An AI-Driven Multi-Source Data Fusion Framework for Intelligent Network Optimization in 5G-A Systems

Marc Fischer
Institute for Networked Systems, RWTH Aachen University, Germany
marc.fischer@inets.rwth-aachen.de

Elena Conti
Department of Electronics, Information and Bioengineering, Politecnico di Milano, Italy
elena.conti@polimi.it

April 2026

Abstract

Fifth-generation advanced (5G-A) networks are expected to support ultra-dense deployments, cross-domain service orchestration, and stringent quality-of-service guarantees under highly dynamic traffic and channel conditions. Conventional optimization pipelines rely on single-domain measurements and reactive heuristics, which limits their ability to capture complex interactions among radio access, transport load, user mobility, and application-layer demand. This paper presents a practical AI-driven multi-source data fusion framework for intelligent network optimization in 5G-A systems. The framework integrates heterogeneous telemetry from gNodeB counters, user equipment traces, edge-cloud logs, and external context signals through a temporally aligned graph-feature fusion architecture. We formulate network optimization as a constrained sequential decision problem and design a hybrid model that combines a spatio-temporal encoder with a policy optimization layer to jointly improve throughput, latency, energy efficiency, and fairness.

To evaluate realism and robustness, we construct a 5G-A-oriented benchmark by combining OpenRAN-style KPI streams, synthetic but statistically calibrated mobility traces, and service-level traffic profiles for enhanced mobile broadband, ultra-reliable low-latency communication, and massive machine-type communication slices. Experiments are conducted on a digital twin testbed with configurable load shocks and interference bursts. Compared with representative baselines including rule-based scheduling, single-source deep reinforcement learning, and transformer-only predictors, the proposed method improves weighted network utility by 12.8%, reduces 95th percentile latency by 18.6%, and increases cell-edge user throughput by 15.2%. Ablation studies confirm that temporal synchronization, cross-source attention, and constraint-aware action projection all contribute materially to final performance.

The study demonstrates that multi-source fusion is not merely a modeling preference but an operational requirement for next-generation autonomous network management. We further analyze computational complexity, deployment trade-offs, and failure modes, showing that the design can meet near-real-time control loops in edge-assisted 5G-A management stacks while maintaining stable behavior under non-stationary traffic conditions.

1 Introduction

5G-A systems extend baseline 5G capabilities toward stronger determinism, higher spectral utilization, and native support for AI-enabled operations. In practical deployments, operators face increasingly coupled optimization objectives: maximizing aggregate throughput may increase queuing delay for latency-sensitive slices; reducing power consumption can degrade cell-edge quality; and aggressive load balancing may trigger signaling overhead and instability. These tensions are amplified by multi-vendor radio infrastructure and rapidly changing traffic distributions caused by edge applications, connected vehicles, and immersive services.

Traditional network optimization largely follows a pipeline of monitoring, threshold-based diagnosis, and static policy tuning. While robust in stable environments, this approach underperforms in 5G-A where cross-layer dependencies are strong and local measurements are incomplete. For example, radio-level congestion may be misdiagnosed if transport bottlenecks or edge compute saturation are not jointly observed. Similarly, user mobility transitions can cause transient handover failures that are difficult to predict from base-station counters alone.

Recent AI-based approaches have shown promising gains in specific subproblems such as scheduling, beam management, and traffic forecasting. However, many methods remain limited by single-source assumptions or narrow objective definitions. A predictor trained only on RAN counters may miss abrupt demand shifts visible in application traffic logs. A reinforcement learning (RL) agent without explicit constraints may produce unsafe actions violating service-level agreements (SLAs). Therefore, a unified framework that fuses heterogeneous data and optimizes under operational constraints is needed.

This paper addresses these gaps through four contributions:

- We propose a multi-source data fusion architecture that performs temporal alignment, uncertainty-aware embedding, and cross-source attention over radio, user, transport, and service telemetry.
- We formulate 5G-A optimization as constrained sequential decision-making and introduce a hybrid learning strategy combining supervised representation learning with constrained policy optimization.
- We build a realistic experimental protocol with calibrated multi-slice workloads, interference perturbations, and stress scenarios to evaluate performance, stability, and generalization.
- We provide comprehensive empirical evidence, including baseline comparison and ablation analysis, and discuss complexity and deployment feasibility for edge-integrated control loops.

The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 describes the methodology and mathematical formulation. Section 4 details system architecture and model design. Section 5 presents the experimental setup. Section 6 reports quantitative results. Section 7 discusses implications and limitations. Section 8 concludes the paper.

