

Predicting connection quality degradation by radio events in mobile networks using interpretable models

Márton Molnár^{*†}, Attila Báder[†], Aitor Hernandez Herranz[†], Samie Mostafavi[†], Alija Pašić^{*}

^{*}Department of Networked Systems and Services, Faculty of Electrical Engineering and Informatics
Budapest University of Technology

{mmolnar, palija}@hit.bme.hu

[†]Ericsson

{attila.bader, marton.molnar, aitor.hernandez.herranz, samie.mostafavi}@ericsson.com

Abstract—Mobile video traffic constitutes the majority of modern cellular network loads, yet end-to-end encryption creates a black-box environment that complicates client-side Quality of Service (QoS) monitoring. However, radio-layer telemetry remains accessible to network operators regardless of encryption, offering a viable source of real-time observability. This study exploits this opportunity using Cell Trace Record (CTR) data — high-resolution radio-layer events extracted directly from a 5G production environment. Our methodology employs a temporal feature engineering pipeline that transforms raw CTR events into an interpretable feature space. By evaluating tree-based ensemble architectures (XGBoost and Random Forests), we validate how temporal patterns in radio-layer parameters correlate with application-level performance. Results demonstrate that radio-layer records alone enable the prediction of delivery rate, Time-to-Last-Byte (TTLB), and Round-Trip Time (RTT) for uplink traffic, with a lead time of 10 to 15 seconds, depending on the target variable. These findings establish that UE (User Equipment)-level radio parameters — particularly modulation schemes — serve as a high-fidelity, proactive proxy for service quality, enabling early detection of degradations and explainable automation in encrypted 5G infrastructures.

Index Terms—5G NR, Quality of Experience (QoE), ML, Radio Access Network (RAN), Mobile Video Streaming, Cell Trace Record (CTR), Network Traffic Analysis, Real-world Measurement, TCP, xApp, rApp, Mobile Network Operator (MNO). Low Latency, Proactive Network, 6G

I. INTRODUCTION

The rapid evolution of 5G and beyond wireless networks has shifted the paradigm of network management from reactive troubleshooting towards proactive, data-driven optimization. As mobile traffic becomes increasingly dominated by latency-sensitive applications, the ability to anticipate network state transitions is critical. Conventional Quality of Service (QoS) monitoring frameworks often rely on instantaneous or post-event reporting, which inherently introduces a feedback delay that limits the effectiveness of dynamic resource allocation, predictive handover, and congestion control mechanisms.

To address these challenges, predictive modeling has emerged as a vital component of modern telecommunications

infrastructure. By leveraging high-resolution telemetry data network orchestrators can forecast future Key Performance Indicators (KPIs) before degradation impacts the end-user Quality of Experience (QoE). However, the highly dynamic nature of the 5G RAN, characterized by fast-fading channels and rapid fluctuations in cell load, introduces significant stochastic noise into these telemetry streams. This necessitates a robust analytical framework capable of distinguishing between transient outliers and actionable network trends.

This study investigates the predictive modeling of throughput and latency metrics using tree-based machine learning applied to radio-layer data. We evaluate the impact of temporal parameters — specifically the depth of historical input traces and the size of the prediction horizon — on model stability and convergence. By identifying the most influential radio-level features and their cross-layer temporal correlations, we provide actionable results for preemptive network optimization. We demonstrate that while throughput-based metrics exhibit higher temporal persistence, latency-related parameters require optimized aggregation intervals to mitigate the effects of volatility. We also demonstrate that while throughput-based metrics exhibit higher temporal persistence, latency-related parameters require optimized aggregation intervals to mitigate the effects of extreme volatility.

From a deployment perspective, these findings align with the broader industry trend toward network capability exposure, where initiatives such as CAMARA [1] aim to surface network intelligence as standardized APIs for third-party applications. Within this paradigm, the O-RAN [2] architecture offers a natural deployment target: the predictive models developed in this work are intended for integration as a RAN Intelligent Controller Application (rApp), delivering proactive service assurance through vendor-neutral interfaces. We evaluate this in the context of live television broadcast workflows, where network-side inference of per-camera uplink health can drive deterministic feed-switching decisions before degradation reaches the viewer.

