

Performance Evaluation of an AI-Driven Adaptive Optimization Framework for 5G Mobile Networks

ATU Michael Osadebamwen¹, Guiawa Mathurine Guiawa¹, Ukagu Stephen Nwachukwu¹

¹ Department of Electronic and Computer Engineering, Igbinedion University Okada, Edo State, Nigeria

Corresponding Author: ATU Osadebamwen Michael

Abstract - Fifth-generation (5G) networks require adaptive resource management under changing radio and traffic conditions, but performance claims are meaningful only when the optimization procedure is reproducible and evaluated across repeated operating scenarios. This study evaluates an AI-triggered adaptive optimization framework using a controlled synthetic 5G snapshot model. A leakage-controlled XGBoost classifier identifies Critical, Moderate, and Optimal network states; optimization then searches transparent, state-conditioned actions that modify bandwidth allocation, interference level, and resource-utilization pressure. Unlike the earlier single-snapshot demonstration, no output KPI is multiplied by a preset improvement factor. Throughput, latency, packet loss, spectral efficiency, and energy consumption are recomputed from the original network equations using matched stochastic residuals. Across 20 independent evaluation seeds and 40,000 snapshots, throughput increased by 20.47%, latency decreased by 2.60%, packet loss decreased by 5.51%, and spectral efficiency increased by 5.41%. Absolute energy consumption increased by 2.94%, revealing a capacity-energy trade-off, while derived energy efficiency improved by 17.03%. Repeated-seed confidence intervals, state-disaggregated analysis, operating-condition robustness checks, action-bound sensitivity, ablation analysis, and computational timing are reported. The findings provide preliminary simulation-based evidence rather than a deployment-level performance claim.

Keywords: 5G networks; adaptive optimization; XGBoost; network performance evaluation; reproducible simulation; quality of service; energy efficiency

1. INTRODUCTION

Fifth-generation (5G) mobile communication networks support heterogeneous service classes with markedly different requirements for throughput, latency, reliability, spectrum use, and energy efficiency. Traffic demand, mobility, interference, and resource occupancy vary continuously, making purely static resource-management rules difficult to sustain across changing operating conditions. Artificial intelligence (AI) and machine learning (ML) have therefore been applied to traffic prediction, beam management, interference control, slicing, scheduling, and energy-aware resource allocation (Alzubaidi et al., 2022; Brilhante et al., 2023; Ezzeddine et al., 2024).

A central evaluation problem is that predictive accuracy alone does not establish network-level benefit. A classifier may identify degraded conditions correctly while the action triggered by that classification produces little improvement, creates an adverse trade-off, or benefits only a narrow operating regime. Conversely, a modest prediction component can still support useful control if the downstream decision mechanism is transparent, reproducible, and evaluated under matched before-and-after conditions. For this reason, a performance-evaluation paper must disclose not only the AI trigger but also how optimization actions are selected and how post-action QoS values are generated.

The literature already contains multi-objective and multi-KPI 5G optimization studies; therefore, the contribution of the present paper is not that multiple KPIs are considered for the first time. Primary studies have used deep reinforcement learning for 5G-TSN scheduling (Zhang et al., 2025), active reward learning for 5G resource allocation (Shahgholi et al., 2025), and ML-based dynamic resource reconfiguration for energy optimization in 5G NR (Yao & Pérez Yuste, 2025). These studies also illustrate an important comparison problem: reported gains depend strongly on topology, traffic model, action space, optimization objective, and simulator. Numerical superiority claims are therefore defensible only under like-for-like conditions.

This revised study is positioned more narrowly. It asks whether a state-triggered adaptive action-search procedure produces repeatable network-level changes within one fully disclosed synthetic snapshot model, and whether those changes remain consistent across random seeds, predicted network states, traffic/SINR/load regimes, and alternative action bounds. The paper also explicitly reports trade-offs rather than requiring all aggregate KPIs to improve simultaneously.

The study addresses three research questions. RQ1: Under repeated independent simulation seeds, does the disclosed adaptive action-search procedure improve throughput, latency, packet loss, and spectral efficiency relative to a null-action baseline? RQ2: What energy-consumption and energy-efficiency trade-offs accompany those changes? RQ3: How robust are the results across predicted network states, operating-condition tertiles, action-bound sensitivity, and restricted single-lever ablations?

