DEPLOY WHERE YOU WANT: SYNTHESIS OF MOBILE NETWORK DIGITAL TWINS VIA GENERATIVE ARTIFICIAL INTELLIGENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

High-fidelity wireless mobile network simulation is essential for developing and benchmarking network operational algorithms, yet remains constrained by the high cost of 3D scene acquisition and the lack of joint modeling across the radio, channel, and user aspects. In this paper, we present MobiSim, a deployable mobile network digital twin platform that overcomes these barriers by grounding simulation in generative AI. Given a geographic area, MobiSim automatically constructs a realistic urban environment from public map data and drives a hybrid raytracing engine to model complex multipath propagation, reproducing received-power trends across four real-world measurement datasets spanning sub-6 GHz and millimeter-wave bands. Further, MobiSim co-simulates a complete 5G RAN protocol stack alongside a generative user behavior foundation model that synthesizes realistic mobility trajectories and uplink/downlink traffic demands (app-usage and traffic NMAE reduced by 6.0% and 5.0%, respectively, relative to AppGen on Shanghai), producing high-fidelity full-stack records spanning channel states, infrastructure operations, and user behaviors. The platform further serves as a verifiable reward environment for network optimization, supporting operations research, reinforcement learning, and LLM-based autonomous management. We release MobiSim together with benchmark datasets and standardized evaluation interfaces covering network planning, antenna adjustment, and wireless resource scheduling, providing the community with a reproducible foundation for mobile network research. The code and generated datasets can be found at `https://anonymous.4open.science/r/MobiSim-F533/`.

1 INTRODUCTION

With the large-scale deployment of fifth-generation (5G) networks and the evolution toward next-generation architectures, the electromagnetic environment has become increasingly complex due to higher spectrum usage Tao et al. (2025); Yang et al. (2025), while user and network behavior grow more heterogeneous due to denser connections and more complex communication protocols Polese et al. (2024); Shi et al. (2024). Consequently, conventional approaches struggle to meet network planning and optimization demands, which require tightly coupled decisions across base station deployment, spectrum allocation, antenna beamforming, and radio resource scheduling. Large-scale trials on operational networks are generally impractical, as they incur high costs and risk to service quality. Therefore, mobile network digital twins that reproduce network operation, evaluate optimization strategies, and generate reliable performance feedback before deployment have become a key foundation for intelligent networks.

While representative platforms, including ns-3 Riley & Henderson (2010), OpenAirInterface Nikaein et al. (2014), and the Vienna Simulator Müller et al. (2018), have been developed for simulating mobile networks, applying these tools to achieve large-scale and generalizable simulations remains hindered by two fundamental challenges. First, accurate radio propagation simulation relies on fine-grained 3D representations of the built environment. Constructing such scene models typically requires UAV-based aerial surveying, LiDAR scanning, and point cloud processing Tang et al. (2010); Siebert & Teizer (2014). While these pipelines can support high-fidelity ray tracing within localized areas, they do not scale readily to large geographic regions or multi-city deployments, and

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

the overall costs remain prohibitively high for general use. Second, jointly modeling base station protocols, electromagnetic propagation, and user dynamics requires integration across tools with different modeling assumptions and interfaces. Protocol simulators such as ns-3, raytracing platforms such as Sionna RT Hoydis et al. (2023), and user mobility generators provide complementary functionality and can be integrated. Such integration requires aligning geographic coordinates, temporal resolutions, channel representations, and control interfaces so that changes in user mobility or network configuration propagate consistently through the coupled simulation. MobiSim addresses this engineering requirement through a shared workflow linking propagation, user dynamics, and radio-access operations.

Recent advances in generative AI offer a promising path toward overcoming both limitations. Modern generative models have demonstrated remarkable content synthesis capabilities across multiple domains. In spatial scene generation, models such as CityDreamer Xie et al. (2024) and UrbanWorld Shang et al. (2024) can synthesize geometrically plausible three-dimensional urban environments directly from OpenStreetMap data, while time-series foundation models can capture complex spatiotemporal patterns and produce high-fidelity user behaviors, as demonstrated by recent works on mobility trajectories Long et al. (2025) and network activities Chai et al. (2026). More importantly, generative AI is driving AI-Generated Networks (AIGNs) by formulating wireless phenomena as conditional generation tasks. For instance, diffusion models enable robust physical-layer channel modeling Wu et al. (2024) and dynamic network resource optimization Fan et al. (2026). Ultimately, these generative frameworks provide a powerful engine for synthesizing high-fidelity mobile network digital twins.

Motivated by these advancements, this paper presents MobiSim, a mobile network digital twin benchmark platform. To implement accurate electromagnetic simulation, MobiSim leverages a generative 3D urban world model to synthesize high-fidelity building geometries for a target geographic area with available map data and street-view imagery, and designs a hybrid raytracing pipeline that combines Monte Carlo Non-Line-of-Sight (NLOS) tracing with deterministic direct propagation. This approach effectively reduces discontinuities and coverage holes caused by finite ray budgets, especially for far-field regions with small angular support and blocked regions that are unlikely to be reached by random rays. We compare MobiSim with real-world measurements from DeepSense 6G Scenario 6 and AERPAW Datasets 18, 23, and 24 (Alkhateeb et al., 2023; Asokan et al., 2025; Singh, 2025a;b). The smoothed received-power curves show consistent variations, with Pearson's r ranging from 0.736 to 0.831 and Spearman's ρ from 0.770 to 0.878. Then, user behavior is jointly modeled across mobility trajectories, application usage, and network traffic through a generative foundation model that learns the spatiotemporal patterns of human activity in urban environments, producing highly realistic network activity data. On the Shanghai dataset, the model achieves the lowest mean errors among the compared methods on six of the eight reported metrics, while Pategail performs better on Duration and DailyLoc. Building on this foundation, MobiSim reproduces the 5G Radio Access Network (RAN) system, including site and sector topology, antenna-domain propagation, mobility management, and physical radio resource control, which supports cell association and handover, duplexing, link adaptation, and PRB granular scheduling. Crucially, rather than a simple simulator, this RAN stack is designed as a programmable policy container. By abstracting complex protocols into standardized control interfaces, MobiSim establishes a unified benchmark environment for tasks like network planning, coverage optimization, and Quality-of-Experience (QoE) maximization. Through these interfaces, operations research solvers, reinforcement learning (RL) agents, and Large Language Model (LLM)-based autonomous optimizers can interactively submit configuration actions and receive physically grounded performance feedback. The main contributions of this work are summarized as follows.

- We propose a generative AI-based mobile network simulator that synthesizes diverse environmental elements with accurate geometric properties for urban areas with available map data and street-view imagery. It enables scalable and geography-aware construction of mobile network digital twins for high-fidelity electromagnetic simulation and network performance evaluation.
- To the best of our knowledge, MobiSim is the first benchmark that integrates generative urban environments, raytracing-based channel simulation, 5G RAN protocol stacks, and generative user behavior within a unified simulation framework, which characterizes the end-to-end performance of mobile networks, spanning radio propagation, user mobility, and infrastructure-level operations.
- MobiSim provides physically grounded feedback and reward signals for evaluating diverse optimization paradigms, including operations research solvers, reinforcement learning agents, and

Table 1: Comparison of mobile network simulation platforms.

Platform	Language	Simulation Method	Generative 3D Environment	Ray-Tracing Channel	5G RAN Stack	User Behavior	Optimization Interface
ns-3	C++	Event-triggered	✗	✗	✓	✗	✓
OpenAirInterface	C	Event-triggered	✗	✗	✓	✗	✓
Vienna Simulator	MATLAB	Rule-based	✗	✗	✓	✗	✗
Wireless InSite	C++	Rule-based	✗	✓	✗	✗	✗
Sionna	Python + Tensorflow	Rule-based	✗	✓	✓	✗	✗
DeepMIMO	Python / MATLAB	Rule-based	✗	✓	✗	✗	✗
DT-SimNet	Python	Data-driven	✗	✗	✓	✓	✓
MobiSim (Ours)	**Python + PyTorch**	**Data-driven + Rule-based**	✓	✓	✓	✓	✓

LLM-based autonomous optimizers. It supports a broad set of tasks such as network deployment, frequency planning, coverage optimization, user experience optimization, and radio resource scheduling, offering a reproducible benchmark for future mobile network design and optimization.

2 Related Work

Mobile network digital twin spans from protocol-level simulators that focus on RAN stack emulation and wireless channel simulations that prioritize electromagnetic fidelity. We review representative works along each direction below.

Mobile Network Digital Twins. Current simulation platforms range from protocol-level tools to emerging digital twins. Protocol-level simulators such as ns-3/5G-LENA Henderson et al. (2008); Patriciello et al. (2019), OpenAirInterface Kaltenberger et al. (2020), and the Vienna 5G Simulator Müller et al. (2018) offer high protocol fidelity but rely on stochastic channel abstractions that fail to model the impact of urban structures on signal propagation. To incorporate realistic environmental interactions, mobile network digital twins have recently been explored as more advanced evaluation environments. DT-SimNet Liu et al. (2024) introduces a digital twin-enabled simulator for large-scale coverage optimization with multi-agent reinforcement learning, while recent studies also investigate digital twin-assisted network automation and LLM-based network management Hong et al. (2024); Jiang et al. (2024). Nevertheless, existing digital twin platforms still tend to rely on coarse propagation models, fixed deployment scenarios, or single-layer network abstractions. Neither traditional protocol simulators nor current digital twin frameworks jointly address geographic generalizability, electromagnetic fidelity, RAN protocol completeness, and user behavior realism in one deployable system.

Wireless Channel Simulation Platforms. Geometry-aware wireless channel simulators improve wireless propagation fidelity by explicitly modeling how wireless signals interact with the physical environment. Wireless InSite Remcom (2024) is a commercial ray-tracing tool for site-specific wireless channel prediction. Sionna RT Hoydis et al. (2023) provides an open-source differentiable ray-tracing framework for wireless channel modeling, and DeepMIMO Alkhateeb (2019) converts ray-tracing outputs into structured MIMO channel datasets for AI research. However, these platforms are strictly confined to physical-layer channel generation. They are generally decoupled from user mobility, application traffic, and infrastructure-level operations. Besides, they typically require pre-built 3D scenes, whose construction remains costly and difficult to scale to arbitrary locations.

While these efforts have produced valuable tools, none have simultaneously achieved scene generalizability, joint modeling of channel and user behavior, and support for closed-loop optimization. Table 1 summarizes the key differences between MobiSim and representative existing platforms.

3 MobiSim: Generative AI-driven Mobile Network Digital Twin

MobiSim aims to enable efficient, full-stack mobile network digital twin construction at geographic locations with suitable map data and street-view imagery. Realizing this goal poses three fundamental challenges. **(i)** Wireless propagation is critically shaped by the 3D geometry of the surrounding built environment, yet conventional simulation approaches rely on simplified 2D maps or idealized propagation models that cannot capture the reflection, diffraction, and refraction processes governing real signal transmission. **(ii)** Accurately modeling the electromagnetic channel requires addressing two distinct physical phenomena: the near-field and far-field propagation effects that standard 3GPP

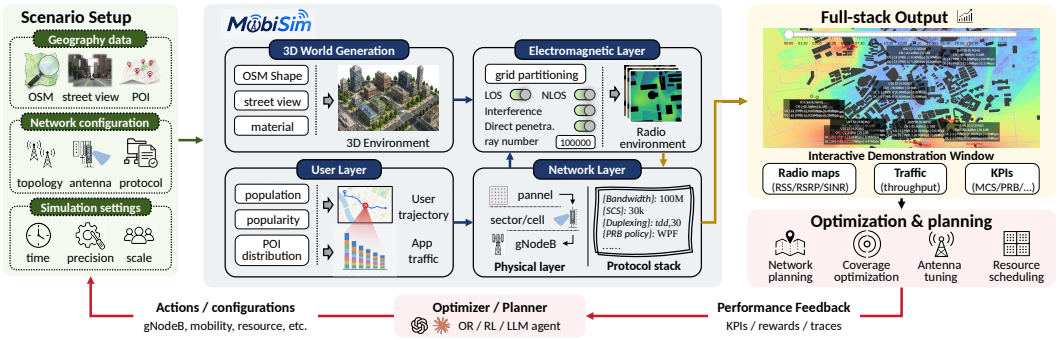

Figure 1: System model and workflow of MobiSim.

path loss formulas fail to reproduce, and the spatial incompleteness inherent in conventional Monte Carlo ray tracing, which under-samples far-field regions with small angular support and blocked regions unlikely to be reached by random rays, resulting in coverage holes and spatial discontinuities in the reconstructed radio map. **(iii)** A digital twin must go beyond static snapshots and reproduce the dynamic evolution of network performance over time, which demands accurate joint modeling of user spatiotemporal behavior, including mobility patterns, application usage, and traffic demands, as well as the interaction rules between users and network infrastructure under varying configurations and environmental conditions.

To address the above challenges, MobiSim is built around 3 core contributions. **(i)** We develop a generative AI-driven 3D world model that reconstructs building geometries and distributions for target regions with available OpenStreetMap data and street-view imagery, providing geographically grounded environments for electromagnetic simulation. **(ii)** We design a hybrid electromagnetic simulation framework that integrates deterministic direct propagation with an importance-sampling-based NLOS ray-tracing strategy, effectively eliminating coverage holes and spatial discontinuities to produce spatially complete radio maps. **(iii)** We construct a user behavior foundation model that jointly captures the spatiotemporal dependencies among mobility trajectories and network traffic demands within urban environments, coupled with a 5G RAN protocol stack comprising antenna configuration, radio resource management, and handover control modules, enabling high-fidelity simulation of network performance under dynamic configurations.

The overall architecture of MobiSim is illustrated in Fig. 1. Taking public geographic data and base station antenna configurations as inputs, MobiSim operates through four tightly integrated modules: 3D world generation, electromagnetic simulation, network layer emulation, and user behavior modeling. Given a target region, MobiSim efficiently generates synchronized outputs spanning user mobility traces, electromagnetic propagation maps, and network performance metrics. Beyond data generation, MobiSim exposes a standardized optimization interface that supports closed-loop network planning and optimization through reinforcement learning agents and LLM-based autonomous optimizers.

3.1 3D World Generation

The fidelity of ray-based electromagnetic simulations is heavily dependent on the geometry and material properties of the underlying 3D environment. Following recent advancements in large-scale 3D world generation Shang et al. (2024); Engstler et al. (2025); Wang et al. (2025); Lu et al. (2026); Huang et al. (2026b) and considering the specific characteristics of electromagnetics, we develop a three-stage framework leveraging 2D and 3D foundation models to construct realistic 3D environments (Figure 2a).

First, we extract geospatial metadata (e.g., 2D building footprints, road network topologies) from OpenStreetMap (OSM) along with their geographic coordinates to retrieve corresponding street-view images via online map APIs. Vision-Language Models (VLMs) are then applied to analyze these images and identify fine-grained elements such as vegetation and street furniture. During this process, street-view images are filtered to select optimal perspectives of target elements, minimizing occlusions from foreground obstacles.

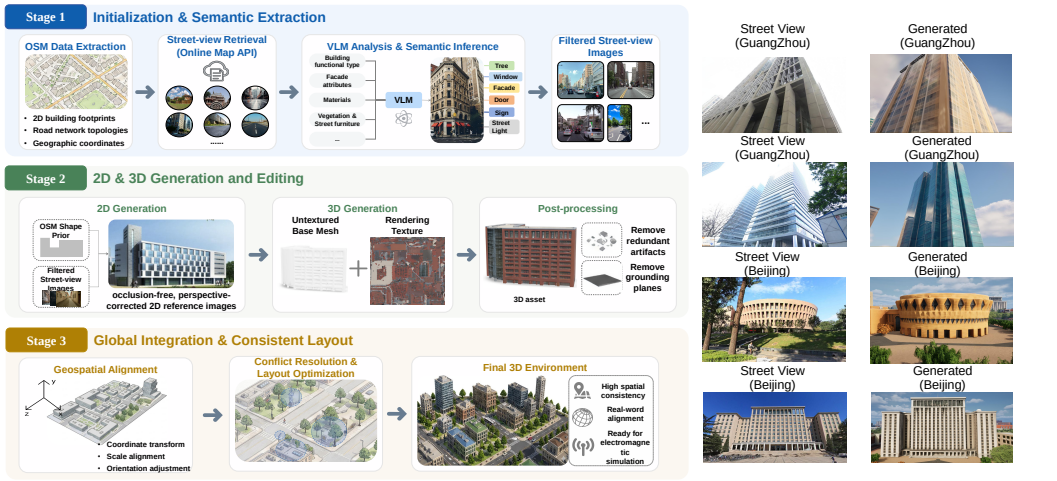

(a) Schematic diagram for 3D world generation

(b) Comparison of street views and generated buildings in Guangzhou and Beijing

Figure 2: The proposed framework of 3D world generation as well as comparisons between street-view images and generated buildings.

Next, a generation module powered by 2D and 3D foundation models produces 3D assets of the environmental elements. The 2D generative model fuses OSM metadata with visual attributes from the filtered images to generate perspective-corrected 2D reference images of urban elements (e.g., buildings). These reference images are then processed by a 3D foundation model to generate untextured base geometric meshes along with high-resolution physically-based rendering texture maps.

Finally, the generated 3D assets are integrated into a unified global coordinate system. We perform precise spatial alignment and procedural positioning to resolve spatial conflicts. Overall, this framework yields a highly consistent 3D environment accurately aligned with the real world, serving as a foundation for subsequent electromagnetic simulation. Details can be found at Appendix A. Implementation details are provided in Appendix A, and quantitative evaluation on Beijing and Guangzhou is reported in Appendix A.4.

3.2 Electromagnetic Layer

The electromagnetic layer characterizes the wireless propagation process from configured base stations to grid cells under a given 3D urban environment, supporting radio map computation, including RSRP and SINR. It adopts a hybrid ray-tracing pipeline that combines Monte Carlo NLOS tracing with deterministic direct components. For each transmitter and receiver grid cell pair (s, c), a path $\pi \in \mathcal{P}_{s,c}$ with K interactions is written as $\pi = (s,\ c,\ \mathbf{x}_{0:K+1},\ \chi_{1:K},\ \mathbf{J}_{1:K},\ \alpha_\pi)$, where $\chi_k \in \{\text{spec}, \text{diffuse}, \text{refrac}, \text{diffraction}\}$ denotes the interaction type, $\mathbf{J}_k \in \mathbb{C}^{2\times 2}$ is the Jones response, and α_π is the path accumulation weight. Detailed formulas are provided in Appendix C.

Monte Carlo NLOS Tracing. This module estimates the NLOS contribution and generates the path set $\mathcal{S}^{\text{mc}}_{s,c}$. Each source launches N rays with initial directions $\omega_0 \sim q_0(\omega_0)$, where q_0 is derived from the antenna pattern and directions with stronger antenna gain. At each ray-surface intersection, one interaction type χ_k is sampled from $\rho_k(\chi)$, which is obtained by normalizing a score combining the power of Jones response and geometric reachability to the receiver grid. The sampling compensations C_1 and C_2 are used to calculate C_{hit} when the ray is projected onto the receiver plane:

$$C_1(\omega_0) = \frac{1}{N q_0(\omega_0)}, \qquad C_2(\pi) = \prod_{k=1}^{K} \frac{1}{\rho_k(\chi_k) p_k^{\text{local}}}. \tag{1}$$

Deterministic Direct Propagation. The MC-NLOS branch can miss distant, blocked, or sparsely sampled grid cells under finite ray budgets and stochastic interaction sampling, resulting in spatially incomplete radio maps. To reduce such artifacts, we add a deterministic direct-propagation branch. **(1) Deterministic LOS** checks the direct segment for each transmitter-cell pair (s, c). If unobstructed, a LOS path is added to $\mathcal{S}^{\text{los}}_{s,c}$. **(2) Deterministic Direct Penetration** handles blocked target cells.

Under a thin-slab transmission assumption, the ray keeps its direction through the surfaces. The blocked direct segment is thus modeled as a pure-refraction path, with each intermediate surface hit treated as a refraction event and added to $\mathcal{S}^{\mathrm{pen}}_{s,c}$. To avoid double-counting, pure-refraction direct paths are excluded from the MC-NLOS set and handled by deterministic penetration. The final path set is decomposed as $\mathcal{S}_{s,c} = \mathcal{S}^{\mathrm{los}}_{s,c} \,\dot\cup\, \mathcal{S}^{\mathrm{pen}}_{s,c} \,\dot\cup\, \mathcal{S}^{\mathrm{nlos}}_{s,c}$, which determines the path weight α_π.
Ray-Based Channel Construction. The channel is modeled by tracing a set of propagation rays and accumulating their complex contributions at the receiver. Each ray carries amplitude and polarization information through a cascade of interaction-dependent transfer matrices, so that the end-to-end path coefficient takes the form

$$a = \mathbf{f}_{\mathrm{rx}}^{H} \mathbf{J}_{\mathrm{accum}} \mathbf{f}_{\mathrm{tx}} \cdot \frac{\lambda}{4\pi r}, \tag{2}$$

where $\mathbf{f}_{\mathrm{tx}}$ and $\mathbf{f}_{\mathrm{rx}}$ are the transmit and receive field vectors, $\mathbf{J}_{\mathrm{accum}}$ encodes the accumulated Jones response along the path, and r is the total path length. All valid ray paths are then integrated using Monte Carlo weighting to obtain an unbiased estimate of the channel transfer function $H(f) = \sum_n v_n$, where v_n is the weighted complex contribution of the n-th ray, and the effective channel gain is calculated as $\Gamma = |H(f)|^2$. The exact weighting rule is given in Appendix C.6.

Channel Quality Metrics. Given the channel gain Γ_m from transmitter m, the received signal strength is $\mathrm{RSS}_m = \Gamma_m P_{\mathrm{tx},m}$. The serving transmitter m^* maximizes RSS, and link quality is quantified by the SINR: $\mathrm{SINR} = \mathrm{RSS}_{m^*}/(I + N_0)$, where I is aggregate inter-cell interference and N_0 is thermal noise. In addition to SINR, the simulator reports RSRP, reflecting the signal strength at individual users and serving as the primary handover trigger in 5G NR.

3.3 Network Layer

Compliant with 3GPP specifications, the network layer reproduces the radio-access functionality of a 5G NR system, covering antenna modelling, mobility management, and physical-layer scheduling.

Physical layer architecture. The simulator mirrors the physical anatomy of a real 5G site through a five-level hierarchy: *antenna element* → *panel* → *panel array* → *sector* → *cell*, with one or more sectors aggregated by a *gNodeB*. The bottom levels capture the radiating geometry: elements follow the TR 38.901 pattern (3GPP, 2024i) and panels are half-wavelength $M \times N$ planar arrays of cross-polarized elements (Appendix B.2), while the upper levels turn this geometry into a steerable physical-layer entity by tiling panels into a coherent aperture exposing the closed-form steering vector of Eq. 9, binding it to an azimuth orientation and mechanical downtilt at the sector, and carrying carrier-specific radio-access state at the cell. This level-by-level encapsulation lets a researcher swap a single radiation pattern, panel layout, or scheduling policy without disturbing the rest of the stack.

Radio Resource and Mobility Management. Mobility follows the 3GPP Release-17 RRM specifications, covering idle-mode access (3GPP, 2024f;g) and connected-mode handover (3GPP, 2024h;b): a UE attaches to the strongest suitable cell via a four-step random-access procedure (Appendix B.3.1), and once connected its serving and neighbour measurements are continuously evaluated against the standard A1–A5 events (Table 6) with a configurable time-to-trigger, then dispatched through a global handover executor that models the Xn signalling latency and flags ping-pongs (Appendix B.3.3, Algorithm 1). On the air interface, each cell schedules uplink and downlink on a per-slot, time-division-duplex basis (3GPP, 2024c) (Appendix B.4.1): the PRB budget set by the TS 38.101-1 BW×SCS table (3GPP, 2024a) (Table 8) is split across UEs by a weighted proportional-fair allocator (Kelly et al., 1998; Capozzi et al., 2013) (Algorithm 2), and each admitted UE transmits a transport block whose size $\mathrm{TBS}(\cdot)$ is set by its PRB share and the MCS choice from the TS 38.214 CQI/MCS chain (3GPP, 2024d) (Eq. 22), yielding the per-slot delivered throughput

$$T_u^{(d)} \;=\; \mathrm{TBS}\big(N^{(d)}_{\mathrm{PRB},u},\, N_d,\, \mathrm{MCS}_u\big) \cdot \rho_s^{(d)} \,/\, T_{\mathrm{slot}}, \tag{3}$$

with $\rho_s^{(d)} = N_d/14$ the symbol fraction of direction $d \in \{\mathrm{DL}, \mathrm{UL}\}$ and T_{slot} the numerology-aware slot duration. Full algorithms, parameter values, and the per-block standards-conformance mapping are deferred to Appendix B.

3.4 User Layer

The user layer provides realistic user dynamics for wireless network simulation, including mobility trajectories, app usage, and traffic demands (including both uplink and downlink traffic). It adopts

a two-stage pipeline where a diffusion-based model first generates coarse-grained user dynamics, followed by map-constrained refinement via the MOSS simulator (Zhang et al., 2024) to ensure consistency with real-world mobility patterns. Detailed implementations are provided in Appendix D (Fig. 18).

For coarse-grained generation, we employ a diffusion-based framework to jointly model spatio-temporal user dynamics. Each location is represented by multi-source features (Yuan et al., 2026), including population density (Tatem, 2017), POI distributions (Haklay & Weber, 2008), and location popularity (Rong et al., 2026), which are encoded into compact embeddings via an autoencoder, while app usage and traffic are encoded using lightweight MLPs. These embeddings are fused and processed by the diffusion model, and decoded to produce coarse-grained user dynamics. The generated trajectories are further refined using the MOSS simulator by constructing a road-network-aware map from OpenStreetMap and simulating user movements at 1-second resolution along valid road segments, ensuring map consistency and physically plausible mobility for reliable wireless network simulation under fine-grained dynamics.

4 Evaluation

In this section, we evaluate the fidelity of MobiSim at the network, electromagnetic, and user behavior layers, as well as benchmarking MobiSim on 3 typical network planning and optimization tasks to demonstrate its utility as a robust evaluation environment for network intelligence research.

4.1 Electromagnetic Computation Fidelity

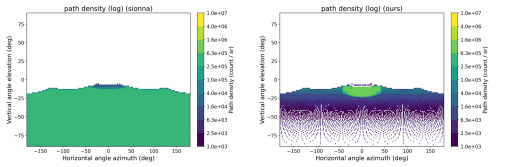

(a) Sionna RT (b) MobiSim

Model	α	β	γ	Weighted RMSE	Valid cells
FSPL	32.44	20.00	20.00	–	–
MobiSim	32.26	19.96	20.04	0.414 dB	99.97%
Sionna RT	31.71	20.70	20.42	5.323 dB	20.22%

(c) Path loss fitting results

Figure 3: Comparison of angular ray distributions and path loss fitting results in a single-wall scenario. The difference in ray distributions is mainly caused by antenna-aware initial ray sampling and the deterministic LOS branch in MobiSim.

To evaluate the electromagnetic fidelity of the proposed ray-tracing simulator, we consider a single-wall LoS benchmark and fit the simulated path loss using the standard logarithmic model

$$PL(d, f_c) = \alpha + \beta \log_{10}(d_{\mathrm{km}}) + \gamma \log_{10}(f_{c,\mathrm{MHz}}), \tag{4}$$

where α, β, and γ are the fitted coefficients. The theoretical FSPL coefficients are used as the reference, where $\alpha = 32.44$, $\beta = 20$, and $\gamma = 20$. The single-wall scenario places a wall to the right of the transmitter ($x > 0$), and the fitting is restricted to receivers in the unobstructed half-plane ($x < 0$), where the LoS paths do not intersect the wall. The fitting is performed using weighted least squares, where cell weights are proportional to d_{3D}^{-2}. Details are reported in Appendix C.8.

As shown in Figure 3c, fitted coefficients of MobiSim exhibit closer agreement with the theoretical FSPL coefficients than those of Sionna RT. MobiSim also achieves a lower weighted RMSE, indicating smaller fitting residuals under the distance-based weighting scheme. These results verify that MobiSim preserves the expected LoS propagation behavior in the unobstructed region.

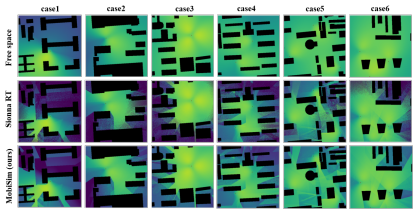

Case	Introduction
Case1	1 BS, 1 Sec.
Case2	1 BS, 2 Sec.
Case3	1 BS, 3 Sec.
Case4	2 BS, 2 Sec.
Case5	3 BS, 3 Sec.
Case6	3 BS, 5 Sec.

