

Machine Learning for Predictive Base Station Sleep Modes to Enhance 5G Greenness

Sahar Ebadinezhad^{1,2}

¹ Department of Computer Information Systems, Faculty of Artificial Intelligence & Informatics, Near East University, Cyprus

² Computer Information Systems Research and Technology Center (CISRTC), Near East University, Cyprus

Email: sahar.ebadinezhad@neu.edu.tr

Manuscript received May 7, 2026; revised June 14, 2026; accepted August 27, 2026; published October 19, 2026

Abstract—The rapid deployment of 5G mobile communications has drastically increased the energy consumption of radio access networks, and Base Stations (BSs) account for 60%–80% of network energy consumption. This work presents a low-complexity predictive base station sleep mode framework based on Long Short-Term Memory (LSTM) traffic forecasting and a threshold-based control loop. It is based on proactive forecasting that is able to automatically transition BSs to sleep mode during low-demand state transitions in the predicted low-demand periods of the network, thereby minimizing frequent state-switching overhead and extending low-power intervals. Based on the real-world Telecom Italia Mobile Big Data Challenge data set (Milan and Trentino), our framework realized energy savings of up to 45.8% over the always-on baseline and 32.0% over the reactive baseline. Most importantly, our system does not breach Quality of Experience (QoE) service at the optimal 100 Mbps speed, and it requires much less computational effort than state-of-the-art reinforcement learning structures. A multi-objective sensitivity analysis presents the Pareto front to capture energy efficiency and service reliability. These results indicate that the lightweight predictive engineering frameworks can be a very practical and scalable way to manage the green 5G networks.

Keywords—base station sleep mode, energy efficiency, Long Short-Term Memory (LSTM), green communications, machine learning, sustainability in Information and Communication Technology (ICT), traffic prediction, 5G networks

I. INTRODUCTION

The rapid introduction of mobile broadband services and widespread use of 5G mobile networks have drastically increased the energy consumption of mobile networks. Base Stations (BSs) constitute about 60%–80% of the total network energy usage and are the main source of operational costs and carbon emissions in cellular networks [1, 2]. Thus, reducing BS energy consumption is a crucial research challenge to ensure a sustainable mobile network. In sleep mode systems, BSs can be made to sleep at low power levels when traffic is low. Traditional reactive sleep mode systems are based on real-time traffic information to trigger sleep and wake-up transitions. To model and evaluate these traffic-driven sleep cycles accurately, empirical spatial base station topologies and real-world cellular data traces are typically gathered from open network platforms [3] and community data archives [4]. While highly effective when traffic is very low and less frequent in off-peak hours, systems like this are not able to react quickly to traffic fluctuations or sleep-wake cycling and are at risk of losing Quality of Experience (QoE) when traffic is high [5]. This has led to the development of predictive and data-driven systems to predict future traffic and to control BS states proactively. By leveraging standardized open Radio Access Network (RAN) network

measurements and trace profiles [6], sophisticated machine learning architectures can effectively capture dynamic demand fluctuations. Deep learning networks, particularly Long Short-Term Memory (LSTM) networks, are proven to be particularly effective for determining time dependencies in cellular traffic time series [7–9].

Many data-based strategies have been proposed to optimize Next Generation NodeB (gNodeB) energy efficiency, such as DRL-based predictive sleep policies [10], bandit algorithms for small cells [11], empirical baseline frameworks [12, 13], and recent federated or slicing-aware traffic estimators [14, 15]. However, the existing approaches either create risks for online convergence or lack lightweight, mathematically bounded control laws to be applied to real-time edge execution.

However, LSTM-based sleep control systems are currently either based on complex optimization algorithms such as the Lyapunov optimization method [16], mixed-integer optimization method [17], and spatiotemporal Deep Reinforcement Learning (DRL) methods [10] that are computationally expensive. All these challenges in practice limit the application of Sleep-Decision Modules (SDMs) under real-time Radio Access Network (RAN) conditions. In this paper, we propose a lightweight predictive BS sleep mode framework that combines LSTM short-term traffic forecasting with adaptive threshold-based sleep control. Instead of presenting a completely new machine learning mathematical approach, our contribution is the development of an architectural synthesis and optimization of a low-complexity engineering framework to work in the context of real-world 5G RAN constraints. In particular, we summarize these contributions as follows:

- Developing a lightweight LSTM traffic forecaster optimized for precise 5G load prediction with minimal computational overhead.
- Designing an actionable threshold-based control loop that converts traffic forecasts into proactive sleep-mode decisions, achieving high energy efficiency without the need for sophisticated AI models.
- Using extensive validation on the real-world Milan and Trentino cellular datasets, we establish a multi-objective Pareto front that can be leveraged to create a model for operators to achieve energy savings between 43.9% and 50.4% with zero QoE service degradation.
- We provide a detailed complexity-performance comparison with state-of-the-art reinforcement learning frameworks, highlighting our framework's better deployment readiness for latency and computational overhead.

II. LITERATURE REVIEW

Energy efficiency in 5G Radio Access Networks (RANs) is a well-known challenge. The base stations consume 70%–80% of total mobile network energy [1, 2], and this level is expected to increase with the advent of ultra-dense heterogeneous deployments, massive Multiple-Input Multiple-Output (MIMO), and millimeter-wave technologies. Reducing this expenditure and not compromising user QoE is a key goal for the wireless industry. The early research into BS energy reduction employed on reactive sleep mode strategies in which BSs are switched to low-power states based on real-time traffic data [3–5]. In this way, BSs are able to save energy during off-peak periods but are not able to predict when traffic will change in real-time. As a result, reactive systems are prone to frequent state transitions, high signaling overhead, and QoE degradation risk [5]. A classical statistical forecasting approach such as AutoRegressive Integrated Moving Average (ARIMA) is employed to predict mobile traffic [6]. However, ARIMA is linear and cannot capture the non-linear, bursty, and spatiotemporally correlated features of real cellular traffic [6].

This has led to the implementation of deep learning models such as LSTM networks to learn long-term temporal dependencies in sequential data. In BS sleep mode control, Jang *et al.* [16] suggested an LSTM-based user traffic prediction model with Lyapunov optimization for BS switching and sleep mode selection. Their approach is effective, but it requires solving a complex optimization problem at each decision step and is computationally complex and not feasible for real-time deployment. Zhu and Wang [17] proposed a joint LSTM traffic prediction and BS sleeping algorithm for energy saving, also using spatiotemporal traffic data. Their work is closely related to ours, though it targets a multi-BS network and does not provide systematic threshold sensitivity analysis. Wu *et al.* [10] proposed the Deep Base Station Control (DeepBSC) framework, integrating graph-based spatiotemporal prediction with a DRL agent operating on the same Telecom Italia Milan dataset. DeepBSC realize approximately 45% energy savings but requires a full DRL training loop, reward engineering, and convergence over thousands of episodes—significant deployment barriers.

More recently, reinforcement learning and multi-agent approaches have been proposed for large-scale BS energy management [13, 14], demonstrating higher energy savings (54–55%) at the cost of very high computational complexity [13, 14]. In contrast, standard architectural frameworks such as 3GPP TS 38.300 prioritize deterministic, low-latency Radio Access Network (RAN) split operations and standardized control signaling [18]. Sun *et al.* [19] propose a deep learning-based traffic-aware sleep mode strategy with Reconfigurable Intelligent Surface (RIS)-aided cell zooming in multi-cell networks. Park and Yoon [20] present a federated LSTM-based traffic prediction framework for BS switching. Jha *et al.* [21] demonstrate the effectiveness of ML techniques for optimizing 5G sustainable network performance in urban environments. Very recent work (2025–2026) on transformer-based cellular traffic prediction [22] and Multi-Agent Reinforcement Learning (MARL)-based mmWave sleep optimization [23] confirms

sustained research interest in lightweight yet accurate sleep control.

Despite this progress, a clear gap remains between complex, high-performance DRL approaches and simple reactive baselines. The present work fills this gap by demonstrating that a well-regularized LSTM model with threshold-based sleep control attain competitive energy savings on a publicly available real-world dataset with zero QoE violations, offering a practically feasible, low-overhead alternative to high-complexity RL frameworks.

III. METHODOLOGY

The proposed framework for reducing the energy consumption of 5G base stations by proactive sleep mode management without compromising QoE is based on four components: (i) traffic data collection and preprocessing; (ii) LSTM-based traffic prediction model development; (iii) threshold-based sleep mode decision algorithm; and (iv) energy consumption and QoE performance evaluation.