II. RELATED WORK

Existing QoS prediction methodologies often suffer from a structural limitation: they are primarily auto-regressive, relying on historical target data while treating the radio network as a black box. Even studies incorporating basic radio metrics like Reference Signal Received Power (RSRP) or Signal-to-Interference-plus-Noise Ratio (SINR) lack the depth needed to capture the physical layer dynamics that dictate data flow. We captured high resolution radio data which proposes a shift toward transparent, radio-centric modeling. We utilize the high-resolution parameters extracted directly from near real-time Cell Trace Records (CTRs), creating a feature space that bridges the gap between raw RAN telemetry and application-layer performance.

The prediction of network performance metrics, particularly throughput and latency, has seen a surge of interest with the maturation of 5G networks. For instance, [3] proposes a closed-loop clustering-based approach for global bandwidth prediction in video streaming, while [4] provides a comparative study of various time-series models implemented in cloud environments for 5G traffic. Similarly, [5] utilizes LSTM models combined with trace similarity clustering to improve prediction accuracy for adaptive video services.

Several works attempt to reduce computational overhead or handle mobility. [6] introduces a low-complexity Kalman Filter-based algorithm specifically for 5G throughput prediction, and [7] focuses on improving QoE through accurate throughput forecasting in mobile adaptive streaming scenarios. Advanced deep learning architectures have also been explored ([8], [9]) for bandwidth and QoE prediction. Furthermore, [10] proposes a hybrid model named HYPER, using ARIMA for intra-cell and Random Forest for cross-cell prediction by incorporating limited network features during handovers.

The authors in [11] created a solution a dynamic 5QI configuration framework, which uses radio connection as primary input to trigger a 5QI change (which might be a possible action of our rApp proposal), however only RSRP is used as the triggering feature, and better results could be acquired by using detailed cell trace level features.

While QoS prediction is well-studied, our approach uniquely leverages high-resolution cell trace radio data. We validate and analyze this methodology in a production-grade industrial 5G environment which reflects the current relevant problems in quality forecasting for the vendors.

A. Motivation and Contributions

Our models use only radio-level parameters extracted from the CTRs. The primary motivation and our technical contributions of this study are summarized as follows:

- **Utilizing RAN CTRs:** Unlike transport and application-layer methods, we utilize exclusively UE and Cell level CTRs which are available near real-time. This enables a direct mapping from the physical layer to the application layer QoS.
- **Eliminating autoregressive features:** The model operates with zero functional dependency on historical traces

of the **target** variables: delivery rate, RTT, or TTLB, and uses CTR-only history. By eliminating autoregressive inputs, we create predictions using solely the historical values of the radio features. Note (CTR-only): e.g. if the target is RTT, the history of TTLB or others are not used.

- **Extended Predictive Horizon for Peak Detection:** The proposed approach identifies critical application-layer fluctuations and peaks with significant lead times (10-15 seconds often). This allows for proactive network adaptation far earlier than reactive methods.

III. MEASUREMENTS AND METHODOLOGY

To model the non-linear relationship between radio metrics and application performance, we implemented XGBoost and RandomForest in a supervised learning context.

A. Tree-Based Ensemble Architectures

Tree-based methods were selected for their robustness to feature scaling and non-linear network data distributions. Our XGBoost implementation generates a point estimate \hat{y}_{t+F} by optimizing a regularized objective function through additive aggregation:

$$\hat{y}_{t+F} = \sum_{m=1}^M f_m(\mathbf{X}_t), \quad f_m \in \mathcal{F} \quad (1)$$

where each f_m is a tree trained via second-order Taylor expansion. This architecture effectively handles missing telemetry through sparsity-aware splitting and prevents overfitting during fluctuations using L_1 and L_2 regularization.

Alternatively, the Random Forest (RF) approach utilizes bootstrap aggregating to average the outputs of M independent learners:

$$\hat{y}_{t+F} = \frac{1}{M} \sum_{m=1}^M f_m(\mathbf{X}_t) \quad (2)$$

By training each tree on randomized data subsets and feature subspaces, RF decorrelates individual learners. This reduces variance and provides a stable baseline against high-frequency noise in signal strength parameters.