Figure 4: Comparison of radio maps generated by the free-space baseline, Sionna RT, and MobiSim.

We compare the RSRP radio maps of MobiSim with the free-space baseline and Sionna RT in Figure 4. The six cases are randomly selected $200\,\mathrm{m} \times 200\,\mathrm{m}$ urban regions, each discretized into a 400×400 receiver grid with a cell size of $0.5\,\mathrm{m}$. The free-space baseline produces smooth radio maps because it does not consider building blockage or electromagnetic interactions. In contrast, both Sionna RT and

MobiSim capture geometry-induced propagation effects, such as shadowing and multipath variations caused by interactions. Compared with Sionna RT, MobiSim yields spatially more complete radio maps, especially in distant or sparse coverage regions. By combining deterministic direct propagation with Monte Carlo NLOS tracing, MobiSim reduces such incompleteness and artifacts induced by ray sampling while preserving the main propagation structures.

Real-world data test. We further validate MobiSim against real millimeter-wave drive test measurements collected in an urban region of Phoenix, Arizona, providing direct evidence that MobiSim is faithful to real-world propagation. Details are reported in Appendix C.11.

4.2 Network-Layer Evaluation

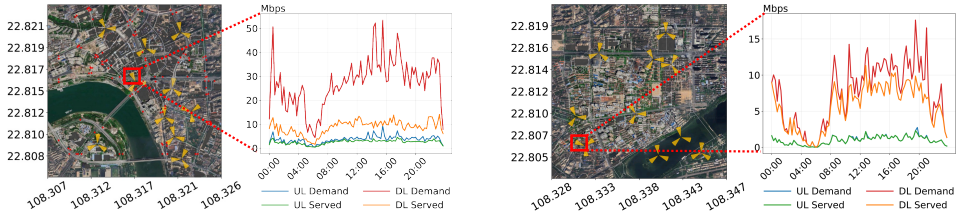

(a) Commercial area (The central urban area of Nanning) (b) Residential Area (Minzu Square)

Figure 5: Daily aggregated traffic for cells in specific urban functional areas generated by MobiSim (yellow sectors denote the simulated base station sectors and red dots indicate the simulated users).

We first evaluate the capability of MobiSim's network layer design to simulate the aggregated traffic patterns in urban environments. We conduct simulations in two representative $2\,\text{km} \times 2\,\text{km}$ regions of Nanning: (a) an area covering the central urban district, and (b) an area around Minzu Square and South Lake. In each region, we deploy 10 tri-sector gNodeB base stations and 2000 users, with sector parameters set to default configurations. Figure 5 presents the satellite images of the two regions, along with the aggregated UL and DL traffic demand and the corresponding served traffic of representative base stations. The selected sector in (a) covers mainly a commercial district, while the selected region in (b) covers mainly residential areas. The aggregated traffic exhibits a clear diurnal pattern. The traffic of these two sectors both remains high from 8:00 to 20:00. In the central commercial area, nighttime traffic remains active until approximately 4:00, reflecting the vibrant nighttime activities in the urban commercial district. In contrast, the traffic in residential areas decreases significantly after 22:00. These observations are consistent with the aggregated traffic patterns observed in real urban networks.

4.3 User Behavior Fidelity

We evaluate the proposed model on a real-world mobile network dataset collected in Shanghai, which includes anonymized user mobility, app usage, and traffic records. Detailed dataset statistics and an additional cross-city evaluation on the Nanchang dataset are provided in Appendix D. As shown in Table 2, our model achieves the lowest mean errors on six of eight metrics: Radius, Distance, and all four app-usage and traffic metrics. Pategail performs better on Duration and DailyLoc; the larger DailyLoc discrepancy suggests that our model may underrepresent occasional visits. Relative to AppGen, app-usage and traffic NMAE decrease by 6.0% and 5.0%, respectively, calculated as $(E_{\text{AppGen}} - E_{\text{ours}})/E_{\text{AppGen}}$. These gains can be attributed to shared latent representations that capture correlations among mobility, app usage, and traffic demand, which prior work typically models independently. The diffusion architecture further supports complex spatiotemporal distributions across modalities.

Table 2: User dynamics generation on Shanghai. Lower is better for all metrics.

(a) Mobility generation.

Method	KS-test			
	Radius	Distance	Duration	DailyLoc
EPR (Simini et al., 2012)	0.6283 ± 0.015	0.9104 ± 0.018	0.8884 ± 0.015	0.7963 ± 0.017
TimeGeo (Jiang et al., 2016)	0.5729 ± 0.013	0.6951 ± 0.013	0.7213 ± 0.012	0.6938 ± 0.014
MoveSim (Feng et al., 2020)	0.4372 ± 0.009	0.2836 ± 0.007	0.3729 ± 0.011	0.6582 ± 0.012
Pategail Wang et al. (2023)	0.7374 ± 0.005	0.2425 ± 0.002	**0.2260 ± 0.001**	**0.3930 ± 0.003**
DiffTraj Zhu et al. (2023)	0.4701 ± 0.004	0.2144 ± 0.003	0.2940 ± 0.003	0.4952 ± 0.005
Ours	**0.4164 ± 0.004**	**0.1449 ± 0.003**	0.2552 ± 0.003	0.7310 ± 0.005

(b) App usage and traffic generation.

Method	App Category		Traffic	
	JSD	NMAE	JSD	NMAE
Time-GAN (Yoon et al., 2019)	0.0285 ± 0.006	0.2939 ± 0.008	0.0325 ± 0.007	0.4829 ± 0.011
LDM (Rombach et al., 2022)	0.0264 ± 0.004	0.2706 ± 0.007	0.0287 ± 0.005	0.4361 ± 0.008
Diffusion-TS Yuan & Qiao (2024)	0.0256 ± 0.003	0.2573 ± 0.006	0.0272 ± 0.005	0.4312 ± 0.009
CSDI (Tashiro et al., 2021)	0.0249 ± 0.003	0.2427 ± 0.003	0.0268 ± 0.004	0.4157 ± 0.005
AppGen Huang et al. (2026a)	0.0229 ± 0.002	0.2196 ± 0.003	0.0265 ± 0.005	0.4091 ± 0.003
Ours	**0.0207 ± 0.002**	**0.2064 ± 0.003**	**0.0255 ± 0.004**	**0.3887 ± 0.002**

5 CASE STUDY: MOBILE PLANNING AND OPTIMIZATION

Problem 1: BS Deployment Planning Tarekegn et al. (2022); Mei & Zhang (2023). We consider a target region divided into regular grids, where the task is to deploy N BSs to improve coverage quality. Given a deployment configuration, MobiSim generates an RSRP map, and a grid is regarded as weakly covered if its center RSRP is below a predefined threshold. The objective is to minimize the average ratio of weakly covered grids by optimizing the deployment coordinates of all BSs. The constraints require BSs to be placed in feasible outdoor areas outside buildings and to maintain a minimum distance from each other.

Problem 2: Network Coverage Optimization (Tekgul et al., 2024; Vannella et al., 2024). We consider a cellular network with multiple deployed BSs, where each BS has 3 sectors with independently configurable antennas. The simulator evaluates the grid-level SINR according to the azimuth and downtilt angles of all sector antennas. The objective is to maximize the overall SINR in dB by optimizing these antenna configurations. The constraints require each angle to remain within its valid physical range and limit the number of antenna adjustments within a time window to avoid excessive hardware reconfiguration.

Problem 3: QoE-Driven Throughput Optimization (Boutiba et al., 2024; Song et al., 2022). We consider a multi-cell TDD network where each cell serves users with both uplink and downlink traffic demands. Given a selected TDD pattern and PRB allocation policy defined in the 3GPP protocol 3GPP (2024c), the simulator evaluates the achieved service rate of each user. The objective is to maximize the overall demand satisfaction ratio, measured by the ratio between achieved downlink throughput and traffic demand. The decision variables for each BS include the standardized TDD uplink/downlink slot pattern and the PRB allocation policy for scheduling radio resources among served users.

Table 3: Comparison of optimization methods for different network optimization tasks.

Optimization Task	Operations Research (Greedy)	Heuristic Algorithm (PSO)	Reinforcement Learning (PPO)	LLM-based Optimization (AlphaEvolve)
BS Deployment(↓)	0.003942	0.001207	/	0.001361
Coverage Optimization(↑)	13.2361	13.3058	15.7006	14.4201
Throughput Maximization(↑)	0.7483	0.7491	0.7906	0.8107

To provide a comprehensive benchmark, MobiSim evaluates each task under multiple optimization paradigms. BS deployment planning is a one-shot combinatorial problem, for which we benchmark a greedy algorithm, Particle Swarm Optimization (PSO), and an LLM-based optimizer. Network coverage optimization and QoE-driven throughput optimization, by contrast, involve sequential decision-making over time, making them natural candidates for Reinforcement Learning (RL); we therefore additionally include RL alongside the greedy, PSO, and LLM-based methods for these two tasks. Table 3 demonstrates MobiSim's use as an optimization environment across tasks and algorithm families. For BS deployment, PSO and the LLM-based optimizer outperform the greedy baseline (0.001207 and 0.001361 vs. 0.003942), supporting population-based and language-model-driven search. RL leads on coverage, while the LLM-based optimizer leads on throughput. These improvements indicate that MobiSim provides informative feedback for search and policy learning, supporting comparisons across optimization paradigms. More detailed results are provided in Appendix E.

6 CONCLUSION

MobiSim is a generative AI-driven mobile network digital twin jointly modeling radio, channel, and user layers. We evaluate its fidelity across 3D cities, wireless propagation, user behavior, and network operations. Benchmarks cover base station deployment, antenna adjustment, and radio resource scheduling, with physically grounded feedback for operations research, reinforcement learning, and LLM-based optimizers. We release the platform, benchmark datasets, and standardized evaluation interfaces to support reproducible mobile network research.

REFERENCES

3GPP. Study on 3D channel model for LTE. Technical Report TR 36.873, 3rd Generation Partnership Project (3GPP), 2017.

3GPP. NR; user equipment (UE) radio transmission and reception; part 1: Range 1 standalone. Technical Specification TS 38.101-1, 3rd Generation Partnership Project (3GPP), 2024a.

3GPP. NR; requirements for support of radio resource management. Technical Specification TS 38.133, 3rd Generation Partnership Project (3GPP), 2024b.

3GPP. NR; physical layer procedures for control. Technical Specification TS 38.213, 3rd Generation Partnership Project (3GPP), 2024c.

3GPP. NR; physical layer procedures for data. Technical Specification TS 38.214, 3rd Generation Partnership Project (3GPP), 2024d.

3GPP. NR; physical layer measurements. Technical Specification TS 38.215, 3rd Generation Partnership Project (3GPP), 2024e.

3GPP. NR; user equipment (UE) procedures in idle mode and in RRC inactive state. Technical Specification TS 38.304, 3rd Generation Partnership Project (3GPP), 2024f.

3GPP. NR; medium access control (MAC) protocol specification. Technical Specification TS 38.321, 3rd Generation Partnership Project (3GPP), 2024g.

3GPP. NR; radio resource control (RRC) protocol specification. Technical Specification TS 38.331, 3rd Generation Partnership Project (3GPP), 2024h.

3GPP. Study on channel model for frequencies from 0.5 to 100 GHz. Technical Report TR 38.901, 3rd Generation Partnership Project (3GPP), 2024i.

Lakshya A. Agrawal, Shangyin Tan, Dilara Soylu, Noah Ziems, Rishi Khare, Krista Opsahl-Ong, Arnav Singhvi, Herumb Shandilya, Michael J. Ryan, Meng Jiang, Christopher Potts, Koushik Sen, Alexandros G. Dimakis, Ion Stoica, Dan Klein, Matei Zaharia, and Omar Khattab. GEPA: Reflective prompt evolution can outperform reinforcement learning. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=RQm2KQTM5r. Oral presentation.

Ahmed Alkhateeb. Deepmimo: A generic deep learning dataset for millimeter wave and massive mimo applications. In *Proceedings of the Information Theory and Applications Workshop (ITA)*, pp. 1–8, San Diego, CA, February 2019. URL https://arxiv.org/abs/1902.06435.

Ahmed Alkhateeb, Gouranga Charan, Tawfik Osman, Andrew Hredzak, Joao Morais, Umut Demirhan, and Nikhil Srinivas. Deepsense 6g: A large-scale real-world multi-modal sensing and communication dataset. *IEEE Communications Magazine*, 61(9):122–128, 2023. doi: 10.1109/MCOM.006.2200730.

Abdulmalik Alwarafy, Mohamed Abdallah, Bekir Sait Ciftler, Ala Al-Fuqaha, and Mounir Hamdi. The frontiers of deep reinforcement learning for resource management in future wireless HetNets: Techniques, challenges, and research directions. *IEEE Open Journal of the Communications Society*, 3:322–365, 2022. doi: 10.1109/OJCOMS.2022.3153226.

Ram Asokan, Amir Hossein Fahim Raouf, Ozgur Ozdemir, İsmail Guvenc, and Mihail Sichitiu. Ericsson 5G NSA network RF and throughput measurements on AERPAW network. Dryad dataset, 2025. URL https://doi.org/10.5061/dryad.wh70rxx06. AERPAW Dataset 18.

Karim Boutiba, Miloud Bagaa, and Adlen Ksentini. Multi-agent deep reinforcement learning to enable dynamic TDD in a multi-cell environment. *IEEE Transactions on Mobile Computing*, 23 (5):6163–6177, 2024. doi: 10.1109/TMC.2023.3315961.

Francesco Capozzi, Giuseppe Piro, Luigi A. Grieco, Gennaro Boggia, and Pietro Camarda. Downlink packet scheduling in LTE cellular networks: Key design issues and a survey. *IEEE Communications Surveys & Tutorials*, 15(2):678–700, 2013. doi: 10.1109/SURV.2012.060912.00100.

Haoye Chai, Xiaoqian Qi, Yibo Ma, Zhaocheng Wang, Lei Yue, and Yong Li. Mobifm: A foundation model for mobile data forecasting. *IEEE Journal on Selected Areas in Communications*, 44: 2494–2509, 2026. doi: 10.1109/JSAC.2025.3642851.

Paul Engstler, Aleksandar Shtedritski, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Syncity: Training-free generation of 3D worlds. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 27585–27595, 2025. doi: 10.1109/ICCV51701.2025.02561.

Dayu Fan, Rui Meng, Xiaodong Xu, Yiming Liu, Guoshun Nan, Chenyuan Feng, Shujun Han, Song Gao, Bingxuan Xu, Dusit Niyato, Tony Q. S. Quek, and Ping Zhang. Generative diffusion models for wireless networks: Fundamental, architecture, and state-of-the-art. *IEEE Communications Surveys & Tutorials*, 28:5632–5677, 2026. doi: 10.1109/COMST.2026.3671110.

Jie Feng, Zeyu Yang, Fengli Xu, Haisu Yu, Mudan Wang, and Yong Li. Learning to simulate human mobility. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 3426–3433, 2020. doi: 10.1145/3394486.3412862.

Mordechai Haklay and Patrick Weber. Openstreetmap: User-generated street maps. *IEEE Pervasive Computing*, 7(4):12–18, 2008. doi: 10.1109/MPRV.2008.80.

Thomas R. Henderson, Mathieu Lacage, George F. Riley, Craig Dowell, and Joseph Kopena. Network simulations with the ns-3 simulator. In *ACM SIGCOMM 2008 Conference on Data Communication Demonstration*, 2008.

Yang Hong, Jun Wu, and Rosario Morello. LLM-Twin: mini-giant model-driven beyond 5G digital twin networking framework with semantic secure communication and computation. *Scientific Reports*, 14(1):19065, 2024. doi: 10.1038/s41598-024-69474-5.

Jakob Hoydis, Fayçal Ait Aoudia, Sebastian Cammerer, Merlin Nimier-David, Nikolaus Binder, Guillermo Marcus, and Alexander Keller. Sionna RT: Differentiable ray tracing for radio propagation modeling. In *2023 IEEE Globecom Workshops (GC Wkshps)*, pp. 317–321, 2023. doi: 10.1109/GCWkshps58843.2023.10465179.

Zihan Huang, Tong Li, and Yong Li. Appgen: Mobility-aware app usage behavior generation for mobile users. *IEEE Transactions on Mobile Computing*, 25(2):1903–1917, 2026a. doi: 10.1109/TMC.2025.3605939.

Zilong Huang, Jun He, Xiaobin Huang, Ziyi Xiong, Yang Luo, Junyan Ye, Weijia Li, Yiping Chen, and Ting Han. Majutsucity: Language-driven aesthetic-adaptive city generation with controllable 3D assets and layouts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 31608–31618, 2026b.

ITU-R. Effects of building materials and structures on radiowave propagation above about 100 MHz. Recommendation ITU-R P.2040-3, International Telecommunication Union (ITU), 2023.

Shan Jiang, Yingxiang Yang, Siddharth Gupta, Daniele Veneziano, Shounak Athavale, and Marta C González. The timegeo modeling framework for urban mobility without travel surveys. *Proceedings of the National Academy of Sciences*, 113(37):E5370–E5378, 2016. doi: 10.1073/pnas.1524261113.

Shufan Jiang, Bangyan Lin, Yue Wu, and Yuan Gao. LINKs: Large language model integrated management for 6G empowered digital twin networks. In *2024 IEEE 100th Vehicular Technology Conference (VTC2024-Fall)*, pp. 1–6, 2024. doi: 10.1109/VTC2024-Fall63153.2024.10757470.

Florian Kaltenberger, Aloizio P. Silva, Abhimanyu Gosain, Luhan Wang, and Tien-Thinh Nguyen. Openairinterface: Democratizing innovation in the 5g era. *Computer Networks*, 176:107284, 2020. doi: 10.1016/j.comnet.2020.107284.

Joseph B. Keller. Geometrical theory of diffraction. *Journal of the Optical Society of America*, 52(2): 116–130, February 1962. doi: 10.1364/JOSA.52.000116.

Frank P. Kelly, Aman K. Maulloo, and David K. H. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49 (3):237–252, 1998. doi: 10.1057/palgrave.jors.2600523.

R.G. Kouyoumjian and P.H. Pathak. A uniform geometrical theory of diffraction for an edge in a perfectly conducting surface. *Proceedings of the IEEE*, 62(11):1448–1461, 1974. doi: 10.1109/PROC.1974.9651.

LAION-AI. LAION aesthetic predictor. https://github.com/LAION-AI/aesthetic-predictor, n.d.

Haoqiang Liu, Tong Li, Fenyu Jiang, Weikang Su, and Zhaocheng Wang. Coverage optimization for large-scale mobile networks with digital twin and multi-agent reinforcement learning. *IEEE Transactions on Wireless Communications*, 23(12):18316–18330, 2024. doi: 10.1109/TWC.2024.3464639.

Qingyue Long, Yuan Yuan, and Yong Li. A universal model for human mobility prediction. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1*, KDD '25, pp. 894–905, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400712456. doi: 10.1145/3690624.3709236. URL https://doi.org/10.1145/3690624.3709236.

Keyang Lu, Sifan Zhou, Hongbin Xu, Gang Xu, Zhifei Yang, Yikai Wang, Zhen Xiao, Jieyi Long, and Ming Li. Yo'city: Personalized and boundless 3d realistic city scene generation via self-critic expansion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3219–3230, 2026.

Weidong Mei and Rui Zhang. Joint base station and irs deployment for enhancing network coverage: A graph-based modeling and optimization approach. *IEEE Transactions on Wireless Communications*, 22(11):8200–8213, 2023. doi: 10.1109/TWC.2023.3260805.

Martin Klaus Müller, Fjolla Ademaj, Thomas Dittrich, Agnes Fastenbauer, Blanca Ramos Elbal, Armand Nabavi, Lukas Nagel, Stefan Schwarz, and Markus Rupp. Flexible multi-node simulation of cellular mobile communications: The vienna 5g system level simulator. *EURASIP Journal on Wireless Communications and Networking*, 2018(1):227, 2018. doi: 10.1186/s13638-018-1238-7.

Navid Nikaein, Mahesh K. Marina, Saravana Manickam, Alex Dawson, Raymond Knopp, and Christian Bonnet. Openairinterface: A flexible platform for 5g research. *SIGCOMM Comput. Commun. Rev.*, 44(5):33–38, October 2014. ISSN 0146-4833. doi: 10.1145/2677046.2677053. URL https://doi.org/10.1145/2677046.2677053.

Alexander Novikov, Ngân Vũ, Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt Wagner, Sergey Shirobokov, Borislav Kozlovskii, Francisco J. R. Ruiz, Abbas Mehrabian, M. Pawan Kumar, Abigail See, Swarat Chaudhuri, George Holland, Alex Davies, Sebastian Nowozin, Pushmeet Kohli, and Matej Balog. AlphaEvolve: A coding agent for scientific and algorithmic discovery. *arXiv preprint arXiv:2506.13131*, 2025. URL https://arxiv.org/abs/2506.13131.

Natale Patriciello, Sandra Lagen, Biljana Bojovic, and Lorenza Giupponi. An e2e simulator for 5g nr networks. *Simulation Modelling Practice and Theory*, 96:101933, 2019. doi: 10.1016/j.simpat.2019.101933.

Michele Polese, Mischa Dohler, Falko Dressler, Melike Erol-Kantarci, Rittwik Jana, Raymond Knopp, and Tommaso Melodia. Empowering the 6g cellular architecture with open ran. *IEEE Journal on Selected Areas in Communications*, 42(2):245–262, 2024. doi: 10.1109/JSAC.2023.3334610.

Qwen Team. Qwen3.5-397B-A17B. Model card, 2026. URL https://modelscope.cn/models/Qwen/Qwen3.5-397B-A17B.

Naina Raisinghani. Nano banana 2: Combining pro capabilities with lightning-fast speed. Google Blog, February 2026. URL https://blog.google/innovation-and-ai/technology/ai/nano-banana-2/.

Remcom. Wireless insite: 3d wireless prediction software. Software documentation, 2024. Commercial ray-tracing software for wireless propagation prediction.

George F. Riley and Thomas R. Henderson. *The ns-3 Network Simulator*, pp. 15–34. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. ISBN 978-3-642-12331-3. doi: 10.1007/978-3-642-12331-3_2. URL https://doi.org/10.1007/978-3-642-12331-3_2.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.

Can Rong, Jingtao Ding, Meng Li, and Yong Li. A global intra-city commuting origin-destination flow dataset for urban sustainable development. *Scientific Data*, 13(1):1269, 2026. doi: 10.1038/s41597-026-07279-z.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. URL https://arxiv.org/abs/1707.06347.

Yu Shang, Yuming Lin, Yu Zheng, Hangyu Fan, Jingtao Ding, Jie Feng, Jiansheng Chen, Li Tian, and Yong Li. UrbanWorld: An urban world model for 3D city generation. *arXiv preprint arXiv:2407.11965*, 2024. URL https://arxiv.org/abs/2407.11965.

Yuan Shen, Wei-Chiu Ma, and Shenlong Wang. SGAM: Building a virtual 3D world through simultaneous generation and mapping. In *Advances in Neural Information Processing Systems*, volume 35, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/8ae9cf363ea625161f885b798c1f1f78-Abstract-Conference.html.

Enyu Shi, Jiayi Zhang, Hongyang Du, Bo Ai, Chau Yuen, Dusit Niyato, Khaled B. Letaief, and Xuemin Shen. Ris-aided cell-free massive mimo systems for 6g: Fundamentals, system design, and applications. *Proceedings of the IEEE*, 112(4):331–364, 2024. doi: 10.1109/JPROC.2024.3404491.

Sebastian Siebert and Jochen Teizer. Mobile 3d mapping for surveying earthwork projects using an unmanned aerial vehicle (uav) system. *Automation in construction*, 41:1–14, 2014. doi: 10.1016/j.autcon.2014.01.004.

Filippo Simini, Marta C González, Amos Maritan, and Albert-László Barabási. A universal model for mobility and migration patterns. *Nature*, 484(7392):96–100, 2012. doi: 10.1038/nature10856.

Simran Singh. Dataset-23: Android-based 4G LTE, 5G NR, and throughput measurements for two sweeps of a UAV near a private AERPAW base station. AERPAW dataset, North Carolina State University, 2025a. URL https://aerpaw.org/dataset/dataset-23-android-based-4g-lte-5g-nr-and-throughput-measurements-for-two-sweeps-

Simran Singh. Dataset-24: Detailed cellular and throughput measurements for horizontal sweeps of a UAV across a private AERPAW base station, using Keysight Nemo and PawPrints. AERPAW dataset, North Carolina State University, 2025b. URL https://aerpaw.org/dataset/dataset-24-android-based-4g-lte-measurements-for-semi-circular-uav-trajectory-arou

Jing Song, Qingyang Song, Ya Kang, Lei Guo, and Abbas Jamalipour. QoE-driven distributed resource optimization for mixed reality in dynamic TDD systems. *IEEE Transactions on Communications*, 70(11):7294–7306, 2022. doi: 10.1109/TCOMM.2022.3208113.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition, 2018.

Pingbo Tang, Daniel Huber, Burcu Akinci, Robert Lipman, and Alan Lytle. Automatic reconstruction of as-built building information models from laser-scanned point clouds: A review of related techniques. *Automation in construction*, 19(7):829–843, 2010. doi: 10.1016/j.autcon.2010.06.007.

Zihan Tao, Haoyu Wang, Hanke Feng, Yijun Guo, Bitao Shen, Dan Sun, Yuansheng Tao, Changhao Han, Yandong He, John E. Bowers, Haowen Shu, Cheng Wang, and Xingjun Wang. Ultrabroadband on-chip photonics for full-spectrum wireless communications. *Nature*, 645:80–87, 2025. doi: 10.1038/s41586-025-09451-8.

Getaneh Berie Tarekegn, Rong-Terng Juang, Hsin-Piao Lin, Yirga Yayeh Munaye, Li-Chun Wang, and Mekuanint Agegnehu Bitew. Deep-reinforcement-learning-based drone base station deployment for wireless communication services. *IEEE Internet of Things Journal*, 9(21):21899–21915, 2022. doi: 10.1109/JIOT.2022.3182633.

Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. Csdi: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in neural information processing systems*, 34:24804–24816, 2021.

Andrew J. Tatem. Worldpop, open data for spatial demography. *Scientific Data*, 4(1):170004, 2017. doi: 10.1038/sdata.2017.4.

Ezgi Tekgul, Thomas Novlan, Salam Akoum, and Jeffrey G. Andrews. Joint uplink–downlink capacity and coverage optimization via site-specific learning of antenna settings. *IEEE Transactions on Wireless Communications*, 23(5):4032–4048, 2024. doi: 10.1109/TWC.2023.3313916.

Filippo Vannella, Alexandre Proutiere, Yassir Jedra, and Jaeseong Jeong. Learning optimal antenna tilt control policies: A contextual linear bandits approach. *IEEE Transactions on Mobile Computing*, 23(12):12666–12679, 2024. doi: 10.1109/TMC.2024.3424192.

Huandong Wang, Changzheng Gao, Yuchen Wu, Depeng Jin, Lina Yao, and Yong Li. Pategail: a privacy-preserving mobility trajectory generator with imitation learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pp. 14539–14547, 2023.

Xinliang Wang, Yijie Kang, Zhenyu Wu, and Yifeng Shi. Sat2RealCity: Geometry-aware and appearance-controllable 3D urban generation from satellite imagery. *arXiv preprint arXiv:2511.11470*, 2025. URL https://arxiv.org/abs/2511.11470.

Zhou Wang, Eero P. Simoncelli, and Alan C. Bovik. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems and Computers*, volume 2, pp. 1398–1402, 2003. doi: 10.1109/ACSSC.2003.1292216.

Tong Wu, Zhiyong Chen, Dazhi He, Liang Qian, Yin Xu, Meixia Tao, and Wenjun Zhang. Cddm: Channel denoising diffusion models for wireless semantic communications. *IEEE Transactions on Wireless Communications*, 23(9):11168–11183, 2024. doi: 10.1109/TWC.2024.3379244.

Jianfeng Xiang, Xiaoxue Chen, Sicheng Xu, Ruicheng Wang, Zelong Lv, Yu Deng, Hongyuan Zhu, Yue Dong, Hao Zhao, Nicholas Jing Yuan, and Jiaolong Yang. Native and compact structured latents for 3D generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14419–14429, 2026.

