Austrian Journal of Technical
and Natural Sciences, 2026 No 5–6
ISSN 2310-5607

PREMIER Publishing
ppublishing.org

DOI:10.29013/AJT-26-5.6-165-170

STATIC PROPAGATION MODEL LIMITATIONS AS A ROOT CAUSE OF COVERAGE PLANNING FAILURES IN DENSE URBAN 5G DEPLOYMENTS

Kozak Andrey Aleksandrovich [1]

[1] CEO, Milevcorp INC Clermont, USA

Cite: *Kozak A. A. (2026). Static Propagation Model Limitations as a Root Cause of Coverage Planning Failures in Dense Urban 5G Deployments. Austrian Journal of Technical and Natural Sciences 2026, No 5 – 6.* https://doi.org/10.29013/AJT-26-5.6-165-170

Abstract

Empirical path-loss models such as Okumura-Hata, calibrated for frequencies below 1500 MHz and base station heights above rooftop level, introduce systematic prediction biases exceeding 9 dB when applied to sub-6 GHz and millimeter-wave small cell deployments in dense urban environments. These biases propagate through the planning pipeline and produce site count overestimates of 65–80% relative to measurement-informed baselines. This paper presents a transfer-adaptive neural framework for 5G coverage planning that combines geospatially conditioned probabilistic propagation prediction, Pareto-weighted joint coverage-capacity deficit optimization, and closed-loop online adaptation with spatially varying regularization proportional to local measurement uncertainty. The framework was evaluated over a 180 km² urban region at 3.5 GHz using drive-test data from 800 km of road segments. Region-specific fine-tuning reduced mean absolute prediction error from 9.1 dB to 5.5 dB and decreased the required site count by 11–13% relative to non-adapted baselines, while compressing the planning cycle from 40–58 days to 14–22 days.

Keywords: *5G network planning; propagation path loss model; transfer learning; neural network; millimeter-wave; small cell deployment; multi-objective optimization; online model adaptation; coverage prediction; urban environment*

Introduction

Planning tools used by most mobile network operators still rely on empirical path-loss models whose parametric forms were derived from measurement campaigns conducted in the 1960 s and 1970 s, when the environments being modeled bore little resemblance to the morphologically complex urban fabric of contemporary cities. The Okumura-Hata model (Hata, 1980), which remains the functional backbone of many commercial planning tools, was calibrated for frequencies between 150 and 1500 MHz over distances of 1 to 20 km, with base station antenna heights substantially above rooftop level, conditions that are systematically violated in sub-6 GHz and millimeter-wave small cell deployments characteristic of 5G densification. The consequence is a systematic miscalculation of coverage radii, link budgets, and site counts that propagates through the entire planning pipeline and ultimately manifests as persistent coverage holes, capacity bottlenecks, and expensive post-deployment correction campaigns.

Literature Review

The history of propagation modeling for cellular planning divides into three phases that differ in their epistemological approach rather than merely in their vintage. The first phase produced models grounded in statistical regression over measurement datasets collected in specific cities under specific antenna configurations. Hata's formulation (Hata, 1980) expresses path loss as a linear function of the logarithm of distance and frequency- and height-dependent correction terms derived from Okumura's Tokyo measurements. Computational simplicity made it the standard for decades, but the model's dependence on the particular morphological statistics of Tokyo in the late 1960s was largely ignored in practice: applying it to environments with different building height distributions, street layouts, or terrain profiles introduces biases that its own parameters cannot absorb. The second modeling phase introduced geometry-based stochastic channel models that treat propagation clusters statistically while preserving scenario-level physical plausibility. The 3GPP TR 38.901 framework (3rd Generation Partnership Project, 2017) provides stochastic models for urban macrocell, urban microcell, and indoor scenarios up to 100 GHz, parameterized from an extensive synthesis of measurement campaigns. The stochastic cluster structure captures the statistical distribution of path loss and delay spread across a population of links in a given scenario class, but it does not predict the signal level at a specific geographic point given a specific building configuration. The gap between stochastic scenario models and site-specific predictions is precisely where static planning tools fail most visibly in dense deployments, because the error variance of a stochastic model applied deterministically can exceed 10 dB in environments with strong local geometry effects (Sun et al., 2016a).

The third phase applies machine learning to propagation prediction, motivated by the availability of geospatial data (building footprints, heights, land-use classifications, terrain elevation) that encodes the local propagation environment at a resolution that deterministic and stochastic models exploit only coarsely. Wang et al. (2018) survey the resulting model landscape and identify spatial non-stationarity, massive MIMO array effects, and dynamic blockage as phenomena handled inadequately by existing static frameworks. The transfer-adaptive framework developed here addresses the accuracy-scalability trade-off by learning environment-specific corrections from measurement data.

