mathrm{edge}}$ and ρ_{edge} is the maximum local edge density. Since ρ_{edge} is bounded by the building geometry and R_{search} is a fixed constant, k is bounded by a constant independent of m. Each ray-edge intersection test costs $O(1)$. The total cost of geometric post-processing is therefore $O(M \cdot k) = O(M)$.

Step 3: LoS boundary rasterisation. The M boundary segments are sorted angularly with respect to $\mathbf{p}_s$ in $O(M \log M)$. The sorted boundary is then rasterised onto the $H \times W$ grid by scanline traversal in $O(n)$, where $n = H \times W$. The total rasterisation cost is $O(M \log M + n)$.

Total cost. Combining the three steps:

$$T_{\mathrm{D^2LoS}} = C_{\mathrm{CNN}} + O(M) + O(M \log M + n) = O(M \log M + n). \tag{S1}$$

In typical urban scenarios, $m \approx 1500$–3000 and $M \approx 20$–80 (i.e., only 2–5% of building vertices lie on the LoS boundary for a given transmitter). Therefore $M \log M \leq 80 \times \log_2(80) \approx 506$, while $n = 257^2 = 66{,}049$. Since $M \log M \ll n$, the effective complexity simplifies to $O(n)$, compared with $O(m \log m + n)$ for the classical method. The improvement factor in the polygon construction step is $m \log m / (M \log M) \geq 30\times$ for typical values.

GPU parallelisation. The three steps parallelise as follows. CNN inference maps to tensor-core parallelism with wall-clock time determined by network depth, not spatial resolution. This is a constant T_{CNN} on any modern GPU. The M edge-snapping and ray-edge intersection operations are independent and execute in $O(1)$ wall-clock time with M parallel threads. Angular sorting of M elements can be performed via parallel merge sort in $O(\log^2 M)$ wall-clock time. The n point-in-polygon evaluations are independent; each tests membership against the sorted angular boundary in $O(\log M)$ via binary search. With n parallel threads, this step has $O(\log M)$ wall-clock time.

The total wall-clock time under GPU parallelisation is:

$$T_{\mathrm{wall}} = T_{\mathrm{CNN}} + O(1) + O(\log^2 M) + O(\log M) = O(T_{\mathrm{CNN}} + \log^2 M). \tag{S2}$$

Since both T_{CNN} and $\log^2 M$ are bounded constants independent of m and n, the wall-clock time per transmitter is $O(1)$ with respect to the scenario size. □

Supplementary Theorem S2: UTD Numerical Stability and Field Continuity

Theorem 2 (Pointwise diffracted field bound and shadow boundary continuity) *Consider the Luebbers UTD diffraction coefficient D for a lossy dielectric wedge with exterior angle $n\pi$. The following properties hold:*

(a) At the shadow boundary ($\phi - \phi' = \pi$), the diffracted field satisfies $|E_d| = 0.5 \cdot |E_i|$, ensuring that the total field $E_{\mathrm{total}} = E_i + E_d$ is continuous across the boundary.

(b) For all observation angles, $|D|$ is bounded above by its value at the shadow boundary. Therefore, the envelope clamp $|D|_{\max} = 0.5$ does not alter the physical solution and serves only as a numerical safeguard.

Proof Part (a): Shadow boundary limit.

The Luebbers UTD diffraction coefficient consists of four terms:

$$D = \sum_{i=1}^{4} D_i, \quad D_i = \frac{-e^{-j\pi/4}}{2n\sqrt{2\pi k}}\, C_i \cdot \cot\!\left(\frac{\pi + \beta_i}{2n}\right) F\Big(kL_i\, a^{+}(\beta_i)\Big), \tag{S3}$$

where $\beta_i \in \{\phi - \phi',\ \phi + \phi'\}$ with appropriate sign combinations, $C_i \in \{1, 1, R_0, R_n\}$ are the face reflection coefficients, and $F(X) = 2j\sqrt{X}\, e^{jX} \int_{\sqrt{X}}^{\infty} e^{-jt^2}\, dt$ is the Fresnel transition function.

At the incident shadow boundary where $\beta_1 = \phi - \phi' \to \pi$, the cotangent factor in D_1 has a pole: $\cot\big(\frac{\pi+\beta_1}{2n}\big) \to \infty$. Simultaneously, the transition function argument satisfies $a^{+}(\beta_1) \to 0$, so $F(kL_1\, a^{+}(\beta_1)) \to 0$. Applying L'Hôpital's rule to the product:

$$\lim_{\beta_1 \to \pi} \cot\!\left(\frac{\pi + \beta_1}{2n}\right) F\Big(kL_1\, a^{+}(\beta_1)\Big) = n\sqrt{\frac{2\pi k}{L_1}}\, e^{j\pi/4}. \tag{S4}$$

Substituting back into D_1 and multiplying by the free-space spreading factor $\sqrt{L_1}/(4\pi d)$ yields $|E_d^{(1)}| \to 0.5 \cdot |E_i|$ at the shadow boundary. The remaining terms D_2, D_3, D_4 involve $\phi + \phi'$ combinations that do not produce poles at $\phi - \phi' = \pi$ and contribute finite corrections. Their sum preserves the total limit of $0.5 \cdot |E_i|$ by the standard UTD continuity condition, which requires the total field to equal $0.5 \cdot E_i$ at the shadow boundary by construction of the Fresnel transition function.

Part (b): Global bound.