The contribution is therefore an evaluation and reproducibility contribution: (i) full disclosure of the snapshot generator and stochastic terms; (ii) concise disclosure of the XGBoost state trigger and leakage control; (iii) an explicit objective and state-conditioned action space; (iv) repeated-seed paired evaluation using common random numbers; and (v) uncertainty, disaggregation, sensitivity, ablation, and runtime analyses. The study does not claim live-network validation, packet-level 5G compliance, or a new learning algorithm.

2. RELATED WORK

2.1 AI-Driven Resource Allocation and Optimization in 5G

AI-assisted 5G optimization spans several network functions. Alzubaidi et al. (2022) review interference-management challenges in beyond-5G design and show why interference, resource allocation, and heterogeneous deployment must be considered jointly. Brillhante et al. (2023) survey AI-aided beamforming and beam management and highlight the difficulty of cross-study benchmarking when datasets and evaluation procedures differ. Ezzeddine et al. (2024) similarly emphasize that energy-saving mechanisms should be evaluated together with QoS because energy reduction can degrade throughput, latency, or user experience.

Recent primary experimental work provides more direct comparators for the general problem, although not for the exact system model used here. Zhang et al. (2025) formulate and evaluate a deep-reinforcement-learning scheduling method for integrated 5G-TSN industrial networks. Shahgholi et al. (2025) propose active reward learning for dynamic 5G network-slicing resource allocation using multiple network indicators in the reward structure. Yao and Pérez Yuste (2025) use ML-based resource prediction to drive dynamic 5G NR cell-resource reconfiguration with explicit energy-oriented objectives. These studies confirm that multi-KPI optimization is established; they also use different topologies, actions, traffic models, and performance definitions, which precludes direct percentage-by-percentage comparison with the present abstract snapshot model.

2.2 Reproducibility and Evaluation Gap

The methodological gap addressed here is therefore not the absence of multi-objective optimization. It is the need to connect a state prediction, a disclosed action space, a disclosed objective, and recomputed network outputs in a way that another researcher can reproduce. A before/after table is insufficient if the post-action values are generated by undocumented adjustment factors, if only a single observation is shown, or if stochastic variation is not separated from the effect of the action.

The revised methodology addresses this by retaining the original analytical equations, storing the random residual terms, changing only the permitted control variables, and recomputing every post-action KPI under the same exogenous snapshot. Repeated independent seeds then characterize variability at the trial level. This design does not transform the abstract simulator into a packet-level network emulator, but it makes the evidence class and the source of the reported changes auditable.

2.3 Study Positioning

The paper is scoped to controlled engineering evaluation. The formal six-layer closed-loop architecture and broader policy-learning formulation belong to a companion system-design study, while classifier development is treated in a companion predictive-classification study. Enough detail from both components is reproduced here to make the performance experiment independently understandable: the classifier features, split and hyperparameters are stated; the optimization objective and action search are specified; and all output-generation equations are disclosed. Companion studies provide additional depth but are not required to understand how the revised results in this paper were obtained.

3. Materials and Methods

3.1 Study Design and Simulation Scope

The experiment uses a controlled analytical snapshot model rather than a packet-level simulator. Each observation represents one independent cell/network-state snapshot defined by user density, traffic demand, interference, received signal power, bandwidth, resource utilization, and mobility. There is no explicit multi-cell geometry, carrier-frequency-specific propagation model, fading trace, packet scheduler implementation, simulation duration, or time step. Channel and scheduling effects are represented at the KPI level through SINR, congestion, resource-utilization variables, scheduler-efficiency features, and stochastic residuals. This boundary is stated explicitly because inserting unimplemented topology or fading parameters would create false reproducibility rather than improve it.

A development dataset of 10,000 snapshots is generated with seed 42 for target construction and classifier training. The revised network-performance experiment is then conducted on 20 independent evaluation datasets generated with seeds 101–120, each containing 2,000 snapshots, for a total of 40,000 evaluation observations. The classifier is fixed after development; the independent evaluation seeds are not used for model selection.

3.2 Synthetic Snapshot Generation

The primary input distributions and ranges are summarized in Table 1. These values are controlled study ranges selected to span low-to-high load, signal, interference, bandwidth, and mobility conditions; they are not presented as universal 3GPP bounds.

The mobility range up to 120 km/h covers the vehicular range commonly discussed in IMT-2020 contexts, but the present simulator does not implement an IMT evaluation environment (ITU-R, 2017).