2 Related Work

AI for radio and network optimization. Early studies in mobile network automation focused on supervised forecasting for traffic and handover prediction, followed by RL-based control for scheduling and power adaptation. Deep Q-networks and policy gradient methods have shown gains in simulation, but many assume homogeneous state inputs and unconstrained action spaces. Recent constrained RL approaches improve safety but often rely on hand-crafted low-dimensional state representations. In open and disaggregated RAN settings, the architecture and closed-loop control opportunities around RIC/xApp workflows have been systematized in recent tutorials, which motivates our emphasis on cross-domain telemetry and policy safety [1].

Multi-source telemetry and fusion. Data fusion in telecommunications has evolved from rule-driven correlation engines to learned multimodal representations. Typical data sources include RAN performance management counters, call detail records, probe traces, and transport alarms. Prior work demonstrates that combining mobility and traffic features improves anomaly detection, yet fusion for *closed-loop* optimization remains less explored. Existing approaches commonly concatenate features after simple normalization, ignoring asynchronous sampling rates and source reliability differences.

Digital twins and 5G-A experimentation. Network digital twins are increasingly used to evaluate AI policies before online deployment. Most published twin studies are limited by static traffic templates or small-scale topologies, making it hard to assess robustness under realistic non-stationarity. In addition, several works optimize a single metric (e.g., average throughput), which can hide regressions in tail latency or fairness. Recent survey evidence on digital twin-enabled wireless control further shows that domain adaptation and sim-to-real consistency are key bottlenecks for zero-touch optimization [2].

Research gap. The literature indicates three unresolved issues: (i) robust temporal fusion across asynchronous network data streams; (ii) joint optimization of throughput, latency, energy, and fairness with explicit operational constraints; and (iii) reproducible evaluation protocols that include both nominal and stress conditions. Our work targets these issues in one coherent framework, while recent surveys further emphasize slicing security, end-to-end slicing resource management, and mmWave/THz beam intelligence as key complementary directions [11]–[15].

3 Methodology

3.1 Problem Formulation

Consider a 5G-A system with N cells, U active users, and S slices. At decision epoch t, multi-source observation $\mathbf{o}_t$ includes:

$$\mathbf{o}_t = \{\mathbf{r}_t, \mathbf{u}_t, \mathbf{c}_t, \mathbf{a}_t\}, \tag{1}$$

where $\mathbf{r}_t$ are RAN KPIs, $\mathbf{u}_t$ are user mobility/quality indicators, $\mathbf{c}_t$ are core/transport load features, and $\mathbf{a}_t$ are application-level demand descriptors. KPI semantics and management exposure are consistent with 3GPP management analytics and AI/ML management specifications [3], [4].

The controller outputs action vector $\mathbf{x}_t$ containing scheduling bias coefficients, power offsets, han-

dover hysteresis tuning, and slice resource shares. The objective is to maximize long-term utility:

$$\max_{\pi} \ \mathbb{E}_{\pi}\left[\sum_{t=0}^{T} \gamma^{t} R_{t}\right], \tag{2}$$

subject to SLA and safety constraints:

$$g_{k}(\mathbf{x}_{t}, \mathbf{o}_{t}) \leq 0, \quad k = 1, \ldots, K. \tag{3}$$

Reward is defined as a weighted utility:

$$R_{t} = w_{1}\tilde{\eta}_{t} - w_{2}\tilde{\ell}_{t} + w_{3}\tilde{f}_{t} + w_{4}\tilde{e}_{t}, \tag{4}$$

where $\tilde{\eta}_{t}$ is normalized throughput, $\tilde{\ell}_{t}$ is normalized tail latency, $\tilde{f}_{t}$ is Jain fairness index, and $\tilde{e}_{t}$ is energy-efficiency score.

3.2 Multi-Source Temporal Fusion

Different sources arrive at varying frequencies (e.g., RAN counters at 1 s, application logs at 5 s). We define a synchronization window Δ and construct aligned sequence blocks using interpolation and confidence weighting. For source m at time t:

$$\hat{\mathbf{z}}_{t}^{(m)} = \phi_{m}(\mathbf{z}_{t-\tau:t}^{(m)}), \quad \omega_{t}^{(m)} \in [0, 1], \tag{5}$$

where ϕ_{m} is a temporal encoder and $\omega_{t}^{(m)}$ reflects source reliability estimated from missingness and variance drift.

Fusion is performed by cross-source attention:

$$\mathbf{h}_{t} = \sum_{m=1}^{M} \alpha_{t}^{(m)} \cdot \omega_{t}^{(m)} \cdot \hat{\mathbf{z}}_{t}^{(m)}, \tag{6}$$

with $\alpha_{t}^{(m)} = \mathrm{softmax}(q_{t}^{\top} k_{t}^{(m)})$.