A. Mathematical Formalization of the Traffic Regime and Predictor State Optimization

To lay the foundation for predictive sleep-mode control, we model the traffic in the mobile network as a discrete-time non-stationary stochastic process $X(t)$ at regular intervals $t \in \{1, 2, \dots, T\}$ and the total demand $X(t)$ over time can be divided into a deterministic, time-varying diurnal trend component $D(t)$, which represents human activity, and a high-frequency (zero-mean) residual component $Z(t)$ as a localized burstiness:

$$X(t) = D(t) + Z(t) \quad (1)$$

where $D(t) = D(t + \tau_{period})$ determines the long-term structural seasonality over a 24-hour period τ_{period} and $E[Z(t)] = 0$. The goal of the predictive front-end is to construct a mapping function $f: X_{history} \rightarrow \hat{X}(t + \Delta t)$ using a lookback window of length k to predict the future demand volume over a lookahead horizon Δt . Let $\mathbf{x}_t = [X(t-k+1), X(t-k+2), \dots, X(t)]^T$ define the historical feature vector. The network maps $\mathbf{x}_t$ to hidden state representations $\mathbf{h}_t \in \mathbb{R}^d$ through a set of deterministic recurrent gating.

$$\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{x}_t + \mathbf{U}_f \mathbf{h}_{t-1} + \mathbf{b}_f) \quad (2)$$

$$\mathbf{i}_t = \sigma(\mathbf{W}_i \mathbf{x}_t + \mathbf{U}_i \mathbf{h}_{t-1} + \mathbf{b}_i) \quad (3)$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_o \mathbf{x}_t + \mathbf{U}_o \mathbf{h}_{t-1} + \mathbf{b}_o) \quad (4)$$

where $\mathbf{f}_t$, $\mathbf{i}_t$, and $\mathbf{o}_t$ represent the forget, input, and output gates, respectively, and $\sigma(\cdot)$ denotes the element-wise sigmoid activation function. The updated memory cell state $\mathbf{c}_t$ and hidden state vector $\mathbf{h}_t$ are defined as:

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tanh(\mathbf{W}_c \mathbf{x}_t + \mathbf{U}_c \mathbf{h}_{t-1} + \mathbf{b}_c) \quad (5)$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t) \quad (6)$$

The predicted estimator $\hat{X}(t + \Delta t)$ is then obtained by a

linear projection layer $\mathbf{W}_y \mathbf{h}_t + \mathbf{b}_y$. To ensure the estimator can track the velocity in a fluid traffic regime, the internal parameter space $\boldsymbol{\theta} = \{\mathbf{W}, \mathbf{U}, \mathbf{b}\}$ is optimized by minimizing the structure of the MSE objective function:

$$\min_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}) = \frac{1}{N} \sum_{n=1}^N (X_n(t + \Delta t) - \hat{X}_n(t + \Delta t))^2 + \lambda \|\boldsymbol{\theta}\|_2^2 \quad (7)$$

where $\lambda \|\boldsymbol{\theta}\|_2^2$ is the L_2 regularization penalty that limits the high-frequency noise component $Z(t)$. The predicted estimator $\hat{X}(t + \Delta t)$ is the baseline control law for the 5G base station $S(t + \Delta t) \in \{\text{Active}, \text{Sleep}\}$, and if we have T_{thr} , the prediction of the predicted estimator is that $S(t + \Delta t)$ is a deterministic function controlled by the safety threshold T_{thr} :

$$S(t + \Delta t) = \begin{cases} \text{Sleep}, & \text{if } \hat{X}(t + \Delta t) \leq T_{thr} \\ \text{Active}, & \text{if } \hat{X}(t + \Delta t) > T_{thr} \end{cases} \quad (8)$$

Thus, by binding the control law $S(t + \Delta t)$ to the robustly optimized time-series estimator rather than real-time reactive signals, we can establish optimal sleep windows. The technical novelty of the new model is that this unified model replaces unstable online reinforcement exploration with a mathematically bounded and computationally inexpensive predictive policy.

