

Transformer-Based Rate Prediction for Multi-Band Cellular Handsets

Ruibin Chen^{*1}, Haozhe Lei^{*†1}, Hao Guo^{1,2},
 Marco Mezzavilla³, Hitesh Poddar⁴, Tomoki Yoshimura⁴, Sundeep Rangan¹

¹NYU WIRELESS, Tandon School of Engineering, New York University, Brooklyn, NY 11201, USA

²Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden

³Dipartimento di Elettronica, Informazione e Bioingegneria (DEIB), Politecnico di Milano, Milan, Italy

⁴Sharp Laboratories of America (SLA), Vancouver, Washington, USA

{rc5018, hl4155, hg2891, srangan}@nyu.edu; marco.mezzavilla@polimi.it; {poddarh, yoshimurat}@sharplabs.com

Abstract—Cellular wireless systems are witnessing the proliferation of frequency bands over a wide spectrum, particularly with the expansion of new bands in FR3. These bands must be supported in user equipment (UE) handsets with multiple antennas in a constrained form factor. Rapid variations in channel quality across the bands from motion and hand blockage, limited field-of-view of antennas, and hardware and power-constrained measurement sparsity pose significant challenges to reliable multi-band channel tracking. This paper formulates the problem of predicting achievable rates across multiple antenna arrays and bands with sparse historical measurements. We propose a transformer-based neural architecture that takes asynchronous rate histories as input and outputs per-array rate predictions. Evaluated on ray-traced simulations in a dense urban micro-cellular setting with FR1 and FR3 arrays, our method demonstrates superior performance over baseline predictors, enabling more informed band selection under realistic mobility and hardware constraints.

Index Terms—UE handset modeling, Multi-Band cellular, Transformer neural network, Rate prediction, Ray tracing

I. INTRODUCTION

Modern commercial handsets are increasingly required to support a growing number of frequency bands with multiple distinct antennas in a compact form factor [1]. Continuous and reliable coverage requires monitoring these bands and initiating frequent handover and serving cell decisions as the UEs move in the environment. There are several factors in emerging systems that are now making this multi-band channel tracking challenging.

Proliferation of bands with rapidly varying coverage: Most importantly, the number of frequency bands is exploding – for example, 3GPP identifies more than 100 bands in FR1 alone [2]. The number of bands is likely to accelerate with the expansion into FR3 [3]–[5]. The coverage from each of these bands can vary rapidly over space. For example, the top of Figure 1 shows the predicted capacity (or Rate) based on ray tracing at two bands – 3.5 GHz and 15 GHz in an urban micro-cellular setting. Details of the simulation are presented below. We observe that the regions where one band is better than

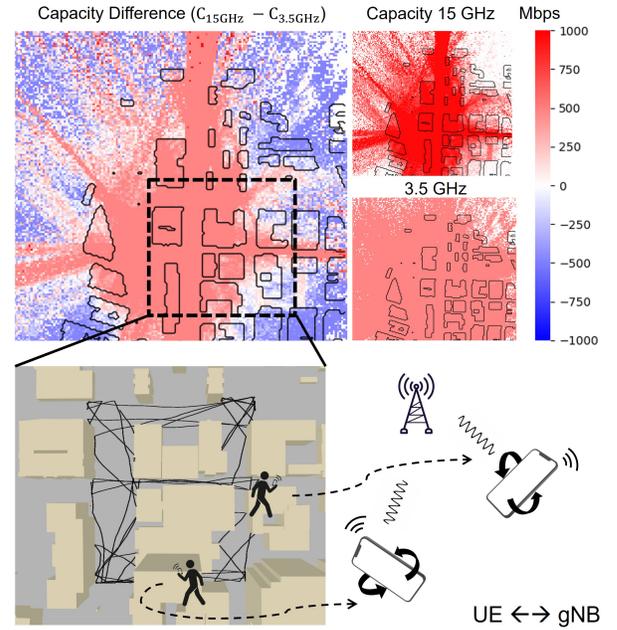


Fig. 1: Top: Ray tracing capacity in an urban micro-cell. Top-left: difference $\Delta C = C_{15} - C_{3.5}$ (red $>0 \rightarrow 15$ GHz, blue $<0 \rightarrow 3.5$ GHz, Mbps). Top-right: per-band capacity at 15 GHz and 3.5 GHz (shared color bar, Mbps). Bottom: zoom of the dashed area with the site model and simulated pedestrian trajectories with handset rotations, for UE mobility and UE-gNB interaction, yielding time-varying band/array rates.

the other vary rapidly due to complex propagation differences between the bands. This fact necessitates that UEs will need to switch bands rapidly to maintain optimal throughput.