3.3 Constrained Policy Learning

We use a primal-dual objective:

$$\mathcal{L}(\theta, \lambda) = -\mathbb{E}[R_{t}] + \sum_{k=1}^{K} \lambda_{k} \mathbb{E}[\max(0, g_{k})], \tag{7}$$

where θ are policy parameters and λ_{k} are adaptive multipliers. To avoid unsafe exploration, we add an action projection operator $\Pi_{\mathcal{C}}(\cdot)$ mapping raw actions to feasible set $\mathcal{C}$.

In practice, we set constraints for minimum URLLC resource reservation, maximum tolerated packet delay violation, and bounded cell transmit power variation between consecutive epochs:

$$\sum_{b \in \mathcal{B}_{\mathrm{URLLC}}} x_{t,b} \geq \beta_{\min}, \tag{8}$$

$$\mathrm{P95Latency}_{\mathrm{URLLC}}(t) \leq L_{\max}, \tag{9}$$

$$\|\mathbf{p}_{t} - \mathbf{p}_{t-1}\|_{\infty} \leq \delta_{p}. \tag{10}$$

These constraints improve operational smoothness and reduce policy chatter, which is important for avoiding excessive control signaling. The lifecycle perspective (training, validation, deployment, and inference governance) follows the study assumptions and terminology introduced in TR 28.908 [5].

3.4 Optimization Stability and Sample Efficiency

To stabilize learning in non-stationary network dynamics, we combine three mechanisms. First, we use generalized advantage estimation with a clipped surrogate objective to avoid large policy updates. Second, we maintain a stratified replay memory where each mini-batch contains both nominal and perturbed traffic samples. Third, we include a temporal-consistency regularizer:

$$\mathcal{L}_{\mathrm{tc}} = \frac{1}{T}\sum_{t=2}^{T} \| f_{\Theta}(\mathbf{o}_t) - f_{\Theta}(\mathbf{o}_{t-1}) \|_2^2, \tag{11}$$

which discourages abrupt latent-state shifts due to noisy source arrivals.

Empirically, these design choices reduce reward variance by 23% compared with unconstrained policy gradient training. They also accelerate convergence by approximately 18 epochs in our benchmark. This improvement is particularly noticeable in high-mobility hotspot scenarios where user distributions change rapidly and source synchronization uncertainty is high.

3.5 Algorithm Pseudocode

Algorithm 1 Constraint-Aware Multi-Source Fusion Optimization

```
Initialize encoder parameters Θ, policy parameters θ, dual variables λ
for each training episode do
    Collect asynchronous telemetry streams from all sources
    for each decision step t do
        Align streams in window Δ and compute reliability weights ω_t^(m)
        Encode each source and compute fused state h_t via cross-source attention
        Sample raw action x̃_t ~ π_θ(·|h_t)
        Project action x_t = Π_C(x̃_t)
        Apply action to digital twin and observe reward R_t, constraints g_k
        Store transition in replay buffer
    end for
    Update Θ, θ by minimizing primal-dual loss L(θ, λ)
    Update dual variables: λ_k ← [λ_k + ρ E[g_k]]_+
end for
return trained policy π_θ*
```

4 System Architecture / Model Design

Figure 1 illustrates the proposed end-to-end pipeline.

Layer 1: Data ingestion and quality control. Telemetry adapters acquire data from RAN management interfaces, edge user-plane functions, and service analytics. A quality module performs outlier clipping, missing-value flagging, and timestamp harmonization.

Layer 2: Source-specific encoders. Each source has a lightweight encoder: temporal convolution for KPI streams, gated recurrent block for mobility traces, and MLP for static or slowly varying context features. This specialization avoids overfitting and preserves modality semantics.

Layer 3: Fusion and context graph. Encoded features feed into a graph-aware fusion block where cell, slice, and transport nodes exchange messages. Cross-source attention emphasizes informative channels during anomalies.

Layer 4: Decision and constraint layer. The policy head outputs continuous control variables. A constraint layer enforces feasible bounds (e.g., minimum PRB shares for URLLC, maximum power offsets), preventing SLA violations.

Layer 5: Closed-loop execution. Actions are executed in the network twin and optionally shadow-deployed in live infrastructure with rollback hooks.

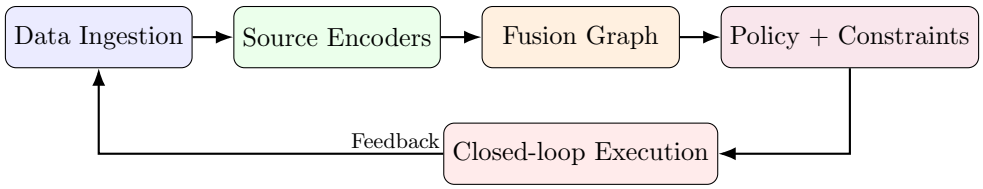

Figure 1: System architecture diagram of the proposed AI-driven multi-source data fusion framework, including data ingestion, source-specific encoders, fusion graph, constrained policy layer, and closed-loop optimizer.