Haozhe Xie, Zhaoxi Chen, Fangzhou Hong, and Ziwei Liu. Citydreamer: Compositional generative model of unbounded 3d cities. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9666–9675, 2024. doi: 10.1109/CVPR52733.2024.00923.

Haozhe Xie, Zhaoxi Chen, Fangzhou Hong, and Ziwei Liu. Compositional generative model of unbounded 4D cities. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(1): 312–328, 2026. doi: 10.1109/TPAMI.2025.3603078.

Han Qing Yang, Jun Yan Dai, Hui Dong Li, Lijie Wu, Zi Hang Shen, Qun Yan Zhou, Meng Zhen Zhang, Si Ran Wang, Zheng Xing Wang, Jun Wei Wu, Shi Jin, Wankai Tang, Qiang Cheng, and Tie Jun Cui. Adaptively programmable metasurface for intelligent wireless communications in complex environments. *Nature Communications*, 16(1):6070, 2025. doi: 10.1038/s41467-025-61409-6. URL https://www.nature.com/articles/s41467-025-61409-6.

Jinsung Yoon, Daniel Jarrett, and Mihaela Van der Schaar. Time-series generative adversarial networks. *Advances in neural information processing systems*, 32, 2019.

Donghan Yu, Yong Li, Fengli Xu, Pengyu Zhang, and Vassilis Kostakos. Smartphone app usage prediction using points of interest. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(4):1–21, 2018. doi: 10.1145/3161413.

Xinyu Yuan and Yan Qiao. Diffusion-TS: Interpretable diffusion for general time series generation. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=4h1apFjO99.

Yuan Yuan, Yuheng Zhang, Jingtao Ding, and Yong Li. Worldmove, a global open data for human mobility. *Scientific Data*, 13(1):549, 2026. doi: 10.1038/s41597-026-06555-2.

Jun Zhang, Wenxuan Ao, Junbo Yan, Can Rong, Depeng Jin, Wei Wu, and Yong Li. MOSS: A large-scale open microscopic traffic simulation system. *arXiv preprint arXiv:2405.12520*, 2024. URL https://arxiv.org/abs/2405.12520.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 586–595, 2018. URL https://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_The_Unreasonable_Effectiveness_CVPR_2018_paper.html.

Yuanshao Zhu, Yongchao Ye, Shiyao Zhang, Xiangyu Zhao, and James Yu. Difftraj: Generating GPS trajectory with diffusion probabilistic model. In *Advances in Neural Information Processing Systems*, volume 36, pp. 65168–65188, 2023. doi: 10.52202/075280-2843.

APPENDIX CONTENTS

A 3D World Generation: Detailed Implementation **18**

A.1 Initialization & Semantic Extraction 18

A.2 2D & 3D Generation and Editing 18

A.3 Global Integration & Consistent Layout 18

A.4 Quantitative Evaluation of 3D World Generation 19

B Network Layer: Detailed Implementation **20**

B.1 Overview 20

B.2 Physical Layer: Hierarchical Site and Antenna Modelling 21

B.3 Mobility: Access, Reselection and Handover 22

B.4 Per-Slot Resource Scheduling 23

B.5 Exposed Outputs and Interfaces 27

C Electromagnetic Layer: Detailed Implementation **28**

C.1 Overview 28

C.2 Radio Material 28

C.3 Antenna-aware Initial Ray Sampling 29

C.4 Importance Sampling in Monte Carlo NLOS Tracing 30

C.5 Deterministic Direct Propogation 33

C.6 Ray-Based Channel Construction 35

C.7 Channel Quality Metrics. 36

C.8 Electromagnetic Fidelity Validation with Sionna RT 37

C.9 Radio Map Comparison and visualization with Sionna RT 38

C.10 Radio Map Comparison under Different Sampling Budgets 38

C.11 Real-World Validation Test 40

C.12 Ablation on Material Properties 42

D User Layer: Detailed Implementation **42**

D.1 Overview 42

D.2 Model Architecture Details 43

D.3 Evaluation Methodology 43

D.4 Additional Results 44

E Mobile Planning and Optimization: Detailed Implementation **46**

E.1 Optimization Algorithms 46

E.2 BS Deployment Planning 47

E.3 Network Coverage Optimization 47

E.4 QoE-Driven Throughput Optimization 48

E.5 Optimization Results of RL and LLM-based Schemes 49

F Experiment Setting **52**

F.1 Hardware Platform 52

F.2 Default Global Simulation Parameters 52

F.3 Default Network Configurations 52

F.4 Computation Time Complexity 53

F.5 Interactive UI 54

G Generated Dataset Description **54**

G.1 Dataset Overview 54

G.2 Radio-Map Collection 56

G.3 Full-Stack Network Collection 56

G.4 3D City Collection 57

H Potential Impacts and Limitation **57**

A 3D World Generation: Detailed Implementation

This appendix describes the three stages of the 3D world generation framework illustrated in Figure 2: (1) Initialization & Semantic Extraction, (2) 2D & 3D Generation and Editing, and (3) Global Integration & Consistent Layout. We then quantitatively evaluate the generated scenes in Beijing and Guangzhou in Appendix A.4.

A.1 Initialization & Semantic Extraction

The first stage focuses on extracting geospatial metadata from OpenStreetMap (OSM), including 2D building footprints, road network topologies, and geographic coordinates. These coordinates serve as spatial anchors to retrieve corresponding street-view images via online map APIs. Subsequently, we leverage Qwen3.5 Qwen Team (2026), a large-scale Vision-Language Model, to analyze these images, inferring semantic information absent from the raw OSM data and identifying fine-grained environmental elements (e.g., vegetation and street furniture). During this process, the street-view images are systematically filtered to select optimal perspectives of the target environmental elements, thereby minimizing occlusions caused by foreground obstacles.

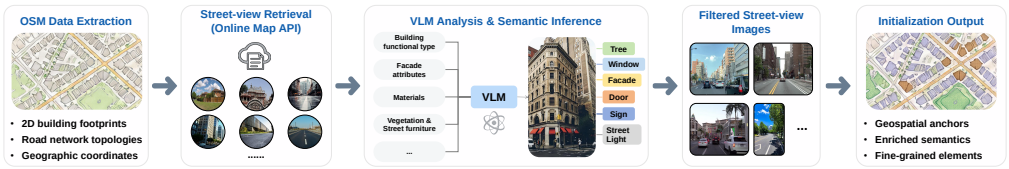

Figure 6: Diagram for the stage 1 of 3D world generation.

A.2 2D & 3D Generation and Editing

Then, based on the obtained metadata and street-view images, we incorporate a generation module powered by 2D and 3D foundation models to generate 3D assets of environmental elements. First, we employ Nano Banana 2 Raisinghani (2026), and feed the filtered street-view images into it. Within its visual latent space, the geospatial metadata obtained from OSM are fused with the physical attributes captured in the street views to synthesize occlusion-free perspective-corrected 2D reference images of the urban element (e.g., buildings). These reference images are then processed by a large flow-matching 3D foundation model based on the O-Voxel representation Xiang et al. (2026), which generates untextured base geometric meshes alongside their corresponding high-resolution physically-based rendering texture maps. Finally, an automated geometric post-processing is adopted to remove redundant artifacts and ground planes, producing topologically sound and materially accurate 3D assets.

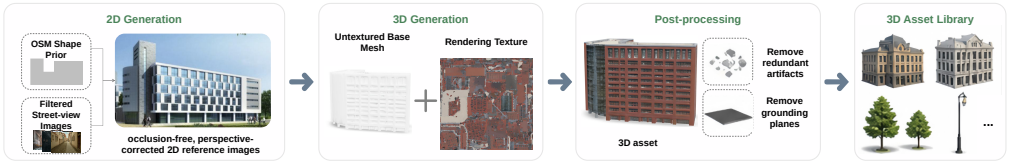

Figure 7: Diagram for the stage 2 of 3D world generation.

A.3 Global Integration & Consistent Layout

The final stage integrates the generated 3D assets into a unified global coordinate system. Specifically, we perform spatial transformations on each 3D model to ensure accurate geometric registration in the physical space. Concurrently, the obtained 3D assets are procedurally positioned to resolve potential spatial conflicts and layout intersections. Overall, this framework yields a highly consistent 3D environment accurately aligned with the real world, serving as a foundation for subsequent electromagnetic simulation.

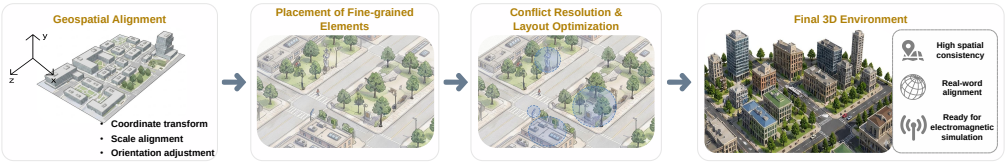

Figure 8: Diagram for the stage 3 of 3D world generation.

Table 4: Quantitative evaluation of 3D world generation on Beijing and Guangzhou. Higher is better for all metrics; bold denotes the best score within each city. Baselines use OSM inputs, whereas MobiSim additionally uses street-view images. These image-based scores assess rendered appearance and structure, rather than direct 3D geometric accuracy.

City	Method	Perceptual Fidelity ↑	Boundary Alignment (%) ↑	MS-SSIM ↑	Aesthetic Fitness ↑
Beijing	SGAM (Shen et al., 2022)	0.282	3.143	0.218	5.143
	CityDreamer (Xie et al., 2024)	0.395	6.755	0.085	4.741
	CityDreamer4D (Xie et al., 2026)	0.399	**7.948**	0.089	5.372
	UrbanWorld (Shang et al., 2024)	0.447	6.813	0.281	4.730
	MobiSim	**0.451**	7.842	**0.302**	**5.983**
Guangzhou	SGAM (Shen et al., 2022)	0.175	1.586	0.091	2.665
	CityDreamer (Xie et al., 2024)	0.145	2.089	0.116	3.785
	CityDreamer4D (Xie et al., 2026)	0.152	2.320	0.042	4.213
	UrbanWorld (Shang et al., 2024)	0.114	2.133	0.189	2.824
	MobiSim	**0.880**	**7.717**	**0.806**	**6.682**

A.4 Quantitative Evaluation of 3D World Generation

A.4.1 Evaluation Setting

We evaluate the rendered appearance and structural consistency of generated urban scenes in Beijing and Guangzhou. The reference images for the image-based fidelity evaluation are the latest satellite imagery available at evaluation time. This quantitative evaluation complements the representative street-view and generated-asset examples in Figure 2b, which provide a qualitative illustration of the generation pipeline.

We compare MobiSim with SGAM (Shen et al., 2022), CityDreamer (Xie et al., 2024), CityDreamer4D (Xie et al., 2026), and UrbanWorld (Shang et al., 2024). In this comparison, the baselines use OSM inputs, while MobiSim additionally uses street-view images retrieved through the Baidu Maps API. The results therefore compare the complete generation pipelines under their respective input conditions; they do not isolate the contribution of the generation method from that of additional visual conditioning.

A.4.2 Evaluation Metrics

We report four complementary image-based metrics, all with higher values indicating better performance. **Perceptual Fidelity** measures similarity to the reference image using the transformed LPIPS (Zhang et al., 2018) score $\max\{0, 1 - \mathrm{LPIPS}\}$, on a scale from 0 to 1. **Boundary Alignment** reports the overlap between generated and reference boundary maps as a percentage, capturing agreement in visible building outlines, roof contours, roads, and facade edges. **MS-SSIM** (Wang et al., 2003) measures structural similarity across image scales. **Aesthetic Fitness** uses the LAION Aesthetic Predictor (LAION-AI, n.d.) to score visual quality on a scale from 0 to 10; it complements the reference-based metrics but does not measure correspondence to the target scene.

A.4.3 Results and Interpretation

Table 4 shows that MobiSim obtains the highest scores on three of the four metrics in Beijing and all four metrics in Guangzhou. In Beijing, its Perceptual Fidelity is only slightly higher than UrbanWorld's, while MS-SSIM increases from 0.281 to 0.302. CityDreamer4D achieves the highest Boundary Alignment. In Guangzhou, MobiSim's Perceptual Fidelity and MS-SSIM reach 0.880 and 0.806, respectively, compared with the strongest baseline scores of 0.175 and 0.189.

The improvements thus vary substantially across the two cities. The temporal alignment of the input and reference imagery is a possible contributor to this difference. The available Beijing street-view images were captured four years earlier than those for Guangzhou, whereas the reference satellite imagery is more recent. Differences in imagery age and city-specific architecture may affect the value of street-view conditioning. However, the present comparison does not control these factors independently and cannot establish which factor causes the larger gains in Guangzhou.

Overall, the results provide evidence of improved rendered-image fidelity under the evaluated input conditions. They do not directly establish accurate building heights, three-dimensional surface geometry, or material properties, and aesthetic quality alone does not establish physical fidelity. These image-based evaluations therefore complement the propagation validation in Appendix C.11, which assesses a different aspect of the generated environment's suitability for simulation.

B Network Layer: Detailed Implementation

B.1 Overview

This appendix gives a self-contained description of the network layer of MobiSim, the part of the simulator that turns the propagation channel produced by the electromagnetic layer into the radio-access metrics that a 5G NR system would actually expose to its core network and to user equipment (UE). The layer reproduces, in line with 3GPP Release-17, four classes of behaviour, organised into four mutually decoupled blocks:

- **Hierarchical site and antenna geometry** (§B.2). The radiating geometry of a base station, from individual antenna elements up to a complete sectorised gNodeB, together with the closed-form steering vector through which the electromagnetic layer enters the physical-layer scheduler.
- **Mobility** (§B.3). Idle-mode cell selection and 4-step random access, idle-mode reselection across intra- and inter-frequency neighbours, and connected-mode handover driven by the standard A1–A5 measurement-event family.
- **Per-slot resource scheduling** (§B.4). TDD slot-format catalogue, bandwidth-to-PRB mapping, and a weighted proportional-fair allocator that selects modulation-and-coding, transport-block size, and uplink transmit power on a per-slot basis.
- **Exposed outputs** (§B.5). A description of the metric trace that benchmark tasks consume from the simulator.

The four blocks are intentionally decoupled. The closed-form steering vector of §B.2 is the only contract through which the electromagnetic layer enters the physical-layer scheduler; the mobility stack of §B.3 only exchanges measurement reports and attach/detach signalling with the per-slot scheduler of §B.4. As a consequence, a researcher can replace a radiation pattern, a panel layout, an A1–A5 threshold profile, or the proportional-fair metric in isolation, without recompiling or revalidating the rest of the stack.

Every functional block of the network layer is implemented in conformance with a specific 3GPP Release-17 technical specification. Table 5 consolidates this mapping, listing for each block the governing document together with the clause and the role it plays in MobiSim; the relevant clauses are also cited inline at their point of use in the subsections below. Anchoring each block to a single, uniquely identifiable specification clause makes it unambiguous which behaviour MobiSim reproduces, and which deviations (if any) are deliberate modelling choices rather than implementation artefacts.

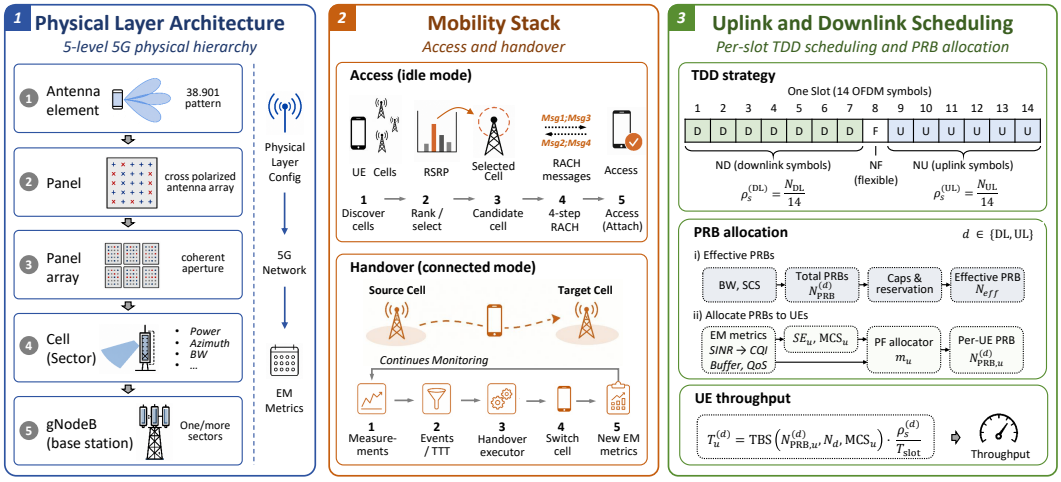

Figure 9: Architecture of the network layer.

Table 5: Mapping of MobiSim network-layer functional blocks to 3GPP technical specifications. Each "Standard" cell names the document on the first line and the relevant clause and role on the second line.

Functional block	Standard
Antenna element and panel array	TR 38.901 (3GPP, 2024i) *Table 7.3-1 / Fig. 7.3-1: element pattern, array layout*
Resource grid sizing	TS 38.101-1 (3GPP, 2024a) *Table 5.3.2-1: BW × SCS ⇒ N_{PRB}*
Per-slot scheduler (PF/RR/MaxT)	TS 38.214 (3GPP, 2024d) *§5: link-to-system mapping, MCS*
Slot-format catalogue	TS 38.213 (3GPP, 2024c) *Table 11.1.1-1: Slot Formats 0–40*
Uplink power control	TS 38.213 (3GPP, 2024c) *§7.1: open-loop PUSCH power control*
UE measurements (RSRP/RSRQ/SINR)	TS 38.215 (3GPP, 2024e) *§5.1: UE measurements*
Measurement events A1–A5	TS 38.331 (3GPP, 2024h) *§5.5.4: measurement events*
Handover execution	TS 38.133 (3GPP, 2024b) *§6: handover delay budget*
Initial cell selection (S-criterion)	TS 38.304 (3GPP, 2024f) *§5.2.3: initial cell selection*
Idle-mode reselection	TS 38.304 (3GPP, 2024f) *§5.2.4: idle-mode reselection*
Random access (4-step contention)	TS 38.321 (3GPP, 2024g) *§5.1: random access*

B.2 Physical Layer: Hierarchical Site and Antenna Modelling

The radiation pattern of a single element follows TR 38.901 Table 7.3-1 (3GPP, 2024i): the vertical and horizontal cuts are

$$A_{E,V}(\theta) = -\min\left(12\left(\frac{\theta-90^\circ}{\theta_{3\mathrm{dB}}}\right)^2, \mathrm{SLA}_V\right), \tag{5}$$

$$A_{E,H}(\phi) = -\min\left(12\left(\frac{\phi}{\phi_{3\mathrm{dB}}}\right)^2, A_{\max}\right), \tag{6}$$

combined as

$$A_E(\theta,\phi) \;=\; -\min\big(-(A_{E,V}+A_{E,H}),\, A_{\max}\big) + G_{E,\max}, \tag{7}$$

with $\theta_{3\mathrm{dB}} = \phi_{3\mathrm{dB}} = 65°$, $\mathrm{SLA}_V = A_{\max} = 30$ dB, $G_{E,\max} = 8$ dBi. Cross-polarised elements share the same magnitude pattern but rotate the field by $\pm 45°$, with field components

$$f_\theta = \sqrt{A_E}\cos(\zeta), \qquad f_\phi = \sqrt{A_E}\sin(\zeta) \tag{8}$$

for slant angle ζ. A panel array is parametrised by (M, N, M_g, N_g), the per-panel rows/columns and the inter-panel grid, plus horizontal and vertical inter-element spacings (d_H, d_V, defaults 0.5λ) and inter-panel spacings ($d_{g,H}, d_{g,V}$). Element positions are first laid out symmetrically in wavelength units and then scaled by λ_0 to absolute metres. With $\mathbf{p}_n$ denoting the resulting position of the n-th element and $\hat{\mathbf{k}}(\theta,\phi)$ the unit propagation direction, the array exposes the closed-form steering vector

$$a_n(\theta,\phi) \;=\; \exp\left(j\,\tfrac{2\pi}{\lambda}\,\mathbf{p}_n \cdot \hat{\mathbf{k}}(\theta,\phi)\right), \tag{9}$$

which is the explicit interface between the electromagnetic layer and the physical-layer scheduler: ray-traced multipath components are projected onto $a(\theta,\phi)$ to assemble the per-link channel matrix consumed by the cell. Equation 9 is evaluated in closed form for arbitrary (θ,ϕ), after which the sector applies a global-to-local rotation matrix which is built from azimuth orientation and mechanical downtilt, and returns the per-element gain plus the coherent array gain $\approx 10\log_{10}(N_{\mathrm{ant}})$ in the steered direction. Sectors are aggregated under a gNodeB, which exposes one or more co-located cells (one per carrier) to the rest of the stack.

B.3 Mobility: Access, Reselection and Handover

The mobility stack reproduces three concerns: how an unattached UE first joins the network (§B.3.1), how an attached but idle UE re-evaluates its serving cell (§B.3.2), and how a connected UE is migrated between cells when its radio conditions warrant it (§B.3.3). All three concerns share the same per-UE measurement primitives, including RSRP, RSRQ, and SINR, which are sampled at the UE position from the radio map produced by the electromagnetic layer.

B.3.1 Initial Access and Random Access

Cell selection (TS 38.304 §5.2.3, (3GPP, 2024f)). A candidate cell is suitable if and only if the level and quality criteria:

$$S_{\mathrm{rxlev}} \;=\; Q_{\mathrm{rxlevmeas}} - (Q_{\mathrm{rxlevmin}} + Q_{\mathrm{rxlevminoffset}}) - P_{\mathrm{compensation}} \;>\; 0, \tag{10}$$

$$S_{\mathrm{qual}} \;=\; Q_{\mathrm{qualmeas}} - (Q_{\mathrm{qualmin}} + Q_{\mathrm{qualminoffset}}) \;>\; 0 \tag{11}$$

both hold. Among suitable cells, the one with the highest RSRP is selected; default thresholds are $Q_{\mathrm{rxlevmin}} = -124$ dBm and $Q_{\mathrm{qualmin}} = -20$ dB.

4-step random access (TS 38.321, (3GPP, 2024g)). The four-step contention-based procedure is modelled explicitly: Msg1 (preamble) draws from a 64-preamble pool with collision probability $1-(1-1/M)^{N_{\mathrm{ue}}-1}$; Msg2 (RAR) is gated by an RAR-window timeout (default 10 ms); Msg3 collisions are resolved by uniform-random tie-breaking; Msg4 verifies contention resolution within a 64 ms timer. Failed attempts re-arm with a 2 dB power-ramping step, a back-off window drawn from $\{0, B, 2B, \ldots\}$, and a maximum of $\mathrm{preambleTransMax} = 10$ retries, after which the access is reported as failed and surfaced to the KPI trace as a RACH failure.

B.3.2 Idle-Mode Reselection

Reselection is invoked only on UEs already camped on a cell in idle mode and is governed by TS 38.304 §5.2.4 (3GPP, 2024f).

Intra-frequency. When $S_{\mathrm{rxlev}} < S_{\mathrm{IntraSearchP}}$ (default 10 dB), neighbours are ranked against the serving cell using a hysteresis-protected RSRP metric,

$$R_n \;=\; \mathrm{RSRP}_n, \qquad R_s \;=\; \mathrm{RSRP}_{\mathrm{serv}} - Q_{\mathrm{hyst}}, \tag{12}$$

with default $Q_{\mathrm{hyst}} = 4$ dB. The highest-ranked neighbour is reselected if it has held the lead for $T_{\mathrm{reselection}}$ (default 1 s).

Inter-frequency. Targets at higher priority are reselected as soon as $S^{(n)}_{\mathrm{rxlev}} > \mathrm{Thresh_X_HighP}$; equal-priority targets fall back to the ranking rule of Eq. 12; lower-priority targets are considered only when the serving cell drops below $\mathrm{Thresh_Serving_LowP}$ and the target exceeds $\mathrm{Thresh_X_LowP}$. A speed-dependent scaling factor on $T_{\mathrm{reselection}}$ extends the timer at medium/high mobility, suppressing oscillations on fast trajectories.

B.3.3 Connected-Mode Handover: A1–A5 Events

For every connected UE, the simulator continuously evaluates the five measurement events of TS 38.331 §5.5.4 (3GPP, 2024h), listed in Table 6. The leaving condition of each event is the negation of its entering condition with the hysteresis Hys applied with the opposite sign, and an event only fires after its triggering condition has held continuously for the configured time-to-trigger (TTT). All thresholds are configurable per UE through a per-UE measurement profile.

Table 6: Entering conditions for the A1–A5 measurement events implemented in MobiSim. M_s, M_n are the serving and neighbour measurement quantities (RSRP by default); Of, Oc are frequency- and cell-specific offsets; Hys is the event hysteresis; Off is the A3 offset; T, T_1, T_2 are configurable thresholds. The leaving condition is the negation of the entering condition.

Event	Operational role	Entering condition
A1	Stop the inter-frequency measurement gap	$M_s - \mathrm{Hys} > T$
A2	Start the inter-frequency measurement gap	$M_s + \mathrm{Hys} < T$
A3	Intra-frequency handover	$M_n + \mathrm{Of}_n + \mathrm{Oc}_n - \mathrm{Hys} > M_s + \mathrm{Of}_s + \mathrm{Oc}_s + \mathrm{Off}$
A4	Priority-driven inter-frequency handover	$M_n + \mathrm{Of}_n + \mathrm{Oc}_n - \mathrm{Hys} > T$
A5	Coverage-driven inter-frequency handover	$M_s + \mathrm{Hys} < T_1$ and $M_n + \mathrm{Of}_n + \mathrm{Oc}_n - \mathrm{Hys} > T_2$

Default parameters and event ordering. $\mathrm{Hys} = 2$ dB, $\mathrm{Off} = 3$ dB, $\mathrm{TTT} = 160$ ms (A3) / 640 ms (A2/A5); A2 threshold -110 dBm, A1 threshold -95 dBm. For each UE, the events are evaluated in the order A1→A2→A3→A4→A5; A1 disables the inter-frequency measurement gap and resets the A4/A5 timers, while A2 enables it, thus A4/A5 reports can only occur in the A2-active state. Generated handover requests are arbitrated by an event priority A4 > A3 > A5 with ties broken by neighbour RSRP, and submitted to a global handover executor with a configurable Xn-interface delay (default 50 ms, (3GPP, 2024b)) before the actual source-to-target detach/attach transition; ping-pong handovers (return to source within T_{pp}, default 5 s) are flagged for downstream statistics. The full per-slot evaluation loop is summarised in Algorithm 1.

B.4 Per-Slot Resource Scheduling

The scheduler runs once per OFDM slot at every cell. It partitions the cell's PRB budget across attached UEs and produces, per slot and per direction $d \in \{\mathrm{DL}, \mathrm{UL}\}$, an admitted UE set, a PRB allocation, an MCS choice, a transport-block size, and a delivered throughput. Section B.4.1 fixes the slot timing structure on which the budget is built; Section B.4.2 maps a cell's bandwidth and subcarrier spacing to its nominal PRB count; Section B.4.3 defines the joint UL/DL allocation algorithm and its supporting equations. Throughout this subsection, we use the symbol fraction

$$\rho_s^{(d)} \;=\; N_d/14, \qquad d \in \{\mathrm{DL}, \mathrm{UL}\}, \tag{13}$$

where N_{DL} and N_{UL} are the effective per-slot symbol counts of the active TDD slot format after the flexible-symbol reassignment defined in Section B.4.1, and the numerology-aware slot duration $T_{\mathrm{slot}} = 10^{-3}/(\mathrm{SCS}/15)$ s (1 ms at 15 kHz, 0.5 ms at 30 kHz, etc.).

B.4.1 Slot Format and Numerology

All 41 slot formats from TS 38.213 Table 11.1.1-1 (3GPP, 2024c) are stored as (D, U, F)-encoded 14-symbol patterns. The triplet (N_D, N_U, N_F) defines the static skeleton of the slot, while the flexible

Algorithm 1 Per-slot mobility evaluation for a connected UE u.