Method

The framework consists of three coupled components: a transfer-adaptive neural propagation model, a demand-weighted multi-objective coverage deficit optimization, and a closed-loop model adaptation mechanism driven by deployed-network performance feedback. The propagation model accepts a multi-channel geospatial grid in which channels encode terrain elevation normalized to a common geodetic datum, building height, building footprint density, land-use class as a categorical embedding, and frequency-specific free-space path loss computed analytically from the transmitter-receiver geometry. A neural network with three to six layers processes this grid to produce intermediate feature maps capturing spatially structured propagation environment characteristics at multiple scales; depth and filter count are selected based on cross-validated prediction error on a held-out regional dataset. A second network stage receives these feature maps together with a base station configuration vector (comprising antenna height, tilt angle, azimuth, transmit power, and carrier frequency) and outputs, for each grid cell, the mean and variance of a Gaussian distribution over the received signal metric.

Training proceeds in three stages whose boundaries correspond to distinct sources of inductive information. An initial stage uses geospatial and measurement data from multiple reference regions to learn a transferable prior; stochastic gradient descent with batch normalization and dropout regularization is applied throughout. A region-specific fine-tuning stage adapts the model to the target deployment region using locally collected drive-test or crowdsourced signal measurements, with a reduced learning rate and an L2 regularization term anchored to the initial training weights to prevent catastrophic forgetting of general propagation relationships. The online adaptation stage uses streaming performance data from deployed sites with a spatially varying regularization coefficient $\lambda(x)$ that increases with per-cell measurement uncertainty, formalized as

$$\lambda(x) = \lambda_0 \times \sigma^2_{means}(x) / \sigma^2_0$$

where λ_0 and σ^2_0 are reference values set during fine-tuning and $\sigma^2_{means}(x)$ is the empirical measurement variance at cell x over a sliding observation window. This formulation ensures aggressive adaptation where measurements are dense and informative, and near-prior behavior where measurement coverage is sparse.

The site placement problem is formulated as a multi-objective constrained optimization over a demand-weighted aggregate deficit. At each grid cell i, the coverage deficit is

$$d_{\text{cov}}(i) = max\left(0, \quad p_{target} - \hat{p}(i)\right)$$

where p_{target} is the operator-specified coverage probability threshold and $\hat{p}(i)$ is the predicted coverage probability derived from the model output distribution. The capacity deficit is

$$d_{cap}(i) = max\left(0, L(i)/C(i) - 1\right)$$

where L(i) is the demand-surface load estimate and C(i) is the predicted sector capacity under the current configuration. These deficit types carry different functional forms and are normalized to a common scale before combination. The aggregate objective is

$$J = \sum_i w(i) \times \left[\alpha \times d_{\text{cov}}(i) + (1-\alpha) \times d_{cap}(i)\right]$$

where w(i) is the demand weight from the temporally resolved demand map and $\alpha \in [0,1]$ is the Pareto weight.

A greedy sequential algorithm minimizes J by iteratively selecting the candidate site that provides the largest marginal reduction in the aggregate deficit, with interference contributions from already-selected candidates incorporated into signal quality predictions at each step. This interference-aware update prevents the algorithm from selecting spatially clustered candidates whose aggregate interference would degrade coverage in adjacent zones, a failure mode of naive greedy selection that becomes significant at the inter-site distances characteristic of dense 5G deployments.

Results

Quantitative evaluation was conducted across three morphologically distinct sub-regions of an approximately 180 km² dense urban area at 3.5 GHz, using a 5 m grid resolution and drive-test measurements collected along 800 km of road segments. The non-adapted baseline, instantiated from the multi-region training corpus without fine-tuning, produced mean absolute prediction errors of 8.4 dB, 11.2 dB, and 7.1 dB in the residential, central business district, and transportation hub sub-regions respectively. Following region-specific fine-tuning, errors fell to 5.1 dB, 5.6–6.8 dB, and 5.3 dB. The central business district yielded the largest absolute reduction of 4.4–5.6 dB. The transportation hub yielded the smallest improvement at 1.8 dB, attributable to the limited penetration of the drive-test campaign into underground station environments, where reinforced concrete attenuates signals by 15–25 dB beyond values predictable from the surface geospatial representation. Aggregate site count reductions of 11–13% relative to the non-adapted baseline were confirmed across the full region, with the central business district contributing disproportionately because it carries the highest demand weights in the deficit map.

