
Inductive Latent Context Persistence: Closing the Post-Handover Cold Start in 6G Radio Access Networks

Anonymous Authors¹

Abstract

In modern radio access networks (RANs) rule-based handover (HO) decisions (e.g., A3/A5) depend on user equipment (UE) measurements only, resulting in different HO decisions for UEs who are in the same location. To overcome it, graph neural networks (GNNs) based methods have been proposed to improve HO key performance indicators (KPIs) using more information than just the measurements. However, existing recurrent or graph-based methods discard the per-UE recurrent state at HO and reinitialize it at the target next-generation Node B (gNB) – losing useful measurement and mobility history as part of the overall context, forcing the target-side model to rebuild the state from post-HO measurements only. We address this post-HO cold start problem with Inductive Latent Context Persistence (ILCP), a learned latent synchronization mechanism that compresses the source-side per-UE recurrent state, transports it over the standard 3GPP Xn interface (Xn) as a 128-byte payload, and adapts it to the target gNB state space at HO. We model the RAN as a dynamic heterogeneous graph, allowing the model to treat UE nodes, gNB nodes, measurement edges, and Xn neighbor edges separately. On the Vienna 4G/5G drive-test, ILCP eliminates ping-pong HOs in the test split with 0.0% compared with 6.5% for the otherwise identical but no-transfer baseline and 22.6% for a Transformer baseline. It also achieves a +5.1 pp average post-HO accuracy gain, with a peak gain of +13.3 pp, over the no-transfer baseline in the 50–250 ms post-HO window. On a single NVIDIA GTX 1080 (8 GB), ILCP runs end-to-end at 7.7 ms p99 per handover decision. Under measurement perturbations including shadow fading, non-line-of-sight (NLOS) blockage, and

synchronization signal block (SSB)-burst sparsity, robustly trained ILCP keeps handover failure (HOF) in the 10–13% and exposes sub optimality of relying only on measurements. In the same fixed-reference-label setting, the 3GPP A3/A5 rule which relies only on measurements increases from 1.1% HOF on the unperturbed trace to 57–65% under perturbed measurements.

1. Introduction

Modern dense RANs deployments may cover a relatively small area with many overlapping capacity and coverage cells, and a UE moving through these should be handed over from one cell to another to maintain a strong radio connection with the network. Obviously, the quality of these HOs is an important KPI of the network deployment and HO control. In dense networks, with fast moving UEs, mobility performance is critical. A missed HO may cause a radio link failure (RLF) and possibly a multi-second outage, while an error-prone HO between two adjacent cells (e.g., a *ping-pong* HO) wastes control-plane signaling in addition to data loss (tgp, 2024a). Due to these reasons, work in 3GPP has produced a family of event-triggered rules (A3/A5, conditional handover (CHO), dual active protocol stack (DAPS), L2-triggered mobility (LTM)) that generally decide HOs from two simple parameters: a hysteresis margin, and a time-to-trigger (TTT), applied to filtered or averaged reference-signal received power (RSRP) measurements (tgp, 2024b). These rules are simple, robust and may be tuned per deployment. However, they do not do predictive extrapolations and they make decisions based on current measurement evidence only which are often imperfect due to a variety of reasons.

With increasing usage of DNNs and large computing, past work studied recurrent and graph-attention models trained on real drive-test traces and network topologies which match or beat the rule-based algorithm on next-cell accuracy at the HO instant, while also extrapolating into the future where rule-based events fail (Ozturk et al., 2019; Wickramasuriya et al., 2017; Lin et al., 2016; Hu et al., 2020; Veličković et al., 2018). The learning process consumes a per-step network information stream including topologies and measurements,

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

a recurrent or self-attentive temporal module integrating the stream into a hidden or embedded state $h_u(t)$ for each UE u , and a downstream head using that hidden state to produce a probability-type estimate over next possible cell(s). In short, a GNN captures the spatial geometry, the recurrent network module captures the temporal dynamics, and the head network ranks the next-cell candidates.

The hidden state $h_u(t)$ summarizes a UE’s local mobility-relevant context comprising of signal measurements from the serving gNB. Now, let us consider two critical issues: (i) a HO transfers the UE to a different gNB that has its own model state and local context. Then, the embedding $h_u(t)$, although representative of the context in the source gNB, may be seen as no longer too relevant for the target gNB after the HO. (ii) currently, the standardized Xn HANDOVER REQUEST carries radio resource control (RRC) state, security context, and UE capabilities (tgp, 2024c) as part of knowledge-transfer before HOs. Therefore, the target gNB will start $h_u(t)$ from scratch and must rebuild a useful h_u from the few measurements it has since the HO. We call this cold restart of the embedding as *Zero-Shot Cold Start* that wastes the historical learned representation of UE context and its usefulness in predicting the next best cell and avoiding an RLF or ping-pong.

To tackle these two issues above and drawbacks of the state-of-the-art, we treat the per-UE recurrent state as a portable network context which can be transferred (as part of overall knowledge-transfer for the HO) between gNBs. To address the practical issue of Xn message limits, we show that a 128-byte differential update is sufficient to preserve the predictive quality of a 128-dimensional gated recurrent unit (GRU) state across the HO boundary. To achieve this, our proposed ILCP protocol compresses the hidden state with a β -variational autoencoder (VAE), transports it on the standard 3GPP Xn interface, and projects it onto the target gNB’s state space at the moment of handover via a learned, gated multilayer perceptron (MLP).

We validate ILCP on a real drive-test traces from Vienna 4G/5G against five baselines using the original rule-based based next-cell decisions as the ground-truth “best cell”. Against the 3GPP A3/A5 rule, we do a comparison under measurement impairments expected of a real network to argue that relying on measurements and fixed rules alone results in sub optimal HO decisions.

Contributions: (i) We formalize the Zero-Shot Handover Cold Start as inductive domain shift on a dynamic heterogeneous graph in which the target gNB must infer a temporal hidden state without access to the measurement history (Section 3). (ii) We introduce ILCP (Section 4), a differential latent synchronization protocol whose 128-byte payload piggy-backs on the existing 3GPP Xn handover messages. (iii) On the Vienna trace ILCP improves post-handover next-

cell accuracy by +5.1 pp average (peak +13.3 pp), eliminates ping-pongs entirely (0.0% vs. 6.5% for the Zero-Knowledge transfer baseline and 22.6% for a Transformer baseline). (iv) Under realistic measurement impairments robustly trained ILCP holds HOF at 10–13% while the 3GPP A3/A5 rule collapses from 1.1% to 57–65% showing the rule-based limitation of idealized noise-free traces.