Away from the shadow boundary, $a^{+}(\beta_i) > 0$ and the Fresnel transition function satisfies $|F(X)| \leq 1$ for all $X \geq 0$, with $|F(X)| \to 1$ as $X \to \infty$. Meanwhile, the cotangent factors remain bounded when β_i is away from $\pm n\pi$. The product $|\cot(\cdot) \cdot F(\cdot)|$ achieves its maximum at the shadow boundary where the L'Hôpital limit applies, and decreases monotonically as the observation angle moves away from the boundary. This monotonic decay is a well-known property of the UTD diffraction coefficient for convex wedges with $n \geq 1$ (see Kouyoumjian and Pathak, Proc. IEEE, 1974).

Therefore, $|D(\phi)| \leq |D(\phi = \phi' + \pi)|$ for all observation angles ϕ. The envelope clamp $|D|_{\max} = 0.5$ does not clip any physically correct solution. It acts solely as a guard against floating-point overflow when $\cot(\cdot)$ and $F(\cdot)$ are evaluated independently near the shadow boundary, where their product is finite but each factor individually diverges.

Numerical implementation. In floating-point arithmetic, finite-precision cancellation in the transition function $F(X)$ for small X can produce spurious large values before the product is formed. The envelope clamp provides a post-hoc correction:

$$|E_d^{\text{clamped}}| = \min\big(|E_d^{\text{computed}}|,\ 0.5 \cdot |E_i|\big). \tag{S5}$$

This ensures that numerical artifacts do not violate the physical bound established in Part (b). □

Remark 1 (Forward-scatter smoothing) For near-grazing forward scattering with deflection angle $0 < \theta < 30^{\circ}$, the four UTD cotangent terms partially cancel, producing a small net diffraction coefficient with rapid oscillation due to the competing signs of D_1–D_4. These oscillations are sensitive to small perturbations in wedge geometry and are not reliably reproduced in discretised building models where edge positions are known only to pixel-level precision ($\Delta r = 1$ m). We replace the oscillatory solution with a monotonic linear attenuation in the log domain:

$$L_{\text{smooth}}(\theta) = 6.02 + \frac{30 - 6.02}{30^{\circ}} \cdot (30^{\circ} - \theta) \quad [\text{dB}], \quad \theta \in [0^{\circ}, 30^{\circ}]. \tag{S6}$$

At $\theta = 30^\circ$, $L_{\text{smooth}} = 6.02$ dB, matching the UTD envelope clamp and ensuring C^0 continuity. At $\theta = 0^\circ$, $L_{\text{smooth}} = 30$ dB, reflecting the physical expectation of strong attenuation for forward-diffracted rays. This smoothing is a modelling choice validated by the end-to-end RSS accuracy reported in the main text.

Supplementary Theorem S3: Geometric Validity of 2D×2D Decomposition

Theorem 3 (Exactness of 2D×2D LoS decomposition) *Consider a set of buildings $\{B_1, \ldots, B_N\}$ under the 2.5D model, where each building is an extruded right prism $B_i = P_i \times [0, h_i]$ with non-overlapping footprints ($P_i \cap P_j = \emptyset$ for $i \neq j$), polygonal footprint P_i, and height $h_i > 0$. For any transmitter $\mathbf{p}_s = (x_s, y_s, h_{\text{tx}})$ and receiver $\mathbf{p}_r = (x_r, y_r, h_{\text{rx}})$ located outside all buildings, the full 3D direct-path LoS test is equivalent to:*

$$\text{LoS}_{\text{3D}}(\mathbf{p}_s, \mathbf{p}_r) = \text{LoS}_{\text{2D-H}}(\mathbf{p}_s, \mathbf{p}_r) \ \wedge \ \text{LoS}_{\text{2D-V}}(\mathbf{p}_s, \mathbf{p}_r), \tag{S7}$$

where $\text{LoS}_{\text{2D-H}}$ is the horizontal intersection test on the 2D projected ray and $\text{LoS}_{\text{2D-V}}$ is the vertical elevation test against building heights. This equivalence holds with zero approximation error.

Proof Setup. Let $R_{\text{3D}}(t) = \big(r_{\text{2D}}(t),\, z(t)\big)$ parameterise the direct ray from $\mathbf{p}_s$ to $\mathbf{p}_r$ with $t \in [0, 1]$, where $r_{\text{2D}}(t) = (1-t)(x_s, y_s) + t(x_r, y_r)$ and $z(t) = (1-t)\, h_{\text{tx}} + t\, h_{\text{rx}}$. Note that $z(t)$ is linear and therefore monotonic or constant over any sub-interval.

Necessary and sufficient condition for occlusion. The ray is occluded by building B_i if and only if there exists $t^* \in (0, 1)$ such that $R_{\text{3D}}(t^*) \in B_i$. By the Cartesian product structure $B_i = P_i \times [0, h_i]$, this decomposes as:

$$\exists\, t^* \in (0,1): \quad r_{\text{2D}}(t^*) \in P_i \ \wedge \ 0 \leq z(t^*) \leq h_i. \tag{S8}$$

Case 1: No horizontal intersection. If $r_{\text{2D}}(t) \notin P_i$ for all $t \in (0, 1)$, then no t^* satisfying Eq. (S8) exists, and the ray clears B_i. No vertical check is needed.