B. Measurement Description and Scenarios

We conducted 10 measurements, each lasting at least 10 minutes, within the Kista Innovation Park (KIP) in the Husky Testbed [12] using the goDASH framework [13]. The campaign focused on mobile uplink (UL) video streaming. We chose uplink streaming due to UE power constraints, asymmetric resource allocation.

To evaluate model performance across various network states, we executed 10 scenarios spanning different background traffic intensities and QoS configurations:

- **No load:** A 0 kbps background traffic scenario in an idle cell.
- **Congested States:** Nine scenarios featuring constant background cross-traffic at 2, 4, and 8 Mbps (3-3-3).

- **5QI Classes:** Each load level was tested across three 3GPP 5G QoS Identifier (5QI) classes: 4, 6, and 9 (default) [14].

This measurement set captures the signature patterns of bitrate degradation across a wide spectrum of physical and transport-layer dynamics. This provides a narrow, but valuable foundation for training our predictive models. Our measurement campaign is continuously ongoing, expanding the number of scenarios and measurements for more evidence.

C. Problem formulation

To formalize the prediction problem, we define the following parameters governing the data structure:

- **Window Size (W):** The temporal resolution. Each data point aggregates raw metrics over W seconds using the mean operator.
- **History (H):** The look-back period: the number of historical points used as input features. The feature vector \mathbf{X}_t at time t is:

$$\mathbf{X}_t = [x_{t-1}, x_{t-2}, \dots, x_{t-H}, \mathbf{z}_t] \quad (3)$$

where \mathbf{z}_t represents exogenous variables. This obviously requires a warm-up period of H observations at the start of each sequence before operating.

- **Prediction Horizon (F):** The look-ahead distance to the target time $t+F$. The model maps \mathbf{X}_t to a specific future value, bypassing intermediate steps:

$$\hat{y}_{t+F} = f(\mathbf{X}_t), \quad \text{s.t. } \hat{y}_{t+k} \notin \text{output}, \forall k \in \{1, \dots, F-1\} \quad (4)$$

The total temporal scope of a single training sample is $(H + F) \times W$ seconds.

Measurement Set: The data was recorded at 1-second resolution and resampled to the given interval W . All measurements followed an identical trajectory to ensure spatial consistency (Figure 1).

Lon/Lat Coordinates through a set of measurements




Fig. 1: Longitude vs. latitude coordinates showing the consistent trajectory across all measurements.

A specific model configuration is denoted by the tuple $\mathcal{M} = \langle W, H, F \rangle$. E.g. a model which uses 10 second W window, 30 second H history and 10 second F prediction horizon is represented as $\langle W = 10, H = 3, F = 1 \rangle$.

Note that the 5QI class identifier was a regular feature in the input feature list, however, the models did not leverage it directly as it was a constant parameter throughout the full measurement. The effect of using a default or a non-default QoS class might be visible in other CTR parameters such as delay or throughput.

D. Evaluation metrics

The used metrics to evaluate the performance are:

$$MSE = \frac{1}{n} \sum (y_i - \hat{y}_i)^2 \quad RMSE = \sqrt{MSE}$$

$$MAE = \frac{1}{n} \sum |y_i - \hat{y}_i| \quad R^2 = 1 - \frac{\sum (y_i - \bar{y})^2}{\sum (y_i - \bar{y})^2}$$

Where y_i is the observed value, \hat{y}_i is the predicted value, \bar{y} is the mean of y , and n is the number of samples.

E. Feature selection

The input features for the predictive models were categorized into three hierarchical levels to capture distinct network dynamics:

- **Baseline Model:** Relied exclusively on downlink signal quality indicators measured at the router: RSRP, RSRQ and SINR.
- **UE-Level Models:** Extended feature space with device-specific metrics, including HARQ ACK/NACK counters for modulation scheme utilization (16-QAM, 64-QAM, and 256-QAM or QPSK), PUSCH (Physical Uplink Shared Channel) SINR, and UE power headroom.
- **Cell-Level Models:** Incorporated aggregate network performance and congestion indicators: resource block utilization, total throughput, cell-level HARQ statistics, and delays.