Require: Serving cell s, neighbour set $\mathcal{N}(u)$, measurements $\{M_s, \{M_n\}_{n\in\mathcal{N}(u)}\}$, per-event timers and active flags $\{(\tau_k, a_k)\}_{k\in\{A1,\dots,A5\}}$, hysteresis Hys, offsets $\{\mathrm{Of.}, \mathrm{Oc.}, \mathrm{Off}\}$, thresholds $\{T, T_1, T_2\}$, time-to-trigger TTT_k, slot duration Δt, Xn delay τ_{Xn}, ping-pong window T_{pp}.
Ensure: Handover request (or none) submitted to the global executor.

```
Sample M_s and {M_n} from the radio map at the current UE position.
for event k in the order A1, A2, A3, A4, A5 do
    Evaluate the entering condition C_k^enter from Table 6.
    if C_k^enter and ¬a_k then
        τ_k ← τ_k + Δt.
        if τ_k ≥ TTT_k then
            a_k ← true; emit candidate event for k.
        end if
    else if the leaving condition ¬C_k^enter |_{Hys→−Hys} holds then
        a_k ← false, τ_k ← 0.
    end if
end for
Gap management: if a_A1 just became true, reset τ_A4, τ_A5 and disable the inter-frequency gap; if a_A2 just became true, enable it.
Among emitted events, arbitrate by priority A4 > A3 > A5 with ties broken by neighbour M_n.
Submit the chosen event to the global handover executor.
After delay τ_Xn, detach u from s and attach to the chosen target n*; if u left n* within the past T_pp, flag the transition as a ping-pong.
return the event (if any).
```

budget N_F is reassigned at runtime to whichever direction currently has the lower load satisfaction ratio, so that idle symbols are never wasted when one direction is undersubscribed. Table 7 lists the full enumeration of formats 0–40; the default Format 28 corresponds to the special slot of the widely deployed $n78$ DDDSU period.

Concretely, let $T_{c,t}^{(d)}$ denote the throughput delivered by cell c in direction $d \in \{\mathrm{DL}, \mathrm{UL}\}$ during the previous EMA window and $D_{c,t}^{(d)}$ the corresponding buffered demand aggregated over its attached UEs. The instantaneous load satisfaction ratio for direction d at slot t is

$$\eta_{c,t}^{(d)} \;=\; \min\left\{ \frac{T_{c,t}^{(d)}}{\max\left(D_{c,t}^{(d)},\, D_{\mathrm{floor}}\right)},\; 1 \right\}, \tag{14}$$

with a small floor D_{floor} that disables reassignment when both queues are essentially empty. The flexible budget is then routed in full to the lagging direction,

$$N_{c,t}^{F,\mathrm{DL}} \;=\; N_F \cdot \mathbb{1}\left[\eta_{c,t}^{\mathrm{DL}} \le \eta_{c,t}^{\mathrm{UL}}\right], \qquad N_{c,t}^{F,\mathrm{UL}} \;=\; N_F - N_{c,t}^{F,\mathrm{DL}}, \tag{15}$$

so that the effective per-slot symbol counts entering Eq. 13 become $N_{\mathrm{DL}} = N_D + N_{c,t}^{F,\mathrm{DL}}$ and $N_{\mathrm{UL}} = N_U + N_{c,t}^{F,\mathrm{UL}}$, and the throughput-scaling fractions are updated accordingly to $f_D = N_{\mathrm{DL}}/14$ and $f_U = N_{\mathrm{UL}}/14$. Ties (including the empty-buffer case where $\eta_{c,t}^{\mathrm{DL}} = \eta_{c,t}^{\mathrm{UL}} = 1$) are broken toward downlink, matching the typical traffic asymmetry of mobile networks; one symbol on each side of the flexible region is always retained as guard period to absorb DL-to-UL switching transients before the reassignment is applied.

B.4.2 Bandwidth, Subcarrier Spacing and PRB Count

Table 8 reproduces the full TS 38.101-1 Table 5.3.2-1 (FR1) and TS 38.101-2 Table 5.3.2-1 (FR2) mapping from channel bandwidth and subcarrier spacing to the nominal number of PRBs N_{PRB}; the cell's resource grid is sized from this table at construction time. A misconfigured deployment that picks a $(\mathrm{BW}, \mathrm{SCS})$ pair that the standards do not define raises an explicit error rather than silently rounding to the nearest legal numerology.

Table 7: Complete enumeration of TS 38.213 slot formats implemented in MobiSim. Each pattern is the 14-symbol sequence (D = downlink, U = uplink, F = flexible/guard period) of a single slot.

Idx	Pattern	N_D	N_U	N_F	Idx	Pattern	N_D	N_U	N_F
0	DDDDDDDDDDDDDD	14	0	0	21	DDDFFFFFFFFFFU	3	1	10
1	UUUUUUUUUUUUUU	0	14	0	22	DFFFFFFFFFFFUU	1	2	11
2	FFFFFFFFFFFFFF	0	0	14	23	DDFFFFFFFFFFUU	2	2	10
3	DDDDDDDDDDDDDF	13	0	1	24	DDDFFFFFFFFFUU	3	2	9
4	DDDDDDDDDDDDFF	12	0	2	25	DFFFFFFFFFFUUU	1	3	10
5	DDDDDDDDDDDFFF	11	0	3	26	DDFFFFFFFFFUUU	2	3	9
6	DDDDDDDDDDFFFF	10	0	4	27	DDDFFFFFFFFUUU	3	3	8
7	DDDDDDDDDFFFFF	9	0	5	28	DDDDDDDDDDDDFU	12	1	1
8	FFFFFFFFFFFFFU	0	1	13	29	DDDDDDDDDDDFFU	11	1	2
9	FFFFFFFFFFFFUU	0	2	12	30	DDDDDDDDDDFFFU	10	1	3
10	FUUUUUUUUUUUUU	0	13	1	31	DDDDDDDDDDDFUU	11	2	1
11	FFUUUUUUUUUUUU	0	12	2	32	DDDDDDDDDDFFUU	10	2	2
12	FFFUUUUUUUUUUU	0	11	3	33	DDDDDDDDDFFFUU	9	2	3
13	FFFFUUUUUUUUUU	0	10	4	34	DFUUUUUUUUUUUU	1	12	1
14	FFFFFUUUUUUUUU	0	9	5	35	DDFUUUUUUUUUUU	2	11	1
15	FFFFFFUUUUUUUU	0	8	6	36	DDDFUUUUUUUUUU	3	10	1
16	DFFFFFFFFFFFFF	1	0	13	37	DFFUUUUUUUUUUU	1	11	2
17	DDFFFFFFFFFFFF	2	0	12	38	DDFFUUUUUUUUUU	2	10	2
18	DDDFFFFFFFFFFF	3	0	11	39	DDDFFUUUUUUUUU	3	9	2
19	DFFFFFFFFFFFFU	1	1	12	40	DFFFUUUUUUUUUU	1	10	3
20	DDFFFFFFFFFFFU	2	1	11					

Table 8: Channel bandwidth (MHz) $\times$ subcarrier spacing (kHz) $\Rightarrow$ nominal N_{PRB}, after TS 38.101-1 Table 5.3.2-1 (FR1) and TS 38.101-2 Table 5.3.2-1 (FR2). Entries marked "—" are not defined by the standard.

	N_{PRB} **at SCS (kHz)**			
BW (MHz)	**15**	**30**	**60**	**120**
FR1 (sub-7 GHz)				
5	25	11	—	—
10	52	24	11	—
15	79	38	18	—
20	106	51	24	—
25	133	65	31	—
30	160	78	38	—
40	216	106	51	—
50	270	133	65	—
60	—	162	79	—
70	—	189	93	—
80	—	217	107	—
90	—	245	121	—
100	—	273	135	—
FR2 (mmWave)				
50	—	—	66	32
100	—	—	132	66
200	—	—	264	132
400	—	—	—	264

B.4.3 Per-Slot UL/DL Scheduling Algorithm

A single slot of joint downlink–uplink scheduling, summarised in Algorithm 2, invokes the same weighted proportional-fair PRB allocator once per direction $d \in \{\mathrm{DL}, \mathrm{UL}\}$, with the symbol fraction $\rho_s^{(d)}$ of Eq. 13 entering as a sector-level scaling factor on the resulting throughput. Before each

per-direction call, the cell shrinks its nominal PRB pool to the effective budget

$$N_{\text{eff}}^{(d)} \;=\; \big\lfloor N_{\text{PRB}} \cdot s \cdot c^{(d)} \cdot (1-r) \big\rfloor, \tag{16}$$

where $s \in [0,1]$ is a global utilisation scale, $c^{(d)} \in [0,1]$ is the per-direction utilisation cap, and $r \in [0, 0.95]$ is the reservation fraction; all three are configurable per cell. The effective budget is then split across UEs in proportion to a weighted proportional-fair score that augments the canonical $\mathrm{SE}/\bar{R}$ ratio with QoS, buffer-occupancy, and starvation terms,

$$N_{\text{PRB},u}^{(d)} \;\approx\; N_{\text{eff}}^{(d)} \cdot \frac{m_u^{(d)}}{\sum_{v \in \mathcal{S}^{(d)}} m_v^{(d)}}, \quad m_u^{(d)} \;=\; w_u^{\text{QoS}} \, \frac{\mathrm{SE}_u^{\beta}}{\bar{R}_u^{\gamma}} \left(\frac{B_u^{(d)}}{B_{\text{ref}}} \right)^{\delta} \left(1 + \kappa \, \frac{s_u}{s_{\text{ref}}} \right), \tag{17}$$

where $\mathcal{S}^{(d)}$ is the admitted UE set, SE_u is the spectral efficiency obtained from the CQI/MCS chain of TS 38.214 (3GPP, 2024d), $\bar{R}_u$ is the EMA-smoothed past throughput, $B_u^{(d)}$ is the buffered demand and s_u counts slots since u was last served. Each admitted UE then transmits one transport block whose size $\mathrm{TBS}(\cdot)$ is determined by its PRB share, the direction's symbol count N_d, and the modulation-and-coding choice MCS_u derived from its SINR via the TS 38.214 CQI/MCS chain (3GPP, 2024d). The per-slot delivered throughput in either direction is

$$T_u^{(d)} \;=\; \mathrm{TBS}\big(N_{\text{PRB},u}^{(d)},\, N_d,\, \mathrm{MCS}_u\big) \cdot \rho_s^{(d)} \,/\, T_{\text{slot}}, \tag{18}$$

with the standard DM-RS / control-channel overhead deductions baked into the TBS quantisation defined below. Algorithm 2 summarises the resulting per-slot processing.

Weighted PF metric and EMA update. The historical rate $\bar{R}_u^{(d)}$ used in Eq. 17 follows the canonical exponential moving average

$$\bar{R}_u^{(d)} \;\leftarrow\; (1-\alpha)\, \bar{R}_u^{(d)} + \alpha \, T_u^{(d),\text{served}}, \tag{19}$$

with $\alpha = 0.01$ and a small floor $\bar{R}_{\min} = 10^{-3}$ Mbps that prevents division-by-zero before a UE has ever been served. The buffer term $B_u^{(d)}/B_{\text{ref}}$ is clipped to $[10^{-9}, 1]$ so that a UE with at least B_{ref} Mbits queued (default $B_{\text{ref}} = 10$ Mbits) is treated as fully backlogged, while the QoS weight w_u^{QoS} is read from the UE's traffic profile and the starvation guard $\kappa \, s_u / s_{\text{ref}}$ recovers strict PF when $\kappa = 0$. Plain PF corresponds to $w_u^{\text{QoS}} = \delta = \kappa = 0$, $\beta = \gamma = 1$; round-robin and max-throughput are obtained by zeroing the appropriate exponents (Kelly et al., 1998; Capozzi et al., 2013).

Integer PRB allocation. Once $\mathcal{S}^{(d)}$ has been chosen, raw PRB shares are

$$\tilde{n}_u \;=\; N_{\text{eff}}^{(d)} \cdot \frac{m_u^{(d)}}{\sum_{v \in \mathcal{S}^{(d)}} m_v^{(d)}}, \qquad n_u^{(0)} \;=\; \mathrm{clip}\big(\lfloor \tilde{n}_u \rfloor,\; n_{\min},\; n_{\max}\big), \tag{20}$$

with a per-UE floor $n_{\min}$ (default 1 PRB) and ceiling $n_{\max} = \lfloor N_{\text{eff}}^{(d)} \cdot \eta \rfloor$ (default $\eta = 1$, i.e. no ceiling). When $\sum_u n_u^{(0)} > N_{\text{eff}}^{(d)}$ the allocator iteratively decrements the lowest-metric UE that still exceeds the floor; the remaining $N_{\text{eff}}^{(d)} - \sum_u n_u^{(0)}$ PRBs are then distributed to the UEs with the largest fractional residue $\tilde{n}_u - \lfloor \tilde{n}_u \rfloor$, capped by $n_{\max}$. The procedure is mass-conservative by construction.

Uplink power control. The TS 38.213 §7.1 (3GPP, 2024c) open-loop PUSCH power-control law

$$P_{\text{UE}} \;=\; \min\left(P_{\text{MAX}},\; P_0 + \alpha_{\text{PC}} \cdot \text{PL} + 10 \log_{10} M_{\text{RB}}\right) \quad [\text{dBm}], \tag{21}$$

is evaluated for every admitted UL UE, with $P_{\text{MAX}} = 23$ dBm (Power Class 3), $P_0 = -80$ dBm and a fractional-compensation factor $\alpha_{\text{PC}} = 0.8$. The resulting power-budget headroom is folded into the configurable bias Δ_{UL} that maps the downlink SINR (obtained by TDD reciprocity (3GPP, 2017)) to its uplink counterpart in step 3.2 of Algorithm 2.

Throughput and TBS quantisation. Equation 18 replaces the heuristic $(1-\eta)$ overhead factor with the canonical NR transport-block-size quantisation of TS 38.214 §5.1.3.2 (3GPP, 2024d): for $N_{\text{PRB},u}^{(d)}$ PRBs and $N_{\text{sym}} = N_d$ symbols per slot we form

$$N_{\text{RE}} \;=\; \min\big(156,\; 12\, N_{\text{sym}} - N_{\text{DMRS}} - N_{\text{oh}}\big) \cdot N_{\text{PRB},u}^{(d)}, \quad N_{\text{info}} \;=\; N_{\text{RE}} \cdot R \cdot Q_m \cdot L, \tag{22}$$

Algorithm 2 Per-slot UL/DL scheduling at a single cell at slot index t.

Require: Attached UE set $\mathcal{U}$, per-UE measurements $\{(\mathrm{RSRP}_u, \mathrm{SINR}_u^{DL})\}_{u\in\mathcal{U}}$, buffered demand $\{B_u^{DL}, B_u^{UL}\}_{u\in\mathcal{U}}$, slot format (N_D, N_U, N_F), EMA throughput $\{\bar{R}_u^{(d)}\}$, starvation counters $\{s_u^{(d)}\}$.

Ensure: Per-UE PRB allocation $\{N_{\mathrm{PRB},u}^{(d)}\}$, MCS choice $\{\mathrm{MCS}_u\}$, throughput $\{T_u^{DL}, T_u^{UL}\}$ pushed to the user-equipment state, and a cell-level result record.

Reset the per-slot PRB grid and clear per-UE PRB counts.
Compute the load satisfaction ratios $\eta_{c,t}^{(\mathrm{DL})}, \eta_{c,t}^{(\mathrm{UL})}$ via Eq. 14 and route the flexible budget N_F to the lagging direction via Eq. 15, yielding the effective symbol counts $N_{\mathrm{DL}} = N_D + N_{c,t}^{F,\mathrm{DL}}$ and $N_{\mathrm{UL}} = N_U + N_{c,t}^{F,\mathrm{UL}}$.
Compute symbol fractions $\rho_s^{(d)} = N_d/14$ for $d \in \{\mathrm{DL}, \mathrm{UL}\}$ via Eq. 13.
for direction $d \in \{\mathrm{DL}, \mathrm{UL}\}$ with $N_d > 0$ **do**
    Compute the effective PRB budget $N_{\mathrm{eff}}^{(d)}$ via Eq. 16.
    for $u \in \mathcal{U}$ with $B_u^{(d)} > 0$ and finite SINR **do**
        $\mathrm{SINR}_u^{(d)} \leftarrow \mathrm{SINR}_u^{DL} + \mathbb{1}[d{=}\mathrm{UL}] \cdot \Delta_{\mathrm{UL}}$ dB.
        $\mathrm{CQI}_u \leftarrow \textsc{SinrToCqi}(\mathrm{SINR}_u^{(d)})$; $\mathrm{MCS}_u \leftarrow \textsc{CqiToMcs}(\mathrm{CQI}_u)$; $\mathrm{SE}_u \leftarrow \textsc{McsToSe}(\mathrm{MCS}_u)$.
        Assemble the weighted-PF metric $m_u^{(d)}$ of Eq. 17.
    end for
    Sort active UEs by $m_u^{(d)}$ descending; admit the top $\min(|\mathcal{U}_{\mathrm{act}}|, N_{\mathrm{eff}}^{(d)})$ as $\mathcal{S}^{(d)}$.
    Compute integer PRB counts $N_{\mathrm{PRB},u}^{(d)}$ via Eq. 20.
    for each admitted u **do**
        Compute $T_u^{(d)}$ from Eq. 18 with $N_{\mathrm{sym}} = N_d$ and $\rho_s^{(d)}$.
    end for
    Update EMA $\bar{R}_u^{(d)}$ via Eq. 19; reset $s_u^{(d)} \leftarrow 0$ if served, else $s_u^{(d)}$ += 1.
end for
for each admitted u **do**
    Push $(N_{\mathrm{PRB},u}^{(d)}, \mathrm{MCS}_u, T_u^{DL}, T_u^{UL})$ to the user-equipment state.
end for
Aggregate cell-level totals $\sum_u T_u^{DL}$ and $\sum_u T_u^{UL}$ into the result record.
return the result record.

with $N_{\mathrm{DMRS}} = 12$ RE/PRB (one front-loaded DM-RS symbol), $N_{\mathrm{oh}} = 0$ by default (configurable for PDCCH/PT-RS/CSI-RS), R the target code rate, Q_m the modulation order from TS 38.214 Table 5.1.3.1, and L the MIMO layer count (currently $L = 1$, with multi-layer support enabled by the steering vector of Eq. 9). N_{info} is then quantised to a TBS through the 93-entry small-block table for $N_{\mathrm{info}} \le 3824$ and through the closed-form large-block rule (with the $R \le 1/4$, $N'_{\mathrm{info}} > 8424$ and default branches) of TS 38.214 §5.1.3.2 otherwise. An optional goodput multiplier $(1 - \mathrm{BLER}_{\mathrm{tgt}})$ may be enabled to convert nominal throughput into expected post-HARQ goodput.

B.5 Exposed Outputs and Interfaces

The network layer exposes its state to the rest of the simulator (and to benchmark optimisation tasks) through three families of traces:

1. *Per-link* measurements, RSRP, RSRQ, and SINR, are produced by sampling the radio map of the electromagnetic layer at the UE position via the steering vector of Eq. 9;
2. *Per-slot* scheduling outputs are emitted by Algorithm 2 at every cell, including the admitted UE set, the per-UE PRB allocation, the MCS index, the served transport-block size, and the downlink/uplink throughput of Eq. 18;
3. *Per-event* mobility outputs, including triggered A1–A5 events, executed handovers (including the source/target pair and the Xn delay actually incurred), ping-pong flags, and RACH

success/failure with attempt count, are emitted by Algorithm 1 and the random-access procedure of §B.3.1.

All three families share a common time index (slot or sub-slot) so that throughput, mobility, and access traces of the same UE can be correlated without additional bookkeeping. No additional learned or calibrated steps are inserted between the simulator's electromagnetic layer and these traces.

C Electromagnetic Layer: Detailed Implementation

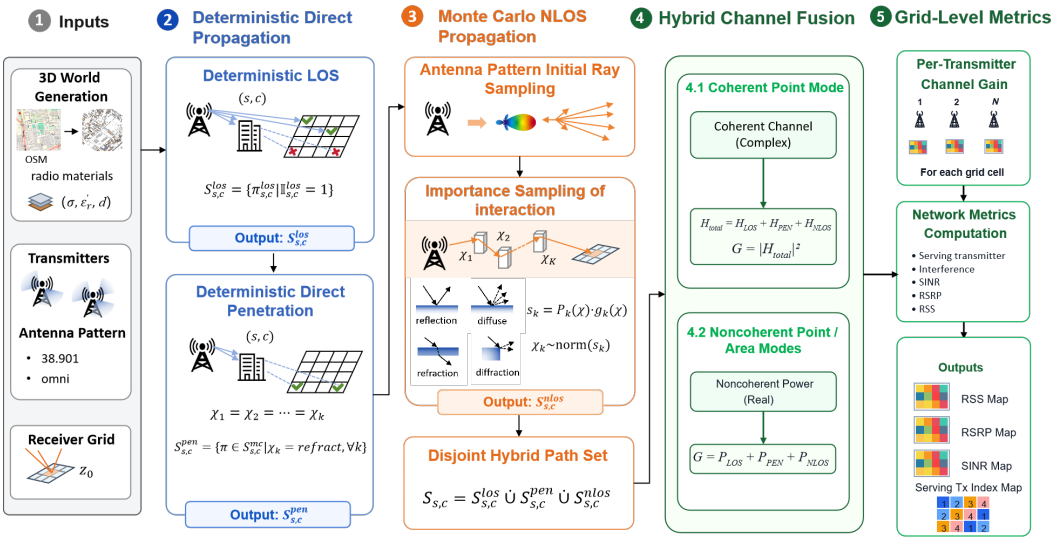

Figure 10: Architecture of electromagnetic layer.

C.1 Overview

This appendix provides the implementation details of the electromagnetic layer. The goal of this layer is to generate propagation paths between each transmitter sector s and receiver grid cell c, evaluate their electromagnetic responses, and accumulate them into grid-level channel gains and radio metrics. The architecture is demonstrated in Figure 10. The detailed implementation is organized as follows:

- **Radio material** (Appendix C.2) defines how scene objects are associated with radio materials, and how frequency-dependent permittivity, conductivity, and complex permittivity are obtained for electromagnetic response calculation.
- **Hybrid ray tracing** (Appendices C.3, C.4, C.5) describes how propagation paths are generated. This includes antenna-aware initial ray sampling, Monte Carlo NLOS tracing with importance sampling, and deterministic LOS and direct penetration.
- **Radio map calculation** (Appendices C.6 and C.7) explains how path-level responses are converted into grid-level channel gains, and how radio maps of RSS, RSRP, and SINR are computed.
- **Visualization and comparison** (Appendix C.9) presents the visualization and comparison results with Sionna RT.

C.2 Radio Material

The electromagnetic response of a propagation ray is determined by the radio material assigned to the surface it intersects. The urban environment is represented as a collection of triangular meshes. A building object is indexed by an object identifier o, and a triangular primitive within the object is indexed by m, so a primitive can be uniquely identified by the pair (o, m). For edge-related interactions, such as diffraction, the three edges of a triangle are further indexed by $t \in \{0, 1, 2\}$,

resulting in an edge token (o, m, t). For a material class u, we characterize its radio response by the real part of the relative permittivity $\epsilon'_r(u, f)$ and the conductivity $\sigma(u, f)$. Their numerical values are dependent on the carrier frequency of the ray. Following Recommendation ITU-R P.2040-3 (ITU-R, 2023), each material class is associated with empirical coefficients (a_u, b_u, c_u, d_u) over a supported frequency range. Given the carrier frequency f_c in Hz and f in GHz, we compute

$$\epsilon'_r(u, f) = a_u f^{b_u}, \tag{23}$$

$$\sigma(u, f) = c_u f^{d_u}. \tag{24}$$

In implementation, when a ray intersects a primitive (o, m), we identify the associated radio material u and select the ITU coefficients (a_u, b_u, c_u, d_u) according to the carrier frequency f. These coefficients are then used to compute ϵ'_r and σ. If the requested frequency is outside the valid frequency range for the material, the nearest supported frequency is used as a fallback.

The lossy dielectric behavior is represented through the complex relative permittivity

$$\epsilon^*_r(u, f_c) = \epsilon'_r(u, f) - j\frac{\sigma(u, f)}{2\pi f_c \epsilon_0}, \tag{25}$$

where ϵ_0 is the vacuum permittivity. The resulting complex permittivity is used to compute the Jones matrix of each electromagnetic interaction.

C.3 Antenna-aware Initial Ray Sampling

To improve the efficiency of NLOS ray tracing, we use the antenna radiation pattern as an importance sampling prior for initial ray directions. The intuition is that directions with larger antenna gain are more likely to generate effective high-contribution feasible paths, and therefore should receive more ray samples. This sampling method changes only the proposal distribution of launched rays, and the resulting bias is corrected by the initial compensation term C_1 in Equation 1.

We sample directions in the local coordinate frame of each transmitter sector. Let $\omega = (e, \phi)$, where $e \in [-\pi/2, \pi/2]$ is the total elevation angle and $\phi \in [-\pi, \pi)$ is the azimuth angle. The polar angle used in the antenna pattern of Appendix B.2 is

$$\theta = \frac{\pi}{2} - e. \tag{26}$$

Given the antenna pattern $A_E(\theta, \phi)$ in dB, we convert it to a linear power gain

$$G_{\text{tx}}(e, \phi) = 10^{A_E(\frac{\pi}{2}-e,\phi)/10}. \tag{27}$$

The local angular domain is discretized into N_e elevation bins and N_ϕ azimuth bins, with bin widths $\Delta e = \frac{\pi}{N_e}$ and $\Delta\phi = \frac{2\pi}{N_\phi}$. Here, the solid angle measures the area occupied by a direction region on the unit sphere. Therefore, for the bin centered at (e_i, ϕ_j), the solid angle is approximated by

$$\Delta\Omega_{ij} = \cos(e_i)\Delta e\Delta\phi. \tag{28}$$

We assign each bin a sampling mass $m_{ij} = G_{\text{tx}}(e_i, \phi_j)\Delta\Omega_{ij}$, and normalize it to obtain the discrete bin probability Q_{ij}. A bin (i, j) is first sampled according to Q_{ij}, and the final direction is uniformly jittered within the selected bin

$$e = e_i + (u_e - \tfrac{1}{2})\Delta e, \qquad \phi = \phi_j + (u_\phi - \tfrac{1}{2})\Delta\phi, \tag{29}$$

where $u_e, u_\phi \sim \mathcal{U}(0, 1)$. Therefore, the local unit direction of each ray is

$$\hat{\mathbf{d}}_{\text{loc}} = [d_x, d_y, d_z]^\top = [\cos e \cos\phi,\ \cos e \sin\phi,\ \sin e]^\top. \tag{30}$$

Since Q_{ij} is a bin probability, the continuous proposal density with respect to solid angle is

$$q_0(\omega) = \frac{Q_{ij}}{\Delta\Omega_{ij}}. \tag{31}$$

Equivalently, this gives a piecewise constant approximation to the weighted angular density,

$$q_0(\omega) \approx \frac{G_{\text{tx}}(\omega)}{\int_{\mathbb{S}^2} G_{\text{tx}}(\omega')\, d\omega'}. \tag{32}$$

In our setting, each base station has one to three sectors with different azimuth orientations and total downtilts. For each sector, rays are first sampled in its local coordinate frame according to the antenna-aware proposal distribution, and are then rotated into the world frame using the sector orientation. Let $\mathbf{f}_s, \mathbf{r}_s, \mathbf{u}_s$ denote the forward, right, and up axes of sector s, determined by its azimuth orientation, mechanical downtilt, and electrical downtilt. The world-frame direction is

$$\hat{\mathbf{d}}_0 = d_x \mathbf{f}_s + d_y \mathbf{r}_s + d_z \mathbf{u}_s. \tag{33}$$

Since this transformation is a rigid rotation, the proposal density $q_0(\omega)$ with respect to solid angle is not changed by the rotation. Using the TR 38.901 antenna pattern, Figure 11 illustrates the resulting initial ray sampling for the three sectors of a centered base station.

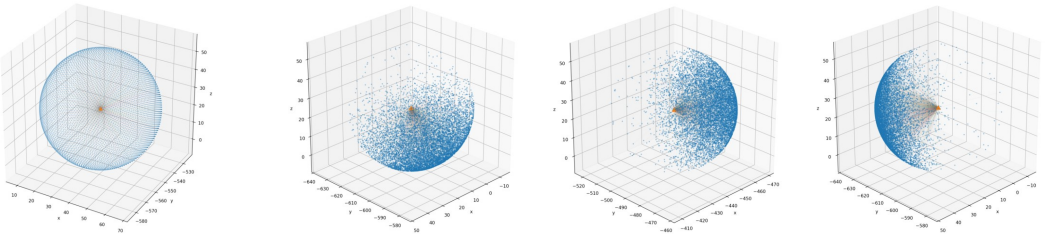

(a) Sionna RT (b) Ours (Sector 1, 2, and 3)

Figure 11: Visualization of initial ray sampling.