Case 2: Horizontal intersection exists. Suppose the 2D projected ray intersects P_i. The intersection consists of one or more intervals $[t_{\text{in}}, t_{\text{out}}] \subset (0, 1)$. The ray is inside B_i in the vertical dimension if and only if $z(t^*) \leq h_i$ for some $t^* \in [t_{\text{in}}, t_{\text{out}}]$. Since $z(t)$ is linear, its minimum over $[t_{\text{in}}, t_{\text{out}}]$ is $\min\big(z(t_{\text{in}}),\, z(t_{\text{out}})\big)$. The condition $z(t^*) \geq 0$ is always satisfied since $h_{\text{tx}} > 0$ and $h_{\text{rx}} > 0$. Therefore, the ray clears B_i vertically if and only if:

$$\min\big(z(t_{\text{in}}),\, z(t_{\text{out}})\big) > h_i. \tag{S9}$$

Global LoS condition. The ray has a clear LoS if and only if it clears every building. Since the footprints are non-overlapping, the horizontal intersection tests are independent across buildings. The global condition is:

$$\text{LoS}_{\text{3D}} = \bigwedge_{i=1}^{N} \Big[\underbrace{\big(r_{\text{2D}}(t) \notin P_i,\ \forall t \in (0,1)\big)}_{\text{no horizontal intersection}} \ \vee \ \underbrace{\min\big(z(t^i_{\text{in}}),\, z(t^i_{\text{out}})\big) > h_i}_{\text{clears rooftop}} \Big]. \tag{S10}$$

This is precisely the conjunction of $\text{LoS}_{\text{2D-H}}$ and $\text{LoS}_{\text{2D-V}}$. The decomposition introduces no approximation because:

1. Each building B_i is a Cartesian product of a 2D footprint and a 1D height interval, making horizontal and vertical dimensions separable.

2. The ray height $z(t)$ is linear, so its extremum over any interval is attained at an endpoint.
3. Building footprints are non-overlapping, so per-building tests are independent.

Therefore Eq. (S7) holds exactly under the stated assumptions. □

Remark 2 (Non-convex footprints) When a building footprint P_i is non-convex, the horizontal ray-footprint intersection may consist of multiple disjoint intervals $\{[t_{\text{in}}^{(k)}, t_{\text{out}}^{(k)}]\}_{k=1}^{K}$. The vertical check in Eq. (S9) is then applied independently to each interval. The decomposition remains exact because the linearity of $z(t)$ and the Cartesian product structure hold regardless of footprint convexity.

Remark 3 (Limitations of the 2.5D assumption) The 2.5D prism assumption requires each building to have a uniform height h_i. This does not hold for buildings with complex rooftop geometries such as pitched roofs, domes, or rooftop equipment. In such cases, the actual rooftop surface is a function $h_i(x, y)$ rather than a constant, and the vertical check in Eq. (S9) must be replaced by:

$$\min_{t\in[t_{\text{in}},t_{\text{out}}]} \left[z(t) - h_i\big(r_{2\text{D}}(t)\big)\right] > 0, \tag{S11}$$

which cannot be evaluated from endpoint values alone when $h_i(x, y)$ is non-linear. The approximation error is bounded by $\max_{(x,y)\in P_i} |h_i(x, y) - \bar{h}_i|$, where $\bar{h}_i$ is the uniform height used in the 2.5D model.

Supplementary Results: Beam Pattern Integration

The main text demonstrates post-hoc beam pattern integration with a 15° half-power beamwidth. Here we provide the complete set of results across four beamwidths (15°, 30°, 45°, 60°) and supplementary qualitative comparisons. In all cases, the antenna gain pattern $G_{\text{ant}}(\phi_{\text{AoD}}, \theta_{\text{AoD}})$ is applied to per-ray AoD after the ray tracing computation, without re-executing visibility determination or path search.

15° beamwidth: qualitative APS and PDP

The main text presents 15° beam RSS maps and quantitative violin distributions. Supplementary Figs. 1 and 2 show the corresponding APS and PDP qualitative comparisons. The narrow 15° beam acts as a stringent test of per-ray AoD accuracy: small angular errors in the predicted rays are amplified by the sharp beam roll-off, causing misclassified rays to be either strongly attenuated or incorrectly retained. D^{2}LoS reproduces the beam-filtered APS profiles with correct dominant peak positions and sidelobe suppression. In the PDP domain, the narrow beam removes most off-axis multipath, leaving only a few dominant delay components. D^{2}LoS captures both their positions and relative power levels, whereas RadioUNet and RMTransformer produce delay profiles that bear no resemblance to the ground truth.

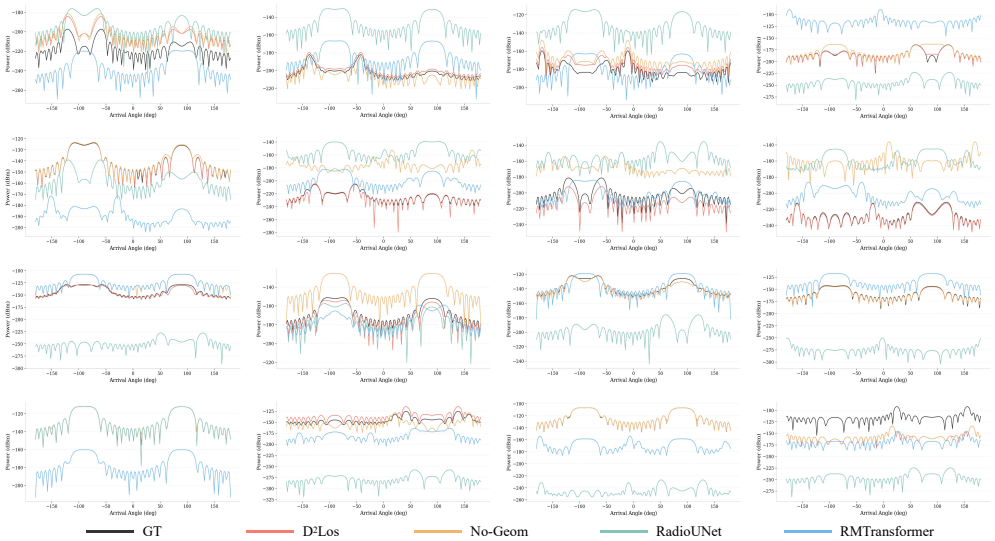

Fig. 1: **APS comparison with a 15° beam pattern.** The directional pattern narrows the angular spread and amplifies the dominant arrival direction. D^{2}LoS reproduces the beam-filtered angular profile. No-Geom preserves the general shape but introduces secondary peak errors. RadioUNet and RMTransformer produce profiles unrelated to the ground truth.

