

Data-Driven Design of 3GPP Handover Parameters with Bayesian Optimization and Transfer Learning

Mohamed Benzaghta^{*}, Sahar Ammar^b, David López-Pérez[#], Basem Shihada^b, and Giovanni Geraci^{†*}

^{*}Univ. Pompeu Fabra, Barcelona, Spain [†]Telefónica Scientific Research, Spain

^bKing Abdullah Univ. of Science and Technology, Saudi Arabia [#]Univ. Politècnica de València, Spain

Abstract—Mobility management in dense cellular networks is challenging due to varying user speeds and deployment conditions. Traditional 3GPP handover (HO) schemes, relying on fixed A3-offset and time-to-trigger (TTT) parameters, struggle to balance radio link failures (RLFs) and ping-pongs. We propose a data-driven HO optimization framework based on high-dimensional Bayesian optimization (HD-BO) and enhanced with transfer learning to reduce training time and improve generalization across different user speeds. Evaluations on a real-world deployment show that HD-BO outperforms 3GPP set-1 and set-5 benchmarks, while transfer learning enables rapid adaptation without loss in performance. This highlights the potential of data-driven, site-specific mobility management in large-scale networks.

I. INTRODUCTION

To meet rising mobile data demands and support new applications, operators are densifying networks to improve coverage, spectral efficiency, and spatial reuse [1]. However, this densification increases handover (HO) frequency, complicating mobility management. HOs allow user equipment (UE) to maintain quality of service (QoS) while switching between serving and target cells, but frequent transitions pose significant challenges [2].

Given the critical role of mobility management, it has been a focal point for research, industry, and standardization bodies. In [2], the authors propose an HO algorithm designed for small-cell networks that incorporates UE mobility direction to mitigate frequent handovers. Their approach analyzes the reference signal received power (RSRP) pattern during the time-to-trigger (TTT) to determine whether a UE is approaching or moving away from a target cell. Similarly, [3] introduces a federated learning-based handover algorithm that dynamically adjusts HO thresholds based on predicted signal strength and historical performance. The work in [4] presents an online learning-based mobility management framework that estimates posterior probabilities of

RSRP values to optimize target cell selection. Analytical models have also been explored, such as the mathematical framework in [5], which derives UE-to-base station (BS) association probabilities and HO rates for hybrid radio frequency–visible light communications. Likewise, [6] studies HO triggering estimation based on speed, distance, and predefined mobility models.

Mobility management efficiency depends heavily on the configuration of A3-offset and TTT, which are critical for minimizing radio link failures (RLFs) and ping-pongs. Handover parameters are typically tuned per cell or within small clusters, but this localized approach often overlooks inter-cell dependencies and fails to deliver optimal network-wide performance. While joint optimization could address this, it quickly becomes intractable at scale. Moreover, parameters tuned for specific mobility profiles may not generalize—for instance, high-speed UEs may experience RLFs if TTT delays handover beyond acceptable SINR levels [7]. These challenges highlight the need for cell-specific, data-driven optimization that adapts to local traffic and mobility patterns.

To overcome these challenges, we propose a high-dimensional Bayesian optimization (HD-BO) framework for scalable, robust mobility management across diverse network deployments and UE speeds. To improve convergence and align with 3GPP’s vision for model generalization [8], we also explore transfer learning across varying mobility conditions.

II. SYSTEM MODEL AND PROBLEM FORMULATION

A. Network Topology and Site-Specific Propagation Modeling

We consider a real-world radio network deployment operated by a leading commercial mobile provider in the UK.