C.4 Importance Sampling in Monte Carlo NLOS Tracing

When a ray intersects a building surface, four main electromagnetic interactions in $\mathcal{X}$ may physically coexist, including specular reflection, diffuse scattering, refraction, and diffraction. However, explicitly branching a ray into all possible interaction types at every hit would lead to an exponential growth in the number of path candidates. Therefore, following a stochastic path-tracing strategy similar to Sionna RT (Hoydis et al., 2023), we sample a single interaction type and compensate for this stochastic choice using specifically designed importance weights.

Let $\mathbf{x}_k$ be the current ray intersection point, $\hat{\mathbf{k}}_i$ be the incident propagation direction, and $\mathbf{n}_k$ be the oriented surface normal. For each candidate interaction $\chi \in \mathcal{X}$, we construct an outgoing direction $\hat{\mathbf{k}}_o(\chi)$, a Jones response $\mathbf{J}_k(\chi)$, and a local sampling density $p_k^{\text{local}}(\chi)$. Motivated by the intuition that electromagnetically stronger and geometrically more reachable interactions are more likely to contribute to the receiver grid plane, we assign higher sampling probability to interactions with larger post-interaction field power and shorter projected distance to the grid plane. Each candidate interaction is assigned a nonnegative proposal score:

$$s_k(\chi) = P_{\mathbf{J},k}(\chi)\, g_k(\chi), \qquad P_{\mathbf{J},k}(\chi) = \|\mathbf{J}_k(\chi)\|_F^2, \tag{34}$$

where $P_{\mathbf{J},k}(\chi)$ measures the electromagnetic strength of the candidate interaction using the squared Frobenius norm of its Jones response, and $g_k(\chi)$ measures its geometric reachability to the receiver grid. Specifically, given a candidate outgoing direction $\hat{\mathbf{k}}_o(\chi)$, we compute its forward intersection with the receiver grid plane $z = z_0$:

$$t_k(\chi) = \frac{z_0 - x_{k,z}}{\hat{k}_{o,z}(\chi)}, \qquad \mathbf{u}_k(\chi) = \mathbf{x}_k + t_k(\chi)\hat{\mathbf{k}}_o(\chi), \tag{35}$$

where $t_k(\chi)$ is the distance to the grid plane and $\mathbf{u}_k(\chi)$ is the projected hit point. The geometric reachability score is defined as:

$$g_k(\chi) = \mathbf{1}\left[t_k(\chi) > 0,\ |\hat{k}_{o,z}(\chi)| > \epsilon,\ \mathbf{u}_k(\chi) \in \mathcal{G}\right] \frac{1}{1 + t_k(\chi)}, \tag{36}$$

where $\mathcal{G}$ is the receiver-grid extent and ϵ is a numerical threshold. Thus, only candidates whose outgoing ray intersects the grid plane in the forward direction and inside the grid boundary receive positive score, and shorter projected distances are sampled with higher probability.

The interaction proposal distribution is then obtained by normalizing these scores:

$$\rho_k(\chi) = \frac{s_k(\chi)}{\sum_{\chi' \in \mathcal{X}} s_k(\chi')}. \tag{37}$$

We next specify how the electromagnetic strength term $P_{\mathbf{J},k}(\chi)$ and the local sampling density $p_k^{\text{local}}(\chi)$ in Equation 34 are computed for different interaction types. Here, $p_k^{\text{local}}(\chi)$ denotes the probability density of the local outgoing direction or edge related sample, and is used in the compensation coefficient C_2 in Equation 1.

Specifically, for reflection and refraction, a single-layer slab is used to model the building objects. Each intersected building surface is treated as a homogeneous slab with finite thickness d and material-dependent complex relative permittivity $\epsilon_r^*(u, f_c)$ defined in Equation 25.

Specular reflection. The outgoing direction is deterministic:

$$\hat{\mathbf{k}}_o = \hat{\mathbf{k}}_i - 2(\hat{\mathbf{k}}_i \cdot \mathbf{n}_k)\mathbf{n}_k, \qquad p_k^{\text{local}} = 1. \tag{38}$$

The Jones matrix is computed in the TE/TM basis using the single-layer slab reflection coefficients and then rotated back to the world polarization basis:

$$\mathbf{J}_k^{\text{spec}} = \sqrt{1 - S^2}\, \mathbf{W}_{\text{out}} \begin{bmatrix} r_{\text{TE}} & 0 \\ 0 & r_{\text{TM}} \end{bmatrix} \mathbf{W}_{\text{in}}, \tag{39}$$

where $S \in [0, 1]$ is the material scattering coefficient. The factor $\sqrt{1 - S^2}$ reduces the amplitude of specular reflection when part of the surface response is assigned to diffuse scattering.

Refraction. We use a single-layer slab approximation. The transmitted ray keeps its propagation direction,

$$\hat{\mathbf{k}}_o = \hat{\mathbf{k}}_i, \qquad p_k^{\text{local}} = 1, \tag{40}$$

while the slab transmission coefficients account for attenuation and phase change through the material. The Jones response is

$$\mathbf{J}_k^{\text{refrac}} = \sqrt{1 - S^2}\, \mathbf{W}_{\text{out}} \begin{bmatrix} t_{\text{TE}} & 0 \\ 0 & t_{\text{TM}} \end{bmatrix} \mathbf{W}_{\text{in}}. \tag{41}$$

The single-layer slab coefficients are computed as follows:

$$a = \sqrt{\epsilon_r^* - \sin^2 \theta_i}, \qquad q = \frac{2\pi d}{\lambda} a, \tag{42}$$

where θ_i is the incidence angle, d is the slab thickness. The Fresnel coefficients from free-space to slab are

$$r_{\text{TE}}^{(0)} = \frac{\cos\theta_i - a}{\cos\theta_i + a}, \qquad r_{\text{TM}}^{(0)} = \frac{\epsilon_r^* \cos\theta_i - a}{\epsilon_r^* \cos\theta_i + a}. \tag{43}$$

The resulting reflection and transmission coefficients are

$$r_{\text{TE}} = r_{\text{TE}}^{(0)} \frac{1 - e^{-2jq}}{1 - \left(r_{\text{TE}}^{(0)}\right)^2 e^{-2jq}}, \qquad r_{\text{TM}} = r_{\text{TM}}^{(0)} \frac{1 - e^{-2jq}}{1 - \left(r_{\text{TM}}^{(0)}\right)^2 e^{-2jq}}, \tag{44}$$

$$t_{\text{TE}} = \frac{\left(1 - \left(r_{\text{TE}}^{(0)}\right)^2\right) e^{-jq}}{1 - \left(r_{\text{TE}}^{(0)}\right)^2 e^{-2jq}}, \qquad t_{\text{TM}} = \frac{\left(1 - \left(r_{\text{TM}}^{(0)}\right)^2\right) e^{-jq}}{1 - \left(r_{\text{TM}}^{(0)}\right)^2 e^{-2jq}}. \tag{45}$$

Diffuse scattering. The outgoing direction is sampled from a cosine-weighted hemisphere:

$$p_k^{\text{local}} = p_{\text{diff}}(\hat{\mathbf{k}}_o) = \frac{\cos\theta_s}{\pi}, \qquad \cos\theta_s = \max(\hat{\mathbf{k}}_o \cdot \mathbf{n}_k, 0). \tag{46}$$

The corresponding Lambertian scattering response is

$$\mathbf{J}_k^{\text{diffuse}} = \alpha \begin{bmatrix} \sqrt{1 - K_x} & -\sqrt{K_x} \\ \sqrt{K_x} & \sqrt{1 - K_x} \end{bmatrix}, \qquad \alpha = S\sqrt{\frac{\cos\theta_s}{\pi} \cos\theta_i}, \tag{47}$$

where $\cos\theta_i = \max(-\hat{\mathbf{k}}_i \cdot \mathbf{n}_k, 0)$, S is the scattering coefficient, and K_x controls cross-polarization mixing.

Diffraction. We model edge diffraction using the uniform theory of diffraction (UTD) (Kouyoumjian & Pathak, 1974). Unlike specular reflection or refraction, the full UTD response depends on the selected wedge geometry and the outgoing direction on the Keller cone (Keller, 1962). Therefore, using the full UTD response directly as the proposal score would require evaluating wedge geometry and cone-dependent diffraction coefficients for every candidate hit, which is computationally expensive for candidates that are not ultimately selected. To reduce this overhead, we assign diffraction a proxy proposal score before interaction selection:

$$s_k(\text{diffraction}) = \bar{P}_{\text{non-diff},k}\, g_k^{\text{diff}}, \tag{48}$$

where $\bar{P}_{\text{non-diff},k}$ is the mean Jones-response power of the non-diffraction candidates at the same hit, and g_k^{diff} is the geometric reachability proxy for diffraction. This allows diffraction to participate in the same interaction proposal distribution $\rho_k(\chi)$.

Once diffraction is selected, we check whether the current triangle hit is close to an edge. If so, the corresponding wedge is obtained from the precollected wedge cache. The diffraction point is taken on this wedge, and the outgoing direction is sampled on the Keller cone, which is the cone of possible outgoing diffracted directions around the wedge edge. Its axis is the wedge-edge direction $\hat{\mathbf{e}}$, and its opening angle is determined by the incident ray:

$$\beta_0 = \arccos\left(|\hat{\mathbf{k}}_i \cdot \hat{\mathbf{e}}|\right). \tag{49}$$

Thus, the sampled outgoing direction $\hat{\mathbf{k}}_o$ satisfies

$$|\hat{\mathbf{k}}_o \cdot \hat{\mathbf{e}}| = \cos\beta_0. \tag{50}$$

In the local wedge frame, let $\mathbf{t}_0$ and $\mathbf{n}_0$ denote two orthogonal directions perpendicular to the edge. The variable φ parameterizes the azimuthal position around the edge of the wedge on the Keller cone. Given a sampled φ, the outgoing direction is constructed as

$$\hat{\mathbf{k}}_o = \sin\beta_0 \cos\varphi\, \mathbf{t}_0 + \sin\beta_0 \sin\varphi\, \mathbf{n}_0 + \cos\beta_0\, \hat{\mathbf{e}}. \tag{51}$$

The local sampling density is

$$p_k^{\text{local}} = p_{\text{edge}}\, p_{\text{cone}}(\varphi), \qquad p_{\text{cone}}(\varphi) = \frac{1}{\Delta\varphi}, \tag{52}$$

where $\Delta\varphi$ is the valid angular interval on the Keller cone.

After the diffraction point and outgoing direction are determined, the Jones response is computed using the UTD wedge coefficient. Let s and s' be the distances from the previous point to the diffraction point and from the diffraction point to the next point, respectively. Let α_w be the exterior wedge angle, $n = \alpha_w/\pi$, and let $k_0 = 2\pi/\lambda$ denote the wavenumber. We define

$$\ell = \frac{ss'}{s+s'} \sin^2\beta_0. \tag{53}$$

With $\Delta\varphi = \varphi - \varphi'$, $\Sigma\varphi = \varphi + \varphi'$, and transition function $F(\cdot)$, the four scalar UTD terms are

$$D_1 = \Gamma \cot\left(\frac{\pi + \Delta\varphi}{2n}\right) F\left(k_0 \ell A_1\right), \qquad D_2 = \Gamma \cot\left(\frac{\pi - \Delta\varphi}{2n}\right) F\left(k_0 \ell A_2\right), \tag{54}$$

$$D_3 = \Gamma \cot\left(\frac{\pi + \Sigma\varphi}{2n}\right) F\left(k_0 \ell A_3\right), \qquad D_4 = \Gamma \cot\left(\frac{\pi - \Sigma\varphi}{2n}\right) F\left(k_0 \ell A_4\right), \tag{55}$$

where

$$\Gamma = -\frac{e^{-j\pi/4}}{2n\sqrt{2\pi k_0}\sin\beta_0}, \tag{56}$$

$$A_1 = A^+(\Delta\varphi), \quad A_2 = A^-(\Delta\varphi), \quad A_3 = A^+(\Sigma\varphi), \quad A_4 = A^-(\Sigma\varphi),$$

with

$$A^{\pm}(\psi) = 2\cos^2\left(\frac{\alpha_w N^{\pm}(\psi) - \psi}{2}\right), \qquad N^{\pm}(\psi) = \text{round}\left(\frac{\psi \pm \pi}{2\alpha_w}\right). \tag{57}$$

Here $A^{\pm}(\cdot)$ are the UTD transition parameters following the standard definition. The first two terms form the direct wedge-diffraction component, while D_3 and D_4 are combined with the single-layer slab reflection responses of the two wedge faces:

$$\mathbf{J}_k^{\text{diffraction}} = \mathbf{W}_{\text{out}} \left[-(D_1 + D_2)\mathbf{I} + D_4\mathbf{R}_0 + D_3\mathbf{R}_n\right] \mathbf{W}_{\text{in}}, \tag{58}$$

where $\mathbf{R}_0$ and $\mathbf{R}_n$ are the TE/TM reflection matrices of the two wedge faces, and $\mathbf{W}_{\text{in}}, \mathbf{W}_{\text{out}}$ rotate the diffraction response between the local wedge polarization basis and the world polarization basis.

Given the selected interaction sequence $\chi_{1:K}$, the local interaction compensation is

$$C_2(\pi) = \prod_{k=1}^{K} \frac{1}{\rho_k(\chi_k)\, p_k^{\text{local}}(\chi_k)}, \tag{59}$$

which is the same compensation term introduced in Eq. 1. The probability $\rho_k(\chi_k)$ compensates the stochastic choice among interaction types, while p_k^{local} compensates the local continuous sampling within the selected interaction, such as the diffuse outgoing direction or the diffraction cone parameter.

Finally, because ray contributions are accumulated on a receiver grid plane, the interaction compensation is combined with a grid-projection factor C_{hit}. The specific form of C_{hit} depends on the interaction type of the last segment that reaches the grid. For ordinary ray-direction samples, including specular, refraction, and diffuse, the field compensation is

$$C_{\text{hit}}^{\text{ord}}(\pi) = \sqrt{\frac{L^2}{A_b|\cos\theta_g|} C_1(\omega_0) C_2(\pi)}, \tag{60}$$

where L is the path length to the grid hit, A_b is the receiver bucket area, and θ_g is the angle between the incoming ray and the grid-plane normal. For diffraction, the cone-to-grid projection uses the diffraction specific compensation

$$C_{\text{hit}}^{\text{diff}}(\pi) = \sqrt{\frac{s'^2 \sin\beta_0}{A_b|\cos\theta_g|} C_1(\omega_0) C_2(\pi)}. \tag{61}$$

Therefore, the path-level accumulation coefficient is

$$\alpha_\pi = \mathbb{I}\left[\pi \in \mathcal{S}_{s,c}^{\text{los}} \cup \mathcal{S}_{s,c}^{\text{pen}}\right] + C_{\text{hit}}\, \mathbb{I}\left[\pi \in \mathcal{S}_{s,c}^{\text{nlos}}\right]. \tag{62}$$

C.5 Deterministic Direct Propogation

Although the MC-NLOS branch can estimate complex multipath propagation, it is constrained by finite ray budgets and stochastic interaction sampling. Distant grid cells have small angular support and are difficult to hit by random rays, while complex building layouts make some valid propagation branches less likely to be sampled. These uncovered cells are finite sampling artifacts rather than necessarily physical zero power regions, and may lead to discontinuous or incomplete radio maps. To address this issue, we introduce a deterministic direct propagation branch, including target-cell LOS and direct penetration components, to explicitly complete dominant direct contributions without relying on random ray hits.

$$\mathbb{I}_{s,c}^{\text{los}} = \mathbf{1}\left[\text{RayIntersect}\left(\mathbf{x}_s + \epsilon\mathbf{d}_{s,c}, \mathbf{d}_{s,c}\right) \text{ has no hit before } \mathbf{x}_c\right]. \tag{63}$$

If $\mathbb{I}_{s,c}^{\text{los}} = 1$, a deterministic LOS path is added to the path set $\mathcal{S}_{s,c}^{\text{los}}$. If $\mathbb{I}_{s,c}^{\text{los}} = 0$, the direct segment between (s, c) is blocked by one or more building surfaces. The direct contribution from the source to this cell is then modeled as a penetration component through the intersected buildings. According to the four electromagnetic interactions described in Appendix C.4, refraction is the only one that transmits the ray through a building surface along the direct segment. Therefore, we introduce a deterministic direct penetration branch, where each surface hit on the blocked direct segment is treated as a refraction event.

This construction relies on the thin-slab transmission approximation used for each penetrated wall. Under this approximation, the geometric exit point is treated as the incident surface interaction point, and the ray keeps its propagation direction across each penetrated wall. Therefore,

$$\hat{\mathbf{k}}_{o,k} = \hat{\mathbf{k}}_{i,k} = \mathbf{d}_{s,c}. \tag{64}$$

Algorithm 3 Hybrid deterministic direct propagation and MC-NLOS tracing.

Require: Scene geometry $\mathcal{M}$, transmitter sectors $\mathcal{B}$, receiver grid cells $\mathcal{C}$, maximum tracing depth $K_{\max}$, number of initial MC rays N
Ensure: Path sets $\{\mathcal{S}_{s,c}\}_{s\in\mathcal{B},c\in\mathcal{C}}$

for all $s \in \mathcal{B}$, $c \in \mathcal{C}$ **do**
  $\mathcal{S}^{\mathrm{los}}_{s,c} \leftarrow \emptyset$, $\mathcal{S}^{\mathrm{pen}}_{s,c} \leftarrow \emptyset$, $\mathcal{S}^{\mathrm{nlos}}_{s,c} \leftarrow \emptyset$
  $\mathbf{d}_{s,c} \leftarrow (\mathbf{x}_c - \mathbf{x}_s)/\|\mathbf{x}_c - \mathbf{x}_s\|$
  $\mathcal{H}_{s,c} \leftarrow \textsc{RayIntersect}(\mathbf{x}_s, d_{s,c}; \mathcal{M})$
  if $\mathcal{H}_{s,c} = \emptyset$ **then**
    $\pi^{\mathrm{los}}_{s,c} \leftarrow (\mathbf{x}_s, \mathbf{x}_c)$
    $\mathcal{S}^{\mathrm{los}}_{s,c} \leftarrow \mathcal{S}^{\mathrm{los}}_{s,c} \cup \{\pi^{\mathrm{los}}_{s,c}\}$
  else
    $\pi^{\mathrm{pen}}_{s,c} \leftarrow (\mathbf{x}_s, \mathbf{x}_1, \ldots, \mathbf{x}_{K_{s,c}}, \mathbf{x}_c;\ \chi_k = \mathrm{refrac},\ \forall k = 1, \ldots, K_{s,c})$
    $\mathcal{S}^{\mathrm{pen}}_{s,c} \leftarrow \mathcal{S}^{\mathrm{pen}}_{s,c} \cup \{\pi^{\mathrm{pen}}_{s,c}\}$
  end if
end for
for all $s \in \mathcal{B}$ **do**
  for $n = 1$ to N **do**
    Sample $\omega_0 \sim q_0(\omega)$ and launch ray $(\mathbf{x}_s, \hat{\mathbf{d}}_0)$
    Initialize path $\pi \leftarrow (\mathbf{x}_s)$ and ray origin $\mathbf{o}_0 \leftarrow \mathbf{x}_s$
    for $k = 1$ to $K_{\max}$ **do**
      $(\mathbf{x}_k, h_k) \leftarrow \textsc{RayIntersect}(\mathbf{o}_{k-1}, \hat{\mathbf{d}}_{k-1}; \mathcal{M})$
      $c \leftarrow \textsc{GridHit}(\mathbf{o}_{k-1}, \hat{\mathbf{d}}_{k-1}, \mathbf{x}_k)$
      if $c \neq \emptyset$ **then**
        if $k > 1$ **then**
          $\mathcal{S}^{\mathrm{mc}}_{s,c} \leftarrow \mathcal{S}^{\mathrm{mc}}_{s,c} \cup \{\pi\}$
        end if
        break
      else if $h_k = 0$ **then**
        break
      end if
      for all $\chi \in \mathcal{X}$ **do**
        Construct $\hat{\mathbf{k}}_o(\chi)$, $\mathbf{J}_k(\chi)$, and $p^{\mathrm{local}}_k(\chi)$
        Compute score $s_k(\chi)$
      end for
      $\rho_k(\cdot) \leftarrow \textsc{Normalize}\big(\{s_k(\chi)\}_{\chi\in\mathcal{X}}\big)$
      Sample interaction type $\chi_k \sim \rho_k(\cdot)$
      Append $(\mathbf{x}_k, \chi_k)$ to π
      $\hat{\mathbf{d}}_k \leftarrow \hat{\mathbf{k}}_o(\chi_k)$, $\mathbf{o}_k \leftarrow \mathbf{x}_k + \epsilon\hat{\mathbf{d}}_k$
    end for
  end for
end for
for all $s \in \mathcal{B}$, $c \in \mathcal{C}$ **do**
  $\mathcal{S}^{\mathrm{nlos}}_{s,c} \leftarrow \mathcal{S}^{\mathrm{mc}}_{s,c} \setminus \{\pi \in \mathcal{S}^{\mathrm{mc}}_{s,c} : \chi_k(\pi) = \mathrm{refrac}, \forall k\}$
  $\mathcal{S}_{s,c} \leftarrow \mathcal{S}^{\mathrm{los}}_{s,c} \dot{\cup} \mathcal{S}^{\mathrm{pen}}_{s,c} \dot{\cup} \mathcal{S}^{\mathrm{nlos}}_{s,c}$
end for
return $\{\mathcal{S}_{s,c}\}$

The attenuation and phase shift caused by the finite material thickness are incorporated into the refraction Jones matrix $\mathbf{J}^{\mathrm{refrac}}_k$ in Eq. 41.

Specifically, let $\mathcal{H}_{s,c} = \{\mathbf{x}_1, \ldots, \mathbf{x}_{K_{s,c}}\}$ denote the ordered surface intersections along the direct segment before reaching $\mathbf{x}_c$. The blocked direct segment is then modeled as a pure-refraction path,

$$\pi^{\mathrm{pen}}_{s,c} = \big(\mathbf{x}_s, \mathbf{x}_1, \ldots, \mathbf{x}_{K_{s,c}}, \mathbf{x}_c;\ \chi_k = \mathrm{refrac},\ \forall k = 1, \ldots, K_{s,c}\big)\,. \tag{65}$$

This path is added to $\mathcal{S}_{s,c}^{\mathrm{pen}}$, and its electromagnetic response is obtained by accumulating the slab transmission responses along the direct segment,

$$\mathbf{J}_{s,c}^{\mathrm{pen}} = \prod_{k=1}^{K_{s,c}} \mathbf{J}_k^{\mathrm{refrac}}. \tag{66}$$

The purpose of the deterministic direct penetration branch is to replace a sampling-induced missing contribution with a deterministic and physically attenuated penetration contribution. If the direct segment crosses multiple building surfaces or highly lossy materials, the accumulated transmission loss still results in weak received power. However, this weak value is then caused by material penetration attenuation, rather than by the failure of random rays to reach the cell.

To avoid double counting, the pure-refraction paths handled by deterministic direct penetration must be removed from the MC-NLOS path set. We first collect the complete MC hit set as $\mathcal{S}_{s,c}^{\mathrm{mc}}$, and then remove paths for which all interactions are refraction:

$$\mathcal{S}_{s,c}^{\mathrm{nlos}} = \mathcal{S}_{s,c}^{\mathrm{mc}} \setminus \left\{ \pi \in \mathcal{S}_{s,c}^{\mathrm{mc}} \mid \chi_k(\pi) = \mathrm{refrac},\ \forall k \right\}. \tag{67}$$

The final path set is therefore composed of deterministic LOS, deterministic direct penetration, and the residual MC-NLOS paths:

$$\mathcal{S}_{s,c} = \mathcal{S}_{s,c}^{\mathrm{los}} \cup \mathcal{S}_{s,c}^{\mathrm{pen}} \cup \mathcal{S}_{s,c}^{\mathrm{nlos}}. \tag{68}$$

The detailed procedure is summarized in Algorithm 3. In implementation, transmitter-cell pairs are processed in chunks with tensor operations and accelerated by multi-GPU execution. Ray intersections are evaluated in parallel rather than loops. This keeps the deterministic LOS and deterministic direct penetration branches scalable for large receiver grids.

C.6 Ray-Based Channel Construction

Transmit Field Vector and Path Polarization Transfer. The transmit field vector $\mathbf{f}_{\mathrm{tx}} \in \mathbb{C}^2$ represents the radiated field amplitude and polarization state in the departure direction (θ, ϕ):

$$\mathbf{f}_{\mathrm{tx}} = \sqrt{G_{\mathrm{elem}}(\theta,\phi)}\, G_{\mathrm{array}} \begin{bmatrix} \cos\psi \\ \sin\psi \end{bmatrix}, \tag{69}$$

where $G_{\mathrm{elem}}(\theta,\phi)$ and G_{array} denote the element gain and array gain, respectively, and ψ is the polarization slant angle, with $\psi = 0$ and $\psi = \pi/2$ corresponding to pure $\hat{\theta}$- and $\hat{\phi}$-polarization.

Each propagation interaction at bounce k is modeled by a 2×2 complex Jones matrix $\mathbf{J}_k$. For specular reflection and transmission, it is constructed from the Fresnel TE/TM coefficients with polarization basis rotation:

$$\mathbf{J}_k = \mathbf{W}_{\mathrm{out},k} \begin{bmatrix} r_{\mathrm{TE},k} & 0 \\ 0 & r_{\mathrm{TM},k} \end{bmatrix} \mathbf{W}_{\mathrm{in},k}, \tag{70}$$

where $\mathbf{W}_{\mathrm{in},k}$ and $\mathbf{W}_{\mathrm{out},k}$ are the basis rotation matrices between the global and local TE/TM frames. Diffuse scattering is modeled analogously using the scattering and cross-polarization coefficients. The accumulated path polarization transfer over K interactions is then

$$\mathbf{J}_{\mathrm{accum}} = \mathbf{J}_K \cdots \mathbf{J}_2 \mathbf{J}_1. \tag{71}$$

Polarization Projection and Path Coefficient. The electric field vector arriving at the receiver before antenna projection is given by

$$\mathbf{e}_{\mathrm{out}} = \mathbf{J}_{\mathrm{accum}} \mathbf{f}_{\mathrm{tx}}, \tag{72}$$

where $\mathbf{J}_{\mathrm{accum}}$ denotes the accumulated Jones matrix along the propagation path and $\mathbf{f}_{\mathrm{tx}}$ is the transmit antenna field vector. The polarization-matched channel amplitude is obtained by projecting $\mathbf{e}_{\mathrm{out}}$ onto the receive antenna field vector $\mathbf{f}_{\mathrm{rx}} \in \mathbb{C}^2$:

$$a_{\mathrm{pol}} = \mathbf{f}_{\mathrm{rx}}^H \mathbf{e}_{\mathrm{out}} = \mathbf{f}_{\mathrm{rx}}^H \mathbf{J}_{\mathrm{accum}} \mathbf{f}_{\mathrm{tx}}. \tag{73}$$

Two receive modes are considered: a fixed-θ mode with $\mathbf{f}_{\mathrm{rx}} = [1, 0]^T$, which extracts the $\hat{\theta}$-polarized component, and an isotropic-power mode with $a_{\mathrm{pol}} = \|\mathbf{e}_{\mathrm{out}}\|_2$, which captures the total received power irrespective of polarization. Incorporating free-space spreading loss, the physical path coefficient is expressed as

$$a = a_{\mathrm{pol}} \frac{\lambda}{4\pi r}, \tag{74}$$

where $\lambda = c/f_c$ is the carrier wavelength and $r = \sum_{k=0}^{K} \|\mathbf{p}_{k+1} - \mathbf{p}_k\|$ denotes the total length of the propagation path.

Monte Carlo Weighted Ray Contribution. The ray tracer employs a two-level importance sampling scheme. The bounce-level proposal density is accumulated as

$$q_{\mathrm{bounce}} = \prod_{k=1}^{K} p_{\mathrm{comp},k} p_{\mathrm{dir},k}, \tag{75}$$

where $p_{\mathrm{comp},k} = s_k / \sum_j s_j$ is the interaction-type selection probability at bounce k and $p_{\mathrm{dir},k}$ is the corresponding conditional directional pdf. The initial directional sampling term is absorbed into the ray-tube solid angle $\Omega = 1/(N_{\mathrm{ray}} q_{0,\mathrm{tx}})$, where N_{ray} is the total number of emitted rays. At propagation distance r, the receiver bin plane has a footprint area $A_{\mathrm{tube}} = r^2\Omega/|\cos\theta_i|$, where θ_i is the incident angle with respect to the receiver-plane normal. The resulting Monte Carlo weighted contribution of a single ray to the complex channel response is

$$v = a \cdot \alpha_\pi \cdot e^{-j2\pi f \tau}, \tag{76}$$

where α_π is the path accumulation coefficient defined in equation 62, and $\tau = r/c$ is the propagation delay.