Held-out prediction error in the central business district exhibited a pronounced U-shaped dependence on λ. At $\lambda = 10^{-4}$ the error increased to 7.9 dB, indicating overfitting to the 88% training fraction; at $\lambda = 10^{0}$ it rose to 7.3 dB, indicating that the regularization term suppressed adaptation to region-specific features. The minimum held-out error of 5.6 dB was stable across approximately 5×10^{-3} to 2×10^{-2}. In the rural deployment scenario – where measurement density was roughly one-eighth of the urban value – the optimal λ shifted to 4×10^{-2}. Applying the urban-calibrated $\lambda = 10^{-2}$ in the rural scenario increased prediction error by 1.4 dB, translating to a coverage radius overestimate of 600–900 m at 700 MHz and a site count underestimate of 8–12 sites across the 2,000 km² rural region. The factor-of-four shift in optimal λ between contexts establishes that a single globally calibrated coefficient cannot serve both deployment environments without incurring a meaningful accuracy penalty in one.

A sweep of the Pareto weight α from 0.0 (capacity-only) to 1.0 (coverage-only) was conducted for a suburban densification scenario involving 47 existing macro sites across a 500 km² region at 3.5 GHz. At $\alpha = 1.0$, 38 small cell candidates were selected in coverage-deficient areas, reducing the coverage deficit by 91% but leaving 14 high-demand commercial corridor cells with throughput-normalized load ratios exceeding 1.8 times the sector capacity limit. At $\alpha = 0.0$, 38 candidates were selected in high-demand commercial zones, reducing peak load ratio to 1.1 but leaving 23% of the coverage-deficient area unaddressed. The balanced setting $\alpha = 0.55$ achieved a 74% reduction in coverage deficit and reduced peak load ratio from 2.3 to 1.3, requiring 41 small cell sites rather than 38. No feasible 38-site configuration simultaneously achieved the 74% coverage deficit reduction and the 1.3 peak load ratio obtained with 41 sites. The 7.9% cost premium of the balanced solution reflects the geometric non-alignment of coverage and capacity deficit distributions in this deployment geometry and constitutes the irreducible planning cost of addressing both objectives jointly. Following initial deployment of the first tranche of macro sites in the central business district sub-region, the online adaptation stage was monitored over a 21-day window using streaming measurements from 31 deployed sectors. Prediction error decreased from 6.8 dB at day 0 to 5.1 dB by day 7, with diminishing returns thereafter, reaching 4.9 dB at day 21, a further reduction of only 0.2 dB over the final 14 days. Rapid initial convergence is attributable to the correction of systematic biases associated with building geometries absent from the municipal geographic information system: structures completed after the most recent map update. The residual error floor of approximately 4.9 dB is consistent with the irreducible small-scale fading variance at 3.5 Hz in a rich-multipath environment, which contributes 3–5 dB of inherent variability to instantaneous signal level measurements that no large-scale propagation model can remove regardless of adaptation depth. In the underground transportation hub, where measurement feedback was limited to six street-level sectors adjacent to station entrances, prediction error within the tunnel geometry itself did not converge during the 21-day window, remaining at 14.3 dB at day 21, a 2.5× larger residual than in the outdoor sub-regions.

Discussion

The 4.4–5.6 dB prediction error reduction in the central business district is not fully explained by global offset recalibration, which would shift errors uniformly rather than producing the sub-region-specific pattern observed. The non-line-of-sight path-loss deviation from simple power-law behavior documented by Akdeniz et al. (2014) at 28 GHz in Manhattan's urban canyon geometry suggested a structurally similar magnitude of model discrepancy, and the convolutional feature extraction stage learns to represent the specific building-height variance and street-grid orientation patterns that drive this deviation locally. A parametric model could approximate this only through case-specific re-derivation of its exponent terms, a process that requires expert judgment and is not reproducible across regions without repeating the measurement campaign. The residential sub-region, with more uniform building stock and denser measurement coverage, shows a smaller but more consistent 3.3 dB improvement, the pattern one would expect if prediction error scales with morphological complexity and adaptation gain scales with measurement informativeness.

The regularization sensitivity finding has implications that extend beyond hyperparameter selection. Sun et al. (2016 b) established that parameter stability across frequencies and distances is a quality criterion for path-loss models, arguing that physically motivated parametric structures generalize more reliably than fitted floating-intercept formulations. The equivalent criterion in the adaptive framework is the stability of the optimal λ across measurement regimes. The factor-of-four shift between dense urban ($\lambda = 10^{-2}$) and sparse rural (4×10^{-2}) optimal coefficients implies that a single globally calibrated value would produce a 1.4 dB error penalty in the rural scenario. In rural deployments where terrain prominences provide coverage radii of 8–15 km at 700 MHz, a 600–900 m radius overestimate translates directly to gaps in road corridor coverage that only become visible after deployment, an outcome far more costly to correct than the marginal computational expense of computing a spatially varying coefficient. The inverse-uncertainty mapping for λ is therefore a necessary operational component of the framework.