30° beamwidth

Increasing the beamwidth to 30° admits more off-axis multipath into the received signal. This relaxes the angular selectivity compared to the 15° case, making the beam-filtered channel closer to the omnidirectional baseline. As shown in Fig. 3, D^{2}LoS maintains compact error distributions across all three channel dimensions. The RSS MAE remains within 1 dB of the omnidirectional result, and the APS shape cosine stays above 0.85. No-Geom shows moderate degradation, while RadioUNet and RMTransformer continue to produce heavy-tailed error distributions. The qualitative RSS maps (Fig. 4) confirm that D^{2}LoS preserves the wider beam footprint and shadow boundaries. The APS and PDP curves (Figs. 5 and 6) show that D^{2}LoS tracks both the dominant and secondary multipath components under the 30° beam.

45° beamwidth

At 45° beamwidth, the directional filtering becomes mild and the beam-specific channel approaches the omnidirectional case. Fig. 7 shows that the error distributions for D^{2}LoS are nearly identical to the omnidirectional violin plots in the main text. This is expected: when the beam is wide enough to capture most multipath components, angular errors in individual rays have limited impact on the aggregated metrics. The RSS maps (Fig. 8) show only subtle differences from the omnidirectional maps, primarily at the beam edges where slight power attenuation is visible. The APS and PDP curves (Figs. 9 and 10) closely match the omnidirectional results for D^{2}LoS,

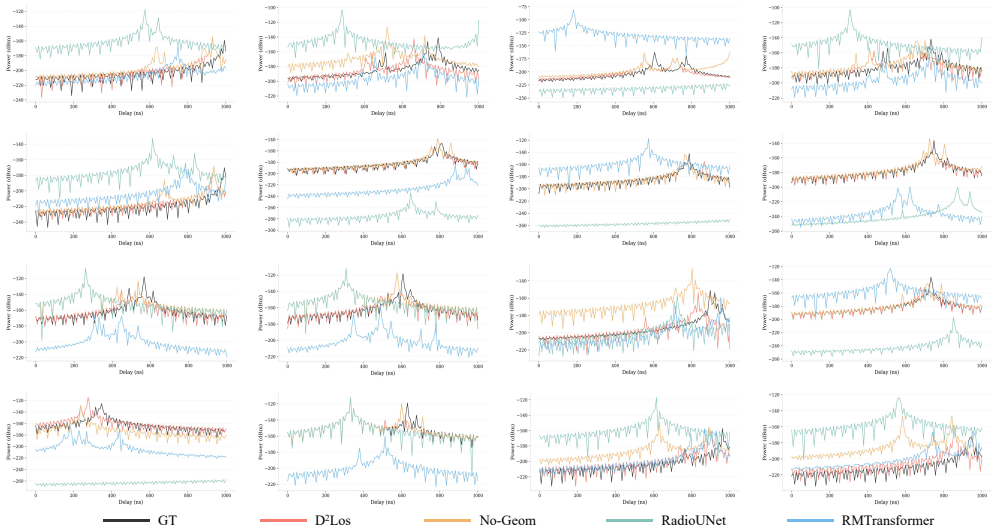

Fig. 2: **PDP comparison with a 15° beam pattern.** The beam pattern suppresses off-axis multipath, reducing the number of resolvable delay components. D^{2}LoS correctly captures this filtering effect. RadioUNet and RMTransformer produce delay profiles with incorrect power levels and spurious peaks.

while RadioUNet and RMTransformer remain unable to produce physically meaningful outputs regardless of beamwidth.

60° beamwidth

The 60° beamwidth represents the widest tested configuration and the mildest directional filtering. As shown in Fig. 11, D^{2}LoS retains low-error, compact distributions that are effectively indistinguishable from the omnidirectional case. The RSS maps (Fig. 12) differ from the omnidirectional maps primarily in a slight overall power offset due to the finite antenna gain. The APS and PDP curves (Figs. 13 and 14) confirm that D^{2}LoS accuracy is independent of beamwidth. Taken together with the 15° results, these findings demonstrate that the per-ray angular information produced by D^{2}LoS is sufficiently accurate to support beam-specific analysis from narrow pencil beams to near-omnidirectional patterns, all from a single ray tracing computation.

Summary

Across all four beamwidths, D^{2}LoS consistently produces compact, low-error distributions in RSS, APS, and PDP metrics. The performance ranking among methods is preserved regardless of beamwidth, confirming that the per-ray angular information predicted by D^{2}LoS is robust to directional filtering. Narrower beams amplify the impact of LoS boundary errors on angular metrics, yet D^{2}LoS maintains stable accuracy even at 15°, validating the geometric precision of the vertex-level prediction and post-processing pipeline.