2. Related Work

Predictive mobility management: Classical machine learning (ML) approaches to handover prediction range from random forests on RSRP features to recurrent networks consuming time-step measurements (Ozturk et al., 2019; Wickramasuriya et al., 2017; Lin et al., 2016; Banerjee et al., 2022; 2021; 2020). More recent work applies GNNs to model the UE–cell topology explicitly (Veličković et al., 2018; Hu et al., 2020; Hamilton et al., 2017), and CHO / DAPS variants (tgp, 2024a;b) add procedural support for early target preparation. None of these works persists learned representations across the handover boundary; the predictor at the target gNB starts from a re-initialized state.

In-context transfer of learned state: Federated and split-learning systems (McMahan et al., 2017; Kairouz et al., 2021) share *parameter* updates between training rounds; our work differs in that the unit of transfer is a single UE’s recurrent state at the moment of handover, the transport latency budget is on the order of the handover preparation window (≈ 100 ms), and the target-side projection is a learned operator. ILCP is therefore closer to *in-context* transfer of episodic memory than to federated optimization, and is complementary to continual-learning methods (Kirkpatrick et al., 2017) that mitigate forgetting in a single per-gNB model. The compressor is a β -VAE (Higgins et al., 2017; Kipf & Welling, 2016) optimized end-to-end with the downstream candidate-set scoring loss, keeping with the information-bottleneck view of representation learning (Tishby & Zaslavsky, 2015). The realignment of source-side and target-side state spaces is reminiscent of CORAL and Subspace Alignment (Sun & Saenko, 2016; Fernando et al., 2013), but specialized to the heterogeneous-graph plus recurrent-state setting.

Robust learning for telecom data: For comparison with A3/A5 under measurement inaccuracies, our mixed-perturbation training recipe is closely related to AugMix and consistency training (Hendrycks et al., 2020; Xie et al., 2020; Miyato et al., 2018).

3. System and Problem Formulation

In this section we formulate the overall system and learning problem. See Figure 2 for an end-to-end system diagram. For readability purposes and because of page constraint,

Figure 2 has been put in Appendix B.1.

Dynamic heterogeneous RAN graph: At any time instant t , we represent the radio access network as a dynamic heterogeneous graph

$$G(t) = (V_{\text{UE}} \cup V_{\text{gNB}}, E_{\text{meas}}(t) \cup E_{\text{Xn}}) \quad (1)$$

Here, V_{UE} denotes the set of currently active UEs, and V_{gNB} denotes the set of deployed base stations. The edge set contains two relation types:

- $E_{\text{meas}}(t) \subseteq V_{\text{UE}} \times V_{\text{gNB}}$, the time-varying set of UE-cell measurement edges, with per-edge attributes such as RSRP, reference-signal received quality (RSRQ), and signal-to-interference-plus-noise ratio (SINR);
- $E_{\text{Xn}} \subseteq V_{\text{gNB}} \times V_{\text{gNB}}$, the deployment-static set of neighboring gNBs connected via the Xn interface.

For each UE $u \in V_{\text{UE}}$, we denote its serving cell at time t by $c_u(t) \in V_{\text{gNB}}$.

Prediction model: Our model comprises following:

1. Given the graph snapshot $G(t)$, a heterogeneous-attention encoder f_θ produces spatial embeddings for the UEs and candidate cells. In particular, for each UE u we compute

$$x_u(t) = f_\theta(G(t), u) \in \mathbb{R}^{d_x}, \quad (2)$$

and for each candidate cell c_k we denote its embedding by $e_{c_k} \in \mathbb{R}^{d_x}$. The embedding $x_u(t)$ summarizes the UE's current local radio and graph context, including measured link quality to visible cells, the current serving-cell relation, neighboring candidate cells, and the connectivity structure induced by measurement and Xn edges.

2. A recurrent module r_ϕ aggregates the per-step UE embeddings over time to form a temporal state

$$h_u(t) = r_\phi(h_u(t-1), x_u(t)) \in \mathbb{R}^{d_h}, \quad (3)$$

which encodes recent temporal information from the UE's measurement and mobility evolution, such as changes in link qualities, visible candidates, and serving-cell associations over the recent past.

3. At time t , we consider only the K visible candidate cells with the highest current RSRP to be the serving cell at horizon Δ . For each candidate cell c_k , a candidate-set scoring function $g_\psi^{(\Delta)}$ assigns the scalar score

$$s_{u,k}^{(\Delta)}(t) = g_\psi^{(\Delta)}([h_u(t), e_{c_k}, h_u(t) \odot e_{c_k}]), \quad (4)$$

where \odot denotes element-wise multiplication. The term $h_u(t) \odot e_{c_k}$ gives the scorer an explicit dimension-wise interaction between the UE temporal state and the candidate-cell embedding. The predicted serving cell at horizon Δ is then

$$\hat{c}_u(t + \Delta) = \arg \max_{k \in \{1, \dots, K\}} s_{u,k}^{(\Delta)}(t). \quad (5)$$

Zero-shot handover cold start problem: Consider a handover of UE u at time t^* , where the serving cell changes from the source cell $c_u(t^*-)$ to the target cell $c_u(t^*+)$. Immediately before handover, the recurrent state $h_u(t^*-)$ has accumulated information from the pre-handover observation sequence. Immediately after handover, the target side only has access to the post-handover measurement edges. If we do not transfer state across the handover, the recurrent module must restart from an initialization h_u^{init} , and therefore loses the useful temporal context accumulated before handover. We quantify the resulting cold-start penalty by comparing a *warm* predictor, which receives oracle access to the pre-handover hidden state, with a *cold* predictor, which restarts from h_u^{init} :

$$\mathcal{L}_{\text{cold}} = \sum_{\substack{(t^*, u) \\ \text{handovers}}} \left[\ell(\hat{c}_u^{\text{cold}}(t^*), c_u^*(t^*)) - \ell(\hat{c}_u^{\text{warm}}(t^*), c_u^*(t^*)) \right], \quad (6)$$

where $c_u^*(t^*)$ denotes the observed serving cell in the unperturbed reference trace, ℓ is the 0/1 prediction loss, \hat{c}_u^{cold} denotes prediction with re-initialized target-side state, and \hat{c}_u^{warm} denotes prediction with oracle access to the pre-handover hidden state. This gap captures the cost of the inductive domain shift introduced by handover.