UE antenna constraints: Many bands may require distinct antennas, yet the available real estate within handsets is severely limited [1]. Moreover, individual antennas often exhibit narrow steerable ranges and single polarization, implying that a coverage for any given band will have a limited field of view. Coverage can also be limited by hand blocking [6], [7]. As a result, natural changes in the orientation of the phone

* These authors contributed equally to this work.

† Corresponding author: Haozhe Lei (hl4155@nyu.edu).

will necessitate further rapid band switching.

Limited measurements: To reduce cost and energy consumption, multiple antennas may be switched to a single RF chain. In addition, it may be too costly for the power to operate and monitor multiple bands simultaneously. Consequently, UEs are forced to make band selection decisions on limited measurements.

Problem: Motivated by these developments, we formulate a general problem of tracking and predicting channel rates from multiple antenna arrays. Different arrays may be associated with different frequency bands. For each array, the UE has limited prior measurements of the available rate since the array may have been switched off to save power, or the RF chain for that array may have been switched to a different array. The problem is to estimate the achievable rate on each array and to predict which array offers the highest rate.

Our contributions: For this problem, we propose a novel transformer neural network architecture that takes as inputs the past measured channel quality indicator (CQI) on each link, a link being a potential connection from a serving cell to a particular UE array in some carrier frequency. The measured CQIs may have been at arbitrary time points in the past. This multi-dimensional CQI history is fed to a temporal convolutional encoder followed by a single transformer layer.

The proposed method is simulated on a realistic antenna layout with two antennas in FR1 (3.5 GHz) and two antenna arrays in FR3 (15 GHz). Ray tracing simulation is used to estimate rates across the antenna arrays with both UE orientation and translational motion in a dense urban setting. It is shown that the method outperforms simple baselines.

Related Work

Measurement-driven studies with form-factor UEs at 28 GHz show that hand/body occlusion and device rotation dominate coverage and beam management, challenging simplified blockage models [6], [7]. Multi-band campaigns (11/16/28/32 GHz) quantify ~ 20 dB human-blockage loss and validate KED/UTD abstractions [8], [9], while FR1/FR3 results at 6.75/16.95 GHz provide calibrated outdoor models and indoor-hotspot angular-spread statistics [10], [11].

EM/RT pipelines benchmark KED/UTD/PO against 60 GHz measurements to expose accuracy–complexity trade-offs [12]; device-as-scatterer CAD reveals panel-dependent pattern loss and reflection beyond fixed-angle masks [13]. Site-specific calibration/validation of NYURay at upper mid-band tightens agreement with measurements, and NYUSIM offers spatially consistent, blockage-aware corpora for synthesis and analysis [14], [15].

Predictors enable proactive switching and lightweight channel inference: AE+LSTM sequence models improve Top- C ranking and reduce misalignment for next-slot beam/BS prediction [16]; cross-band learning maps sub-6 GHz features to mmWave beam/blockage; vision/semantics aid LoS/NLoS forecasting and candidate selection [17]–[19]. Physics-informed Reinforcement Learning (RL) further em-

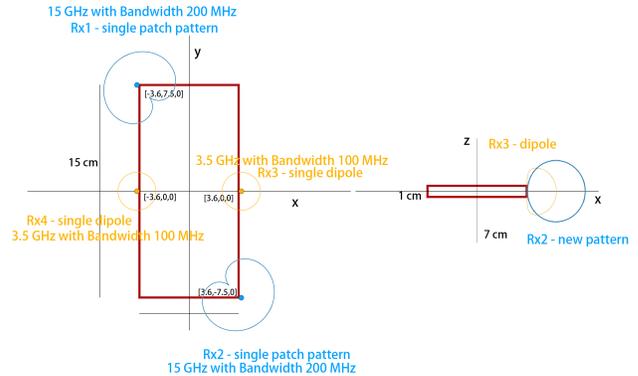


Fig. 2: Illustration of UE radian pattern. The red rectangular shows the UE body shape as 7cm \times 15cm \times 1cm. New directive 15GHz FR3 radian pattern is shown in blue while the traditional 3.5GHz dipole pattern is shown as in yellow. The figure demonstrates the horizontal cut of the UE radian pattern in left with vertical cut in right.

beds channel dynamics, demonstrating zero-shot and DT-in-the-loop gains with growing industrial traction [20]–[22].