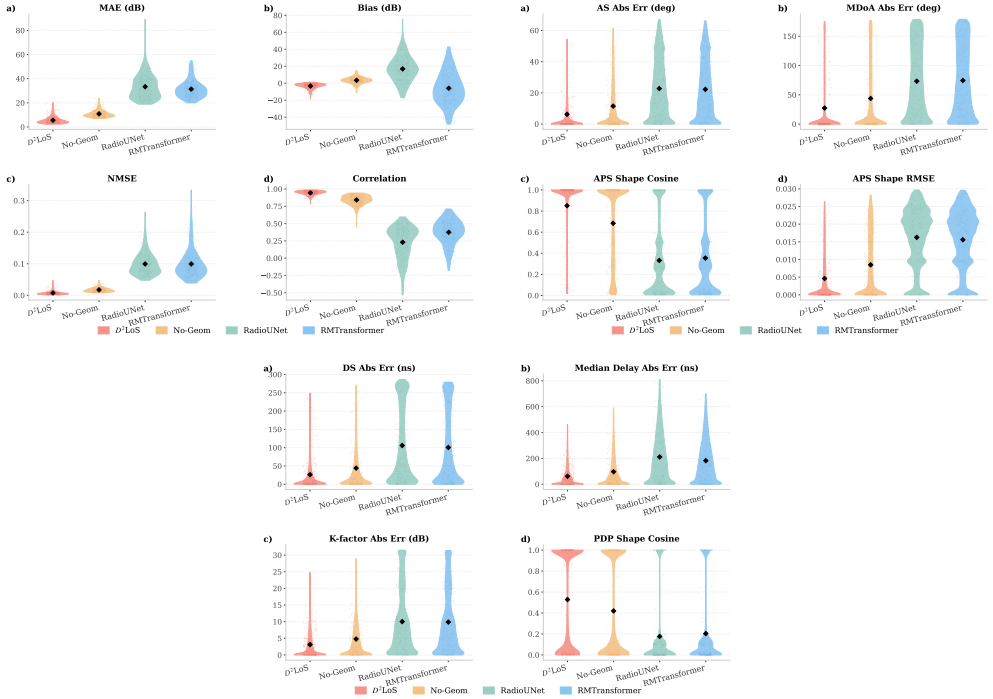

Fig. 3: **Prediction accuracy with a 30° beam pattern.** D^{2}LoS maintains compact, low-error distributions across RSS, APS, and PDP metrics, consistent with the 15° case.

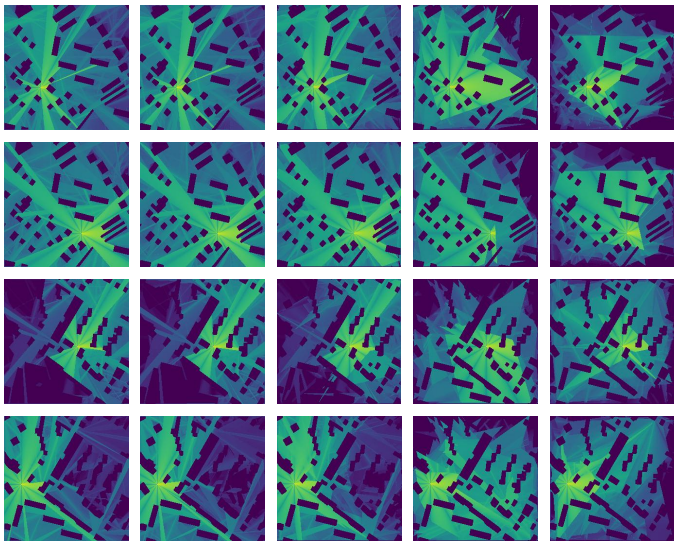

Fig. 4: **RSS radio maps with a 30° beam pattern.** The wider beam produces broader spatial coverage compared to the 15° case. D^{2}LoS preserves both the beam footprint and shadow structure.

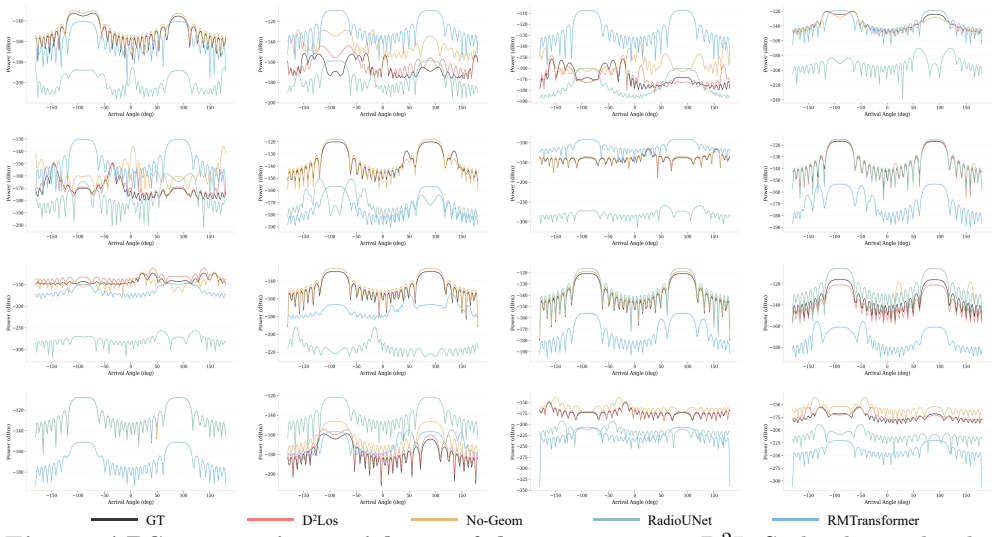

Fig. 5: **APS comparison with a 30° beam pattern.** D^{2}LoS closely tracks the ground truth angular profile under the wider beam.