Grid Accumulation and Channel Gain. The received ray contributions are accumulated at each grid bucket either coherently or non-coherently. In the coherent mode, the complex path contributions are summed to form the channel transfer function

$$H(f) = \sum_{n=1}^{N} v_n, \tag{77}$$

and the effective channel gain is evaluated as $\Gamma = |H(f)|^2$, preserving inter-path phase interference and capturing small-scale fading effects. In the non-coherent mode, the path powers are accumulated directly as

$$P = \sum_{n=1}^{N} |v_n|^2, \tag{78}$$

with the effective channel gain taken as $\Gamma = P$ for point receivers and $\Gamma = \langle P \rangle_A$ for area receivers, where $\langle \cdot \rangle_A$ denotes spatial averaging over sub-bins. Compared to coherent accumulation, the non-coherent mode produces a spatially smoother power estimate that is more suitable for large-scale coverage prediction.

C.7 Channel Quality Metrics.

This section describes the channel quality metrics computed by MobiSim. All metrics are derived from the channel gain Γ_m obtained from the ray-tracing accumulation procedure described in Eq. 78.

Received Signal Strength. The received signal strength (RSS) quantifies the total power delivered to a receiver from a given transmitter after accounting for the physical propagation channel. For transmitter m, the RSS is defined as

$$\mathrm{RSS}_m = \Gamma_m P_{\mathrm{tx},m}, \tag{79}$$

where $P_{\mathrm{tx},m}$ denotes the transmit power of transmitter m and Γ_m is the channel gain obtained from the ray-tracing accumulation procedure described in Eq. 78. Since multiple base stations may be visible from a given receiver location in a multi-cell deployment, the aggregate received signal strength over all M transmitters is

$$\mathrm{RSS} = \sum_{m=1}^{M} \mathrm{RSS}_m. \tag{80}$$

Table 9: Channel quality metrics

Symbol	Description
Γ_m	Channel gain from transmitter m
RSS_m	Received signal strength from transmitter m
RSS	Aggregate received signal strength over all transmitters
m^*	Serving transmitter index
RSRP	Reference signal received power of the serving cell
I	Aggregate inter-cell interference power
N_0	Thermal noise power
SINR	Signal-to-interference-plus-noise ratio

Serving Cell Selection. In a multi-cell network, each user is associated with the transmitter that provides the strongest received signal. The serving transmitter index is determined by

$$m^* = \arg \max_{m \in \{1,\ldots,M\}} \mathrm{RSS}_m. \tag{81}$$

This max-RSS association criterion is consistent with the reference signal-based cell selection procedure specified in 3GPP TS 38.304, where the UE selects the cell with the highest measured RSRP.

Reference Signal Received Power. The reference signal received power (RSRP) is a standardized metric defined in 3GPP TS 38.215 to characterize the downlink signal quality of the serving cell. In MobiSim, RSRP is computed as the received signal strength from the serving transmitter:

$$\mathrm{RSRP} = \mathrm{RSS}_{\mathrm{serving}} = \mathrm{RSS}_{m^*}. \tag{82}$$

RSRP reflects the signal strength experienced by individual users across the coverage area and serves as the primary metric for handover decisions, cell reselection, and coverage assessment in 5G NR.

Inter-Cell Interference. In a multi-cell deployment, transmissions from non-serving base stations contribute to inter-cell interference at the receiver. The aggregate interference power is computed as

$$I = \sum_{m \neq m^*} \mathrm{RSS}_m, \tag{83}$$

where the summation excludes the serving transmitter m^*. In practice, the received power from distant transmitters may fall below the noise floor; accordingly, negative values arising from numerical computation are clipped to zero for stability.

Signal-to-Interference-plus-Noise Ratio. The signal-to-interference-plus-noise ratio (SINR) is the primary indicator of link quality and determines the achievable data rate through modulation and coding scheme (MCS) selection. At each receiver location, SINR is defined as

$$\mathrm{SINR} = \frac{\mathrm{RSS}_{m^*}}{I + N_0} = \frac{\Gamma_{m^*} P_{\mathrm{tx},m^*}}{\sum_{m \neq m^*} \Gamma_m P_{\mathrm{tx},m} + k_B T_0 B}, \tag{84}$$

where $N_0 = k_B T_0 B$ is the thermal noise power computed following the Johnson-Nyquist model, with k_B denoting the Boltzmann constant, $T_0 = 290\,\mathrm{K}$ the reference noise temperature, and B the system bandwidth. A higher SINR indicates a cleaner received signal with less interference contamination, enabling higher-order modulation and yielding greater spectral efficiency.

C.8 Electromagnetic Fidelity Validation with Sionna RT

In Section 4.1, we conduct path loss fitting experiments under a single-wall scenario using the released Sionna RT RadioMapSolver as the reference. Both simulators use identical configurations, including a $5\,\mathrm{km} \times 5\,\mathrm{km}$ region, a 500×500 receiver grid, one isotropic transmitter, carrier frequencies of 1, 1.2, and 1.5 GHz, transmitter heights ranging from 30 to 80 m, and 10^6 rays per source. In total, path loss spatial data from 33 configurations are collected for fitting. During preprocessing, we retain

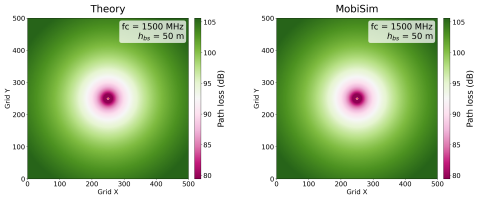

(a) Theoretical (b) MobiSim (c) Sionna

Model	α	β	γ
FSPL	32.44	20.00	20.00
Ours	32.44	20.00	20.00
Sionna	30.66	17.78	20.00

(d) Fitted coefficients

Figure 12: Path loss heatmaps and fitted FSPL coefficients under the free-space LoS benchmark. MobiSim exhibits a propagation pattern and fitted coefficients consistent with the free-space attenuation law, while Sionna shows noticeable deviation in distance-dependent attenuation behavior.

grid cells with valid path loss and positive path gain. Due to finite ray sampling, Sionna RT may not identify valid propagation paths for some distant grid cells. Therefore, each simulator is evaluated over its own valid grid cells.

In addition to the single-wall experiment, we further evaluate MobiSim under an ideal free-space LoS benchmark. The fitting procedure remains unchanged. We fit Equation 4 using weighted least squares, where the cell weights are proportional to d_{3D}^{-2}, with d_{3D} denoting the 3D distance between the receiver cell and the transmitter. The weights are normalized so that each configuration contributes equally.

As shown in Table 12d and Fig. 12, MobiSim closely matches the theoretical FSPL coefficients and reproduces the expected distance-dependent attenuation pattern, whereas Sionna RT exhibits deviations in the distance term. This agreement is attributed to MobiSim's deterministic LoS branch, which follows the analytical free-space propagation model and recovers the theoretical coefficients under ideal LoS conditions.

C.9 Radio Map Comparison and visualization with Sionna RT

We compare the radio maps generated by MobiSim with those produced by Sionna RT in a $2\,\mathrm{km} \times 2\,\mathrm{km}$ urban region, which is discretized into 500×500 grid cells. The scenario contains 10 base stations, each deployed at a height of $25\,\mathrm{m}$ and equipped with three sectors. All sectors use the same transmit power of $46\,\mathrm{dBm}$. The receiver grid plane has a height of $1.5\,\mathrm{m}$. For a fair comparison, both MobiSim and Sionna RT use 10^6 sampled rays per source. Figure 13 shows the spatial distributions of RSS, RSRP, and SINR obtained from Sionna RT and MobiSim, together with their point-wise differences.

In the central region, the radio maps generated by MobiSim are visually consistent with those of Sionna RT across all three metrics, indicating that MobiSim preserves the main propagation characteristics captured by Sionna RT in well-sampled areas. However, noticeable differences appear in some peripheral, distant, and locally shadowed regions. In these regions, the angular support of a grid cell with respect to the transmitter can be very small, and stochastic ray tracing with a finite ray budget may miss some valid propagation paths. This can lead to discontinuous low-power holes or uncovered grid cells in the Sionna RT radio maps. In contrast, MobiSim produces more spatially complete radio maps while maintaining consistency with Sionna RT in the central region. As a result, the deviations in peripheral and regions with sparse ray coverage can exceed $20\,\mathrm{dB}$ for RSS, RSRP, and SINR.

To further quantify the consistency between the two simulators, Figure 14 shows the CDFs and PDFs of RSS, RSRP, and SINR within the central region. The distributions obtained by MobiSim closely match those of Sionna RT. The JSD between the corresponding PDFs is 0.0193 for RSS, 0.0188 for RSRP, and 0.0216 for SINR, indicating good distributional agreement in the central evaluation region.

C.10 Radio Map Comparison under Different Sampling Budgets

To compare radio maps generated by our method and Sionna RT under different sampling budgets, we select a real-world urban region. Specifically, a 200×200 m^2 area is extracted from the scene with a

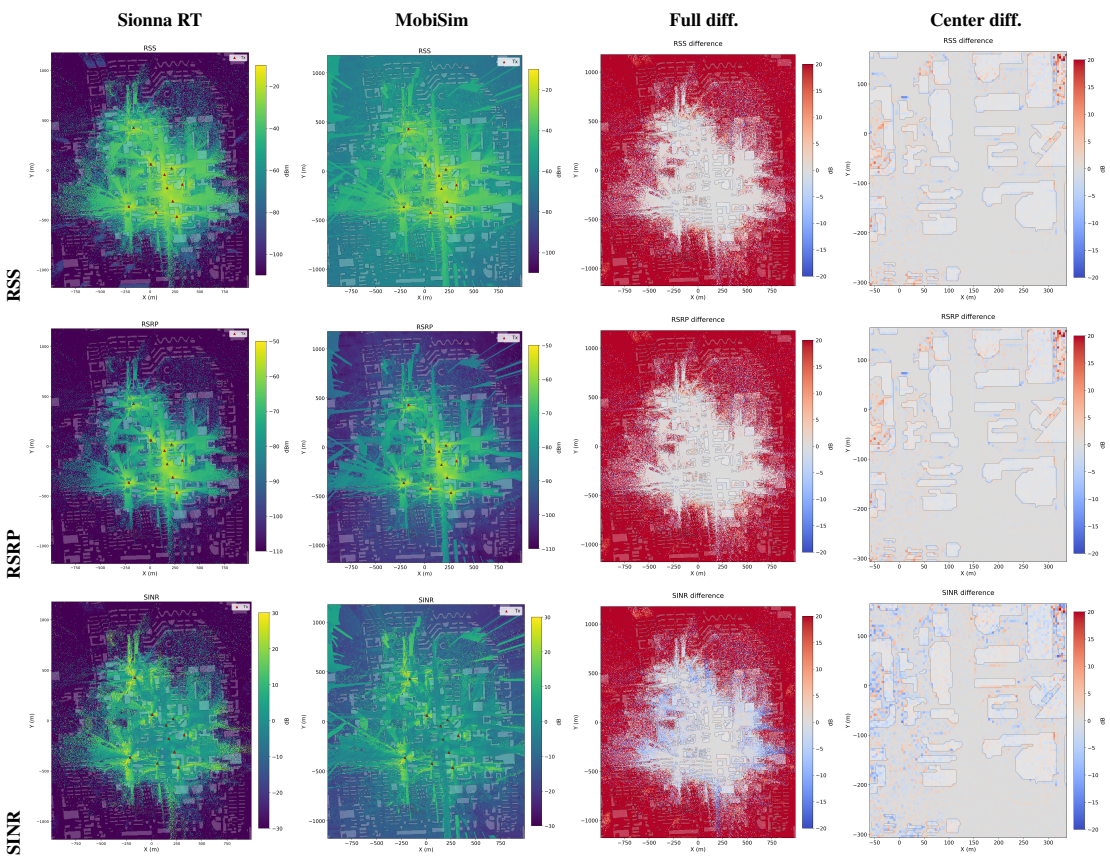

Figure 13: Comparison of radio maps between Sionna RT and MobiSim. Rows correspond to RSS, RSRP, and SINR. Columns show the results of Sionna RT, MobiSim, the full-map difference (MobiSim − Sionna RT, $2\,\mathrm{km} \times 2\,\mathrm{km}$ region), and the difference in the central area.

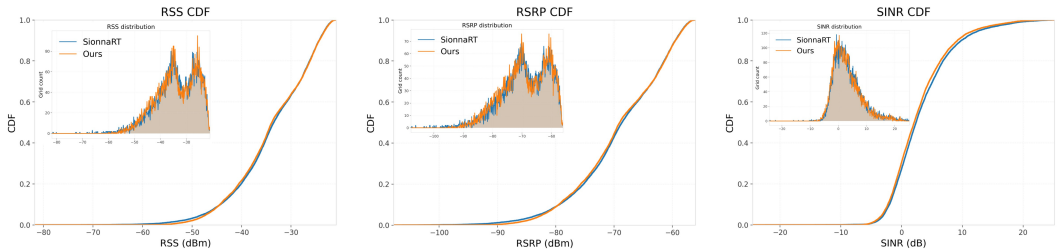

(a) RSS (b) RSRP (c) SINR

Figure 14: CDF and PDF of the metrics in the central region.

grid resolution of 1 m. The transmitter and receiver heights are set to $h_{\mathrm{tx}} = 25$ m and $h_{\mathrm{rx}} = 1.5$ m, respectively.

Four representative receiver regions are selected around the base station, including the near, front, side, and back regions, which cover diverse propagation characteristics under both LOS and NLOS conditions. We vary the number of ray samples in both Sionna RT and MobiSim, and evaluate the resulting radio maps using three metrics: (1) the spatial distribution consistency measured by the Jensen-Shannon divergence (JSD) between RSRP maps, (2) the average RSRP values, and (3) the 95th-percentile RSRP values. These metrics characterize the convergence behavior of Sionna RT and the agreement between the two simulation approaches under different sampling budgets.

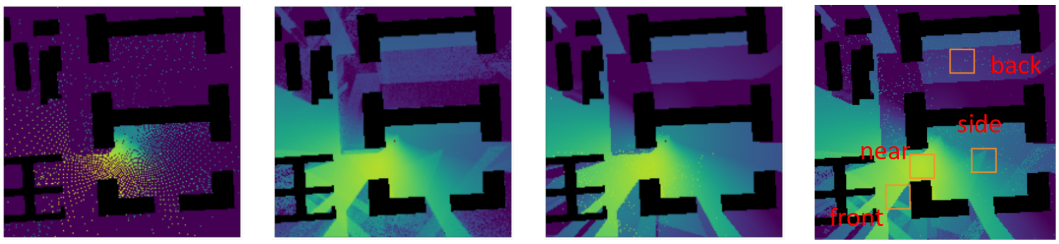

(a) Sionna, $N = 10^4$ (b) Sionna, $N = 4 \times 10^6$ (c) MobiSim, $N = 10^4$ (d) MobiSim, $N = 4 \times 10^6$

Figure 15: RSRP radio maps in the real-world scene under different numbers of sampled rays.

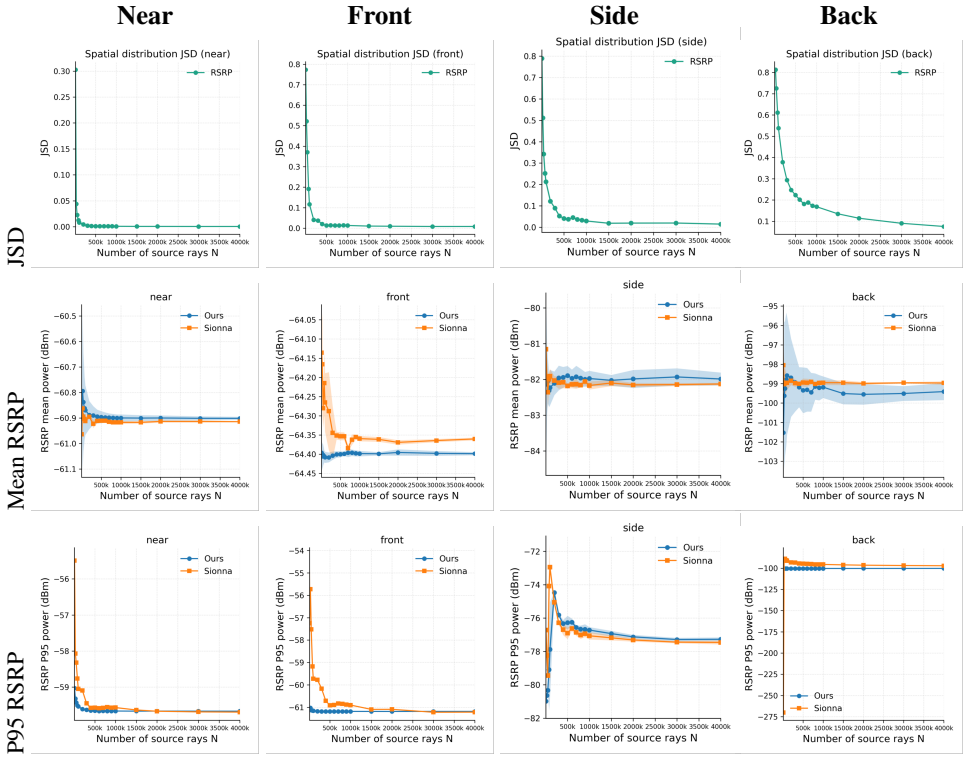

Figure 16: Variation of radio map statistics with the number of sampled rays in the real-world scene. Columns correspond to the near, front, side, and back regions. Rows show JSD, mean RSRP, and the 95th-percentile RSRP, respectively.

The four regions are illustrated in Figure 15. As shown in Figure 16, compared with Sionna RT, our method achieves comparable radio map generation results with a much smaller ray sampling budget, demonstrating higher sampling efficiency.

C.11 Real-World Validation Test

We validate the electromagnetic layer of MobiSim on four public real-world measurement datasets: DeepSense 6G Scenario 6 and AERPAW Datasets 18, 23, and 24. Table 10 summarizes their locations, evaluated frequency bands, and numbers of base stations. Together, these datasets cover vehicle-to-infrastructure measurements in the Phoenix metropolitan area and air-to-ground measurements in Raleigh, spanning both the 60 GHz millimeter-wave band and the 2.1575 GHz LTE band.

DeepSense 6G Scenario 6 comprises a stationary receiver equipped with a phased array and a vehicle-mounted transmitter operating at 60 GHz. The three AERPAW datasets contain measure-

Table 10: Real-world datasets used for electromagnetic-layer validation. Frequency and BSs refer to the carrier and number of base-station sites evaluated in this experiment.

Dataset	City	Frequency	BSs
DeepSense 6G S6 (Alkhateeb et al., 2023)	Phoenix, USA	mmWave, 60 GHz	1
AERPAW 18 (Asokan et al., 2025)	Raleigh, USA	LTE, 2.1575 GHz	1
AERPAW 23 (Singh, 2025a)	Raleigh, USA	LTE, 2.1575 GHz	1
AERPAW 24 (Singh, 2025b)	Raleigh, USA	LTE, 2.1575 GHz	1

ments collected by UAV-mounted devices communicating with a fixed private Ericsson base station. Dataset 18 uses a Quectel modem along a zigzag flight path; Dataset 23 uses an Android smartphone along a rectangular sweep trajectory; and Dataset 24 records horizontal UAV sweeps at increasing distances from the base station using Keysight Nemo and PawPrints. Although the AERPAW releases include both LTE and 5G NR measurements, we evaluate their LTE measurements at the frequency listed in Table 10.

For each dataset, we reconstruct the measurement environment in MobiSim, place the fixed terminal at its recorded location, and configure the carrier frequency to match the evaluated measurements. We then apply the hybrid ray-tracing pipeline described in Appendices C.3–C.5 at the recorded mobile-terminal positions to obtain the corresponding received-power sequence.

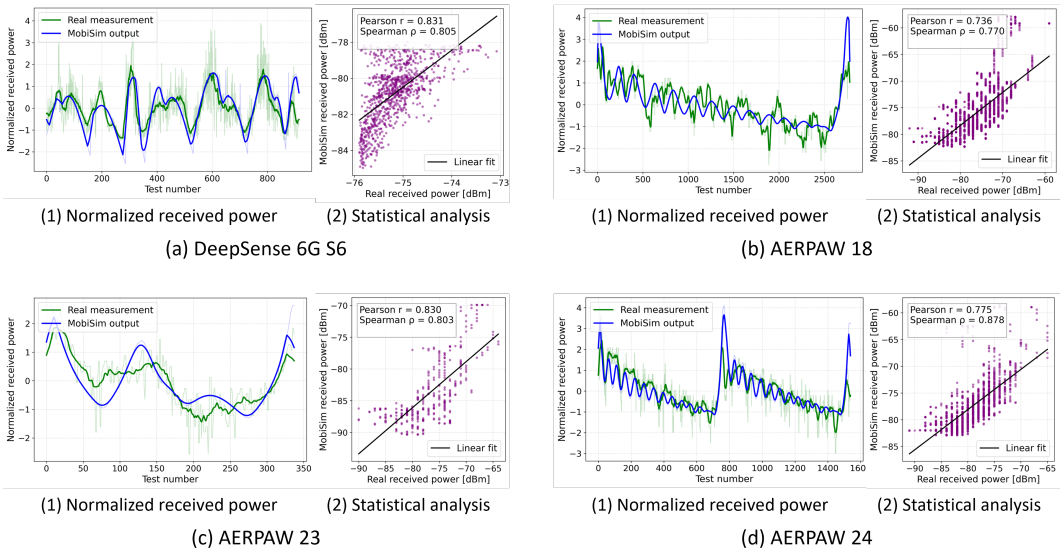

Figure 17: Real-world validation of the MobiSim electromagnetic layer on DeepSense 6G S6 (top left), AERPAW 18 (top right), AERPAW 23 (bottom left), and AERPAW 24 (bottom right). Each dataset is represented by normalized received-power curves comparing real measurements (green) with MobiSim outputs (blue), and a measured-versus-simulated received-power scatter plot with a linear fit. Pearson and Spearman correlation coefficients are annotated in the corresponding panels.

Figure 17 compares the measured and simulated received-power sequences. The smoothed curves exhibit consistent rises and falls along the measurement trajectories. The Pearson correlation coefficients range from 0.736 to 0.831, while the Spearman rank correlation coefficients range from 0.770 to 0.878 across the four datasets. In DeepSense 6G S6, MobiSim captures the repeated peaks and troughs along the vehicle trajectory. In the AERPAW datasets, it reproduces the broad attenuation and recovery patterns as the UAV moves relative to the fixed base station, including the pronounced mid-sequence power increase in Dataset 24. The corresponding scatter plots further show a positive relationship between measured and simulated received power. These results demonstrate consistent propagation modeling across two geographic areas, two frequency bands, and both ground-vehicle and aerial trajectories, supporting the use of MobiSim's electromagnetic layer in the downstream network-level experiments.

C.12 Ablation on Material Properties

To evaluate the impact of material properties on electromagnetic propagation, we conduct additional experiments on DeepSense 6G Scenarios 31 and 32 (Alkhateeb et al., 2023). We keep the scene geometry and other simulation settings unchanged, and replace the surface material with concrete, glass, wood, metal, and plasterboard. As shown in Table 11, different material settings lead to noticeable variations in the simulation results, confirming that material properties affect electromagnetic propagation and the resulting received power.

Table 11: Ablation results with different material settings on DeepSense 6G Scenarios.

Scenario	Material	Pearson r	Spearman ρ
DeepSense 6G 31	Concrete	**0.649**	**0.818**
	Glass	0.588	0.747
	Wood	0.630	0.752
	Metal	0.615	0.767
	Plasterboard	0.645	0.784
DeepSense 6G 32	Concrete	**0.561**	**0.775**
	Glass	0.509	0.707
	Wood	0.553	0.737
	Metal	0.519	0.700
	Plasterboard	0.560	0.751

D User Layer: Detailed Implementation

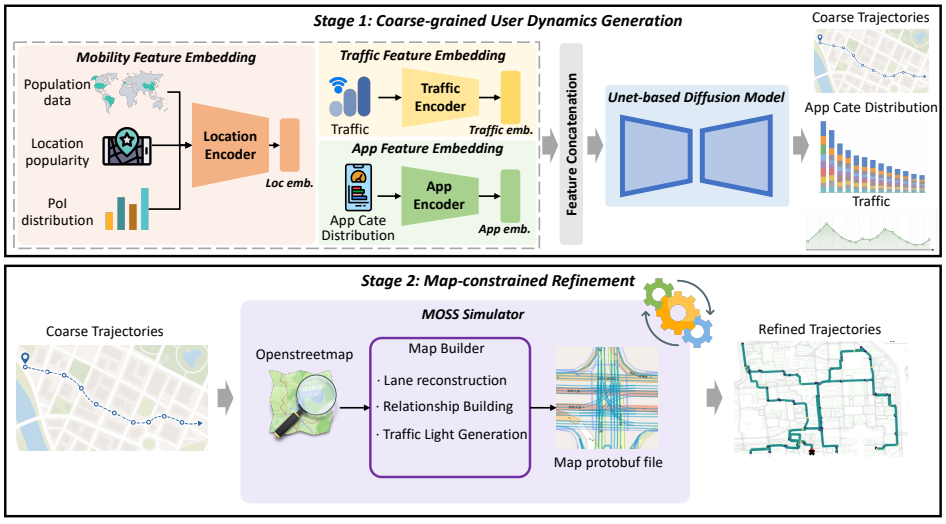

Figure 18: Workflow of user layer for generating user dynamics.

D.1 Overview

This appendix provides comprehensive details on the user layer, with the overall workflow illustrated in Figure 18. We first introduce the architecture and implementation of the coarse-grained user dynamics generation model in D.2, followed by a description of the evaluation methodology in D.3. Finally, we present supplementary visualization results in D.4 to further demonstrate the realism and effectiveness of the generated user dynamics.

Table 12: Dataset Statistics.

Cities	Collection Date	# Users	# Apps	# Records	# App Categories
Shanghai	19 April 2016 - 25 April 2016	1,000	2,000	4,171,950	20
Nanchang	19 May 2022 - 25 May 2022	2,000	1,160	19,929,538	20

D.2 Model Architecture Details

This section describes the input representations and encoding schemes used in the coarse-grained user dynamics generation model. For location features, each location is represented by a 36-dimensional feature vector $\mathbf{l_i} \in \mathbb{R}^{36}$, including 1 dimension for population density, 34 dimensions for POI distributions across POI categories, and 1 dimension for location popularity derived from global high-resolution origin-destination commuting flow data. These heterogeneous features are encoded into an 8-dimensional latent vector via the autoencoder to obtain compact semantic representations of urban regions:

$$\mathbf{z_i} = \text{Encoder}(\mathbf{l_i}), \mathbf{l_i} = Decoder(\mathbf{z_i}). \tag{85}$$

The autoencoder is pretrained to reconstruct multi-source location attributes. In addition, app category usage distribution at each timestep is represented as a 20-dimensional vector over app categories and projected into an 8-dimensional latent representation using an MLP-based encoder. Traffic demand, including uplink and downlink traffic, is modeled as a 2-dimensional sequence and similarly encoded into an 8-dimensional latent representation using another MLP-based encoder. All resulting latent representations are jointly fed into the UNet-based diffusion model for spatio-temporal generation of mobility, app usage, and traffic dynamics.

After generating sequences of latent embeddings using the diffusion-based model, we map them back to the original space. Specifically, for the generated location embeddings $\{\hat{\mathbf{z}}_1, \hat{\mathbf{z}}_2, \ldots, \hat{\mathbf{z}}_n\}$, first compute the embedding set $\mathbf{Z} = \{\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_m\}$ for all candidate locations in the target city using the autoencoder. Each generated embedding $\hat{\mathbf{z}}_i$ is then mapped to its nearest neighbor in $\mathbf{Z}$ based on Euclidean distance, yielding trajectories grounded in real-world geography. In contrast, app category usage distributions and traffic sequences are directly reconstructed from the generated latent representations using their respective decoders.