II. FREQUENCY HOPPING PROBLEM FORMULATION

Consider a UE that can be served in the downlink from one of N_{BS} base station (gNB) cells with reception on one of N_{arr} antenna arrays. We define a *link* as a pair (i, j) from gNB cell i to antenna array j , and we assume there are M potential links. Each link $i = 1, \dots, M$ is associated with some carrier frequency $f^{(i)}$. Two different links may use the same frequency or the same serving cell. Also, depending on the bandwidth of the antenna, two different serving cells could be received in the same antenna array.

As one simple example that we will use in the simulations below, we consider a UE being served by $N_{BS} = 2$ potential cells that are co-located, so they are from a single gNB site. The two cells operate in two distinct frequencies: BS1 at 3.5 GHz in FR1 and BS2 15 GHz in FR3 (details in Table II.). The UE antenna layout that we will use in the simulations below is Figure 2. There are $N_{arr} = 4$ antennas – two dipole antennas at 3.5 GHz and two patch antennas at 15 GHz. Four antennas are labeled as RX1 to RX4. In this case, there are $M = 4$ potential links:

- Link 1: From BS1 to RX3 at $f^{(1)} = 3.5$ GHz
- Link 2: From BS1 to RX4 at $f^{(2)} = 3.5$ GHz
- Link 3: From BS2 to RX1 at $f^{(3)} = 15$ GHz
- Link 4: From BS2 to RX2 at $f^{(4)} = 15$ GHz

Each antenna has a directivity pattern, so the link quality from the same serving cell may be different for two different antennas. For example, the link quality on link 1 may be different than link 2 depending on the orientation of the different antennas relative to the angles of arrivals of the paths. In Figure 2, the antenna placement, orientation, and antenna pattern follow the 3GPP simulation model [23]. In

the simulations below, we will also consider blockage from the metal UE body itself that influences coverage depending on orientation.

Our broad problem is to track the link quality on the M links. Time is divided into measurement periods of length T . In the simulations below, we take $T = 50$ ms, which is sufficient for several measurements of the channel quality in a given antenna. We let $\gamma_t^{(i)}$ denote the CQI on link i in time period t if the UE were to receive downlink transmissions in that link in that time period. For simplicity, we will take $\gamma_t^{(i)}$ as the average SNR on the link. We assume that the rate of a practical model of spectral efficiency ρ is applied according to [24] [25], modifying from the 3GPP Shannon capacity [26].

$$R_t^{(i)} = \phi_i(\gamma_t^{(i)}) = B_i \min\{\rho_{\max}, \beta \log_2(1 + \gamma_t^{(i)})\}, \quad (1)$$

where $\phi(\gamma)$ is the CQI-rate mapping function; B_i is the bandwidth of link i ; β is the coding efficiency with a typical value of 0.6; ρ_{\max} is the practical channel capacity subject to the 3GPP MCS table, which is commonly selected from 4.8 - 7.5 based on different QAM strategies. The model below can be extended to more complex CQI descriptions and multi-stream transmissions.

As discussed in the Introduction, a critical challenge is that the CQI across all links may not be observable. First, two or more antennae may share a common RF receiver chain via a switch implying that only one of those antennas can be used. Second, even if the two antennas can be activated at the same time, it may not be desirable to turn both of them on to save energy. To model this effect, we let $A_t \subseteq \{1, \dots, M\}$ denote the set of link indices that were active in period t . The link being active requires that the antenna for that link is switched to an active RF chain and that RF chain is tuned to the appropriate frequency of that link. We assume that the CQI can only be measured on active links. We describe the observations as:

$$o_t^{(i)} = \begin{cases} \widehat{\gamma}_t^{(i)}, & i \in A_t, \\ \perp, & i \notin A_t, \end{cases} \quad (2)$$

where $\widehat{\gamma}_t^{(i)}$ denotes an estimate of the CQI $\gamma_t^{(i)}$ and \perp denotes “no measurement” since the link was not active. We let $\mathbf{o}_t = (o_t^{(1)}, \dots, o_t^{(M)})$ denote the vector of observations across links. Our goal at each time t is to estimate the rates $R_t^{(i)}$ for all links $i = 1, \dots, M$ based on the past observations \mathbf{o}_s for $s < t$. Note that since we are using observations for $s < t$, our model is predicting the rates one step in the future.

Although not simulated in this work, such rate predictions can be used to guide selection of measurements, cell selection, and handover decisions.