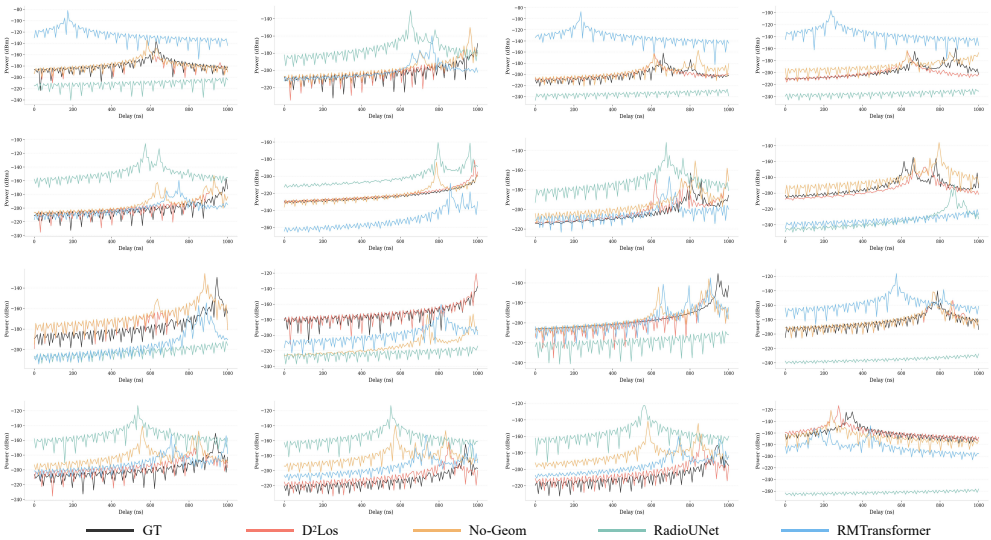

Fig. 6: **PDP comparison with a 30° beam pattern.** The wider beam retains more multipath components than the 15° case. D^{2}LoS reproduces the delay structure accurately.

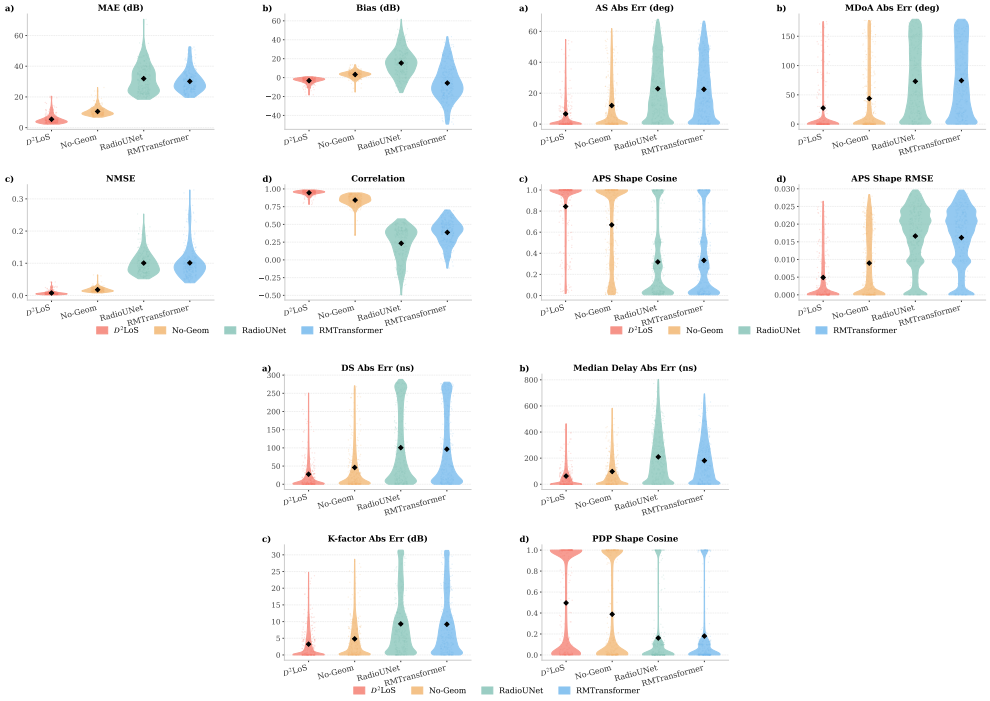

Fig. 7: **Prediction accuracy with a 45° beam pattern.** Error distributions are nearly identical to the omnidirectional results, confirming that the wide beam captures most multipath energy.

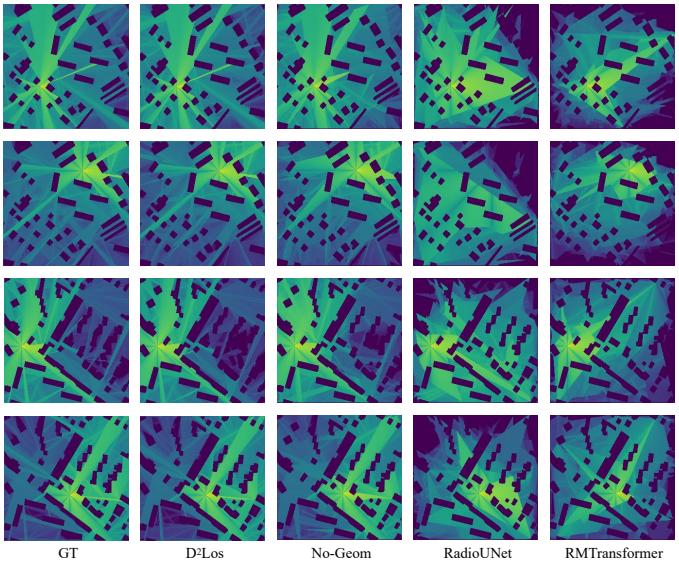

Fig. 8: **RSS radio maps with a 45° beam pattern.** The coverage pattern is close to omnidirectional with mild directional shaping at the beam edges.