For each snapshot, SINR is computed as $\text{SINR} = S - I$, where S and I are signal and interference powers in dBm. Congestion is $C = 0.45(U/1000) + 0.35(D/1000) + 0.20(R/100)$, where U is user density, D is traffic demand in Mbps, and R is resource utilization in percent. Spectral efficiency is $\text{SE} = \text{clip}\{\log_2[1 + 10^{(\text{SINR}/10)}](1 - 0.35C), 0.1, 8.0\}$.

The baseline QoS outputs are then generated as follows: throughput $T = \text{clip}\{B \cdot \text{SE} \cdot (1 - 0.40C) + \varepsilon_T, 1, \infty\}$; latency $L = \text{clip}\{5 + 75C + 0.12V - 0.35\text{SINR} + \varepsilon_L, 1, \infty\}$; packet loss $\text{PL} = \text{clip}\{0.2 + 7C - 0.08\text{SINR} + 0.015V + \varepsilon_{\text{PL}}, 0, 20\}$; and energy $E = \text{clip}\{40 + 0.06U + 0.09D + 0.45B + \varepsilon_E, 20, \infty\}$. The residuals are $\varepsilon_T \sim N(0, 8^2)$, $\varepsilon_L \sim N(0, 4^2)$, $\varepsilon_{\text{PL}} \sim N(0, 0.5^2)$, and $\varepsilon_E \sim N(0, 5^2)$. The same residual realization is retained for the matched post-action recomputation of each snapshot.

Seven additional radio/resource features are generated for the classifier: RSRP, RSRQ, CQI, PRB utilization, cell load, handover success rate, and scheduler efficiency, using the same equations and fixed generator sequence as the companion classification study. The complete executable generator is supplied in the accompanying reproducibility notebook.

3.3 Network-State Classification Trigger

Network state is represented by Critical, Moderate, and Optimal labels derived from the Adaptive Optimization Score (AOS). AOS is computed on a 0–100 scale using weights 0.30 throughput, 0.25 latency, 0.15 packet loss, 0.20 spectral efficiency, and 0.10 energy consumption, after normalization by development-dataset reference maxima. $\text{AOS} \geq 80$ is labelled Optimal, $60 \leq \text{AOS} < 80$ Moderate, and $\text{AOS} < 60$ Critical.

To reduce direct target leakage, the five QoS quantities used to construct AOS and AOS itself are removed from classifier inputs, leaving 15 radio/resource predictors. XGBoost is trained on a stratified 80:20 development split with 300 estimators, maximum depth 4, learning rate 0.05, subsample 0.90, column subsample 0.90, and `random_state = 42` (Chen & Guestrin, 2016). The held-out accuracy is 96.45%. The corresponding confusion matrix is reproduced in Table 2. In the companion classifier validation, stratified five-fold mean accuracy is $95.36 \pm 0.44\%$, macro F1 is $90.96 \pm 0.87\%$, and balanced accuracy is $89.43 \pm 0.97\%$. This paper does not use classifier accuracy as evidence of optimization benefit; it uses the fixed classifier only to trigger the subsequent action search.

3.4 Optimization Objective, Decision Variables, and Action Search

The revised optimization stage no longer alters output KPIs directly. It operates on three quantities already present in the network model: allocated bandwidth B , interference power I , and resource-utilization pressure R . User density, traffic demand, mobility, and received signal power are held fixed within each before/after pair. Candidate changes are capped at the original generator boundaries: bandwidth ≤ 100 MHz, interference ≥ -110 dBm, and resource utilization $\geq 20\%$.

Candidate actions are evaluated with the same normalized AOS objective used to summarize network condition. For each candidate, the modified B , I , and R are inserted into the original SINR, congestion, spectral-efficiency, throughput, latency, packet-loss, and energy equations; no fixed target improvement percentage appears in the code. The candidate with the highest AOS is selected. If no candidate improves AOS, the null action is retained.

The primary state-conditioned action bounds are given in Table 3. Critical states permit a +20 MHz bandwidth option, up to 2 dB interference reduction, and up to 5 percentage points of resource-utilization relief. Moderate states use smaller interference and resource-relief actions and do not expand bandwidth. Optimal states use the null action. These are study-defined control bounds rather than standardized limits; Section 3.6 tests conservative and aggressive alternatives to quantify sensitivity to this design choice.