Network deployment: The cellular network considered in this study comprises 10 deployment sites, with antenna heights ranging from 22 to 56 meters. Each site hosts three sector antennas, resulting in a total of 30 cells. The deployment covers an urban area of approximately 1400×1275 m in London, bounded by latitudes $[51.5087, 51.5215]$ and longitudes $[-0.1483, -0.1296]$. A 3D model of the area is generated using OpenStreetMap data, incorporating both terrain elevation and building structures. Base

This work was supported by *a)* HORIZON-SESAR-2023-DES-ER-02 project ANTENNAE (101167288), *b)* the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union NextGenerationEU through actions CNS2023-145384 and CNS2023-144333, *c)* the Spanish State Research Agency through grants PID2021-123999OB-I00 and CEX2021-001195-M, *d)* the Generalitat Valenciana, Spain, through the CIDEAGENT PlaGenT, Grant CIDEXG/2022/17, Project iTENTE, and *e)* the UPF-Fractus Chair.

stations (BSs) are placed and configured to reflect the actual layout and topology of the real-world cellular network.

Channel propagation modeling: The channel between BS b and UE k is modeled using Sionna RT [9], a widely used 3D ray-tracing tool for site-specific radio wave propagation analysis. Simulations are performed at a carrier frequency of 2 GHz over a 10 MHz bandwidth with 50 physical resource blocks (PRBs). Buildings are modeled using the `itu_concrete` material to represent their permittivity and conductivity. Ray-tracing is configured to allow up to 5 reflections and 1 diffraction, capturing key propagation effects in the urban environment.

SINR formulation: The downlink SINR in dB experienced by UE k from its serving BS b_k is computed as:

$$\text{SINR}_{\text{dB},k} = 10 \log_{10} \left(\frac{p_{b_k} \cdot G_{b_k,k}}{\sum_{b \in \mathcal{B} \setminus b_k} p_b \cdot G_{b,k} + \sigma_T^2} \right), \quad (1)$$

where $G_{b,k}$ represents the square magnitude of the channel gain, capturing both small-scale and large-scale fading, averaged over the 50 PRBs (each with a bandwidth of 180 kHz). The thermal noise power σ_T^2 over 10 MHz is obtained from a power spectral density of -174 dBm/Hz , while the transmit power of BS b across the entire bandwidth is $p_b = 46 \text{ dBm}$ [10].

Challenges in cellular mobility management: We select five main streets in a central London urban area as the focus of our mobility study. Fig. 1 presents a 2D map where circled triangles mark BS sites and colored dots represent UE locations along the selected streets, which are highlighted in black and labeled by ID. The dot colors indicate the serving cell site with the highest average received power, as shown in the background heatmap. This spatial distribution illustrates the complexity of real-world mobility management, where rapid signal variations may lead to frequent handovers—highlighting the need for a robust, data-driven framework capable of supporting large-scale deployments and ensuring seamless connectivity.

B. Handover Procedures in Cellular Networks

In the following, we provide a detailed overview of the cellular mobility management problem and introduce key concepts, including HO, RLF, and ping-pongs.

Handover process: In cellular networks, handovers (HOs) can occur between cells of different radio access technologies (RATs), between cells of the same RAT operating on different frequencies, or between cells of the same RAT and frequency [11]. This study focuses on intra-RAT, intra-carrier HOs and specifically on hard HOs, where the UE disconnects from the source cell before establishing a connection with the target cell.

The UE performs HO measurements at both the physical (L1) and network (L3) layers. RSRP values are estimated for all neighboring cells, collected every 40 ms and averaged over five




Fig. 1: 2D view of the selected urban area showing UE positions (colored dots) and BS sites (triangles). Black lines indicate the five selected streets; colors represent the cell providing the strongest received power.

intervals, a process known as *L1 filtering*. The resulting L1-filtered measurements are further smoothed via a first-order IIR filter at L3—referred to as *L3 filtering*—with a typical period of 200 ms. This two-stage filtering reduces the impact of fading and measurement noise [12]. For a visual illustration of the L1 and L3 filtering processes, we refer the reader to Fig. 1 in [12].

A HO is triggered when the L3-filtered measurement satisfies a predefined event condition. This study focuses on *Event A3*, which governs intra-RAT, intra-carrier HOs [11], [12]. Event A3 is met when the target cell's L3-filtered RSRP exceeds that of the serving cell by a hysteresis margin (A3 offset). Once triggered, the UE starts a Time-to-Trigger (TTT) timer. If the condition holds throughout the TTT, the UE sends a measurement report to the serving cell, initiating HO preparation.