Objective and Method: Our goal is to reduce the gap between the cold and warm predictors at the target gNB, using only a constant-size message delivered over the existing Xn interface under realistic latency constraints, in the next section we present our proposed ILCP method. ILCP maps the pre-handover recurrent state $h_u(t^*-)$ to a compact latent z_u , transports z_u over Xn, and reconstructs an adapted target-side state h_u^{new} that can act as a substitute for the re-initialized state h_u^{init} after handover. We fix the candidate-set size to $K = 8$ as a typical operational value. Note that the scoring function $g_\psi^{(\Delta)}$ is applied independently to each candidate and shared across the K candidates. The method therefore does not depend on a fixed global parameterization of cell identifiers and can handle candidate-set changes caused by topology evolution.

4. Inductive Latent Context Persistence

We implement ILCP as a compact (128-Byte), learned synchronization protocol that can be carried over the existing 3GPP Xn interface. We train the full model for the downstream candidate-set prediction task. For each UE u at time t , we convert the candidate scores from Section 3 into a probability distribution over the visible candidate set:

$$p_\Theta^{(\Delta)}(c_k | u, t) = \frac{\exp(s_{u,k}^{(\Delta)}(t))}{\sum_{j=1}^K \exp(s_{u,j}^{(\Delta)}(t))}, \quad (7)$$

where $\{c_1, \dots, c_K\}$ is the visible candidate set at time t , and $s_{u,k}^{(\Delta)}(t)$ is defined in Eq. (4). We then optimize the downstream prediction objective

$$\mathcal{L}_{\text{pred}}^{(\Delta)} = - \sum_{(u,t)} \log p_{\Theta}^{(\Delta)}(c_u(t + \Delta) | u, t), \quad (8)$$

where $c_u(t + \Delta)$ is the ground-truth serving cell at horizon Δ . To this end, we jointly train the heterogeneous-graph encoder f_{θ} , the recurrent module r_{ϕ} , the candidate-set scorer $g_{\psi}^{(\Delta)}$, the latent compressor \mathcal{E}_{ϕ} , the latent decoder \mathcal{D}_{ψ} , and the target-side projection block comprising the gate network and the MLP in Eq. (11). See Figure 2 for an end-to-end system diagram.

Heterogeneous-graph backbone: We encode each graph snapshot $G(t)$ with three heterogeneous graph transformer (HGT) layers (Hu et al., 2020) and use 4 attention heads per layer. We use type-specific attention parameters for UE \leftrightarrow gNB measurement edges and gNB \leftrightarrow gNB Xn edges, which preserves the asymmetry between instantaneous radio-measurement relations and deployment-static topology relations. By contrast, a relation-agnostic GAT (Veličković et al., 2018) averages across both edge classes and loses about 30 pp on our main accuracy metric (see GAT-Temporal in Table 1).

Per-UE recurrent module: We instantiate the recurrent module r_{ϕ} as a GRU (Cho et al., 2014) with hidden dimension $d_h = 128$, and update the temporal state as

$$h_u(t) = r_{\phi}(h_u(t-1), x_u(t)) = \text{GRU}(h_u(t-1), x_u(t)). \quad (9)$$

This state encodes recent temporal information from the UE’s measurement and mobility evolution. In state-of-the-art, $h_u(t^*-)$ is discarded after every handover whereas we preserve and transfer it through the steps below.

β -VAE compressor: We compress each recurrent state $h_u \in \mathbb{R}^{128}$ into a 32-dimensional latent vector

$$z_u = \mathcal{E}_{\phi}(h_u)$$

with an end-to-end-trained β -VAE (Higgins et al., 2017; Kipf & Welling, 2016), which yields a 128-byte FP32 payload. We train the compressor with

$$\mathcal{L}_{\text{VAE}} = \mathbb{E}_{q_{\phi}(z|h)}[\log p_{\psi}(h | z)] - \beta \text{KL}(q_{\phi}(z | h) \| p(z)), \quad (10)$$

with a suitable value of β , and optimize it jointly with Eq. (8). Joint training makes the compressor task-aware such that it preserves those dimensions of h_u that matter for the candidate scorer after an HO has taken place recently, and compresses the rest.

Latent transport over the Xn interface: At handover time t^* , we append z_u as an optional information element to the HANDOVER REQUEST message defined in TS 38.423 (tgp, 2024c). We note that the resulting 128-byte payload fits within existing message-size budgets.

Gated projection at the target gNB: When the target gNB receives the payload, we first decode it into

$$\tilde{h}_u = \mathcal{D}_{\psi}(z_u).$$

We then compute the target-side spatial embedding x_u^{new} from the new local graph context and combine them through

$$h_u^{\text{new}} = \text{LN}\left(\tilde{h}_u + \gamma \odot \text{MLP}([\tilde{h}_u, x_u^{\text{new}}])\right), \quad (11)$$

where,

$$\gamma = \sigma(g_{\theta}([\tilde{h}_u, x_u^{\text{new}}]))$$

is a learned sigmoid gate and LN denotes LayerNorm. We treat \tilde{h}_u as a transferred prior and use the gated MLP term to adapt that prior to the new serving-cell context. The gate γ controls how strongly we correct the transferred state using the new target-side observation, while LayerNorm keeps the combined representation well scaled and prevents either from dominating purely because of norm drift between source-side and target-side embedding spaces. During training, we apply inbound dropout to \tilde{h}_u with rate 0.2. This keeps the projection block useful even when the transferred state is incomplete or partially corrupted.