F. Training and Testing

Data partitioning was implemented on measurement level to ensure temporal integrity and prevent data leakage. Eight of the 10 total measurement sessions were training measurements and two were reserved as a final test set. The eight measurements formed the training set, which was partitioned using 4-Fold Cross-Validation. Within each fold, 6 measurements of the training set were used for model fitting, and two measurements served as a validation set for hyperparameter optimization using Grid Search Optimization. The number of trees were 200 for Baseline and 300 for UE and CELL models, to prevent overfitting for the few input features in Baseline.

IV. MODEL ANALYSIS AND RESULTS

The predictive performance of the proposed model is heavily dependent on the H value. Because the model requires a history of p preceding observations to populate its input vector, the initial p samples of any new measurement contain incomplete temporal information. During this interval, $t \in [t_0, t_0 + p)$, where t_0 is the start index of a sequence, the lagged features are zeros and are ignored in evaluation metrics. In our visualizations, this region is colored by a gray shaded area to distinguish these states from the *steady-state*



Fig. 2: Cross-metric correlation analysis. The top plot illustrates absolute signal strength values (RSRP, SINR), while the bottom panels show normalized relationships between connection quality (Delivery Rate, TTLB, RTT), Buffer state, and QoE metrics (Stall, ITU-T P.1203 [15]). [QoS 5QI = default, $W = 10$]

prediction phase where $t \geq t_0 + p$. This warm-up zone is displayed with a grey area in Fig 3, 6 and 7.

The cross-metric correlation in Fig. 2 ($W = 10$) highlights the causal chain between radio-layer performance and user-perceived quality. The observed sequence illustrates three key stages:

- **Event (1):** Radio conditions (SINR and RSRP) begin to decline, triggering a transition from 256-QAM to 64-QAM. At this point, the playback buffer remains stable and the delivery rate is not yet significantly affected.
- **Event (2):** The playback buffer begins to deplete. Notably, the ITU-T P.1203 QoE score remains high and no stalls have occurred, marking this as a critical window for proactive network or application-layer adjustments.
- **Event (3):** The buffer is fully exhausted, leading to inevitable playback stalls and QoE degradation.

Between Events (2) and (3), a transient improvement in RSRP triggered a slight return to 256-QAM modulation, which later proved insufficient to sustain the link stability.

This analysis demonstrates that the playback buffer acts

as a vital bridge between QoS and QoE. Identifying early fluctuations in radio and connection parameters enables the proactive forecasting of QoE shifts before they impact the end-user, a relationship fundamental to our target parameter selection.

It is an important observation before modeling that HARQ counters indicate the upcoming radio environmental changes. Our primary findings during the exploratory data analysis was the importance of HARQ counters especially for 16QAM to 256QAM.

A. Bitrate Prediction




Fig. 3: Delivery Rate Prediction using RandomForest [$W = 5, H = 5, F = 3$], RMSE = 1666.94 [UE UL], 1557.86 [CELL UL], MAE = 1373.01 [UE UL], 1226.74 [CELL UL]

As illustrated in Fig. 3, the proposed algorithms effectively characterize transient degradation and negative peaks in the delivery rate for a forecast horizon of $F = 3$ (and 15 seconds, with $W = 5$). Comparative analysis demonstrates that both cell-level and UE-level modeling approaches achieve high predictive accuracy. The feature importance analysis, detailed in Fig. 4, identifies modulation ACK and NACK signaling as critical determinants of the predictive output. Notably, several features exhibit higher importance at a temporal lag of $lag = 2$ compared to their instantaneous values. This lead-lag relationship confirms that bitrate degradation originates in the radio layer and it is observable. The HARQ counters provide a predictive time window, which is dependent on the buffer size as well. This allows the models to preemptively identify performance drops, and possibly notify other network components accordingly.

B. TTLB and RTT prediction

Predicting TTLB and RTT necessitates a slightly different modeling approach due to their fluctuations compared to the more persistent trends observed in delivery rates. For these parameters we used $W = 10$ against stochastic variances and to ensure operational reliability in the resulting forecasts.