Radio link failure: Radio link failure (RLF) is triggered when a UE experiences sustained signal degradation. It is governed by the SINR thresholds Q_{out} and Q_{in} , along with the T_{310} timer. If the downlink SINR ($\text{SINR}_{\text{dB},k}$) falls below Q_{out} , a synchronization issue is detected, activating T_{310} . If SINR recovers above Q_{in} before the timer expires, the UE is considered re-synchronized. Otherwise, expiration of T_{310} results in RLF [11].

Handover ping-pongs: A HO ping-pong occurs when a UE briefly connects to a cell before handing over again. This *time-of-stay* starts with the HO complete message (i.e., indication that the UE has successfully completed the handover) to the target cell and ends with a subsequent HO complete message. A ping-pong is recorded if the time-of-stay is below a threshold T_p (e.g., 1 s) and the new target cell matches the original source cell. These unnecessary HOs increase signaling overhead and reduce network efficiency.

Performance trade-off: Small A3-offset and TTT values may trigger premature handovers, increasing ping-pongs, while larger values can delay HOs and raise the risk of RLF. Effective mobility management requires tuning these parameters based on UE velocity and site-specific radio conditions. In 3GPP networks, HO decisions are governed by predefined threshold sets. We focus on benchmark configurations set-1 and set-5, which represent the two extremes in HO performance trade-offs. Set-1 delays HOs to reduce ping-pongs (TTT = 480 ms, A3-offset = 3 dB), while set-5 accelerates HOs to minimize RLF (TTT = 40 ms, A3-offset = -1 dB) [11].

C. Problem Formulation

Our objective is to determine the optimal configuration of A3-offset and TTT parameters that minimizes conflicting HO key performance indicators (KPIs), specifically the trade-off between RLFs and ping-pongs. The optimization problem is formally defined as follows:

$$\min_{\mathbf{A3}, \mathbf{TTT}} w_{PP} \cdot \frac{\sum_t \mathbb{1}_{PP_t}}{\sum_t \mathbb{1}_{HO_t} - \mathbb{1}_{HOF_t}} + w_{RLF} \cdot \frac{\sum_t \mathbb{1}_{RLF_t}}{\sum_t \mathbb{1}_{HO_t}}, \quad (2)$$

$$\text{s.t. } \mathbf{A3}_b \in (\underline{\mathbf{A3}}, \overline{\mathbf{A3}}), \quad b = 1, \dots, \mathcal{B} \quad (2a)$$

$$\mathbf{TTT}_b \in (\underline{\mathbf{TTT}}, \overline{\mathbf{TTT}}), \quad b = 1, \dots, \mathcal{B} \quad (2b)$$

where $\mathbb{1}_{(\cdot)}$ is an indicator function that evaluates whether a handover, RLF, or ping-pong occurs at time step t . It returns 1 if the corresponding event is observed, and 0 otherwise. The objective function minimizes a weighted sum of events, with w_{RLF} and w_{PP} representing the relative importance of reducing RLF and ping-pongs, respectively. The vectors $\mathbf{A3}$ and $\mathbf{TTT}$ contain the A3-offsets $\mathbf{A3}_b$ and time-to-trigger values $\mathbf{TTT}_b$ for all base stations $b \in \mathcal{B}$. The optimization is subject to bounds, where $\underline{\mathbf{A3}}$ and $\underline{\mathbf{TTT}}$ denote the minimum allowed values, and $\overline{\mathbf{A3}}$, $\overline{\mathbf{TTT}}$ the maximum.

The optimization problem (2) is nonconvex due to the non-concavity of the objective function. Additionally, with 30 cells, the problem involves 60 optimization variables, necessitating efficient techniques to navigate the large search space.

III. HIGH-DIMENSIONAL BAYESIAN OPTIMIZATION

A. Introduction to Bayesian Optimization

Bayesian Optimization (BO) iteratively builds a probabilistic *surrogate model* of the objective function $f(\cdot)$ using prior evaluations [13]. This surrogate, which is cheaper to evaluate than $f(\cdot)$, is updated as new data becomes available. An acquisition function $\alpha(\cdot)$ then guides the selection of the next point by balancing exploration (searching for better regions) and exploitation (refining known good solutions).