Supervised and Robust training: Before moving to the experiments, we clarify the supervised training used in this paper. We use the serving-cell sequence observed in the trace as the reference label and assume, for the purpose of supervised training, that this sequence represents the desired next-cell target. This is an experimental assumption as the “best cell” cannot be observed in the network trace. In fact, the logged serving cell is the decision realized by the deployed handover logic, which is typically A3/A5-like, and it need not always be the operationally best cell. This assumption also affects the comparison with A3/A5 on the unperturbed trace. If the label is itself aligned with A3/A5 behavior, then A3/A5 has an inherent advantage on exact-match metrics. For this reason, we interpret the unperturbed-trace comparison with care and use additional metrics, such as ping-pong rate and post-handover recovery, to evaluate whether a method produces stable mobility behavior. To model unavoidable variability in UE measurements, and to expose limitations of A3/A5 rules in recovering an optimal best serving cell consistently, we also train robust versions of both ILCP and the Zero-Knowledge baseline. At each step, we sample the input measurement stream from a mixture of the unperturbed reference trace, shadow-fading-corrupted RSRP ($\sigma_s \sim \mathcal{U}\{6, 12\}$ dB), random NLOS blockage, and SSB-burst sub-sampling. We keep the supervision target fixed to the serving cell observed on the unperturbed reference trace, so the model learns to recover the same reference decision under impaired measurements rather than imitate impairment-induced fluctuations. Finally, we remark that the experiment results are somewhat limited due to the limitations of the available trace data such as fixed topology (so that denser networks cannot be considered) and a limited number of actual HOs. In the future extension of this work, we will use network simulators to produce traces with opti-

mal serving cells (in some sense), i.e., not relying on A3/A5 decisions, and demonstrate the full potential of our method in different network topologies. We will also use a multitude of node features that help predict the “operationally best cell”.

5. Experiments

We evaluate the setup of Section 4 on the Vienna trace (Anonymized, 2024), using the baselines and metrics below. It is a multi-cell, multi-tier urban trace with dense cell overlap, 31 handover events in the held-out test split, and per-step measurements at 100 Hz. Unless noted, main-body results use the unperturbed trace; measurement-impairment studies follow Appendix C under the fixed-reference-label protocol described in Section 4.

Baseline Methods: We compare ILCP against five baseline methods and one rule-based handover rule.

1. *Zero-Knowledge* HGT is an otherwise identical HGT+GRU prediction method but which re-initializes the per-UE recurrent state at every HO. This baseline isolates the effect of cross-handover state persistence because it shares the same graph backbone, recurrent module, and candidate-set scorer as ILCP.
2. *GAT-Temporal* uses a single-relation graph attention network (GAT) (Veličković et al., 2018) with the same hidden dimension, followed by the same GRU. This baseline tests the value of relation-specific heterogeneous attention.
3. *Transformer-Temporal* uses the same HGT backbone followed by a Transformer encoder (Vaswani et al., 2017) over a sliding window of length $T = 16$. This baseline compares recurrent state accumulation with windowed self-attention.
4. *LSTM* is a direct sequence model over the raw per-step measurement vector (Hochreiter & Schmidhuber, 1997), without the heterogeneous graph backbone.
5. *A3/A5* is the standard 3GPP event-triggered handover rule (tgp, 2024a;b), with hysteresis H and TTT T set to default values.

All neural-network-based methods use the architectural hyperparameters in Table 2 of Appendix B.3. Unless stated otherwise, we report the mean over 1000 bootstrap resamples of the test split, together with 95% confidence intervals (CIs); the resampling procedure is given in Appendix D.

Metrics: In the following, $\text{Acc}@t = 0$ denotes next-cell accuracy when we evaluate the model at the HO instant t^* . The target is the future serving cell (selected using fixed network logic) at the prediction horizon Δ , i.e., $c_u(t^* + \Delta)$. After an HO, $\text{Acc}@t = \delta$ denotes next-cell accuracy when we evaluate the model δ steps after the HO instant. For an HO at

Table 1. Headline mobility-prediction metrics on the Vienna 4G/5G drive-test (test split, 31 handovers)

Method	Acc@t=0	HOF	PP	Ovr
ILCP (ours)	83.9 (71–94)	16.1 (3–32)	0.0 (0–0)	74.1
ZK-HGT	87.1 (74–97)	12.9 (3–26)	6.5 (0–16)	75.6
GAT-Temporal	22.6 (10–39)	77.4 (61–90)	61.3 (45–77)	19.3
Transformer	77.4 (61–90)	22.6 (10–39)	22.6 (10–39)	66.9
LSTM	12.9 (0–26)	87.1 (74–97)	83.9 (71–94)	6.3
A3/A5	100.0 (100–100)	0.0 (0–0)	3.2 (0–10)	72.4

time t^* , this means that the model is evaluated at $t^* + \delta$ and its output is compared with the ground-truth future serving cell $c_u(t^* + \delta + \Delta)$. Plotting $\text{Acc}@t = \delta$ as a function of δ shows how quickly the target-side model recovers useful context after HO, and therefore makes the post-HO cold-start gap visible. In the following, HOF is the fraction of HO events for which the predicted target cell at the HO instant is incorrect, while Ping-pong rate (PP) represents problematic HO that are reversed within a short time-frame (for example, 500 ms).

5.1. Overall Mobility-Prediction Performance

Table 1 summarizes the main mobility-prediction metrics on Vienna 4G/5G with bootstrap 95% CIs. We observe that ILCP eliminates ping-pongs entirely with PP = 0.0%, compared with 6.5% for the otherwise identical Zero-Knowledge baseline and 22.6% for the Transformer baseline. At the handover instant, ILCP reaches 83.9% $\text{Acc}@t = 0$, while the Zero-Knowledge baseline reaches 87.1%. This difference is within the reported confidence intervals and, as shown in Section 5.2, the ordering changes in the early post-handover window. The sequence-only baselines perform substantially worse. LSTM reaches only 12.9% $\text{Acc}@t = 0$, and GAT-Temporal reaches 22.6%. This indicates that the heterogeneous graph structure carries important predictive information. In particular, relation-specific modeling of UE-to-cell measurement edges and Xn neighbor edges is more useful than recurrence alone. The A3/A5 rule reaches 100% on the clean trace trivially because the HO labels are tied to the rule-based HO behavior so the comparison here is misleading and only shown for sanity check. The ILCP does not match the A3/A5 rule even if it is trained to predict the observed serving cells because, in contrast to A3/A5, ILCP uses UE’s network graph context summarized by its embedding which also includes measurements but does not solely rely on them. We show this advantage in Appendix C where we show the limitation of A3/A5 and other baselines under practical measurement variations and impairments.