III. NEURAL NETWORK PREDICTOR

The predictor in Figure 3 follows a physically grounded pipeline. At the input, we first transform each observation o_t^i in (2) to a 3-dimensional vector:

$$\mathbf{y}_t^{(i)} = \begin{cases} (1, 10 \log_{10}(\widehat{\gamma}_t^{(i)}), \phi(\widehat{\gamma}_t^{(i)})) & i \in A_t \\ (0, 0, 0) & i \notin A_t \end{cases} \quad (3)$$

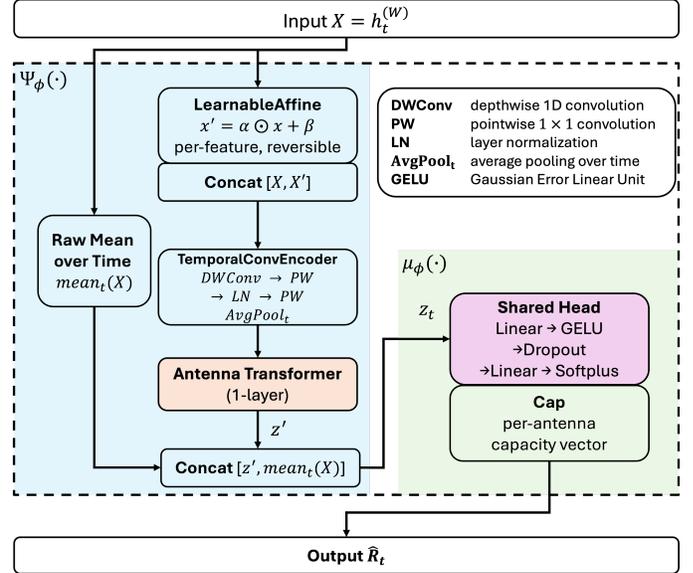


Fig. 3: Overall rate estimation network. Each antenna’s recent history is passed through a learnable per-feature affine scaling that preserves units for rate and SNR while improving numerical conditioning. The scaled and raw features are concatenated. A lightweight temporal encoder summarizes each antenna’s W -sample window into one embedding. An antenna-level Transformer with one encoder layer mixes cross-antenna dependencies. The result is concatenated with the raw mean over time, which averages only along the time dimension W and keeps the original antenna count N and feature dimension F unchanged. A shared head with Softplus produces non-negative per-antenna rate estimates. A per-antenna capacity cap conditioned on bandwidth and hardware is applied.

where the first component is 1 or 0 indicating if the a measurement on link i was made or not (that is, if $i \in A_t$ or $i \notin A_t$). If a measurement was made, the second component is the SNR in dB scale and the third component is the rate corresponding to the estimated SNR where $\phi(\gamma)$ is the CQI-rate mapping function in (1). If no measurement was made, the two other components are set to zero. We then compute a history of the last W such measurements:

$$\mathbf{h}_t^{(w)} = \{\mathbf{y}_t^{(i)}, i = 1, \dots, M, s \in [t - W, t]\}. \quad (4)$$

Since each component $\mathbf{y}_t^{(i)}$ is 3-dimensional, this history window has length $3MW$ where M is the total number of links and W is the window size.

The W -step history is first processed by a learnable per-feature affine scaling that preserves units for rate and SNR. We keep both the raw and the scaled versions to retain absolute levels and to ease optimization, then a lightweight temporal encoder compresses the short window into one embedding per antenna. Before prediction we append the raw mean over time computed by averaging only along the time dimension W . This operation removes the window axis yet keeps the original antenna count M and feature dimension F unchanged, which

preserves absolute scale information useful for calibration. A shared head with Softplus enforces non-negativity and a per-antenna capacity cap reflects bandwidth and hardware limits.

Cross-antenna interaction is modeled on the antenna index rather than on time. We use a one-layer antenna-level Transformer, namely a single encoder block with pre-layer normalization, multi-head self-attention, and a position-wise feed-forward network with residual connections. This mixes dependencies among antennas while keeping capacity controlled for moderate $|\mathcal{I}_{\text{TX}}|$ and limited data. No positional encoding along antennas is required because antenna identities are fixed and are implicitly captured by the learned embeddings.

IV. NUMERICAL EXPERIMENTS

A. Experimental Set-up

The overall procedures for offline dataset generation can be divided into the ray tracing and Data-preparing parts.

TABLE I: UE mobility pattern

Component	Parameter [unit]	Low	Medium	High
Pedestrian	Max velocity [m/s]	2	5	10
	Heading rate [rad/s]	$\pm\pi/2$	$\pm\pi/3$	$\pm\pi/6$
Handset rotations	Yaw increment [rad/50 ms]	$\pm\pi/10$	$\pm 2\pi/10$	$\pm 3\pi/10$
	Pitch increment [rad/50 ms]	$\pm\pi/20$	$\pm 2\pi/20$	$\pm 3\pi/20$

Note: Uniform sampling within the listed intervals; “heading rate” is pedestrian orientation change.

