

Jamming-Resilient Handover Triggering for Programmable 6G Radio Access Networks using Reinforcement Learning

George Amponis, *Graduate Student Member, IEEE*, Panagiotis Radoglou-Grammatikis, *Member, IEEE*, Antonios Sarigiannidis, Georgios Kakamoukas, Thomas Boufikos, Thomas Lagkas, *Senior Member, IEEE*, Vasileios Argyriou, Theofano Kollatou, and Panagiotis Sarigiannidis, *Member, IEEE*

Abstract—Resilient operation under abrupt Radio Frequency (RF) disruption is critical for next-generation cellular networks. Conventional static Event-trigger hand-over (HO) logic creates a fatal race condition under jamming: the handover command, triggered by slow-moving Reference Signal Received Power (RSRP) metrics, arrives too late to be decoded by a User Equipment (UE) whose Signal-to-Interference-plus-Noise Ratio (SINR) has already collapsed, resulting in radio-link failure. This paper proposes a 3rd Generation Partnership Project (3GPP)-conformant HO controller that instead of the fixed rule at the base-station mobility layer uses a Reinforcement Learning (RL) policy. The agent examines a reduced state vector, discretized SINR, serving-to-neighbour signal difference, and Hybrid Automatic Repeat reQuest (HARQ) error density, and outputs a binary trigger-or-defer action; no Radio Resource Control (RRC) or core-network signalling is changed. Realized in the LENA extension of ns-3 and tested in a multi-cell scenario with on-demand interference, the controller maintains connection continuity without extra signalling, computational overhead or ping-pong behaviour. Since its interface is restricted to vendor-agnostic Key Performance Indicators (KPIs) and a one-bit action, the mechanism can be encapsulated as an Open Radio Access Network (O-RAN) near-real-time xApp and migrated as-is to AI-native mobility functions anticipated in the future radio architectures.

Index Terms—O-RAN; Programmable Wireless Networks; RAN Intelligent Controller (RIC); xApp; 6G Architecture; Reinforcement Learning.

George Amponis is with K3Y Ltd, Studentski District, Vitosha Quarter, Bl. 9, 1700 Sofia, Bulgaria, and the Dept. of Informatics, Democritus University of Thrace, Kavala, Greece (e-mail: gamponis@{k3y.bg, cs.duth.gr}). Panagiotis Radoglou-Grammatikis is with K3Y Ltd, Studentski District, Vitosha Quarter, Bl. 9, 1700 Sofia, Bulgaria, and the Dept. of Electrical and Computer Engineering, University of Western Macedonia, Campus ZEP Kozani, Kozani, 50100, Greece (e-mail: pradoglou@{k3y.bg, uomw.gr}). Antonios Sarigiannidis is with K3Y Ltd, Studentski District, Vitosha Quarter, Bl. 9, 1700 Sofia, Bulgaria (e-mail: asarigia@k3y.bg). Georgios Kakamoukas is with K3Y Ltd, Studentski District, Vitosha Quarter, Bl. 9, 1700 Sofia, Bulgaria (e-mail: gkakamoukas@k3y.bg). Thomas Boufikos is with K3Y Ltd, Studentski District, Vitosha Quarter, Bl. 9, 1700 Sofia, Bulgaria (e-mail: tboufikos@k3y.bg). Thomas Lagkas is with the Dept. of Informatics, Democritus University of Thrace, Kavala Campus, Greece (e-mail: tlagkas@cs.duth.gr). Vasileios Argyriou is with the Dept. of Networks and Digital Media, Kingston University, London, UK (e-mail: vasileios.argyriou@kingston.ac.uk). Theofano Kollatou is with the Dept. of Electrical and Computer Engineering, University of Western Macedonia, Campus ZEP Kozani, Kozani, 50100, Greece (e-mail: tkollatou@uowm.gr). Panagiotis Sarigiannidis is with the Dept. of Electrical and Computer Engineering, University of Western Macedonia, Campus ZEP Kozani, Kozani, 50100, Greece (e-mail: psarigiannidis@uowm.gr).