D.3 Evaluation Methodology

Dataset Description. We evaluate the proposed model on two real-world mobile network datasets collected from Shanghai and Nanchang, respectively. The Shanghai dataset is a public research dataset collected by an Internet service provider (Yu et al., 2018), covering the entire Shanghai metropolitan region over one week, from 19 April 2016 to 25 April 2016. The Nanchang dataset was collected by China Mobile and contains 19,929,538 records from 2,000 users across the entire metropolitan area of Nanchang, a provincial capital city in China. It covers seven consecutive days, from 19 May 2022 to 25 May 2022. Each record in both datasets contains an anonymized user ID, application category ID, location ID, traffic volume, and timestamp. Dataset statistics are summarized in Table 12. For the Nanchang dataset, the raw records are not publicly redistributed due to privacy considerations. All user identifiers are anonymized, and the data are stored on protected servers with access control and firewall protection. We only report aggregated dataset statistics and evaluation results, without releasing individual-level trajectories, app-usage traces, or raw traffic records.

Evaluation Metrics. We assess the quality of generated user dynamics from three perspectives: mobility patterns, app usage, and traffic dynamics. For mobility evaluation, we use the Kolmogorov–Smirnov (KS) test to measure the distributional discrepancy between real and generated trajectories over several trajectory statistics, including radius of gyration (Radius), travel distance (Distance), trajectory duration (Duration), and daily visited locations (DailyLoc). Radius measures the spatial range of a user's movements; Distance measures the total travel length along a user's trajectory; Duration measures the temporal span of a trajectory; and DailyLoc measures the number of distinct locations visited by a user within a day. For app usage and traffic demand, we evaluate both distributional similarity and point-wise accuracy. Jensen–Shannon divergence (JSD) is used to quantify distributional similarity, while normalized mean absolute error (NMAE) is used to measure sequence-level deviations between real and generated values. Overall, all reported values in the

quantitative evaluation tables are lower-is-better: Radius, Distance, Duration, and DailyLoc report KS-test distances over the corresponding mobility statistics, while JSD and NMAE measure distributional discrepancy or prediction error.

Training Details. The Shanghai and Nanchang datasets are both split into training, validation, and testing sets with a ratio of 7:2:1. The split is conducted at the user level, so that there is no overlap of users among the three sets, thereby avoiding user-level data leakage. The model is trained for a maximum of 500 epochs with early stopping based on validation performance. Specifically, training is stopped when the validation performance does not improve for 10 consecutive epochs. The training batch size is 256, and the learning rate is set to 0.0005. All experiments are conducted on a single NVIDIA A100-SXM4-80GB GPU.

Baselines. We compare the proposed model with representative trajectory and user behavior generation methods, which are outlined below:

EPR (Simini et al., 2012) is a classical human mobility model that captures individual movement patterns through exploration of new locations and preferential return to frequently visited places.

TimeGeo (Jiang et al., 2016) is a model-based trajectory generation method that captures temporal activity patterns via dwell rate, burst rate, and home-based tour frequency, while modeling spatial movement using the exploration and preferential return (EPR) mechanism.

MoveSim (Feng et al., 2020) is a generative adversarial framework for trajectory generation that incorporates domain knowledge of human mobility regularities to improve the realism of generated movement patterns.

Pategail (Wang et al., 2023) is a generative adversarial imitation learning model that replicates human decision-making processes for trajectory generation.

DiffTraj (Zhu et al., 2023) is a diffusion-based model for trajectory generation that conditions on spatio-temporal features to capture realistic mobility patterns.

TimeGAN (Yoon et al., 2019) is a generative adversarial network for time-series data that incorporates step-wise supervised loss to better preserve temporal dynamics.

LDM (Rombach et al., 2022) performs diffusion in a latent space using a pre-trained autoencoder, enabling more efficient and expressive generation.

Diffusion-TS (Yuan & Qiao, 2024) is a diffusion-based generative framework for time-series data that employs an encoder-decoder transformer with disentangled temporal representations to capture complex temporal dependencies.

CSDI (Tashiro et al., 2021) is a conditional score-based diffusion model, originally developed for time-series imputation, that learns to recover missing values by modeling the data distribution conditioned on observed inputs.

AppGen (Huang et al., 2026a) is an autoregressive diffusion-based model for app usage generation conditioned on user mobility, capturing sequential dependencies via latent spatio-temporal representations.

D.4 Additional Results

We provide additional quantitative and visualization results to evaluate the realism of the generated user dynamics.

Quantitative Results on the Nanchang Dataset. Table 13 reports the additional evaluation results on the Nanchang dataset. For mobility generation, our model achieves the best performance on Distance and Duration and the second-best performance on Radius, showing strong performance in capturing stable and recurrent mobility patterns. Its relatively weaker performance on DailyLoc is mainly because DailyLoc measures the number of distinct locations visited by a user within a day, reflecting the diversity of daily spatial exploration, while model tends to generate regular daily routines and may underrepresent occasional or low-frequency visits to peripheral locations, making DailyLoc more challenging than other mobility metrics. For app usage and traffic generation, our

Table 13: User behavior generation results on the Nanchang dataset. The best result is highlighted in bold, and the second-best result is underlined. Our model outperforms existing baselines across most metrics in mobility, app usage, and traffic generation.

(a) Mobility generation.

Method	Radius	Distance	Duration	DailyLoc
EPR	0.6417 ± 0.004	0.9285 ± 0.021	0.8736 ± 0.016	0.7914 ± 0.018
TimeGeo	0.5893 ± 0.004	0.7328 ± 0.018	0.7249 ± 0.014	0.6847 ± 0.015
MoveSim	0.4426 ± 0.003	0.3017 ± 0.009	0.3864 ± 0.012	0.6525 ± 0.013
Pategail	0.8131 ± 0.005	0.6869 ± 0.003	0.1585 ± 0.0024	0.4551 ± 0.005
DiffTraj	**0.3602 ± 0.013**	0.4424 ± 0.016	0.1625 ± 0.0080	**0.1795 ± 0.010**
Ours	0.3858 ± 0.005	**0.1526 ± 0.004**	**0.1504 ± 0.004**	0.6802 ± 0.003

(b) App usage and traffic generation.

Method	App Category		Traffic	
	JSD	NMAE	JSD	NMAE
Time-GAN	0.0316 ± 0.004	0.29354 ± 0.007	0.0351 ± 0.007	0.4842 ± 0.010
LDM	0.0299 ± 0.002	0.2762 ± 0.007	0.0314 ± 0.006	0.4601 ± 0.007
Diffusion-TS	0.0284 ± 0.003	0.2702 ± 0.006	0.0295 ± 0.006	0.4488 ± 0.008
CSDI	0.0268 ± 0.003	0.2579 ± 0.004	0.0282 ± 0.005	0.4273 ± 0.004
AppGen	0.0237 ± 0.002	0.2368 ± 0.004	0.0270 ± 0.003	0.4102 ± 0.003
Ours	**0.0210 ± 0.002**	**0.2106 ± 0.002**	**0.0251 ± 0.003**	**0.3896 ± 0.002**

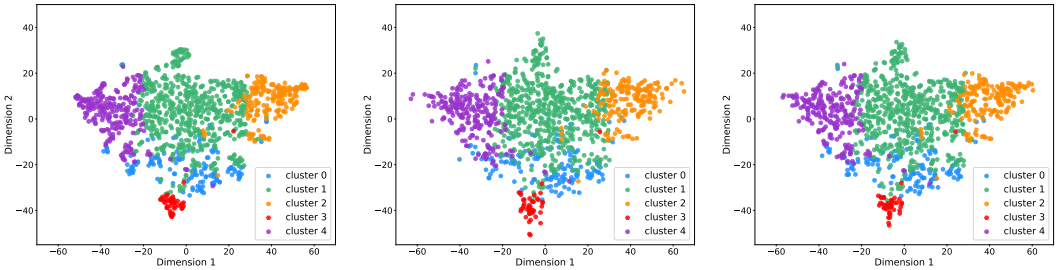

(a) Real (b) LDM (c) Ours

Figure 19: App category usage distribution across locations for real data, LDM-generated data, and our proposed model in the Shanghai dataset. Our model shows better alignment with real data clusters, indicating improved modeling of location-dependent app usage patterns.

model achieves the best performance across all JSD and NMAE metrics. These results demonstrate that the proposed model can jointly generate mobility trajectories, app usage, and traffic demand with balanced overall performance across multiple user behavior generation tasks.

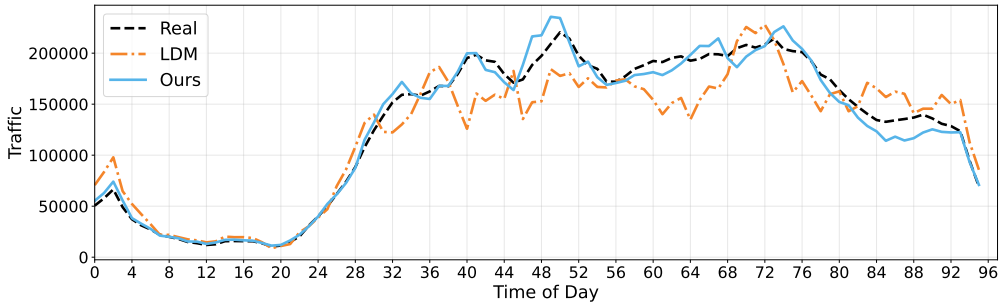

Figure 20: Traffic demand over time comparison between real data, LDM-generated data, and our proposed model in the Shanghai dataset. The proposed model more closely follows the temporal trends and patterns of real data.

App Usage Distribution across Locations. We analyze spatial patterns of app usage on the Shanghai dataset by aggregating app-category distributions at each location and applying K-means clustering. The clustered representations are further projected into a two-dimensional space using t-SNE for visualization. The results are demonstrated in Figure 19. Compared with LDM, our model shows better alignment with real data clusters, indicating a more accurate modeling of location-dependent app usage patterns.

Traffic over Time. We further analyze the temporal evolution of traffic demand over a day on the Shanghai dataset. As shown in Figure 20, the proposed model better aligns with the overall trend and patterns of real data compared to LDM-generated results. This demonstrates its effectiveness in capturing both global temporal structure and fine-grained variations.

E MOBILE PLANNING AND OPTIMIZATION: DETAILED IMPLEMENTATION

E.1 OPTIMIZATION ALGORITHMS

Reinforcement Learning (RL) RL offers a complementary optimization paradigm, in which an agent learns to make sequential decisions by interacting with an environment and maximizing a cumulative reward signal Sutton & Barto (2018). Unlike gradient-based methods, RL makes no assumptions about the differentiability or structure of the objective, and the policy is refined purely through trial-and-error feedback. This makes RL particularly well-suited for wireless network optimization tasks, where the reward landscape is non-convex, the state space is high-dimensional, and the environment dynamics are stochastic Alwarafy et al. (2022).

In this work, we adopt Proximal Policy Optimization (PPO) Schulman et al. (2017) as a representative deep RL algorithm. PPO is a policy gradient method that stabilizes training by constraining the magnitude of each policy update via a clipped surrogate objective, preventing destructively large steps while still enabling efficient multi-epoch updates on each collected trajectory. Its favorable balance between implementation simplicity, sample efficiency, and training stability has made it one of the most widely adopted algorithms in both continuous control and discrete decision-making settings. Within MobiSim, the agent observes the full network state and receives dense feedback from the simulator, providing a rich learning signal that enables PPO to converge reliably without additional reward shaping.

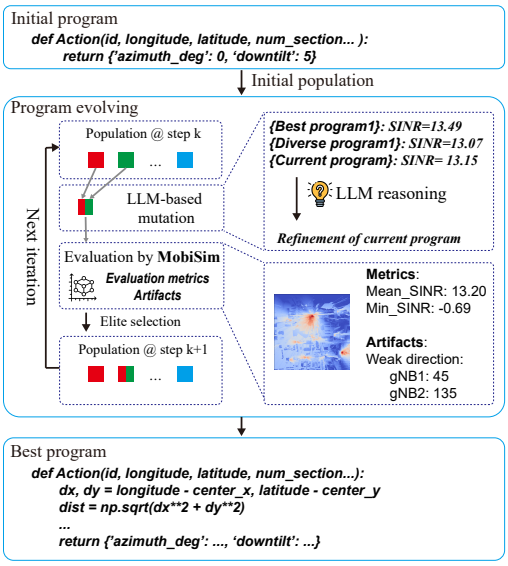

Figure 21: Adapting LLM-based evolving algorithm to the MobiSim platform. Here, we take the base station coverage optimization as an example.

LLM-based evolving algorithm. Recent advances in large language models (LLMs) have enabled a new class of program-based optimization methods, where solutions are represented as executable code and iteratively refined through LLM reasoning. These approaches have shown strong potential in tackling combinatorial optimization and control problems such as CPU scheduling Novikov et al. (2025); Agrawal et al. (2026). Rather than relying on gradient information or task-specific parameterizations, they operate directly on functional programs, making them particularly suitable for complex environments with heterogeneous inputs and non-differentiable objectives. In this work, we adopt this paradigm as a representative class of general-purpose optimizers to demonstrate the capability of MobiSim to support advanced, program-driven decision-making. We find that such methods can be seamlessly instantiated within MobiSim and benefit from its high-fidelity, full-stack simulation, and rich feedback signals.

Following Alphaevolve Novikov et al. (2025), we instantiate a program-evolving optimization loop within MobiSim (Figure 21). The procedure starts from an initial population of simple, task-agnostic programs such as constant or heuristic mappings from inputs to actions, which serve as seeds for subsequent refinement. At each iteration, new candidate programs are generated via LLM-guided mutation and recombination, where the model edits existing programs conditioned on their past performance and structural diversity. Each candidate program is then evaluated through the MobiSim environment. Importantly, the simulator provides two complementary forms of feedback. First, it returns task-specific scalar metrics (SINR, throughput, or coverage), which are used as fitness signals for selection. Second, it allows extra structured artifacts derived from the simulation process, such as spatial performance maps or identified weak coverage directions. These artifacts go beyond standard reward signals and offer interpretable, fine-grained feedback about the environment and system behavior. Based on the scalar metrics, top-performing programs are selected and used to form the next generation. This process iterates, progressively improving program quality through the interplay between LLM-based program refinement and simulator-in-the-loop evaluation.

E.2 BS Deployment Planning

Problem formulation We consider a target geographic region Ω that is discretized into a set of regular grids $\mathcal{G} = \{1, 2, \ldots, M\}$. The center coordinate of grid m is denoted by $\mathbf{q}_m$. The task is to deploy N BSs within this region to improve the overall coverage quality. The deployment location of BS i is represented as $\mathbf{p}_i = [x_i, y_i]^{\mathrm{T}}$, and the overall deployment configuration is defined as:

$$\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^{N}. \tag{86}$$

Given a deployment configuration $\mathbf{P}$, MobiSim evaluates the corresponding electromagnetic propagation condition and generates an RSRP map over all grid cells:

$$\boldsymbol{\rho}(\mathbf{P}) = \mathcal{S}(\mathbf{P}, \Omega), \tag{87}$$

where $\mathcal{S}(\cdot)$ denotes the simulation function of MobiSim, and $\boldsymbol{\rho}(\mathbf{P}) = \{\rho_m(\mathbf{P})\}_{m\in\mathcal{G}}$ denotes the simulated RSRP values of all grids under deployment configuration $\mathbf{P}$. Here, $\rho_m(\mathbf{P})$ represents the RSRP at the center of grid m.

A grid is regarded as weakly covered if its simulated RSRP is lower than a predefined coverage threshold ρ_{th}:

$$z_m(\mathbf{P}) = \mathbb{I}\left(\rho_m(\mathbf{P}) < \rho_{\mathrm{th}}\right), \tag{88}$$

where $\mathbb{I}(\cdot)$ is the indicator function. Accordingly, the weak coverage ratio over the entire region is defined as:

$$\xi(\mathbf{P}) = \frac{1}{|\mathcal{G}|} \sum_{m\in\mathcal{G}} z_m(\mathbf{P}). \tag{89}$$

This metric can also be reported as a percentage, i.e., $100\xi(\mathbf{P})\%$.

The BS deployment planning problem aims to minimize the network-wide weak coverage ratio by optimizing the deployment coordinates of all BSs. It is formulated as the following constrained optimization problem:

$$\begin{aligned} \min_{\mathbf{P}} \quad & \xi(\mathbf{P}) \\ \text{s.t.} \quad & \mathbf{p}_i \in \Omega_{\mathrm{free}}, \quad \forall i \in \mathcal{N}, \\ & \|\mathbf{p}_i - \mathbf{p}_j\|_2 \geq d_{\min}, \quad \forall i, j \in \mathcal{N},\ i \neq j, \end{aligned} \tag{90}$$

where $\mathcal{N} = \{1, 2, \ldots, N\}$ denotes the BS set, $\Omega_{\mathrm{free}} = \Omega \setminus \Omega_{\mathrm{bld}}$ denotes the feasible deployment region after excluding building areas, Ω_{bld} denotes the set of building footprints, and $d_{\min}$ is the minimum distance required between any two deployed BSs.

E.3 Network Coverage Optimization

Problem formulation We consider a cellular network comprising N BSs, where each BS i is deployed with $S = 3$ sectors and each sector s is equipped with a dedicated antenna. The antenna at sector (i, s) is characterized by two configurable angular parameters: the azimuth angle ϕ and the

downtilt angle θ. The network is spatially discretized into a grid of $|\mathcal{G}|$ cells, and the SINR at each grid point g is determined by the joint antenna configuration across all BSs and sectors. The objective is to maximize the average SINR across all grid cells:

$$\max_{\{\phi_{i,s},\,\theta_{i,s}\}} \quad \frac{1}{|\mathcal{G}|}\sum_{g\in\mathcal{G}} \mathrm{SINR}_g(\{\phi_{i,s},\theta_{i,s}\}) \tag{91}$$

subject to:

$$\phi_{i,s}\in[0,\,2\pi), \quad \theta_{i,s}\in\left[-\frac{\pi}{2},\,\frac{\pi}{2}\right], \quad \forall\, i\in\mathcal{N},\ s\in\mathcal{S} \tag{C1}$$

$$\Delta_{i,s}^{(t)} \le K, \quad \forall\, i\in\mathcal{N},\ s\in\mathcal{S} \tag{C2}$$

where $\Delta_{i,s}^{(t)}$ denotes the total number of antenna parameter adjustments for sector s of BS i within a sliding window of W time steps up to time t, and K is the maximum allowable reconfiguration count to prevent excessive hardware wear and signaling overhead. To address this non-convex, high-dimensional problem with a black-box simulator, we formulate it as a Markov Decision Process, where the reward is defined as the instantaneous network-wide average SINR $r_t = \frac{1}{|\mathcal{G}|}\sum_{g\in\mathcal{G}} \mathrm{SINR}_g^{(t)}$. The state at time t is defined as:

$$\mathbf{s}_t = \left\{ \mathbf{q}_u,\ \mathbf{l}_i,\ \phi_{i,s}^{(t-1)},\ \theta_{i,s}^{(t-1)},\ \mathrm{SINR}_g^{(t-1)} \;\middle|\; u\in\mathcal{U},\ i\in\mathcal{N},\ s\in\mathcal{S},\ g\in\mathcal{G} \right\} \tag{92}$$

where $\mathbf{q}_u$ denotes the position of user u, $\mathbf{l}_i$ denotes the location of BS i, $\phi_{i,s}^{(t-1)}$ and $\theta_{i,s}^{(t-1)}$ are the antenna angles from the previous time step, and $\mathrm{SINR}_g^{(t-1)}$ is the per-grid SINR observed at the previous step. The action $\mathbf{a}_t = \{\phi_{i,s}^{(t)}, \theta_{i,s}^{(t)}\}$ specifies the updated azimuth and downtilt angles for each sector antenna.

E.4 QoE-Driven Throughput Optimization

Problem formulation We consider a multi-cell network where each BS i serves a set of users $\mathcal{U}_i$ under a TDD frame structure. The system has two independent decision variables per BS: (i) the TDD UL/DL time-slot pattern τ_i, selected from $|\mathcal{T}| = 41$ standardized TDD configurations as specified in Table 7, and (ii) the PRB allocation policy for distributing radio resources among served users, parameterized by a per-user importance weight vector $\mathbf{w}_i^d \in \mathbb{R}^{|\mathcal{U}_i|}$ for each link direction $d\in\{\mathrm{UL},\mathrm{DL}\}$, from which a normalized allocation probability is derived via softmax:

$$\pi_{i,u}^d = \frac{\exp(w_{i,u}^d)}{\sum_{u'\in\mathcal{U}_i}\exp(w_{i,u'}^d)} \tag{93}$$

so that the number of PRBs allocated to user u in direction d is $b_{i,u}^d = \lfloor \pi_{i,u}^d \cdot B_{\max}^d \rfloor$. We adopt a QoE-centric objective defined by the demand satisfaction ratio. For user u at time t, the DL and UL satisfaction ratios are:

$$\eta_u^{\mathrm{DL}}(t) = \frac{R_u^{\mathrm{DL}}(t)}{D_u^{\mathrm{DL}}(t)}, \qquad \eta_u^{\mathrm{UL}}(t) = \frac{R_u^{\mathrm{UL}}(t)}{D_u^{\mathrm{UL}}(t)} \tag{94}$$

where $R_u^{\mathrm{DL}}(t)$, $R_u^{\mathrm{UL}}(t)$ are the achieved downlink and uplink service rates, and $D_u^{\mathrm{DL}}(t)$, $D_u^{\mathrm{UL}}(t)$ are the corresponding traffic demands. The optimization problem is:

$$\max_{\{\tau_i,\,\mathbf{w}_i^{\mathrm{UL}},\,\mathbf{w}_i^{\mathrm{DL}}\}_{i=1}^N} \quad \frac{1}{2*|\mathcal{U}|}\sum_{u\in\mathcal{U}}\Big(\eta_u^{\mathrm{DL}}(t)+\eta_u^{\mathrm{UL}}(t)\Big) \tag{95}$$

To solve this problem, we formulate it as an MDP. The state at time t is defined as:

$$\mathbf{s}_t = \left\{ R_u^{d,(t-1)},\ D_u^{d,(t-1)},\ \tau_i^{(t-1)},\ \mathbf{w}_i^{d,(t-1)} \;\middle|\; u\in\mathcal{U}_i,\ i\in\mathcal{N},\ d\in\{\mathrm{UL},\mathrm{DL}\} \right\} \tag{96}$$

where $R_u^{d,(t-1)}$ and $D_u^{d,(t-1)}$ denote the achieved service rate and traffic demand of user u in direction d at the previous time step, and $\tau_i^{(t-1)}$, $\mathbf{w}_i^{d,(t-1)}$ are the TDD pattern and PRB importance weights applied at the previous step. The action $\mathbf{a}_t = \{\tau_i^{(t)}, \mathbf{w}_i^{\mathrm{UL},(t)}, \mathbf{w}_i^{\mathrm{DL},(t)}\}_{i=1}^N$ comprises the TDD pattern selection and per-user PRB importance weights for each BS, and the reward directly reflects the network-wide QoE demand satisfaction.

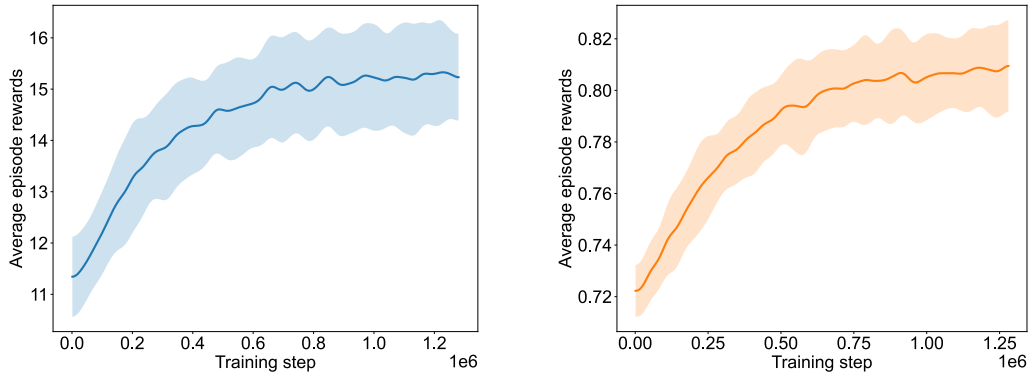

(a) Network coverage optimization

(b) QoE-driven throughput optimization

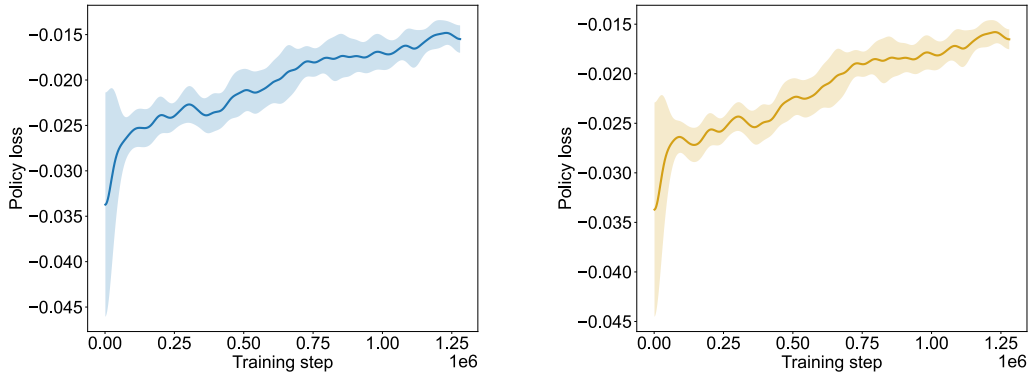

(c) Network coverage optimization

(d) QoE-driven throughput optimization

Figure 22: Results of the proposed optimization.

E.5 Optimization Results of RL and LLM-based Schemes

RL Figure 22 presents the training dynamics of the PPO agent across both optimization tasks. In the upper row, the average episode reward increases monotonically throughout training for both the network coverage task (Figure 22a) and the QoE-driven throughput task (Figure 22b), demonstrating that the agent consistently improves its policy within the MobiSim environment. For network coverage optimization, the reward rises from approximately 11.2 to 15.7 over 1.28×10^6 training steps, while the QoE-driven task yields a demand satisfaction ratio improving from 0.72 to above 0.79, indicating that the learned PRB allocation and TDD pattern selection effectively adapt to user traffic demands. The shaded regions, representing the variance across independent runs, remain well-bounded throughout training, suggesting stable convergence behavior under both problem settings.

The corresponding policy loss curves are shown in the lower row of Figure 22. In both tasks, the policy loss remains strictly negative across the entire training horizon, which, under the PPO clipped surrogate objective Schulman et al. (2017), confirms that each gradient update yields a net improvement over the previous policy. The loss magnitude gradually diminishes and stabilizes, consistent with the reward plateauing in later training stages. Taken together, these results validate that PPO can be seamlessly instantiated within MobiSim and reliably optimizes both antenna configuration and radio resource allocation without any task-specific reward shaping or manual tuning.

LLM-based evolving algorithm

- **BS Deployment Planning.** In our AlphaEvolve framework, we optimize the locations of $N = 5$ BSs to minimize the weak coverage ratio over the entire area. At each iteration, the

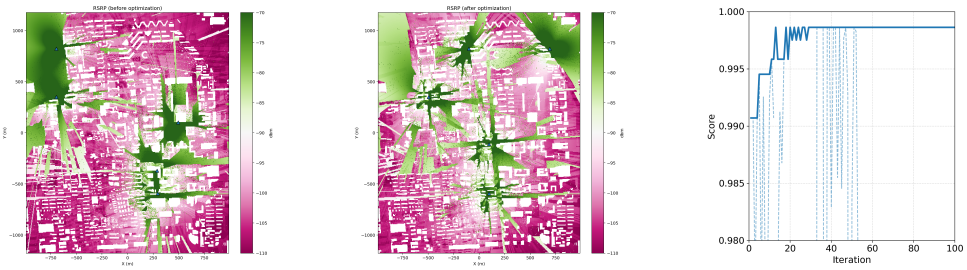

(a) Before (b) After (c) Score

Figure 23: Results of AlphaEvolve for BS deployment. (a) RSRP map before optimization. (b) RSRP map after optimization. (c) Learning curve of the evolving process.