3.5 Baseline, Repeated Trials, and Paired Evaluation

The conventional baseline is defined operationally as the null action under the same synthetic snapshot: the original B , I , R and all exogenous quantities are retained. The optimization comparison is paired using common random numbers; the same ε_T , ε_L , ε_{PL} , and ε_E values that generated the baseline are reused after the action. Consequently, the before/after difference for a snapshot is attributable to the permitted control change rather than to a new random draw.

For each of the 20 evaluation seeds, XGBoost predicts the state of all 2,000 snapshots, the state-conditioned candidate search is executed, and trial-level mean values are calculated. The primary unit for uncertainty reporting is therefore the independent simulation seed, not the individual snapshot. Mean percentage change, standard deviation across the 20 trials, 95% t confidence intervals, and two-sided paired t-tests on per-seed before/after means are reported.

3.6 Robustness, Sensitivity, Alternative Optimizer, Ablation, and Instance-Level Analysis

Five additional analyses are used to avoid relying on aggregate means alone. First, results are disaggregated by predicted Critical, Moderate, and Optimal state. Second, traffic demand, SINR, and cell load are each divided into Low, Medium, and High within-trial tertiles to test whether the direction of the result is confined to one operating-condition band. Third, action-bound sensitivity compares conservative, primary, and aggressive control sets. Fourth, a controlled alternative-optimizer benchmark uses the identical state labels, candidate action space, constraints, and common-random-number recomputation as the primary method,

but selects the feasible action that maximizes throughput alone rather than the multi-objective AOS. Fifth, an ablation comparison restricts the AOS optimizer to bandwidth only, interference only, or resource-relief only under the same data and objective.

At the individual-snapshot level, each KPI is also classified as improved, unchanged, or deteriorated. This prevents the aggregate mean from being interpreted as evidence that every user or network state benefits. Energy efficiency is derived as $EE = \text{throughput} / \text{energy consumption (Mbps/W)}$.

3.7 Computational Timing and Standards Context

Software execution time is measured over 30 repeated batches of 2,000 snapshots after warm-up, separately for XGBoost inference and the optimization action search. The resulting timing is an implementation benchmark on the reported Python/CPU environment, not an end-to-end RAN control-loop latency claim; live deployment would additionally include telemetry, transport, controller, and actuation overhead.

Report ITU-R M.2410-0 specifies IMT-2020 minimum technical requirements including 100 Mbit/s downlink user-experienced data rate and 4 ms eMBB user-plane latency under defined evaluation conditions (ITU-R, 2017). The throughput and latency variables in this study are abstract snapshot outputs and are not measured using the ITU evaluation methodology; consequently, the revised paper does not claim IMT-2020 compliance or compute superiority percentages against those targets.

4. Results

4.1 Network-State Distribution

Across the 40,000 independent evaluation snapshots, XGBoost classified 26,947 (67.37%) as Critical, 11,687 (29.22%) as Moderate, and 1,366 (3.42%) as Optimal. Figure 3 summarizes this distribution. The imbalance is consistent with the AOS-derived development target and motivates reporting state-specific results rather than relying only on overall means.

4.2 Repeated-Seed Network Performance

Table 4 reports the principal revised results. Mean throughput increased from 231.02 to 278.28 Mbps, corresponding to a 20.47% mean increase across the 20 independent seeds (95% CI 20.18–20.75%). Mean latency decreased from 45.65 to 44.47 ms, a 2.60% reduction (95% CI 2.59–2.60%). Packet loss decreased from 3.312% to 3.129%, a 5.51% reduction (95% CI 5.50–5.53%), and spectral efficiency increased from 4.866 to 5.129 bps/Hz, a 5.41% increase (95% CI 5.33–5.48%). All four changes were consistent across seeds and the paired seed-level tests gave $p < 0.001$.

Absolute energy consumption moved in the opposite direction: 145.76 to 150.05 W, a 2.94% increase rather than a reduction (95% CI for 'improvement' –2.97 to –2.91%; $p < 0.001$). This is an explicit capacity-energy trade-off caused primarily by bandwidth expansion in Critical states. Despite higher absolute energy use, derived energy efficiency increased from 1.585 to 1.855 Mbps/W, a mean improvement of 17.03% (95% CI 16.78–17.28%), because throughput grew substantially more than energy consumption. Figures 1 and 2 visualize the normalized before/after values and trial-level changes.