Objective function evaluation: We define a query point $\mathbf{x} = [\mathbf{A3}, \mathbf{TTT}]$ as the configuration of A3-offset and TTT values for each BS $b \in \mathcal{B}$, and obtain the corresponding objective value $f(\mathbf{x})$ from (2). The objective function $f(\cdot)$ is a mathematically intractable stochastic function that reflects the system model in Section II and accounts for the randomness in UE locations and wireless channels. We evaluate $f(\cdot)$ via system-level simulations, with each evaluation at point $\mathbf{x}$ producing a noisy sample $\tilde{f}(\mathbf{x})$. In real-world settings, these samples could also be obtained from measurements. Let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]$ denote a set of N query points, and $\mathbf{f}(\mathbf{X}) = [f_1, \dots, f_N]^\top$ the corresponding evaluations, where $f_i = f(\mathbf{x}_i)$ for $i = 1, \dots, N$.

Gaussian Process prior distribution: We use a Gaussian Process (GP) prior, $\tilde{f}(\cdot)$, to construct a surrogate model (i.e., the posterior) that approximates the objective function $f(\cdot)$ [13]. The resulting GP model enables the prediction of $\tilde{f}(\mathbf{x})$ at a query point $\mathbf{x}$ based on previously observed values, $\tilde{\mathbf{f}}(\mathbf{X}) = \tilde{\mathbf{f}}$, over which the model is trained. Formally, the GP prior on the objective function $\tilde{f}(\mathbf{x})$ assumes that for any set of input points $\mathbf{X}$, the corresponding function values $\tilde{\mathbf{f}}$ are jointly distributed as

$$p(\tilde{\mathbf{f}}) = \mathcal{N}(\tilde{\mathbf{f}} \mid \boldsymbol{\mu}(\mathbf{X}), \mathbf{K}(\mathbf{X})), \quad (3)$$

where $\boldsymbol{\mu}(\mathbf{X}) = [\mu(\mathbf{x}_1), \dots, \mu(\mathbf{x}_N)]^\top$ is the $N \times 1$ mean vector, and $\mathbf{K}(\mathbf{X})$ is the $N \times N$ covariance matrix, with each entry (i, j) given by the covariance function $k(\mathbf{x}_i, \mathbf{x}_j)$. For a given point $\mathbf{x}$, the mean function $\mu(\mathbf{x})$ provides prior knowledge about $f(\mathbf{x})$, while the kernel function $\mathbf{K}(\mathbf{X})$ captures the uncertainty between different input values $\mathbf{x}$.

Gaussian Process posterior distribution: Given a set of observed noisy samples $\tilde{\mathbf{f}}$ at previously sampled points $\mathbf{X}$, the posterior distribution of $\hat{f}(\mathbf{x})$ at a new query point $\mathbf{x}$ can be expressed as [14]:

$$p(\hat{f}(\mathbf{x}) = \hat{f} \mid \mathbf{X}, \tilde{\mathbf{f}}) = \mathcal{N}(\hat{f} \mid \mu(\mathbf{x} \mid \mathbf{X}, \tilde{\mathbf{f}}), \sigma^2(\mathbf{x} \mid \mathbf{X}, \tilde{\mathbf{f}})), \quad (4)$$

where the posterior mean and variance are given by:

$$\mu(\mathbf{x} \mid \mathbf{X}, \tilde{\mathbf{f}}) = \mu(\mathbf{x}) + \tilde{\mathbf{k}}(\mathbf{x})^\top (\tilde{\mathbf{K}}(\mathbf{X}))^{-1} (\tilde{\mathbf{f}} - \boldsymbol{\mu}(\mathbf{X})), \quad (5)$$