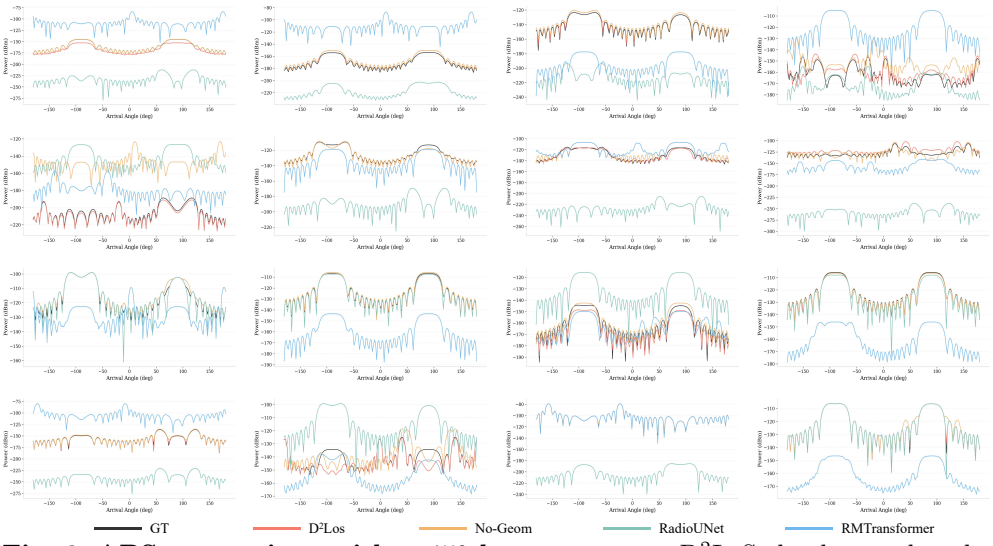

Fig. 9: **APS comparison with a 45° beam pattern.** D^{2}LoS closely matches the ground truth. The wider beam preserves most angular features from the omnidirectional case.

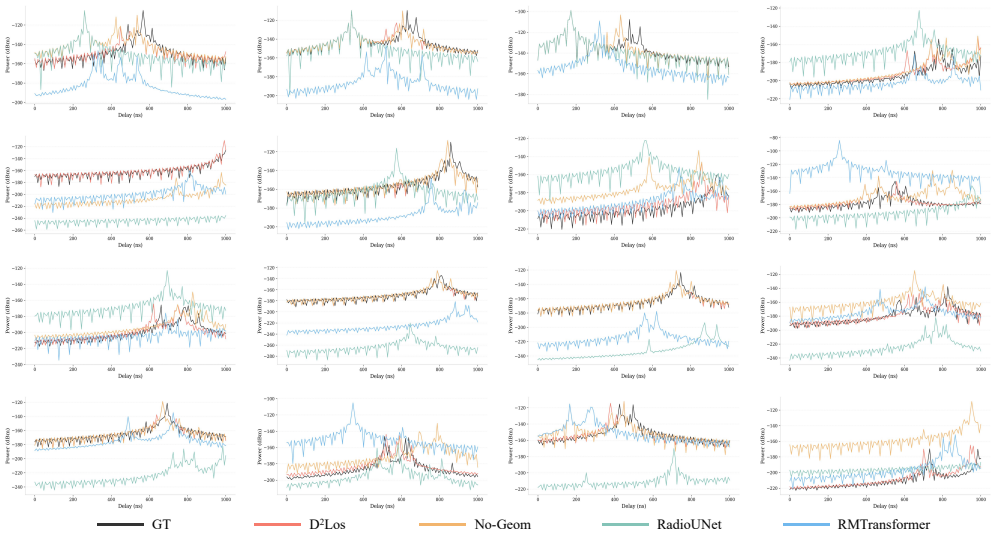

Fig. 10: **PDP comparison with a 45° beam pattern.** The delay profile is close to the omnidirectional case. D^{2}LoS maintains accurate reproduction of multipath structure.

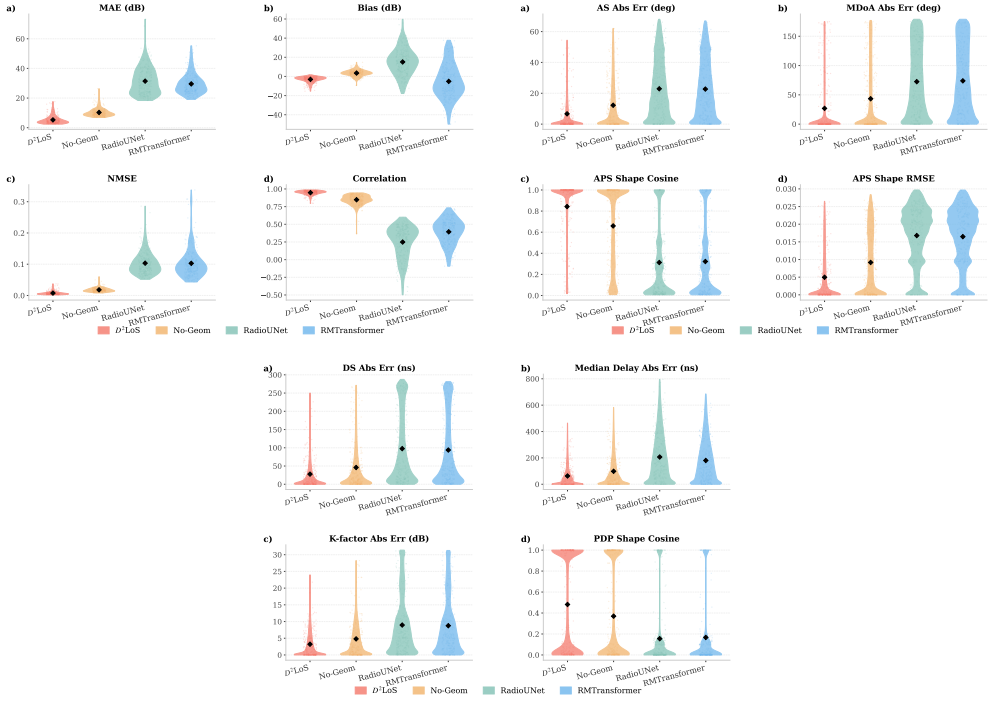

Fig. 11: **Prediction accuracy with a 60° beam pattern.** Error distributions are effectively indistinguishable from the omnidirectional case, confirming robustness across all tested beamwidths.