4.3 State-Disaggregated Performance

Table 5 shows that the aggregate result is dominated by Critical-state interventions. Critical snapshots gained approximately 50.58% throughput, 2.77% lower latency, 5.37% lower packet loss, and 10.28% higher spectral efficiency, with a 4.25% increase in energy consumption. Moderate states obtained smaller throughput and spectral gains but still reduced latency and packet loss, while energy was unchanged because the primary Moderate action set did not permit bandwidth expansion. Optimal states retained the null action and therefore remained unchanged by design.

4.4 Robustness Across Traffic, SINR, and Cell Load

The operating-condition analysis in Table 6 shows that throughput, latency, packet loss, and spectral efficiency improved in every Low/Medium/High tertile of traffic demand, SINR, and cell load. The magnitude varied materially. For example, throughput gain was largest in the low-SINR tertile and smallest in the high-SINR tertile, which is consistent with more headroom for interference-control and capacity actions under poorer radio conditions. Energy consumption increased in each tertile because a subset of Critical snapshots selected the bandwidth-expansion action. These results support robustness of the direction of the principal QoS changes while also showing that effect size depends on operating condition.

4.5 Individual-Snapshot Trade-Offs

Table 7 prevents aggregate improvement from being interpreted as universal benefit. Throughput improved in 94.23% of snapshots and was unchanged in 5.77%; latency improved in 96.57% and was unchanged in 3.43%; packet loss improved in 93.13%; and spectral efficiency improved in 64.97%, with the remainder unchanged. None of these four metrics deteriorated under the selected action search because the candidate action had to increase AOS and the permitted controls are monotonic for those equations within the tested bounds. Energy consumption, however, deteriorated in 47.60% of snapshots and was unchanged in 52.40%; it never decreased under the current energy model because interference and utilization controls carry no explicit energy term while bandwidth expansion increases E. Derived energy efficiency improved in 94.23% of snapshots and did not deteriorate.

4.6 Action-Bound Sensitivity, Alternative Optimizer, and Ablation

Table 8 demonstrates that performance magnitude depends strongly on control authority. Under the conservative action set, throughput improved by only about 2.9%, whereas the aggressive set produced a gain above 50% together with a larger energy penalty. The primary policy lies between these extremes. This sensitivity is important: the chosen action bounds are transparent design parameters, not uniquely optimal constants.

The controlled alternative-optimizer comparison in Table 9 shows the trade-off created by the objective function itself. With the same feasible actions, the throughput-greedy strategy achieved a slightly larger throughput gain (20.59% versus 20.47%) but smaller latency reduction (2.25% versus 2.60%), smaller packet-loss reduction (4.55% versus 5.51%), a larger energy penalty (3.41% versus 2.94%), a lower AOS gain (2.782 versus 2.881 points), and a smaller energy-efficiency gain (16.62% versus 17.03%). Thus, the multi-objective AOS selector gives up approximately 0.13 percentage points of throughput gain in exchange for a more balanced outcome across the other modeled objectives. This is a controlled internal algorithmic benchmark rather than a claim of superiority over externally published systems with different action spaces and network models.

The restricted-strategy ablation in Table 10 shows that the full three-lever AOS action search yields a larger AOS gain than bandwidth-only, interference-only, or resource-relief-only variants. Bandwidth-only control produces most of the capacity gain but also the energy penalty; interference-only control improves throughput, latency, packet loss, and spectral efficiency without changing the present energy equation; resource-relief-only control has smaller but consistent effects. This decomposition makes the source of the aggregate result auditable.

4.7 Computational Timing

On the execution environment used for the reproducibility run, median XGBoost inference time for a 2,000-snapshot batch was approximately 10.73 ms and the vectorized action search required approximately 4.69 ms, for a combined batch time of 15.42 ms (Table 11). The corresponding arithmetic per-snapshot time is about 0.0077 ms, but this should not be interpreted as deployed decision latency because batch vectorization and the absence of network I/O make the computational benchmark fundamentally different from an operational RAN control loop.

5. Discussion

5.1 What the Revised Experiment Supports

The corrected experiment supports a narrower but stronger engineering conclusion than the earlier single-snapshot demonstration. The reported percentages now emerge from state-conditioned changes to underlying network/resource variables followed by recomputation of the QoS equations; no output KPI is assigned a fixed improvement multiplier. Repeated seeds show that the direction of the principal throughput, latency, packet-loss, and spectral-efficiency changes is stable in this controlled model.