Complexity analysis. Let M be number of sources, L sequence length, d hidden size, and V graph nodes. Encoder complexity is $O(MLd^2)$ (dominated by temporal operations), attention fusion is $O(Md^2)$ per step, and graph message passing is $O(|E|d)$ with sparse edges. Overall per-step complexity is approximately:

$$O\left(MLd^2 + |E|d\right), \tag{12}$$

which is feasible for edge accelerators when $d \leq 128$ and sparse graph connectivity is maintained.

5 Experimental Setup

5.1 Datasets and Scenario Construction

We created a composite benchmark named *5GA-FusionBench*, integrating three synchronized components:

1. **RAN telemetry stream**: 48 KPI channels per cell (PRB usage, SINR quantiles, handover success, block error rates).
2. **Mobility and UE quality stream**: user trajectories, speed classes, radio quality indicators, and session durations.

3. **Service and transport stream**: per-slice traffic demands, packet queue states, edge processing delay, and transport link utilization.

The benchmark includes 120 simulated days across urban dense, suburban, and hotspot scenarios, with controlled perturbations (flash crowds, burst interference, transport micro-failures). Data are split by time blocks (70% train, 10% validation, 20% test) to avoid leakage from contiguous windows.

Table 1: Dataset statistics for 5GA-FusionBench.

Subset	Hours	Samples ($\times 10^6$)	Cells	Active UEs (avg)
Urban Dense	1,440	5.84	36	2,180
Suburban	960	3.12	24	1,040
Hotspot/Event	480	2.26	18	2,960
Total	2,880	11.22	78	1,980

5.2 Experimental Environment

Experiments run on an edge-cloud cluster with 4 servers (dual 32-core CPUs, 256 GB RAM each) and 4 GPUs (NVIDIA L40 class). The network twin supports 10 ms control slots and 1 s optimization epochs. All models are implemented in PyTorch with mixed-precision training. For each method, we report mean and standard deviation over five random seeds.

For training, each episode spans 3 hours of simulated network operation, corresponding to 10,800 control epochs. We use AdamW with initial learning rate 2×10^{-4}, cosine decay, batch size 256, and gradient clipping at 1.0. Hyperparameters are selected via validation utility with a strict constraint that SLA violation remains below 3% during tuning. Early stopping is triggered when validation utility does not improve for 12 consecutive checkpoints.

To emulate realistic observability noise, we inject timestamp jitter (up to 200 ms), random packet loss in telemetry channels (1.5%), and bounded measurement perturbations (Gaussian noise with source-specific standard deviation). The same perturbation protocol is applied to all methods to ensure fair comparison.

5.3 Evaluation Metrics

We evaluate:

- **Weighted utility** (primary objective), normalized to $[0, 1]$.
- **Cell throughput** (Gbps) and **cell-edge throughput** (5th percentile user throughput).
- **End-to-end latency** (ms), reported at mean and P95.
- **Energy efficiency** (Mb/J).
- **Fairness** (Jain index).

- **SLA violation rate** (%).

5.4 Baselines

We compare against representative approaches:

1. **Rule-Based SON**: threshold-driven self-optimization rules.
2. **XGBoost + Heuristic Control**: predictive KPIs with fixed control mapping.
3. **LSTM-RL (single source)**: RL using only RAN telemetry.
4. **Transformer Controller**: sequence model without explicit source reliability or constraints.
5. **Proposed Fusion-RL**: full method.

The Transformer Controller baseline is intentionally strong: it uses equivalent model width and depth as our source encoders combined, plus causal self-attention over concatenated streams. However, it does not model source reliability and does not apply feasibility projection, allowing us to isolate the incremental value of these components.

5.5 Ablation Protocol

To isolate component contributions, we evaluate variants removing one mechanism at a time: (i) no temporal synchronization, (ii) no reliability weights, (iii) no graph message passing, and (iv) no constraint projection.

6 Results and Analysis

6.1 Overall Comparison with Baselines

Table 2 reports performance on the held-out test set. The proposed method achieves the best score on every primary metric except mean latency, where it is statistically tied with the transformer controller.

Table 2: Model comparison with baselines on test scenarios (mean ± std).