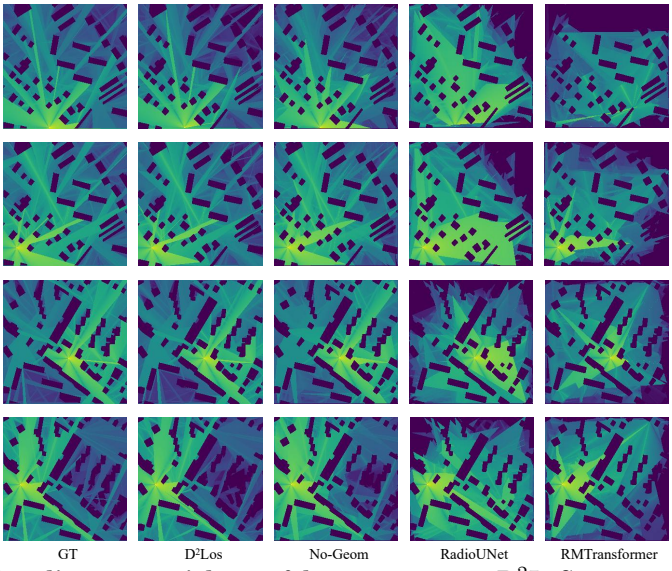

Fig. 12: **RSS radio maps with a 60° beam pattern.** D^{2}LoS accurately reproduces the near-omnidirectional coverage with mild directional shaping.

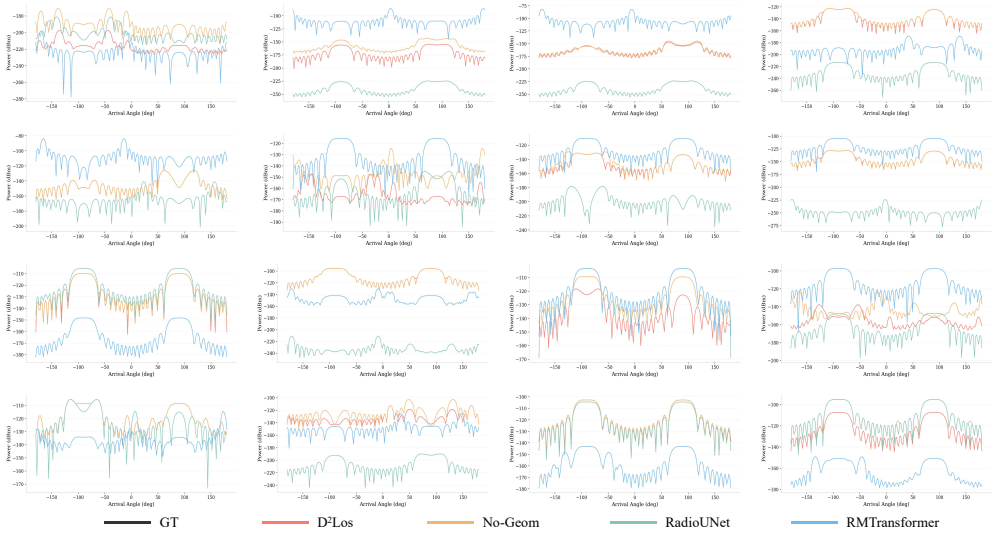

Fig. 13: **APS comparison with a 60° beam pattern.** The angular profile is close to the omnidirectional case. D^{2}LoS maintains full accuracy.

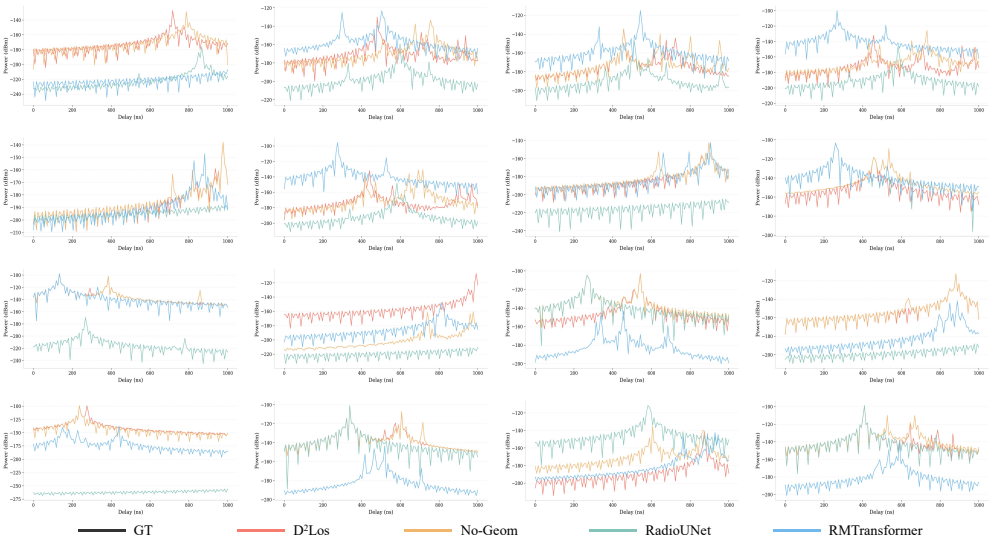

Fig. 14: **PDP comparison with a 60° beam pattern.** Delay structure matches the omnidirectional case closely. D^{2}LoS reproduces it accurately.