LLM generates an executable deployment program, which is then evaluated by the simulator. The evolved program receives an instance containing the grid size, the number of BSs, the candidate grid IDs, grid resolution, and the constraints including the minimum inter-BS distance, and the building mask. It outputs a list of selected N grid IDs, $\mathbf{g} = \{g_1, \ldots, g_N\}$, where each g_i denotes the deployment grid cell of one BS. The evaluator first validates the generated deployment. A candidate solution is regarded as invalid if any selected grid ID is outside the candidate set, if duplicate sites exist, or if a selected site falls inside the building mask. For valid solutions, the evaluator converts the selected grid IDs into simulator configurations and runs the radio map simulation. The aggregated RSRP map returned by the simulator is then used to evaluate the weak coverage ratio:

$$r_{\text{weak}} = \frac{1}{|\Omega_{\text{v}}|} \sum_{c \in \Omega_{\text{v}}} \mathbb{I}\left[\text{RSRP}(c) < \tau_{\text{weak}}\right], \quad \tau_{\text{weak}} = -100 \text{ dBm}, \tag{97}$$

where Ω_{v} denotes the set of grid cells. To avoid degenerate deployments with overly close sites, we further add a soft distance penalty:

$$P_{\text{dist}} = \lambda_{\text{d}} \sum_{1 \le i < j \le N} \frac{\max\left(0, d_{\min} - \|\mathbf{x}_{g_i} - \mathbf{x}_{g_j}\|_2\right)}{d_{\min}}, \tag{98}$$

where $\mathbf{x}_{g_i}$ is the physical coordinate of grid cell g_i, $d_{\min}$ is the minimum allowed inter-site distance, and λ_{d} is the penalty weight. The optimization loss and combined score are defined as

$$\mathcal{L}(\mathbf{g}) = r_{\text{weak}}(\mathbf{g}) + P_{\text{dist}}(\mathbf{g}), \quad S(\mathbf{g}) = 1 - \mathcal{L}(\mathbf{g}). \tag{99}$$

If a constraint is violated, the evaluator directly assigns an invalid score, preventing such programs from being selected in later evolution. In addition to the scalar score, the evaluator also provides spatial feedback, including the selected BS locations, the weak coverage centroid, the top weak coverage blocks, and the local weak ratio around each selected BS. This feedback allows the evolved programs to progressively exploit previously observed weak regions without directly calling the simulator inside the generated program.

We use GPT-4o and DeepSeek-V3.2 as the base LLMs. As shown in Fig. 23, the optimized deployment differs from the initial deployment, indicating that the evolutionary process effectively searches the feasible BS placement space. The optimization score increases over iterations and gradually stabilizes. The corresponding weak coverage ratio decreases to 0.00136, demonstrating that the proposed AlphaEvolve-based deployment optimization framework can effectively improve coverage performance on the simulator platform.

- **Network Coverage Optimization.** In the network coverage optimization problem, we adopt a multi-agent distributed modeling approach, where each base station is managed independently by the decision program. The input state describes the station's physical operation states, including location, the number of sectors, and the initial azimuth orientation of each sector. The output action consists of per-sector adjustments: the adjusted azimuth and the downtilt. This input-output mapping is learned by AlphaEvolve iteratively, enabling each agent to optimize its own station's configuration based on local state information. The

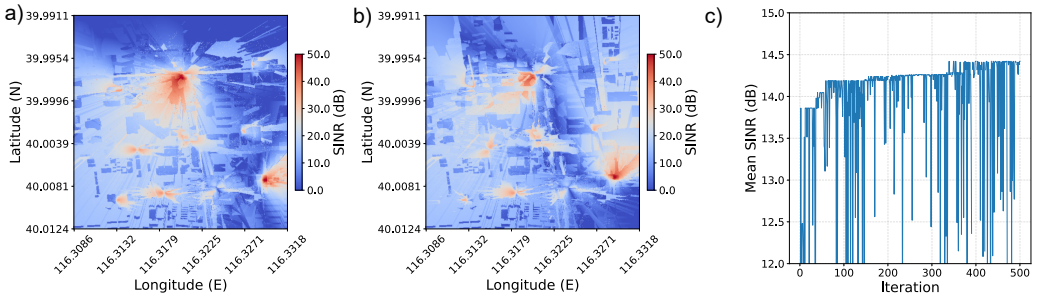

Figure 24: Effectiveness of the base station antenna coverage optimization. a) SINR map before the optimization. b) SINR map after AlphaEvolve optimization. c) Learning curve of the evolving process.

distributed modeling approach offers excellent scalability, as the number of agents can grow with the network size without central coordination overhead.

The overall reward is defined as the average SINR across the entire grid, which reflects the global coverage and interference condition. To facilitate interpretability and per-site analysis, we additionally compute an artifact for each base station. Specifically, for a given station, we extract the SINR grid around its location (clipped at the domain boundaries), partition it into 4×4 equal subregions, and compute the mean SINR within each subregion. We then identify the subregion with the lowest average SINR and calculate the angle from the base station to the center of that subregion. This artifact, while not part of the reward signal, provides insight into the weakest local coverage direction and assists the azimuth adjustment.

All experiments are conducted on our benchmark simulation platform, which faithfully models realistic propagation, antenna patterns, and network constraints. As shown in Figure 24, the average SINR steadily increases from 12.47dB to 14.42dB over the course of optimization, demonstrating consistent convergence. More importantly, the coverage map reveals that most of the low-SINR area in the northern region is effectively eliminated after optimization. These results not only validate the effectiveness of the proposed AlphaEvolve-based optimization, but also confirm the reliability of the simulation platform: it provides reliable, reproducible performance evaluation and informative diagnostic feedback, making it a solid benchmark for future research on base station configuration learning.

- **QoE-Driven Throughput Optimization.** We further apply AlphaEvolve to optimize the QoE under a fixed BS deployment. The evolved program only optimizes the TDD time-domain strategy and the PRB allocation policy. Specifically, it outputs a TDD pattern vector and PRB strategies. At each iteration, the LLM rewrites an executable program based on the problem description and previous evaluation metrics. The evolved program receives an instance containing the number of BSs, the number of users, the valid TDD pattern IDs, and the available PRB policies. The instance may also include simulator feedback from previous evaluations, such as BS-level traffic demand, served traffic, and demand satisfaction ratio, which allows the generated program to adapt the TDD pattern of each BS according to the spatially heterogeneous traffic load.

  The evaluator first validates the generated decision, and calls the simulator to obtain the achieved throughput and the corresponding traffic demand for each user. The downlink and uplink demand satisfaction ratios are $\eta_u^{\mathrm{DL}}(t)$ and $\eta_u^{\mathrm{UL}}(t)$, respectively. The QoE objective $Q(\mathbf{p}, \pi_{\mathrm{PRB}})$ is then defined as the average uplink and downlink demand satisfaction ratio according to Equation 95. Therefore, the final optimization loss and score for AlphaEvolve framework are

$$\mathcal{L}(\mathbf{p}, \pi_{\mathrm{PRB}}) = 1 - Q(\mathbf{p}, \pi_{\mathrm{PRB}}) + P_{\mathrm{con}}, \quad S(\mathbf{p}, \pi_{\mathrm{PRB}}) = Q(\mathbf{p}, \pi_{\mathrm{PRB}}) - P_{\mathrm{con}}, \quad (100)$$

  where P_{con} denotes the constraint penalty for validating the generated TDD pattern IDs and PRB strategy. Thus, a higher score indicates that more users satisfy their uplink and downlink service demands while satisfying the PRB constraints.

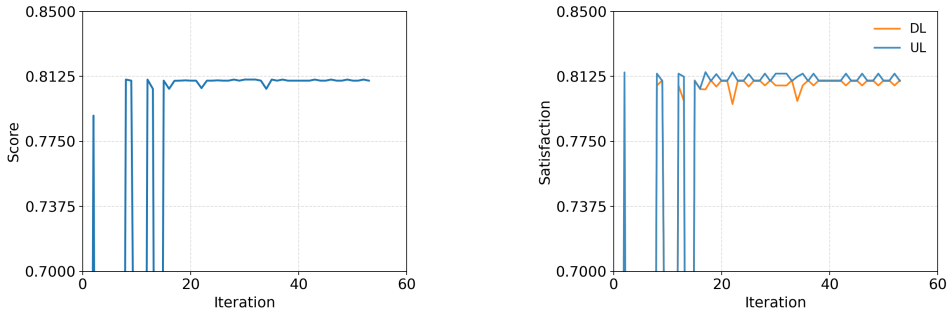

Figure 25: Learning curve and satisfaction of AlphaEvolve for QoE optimization.

In addition to the score, the evaluator provides compact feedback, including BS-level downlink and uplink demand, served traffic, user count, and satisfaction ratios. This enables AlphaEvolve to progressively discover interpretable resource-allocation heuristics. As shown in Fig. 25, the score increases with iterations and gradually converges to 0.8107, demonstrating the effectiveness of the proposed AlphaEvolve-based QoE optimization framework.

F Experiment Setting

This appendix details the hardware platform on which all experiments reported in the main paper are run, and enumerates the default simulator parameters that drive the scene construction, the electromagnetic ray tracing, the network access and handover logic, and the resource scheduling.

F.1 Hardware Platform

Table 14 summarizes the full configuration for the evaluation experiments.

Table 14: Hardware platform used for evaluation experiments.

Component	**Specification**
GPU	NVIDIA A100-SXM4-80GB $\times$ 3
CPU model	Intel Xeon Platinum 8358P @ 2.60 GHz
CPU architecture	x86_64
Physical CPU sockets	2
Cores per socket	32
Total physical cores	64

F.2 Default Global Simulation Parameters

The global panel fixes the high-level simulation budget that is shared by every layer of MobiSim: the random seed used by all stochastic components, the number of base stations sampled into the scene, the maximum reflection depth of the ray tracer, the geographic extent of the default urban region, the receiver-grid resolution, the UE plane height, the per-source ray budget, and the GPU list used for parallel radio-map computation. The default values are listed in Table 15.

F.3 Default Network Configurations

The default network-layer configuration covers the per-sector antenna setup, the initial-access and handover thresholds, and the per-slot scheduling logic, with values chosen to reflect a widely deployed 5G NR mid-band deployment. The per-element radiation pattern follows TR 38.901 and is described in Appendix B.2; the full A1 to A5 measurement-event family is enabled by default and detailed in

Table 15: Default global simulation parameters.

Parameter	Default value	Role
Random seed	72	Controls all stochastic components
Number of base stations	10	Tri-sector gNodeBs sampled in the scene
Reflection depth $K_{\max}$	5	Maximum bounces of the ray tracer
Base-station height	25 m	Antenna mounting height
Receiver grid X	500	Grid columns over the bounding box
Receiver grid Y	500	Grid rows over the bounding box
UE plane height z_0	1.5 m	Receiver-grid altitude
Rays per source	10^5	Monte Carlo budget per transmitter sector
GPU list	0, 1, 2	GPUs used for ray-tracing and radio-map computation

Appendix B.3.3; and the TDD slot formats follow the TS 38.213 catalogue listed in Appendix B.4.1. The full set of default values is consolidated in Table 16.

Table 16: Default network-layer configurations.

Parameter	Default value	Range / set
Sector and Antenna Parameters		
Sectors per gNodeB	3	$[1, 12]$
Sector azimuth orientation	$0, 120, 240^\circ$	$[0^\circ, 359.9^\circ]$
Mechanical downtilt	5°	$[-30^\circ, 30^\circ]$
Electrical downtilt	10°	$[-30^\circ, 30^\circ]$
Carrier frequency	3.5 GHz	$[0.6, 100]$ GHz
Transmit power	46 dBm	$[0, 80]$ dBm
Channel bandwidth	100 MHz	$[5, 400]$ MHz
Subcarrier spacing	30 kHz	$\{15, 30, 60, 120\}$ kHz
Panel rows M	1	$[1, 8]$
Panel columns N	1	$[1, 8]$
Layout random seed (rand. layout)	2026	$[0, 10^5]$
Access and Handover Parameters		
Access RSRP threshold	-120 dBm	$[-160, -60]$ dBm
Handover hysteresis	2.5 dB	$[0, 15]$ dB
Time-to-trigger	1 step	$[1, 20]$ steps
Scheduling Parameters		
Duplex mode	TDD	$\{\text{TDD}, \text{FDD}\}$
Fallback TDD pattern (TS 38.213)	30	$\{0, 1, \ldots, 40\}$

F.4 Computation Time Complexity

Table 17 reports the simulation timing across five geographic area sizes (100 m $\times$ 100 m to 3 km $\times$ 3 km) and one to three sectors per area. Cold-start time, which covers map loading and worker initialization, ranges from 49.62 s for the smallest area to 64.73 s for the largest, reflecting the increased overhead of preparing a larger environment. The cold-start phase additionally performs one-off operations, such as fetching OSM map data online, which contributes extra overhead at first launch; when running optimization tasks, these cold-start operations are not repeated across subsequent runs. Once the system is warmed up, the per-step simulation time remains highly efficient, spanning 5.09 s to 7.91 s across all configurations, where each step corresponds to a single electromagnetic computation in the RL loop. Notably, increasing the number of sectors from 1 to 3 introduces only a modest overhead of less than 1 s per step within any given area size, demonstrating that the simulator scales gracefully with the number of sectors and remains practical for RL training across a wide range of spatial scales.

Table 17: Simulation timing results for the electromagnetic computation step across different geographic area sizes and sector counts. Cold-start time includes map loading and worker initialization. Per-step simulation time is measured after warm-up.

Area Size	Cold-Start Time	Sectors	Per-Step Simulation Time
3 km × 3 km	64.73 s	1	7.18 s
		2	7.54 s
		3	7.91 s
2 km × 2 km	60.31 s	1	6.47 s
		2	6.82 s
		3	7.13 s
1 km × 1 km	56.08 s	1	5.96 s
		2	6.37 s
		3	6.74 s
500 m × 500 m	52.84 s	1	5.43 s
		2	5.71 s
		3	6.02 s
100 m × 100 m	49.62 s	1	5.09 s
		2	5.38 s
		3	5.67 s

F.5 Interactive UI

MobiSim provides an interactive execution interface, as illustrated in Figure 26. The UI is divided into six major modules: Global Simulation Parameters, OSM Region, Base Station Sector Configuration, Access and Handover Parameters, Scheduling Parameters, and Output Settings. The OSM Region module supports either manual input of the coordinate range for an arbitrary geographic area or direct region selection through an online OSM window. The Base Station Sector Configuration module supports both random initialization and manual configuration. After completing all settings, users can start the simulation by clicking the "Simulate" button. The simulation outputs the full-map visualization as well as the RSRP map for each sector (Figure 27). In addition, the UI provides statistics of the aggregated uplink and downlink traffic for each sector. It should be noted that, due to operations such as cold start and interface rendering, the simulation time in the UI is usually longer than that of backend simulator calls, such as those in RL processes.

G Generated Dataset Description

We release the data generated by MobiSim through the Hugging Face repository `MobiSim/MobiSim`. The release is designed to expose the simulator at three complementary levels: dense radio maps for spatial learning tasks, synchronized full-stack traces for spatiotemporal network modeling, and assembled 3D city assets for geometry-aware simulation. Unless explicitly stated otherwise, received-power values, channel-quality values, traffic records, and radio-access-network (RAN) events in this release are outputs of MobiSim rather than measurements collected from operational mobile networks.

G.1 Dataset Overview

The public snapshot used in this work contains 14.777 GiB of data and 24,609 collection files. It comprises 931 radio-map tiles from 11 cities, 222 full-stack tiles from three cities, and two assembled 3D city scenes. A radio map tile and a full-stack tile both cover an approximately $2\,\mathrm{km} \times 2\,\mathrm{km}$ geographic region; tiles appearing in different collections are counted separately because they provide different data products. Table 18 summarizes the three collections.

The collections intentionally overlap in geographic coverage. For example, Cardiff, Shanghai, and Tianjin occur in both the radio-map and full-stack collections. This overlap makes it possible to study

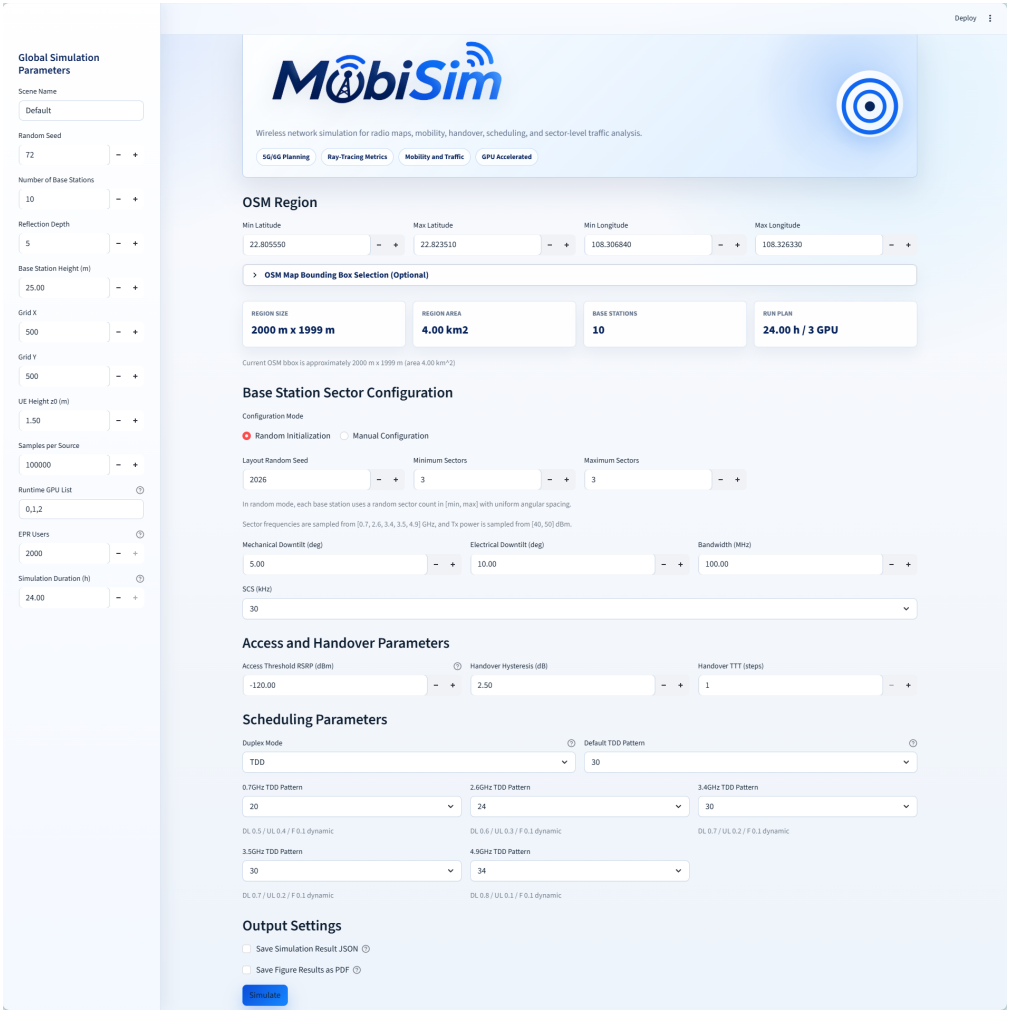

Figure 26: MobiSim UI main page.

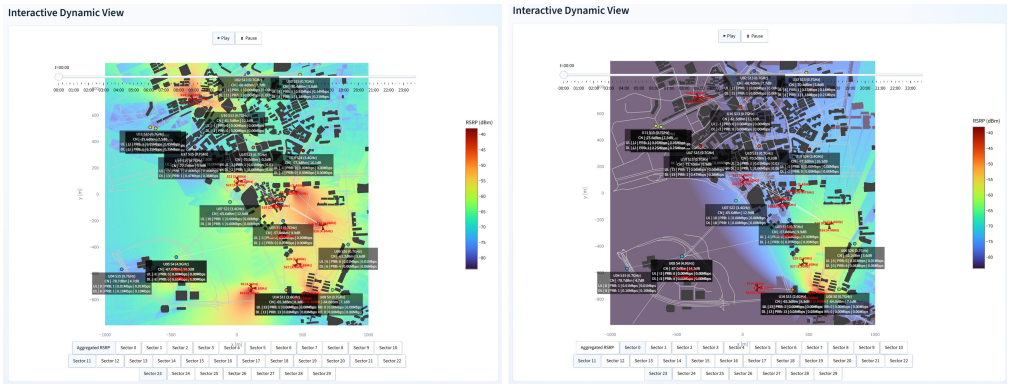

Figure 27: Interactive interface for MobiSim UI simulation results.

the same urban geometry at different levels of abstraction, from a static spatial field to time-dependent UE, cell, and RAN behavior.

Table 18: Summary of the released MobiSim data. File counts and sizes refer to the public snapshot used in this work.

Collection	Cities	Tiles/scenes	Files	Size (GiB)
Radio maps	11	931 tiles	7,459	4.600
Full-stack data	3	222 tiles	16,815	4.448
3D cities	2	2 scenes	335	5.729
Total	11 unique cities	1,153 tiles + 2 scenes	24,609	14.777

G.2 Radio-Map Collection

The `RadioMaps/` collection contains 931 urban tiles distributed across 11 cities. Table 19 gives the number of tiles for each city. Each leaf directory is named by its latitude–longitude bounding box, `<lat_min>_<lat_max>_<lon_min>_<lon_max>`, so that every sample can be associated with its geographic extent without a separate index.

Table 19: Geographic coverage of the radio-map collection.

City	Tiles	City	Tiles	City	Tiles
Beijing	90	Cardiff	42	Chongqing	64
Guangzhou	110	Hong Kong	72	Nanjing	81
Nanning	150	Rome	42	Shanghai	90
Tianjin	90	Xiamen	100	**Total**	**931**

Every tile couples a geographic scene with numerical and visual radio representations. The main files are:

- `buildings.json`, containing building footprints and heights derived from OpenStreetMap (OSM);
- `sector_params.json`, recording base-station locations, sector orientations, antenna parameters, carrier settings, and transmit powers;
- `RSRP.npy` and `SINR.npy`, containing dense two-dimensional Reference Signal Received Power (RSRP, in dBm) and Signal-to-Interference-plus-Noise Ratio (SINR, in dB) arrays;
- `RSRP.png` and `SINR.png`, providing directly viewable renderings of the corresponding numerical arrays;
- `initial_scene.pdf`, visualizing the building and sector layout, and `colorbar.json`, recording the value ranges used for rendering.

This organization keeps the learning targets, geographic priors, transmitter configuration, and human-readable visualization together. It therefore supports tasks including radio-map reconstruction, coverage prediction, best-cell analysis, and geometry-conditioned signal-field generation without requiring the simulator to be rerun.

G.3 Full-Stack Network Collection

The `Full-stack-Data/` collection extends the static radio maps with synchronized network dynamics for Cardiff, Shanghai, and Tianjin. Its 222 tiles and city-level sizes are summarized in Table 20.

For each completed tile, `base_station_config.json` records the tile bounds, random seeds, sector definitions, and resource-allocation settings. The `radiomap.npz` archive stores coordinates, RSRP, SINR, the best-sector identifier, and a building mask. The remaining files expose network behavior at three granularities:

- `ue_timeseries/` stores per-UE mobility, radio measurements, traffic demand, scheduling decisions, modulation and coding scheme (MCS), physical resource block (PRB) allocation, and transport-block size (TBS);

Table 20: Coverage and size of the full-stack collection.

City	Tiles	Files	Size (GiB)
Cardiff	42	2,773	0.616
Shanghai	90	7,021	1.743
Tianjin	90	7,021	2.089
Total	222	16,815	4.448

- `cell_timeseries/` stores per-sector uplink and downlink demand together with delivered traffic;
- `ran_events/` stores access and handover events, their source and target sectors, and success indicators.

Each tile covers 84,600 seconds (23.5 hours) of simulated network operation. Records are aggregated every 60 seconds and partitioned into files spanning at most one hour, which permits partial loading of long simulations. A completion record, `_COMPLETE.json`, provides the tile status and summary statistics. Tiles labeled `complete_no_users` can legitimately omit one or more time-series directories because no UE activity was generated in that region.

The full-stack collection preserves the common tile and time references across propagation, mobility, traffic, scheduling, and RAN events. It can consequently be used for tasks that cannot be formulated from radio maps alone, including traffic forecasting, handover prediction, resource-allocation learning, and the construction of network world models.

G.4 3D City Collection

The `3D_Cities/` collection contains assembled Beijing and Guangzhou scenes in both Blender and glTF 2.0 formats. Table 21 reports the validated scene statistics. These values were obtained by re-importing the exported glTF assets into Blender 3.2.2; both scenes passed the included re-import checks.

Table 21: Released 3D city assets and validated glTF scene statistics.

City	Files	Size (GiB)	Nodes	Meshes	Vertices	Polygons
Beijing	328	2.228	179	175	4,850,613	5,905,186
Guangzhou	7	3.501	1,137	1,135	15,035,292	19,066,420
Total	335	5.729	—	—	—	—

Each scene includes a standalone `combined.blend` project and a glTF package consisting of `combined.gltf`, `combined.bin`, and any referenced external textures. The associated source and conversion manifests record input provenance, file sizes, SHA-256 hashes, conversion settings, external references, and validation results. The released conversion did not apply mesh simplification, emeshing, Draco compression, or requested texture re-encoding. These assets provide reusable geometry for visualization, geometry-conditioned learning, and independent propagation experiments.

H Potential Impacts and Limitation

Potential positive impacts. MobiSim is designed primarily as a mobile network simulation benchmark, but its underlying capabilities, like generative urban scene synthesis, physics-based electromagnetic modeling, and AI-driven user behavior generation, have implications that extend well beyond the wireless networking community. We outline several key potential impacts below.

• *Enabling full-stack and counterfactual mobile network data generation*. (i) A persistent bottleneck in mobile network research is the scarcity of multi-layer, full-stack operational data. Existing datasets

are predominantly confined to a single simulation layer, including channel measurements, protocol traces, or mobility records, and rarely capture the joint dynamics across all three. MobiSim directly addresses this gap by generating synchronized records spanning electromagnetic channel states, RAN protocol operations, and user behavior, providing a data modality that is practically infeasible to collect at scale from operational networks. (ii) MobiSim supports the generation of hard-to-collect but critical data, including counterfactual observations under hypothetical infrastructure configurations and rare anomalous events such as sudden traffic surges, interference spikes, or cascading coverage failures. Such data is fundamental to building mobile network world models and stress-testing network intelligence algorithms, yet remains almost entirely absent from existing public datasets.

• *Advancing AI-driven network optimization.* The evolution toward 6G mandates AI-native network architectures in which intelligent algorithms autonomously manage planning, configuration, and resource allocation at scale. Optimization algorithms such as RLs and LLM-based optimizers (*e.g.,* AlphaEvolve) require accurate, system-level performance feedback for effective training and evaluation. Existing simulation tools, which are typically confined to a single network layer, produce feedback signals that are too coarse and spatially inconsistent to support effective policy learning or high-fidelity algorithm evaluation. MobiSim addresses this by exposing physically grounded, full-stack reward signals that reflect the true consequences of configuration decisions across the electromagnetic, protocol, and user layers simultaneously, making it a suitable training and benchmarking environment for the next generation of network intelligence algorithms.

• *Lowering the barrier to reproducible network research.* Beyond data and optimization, MobiSim has broader implications for the reproducibility and accessibility of mobile network research. The ability to instantiate a high-fidelity simulation at any geographic location from public map data removes the dependency on proprietary datasets, specialized measurement equipment, and operator partnerships that currently limit participation in this field. Researchers without access to live network infrastructure can conduct controlled, reproducible experiments across diverse urban environments, enabling more rigorous cross-environment generalization studies and facilitating fair comparison of algorithms across standardized benchmarks.

• *Supporting smart city and urban computing research.* MobiSim provides joint simulation of human mobility, application-layer traffic, and physical infrastructure, which extends to urban computing problems beyond networking. Researchers in transportation, urban planning, and public health can leverage MobiSim's generative user behavior module to study how populations move through and interact with urban environments under varying infrastructure configurations—questions that are crucial to smart city design and emergency response planning.

Limitations. Despite its contributions, MobiSim has several limitations. First, the generative urban world model currently captures building layouts and geometric shapes but does not yet model material properties, which may limit electromagnetic simulation accuracy in material-sensitive propagation scenarios; we plan to extend the world model to jointly estimate per-surface material attributes from satellite imagery and street-level data. Second, MobiSim currently supports only 5G NR gNB protocol emulation and does not yet incorporate LTE stacks, while most live operational networks run in hybrid LTE/5G configurations; extending the protocol stack to support LTE eNB emulation would enable full heterogeneous network simulation. Third, current ray-tracing engine leverages GPU-accelerated matrix operations to achieve practical simulation speeds, but city-scale deployments with large numbers of rays and reflection paths may impose significant memory and computational overhead. We observe that the resulting channel matrices are highly sparse, as many grid points receive no meaningful signal contribution from distant base stations. Exploiting this sparsity through compressed representations or sparse matrix acceleration presents a promising direction for scaling MobiSim to metropolitan-scale simulations.