$$\sigma^2(\mathbf{x} \mid \mathbf{X}, \tilde{\mathbf{f}}) = k(\mathbf{x}, \mathbf{x}) - \tilde{\mathbf{k}}(\mathbf{x})^\top (\tilde{\mathbf{K}}(\mathbf{X}))^{-1} \tilde{\mathbf{k}}(\mathbf{x}), \quad (6)$$

where $\tilde{\mathbf{k}}(\mathbf{x}) = [k(\mathbf{x}, \mathbf{x}_1), \dots, k(\mathbf{x}, \mathbf{x}_N)]^\top$ is the $N \times 1$ covariance vector, and $\tilde{\mathbf{K}}(\mathbf{X}) = \mathbf{K}(\mathbf{X}) + \sigma^2 \mathbf{I}_N$, with σ^2 representing the observation noise (i.e., the variance of the Gaussian distribution), and $\mathbf{I}_N$ denoting the $N \times N$ identity matrix. Note that (5) and (6) define the mean and variance of the estimated function $\hat{f}(\mathbf{x})$, where the variance quantifies the uncertainty in the prediction.

Initial dataset creation and acquisition function: The BO algorithm starts by constructing a GP prior $\{\mu(\cdot), k(\cdot, \cdot)\}$ using an initial dataset $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_{N_0}, f_1, \dots, f_{N_0}\}$, comprising N_0 initial observations. This dataset is generated via system-level simulations based on the objective in (2) and the model outlined in Section II. For each point $\mathbf{x}_i \in \mathcal{D}$, the A3-offset and TTT values are randomly sampled from $[-1 \text{ dB}, 3 \text{ dB}]$ and $[40 \text{ ms}, 480 \text{ ms}]$, respectively. Subsequently, the algorithm selects the next query point $\mathbf{x}_n$ using an acquisition function $\alpha(\cdot)$. In this work, we adopt Thompson sampling as the acquisition strategy, due to its effectiveness in navigating exploration–exploitation trade-off [15].

Limitations of Bayesian Optimization: While BO [13] has proven effective for optimizing coverage-capacity tradeoffs and radio resource allocation [16]–[19], its scalability is limited to around twenty decision variables in continuous domains [14]. This constraint restricts its use for mobility parameter optimization in large-scale cellular networks [20]. To overcome this, we take the first step toward leveraging High-Dimensional BO (HD-BO) to optimize handover-related KPIs in realistic, large-scale deployments.

B. Introduction to High-dimensional Bayesian Optimization

To address the scalability limitations of BO in high-dimensional settings, we employ Trust Region Bayesian Optimization (TurBO) [15].¹ Unlike traditional BO, which relies on global surrogate models, TurBO manages multiple independent local models, each operating within a distinct region of the search space. It achieves global optimization by maintaining several trust regions (TRs) in parallel and allocating samples using an implicit multi-armed bandit strategy, thus focusing exploration on the most promising areas. Each TR employs a GP surrogate model, which retains BO’s key advantages—such as robustness to noise—while avoiding the limitations of overly simple local models. The TR is defined as a hyperrectangle centered at the current best solution f^* , with initial side length $L \leftarrow L_{\text{init}}$. The side length of each dimension is then scaled according to its corresponding GP length scale λ_i , ensuring adaptive resolution across dimensions. The side length for each dimension is given by:

$$L_i = \lambda_i L \cdot \left(\prod_{j=1}^d \lambda_j \right)^{-1/d} \quad (7)$$

where d is the total number of dimensions (i.e., optimization parameters under consideration). During each local optimization run, an acquisition function selects a batch of q candidates at each iteration, ensuring they remain within the designated

TR. If the TR’s side length L were large enough to cover the entire search space, this method would be equivalent to standard vanilla-BO. Thus, adjusting L is crucial: the TR must be large enough to encompass promising solutions while remaining compact enough to ensure the local model’s accuracy. The TR is dynamically resized based on optimization progress: it is doubled ($L \leftarrow \min\{L_{\text{max}}, 2L\}$) after τ_{succ} consecutive successes and halved ($L \leftarrow L/2$) after τ_{fail} consecutive failures. Success and failure counters are reset after each adjustment. If L falls below L_{min} , the TR is discarded, and a new one is initialized at L_{init} . The TR’s side length is capped at L_{max} . TurBO maintains m trust regions simultaneously, denoted as TR_l , where $l \in \{1, \dots, m\}$, each defined as a hyperrectangle with a base side length $L_l \leq L_{\text{max}}$. Candidate selection involves choosing a batch of q candidates from the union of all TRs. Thompson sampling is used for selecting candidates both within and across trust regions.