While the delivery rate allows for stable, long-term trend tracking, the TTLB and RTT models are specifically tuned for peak detection. This ensures that critical service degradations are captured with precision, providing a robust trigger for proactive network management. The resulting time-series forecasts for TTLB and RTT are illustrated in Figs. 6 and 7, respectively. Furthermore, the SHAP analysis for the UE-level

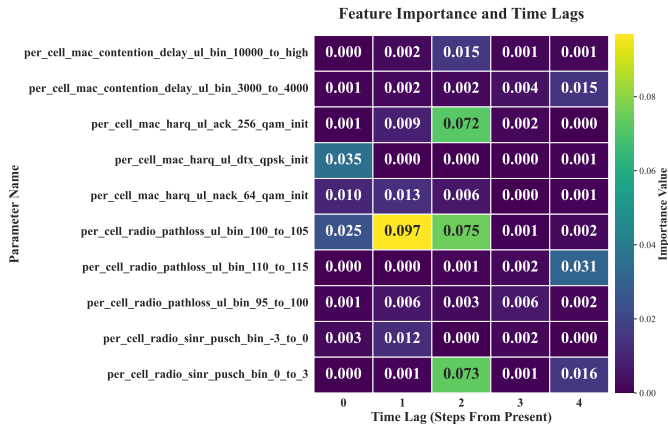


Fig. 4: Top CELL level Delivery Rate prediction feature importance with time lags [$W = 5, H = 5, F = 3$]

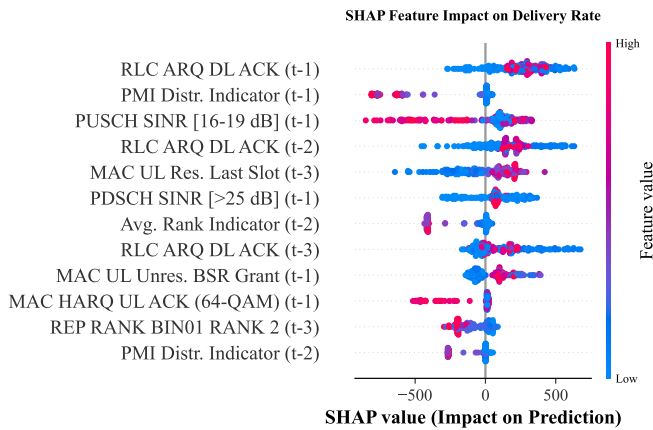


Fig. 5: SHAP values for UE level model on Del Rate Prediction ($W = 5, H = 3, F = 3$)

TTLB prediction, as shown in Fig. 8 confirms the model's reliance on granular ACK messages for other parameters than delivery rate.




Fig. 6: TTLB Prediction using XGBoost [$W = 10, H = 3, F = 1$], RMSE = 97.46 [UE UL], 92.58 [CELL UL]

C. Modeling conclusion

While aggregate performance metrics provide a foundational baseline, the operational integrity of telecommunication forecasting models is truly validated by their performance




Fig. 7: RTT Prediction using XGBoost [$W = 10, H = 3, F = 1$], RMSE = 21.49 [UE UL], 17.63 [CELL UL]

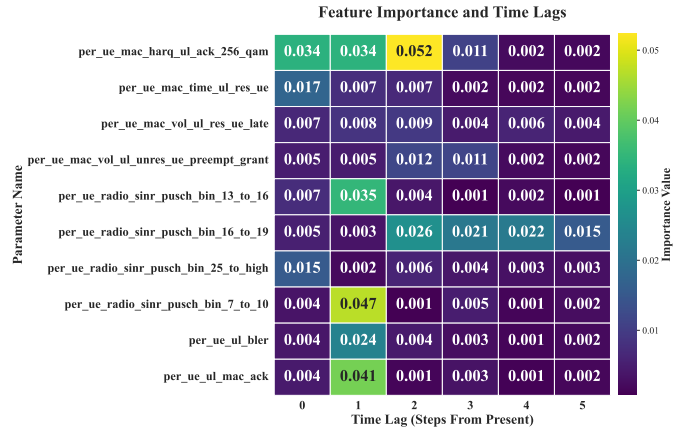


Fig. 8: Top UE Level TTLB feature importance with time lags [$W = 10, H = 3, F = 1$]