It is important to mention that the ping-pong column is very important operationally in current RAN deployments and even more important in future (much denser with many overlapping small cells) cloud deployments in which the




Figure 1. Post-handover next-cell accuracy on Vienna as a function of timesteps after the handover instant, with 1 step \approx 10 ms.

best cell in terms of signal-strength can vary faster than our test-data. In this case, A3/A5 type rules may have an even higher ping-pong rate. In this regard, we remind that the Zero-Knowledge baseline shares the same HGT backbone, recurrent module, and candidate-set scoring function as ILCP, but it re-initializes the per-UE recurrent state at every HO. Its nonzero ping-pong rate therefore isolates the effect of losing cross-HO state. The Transformer baseline also produces a higher ping-pong rate. In both cases, the target-side state has not yet rebuilt enough temporal context when a small RSRP fluctuation appears after HO. In contrast, ILCP carries source-side history into the target side and can therefore identify such events as transient rather than immediately HO-worthy.

5.2. Post-handover Cold-Start Gap

Figure 1 plots $\text{Acc}@_\delta$ for $\delta \in [0, 30]$ post-handover steps. Each point evaluates the model δ steps after the handover instant and compares its output with the ground-truth future serving cell at horizon Δ . The Zero-Knowledge baseline and ILCP start in the same range at the HO instant, but ILCP improves over the next few post-handover steps. This is the interval in which the Zero-Knowledge baseline is still rebuilding temporal context from a re-initialized state, while ILCP already uses an adapted version of the source-side recurrent state. The gap peaks at +13.3 pp on step 3 and averages +5.1 pp over steps 5 to 25. This measured gap is the empirical manifestation of the cold-start penalty defined in Equation (6). In operational terms, the result means that the target-side model becomes useful faster after HO when source-side recurrent context is preserved. This effect is once again vital for denser deployments with more overlapping cells, where the early post-handover period can contain borderline candidate-cell decisions and short-lived measurement fluctuations. Because of the strict page restriction, detailed results about robustness under measurement

impairments could not be included within the main body of the paper; instead, they can be found in Appendix C.

6. Conclusion

On the Vienna 4G/5G trace, ILCP reduces the cold-start gap relative to the no-transfer baseline and eliminates ping-pongs in the test split. This is very critical for dense deployments with overlapping cells and borderline handover conditions. The result suggests that preserving recent source-side measurement and mobility history helps the target side distinguish stable mobility trends from short-lived measurement fluctuations. The comparison with A3/A5 should be interpreted carefully. On the unperturbed trace, A3/A5 is strongly aligned with the reference handover labels, so its clean-trace HOF is not directly comparable to the learned models. The perturbation experiments in Appendix C test a different question, i.e., whether a method can recover the fixed reference optimal serving-cell sequence when the input measurements are degraded. Under this setting, robustly trained ILCP is less sensitive to shadow fading, blockage, and sparse sampling than the reactive A3/A5 rule.

The method is lightweight enough for practical deployment. The 128-byte latent payload adds little Xn overhead. The successful compression indicates that the recurrent state contains a relatively low-dimensional task-relevant subspace.

There are some limitations which still remain including limitations due to usage of a real-life publicly available network trace. First, the dataset contain a limited number of handover events, so larger traces are needed to validate the transfer mechanism more strongly. Secondly, the perturbation experiments model measurement impairments but do not cover cases where the operationally preferred serving cell truly changes, for example due to an outage or a load shift. Our future work will include controlled simulations comprising varying network topologies, large number of traces, optimal serving cells due to varying criteria, and varying UE speeds and mobility patterns.

References

- NR; overall description; stage-2 (release 18). Technical Specification 38.300, 3GPP, 2024a.
- NR; radio resource control (RRC); protocol specification (release 18). Technical Specification 38.331, 3GPP, 2024b.
- NG-RAN; Xn application protocol (XnAP) (release 18). Technical Specification 38.423, 3GPP, 2024c.
- Anonymized. Vienna 4G/5G drive-test dataset for handover and mobility research. Public release; details in supplementary material, 2024.
- Banerjee, A., Mwanje, S. S., and Carle, G. Game theoretic