B. Dataset Description & Traffic Data Preprocessing

Our experiments are conducted on the Telecom Italia Mobile Big Data Challenge dataset [1], which is a collection of aggregated and anonymized data of cellular traffic in the City of Milan and the Province of Trentino, Italy. The publicly available multi-source data have long been utilized for mobile traffic analysis and energy-saving work [9, 10], and the benchmarking process is replicable. The raw traffic time series is preprocessed as follows. First, min-max normalization is applied to the traffic volume at $[0, 1]$ to avoid feature dominance due to magnitude differences [7, 8]. Second, normalized data are transformed into supervised learning samples with an overlapping sliding window of size $W = 10$ time steps, i.e., one step ahead (prediction horizon $H = 1$). Third, the data is processed into 80% training and 20% testing subsets, and no testing data is included in the training to ensure unbiased testing.

C. LSTM Traffic Prediction Model

A Long Short-Term Memory (LSTM) neural network is leveraged in short-term cellular traffic forecasting. LSTM networks have long-term temporal dependence and minimize the vanishing gradient problem that is typical of recurrent neural networks. We use one LSTM layer with 64 units, a dropout layer (rate = 0.2), and a fully connected output layer. To avoid overfitting, we use two regularization steps: (1) dropout (rate = 0.2) after the LSTM layer and (2) L2 weight regularization ($\lambda = 1 \times 10^{-4}$) for LSTM weights. Early stopping with patience = 5 epochs stops training if validation loss is not improved. In this way, training and validation losses converge in tandem, as shown in Fig. 1, demonstrating convergence without overfitting after application of dropout and L2 regularization. The model is trained using the Adam optimizer [8] with learning rate = 0.001 and batch size = 32

for low Mean Squared Error (MSE) loss. Performance is evaluated using MAE, the Root Mean Squared Error (RMSE), and the Mean Absolute Percentage Error (MAPE) on the held-out test set, averaged over five independent runs with 95% confidence intervals.




Fig. 1. Training and validation loss (MSE) of the LSTM model over epochs with early stopping (patience = 5).

D. Proposed Algorithm Structure & Experimental Setup

The model configuration and experimental setup resulting from our technique is summarized in Table 1.

Table 1. Experimental parameters and model configuration

Parameter/Configuration	Value/Description
Dataset	Telecom Italia Milan & Trentino [1]
Data Split	80% Training/20% Testing
Normalization	Min-Max Scaling [0, 1]
Sliding Window (W)	10 time steps
Prediction Horizon (H)	1 step ahead
Model Architecture	LSTM, 1 hidden layer, 64 units
Dropout Rate	0.2
L2 Regularization (λ)	1×10^{-4}
Optimizer	Adam
Learning Rate	0.001
Batch Size	32
Loss Function	Mean Squared Error (MSE)
Early Stopping Patience	5 epochs
Evaluation Metrics	MAE, RMSE, MAPE
Traffic Thresholds (T_{thr})	50, 100, 150, 200 Mbps
Active-mode Power (E_{Active})	130 W [18]
Sleep-mode Power (E_{Sleep})	10 W [18]
Training Environment	Python 3.9, TensorFlow 2.x, Intel Core i7, 16 GB RAM
Statistical Replication	5 independent runs; 95% CI reported

E. Energy Consumption and Base Station State Model

1) Energy consumption model

The total energy consumption of the base station is modelled as:

$$E_{Total} = E_{Active} \times t_{active} + E_{Sleep} \times t_{sleep} \quad (9)$$

where $E_{Active} = 130$ W and $E_{Sleep} = 10$ W represent active-mode and sleep-mode power consumption, respectively, consistent with the 3rd Generation Partnership Project (3GPP) macro-cell power model [18]. The variables

t_{active} and t_{sleep} denote cumulative time durations (in hours) in active and sleep states.

The Energy Saving (ES) percentage is defined as:

$$ES(\%) = \frac{E_{Always-On} - E_{Framework}}{E_{Always-On} \times 100} \quad (10)$$

where $E_{Always-On} = 30.50$ kWh is the reference energy of the always-on baseline.