during critical network transients. In these high-variance regions, predictive accuracy serves as a definitive trigger for proactive network intervention. Consequently, our evaluation framework prioritizes segment-specific analysis, ensuring the model captures transient peaks and throughput drops that are often numerically suppressed by traditional R^2 and $RMSE$ calculations during prolonged stable states. The quantitative results for the complete test set are summarized in Table I. As demonstrated, both the UE-level and Cell-level models tend to outperform the Baseline models in most cases. This confirms the power of high-resolution CTRs, especially HARQ counters, which capture complex physical-layer dependencies that signal strength values fail to represent. The developed predictive capability facilitates two primary network optimization use cases. For real-time traffic, our model provides proactive QoS forecasting, utilizing bitrate as a primary indicator of link stability. Simultaneously, for buffered services, the approach enables the characterization of buffer depletion trends to preempt stall events—demonstrating that the model effectively translates raw RAN data into actionable application-level intelligence.

V. CONCLUSION

This study demonstrated the efficiency of machine learning models in predicting critical network performance metrics, specifically delivery rate, RTT, and TTLB, across both UE

TABLE I: Model Performance Comparison: XGBoost (XGB) vs. Random Forest (RF) across Base, UE, and CELL configurations. ([ms] , for 10 second W time windows)

Target	Metric	Base		UE		CELL	
		XGB	RF	XGB	RF	XGB	RF
TTLB	R^2	0.10	0.08	0.19	0.18	0.11	0.12
	RMSE	105.00	108.30	97.40	99.40	92.60	109.80
RTT	R^2	0.12	0.13	0.22	0.21	0.19	0.21
	RMSE	31.00	29.80	21.50	22.50	17.60	18.40

TABLE II: Final Test Metrics Comparison for Delivery Rate [kbps].

Configuration	R^2	RMSE	MAE	MSE
Baseline	-0.3150	2139.85	1581.33	4.58e+06
UE	0.2020	1666.94	1373.01	2.78e+06
CELL	0.3031	1557.86	1226.74	2.43e+06

and cell-level granularities from CTR records. Our evaluation reveals that predictive accuracy is highly sensitive to the temporal configuration of the model; specifically, a 15 second prediction horizon (F) offers the most robust balance between foresight and precision for delivery rate and 10 second forecast horizon F works best for RTT and TTLB. Furthermore, the results indicate that increasing the historical look-back window (H) consistently yields improvements in model convergence and error reduction.

The feature importance analysis underscored the critical role of modulation-specific signaling, where MAC-layer HARQ ACK/NACK messages for high-order modulation (256-QAM) emerged as the primary predictors. Notably, the significance of lagged (historical) features showcase the temporal relationship between radio layer signaling and client-side reporting, enabling the model to preemptively identify throughput degradation or delays.

These findings benefit both application and network layers. On the network side, the predictions could enable preventive internal actions such as automated traffic shaping and dynamic 5QI policy adjustments, allowing operators to act on forecasted degradations before they manifest in transport KPIs or end-user experience. On the application side, the framework can signal predicted quality changes through in-band mechanisms such as L4S [16] marking and SCONE [17] signaling, or by exporting QoS notifications to northbound interfaces such as CAMARA [1] or Network Exposure Functions (NEF)—enabling applications such as the live broadcast workflow presented in this work to adapt proactively before quality is impacted. As the telecommunications industry moves toward exposing network intelligence as standardized, consumable services, the predictive capabilities demonstrated in this work represent a concrete building block for this vision.

ACKNOWLEDGMENT

This project has been implemented with the support provided by the Ministry of Culture and Innovation of Hungary

from the National Research, Development and Innovation Fund, financed under the KDP-2023 funding scheme. Project no. 153148 have been implemented with the support provided from the National Research, Development and Innovation Fund of Hungary, financed under the NKKP STARTING_25 funding scheme.