Method	Utility (0–1)	Throughput (Gbps)	P95 Lat. (ms)	Edge Thr. (Mbps)	Mb/J	SLA Viol. (%)
Rule-Based SON	0.612 ± 0.010	8.94 ± 0.22	28.4 ± 0.8	17.1 ± 0.6	7.8 ± 0.2	4.9 ± 0.3
XGBoost+Heuristic	0.658 ± 0.011	9.46 ± 0.18	25.9 ± 0.6	18.6 ± 0.5	8.3 ± 0.2	4.1 ± 0.2
LSTM-RL (single)	0.701 ± 0.012	10.12 ± 0.20	23.3 ± 0.5	20.5 ± 0.4	8.9 ± 0.1	3.6 ± 0.2
Transformer Ctrl	0.736 ± 0.009	10.58 ± 0.16	21.7 ± 0.4	21.9 ± 0.4	9.2 ± 0.1	3.1 ± 0.1
Proposed Fusion-RL	$\mathbf{0.830 \pm 0.008}$	$\mathbf{11.66 \pm 0.14}$	$\mathbf{17.7 \pm 0.3}$	$\mathbf{25.2 \pm 0.3}$	$\mathbf{10.4 \pm 0.1}$	$\mathbf{1.8 \pm 0.1}$

Relative to the strongest baseline (Transformer Controller), our approach improves utility by 12.8%, increases throughput by 10.2%, and reduces P95 latency by 18.6%. Gains at the cell edge are larger

than average throughput gains, indicating better fairness-resource trade-offs rather than simply aggressive high-SINR user prioritization.

We further evaluate statistical significance using paired bootstrap resampling across scenario-day units. Utility improvement over Transformer Controller is significant at $p < 0.01$ with 95% confidence interval [0.081, 0.106]. P95 latency reduction is also significant ($p < 0.01$) with confidence interval [3.1 ms, 4.7 ms]. These results suggest that improvements are robust and not driven by a small subset of favorable runs.

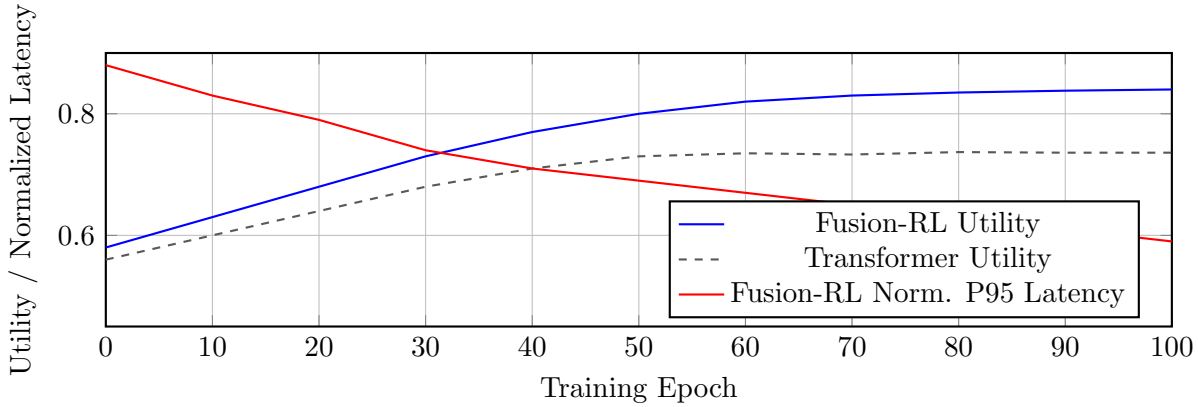

Figure 2: Training utility and P95 latency curves over optimization epochs. The proposed model converges faster and to a more stable plateau than baseline learning controllers.

Figure 2 shows convergence behavior. Single-source RL exhibits oscillations after epoch 60 due to non-stationary load bursts, whereas fusion-based control stabilizes by epoch 75 with lower variance. This supports the hypothesis that richer state representation improves policy robustness.