2) Standardized base station energy and operational state model

We now formalize the energy consumption of the 5G base station (gNodeB) using the standard 3GPP/EARTH linear power model. The total power consumption $P_{total}(t)$ depends on the active state $S(t) \in \{Active, Sleep\}$ and the initial traffic use $\rho(t) \in [0, 1]$:

$$P_{total}(t) = \begin{cases} P_{baseline} + \Delta_{slope} \cdot \rho(t) \cdot P_{max}, & \text{if } S(t) = Active \\ P_{sleep}, & \text{if } S(t) = Sleep \end{cases} \quad (11)$$

where $P_{baseline}$ is the static power consumed by the baseband processing unit, cooling fans, and power supply elements under zero traffic, and Δ_{slope} is the dynamic power amplifier scaling to full load P_{max} . The deep sleep state power consumption is fixed at a very low baseline P_{sleep} , which is a 5G Advanced Sleep Mode (ASM) in which the power-hungry Radio Frequency (RF) components are progressively deactivated. We set $P_{max} = 711$ W, $P_{baseline} = 354$ W, $\Delta_{slope} = 2.15$ and $P_{sleep} = 12.5$ W. We use the same macro-gNodeB footprint for the whole system.

F. Sleep Mode Decision Algorithm

Algorithm 1 summarizes the proposed predictive sleep mode control framework. The algorithm proceeds in four phases. In Phase I, raw traffic data are collected, normalized, and formed into sliding-window supervised samples. In Phase II, the LSTM model is trained with dropout, L2 regularization, and early stopping. In Phase III, the trained LSTM predicts T_{pred} for each time step; if $T_{pred} < T_{thr}$, the BS enters Sleep Mode; otherwise it remains Active. In Phase IV, total energy consumption, energy savings ES, QoE violation count Q_{events} , and Throughput Degradation Ratio (TDR) are computed.

Algorithm 1: Predictive Base Station Sleep Mode Framework

Input: Raw historical cellular traffic time-series T ; Energy consumption parameters $E_{Always-On}$, E_{Sleep} ; *Traffic threshold* (T_{thr})

Output: Predicted Sleep Mode status; Total energy consumption (E_{Total}); Energy Savings (ES); QoE violation count (Q_{events}).

Phase I: Traffic Data Collection and Preprocessing

Collect raw cellular traffic data from the Milan & Trentino datasets.

Apply Min-Max normalization to scale traffic volume data to a range of $[0, 1]$.

Construct supervised learning samples using an overlapped sliding window approach.

Divide the processed data into an 80% training set and a 20% testing set.

Phase II: LSTM-based Traffic Prediction Model Development

Initialize an LSTM network with one hidden layer containing

64 units.

Include a dropout layer with a rate of 0.2 to prevent model overfitting.

Train the model using the Adam optimizer to minimize the Mean Squared Error (MSE).

Evaluate prediction accuracy using MAE, RMSE, and MAPE metrics.

Phase III: Decision Algorithm for Sleep Mode Prediction

Input real-time traffic features into the trained LSTM model to generate the prediction T_{pred} (PredictedTraffic)

If $T_{pred} < T_{thr}$:

Transition the Base Station (BS) to **Sleep Mode** to reduce power consumption.

Else

Maintain the BS in **Active Mode** to meet demand.

End If

Phase IV: Performance Simulation and Evaluation

Calculate the total energy consumption ($E_{Predictive}$) based on the cumulative duration of active and sleep states.

Determine the **Energy Savings (ES)** percentage relative to the Always-On baseline using the following formula:

$$ES = ((E_{Always-On} - E_{Predictive}) / E_{Always-On}) \times 100$$

Evaluate the framework performance, which resulting in a baseline energy saving of 45.8% at the primary operational threshold ($T_{thr} = 100$ Mbps).

Monitor Quality of Experience (QoE) by counting events where traffic demand exceeds BS capacity during sleep intervals.

Validate the Energy-QoE tradeoff using a Pareto Front to ensure Q_{events} remains below the maximum acceptable limit (10).

Return ES (ranging from 43.9% to 50.4% depending on threshold) and Q_{events} (targeting zero violations)

G. QoE Evaluation

A QoE violation is formally defined as any sampling interval in which: (i) the base station is in sleep mode, and (ii) the observed traffic demand $T(t)$ exceeds T_{thr} before the wakeup transition completes. Two supplementary QoE metrics are evaluated: (a) estimated wakeup-induced latency overhead per violation, based on 3GPP NR sleep mode wakeup time of 5–20 ms [18]; and (b) the Throughput Degradation Ratio (TDR), defined as the fraction of total traffic demand not served during sleep intervals. Borderline cases ($T(t)$ within $\pm 5\%$ of T_{thr}) are also analyzed separately.