I. INTRODUCTION

Resilience is one of the major security requirements of next-generation cellular systems because mission-critical services depend on in-service wireless connectivity in the demanding mobility requirements of emerging applications, which in turn, create a resilience gap [1]. Recent work has shown that reactive jammers can collapse the SINR of a serving link within a few milliseconds, forcing Radio Link Failures (RLFs) and disrupting sessions before any recovery procedure is completed [2]. The vulnerability is rooted in the design of 3GPP mobility triggers: the ubiquitous Event A3 initiates an HO only after a neighbour cell's RSRP has exceeded that of the serving cell for a configured Time-to-Trigger (TTT) interval [3]. In the case of rapid SINR collapse, there is a severe resilience gap as the A3/TTT rule responds after it is too late.

We are attempting to fill this gap with having an AI-native HO controller, which logically and functionally replaces the set threshold in the gNodeB mobility layer with RL policy. This paper contains a number of substantial contributions to this challenge. To this end we propose and deploy a new inline RL handover trigger, directly embedded into the standard `DoReportUeMeas` callback, and evaluate it on a reproducible and closed-loop simulation platform integrating the ns-3 NR/LENA stack [4] and an external Python Q-learning agent. The best is that all the abovementioned enhancements of the mechanism of control over HO can be embraced without changing 3GPP protocols. We also give a security-aware assessment against a wideband jamming attacker and demonstrate that our agent performs better in terms of number of successful HOs and link resilience than an optimized static baseline. Finally, we outline a clear deployment path for our agent as a near-real-time xApp within the O-RAN architectural vision, aligning our research with the industry's trajectory towards RICs [5].

When considering future 6G networks, the idea of a handover is also transforming, and it is particularly so in cell free and user centric networks. In these paradigms, the discrete HO event is replaced by a continuous process of dynamically managing a UE's active set of cooperating access points. This becomes exceptionally challenging in high-mobility scenarios,

such as vehicle-to-everything (V2X) or drone communications, where the network must make predictive, ultra-low-latency decisions to maintain a seamless connection. The provided AI-native solution provides a blueprint of the systems in the future. The principle of using a lightweight, learning-based agent to interpret real-time KPIs and issue agile control commands is directly applicable, whether the action is a traditional HO or a dynamic update to a UE's serving cluster in a cell-free environment.

II. BACKGROUND AND RELATED WORK

A. 3GPP Handover Procedures and Limitations

5G NR is based on UE measurement reports; Event A3 is raised when RSRP of a neighbour cell is higher than the serving cell by a configurable offset over a TTT interval [3]. Demonstration of parameters indicate that wrong offset/TTT pairs results in ping-pong HOs or late RLFs in ultra-dense deployments [6]. Since A3 measures filtered RSRP, it is not sensitive to non-stationary interference that is sudden. This forms a critical vulnerability since a rapid SINR stall may lead to the handover command arriving late leading to race condition resulting in RLF. This intrinsic timing discrepancy between the slow RSRP-based trigger and the fast channel degradation is the main issue which our work deals with.

B. Handover Optimization Research

The majority of HO-optimization research is focused upon Quality of Service (QoS) or load balancing. The most recent surveys take into account the AI-based approaches to ultra-dense networks with the attention to fuzzy logic and supervised learning in order to achieve the balance in the load [7]. More recent work has focused on Deep Reinforcement Learning (DRL). For example, the authors of [8] propose a Proximal Policy Optimization (PPO) agent to adapt handover protocols, focusing on improving data rates and reducing failures for UEs at different speeds. Similarly, Kwong et al. [9] employ a Deep Deterministic Policy Gradient (DDPG) agent to dynamically adjust the Handover Margin (HOM) in Ultra-Dense Networks. The other DRL-based methods have been focusing on energy efficiency [10] or reliable operation in harsh mmWave channels [11].

While powerful, these DRL methods often rely on complex architectures that can be difficult to train and act as "black boxes," making their decision-making process opaque. Our measurements are intentionally not in line with such models so that a baseline of jamming resilience can be set in a comprehensible manner. We show that in the context of this problem, a lightweight and highly interpretable tabular Q-learning agent is not only sufficient, but also desirable.