In this study, TurBO runs using an open-source repository [15] with the following hyperparameters: $\tau_{\text{succ}} = 3$, $\tau_{\text{fail}} = 15$, $L_{\text{init}} = 0.8$, $L_{\text{min}} = 2^{-7}$, $L_{\text{max}} = 1.6$.

IV. NUMERICAL RESULTS

To evaluate the HD-BO framework, we jointly optimize A3-offset and TTT to minimize the objective in (2). Our case study examines UE mobility at 3, 30, and 60 km/h across five main streets in London. We also assess HD-BO’s ability to generalize across user speeds, focusing on its performance with transfer learning. For our evaluations, we set $Q_{\text{out}} = -8 \text{ dB}$, $T_{310} = 1 \text{ s}$, and $T_p = 1 \text{ s}$, as per [11].

A. Performance of Data-driven Mobility Management

Fig. 2 shows RLF and ping-pong performance for UEs at 3, 30, and 60 km/h under HD-BO, comparing one-threshold (all cells share the same configuration) and per-cell optimization with $w_{\text{pp}} = 9$ and $w_{\text{RLF}} = 1$. Fig. 3 illustrates how optimizing for one speed affects others. Key observations are as follow:

- Compared to the 3GPP set-1 benchmark (tuned for ping-pong reduction), HD-BO with one-threshold optimization reduces ping-pongs by 20%. Per-cell optimization yields further gains, achieving a 35% reduction at 3 km/h. At 60 km/h, improvements are more significant—57% and 73% reductions for one-threshold and per-cell cases, respectively (see Fig. 2). These results highlight the advantage of per-cell optimization in mobility management.
- Optimizing handover parameters for a specific speed using HD-BO may degrade performance at other speeds. For example, a configuration optimized for 3 km/h leads to an $8\times$ increase in ping-pongs at 30 km/h (Fig. 3), as higher speeds require shorter TTT values to ensure timely handovers. These results highlight the need to apply HD-BO across all speed categories—3, 30, and 60 km/h—to achieve robust and balanced mobility performance.

¹We implemented and tested three HD-BO methods: Sparse Axis-Aligned Subspaces (SAASBO) [21, Section 4], BO via Variable Selection (VSBO) [22, Section 3], and Trust Region BO (TurBO) [15, Section 2]. TurBO demonstrated superior performance and higher suitability for the problem under consideration.




Fig. 2: RLFs and ping-pongs at different speeds under HD-BO (one-threshold and per-cell) and 3GPP baseline. Lower is better.




Fig. 3: Ping-pongs for UEs at different speeds with HD-BO optimized for a single target speed. Lower is better.

Fig. 4 shows the CDF of downlink SINR for UEs at different speeds when the network is jointly optimized across all speed categories using HD-BO. The black curves denote the 3GPP set-5 benchmark (upper bound), which enables immediate HOs without considering ping-pongs. The blue curves correspond to HD-BO optimization with $w_{PP} = 1$ and $w_{RLF} = 9$, prioritizing RLF reduction across five streets. HD-BO achieves SINR performance close to the 3GPP upper bound while reducing ping-pongs by 8% (3 km/h), 28% (60 km/h), and 100% (30 km/h).