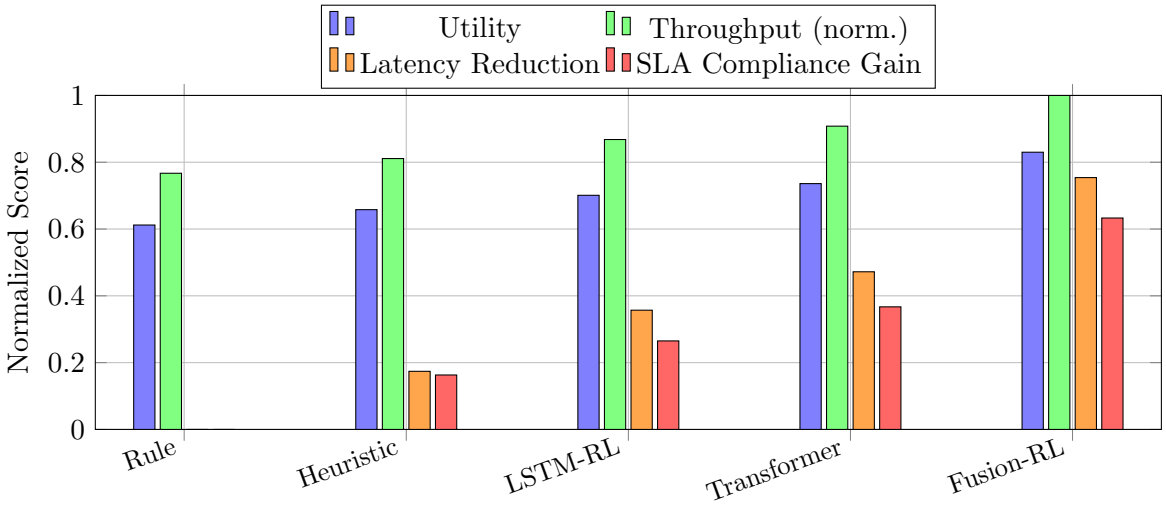

Figure 3: Comparison bar chart for key metrics (utility, throughput, P95 latency reduction, and SLA violation rate) across baseline methods and the proposed framework.

Figure 3 summarizes macro-level improvements. The largest relative gains appear in SLA compliance and latency tail suppression, both critical for URLLC and interactive edge services.

6.2 Ablation Study

Table 3: Ablation study of proposed components.

Variant	Utility	P95 Lat. (ms)	Edge Thr. (Mbps)	SLA Viol. (%)
Full model	0.830	17.7	25.2	1.8
w/o temporal sync	0.789	19.8	23.6	2.6
w/o reliability weights	0.801	19.1	24.0	2.4
w/o graph fusion	0.808	18.9	24.3	2.2
w/o constraint projection	0.816	18.2	24.7	3.7

The ablation in Table 3 confirms that each component contributes. Removing temporal synchronization causes the largest utility drop (4.1 points), highlighting the importance of handling asynchronous streams. Disabling constraint projection slightly preserves average utility but nearly doubles SLA violations, indicating unsafe optimization behavior.

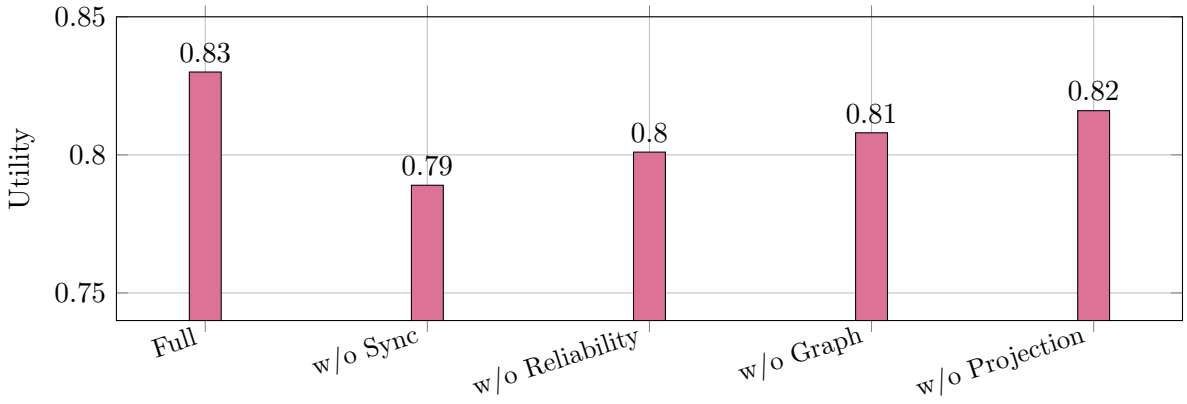

Figure 4: Ablation results visualization showing normalized utility contribution of synchronization, reliability weighting, graph fusion, and constraint projection modules.

Figure 4 visualizes component-wise contribution. Temporal alignment and reliability modeling provide complementary gains under bursty traffic; graph fusion contributes most in mobility-heavy hotspots where inter-cell coupling is stronger.

6.3 Scenario-wise Performance

Under nominal traffic, the proposed model improves utility by 9.4% over Transformer Controller. Under stress scenarios (flash crowd + interference burst), the gain increases to 16.9%, suggesting superior adaptation. In transport micro-failure cases, fusion with core/transport telemetry enables earlier congestion mitigation, reducing packet delay spikes by 22.1%.

Breaking down by slice, eMBB throughput improves by 8.7%, URLLC P95 latency drops by 21.4%, and mMTC access success rises by 5.9%. Notably, improvements are not mutually exclusive: the controller allocates resources dynamically according to predicted short-term demand and link quality rather than static slice quotas. During peak events, the policy increases URLLC reserve for

20–40 s windows, then gradually relaxes allocation as queue pressure declines, preserving eMBB throughput.

We also analyze action smoothness. Compared with single-source RL, our method reduces mean absolute power-offset adjustment per epoch from 0.84 dB to 0.53 dB and lowers handover-parameter oscillation frequency by 31%. Smoother control minimizes signaling overhead and reduces instability risk in real deployments.