C. Security-Oriented Mobility and Research Gap

Active adversarial jamming is normally countered at the PHY layer through such methods as beam-nulling or frequency hopping. Nevertheless, a risk of disruption of the control-plane is considerable. The work by Lichtman et al. [12] provides a foundational threat assessment for 5G NR, identifying the

Physical Broadcast Channel (PBCH) and synchronization signals as key vulnerabilities to jamming and spoofing. They verify that despite the architecture enhancements in 5G, jamming attacks can occur and force UEs that cannot access a cell or decode needed system information to move. Although the threat is well-documented, there is little to no work done on protocol-conformant, learning-based handover schemes that expressly address this form of active interference; this paper goes some way to fill that particular gap.

D. Portability to Open RAN

The evolution towards disaggregated and intelligent radio access networks, standardized by the O-RAN Alliance [5], provides a clear deployment path for our mechanism. Our framework serves as a direct functional prototype of this architecture: the Python agent contains the core logic that would be packaged as a near-real-time xApp, the inline C++ hook in the gNB mirrors the role of a standardized E2 Agent, and our TCP socket communication represents the function of the official E2 interface.

However, our work takes this concept a step further: while much of the O-RAN discussion focuses on generalized optimizations for metrics like load balancing or energy efficiency, we provide a concrete blueprint for a highly specialized, security-focused xApp designed to address the issues introduced by jamming-induced false handover triggering. We demonstrate how the RIC can move beyond QoS improvements to host active, AI-driven defense mechanisms that respond to PHYa layer threats in milliseconds. This provides a tangible example of how the O-RAN architecture enables a new class of resilient applications, contributing to the broader 6G vision of a zero-touch, self-defending, and autonomous network [13].

III. PROPOSED MECHANISM

Adversarial jamming exposes a fatal race condition inherent in the 3GPP handover mechanism [14]. The conventional Event A3 trigger, designed for stability, relies on time-averaged RSRP measurements and a TTT delay. This logic is too slow to react to a wideband jammer, which does not significantly alter the slow-moving RSRP but causes a near-instantaneous collapse in the SINR. Consequently, by the time the static A3 rule is satisfied and the RRCR reconfiguration (handover command) message is sent, the UE's SINR has often fallen below the threshold required for successful decoding, resulting in a Radio Link Failure. Our proposed mechanism mitigates this threat by replacing the static rule with a learned, proactive policy. Instead of relying solely on RSRP, our agent observes a state vector composed of faster, more immediate indicators of link distress: the instantaneous SINR and the density of Hybrid Automatic Repeat reQuest (HARQ) NACKs. Through reinforcement learning, the agent learns to recognize the early signatures of a jamming-induced SINR collapse and to treat them as predictors of an impending link failure. This enables it to issue a handover trigger preemptively, acting within the critical window where the radio link is still viable

enough to successfully deliver the handover command to the UE, thereby winning the race condition and preserving the connection.

A. System Overview

The operational flow of our RL-driven handover mechanism is illustrated in the sequence diagram in Figure 1. The system has been designed as a closed loop with intelligent policy engine built into the 3GPP mobility as it is and without interfering with external signaling.

The sequence begins when the serving gNB receives a standard RRC MeasReport from a UE, containing measurements such as RSRP and SINR (1). Rather than this being readily judged against a fixed rule, a custom C++ hook intercepts the report. It builds a reduced state vector, s , composed of discretized SINR, the serving-to-neighbor RSRP difference, and recent HARQ NACK density [15] (1.2). This state is then serialized and sent via a TCP socket to the external Python agent (2).