B. Transfer Learning with HD-BO

For commercial deployment, machine learning models must perform reliably across diverse operating conditions [23]. This subsection evaluates the generalization capability of HD-BO across different UE speeds via transfer learning. In optimization,




Fig. 4: SINR for UEs at various speeds under HD-BO ($w_{PP} = 1$, $w_{RLF} = 9$) and 3GPP baseline.

transfer learning accelerates convergence on a new task (target) by reusing data or knowledge from a related task (source), particularly useful when generating the initial BO dataset $\mathcal{D}$ is costly (e.g., real-world measurements). Let $\mathcal{D}_{sr}$ and $\mathcal{D}_{tg}$ denote datasets for the source and target scenarios, respectively. We compare three configurations: 100% target data ($\mathcal{D} = \mathcal{D}_{tg}$), 100% source data ($\mathcal{D} = \mathcal{D}_{sr}$), and a 50%/50% mix. The source scenario features UEs moving at 60 km/h, while the target scenario uses the same deployment but at 30 km/h. The objective is to assess whether HD-BO can adapt to new mobility patterns without full retraining.

Convergence and performance of transfer learning: Fig. 5 illustrates the convergence behavior of transfer learning using HD-BO by tracking the best observed objective at each iteration n . To provide a practical performance measure, we plot the min-max normalized KPI, where 0 represents the 3GPP baseline (worst performance), and 1 corresponds to the upper bound—the best achievable KPI when the HD-BO posterior is fully trained on the target scenario ($\mathcal{D} = \mathcal{D}_{tg}$). The initial dataset $\mathcal{D}$ consists of $N_o = 60$ observations.

Fig. 5 shows that when 50% of $\mathcal{D}$ is sourced from the target scenario, convergence occurs at a rate comparable to the case where 100% of $\mathcal{D}$ is from the target scenario, demonstrating that transfer learning can effectively reduce the need for extensive new data collection. Fig. 6 compares the achieved RLF and ping-pong percentages across different initial dataset configurations. The results indicate comparable performance among all three dataset variations. Even without target-specific prior knowledge ($\mathcal{D} = \mathcal{D}_{sr}$), performance degrades by only 3%, underscoring the ability to generalize across related mobility scenarios.




Fig. 5: Convergence of transfer learning with diverse UE speeds.




Fig. 6: Ping-pongs and RLF with transfer learning. Lower is better.

V. CONCLUSION

We proposed a data-driven mobility management framework based on high-dimensional Bayesian optimization (HD-BO) for optimizing handovers in cellular networks. By jointly optimizing the A3-offset and TTT parameters across all cells, our framework effectively balances the trade-off between ping-pongs and radio link failures in site-specific scenarios. Moreover, since we leverage optimizing 3GPP standard-compliant parameters, our approach is readily applicable to real-world networks. We also demonstrated that transfer learning improves model generalization, reducing training overhead while maintaining performance. Practical evaluations across diverse UE speeds confirm the framework's effectiveness in dense urban deployments. Future directions include extending the method to beam-based mobility in mmWave systems and supporting vertical handovers in integrated terrestrial and non-terrestrial networks [24], [25].

REFERENCES

- [1] D. Lopez-Perez, N. Piovesan, and G. Geraci, "Capacity and power consumption of multi-layer 6G networks using the upper mid-band," in *Proc. IEEE ICC*, 2025.
- [2] S. M. Shahid, J.-H. Na, and S. Kwon, "Incorporating mobility prediction in handover procedure for frequent-handover mitigation in small-cell networks," *IEEE Trans. Network Sci. and Eng.*, 2024.