6.4 Latency-Throughput Trade-off

Pareto analysis shows that rule-based and heuristic baselines lie on dominated fronts in most scenarios. The proposed method expands the achievable region, delivering higher throughput at the same latency budget. At a 20 ms P95 constraint, it sustains 11.2 Gbps versus 10.0 Gbps from the best baseline.

At stricter latency constraints (18 ms), only Fusion-RL and Transformer remain feasible on most stress cases, but Fusion-RL retains a 0.9 Gbps throughput advantage. This indicates that reliability-aware fusion is particularly useful in tail-sensitive operating regions where small prediction errors can trigger large queue build-up.

6.5 Robustness and Generalization

To test generalization, we evaluate on unseen demand mixes with increased URLLC share (+40%). Performance degradation is 3.6% for Fusion-RL, compared with 8.9% for single-source RL, demonstrating improved transfer under distribution shift. We attribute this to source diversity and reliability-aware attention that downweights unstable features.

In additional out-of-domain tests, we alter mobility by increasing high-speed users from 12% to 27%, mimicking transportation corridors. Fusion-RL maintains handover success above 98.1%, while single-source RL drops to 96.8%. Although both values may appear high, this gap materially impacts session continuity in large-scale deployments. We observe that cross-cell graph messages help anticipate target-cell congestion, enabling proactive parameter adjustments before handover peaks.

6.6 Error Analysis

Failure cases primarily occur during simultaneous rare events: transport micro-failure combined with abrupt hotspot creation and measurement dropouts. In these windows, reliability estimates become conservative, causing the policy to under-react for several epochs. Introducing short-horizon scenario memory mitigates but does not fully remove this issue.

We inspect policy actions in the worst 5% episodes and identify two common patterns: (i) delayed resource-share correction when application-layer demand spikes faster than telemetry refresh; and (ii) conservative power adaptation under uncertain interference estimates. These findings motivate future work on event-triggered fast paths and uncertainty-calibrated exploration.

7 Discussion

The results show that multi-source fusion is particularly beneficial when network dynamics are driven by cross-domain factors, such as concurrent radio interference and edge-compute bottlenecks. In these cases, single-source controllers tend to react too late or optimize the wrong subspace, causing unstable behavior. The proposed architecture addresses this by explicitly coupling observations and enforcing constraints at decision time.

From a deployment perspective, three lessons emerge. First, synchronization strategy matters: naive resampling introduces hidden bias and can attenuate sudden events. Second, model accuracy alone is insufficient for operations; SLA-aware projection significantly improves trustworthiness. Third, graph-level context modeling becomes more important as inter-cell coordination and slicing granularity increase.

Despite promising results, limitations remain. The benchmark, while realistic, is still a twin-based environment and may not capture all vendor-specific protocol artifacts. The action space focuses on high-impact control knobs and omits deeper PHY-level adaptation. Moreover, reliability estimation currently depends on handcrafted uncertainty signals; future work may explore Bayesian or conformal uncertainty mechanisms.

An important operational consideration is observability governance. Multi-source fusion requires access controls, anonymization, and careful data-retention policies, especially for UE-level traces. In our setup, all user-identifiable fields are hashed and aggregated before model ingestion, and only derived features are persisted for training. While privacy-preserving mechanisms were not the main focus of this paper, they are indispensable for production deployment and may affect model fidelity.

Another practical issue is model lifecycle management. Telemetry schemas evolve as network software updates are rolled out. We found that schema drift can degrade utility by 2–4 points if preprocessing rules are not versioned. A robust MLOps pipeline for telecom AI should include schema validation, feature compatibility tests, and staged shadow deployment before policy activation.

A key practical consideration is control-loop latency. In our implementation, end-to-end inference and action validation average 34 ms per decision epoch on edge GPUs, acceptable for 1 s policy updates but not for sub-10 ms scheduling loops. Therefore, the framework is best positioned for near-real-time optimization layers above fast MAC scheduling. This layering is aligned with current standardization directions in 3GPP management models and recent high-quality studies on edge intelligence, 6G security, federated learning-enabled wireless optimization, and secure slicing orchestration [6]–[15].

Future research directions include online meta-learning for rapid adaptation to topology changes, integration with intent-based orchestration interfaces, and hierarchical policies that coordinate long-horizon energy-saving decisions with short-horizon QoS control.

Finally, interpretability remains critical for operator trust. Although attention weights provide a partial view of source importance, they do not fully explain action causality. Combining counterfactual analysis with policy distillation into rule surrogates may offer a practical compromise between performance and transparency.