The Python agent, which maintains a Q-table, performs a lookup on the received state and immediately returns a binary action (either `defer` or `trigger`) back to the gNB (3). It is based on this decision that further reaction is determined. If the action is `defer`, the gNB takes no mobility action and continues with normal MAC scheduling (4a). If the action is `trigger`, the gNB initiates the standard 3GPP handover procedure (4b) by sending an HO Request to the target gNB over the Xn/NGAP interface (5b) and, upon acknowledgment, issuing the RRC Reconfiguration command to the UE (7); to enable learning during training runs, the framework includes a reward feedback loop. After an action’s outcome is observed (e.g., a successful handover, a radio-link failure, or a ping-pong event), the serving gNB computes a corresponding reward r (10) and sends the new state s' , the reward, and the episode status back to the agent for its Q-table update (11). Most importantly, this whole process of decision-making is within the serving gNB. The air interface or core network does not carry any proprietary messages.

B. State, Action, and Reward Design

State: The agent observes a three-tuple

$$s = \{\text{SINR}_{\text{bin}}, \Delta\text{RSRP}_{\text{bin}}, \text{NACK}_{\text{bin}}\},$$

where each component is discretized into two or three levels. The selection is intentionally vendor-agnostic: all fields are already present in 3GPP measurement reports and require no gNB firmware change.

Action: The action space is minimal,

$$\mathcal{A} = \{\text{defer}, \text{trigger}\},$$

mapping directly to the presence or absence of an HandoverRequest. This binary interface is sufficient to replace the static Event-A3/A5 rule and remains compatible with any future mobility controller.




Fig. 1: Sequence diagram of the RL-driven handover process, showing the interaction between the UE, serving gNB, Python agent, and target gNB.

Reward hierarchy: The reward signal encodes operational priorities:

- **Radio-link failure (RLF):** large negative reward—connectivity preservation is paramount.
- **Ping-pong hand-over:** moderate penalty—stability is desirable but secondary.
- **Successful, stable HO:** positive reward—encourages proactive mobility.
- **Healthy defer decision:** small positive reward—avoids over-triggering when the link remains sound.

A grid search over a $\pm 30\%$ scaling of these weights produced less than five-percent variance in HO-failure rate, indicating policy robustness to hyper-parameter choice.

C. Implementation Details

Inline C++ hook: Listing 1 sketches the modified `DoReportUeMeas()` handler. Only ten lines are required beyond serialization; all RRC, NGAP, and core-network procedures remain intact.

Inter-process communication: The `RlSocketComm` class opens an ephemeral TCP server at simulation start; the Python script connects and exchanges newline-terminated ASCII messages. Throughput never exceeds a few kilobytes per second, so no optimization is required.

Standards compliance: Because the mechanism merely alters when an existing RRC message is generated—and

Algorithm 1 Inline hand-over decision hook

```

0:  $s \leftarrow \text{BUILDSTATE}(\text{MeasReport})$ 
0:  $\text{SENDTOAGENT}(s)$ 
0:  $a \leftarrow \text{RECVACTION}()$ 
0: if  $a = \text{trigger}$  then
0:    $\text{SENDHANDOVERREQUEST}()$ 
0: return  $=0$ 

```

neither its payload nor the NG core workflow—the proposal is fully compliant with 3GPP TS 38.331 and TS 38.413. Consequently, the same binary action can be transported via the O-RAN E2 interface or mapped to a Service-Based Interface in foreseen 6G control planes.

IV. EVALUATION

A. Experimental Setup

The NR/LENA component of *ns-3* is used to run experiments [4]. There are four gNBs at four corners of a square of 1-km length. A jammer node is located in the middle and injects wideband Gaussian noise during a configurable time interval in order to represent a PHY layer attack. The mobility pattern of UE instances is random-walk which repeatedly crosses cell boundaries. Each UE sends constant-bit-rate UDP traffic to a remote host in the EPC core and generates a continuous stream of user-plane packets. The RL is trained through 1,000 episodes, each having 60 seconds. A grid search over learning rate $\alpha \in \{0.05, 0.1, 0.2\}$ and discount factor $\gamma \in \{0.8, 0.9, 0.95\}$ identifies a robust operating point; ϵ -greedy exploration decays linearly from 0.5 to 0.05. All the averages are based on 30 independent seeds; confidence intervals are calculated within the Student *t*-distribution and are shown as areas of gray in the plot.