- [3] W.-C. Chien, Y. Huang, B.-Y. Chang, and W.-Y. Hwang, "Privacy-preserving handover optimization using federated learning and LSTM networks," *Sensors*, 2024.
- [4] R. Karmakar, G. Kaddoum, and S. Chattopadhyay, "Mobility management in 5G and beyond: A novel smart handover with adaptive time-to-trigger and hysteresis margin," *IEEE Trans. Mobile Comput.*, 2022.
- [5] R. Arshad and L. Lampe, "Stochastic geometry analysis of user mobility in RF/VLC hybrid networks," *IEEE Trans. Wireless Commun.*, 2021.
- [6] A. A. Malik, M. A. Jamshed, A. Nauman, A. Iqbal, A. Shakeel, and R. Hussain, "Performance evaluation of handover triggering condition estimation using mobility models in heterogeneous mobile networks," *IET Networks*, 2024.
- [7] E. Gures, I. Shayea, A. Alhammadi, M. Ergen, and H. Mohamad, "A comprehensive survey on mobility management in 5G heterogeneous networks: Architectures, challenges and solutions," *IEEE Access*, 2020.
- [8] 3GPP TR 38.843, "Study on artificial intelligence (AI)/machine learning (ML) for NR air interface (Release 18)," Jun. 2023.
- [9] J. Hoydis, F. A. Aoudia, S. Cammerer, M. Nimier-David, N. Binder, G. Marcus, and A. Keller, "Sionna RT: Differentiable ray tracing for radio propagation modeling," *arXiv:2303.11103*, 2023.
- [10] 3GPP Technical Report 36.814, "Evolved Universal Terrestrial Radio Access (EUTRA); Further advancements for E-UTRA physical layer aspects," Mar. 2017.
- [11] 3GPP TR 36.839, "Mobility enhancements in heterogeneous networks," Aug. 2012.
- [12] D. Lopez-Perez, I. Guvenc, and X. Chu, "Mobility management challenges in 3GPP heterogeneous networks," *IEEE Communications Magazine*, vol. 50, no. 12, pp. 70–78, 2012.
- [13] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, "Taking the human out of the loop: A review of Bayesian optimization," *Proc. IEEE*, vol. 104, no. 1, pp. 148–175, 2015.
- [14] P. I. Frazier, "A tutorial on Bayesian optimization," *arXiv:1807.02811*, 2018.
- [15] D. Eriksson, M. Pearce, J. Gardner, R. D. Turner, and M. Poloczek, "Scalable global optimization via local Bayesian optimization," *NeurIPS*, vol. 32, 2019.
- [16] M. Benzaghta, G. Geraci, D. López-Pérez, and A. Valcarce, "Designing cellular networks for UAV corridors via Bayesian optimization," in *Proc. IEEE Globecom*, 2023, pp. 4552–4557.
- [17] Y. Zhang, O. Simeone, S. T. Jose, L. Maggi, and A. Valcarce, "Bayesian and multi-armed contextual meta-optimization for efficient wireless radio resource management," *IEEE Trans. Cognitive Commun. and Netw.*, 2023.
- [18] L. Maggi, C. Mihailescu, Q. Cao, A. Tetich, S. Khan, S. Aaltonen, R. Koblit, M. Holma, S. Macchi, M. E. Ruggieri *et al.*, "Energy savings under performance constraints via carrier shutdown with Bayesian learning," in *Proc. EuCNC*, 2023, pp. 1–6.
- [19] E. Tekgul, T. Novlan, S. Akoum, and J. G. Andrews, "Joint uplink-downlink capacity and coverage optimization via site-specific learning of antenna settings," *IEEE Trans. Wireless Commun.*, 2024.
- [20] E. de Carvalho, A. Valcarce, and G. Geraci, "Towards mobility management with multi-objective Bayesian optimization," in *Proc. IEEE WCNC*, 2023, pp. 1–6.
- [21] D. Eriksson and M. Jankowiak, "High-dimensional Bayesian optimization with sparse axis-aligned subspaces," in *Uncertainty in Artificial Intelligence*. PMLR, 2021, pp. 493–503.
- [22] Y. Shen and C. Kingsford, "Computationally efficient high-dimensional Bayesian optimization via variable selection," *arXiv:2109.09264*, 2021.
- [23] X. Lin, "An overview of the 3GPP study on artificial intelligence for 5G new radio," *arXiv:2308.05315*, 2023.
- [24] G. Geraci, D. López-Pérez, M. Benzaghta, and S. Chatzinotas, "Integrating terrestrial and non-terrestrial networks: 3D opportunities and challenges," *IEEE Commun. Mag.*, 2023.
- [25] M. Benzaghta, G. Geraci, R. Nikbakht, and D. López-Pérez, "UAV communications in integrated terrestrial and non-terrestrial networks," in *Proc. IEEE Globecom*, 2022, pp. 3706–3711.