Ray tracing: The generator first applies the UE mobility pattern model list in Table I to update the pedestrian trajectory by sampling the movement velocity, acceleration, and heading rate. The UE is navigated within the NYU Tandon campus using a three-level pedestrian mobility model. Each trajectory consists of 1200 time steps, corresponding to 60 s with a step size of 50 ms. The generator uses the generated routes as the UE center positions, and rotates the handset around the body center. This rotation is applied to adjust the position and orientation of each RX relative to the body coordinate frame.

Sionna [27], an open source software that provides powerful ray tracing results based on the open street map, can then simulate the complex downlink channel gain with antenna directivity for multi-path communication between the gNB with different RXs on UE. The capacity and SNR are calculated based on these channels. Table II lists the ray tracing simulation parameters.

The experiment is conducted in the boundary area of a 15 GHz serving cell, where the channel rates of the two frequency bands frequently exhibit rank reversals. In addition, the self-rotation of the user causes the rates of the RXs operating on the same frequency to alternately become dominant.

For each mobility level, 20 random routes are generated. Each route includes 50 handset rotations, resulting in 1000 instances in total. Each route lasts for 60 seconds, corresponding to 1200 steps (50 ms/step) sampling steps with a step size of 50 ms. Example ray tracing data for each RX is shown in Figure 4.

TABLE II: Simulation Parameters

Parameter	Value
Carrier frequency (Bandwidth)	15 GHz (200MHz)
	3.5 GHz (100MHz)
Tx antenna pattern	Tr38901 V polarization
Transmitter Power	51 dBm
Noise Power Spectral Density	-174 dBm/Hz
Noise Figure	7 dB
Measure Period	50 ms
Code Efficiency β	0.6
Max Spectral Efficiency ρ_{max}	4.8 bits/Hz
Sionna ray types	specular, diffuse, refraction
Sionna max interaction depth	5 interactions

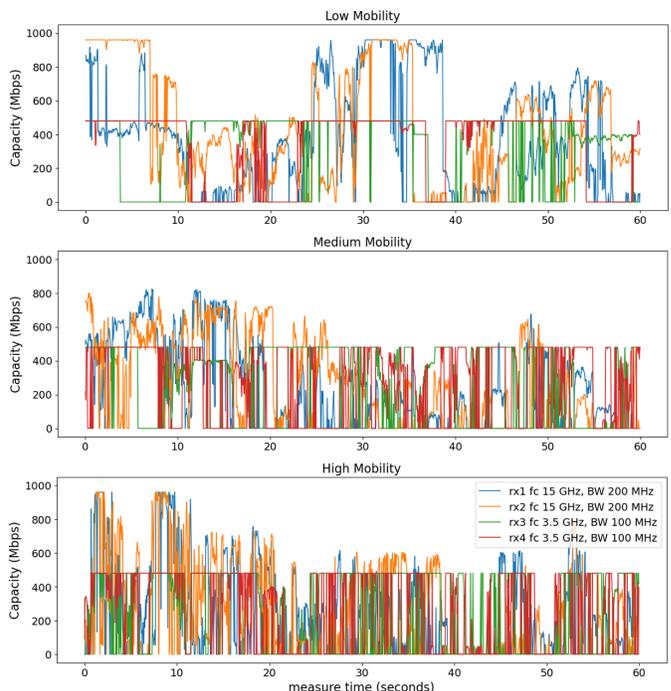


Fig. 4: Channel Capacity at three-level mobility. As the mobility level increases, the capacity at each RX exhibits more frequent variations. This increased temporal variability makes it challenging to maintain high transmission rates using only instantaneous feedback. Therefore, accurate rate prediction becomes essential to anticipate channel changes and enable timely adaptation of transmission strategies.

Data-preparing: Real systems face energy/thermal limits, so not all antennas can be active. To explicitly simulate this missing-measurement regime, we use an adversarial multi-armed bandit policy π_{adv} that defines a probability distribution over subsets,

$$A_t \sim \pi_{\text{adv}}(\cdot | h_t^{(W)}),$$

For logging and a baseline, we use an exponential-weights forecaster (EXP3/Hedge) with ε -greedy exploration [28]. After



Fig. 5: Time-series performance under medium mobility. Top: comparison of the maximum predicted capacity of the proposed model, the masked-previous benchmark, and the maximum ground-truth capacity; the proposed model tracks rapid swings with less lag and fewer overshoots. Bottom: per-antenna MSE over time (gray shading marks timestamps with measurements); the proposed model (orange) is consistently below the masked-previous baseline (blue), indicating lower error both with and without measurements. Overall, per-antenna MSE is reduced by 16–38% (about 28% on average).

collecting ray tracing data, we run an ε -greedy exponential-weights bandit ($\varepsilon = 0.2$) to choose at each step a hopping observation subset of RX antennas based on historical scores. The chosen subset defines a mask: the UE keeps the corresponding ray-traced capacities and treats the rest as missing. We log the masked observations and the chosen subsets, and use the full capacities as labels to train the rate-only predictor. The generated dataset, including 1000 input instances, is shuffled and divided into 70%:15%:15% for training, validation, and testing. The whole project can be found in [GitHub repository](#).