B. Baselines

- **Optimized Static A3:** An exhaustive sweep over Event-A3 offset $O \in \{0, \dots, 6\}$ and Time-to-Trigger TTT $\in \{0, \dots, 320\}$ selects the [offset, TTT] pair that minimises HO failures under our specific wideband noise profile.
- **Reactive Power Ramping:** Upon a sharp SINR drop, the UE increases its uplink power by 3 dB for 200 ms before reverting, representing a non-mobility-based mitigation heuristic.

C. Metrics

- 1) **Handover Failure Probability:** Fraction of HO attempts that result in a Radio Link Failure (RLF) before the handover is complete.
- 2) **UE SINR CDF:** The empirical Cumulative Distribution Function of SINR samples taken at the UE every 10 ms.
- 3) **Ping-Pong Probability:** Fraction of handovers that are followed by a return handover to the previous cell within 2 seconds.




Fig. 2: Handover Resilience vs. Wideband Gaussian Noise Power. The RL Agent shows graceful degradation compared to the brittle failure of static baselines.

D. Results

The resilience of each algorithm against wideband Gaussian noise is presented in Figure 2. The results demonstrate that the static baselines are brittle. The Optimized Static A3 controller, while effective at very low noise power, exhibits a "cliff-edge" behavior, with its failure rate rising sharply beyond 25 dBm. The Reactive Power Ramping heuristic offers a marginal improvement but also fails to cope with increasing interference. In contrast, the RL Agent demonstrates superior robustness, maintaining a low failure rate that degrades gracefully even under high-power (40 dBm) noise injection. The non-overlapping confidence intervals at higher power levels confirm the statistical significance of this resilience gain.

The reason for the RL agent's superior resilience is revealed by analyzing the UE SINR distributions. Figure 3 establishes a baseline in a clean channel environment, showing that all three algorithms maintain a similarly high SINR distribution, confirming the RL agent "does no harm" in normal operation. Under the wideband noise attack, however, Figure 4 shows a stark divergence. The Static A3 controller leaves the UE "stuck" in a poor radio environment, with the majority of its SINR samples falling below 0 dB. The RL agent, by contrast, successfully maintains a much healthier link for the UE by proactively triggering handovers, keeping the SINR in a range where control commands can still be successfully decoded.

Finally, we evaluate whether the agent's proactivity leads to network instability. Figure 5 plots the ping-pong probability against increasing noise power. The results show that the RL Agent's stability is statistically indistinguishable from the highly conservative Optimized Static A3 baseline at low-to-moderate interference levels. The initial dip in ping-pong rate around 15-20 dBm suggests a "decision-freezing" effect, where moderate, uniform interference removes the cell-edge




Fig. 3: User SINR Distribution in a baseline scenario with no external interference, showing comparable performance across all algorithms.




Fig. 4: User SINR Distribution under the Wideband Gaussian Noise scenario. The RL Agent successfully maintains a healthier link quality for the UE.

ambiguity that typically causes ping-pong. As noise power increases further, the RL agent’s stability degrades gracefully and remains well within acceptable operational limits, while the Power Ramping heuristic becomes increasingly unstable. This confirms that the RL agent achieves its resilience gains without a significant stability trade-off.

V. DISCUSSION

Our evaluation demonstrates that a lightweight Reinforcement Learning agent can significantly enhance handover resilience against wideband noise jamming, outperforming even an optimized static baseline. This section interprets these find-




Fig. 5: Handover Stability vs. Wideband Gaussian Noise Power. The RL Agent maintains stability comparable to the static baseline across all noise intensities.

ings, contextualizes the trade-offs, and discusses the broader implications for future network architectures.

A. The Emergence of a Proactive, State-Aware Policy

The core finding of this work is not just that the RL agent performs better, but *how* it does so. The results reveal the emergence of a sophisticated, proactive policy that correctly balances the competing demands of resilience and stability. The SINR distribution in Figure 4 is the key evidence: the RL agent learns to treat a rapid SINR drop, corroborated by HARQ NACKs, as a leading indicator of an impending Radio Link Failure. It learns to win the “race condition” by initiating a handover *before* the link is too degraded to transmit the command—a state the slow, RSRP-based A3 logic is blind to.