The throughput gain is the largest effect because the Critical-state action set includes additional bandwidth and because interference/resource relief raises spectral efficiency and reduces congestion simultaneously. Latency and packet-loss effects are smaller because the permitted interference and resource-relief actions are intentionally modest. The operating-condition analysis also shows that the same policy does not have a uniform effect: low-SINR and high-load cases generally have more improvement headroom than favorable conditions.

5.2 Energy Trade-Off and the Meaning of Balanced Performance

The revised results do not support the previous statement that absolute energy consumption improves simultaneously with all other KPIs. In the disclosed energy equation, additional allocated bandwidth increases power consumption. Critical-state bandwidth expansion therefore produces a modest aggregate energy penalty. This result is scientifically useful because it exposes the multi-objective trade-off that an optimization paper should reveal rather than conceal.

At the same time, throughput per watt improves by about 17%, showing that higher absolute power can coexist with improved energy efficiency. This distinction is consistent with the broader literature's warning that energy optimization and QoS must be evaluated together rather than inferred from a single power metric (Ezzeddine et al., 2024). A future implementation with an explicit transmit-power or sleep-mode control variable could pursue absolute energy reduction directly; the current abstract model does not contain such a control and therefore should not claim it.

5.3 Relationship to Prior Experimental Studies

The direction of the capacity and reliability improvements is consistent with the general premise of recent AI-driven resource-allocation studies, but direct numerical comparison is intentionally avoided. Zhang et al. (2025) evaluate a 5G-TSN industrial scheduling problem with deep reinforcement learning; Shahgholi et al. (2025) use a different network-slicing reward formulation; and Yao and Pérez Yuste (2025) target energy reduction through dynamic cell-resource reconfiguration. Their network models, baselines, action spaces, and objectives differ materially from the present snapshot model. The current evidence therefore establishes improvement relative to the disclosed null-action baseline and restricted ablations, not superiority over those external systems.

5.4 Limitations

First, the experiment is synthetic and snapshot-based. The analytical relationships intentionally encode wireless-network dependencies, so correlations and optimization responses partly reflect the model structure. This is why repeated-seed robustness should be interpreted as reproducibility within the stated model, not as independent field validation.

Second, the simulator does not implement explicit cell topology, carrier-frequency propagation, small-scale fading, packet scheduling, time evolution, or end-to-end protocol stacks. Consequently, quantities such as 44 ms latency cannot be interpreted as an IMT-2020 user-plane latency measurement, and no standards-compliance claim is made.

Third, the action bounds are study-defined. Sensitivity analysis shows that performance magnitude changes substantially with control authority. Real deployments would require action bounds derived from radio configuration, spectrum licensing, interference-coordination constraints, slice policies, and safety limits.

Fourth, the current energy equation responds directly to user density, traffic demand, and bandwidth but not to interference-control or scheduler power. This limits the energy conclusions and explains why absolute energy does not decrease under the primary action set. A more physical power model is required before deployment-level sustainability claims are appropriate.

Fifth, although the XGBoost trigger is strongly validated within the companion classification work, the complete optimization loop has not been tested on live operator data or a packet-level O-RAN testbed. External validation of the optimization policy remains future work.

6. Conclusion

This study provides preliminary simulation-based evidence for an AI-triggered adaptive optimization procedure under a fully disclosed synthetic 5G snapshot model. The revised experiment replaces fixed output multipliers with state-conditioned changes to bandwidth, interference, and resource-utilization variables, followed by recomputation of all QoS outputs using matched stochastic residuals. Across 20 independent seeds and 40,000 evaluation snapshots, throughput increased by 20.47%, latency decreased by 2.60%, packet loss decreased by 5.51%, and spectral efficiency increased by 5.41%. Absolute energy consumption increased by 2.94%, revealing a genuine capacity-energy trade-off, while throughput-per-watt improved by 17.03%. State, operating-condition, sensitivity, controlled alternative-optimizer, ablation, and instance-level analyses show where the benefits occur and where the trade-off arises. These findings should not be interpreted as proof of live-network performance or as evidence that every KPI improves for every snapshot. Full reproducibility is supported by disclosure of the generator, classifier configuration, action search, software environment, repeated seeds, and analysis code. Future work should transfer the policy to a packet-level or O-RAN-compatible testbed with explicit topology, channel, scheduler, power, and temporal dynamics, and should compare against like-for-like established optimization baselines.