B. Model training

The transformer-based predictor in Figure 3 starts at $t = W$ and is trained on the remaining steps using the Adam optimizer with a learning rate of 10^{-3} , a dropout rate of 0.2, and a batch size of 128 for 200 epochs.

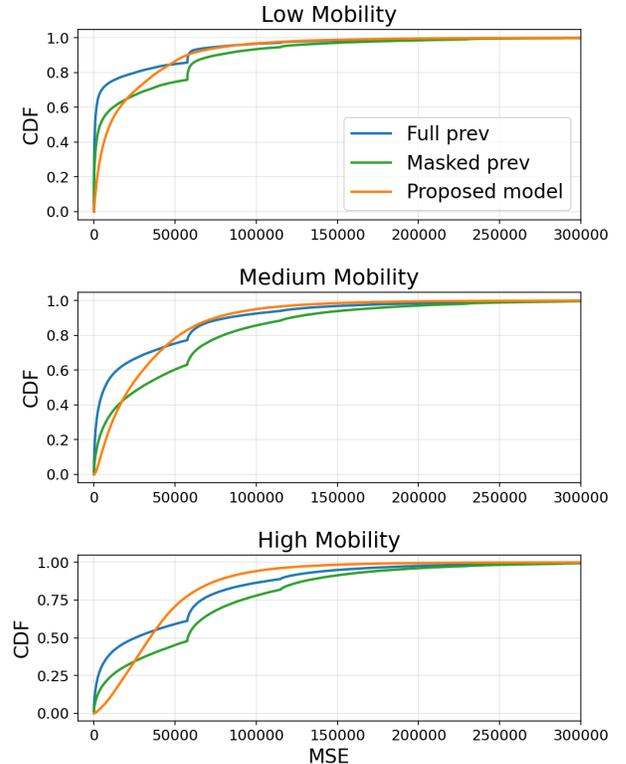


Fig. 6: The CDF plots compare the proposed model with the full- and masked-previous baselines under different mobility levels. Although the proposed model has a slightly lower probability of very small errors, it effectively suppresses large errors, leading to a tighter error distribution. This effect becomes more evident as mobility increases, demonstrating improved robustness under rapid channel variations.

C. Benchmarks

To evaluate the contribution of the learned predictor, we compare against two simple, predictor-free baselines that rely only on past observations.

B1: Full previous observation. Always activate all antennas; predict by a one-step lag of the observed per-antenna rates.

$$A_t \equiv \mathcal{I}_{\text{rx}}, \quad \hat{\mathbf{R}}_0 := \mathbf{R}_0, \quad \hat{\mathbf{R}}_t := \mu^{\text{full}}(o_{t-1}) := \mathbf{R}_{t-1}, \quad t \geq 1. \quad (5)$$

B2: Bandit-masked previous rate. Actions follow the adversarial bandit; an antenna’s prediction is refreshed only if it was active at the previous step, otherwise it is carried forward. After recursion, the first W steps are dropped to align with the predictor’s window.

$$A_t \sim \pi_{\text{adv}}(\cdot), \quad \hat{\mathbf{R}}_0 := \mathbf{R}_0. \quad (6)$$

$$\begin{aligned} \hat{\mathbf{R}}_t &:= \mu^{\text{mask}}(\hat{\mathbf{R}}_{t-1}, \mathbf{R}_{t-1}, \mathbf{m}_{t-1}) \\ &= \mathbf{m}_{t-1} \odot \mathbf{R}_{t-1} + (\mathbf{1} - \mathbf{m}_{t-1}) \odot \hat{\mathbf{R}}_{t-1}, \quad t \geq 1. \end{aligned} \quad (7)$$

D. Experiment Results

Figure 5 reports time-domain errors for a medium-mobility run. The instantaneous Mean Square Error (MSE) of the proposed predictor (orange) stays below the masked-previous baseline (blue) across most timestamps, yielding lower per-antenna averages with an overall reduction of $\approx 28\%$. Gray shading marks measurement times; the gap persists in every antenna, indicating that the predictor exploits temporal structure rather than copying the last observation.