REFERENCES

- [1] CAMARA Project, “CAMARA — The Telco Global API Alliance,” 2024. [Online]. Available: <https://camaraproject.org/>
- [2] O-RAN, [Accessed 01-04-2026]. [Online]. Available: <https://www.o-ran.org/>
- [3] S. Afshar, R. Razavi, and M. Moshirpour, “Closed-loop clustering-based global bandwidth prediction in real-time video streaming,” *IEEE Transactions on Machine Learning in Communications and Networking*, vol. 3, pp. 448–462, 2025.
- [4] I. Alzalam, C. Lipps, and H. D. Schotten, “Time-series forecasting models for 5g mobile networks: A comparative study in a cloud implementation,” in *2024 15th International Conference on Network of the Future (NoF)*, 2024, pp. 54–62.
- [5] A. Biernacki, “Throughput prediction of 5g network based on trace similarity for adaptive video,” *Applied Sciences*, vol. 14, no. 5, 2024. [Online]. Available: <https://www.mdpi.com/2076-3417/14/5/1962>
- [6] M. Biswas, A. Chakraborty, and B. Palit, “A kalman filter based low complexity throughput prediction algorithm for 5g cellular networks,” *IEEE Transactions on Vehicular Technology*, vol. 73, no. 5, pp. 7089–7101, 2024.
- [7] G. Lv, Q. Wu, Q. Tan, W. Wang, Z. Li, and G. Xie, “Accurate throughput prediction for improving qoe in mobile adaptive streaming,” *IEEE Transactions on Mobile Computing*, vol. 23, no. 5, pp. 5799–5817, 2024.
- [8] T. Azmin, M. Ahmadinejad, and N. Shahriar, “Bandwidth prediction in 5g mobile networks using informer,” in *NoF*, 2022, pp. 1–9.
- [9] M. Helmy, M. S. Hassan, M. H. Ismail, and U. Tariq, “Autoformer-based mobility and handoff-aware prediction for qoe enhancement in adaptive video streaming in 4g/5g networks,” *Journal of Network and Computer Applications*, p. 104324, 2025.
- [10] J. Lv, Y. Lin, M. Hou, Y. Li, Y. Gao, and W. Dong, “Accurate bandwidth and delay prediction for 5g cellular networks,” *ACM Trans. Internet Technol.*, vol. 25, no. 2, Apr. 2025. [Online]. Available: <https://doi.org/10.1145/3703629>
- [11] K. Horita, N. Fukumoto, and A. Nakao, “Improving qos of 5g video streaming through network exposure function,” in *GLOBECOM 2023 - 2023 IEEE Global Communications Conference*, 2023, pp. 375–380.
- [12] A. Hernandez and F. S. Barbosa, “An end-to-end testbed for communication, compute, and control co-design: The kista innovation park,” in *Computer Safety, Reliability, and Security: SAFECOMP 2025 Workshops*, M. Törngren, B. Gallina, E. Schoitsch, E. Troubitsyna, and F. Bitsch, Eds. Cham: Springer Nature Switzerland, 2026, pp. 5–16.
- [13] A. Benteleb, S. Akbari, A. C. Begen, and R. Zimmermann, “godash: A headless dash player for advanced experimentation,” in *Proceedings of the 11th ACM Multimedia Systems Conference (MMSys)*, 2020, pp. 307–312.
- [14] “3GPP TS 23.501,” 2015. [Online]. Available: <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>
- [15] ITU, “P.1203 Parametric bitstream-based quality assessment of progressive download and adaptive audiovisual streaming services over reliable transport — itu.int,” <https://www.itu.int/rec/T-REC-P.1203>, [Accessed 24-03-2026].
- [16] K. De Schepper, M. Bagnulo, and G. White, “Rfc 9330: low latency, low loss, and scalable throughput (l4s) internet service: architecture,” 2023.
- [17] M. Ihlar and M. Boucadair, “Standard Communication with Network Elements (SCONE) Protocol,” Working Draft, IETF Secretariate, Internet-Draft draft-ietf-scone-protocol-04, Mar. 2025. [Online]. Available: <https://datatracker.ietf.org/doc/draft-ietf-scone-protocol-04/>