Data Availability and Reproducibility

The complete synthetic-data generator, leakage-controlled state-classification trigger, optimization action-search procedure, repeated-seed evaluation code, statistical analysis, and figure-generation scripts are supplied in the accompanying reproducibility notebook for this revision and are available from the corresponding author. The study uses controlled synthetic network-state snapshots and contains no subscriber or personally identifiable data.

Conflict of Interest

The authors declare no conflict of interest.

References

- 3rd Generation Partnership Project (3GPP). (2025). System architecture for the 5G System (5GS) (3GPP TS 23.501, Version 18.11.0, Release 18).
- Alzubaidi, O. T. H., Hindia, M. N., Dimyati, K., Noordin, K. A., Wahab, A. N. A., Qamar, F., & Hassan, R. (2022). Interference challenges and management in B5G network design: A comprehensive review. *Electronics*, 11(18), 2842. <https://doi.org/10.3390/electronics11182842>
- Brilhante, D. S., Manjarres, J. C., Moreira, R., de Oliveira Veiga, L., de Rezende, J. F., Müller, F., Klautau, A., Mendes, L. L., & de Figueiredo, F. A. P. (2023). A literature survey on AI-aided beamforming and beam management for 5G and 6G systems. *Sensors*, 23(9), 4359. <https://doi.org/10.3390/s23094359>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. <https://doi.org/10.1145/2939672.2939785>
- Ezzeddine, Z., Khalil, A., Zeddini, B., & Ouslimani, H. H. (2024). A survey on green enablers: A study on the energy efficiency of AI-based 5G networks. *Sensors*, 24(14), 4609. <https://doi.org/10.3390/s24144609>
- International Telecommunication Union Radiocommunication Sector (ITU-R). (2017). Minimum requirements related to technical performance for IMT-2020 radio interface(s) (Report ITU-R M.2410-0). International Telecommunication Union.

- Morocho-Cayamcela, M. E., & Lim, W. (2018). Artificial intelligence in 5G technology: A survey. Proceedings of the 2018 International Conference on Information and Communication Technology Convergence, 860–865. <https://doi.org/10.1109/ICTC.2018.8539642>
- Shahgholi, T., Khamforoosh, K., Sheikhamadi, A., & Azizi, S. (2025). Optimization of resource allocations in 5G mobile network using active reward learning. Engineering Science and Technology, an International Journal, 68, 102089. <https://doi.org/10.1016/j.jestch.2025.102089>
- Yao, X., & Pérez Yuste, A. (2025). A ML-based resource allocation scheme for energy optimization in 5G NR. Sensors, 25(16), 4978. <https://doi.org/10.3390/s25164978>
- Zhang, Y., Sun, L., Ma, Z., Wang, J., Fu, M., & Joung, J. (2025). A 5G-TSN joint resource scheduling algorithm based on optimized deep reinforcement learning model for industrial networks. Ad Hoc Networks, 170, 103783. <https://doi.org/10.1016/j.adhoc.2025.103783>

Tables

Table 1. Controlled Input Distributions Used in the Synthetic Snapshot Generator

Variable	Distribution / Range	Role
User density	Discrete uniform: 50–1000 users	Exogenous load
Traffic demand	Uniform: 50–1000 Mbps	Exogenous load
Interference power	Uniform: –110 to –70 dBm	Radio condition / controllable downward
Signal power	Uniform: –95 to –45 dBm	Exogenous radio condition
Bandwidth	{20, 40, 60, 80, 100} MHz	Resource / controllable upward
Resource utilization	Uniform: 20–100%	Load pressure / controllable downward
Mobility	Uniform: 0–120 km/h	Exogenous mobility

Table 2. XGBoost Development and Validation Summary

Item	Value
Development observations	10,000; seed 42
Train/test split	80:20 stratified
Input features	15; five AOS constituent QoS outputs excluded
Selected XGBoost	300 estimators; depth 4; learning rate 0.05; subsample 0.90; colsample 0.90
Held-out accuracy	96.45%
Five-fold training accuracy	95.36 ± 0.44%
Five-fold macro F1	90.96 ± 0.87%
Held-out confusion matrix	Critical: 1322/22/0; Moderate: 25/546/8; Optimal: 0/16/61 (actual rows, predicted columns)

Table 3. Primary State-Conditioned Optimization Action Bounds

Predicted state	Bandwidth change	Interference reduction	Resource-utilization relief
Critical	0 or +20 MHz	0 or 2 dB	0 or 5 percentage points
Moderate	0 MHz	0 or 1 dB	0 or 3 percentage points
Optimal	0 MHz	0 dB	0 percentage points