To aggregate performance across runs, Figure 6 presents the Cumulative Distribution Function (CDF) of the MSE over all test data (150 records, 1180 samples each) under low, medium, and high mobility. Although the probability of very small errors is slightly lower, the proposed model markedly suppresses large-error events, producing a left-shifted CDF in the moderate-to-high error range. The advantage strengthens with mobility, demonstrating robustness to rapid channel variation and safer tails for rate selection.

V. CONCLUSIONS

We formulated a channel tracking problem across multiple bands for a realistic handset with multiple antennas and switching constraints. The tracker uses a transformer-based neural network and, using past measurements, outperforms simple baselines. In this paper, the set of selected antennas is specified. A future step is to consider the *feedback* nature of the problem, where rate predictions are used to select measurements. This setting presents a challenging exploration–exploitation tradeoff and can potentially be explored with reinforcement learning methods.

REFERENCES

- [1] Y. Wang, L. Sun, Z. Du, and Z. Zhang, “Review antenna design for modern mobile phones: A review,” *Electromagnetic Science*, vol. 2, no. 2, pp. 1–36, 2024.
- [2] 3GPP, “TSG RAN NR; User Equipment (UE) radio transmission and reception; Part 1: Range 1 Standalone (Release 18),” 3GPP, Tech. Rep. TS 38.101-1, June 2024, version 18.3.0. [Online]. Available: <https://www.3gpp.org/DynaReport/38101-1.htm>
- [3] S. Kang, M. Mezzavilla, S. Rangan, A. Madanayake, S. B. Venkatakrishnan, G. Hellboug, M. Ghosh, H. Rahmani, and A. Dhananjay, “Cellular wireless networks in the upper mid-band,” *IEEE Open Journal of the Communications Society*, vol. 5, pp. 2058–2075, 2024.
- [4] P. Testolina, M. Polese, and T. Melodia, “Sharing spectrum and services in the 7–24 ghz upper midband,” *IEEE Communications Magazine*, vol. 62, no. 8, pp. 170–177, 2024.
- [5] D. Shakya, M. Ying, T. S. Rappaport, H. Poddar, P. Ma, Y. Wang, and I. Al-Wazani, “Propagation measurements and channel models in Indoor Environment at 6.75 GHz FR1 (C) and 16.95 GHz FR3 Upper-mid band Spectrum for 5G and 6G,” in *GLOBECOM 2024-2024 IEEE Global Communications Conference*. IEEE, 2024, pp. 998–1003.
- [6] V. Raghavan, M. Kovacevic, R. Ratasuk, A. Ghosh, M. Rybalko, S. Subramanian, S. Hong, A. Tarighat, and Z. Pi, “Spatio-temporal impact of hand and body blockage for millimeter-wave user equipment design at 28 GHz,” *IEEE Commun. Mag.*, vol. 56, no. 12, pp. 46–52, 2018.
- [7] V. Raghavan, S. Subramanian, A. Tarighat, Z. Pi, and A. Ghosh, “Hand and body blockage measurements with form-factor user equipment at 28 GHz,” *IEEE Trans. Antennas Propag.*, vol. 70, no. 1, pp. 606–619, 2022.
- [8] W. Qi, J. Huang, J. Sun, Y. Tan, C. Wang, and X. Ge, “Measurements and modeling of human blockage effects for multiple millimeter wave bands,” in *Proc. IEEE IWCNC*, 2017, pp. 1604–1609.
- [9] U. T. Virk and K. Haneda, “Modeling human blockage at 5G millimeter-wave frequencies,” *IEEE Trans. Antennas Propag.*, vol. 68, no. 3, pp. 2256–2266, 2020.
- [10] D. Shakya, M. Ying, T. S. Rappaport, P. Ma, I. Al-Wazani, Y. Wu, Y. Wang *et al.*, “Urban outdoor propagation measurements and channel models at 6.75 GHz FR1 (c) and 16.95 GHz FR3 upper mid-band spectrum for 5G and 6G,” 2024, [Online]. Available: <https://arxiv.org/abs/2410.17539>.
- [11] M. Ying, P. Ma, D. Shakya, and T. S. Rappaport, “Angular spread statistics for 6.75 GHz FR1 (c) and 16.95 GHz FR3 mid-band frequencies in an indoor hotspot environment,” 2024, manuscript, work on angular spread statistics in FR1/FR3 indoor hotspot channels.