IV. RESULTS

The current section introduces the experimental results achieved through the assessment of the predictive base station sleep mode mechanism. All the results are generated based on the Telecom Italia Milan and Province of Trentino dataset. The predictive base station sleep mode mechanism is evaluated in terms of accuracy of predicted traffic, energy efficiency, and robustness to the threshold value. The predictive mechanism is compared with the traditional reactive mechanism as well as the always-on mechanism. The predictive mechanism is based on LSTM.

A. Traffic Prediction Accuracy

The optimization profile and tracking accuracy of the proposed LSTM framework are evaluated on both training

and testing data with the parameters in Table 1. The network shows the convergence of the Mean Square Error (MSE) and the RMSE on consecutive training epochs in Fig. 1. The model decays in the same way in both training and validation loss and achieves the early stopping optimization at epoch 10 with no structural overfitting or high-frequency divergence. Once the model is deployed on the data (held-out test data), the model's absolute tracking accuracy is measured. The LSTM framework achieves MAE of 95.36 Mbps and RMSE of 128.42 Mbps on the Milan macrotrace. To further validate these metrics and provide a relative level of performance, we assess the Mean Absolute Percentage Error (MAPE) and find an error rate of 11.84% and 13.12% for the urban Milan and Trentino regions, respectively. To understand the variance of the errors, Figs. 2 and 3 plot the Probability Density Function (PDF) and frequency distribution of the prediction errors for both against a classical ARIMA baseline. Note that, as shown in Fig. 2, while the baseline ARIMA has a fixed localized variance ($\sigma = 100.17$), the LSTM gives a consistent structural offset. In the regularized LSTM, the error density envelope is minimized around the variance values ($\mu = 41.28$, $\sigma = 123.69$). Since the absolute tracking errors are measured in this Gaussian-like distribution, the threshold control loop downstream is protected from outlier capacity decrease during transition windows, and the QoE violations are zero.

The ARIMA is a linear model that systematically underestimates traffic peaks by regression to the mean [6] as a policy to suppress variance. The LSTM model tracks the steering of traffic load transitions (the rises and falls that are operationally crucial for proactive sleep mode activation). Even with extensive variance, the LSTM predicts the peak of the traffic in a way that makes it possible for early BS sleep transitions to occur and for early wake-ups in the early morning of the day.

Here, we prove the predictive performance of the proposed LSTM architecture by using the three statistical measures: the Mean Absolute Error (MAE), the Root Mean Squared Error (RMSE) and the Mean Absolute Percentage Error (MAPE). Each of these measures is used to identify the specific performance profile needed for secure radio edge scheduling. The results for the measured LSTM performance are given in Table 2.

As shown in Table 2, the model shows a good balance in tracking. In the dense city of Milan, the global MAE is 95.36 Mbps, and the corresponding RMSE is 128.42 Mbps. The

small gap between the MAE and RMSE confirms that the structure-regularized LSTM optimization effectively takes care of high-amplitude outlier errors and prevents the framework from experiencing sudden and unexpected traffic surges that can lead to erroneous sleep cycles.




Fig. 2. Distribution of prediction errors vs probability density for the LSTM and ARIMA models.




Fig. 3. Distribution of prediction errors vs frequency for the LSTM and ARIMA models.

The global MAPE also stays at 11.84% for the Milan trace and 13.12% for the Trentino trace. This relative percentage accuracy shows that our forecasting model remains a tight tracking envelope for both high-volume daytime workloads and ultra-low-capacity late-night intervals that are consistent with the multi-metric criteria claimed throughout the system' topology.

Table 2. Multi-metric traffic prediction accuracy profile

Evaluated Dataset Profile	Mean Absolute Error (MAE)	Root Mean Squared Error (RMSE)	Mean Absolute Percentage Error (MAPE)
Milan Macrotrace (Urban Region)	95.36 Mbps	128.42 Mbps	11.84%
Trentino Trace (Provincial Region)	42.15 Mbps	58.74 Mbps	13.12%

B. Energy Consumption Analysis

To evaluate the operational flexibility of the proposed framework, performance metrics are analyzed across distinct traffic bounds; specifically, the predictive framework yields a net energy saving of 43.9% under a baseline threshold ($T_{thr} = 50$ Mbps), which scales to 45.8% at the optimal operating point ($T_{thr} = 100$ Mbps).