Table 4. Primary Repeated-Seed Performance Results (20 Seeds; 2,000 Snapshots per Seed)

Metric	Before mean	After mean	Mean improvement (%)	SD	95% CI (%)	Paired p
Throughput (Mbps)	231.018	278.280	20.47	0.609	[20.18, 20.75]	<0.001
Latency (ms)	45.655	44.469	2.60	0.011	[2.59, 2.60]	<0.001
Packet loss (%)	3.312	3.129	5.51	0.038	[5.50, 5.53]	<0.001
Spectral efficiency (bps/Hz)	4.866	5.129	5.41	0.164	[5.33, 5.48]	<0.001
Energy consumption (W)	145.762	150.046	-2.94	0.062	[-2.97, -2.91]	<0.001

Table 5. State-Disaggregated Mean Percentage Improvements Across 20 Trials

State	Mean n/trial	Throughput	Latency	Packet loss	Spectral eff.	Energy cons.
Critical	1347.3	50.58	2.77	5.37	10.28	-4.25
Moderate	584.4	1.44	2.21	6.56	0.96	0.00
Optimal	68.3	0.00	0.00	0.00	0.00	0.00

Table 6. Robustness of Percentage Improvements Across Operating-Condition Tertiles

Condition	Tertile	Throughput	Latency	Packet loss	Spectral eff.	Energy cons.
Cell load	High	29.28	2.24	4.62	5.99	-3.13
Cell load	Low	13.31	3.18	7.14	4.82	-2.66
Cell load	Medium	20.90	2.64	5.70	5.45	-2.94
SINR	High	8.19	2.24	6.93	0.36	-1.52
SINR	Low	56.44	2.70	4.63	24.51	-3.28
SINR	Medium	30.65	2.78	6.09	8.58	-4.02
Traffic demand	High	25.80	2.36	4.92	5.77	-2.91
Traffic demand	Low	15.77	2.92	6.39	4.99	-2.97
Traffic demand	Medium	20.67	2.61	5.57	5.48	-2.95

Table 7. Individual-Snapshot Outcome Direction Across 40,000 Evaluations

Metric	Improved (%)	Unchanged (%)	Deteriorated (%)
Throughput (Mbps)	94.23	5.77	0.00
Latency (ms)	96.57	3.43	0.00
Packet loss (%)	93.13	6.87	0.00
Spectral efficiency (bps/Hz)	64.96	35.03	0.00
Energy consumption (W)	0.00	52.40	47.60

Table 8. Action-Bound Sensitivity Analysis

Policy	AOS gain	Throughput	Latency	Packet loss	Spectral eff.	Energy cons.
aggressive	6.763	53.13	4.99	10.69	10.70	-6.82
conservative	0.888	2.85	1.47	2.98	2.73	0.00
primary	2.881	20.47	2.60	5.51	5.41	-2.94

Table 9. Controlled Alternative-Optimizer Comparison Under the Primary Action Bounds

Strategy	AOS gain	Throughput	Latency	Packet loss	Spectral eff.	Energy cons.
AOS multi-objective	2.881	20.47	2.60	5.51	5.41	-2.94
Throughput-greedy	2.782	20.59	2.25	4.55	5.40	-3.41

Table 10. Restricted-Strategy Ablation Comparison

Strategy	AOS gain	Throughput	Latency	Packet loss	Spectral eff.	Energy cons.
Bandwidth only	1.054	13.45	0.00	0.00	0.00	-2.83
Full three-lever	2.881	20.47	2.60	5.51	5.41	-2.94
Interference only	1.396	5.01	1.24	3.81	5.21	0.00
Resource-relief only	0.301	0.53	1.36	1.70	0.19	0.00

Table 11. Computational Timing Benchmark for a 2,000-Snapshot Batch

Stage	Median batch time (ms)	Arithmetic per-snapshot time (ms)
XGBoost inference	10.730	0.005365
Optimization action search	4.688	0.002344
Combined	15.417	0.007709

Figures




Figure 1. Mean post-optimization KPI values normalized to the corresponding pre-optimization baseline (100%) across 20 independent seeds.




Figure 2. Mean percentage KPI change across 20 independent seeds with 95% confidence intervals; negative values indicate deterioration under the sign convention used for lower-is-better metrics.




Figure 3. Predicted Critical, Moderate, and Optimal network-state counts across the 40,000 evaluation snapshots.