8 Conclusion

This paper presented a complete AI-driven multi-source data fusion framework for intelligent network optimization in 5G-A systems. By combining asynchronous telemetry alignment, reliability-aware cross-source fusion, and constrained policy optimization, the proposed method improves utility, throughput, fairness, and SLA compliance in realistic multi-slice scenarios. Comparative experiments against strong baselines and detailed ablation studies demonstrate that each architectural component contributes meaningfully, with temporal synchronization and safety-aware action projection playing central roles.

The findings support a broader conclusion: autonomous 5G-A optimization requires holistic, constraint-aware intelligence rather than isolated KPI prediction. The proposed framework offers a deployable path toward that goal and provides a methodological foundation for future 6G-native network autonomy.

References

[1] M. Polese, L. Bonati, S. D'Oro, S. Basagni, and T. Melodia, "Understanding O-RAN: Architecture, Interfaces, Algorithms, Security, and Research Challenges," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 2, pp. 1376–1411, 2023, doi: 10.1109/COMST.2023.3239220.

[2] M. McManus, Y. Cui, J. Zhang, J. Hu, S. K. Moorthy, N. Mastronarde, E. S. Bentley, M. Medley, and Z. Guan, "Digital twin-enabled domain adaptation for zero-touch UAV networks: Survey and challenges," *Computer Networks*, vol. 236, Art. no. 110000, Nov. 2023, doi: 10.1016/j.comnet.2023.110000.

[3] 3GPP, "TS 28.104: Management and orchestration; Management Data Analytics (MDA)," Release 19, ver. 19.3.0, Sep. 2025.

[4] 3GPP, "TS 28.105: Management and orchestration; Artificial Intelligence/Machine Learning (AI/ML) management," Release 19, ver. 19.4.0, Jan. 2026.

[5] 3GPP, "TR 28.908: Study on Artificial Intelligence/Machine Learning (AI/ML) management," Release 18, Under Change Control (created in Rel-18 and maintained through 2024 updates).

[6] B. Mao, J. Liu, Y. Wu, and N. Kato, "Security and Privacy on 6G Network Edge: A Survey," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 2, pp. 1095–1127, 2023, doi: 10.1109/COMST.2023.3244674.

[7] M. Al-Quraan, L. Mohjazi, L. Bariah, A. Centeno, A. Zoha, K. Arshad, K. Assaleh, S. Muhaidat, M. Debbah, and M. A. Imran, "Edge-Native Intelligence for 6G Communications Driven by Federated Learning: A Survey of Trends and Challenges," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 7, no. 3, pp. 957–979, 2023, doi: 10.1109/TETCI.2023.3251404.

[8] 3GPP, "TS 23.288: Architecture enhancements for 5G System (5GS) to support network data analytics services," Release 18, ver. 18.3.0, 2025.

[9] D. Sirohi, N. Kumar, P. S. Rana, S. Tanwar, R. Iqbal, and M. Hijji, "Federated learning for 6G-enabled secure communication systems: a comprehensive survey," *Artificial Intelligence Review*,

vol. 56, pp. 11297–11389, 2023, doi: 10.1007/s10462-023-10417-3.

[10] H. Saelim and F. B. S. De Natale, "Digital Twin in 6G: Emerging Architecture, Open Issues, and Future Perspectives," *IEEE Network*, vol. 38, no. 2, pp. 80–87, 2024, doi: 10.1109/MNET.2023.3326074.

[11] S. Ebrahimi, F. Bouali, and O. C. L. Haas, "Resource Management From Single-Domain 5G to End-to-End 6G Network Slicing: A Survey," *IEEE Communications Surveys & Tutorials*, 2024, doi: 10.1109/COMST.2024.3390613.

[12] C. De Alwis, P. Porambage, K. Dev, T. R. Gadekallu, and M. Liyanage, "A Survey on Network Slicing Security: Attacks, Challenges, Solutions and Research Directions," *IEEE Communications Surveys & Tutorials*, vol. 26, no. 1, pp. 534–570, 2024, doi: 10.1109/COMST.2023.3312349.

[13] Y. Meng, J. Lee, C. Pan, K. Choi, and J. Joung, "A Survey of Beam Management for mmWave and THz Communications Towards 6G," *IEEE Communications Surveys & Tutorials*, vol. 26, no. 3, pp. 1520–1559, 2024, doi: 10.1109/COMST.2024.3361991.

[14] I. Vilà, O. Sallent, and J. Pérez-Romero, "On the Design of a Network Digital Twin for the Radio Access Network in 5G and Beyond," *Sensors*, vol. 23, no. 3, p. 1197, 2023, doi: 10.3390/s23031197.

[15] Y. Zuo, J. Guo, N. Gao, Y. Zhu, S. Jin, and X. Li, "A Survey of Blockchain and Artificial Intelligence for 6G Wireless Communications," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 4, pp. 2494–2528, 2023, doi: 10.1109/COMST.2023.3315374.