- 330 conflict resolution mechanism for cognitive autonomous
 331 networks. In *2020 International Symposium on Perform-*
 332 *ance Evaluation of Computer and Telecommunication*
 333 *Systems (SPECTS)*, pp. 1–8. IEEE, 2020.
- 334 Banerjee, A., Mwanje, S. S., and Carle, G. Toward control
 335 and coordination in cognitive autonomous networks.
 336 *IEEE Transactions on Network and Service Management*,
 337 19(1):49–60, 2021.
- 338 Banerjee, A., Mwanje, S. S., and Carle, G. Trust and perfor-
 339 mance in future ai-enabled, open, multi-vendor network
 340 management automation. *IEEE Transactions on Network*
 341 *and Service Management*, 20(2):995–1007, 2022.
- 342 Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D.,
 343 Bougares, F., Schwenk, H., and Bengio, Y. Learning
 344 phrase representations using RNN encoder–decoder for
 345 statistical machine translation. In *Proceedings of the 2014*
 346 *Conference on Empirical Methods in Natural Language*
 347 *Processing (EMNLP)*, pp. 1724–1734, 2014.
- 348 Fernando, B., Habrard, A., Sebban, M., and Tuytelaars, T.
 349 Unsupervised visual domain adaptation using subspace
 350 alignment. In *Proceedings of the IEEE International*
 351 *Conference on Computer Vision (ICCV)*, pp. 2960–2967,
 352 2013.
- 353 Hamilton, W. L., Ying, R., and Leskovec, J. Inductive
 354 representation learning on large graphs. In *Advances in*
 355 *Neural Information Processing Systems (NeurIPS)*, 2017.
- 356 Hendrycks, D., Mu, N., Cubuk, E. D., Zoph, B., Gilmer,
 357 J., and Lakshminarayanan, B. AugMix: A simple data
 358 processing method to improve robustness and uncertainty.
 359 In *International Conference on Learning Representations*
 360 *(ICLR)*, 2020.
- 361 Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X.,
 362 Botvinick, M., Mohamed, S., and Lerchner, A. beta-
 363 VAE: Learning basic visual concepts with a constrained
 364 variational framework. In *International Conference on*
 365 *Learning Representations (ICLR)*, 2017.
- 366 Hochreiter, S. and Schmidhuber, J. Long short-term memory.
 367 *Neural Computation*, 9(8):1735–1780, 1997.
- 368 Hu, Z., Dong, Y., Wang, K., and Sun, Y. Heterogeneous
 369 graph transformer. In *Proceedings of The Web Conference*
 370 *2020 (WWW)*, pp. 2704–2710, 2020.
- 371 Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis,
 372 M., et al. Advances and open problems in federated
 373 learning. *Foundations and Trends in Machine Learning*,
 374 14(1–2):1–210, 2021.
- 375 Kipf, T. N. and Welling, M. Variational graph auto-encoders.
 376 *NIPS Workshop on Bayesian Deep Learning*, 2016.
- 377 Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Des-
 378 jardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T.,
 379 et al. Overcoming catastrophic forgetting in neural net-
 380 works. *Proceedings of the National Academy of Sciences*,
 381 114(13):3521–3526, 2017.
- 382 Lin, P.-C., Casanova, L. F. G., and Fatty, B. K. Data-driven
 383 handover optimization in next generation mobile commu-
 384 nication networks. *Mobile Information Systems*, 2016(1):
 2368427, 2016.
- McMahan, B., Moore, E., Ramage, D., Hampson, S., and
 y Arcas, B. A. Communication-efficient learning of deep
 networks from decentralized data. In *Artificial Intelli-*
gence and Statistics (AISTATS), pp. 1273–1282, 2017.
- Miyato, T., Maeda, S.-i., Koyama, M., and Ishii, S. Vir-
 tual adversarial training: A regularization method for
 supervised and semi-supervised learning. 2018.
- Ozturk, M., Gogate, M., Onireti, O., Adeel, A., Hussain, A.,
 and Imran, M. A. A novel deep learning driven, low-cost
 mobility prediction approach for 5g cellular networks:
 The case of the control/data separation architecture (cdsa).
Neurocomputing, 358:479–489, 2019.
- Sun, B. and Saenko, K. Deep CORAL: Correlation align-
 ment for deep domain adaptation. In *European Confer-*
ence on Computer Vision (ECCV) Workshops, pp. 443–
 450, 2016.
- Tibshirani, R. J. and Efron, B. An introduction to the boot-
 strap. *Monographs on statistics and applied probability*,
 57(1):1–436, 1993.
- Tishby, N. and Zaslavsky, N. Deep learning and the infor-
 mation bottleneck principle. In *2015 IEEE Information*
Theory Workshop (ITW), pp. 1–5, 2015.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones,
 L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Atten-
 tion is all you need. In *Advances in Neural Information*
Processing Systems (NeurIPS), 2017.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A.,
 Liò, P., and Bengio, Y. Graph attention networks. In
International Conference on Learning Representations
(ICLR), 2018.
- Wickramasuriya, D. S., Perumalla, C. A., Davaslioglu, K.,
 and Gitlin, R. D. Base station prediction and proactive
 mobility management in virtual cells using recurrent neu-
 ral networks. In *2017 IEEE 18th Wireless and Microwave*
Technology Conference (WAMICON), pp. 1–6, 2017.
- Xie, Q., Dai, Z., Hovy, E., Luong, T., and Le, Q. Un-
 supervised data augmentation for consistency training.
 In *Advances in Neural Information Processing Systems*
(NeurIPS), 2020.

- 385 **A. List of abbreviations**
 386
 387
 388 **CHO** conditional handover
 389
 390 **CI** confidence interval
 391
 392 **DAPS** dual active protocol stack
 393
 394 **GAT** graph attention network
 395
 396 **gNB** next-generation Node B
 397
 398 **GNN** graph neural network
 399
 400 **GRU** gated recurrent unit
 401
 402 **HGT** heterogeneous graph transformer
 403
 404 **HO** handover
 405
 406 **HOF** handover failure
 407
 408 **ILCP** Inductive Latent Context Persistence
 409
 410 **KPI** key performance indicator
 411
 412 **LTE** Long-Term Evolution
 413
 414 **LTM** L2-triggered mobility
 415
 416 **ML** machine learning
 417
 418 **MLP** multilayer perceptron
 419
 420 **NLOS** non-line-of-sight
 421
 422 **RAN** radio access network
 423
 424 **RLF** radio link failure
 425
 426 **RRC** radio resource control
 427
 428 **RRM** radio resource management
 429
 430 **RSRP** reference-signal received power
 431
 432 **RSRQ** reference-signal received quality
 433
 434 **SINR** signal-to-interference-plus-noise ratio
 435
 436 **SSB** synchronization signal block
 437
 438 **TTT** time-to-trigger
 439
 440 **UE** user equipment
 441
 442 **VAE** variational autoencoder
 443
 444 **Xn** Xn interface

B. Details on System Design and Experimental Setup

B.1. System Diagram

Figure 2 summarizes the end-to-end ILCP path from the source gNB to the target gNB at a handover instant: the heterogeneous-graph backbone and GRU produce the pre-handover recurrent state, which the β -VAE maps to a fixed-size latent carried as an optional information element on the Xn HANDOVER REQUEST. On the target side, the latent is decoded and fused with the fresh target-side spatial embedding via the gated MLP and LayerNorm block before candidate-set scoring. Arrows indicate control and data flow through encoding, transport, adaptation, and prediction; dotted boxes group components that share parameters across handovers while remaining logically distinct on source versus target.