This proactive nature, however, is intelligently controlled. As shown in Figure 5, the agent does not become “trigger-happy.” It learns that there is a penalty for instability (ping-pong), and therefore develops a nuanced policy: be aggressive only when the risk of link failure is high and a viable neighbor exists. The initial dip in ping-pong rate under moderate interference further suggests the agent learns a sophisticated behavior: in a chaotic but not yet catastrophic environment, it becomes more conservative to avoid making a bad situation worse, a dynamic that simple, static rules cannot replicate.

B. On the Sufficiency of a Compact State Representation

One of the design decisions was a minimalist state vector with 3 features. The results justify this decision. The selected KPIs (instantaneous link quality (SINR), the availability of a better alternative (Δ RSRP), and immediate user-plane impact (NACKs)) provide an orthogonal basis for making resilience-oriented decisions. KPI minimalism ensures that the inputs are vendor-neutral KPIs that can easily be found in any system that

complies with 3GPP. In case of high-mobility use cases a UE speed feature would be added and in case of managing network congestion a target cell load indicator would be required. Such can be added to the state vector without disturbing the essential protocol-compliant character of the control loop.

C. Security Scope and the Path to Deployment

Our evaluation focuses on a PHY layer availability attack (wideband jamming). The agent's learned policy is to "flee" a compromised radio environment. The strategy would not (by default) work against other threats, like a rogue gNB attack where a cell transmits a powerful unlawful signal. In this situation, the current agent would be misled to make a hazardous handover. Future effort would focus on combining the state vector with security-related KPIs, such as trust scores or authentication flags of higher-layer security functions. In the end, the most practical route of deployment of this mechanism is as an xApp in an O-RAN RIC. Our system's design is a direct analogue to the RIC paradigm: the C++ hook is the E2 Agent, the TCP socket is the E2 interface, and the Python script is the xApp.

VI. CONCLUSION AND FUTURE WORK

Our research has introduced a 3GPP compliant, RL-based handover trigger that mitigates the race condition which can lead to a fatal attack by wideband jamming. The mechanism is trained to predict link degradation based on fast-changing indicators such as SINR and HARQ errors, by replacing the slower, RSRP-based Event A3 logic. This enables it to be proactive over the limited time window before the communication link fails, but still be capable of delivering the handover command successfully. Our experiments demonstrate that the method provides a large decrease in the number of radio-link failures across a variety of realistic attack scenarios including portable, low-power jammers (such as those employed in vehicle-mounted systems), without compromising the handover stability of the network, and at a minimal compute overhead cost.

The tabular Q-policy shows that even a very lightweight, protocol-compatible solution can constitute an initial line of defense against PHY layer availability attacks. The closed-loop ns-3-Python simulation framework, baselines, and the result data are published to the community in order to facilitate future work and reproducibility. Future development will be carried out on three main axes. The tabular Q-agent will be substituted by a lightweight Deep Q-Network to prove its scalability and performance in complex, multi-UE and multi-gNB environments, and the end goal will be to prototype and test on a live O-RAN RIC testbed. We will extend the agent's state space to include UE velocity, cell-load metrics, and higher-layer trust indicators to develop policies that are resilient not only to jamming but also to more sophisticated threats (e.g., rogue gNBs). Finally, we will investigate the policy's effectiveness on emerging 6G propagation regimes, including THz bands, reconfigurable intelligent surfaces, and cell-free deployments.