The main target of the proposed framework is minimizing the energy consumption of the BS while ensuring an acceptable level of QoE. In order to estimate the energy

efficiency, three modes of operation (always on, reactive sleep mode, and the proposed predictive sleep mode framework) have been compared.

The predictive model for sleep mode saves 43.9% relative to the always-on mode, and it clearly outperforms the reactive approach. The improvement is mainly due to the proactive mechanism of the proposed method, where it becomes possible for the base station to turn off earlier and remain in low power states for longer continuous time durations. As a result, it prevents unwanted energy consumption for frequent

transition states among different modes, as has been emphasized in energy efficiency research in the past [3, 5, 10].

Fig. 4 compares total energy consumption under three operational policies. The always-on baseline consumes 30.50 kWh (reference). The reactive sleep mode scheme realize 32.0% energy savings (20.73 kWh) by transitioning the BS to sleep based on observed traffic dropping below $T_{thr} = 100$ Mbps. The proposed predictive (LSTM-based) framework attains 45.8% energy savings (16.52 kWh) compared to the reactive baseline by 43.1%. This is because the LSTM is able to predict the traffic decline and execute sleep mode transitions earlier to provide longer uninterrupted low-power intervals and fewer wake-up events. These findings are consistent with recent literature demonstrating that predictive sleep control strategies outperform reactive approaches in urban cellular environments with strong temporal regularities [16, 17, 19].




Fig. 4. Comparative analysis of total energy consumption (kWh) under three operational policies.

To make sure the observed energy savings and no QoE violations are statistically significant instead of only a feature of the offline evaluation window, a rolling variance analysis was performed over the prediction epochs over the Milan and Trentino trace sets for a total of 10,000 times. The 95% Confidence Intervals (CI_{95}) across all these prediction epochs reveal that under the model ($T_{thr} = 100$ Mbps) the average energy saving of 45.8% is well defined and narrow ($CI_{95} = [45.62\%, 45.98\%]$). In addition, the capacity deficit parameter (i.e. traffic under-predictions that would cause QoE drops) is determined by a mean value of 0.000 Mbps ($\mu + 3\sigma = 0.0002$ Mbps) when traffic levels are high during peak hours and it is clear ($p < 0.001$) that forecasting errors are isolated from critical areas of the network and there is no shortage of capacity.

C. Threshold Sensitivity Analysis

For the purpose of assessing the robustness of the proposed framework, the threshold sensitivity analysis is performed. In the sensitivity analysis, the threshold value T_{min} is varied to see its impact on the framework.

The results shown above suggest that with an increased threshold value, there are higher levels of energy savings because the BS is allowed to sleep for a prolonged time. At a lower threshold value, there is a conserved approach taken by the predictive model with a slight energy saving and no

disruption to services. An increased value of the threshold reveals a remarkable improvement in energy saving with a value above 43.9% without violating any service aspects.

In particular, for all considered threshold levels, it is noticed that the predictive mode framework maintains a late wake-up event counter of zero, thus signifying that the traffic forecasts obtained using the LSTM are accurate enough to avoid degrading service. Such a characteristic differentiates it from a reactive approach, where QoE might take place just before an increase in traffic and result in higher latency and possible violations in QoE [10, 12].

These results for threshold sensitivity also illustrate the significance of threshold selection in predictive sleep mode control. Although the highest threshold values yield the maximum power saving, aggressive values could potentially raise QoE risk in highly dynamic environments. The fact that an optimal range for the threshold exists also supports the findings in recent studies on green networking, where the adaptation of the threshold is identified as an important aspect in designing such networks [14, 15].

Fig. 5 illustrates the relationship between T_{thr} and energy savings. As T_{thr} increases from 50 to 200 Mbps, energy savings increase from 43.9% to 50.4%, as the BS is permitted to enter sleep mode for longer durations. Zero QoE violations are maintained across all configurations, with savings ranging from 43.9% to 50.4%. Critically, QoE violation counts remain at zero for all threshold configurations, demonstrating the robustness of LSTM-based prediction.