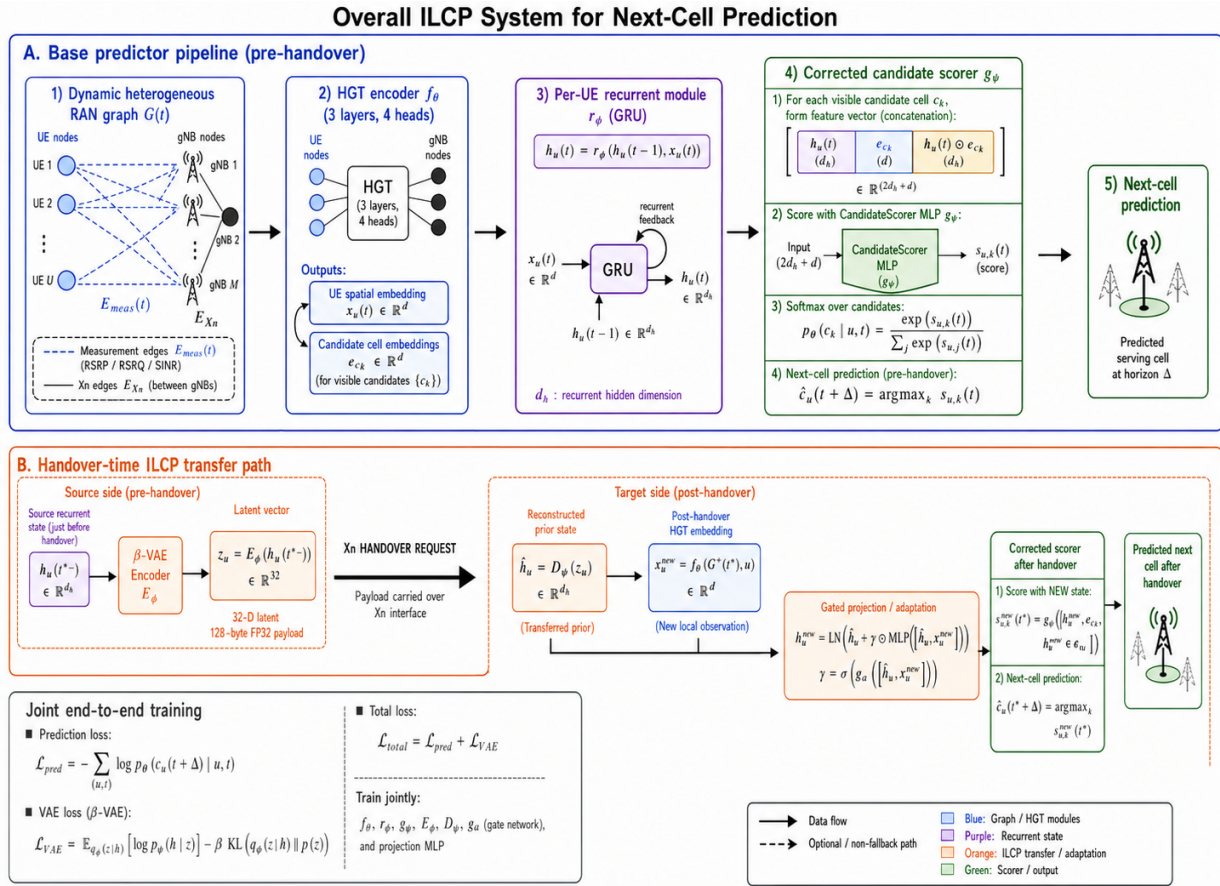


Figure 2. System Diagram

B.2. ILCP inference pseudocode

The pseudocode reflects the implementation of our proposed solution.

Algorithm 1 ILCP inference at a handover instant t^*

```

495 1: Input: source-side recurrent state  $h_u^{\text{src}} \in \mathbb{R}^{d_h}$  for UE  $u$ ; current target-side spatial embedding  $x_u^{\text{new}} \in \mathbb{R}^{d_x}$ ;  $\beta$ -VAE
496 encoder  $\mathcal{E}_\phi$ , decoder  $\mathcal{D}_\psi$ ; gating MLP  $g_\theta$ ; LayerNorm LN; visible candidate cells  $\{c_1, \dots, c_K\}$  with embeddings
497  $\{e_{c_k}\}_{k=1}^K$ ; candidate-set scorer  $g_\psi$ .
498 2: Output: predicted next cell  $\hat{c}_u(t^* + \Delta)$  and updated state  $h_u^{\text{new}}$ .
499 3:
500 4: // Source side, at handover decision
501 5:  $z_u \leftarrow \mathcal{E}_\phi(h_u^{\text{src}})$  {32-dimensional latent,  $\leq 128$  B FP32}
502 6: Append  $z_u$  as an optional IE to the HANDOVER REQUEST on the Xn interface.
503 7:
504 8: // Target side, on receipt of HANDOVER REQUEST
505 9:  $\tilde{h}_u \leftarrow \mathcal{D}_\psi(z_u)$ 
506 10:  $\gamma \leftarrow \sigma\left(g_\theta\left([\tilde{h}_u, x_u^{\text{new}}]\right)\right)$  {learned sigmoid gate}
507 11:  $h_u^{\text{new}} \leftarrow \text{LN}\left(\tilde{h}_u + \gamma \odot \text{MLP}\left([\tilde{h}_u, x_u^{\text{new}}]\right)\right)$ 
508 12:
509 13: // Candidate-set scoring
510 14: for  $k = 1$  to  $K$  do
511 15:    $s_k \leftarrow g_\psi^{(\Delta)}([h_u^{\text{new}}, e_{c_k}, h_u^{\text{new}} \odot e_{c_k}])$ 
512 16: end for
513 17:  $\hat{c}_u(t^* + \Delta) \leftarrow c_{\arg \max_k s_k}$ 
514 18: return  $\hat{c}_u(t^* + \Delta), h_u^{\text{new}}$ 

```

B.3. Experimental setup and hyperparameters

In this section we describe our experimental setup and parameters.

Data split: The Vienna trajectory traces are partitioned into training, validation, and test trajectories. The split contains 2200 measurement steps over 31 handover events in the test partition, and 25 500 steps over 95 events in the train+val partitions.

Preprocessing: Per-edge measurement features ($RSRP$, $RSRQ$, $SINR$) are normalized per-cell to zero mean and unit variance over the train split, with the normalisation parameters fixed on validation and test. Per-step UE identifiers are anonymized.

Hyperparameters: The hyperparameters are shown in Table 2.

Table 2. Hyperparameters used for the experiments. The same configuration is used for ILCP and the baselines except where the architecture differs structurally (e.g., Transformer-Temporal replaces the GRU with a single-layer self-attention encoder over a 16-step sliding window).