ACKNOWLEDGEMENT

This work has received funding from the European Union and the Swiss State Secretariat for Education, Research and Innovation (SERI) under the grant agreement No 101192749 (XTRUST-6G). Disclaimer: Funded by the European Union. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

REFERENCES

- [1] I. Parvez, A. Rahmati, I. Guvenc, A. I. Sarwat, and H. Dai, "A Survey on Low Latency Towards 5G: RAN, Core Network and Caching Solutions," *IEEE Communications Surveys Tutorials*, vol. 20, no. 4, pp. 3098–3130, 2018.
- [2] J. R. Stegmann, M. Gundall, and H. D. Schotten, "Smart PRACH Jamming: A Serious Threat for 5G Campus Networks," in *Proc. IEEE Globecom Workshops*, 2024, arXiv:2410.08729.
- [3] "TS 38.331 V17.10.0: NR; Radio Resource Control (RRC); Protocol Specification," Tech. Rep., Mar. 2024.
- [4] M. Mezzavilla, M. Zhang, M. Polese, R. Ford, S. Dutta, S. Rangan, and M. Zorzi, "End-to-End Simulation of 5G mmWave Networks," *IEEE Communications Surveys Tutorials*, vol. 20, no. 3, pp. 2237–2263, 2018.
- [5] A. Arnaz, J. Lipman, M. Abolhasan, and M. Hiltunen, "Toward Integrating Intelligence and Programmability in Open Radio Access Networks: A Comprehensive Survey," *IEEE Access*, vol. 10, pp. 67747–67770, 2022.
- [6] A. Peltonen, R. Sasse, and D. Basin, "A Comprehensive Formal Analysis of 5G Handover," in *Proc. ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec)*, 2021, pp. 1–14.
- [7] C. Chabira, I. Shayea, G. Nurzhaubayeva, L. Aldasheva, D. Yedilkhan, and S. Amanzholova, "AI-Driven Handover Management and Load Balancing Optimization in Ultra-Dense 5G/6G Cellular Networks," *Technologies*, vol. 13, no. 7, 2025. [Online]. Available: <https://www.mdpi.com/2227-7080/13/7/276>
- [8] P. J. Gu, J. Voigt, and P. M. Rost, "A Deep Reinforcement Learning-based Approach for Adaptive Handover Protocols in Mobile Networks," 2024. [Online]. Available: <https://arxiv.org/abs/2401.14823>
- [9] C. F. Kwong, C. Shi, Q. Liu, S. Yang, D. Chieng, and P. Kar, "Autonomous Handover Parameter Optimisation for 5G Cellular Networks Using Deep Deterministic Policy Gradient," *Expert Systems with Applications*, vol. 246, p. 122871, 2024.
- [10] H. Ju, S. Kim, Y. Kim, and B. Shim, "Energy-Efficient Ultra-Dense Network with Deep Reinforcement Learning," 2021. [Online]. Available: <https://arxiv.org/abs/2112.13189>
- [11] M. Chiputa, M. Zhang, G. G. M. N. Ali, P. H. J. Chong, H. Sabit, A. Kumar, and H. Li, "Enhancing Handover for 5G mmWave Mobile Networks Using Jump Markov Linear System and Deep Reinforcement Learning," *Sensors*, vol. 22, no. 3, 2022. [Online]. Available: <https://www.mdpi.com/1424-8220/22/3/746>
- [12] M. Lichtman, R. Rao, V. Marojevic, J. Reed, and R. P. Jover, "5G NR Jamming, Spoofing, and Sniffing: Threat Assessment and Mitigation," in *2018 IEEE International Conference on Communications Workshops (ICC Workshops)*, 2018, pp. 1–6.
- [13] V. T. Kim Anh, "The Rise of AI in 6G Networks: A Comprehensive Review of Opportunities, Challenges, and Applications," in *2024 International Conference on Advanced Technologies for Communications (ATC)*, 2024, pp. 333–338.
- [14] G. Asemian, M. Kulhandjian, M. Amini, B. Kantarci, C. D'Amours, and M. Erol-Kantarci, "The Impact of Mobility, Beam Sweeping and Smart Jammers on Security Vulnerabilities of 5G Cells," *Wireless World Research and Trends Magazine*, p. 71–78, Dec. 2024. [Online]. Available: <http://dx.doi.org/10.13052/2794-7254.009>
- [15] B. Göktepe, C. Hellge, T. Schierl, and S. Stanczak, "Distributed Machine-Learning for Early HARQ Feedback Prediction in Cloud RANs," *IEEE Transactions on Wireless Communications*, vol. 23, no. 1, pp. 31–44, 2024.