- [12] S. Mukherjee, G. Skidmore, T. Chawla, A. Bhardwaj, C. Gentile, and J. Šenić, “Scalable modeling of human blockage at millimeter-wave: A comparative analysis of knife-edge diffraction, the uniform theory of diffraction, and physical optics against 60 GHz channel measurements,” *IEEE Access*, vol. 10, pp. 121 720–121 731, 2022.
- [13] F. Fernandes, M. Duits, J. van den Biggelaar, H. Wang, M. Timmers, A. Nur, and G.-J. Janssen, “Hand blockage impact on 5G mmwave beam management performance,” *IEEE Access*, vol. 10, pp. 106 032–106 047, 2022.
- [14] M. Ying, P. Ma, D. Shakya, and T. S. Rappaport, “Site-specific location calibration and validation of ray-tracing simulator NYURay at upper mid-band frequencies,” 2024, preprint, validation of NYURay via site-specific location calibration.
- [15] S. Ju, O. Kanhere, Y. Xing, and T. S. Rappaport, “A millimeter-wave channel simulator NYUSIM with spatial consistency and human blockage,” 2019, [Online]. Available: <https://arxiv.org/abs/1908.09762>.
- [16] S. H. A. Shah and S. Rangan, “Multi-cell multi-beam prediction using auto-encoder LSTM for mmwave systems,” *IEEE Open J. Commun. Soc.*, vol. 3, pp. 2194–2209, 2022.
- [17] M. Alrabeiah and A. Alkhateeb, “Deep learning for mmwave beam and blockage prediction using sub-6 GHz channels,” *IEEE Trans. Commun.*, vol. 68, no. 9, pp. 5504–5518, 2020.
- [18] W. Xu, F. Gao, X. Tao, J. Zhang, and A. Alkhateeb, “Computer vision aided mmwave beam alignment in V2X communications,” *IEEE Trans. Wireless Commun.*, vol. 22, no. 4, pp. 2699–2714, 2023.
- [19] Y. Yang, F. Gao, X. Tao, G. Liu, and C. Pan, “Environment semantics aided wireless communications: A case study of mmwave beam prediction and blockage prediction,” 2023, [Online]. Available: <https://arxiv.org/abs/2301.05837>.
- [20] M. Yin, T. Li, H. Lei, Y. Hu, S. Rangan, and Q. Zhu, “Zero-shot wireless indoor navigation through physics-informed reinforcement learning,” in *Proc. 2024 IEEE Int. Conf. Robot. Autom. (ICRA)*, 2024, pp. 5111–5118.
- [21] T. Li, H. Lei, H. Guo, M. Yin, Y. Hu, Q. Zhu, and S. Rangan, “Digital twin-enhanced wireless indoor navigation: Achieving efficient environment sensing with zero-shot reinforcement learning,” *IEEE Open J. Commun. Soc.*, vol. 6, pp. 2356–2372, 2025.
- [22] T. Li, H. Lei, M. Yin, and Y. Hu, “Reinforcement learning with physics-informed symbolic program priors for zero-shot wireless indoor navigation,” in *Proc. Reinforcement Learning Conf. (RLC)*, 2025.
- [23] 3GPP, “TSG RAN WG1#121; Introduction of Rel-19 7-24 GHz channel model enhancements,” 3GPP, Tech. Rep. R1-2504964, May 2025, tR38901, CR0026 version 18.0.0. [Online]. Available: https://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_121/Inbox/
- [24] Y. Hu, M. Yin, S. Rangan, and M. Mezzavilla, “Parametrization and estimation of high-rank line-of-sight MIMO channels with reflected paths,” vol. 23, no. 4, pp. 3808–3822, 2024.
- [25] P. Mogensen, W. Na, I. Z. Kovács, F. Frederiksen, A. Pokhariyal, K. I. Pedersen, T. Kolding, K. Hugl, and M. Kuusela, “LTE capacity compared to the shannon bound,” in *Proc. IEEE Veh. Technol. Conf. (VTC Spring)*, 2007, pp. 1234–1238.
- [26] 3GPP, “3GPP TR 38.901 V16.1.0: Study on channel model for frequencies from 0.5 to 100 GHz,” 3rd Generation Partnership Project (3GPP), Tech. Rep. TR 38.901, Dec. 2020, release 16. [Online]. Available: <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3173>
- [27] J. Hoydis, S. Cammerer, F. Ait Aoudia, M. Nimier-David, L. Maggi, G. Marcus, A. Vem, and A. Keller. (2022) Sionna. Ver. 1.1.0, [Online]. Available: <https://nvlabs.github.io/sionna/>.
- [28] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, “The non-stochastic multiarmed bandit problem,” *SIAM J. Comput.*, vol. 32, no. 1, pp. 48–77, 2002.