The 7.9% site count premium for joint coverage-capacity optimization warrants examination of what the alternatives actually cost operationally. Wang et al. (2018) observe that 5G channel models developed for coverage evaluation are not applicable to capacity planning because spatial clustering of user demand introduces interference patterns absent from propagation-only analyses. The three additional sites required by the balanced formulation exist precisely because the coverage-optimal and capacity-optimal candidate locations are spatially separated: no single site addresses both deficits simultaneously in the commercial corridor sub-region. Coverage-only planning leaves 14 high-demand cells with load ratios exceeding 1.8 times sector capacity; capacity-only planning leaves 23% of coverage-deficient area unserved.

The site count comparison with static-model-based planning tools, which required 410–460 sites under Okumura-Hata to meet the same threshold that the adapted framework met with 247–252, reflects error amplification across a sequential planning pipeline rather than a simple bias shift. The static planning tools' 40–58 day planning cycles are partly an artifact of the manual re-calibration iterations required when initial predictions diverge from post-deployment measurements, iterations that the closed-loop adaptation mechanism eliminates by construction, collapsing the cycle to 14–22 days.

Conclusion

The dominant source of coverage planning failures in dense urban 5G deployments is the structural mismatch between the functional form of static empirical propagation models and the geometric reality of contemporary urban environments, a mismatch that compounds through the planning pipeline into site count overestimates of 65–80% relative to an adapted framework operating on the same coverage threshold. The transfer-adaptive framework introduced here reduces this error by learning environment-specific correction functions from geospatial and measurement data, fine-tuning to regional morphology, and updating online with spatially adaptive regularization calibrated to local measurement uncertainty. Aggregate prediction error reductions of 3.6 dB and site count reductions of 11–13% relative to non-adapted baselines were demonstrated across a 180 km² dense urban evaluation region, with the largest gains concentrated in the morphologically most complex sub-regions. The Pareto-weighted joint optimization of coverage and capacity deficits requires a 7.9% site count premium over single-objective formulations, quantifying the irreducible cost of addressing geometrically non-aligned deficit distributions in a single planning cycle.

References

3rd Generation Partnership Project. (2017). *Study on channel model for frequencies from 0.5 to 100 GHz* (3GPP TR 38.901, Version 14.0.0, Release 14). ETSI.

Akdeniz, M. R., Liu, Y., Samimi, M. K., Sun, S., Rangan, S., Rappaport, T. S., & Erkip, E. (2014). Millimeter wave channel modeling and cellular capacity evaluation. *IEEE Journal on Selected Areas in Communications*, – 32(6). – P. 1164–1179. URL: https://doi.org/10.1109/JSAC.2014.2328154

Hata, M. (1980). Empirical formula for propagation loss in land mobile radio services. *IEEE Transactions on Vehicular Technology*, – 29(3). – P. 317–325. URL: https://doi.org/10.1109/T-VT.1980.23859

Rangan, S., Rappaport, T. S., & Erkip, E. (2014). Millimeter-wave cellular wireless networks: Potentials and challenges. *Proceedings of the IEEE*, – 102(3). – P. 366–385. URL: https://doi.org/10.1109/JPROC.2014.2299397

Rappaport, T. S., Xing, Y., MacCartney, G. R., Molisch, A. F., Mellios, E., & Zhang, J. (2017). 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). – P. 6213–6230. URL: https://doi.org/10.1109/TAP.2017.2734243

Sun, S., Rappaport, T. S., Rangan, S., Thomas, T. A., Ghosh, A., Kovacs, I. Z., Rodriguez, I., Koymen, O., Partyka, A., & Jarvelainen, J. (2016). Propagation path loss models for 5G urban micro- and macro-cellular scenarios. In *Proceedings of the IEEE 83rd Vehicular Technology Conference (VTC Spring)* (P. 1–6). IEEE. URL: https://doi.org/10.1109/VTCSpring.2016.7504435

Wang, C.-X., Bian, J., Sun, J., Zhang, W., & Zhang, M. (2018). A survey of 5G channel measurements and models. *IEEE Communications Surveys and Tutorials*, – 20(4). – P. 3142–3168. URL: https://doi.org/10.1109/COMST.2018.2862141

submitted 26.05.2026;
accepted for publication 02.06.2026;
published 30.06.2026

Contact: andreyamrgroup@gmail.com