Parameter	Value
<i>Architecture</i>	
HGT layers	3
HGT hidden dim.	128
HGT attention heads	4
GRU hidden dim.	128
β -VAE latent dim. (d_z)	32 (= 128 B fp32 payload)
β (KL weight)	0.001
Candidate-set size K	8
ILCP gating MLP	2 layers, 128 hidden, ReLU
Inbound dropout (training)	0.2
<i>Training</i>	
Optimizer	AdamW
Learning rate	3×10^{-4}
Weight decay	10^{-4}
Batch size (UEs \times steps)	32×16
Epochs (max)	80
Early-stopping patience	8
Validation split	15% of trajectories
<i>Mixed-perturbation training</i>	
Mixture weight (clean : impaired)	0.5 : 0.5
Shadow fading σ_s	$\sim \mathcal{U}\{6, 12\}$ dB
Quantization step	1 dB
Reporting delay	40 ms
L3 IIR filter coefficient k	4 (per 3GPP TS 36.331)
<i>Evaluation</i>	
Bootstrap resamples	1000
CI level	95%
Inference hardware	1 \times NVIDIA GTX 1080 (8 GB), Intel i7-8700K, 16 GB RAM

C. Robustness under Measurement Impairments

In a real deployment, the measurements used for HO decisions may be inaccurate or variable due to shadowing, blockage, reporting delay, filtering, and incomplete measurement sampling. We model these effects as measurement perturbations and evaluate how well each method recovers a fixed reference target (which may be assumed as “optimal” in some sense) when the input measurements are degraded. This setting is relevant for operational RAN deployments because UEs should be handed over to optimal serving cells, rather than to cells that appear preferable only because of short-lived measurement fluctuations. To this end, we perturb the test-split measurements along three axes that commonly affect radio measurements in real deployments. We add shadow fading with σ_s from 0 to 12 dB, including quantization, reporting delay, and the L3 Long-Term Evolution (LTE) radio resource management (RRM) filter. We also add blockage events from 0 to 32 per trace and SSB-burst sub-sampling with periods from 1 to 16 steps. As the supervised reference target label, we keep the target fixed to the logged serving cell observed on the original unperturbed reference trace. The experiment therefore measures robustness to measurement perturbations while holding the reference serving-cell sequence fixed.

Table 3 summarizes the change in HOF from the unperturbed trace to the strongest perturbation setting. Figure 3 shows the shadow-fading sweep on the Vienna 4G/5G data-set. On the unperturbed trace, A3/A5 has very low HOF, as expected, because the reference handover labels are aligned with the rule-based handover behavior. When shadow fading is added to the measurement stream, its HOF increases sharply from 1.1% to the 57–65% range. Robustly trained ILCP, in contrast, remains in the 10–13% range across the sweep. The Zero-Knowledge baseline is also relatively stable, but at a higher 20–26% HOF range.

Table 3. HOF (%) under realistic measurement impairments: shadow fading $\sigma_s \in \{0, 3, 6, 9, 12\}$ dB (Noise), NLOS blockage events $\in \{0, 4, 8, 16, 32\}$, and SSB-burst sub-sampling period $\in \{1, 2, 4, 8, 16\}$. Each entry: HOF on clean trace \rightarrow HOF at the worst level.

Dataset	Method	Noise	Blockage	SSB
Vienna	ILCP (ours)	11.7 \rightarrow 12.3	11.7 \rightarrow 13.4	11.7 \rightarrow 14.5
	ZK-HGT	26.3 \rightarrow 25.1	26.3 \rightarrow 27.9	26.3 \rightarrow 29.6
	A3/A5	1.1 \rightarrow 65.4	1.1 \rightarrow 57.5	1.1 \rightarrow 72.6

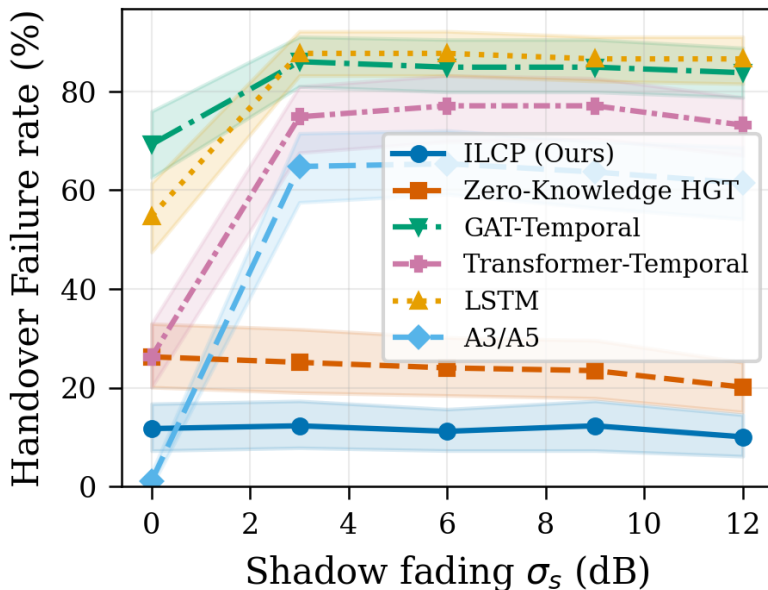


Figure 3. Vienna handover failure rate under shadow-fading perturbations. Robustly trained ILCP remains nearly flat across $\sigma_s \in [0, 12]$ dB, while the A3/A5 rule degrades sharply once nonzero shadow fading is added.

D. Bootstrap procedure

All confidence intervals reported in the paper are non-parametric percentile-bootstrap CIs (Tibshirani & Efron, 1993) over the test-split handover events. Let E be the set of test-split handover events ($|E|=31$ for Vienna). For each of $B=1000$ bootstrap rounds we sample $|E|$ events with replacement from E , recompute every reported metric on the resampled set, and store the values. The reported point estimate is the mean over the resamples; the lower and upper CI edges are the 2.5th and 97.5th percentiles. For metrics that are pure proportions (HOF, ping-pong rate, accuracy at a single step) a binomial CI would also be appropriate; we use the bootstrap uniformly so that derived metrics (averages over post-handover windows, ratios between methods on the same resample) are estimated under the same procedure.