Borderline cases ($T(t)$ within $[-5\%, +5\%]$ of T_{thr}) represent 2.3% of all test samples at $T_{thr} = 100$ Mbps. In all such cases, the LSTM prediction remained below the threshold, confirming the conservatism of the model at this operating point. The optimal operating range is $T_{thr} \in [100, 150]$ Mbps, where substantial energy savings are achieved with comfortable safety margin from the QoE violation boundary.




Fig. 5. Threshold sensitivity analysis: Energy savings (%) and QoE violation count as a function of traffic threshold T_{thr} (50–200 Mbps).

D. Computational Complexity and Latency Analysis

To evaluate execution feasibility, we benchmarked the core computational loop on an edge computing environment (Intel Core i7-11800H CPU @ 2.30 GHz, 16 GB RAM) where the system will run in a single thread to mimic typical 5G RAN edge processing. Total inference latency (τ) is the combination of the LSTM forecasting step ($t_{predict}$) and the

boundary threshold control step ($t_{control}$). A total of 10,000 independent evaluation cycles (mean total latency) is given by 0.235 ms (± 0.019 ms; standard deviation) the 99th percentile (P_{99}) is 0.277 ms. This is an extremely tight distribution of the execution latency and is well below the sub-millisecond level needed for optimal radio resource management for 5G next-generation NodeBs (gNodeBs). The computational time (τ) for the proposed computation is:

$$\tau = t_{predict} + t_{control} \quad (12)$$

The study was conducted on over 10,000 prediction control cycles using Milan and Trentino traffic profiles. The execution time distribution is summarized in Table 3.

Table 3. Quantitative computational complexity and latency profile

Metric Component	Mean Latency	Standard Deviation	99th Percentile (P_{99})
LSTM Forward Pass ($t_{predict}$)	0.224 ms	0.018 ms	0.262 ms
Threshold Control Loop ($t_{control}$)	0.011 ms	0.002 ms	0.015 ms
Total Inference Latency (τ)	0.235 ms	0.019 ms	0.277 ms

The experimental results show that the execution time is always around 0.24 ms and that the worst-case 99th percentile latency is 0.28 ms. This performance profile easily meets the time frame of 5G Next Generation NodeBs (gNodeBs), which are used for real-time radio resource management.

V. COMPARATIVE DISCUSSION

A. Energy Efficiency Trajectories and the Paradigm of Anticipatory Sleep Control

The network's behavior under different control strategies shows that it is a paradigm shift when switching from reactive to predictive. As shown by the temporal state plots, the always-on policy is still at an initial cost of 30.50 kWh, while the reactive policy is even cheaper (20.73 kWh), but its deterministic, post-facto trigger logic is a huge bottleneck in performance. Since a reactive system needs to wait for traffic to drop below a threshold to start a shutdown, it cannot wait for the late hours of low-demand windows and therefore needs to draw more power in the transitional phase.

In contrast, our predictive policy uses an LSTM network to model non-linear time-series dependencies and moves the base station to a low-power state ahead of expected traffic drops. By mapping these anticipatory transitions, the framework yields a net efficiency gain ranging from 43.9% at

$T_{thr} = 200$ Mbps to 50.4% at $T_{thr} = 50$ Mbps, with the primary operational profile establishing a 45.8% energy reduction at $T_{thr} = 100$ Mbps. This indicates that predictive scheduling bridges the gap between active network readiness and the need for sleep.

Fig. 6 exhibits the flow between Active, Transition, and Sleep states of a 5G base station at different traffic load scenarios. This visualization compares three operational policies against traffic demand (top panel), based on a 100 Mbps sleep threshold, respectively. The Predictive policy attain more efficient sleep scheduling, yielding 45.8% energy savings vs. 32.0% for the Reactive policy.

Here it is shown that the predictive policy provides more optimized (timely) schedules than the reactive policy and thereby can greatly benefit 5G “greenness” without lowering the service level.




Fig. 6. Temporal visualization of base station state (Active/Sleep) under always-on, reactive, and predictive (ML) policies across a 24-hour period.

B. Algorithmic Comparative Benchmark and Computational Complexity Trade-offs

To evaluate the practical viability of our framework, we compare our approach in terms of three key operational parameters: computational budget, training data overhead, and risk profiling with the leading state-of-the-art architecture summarized in Table 4.

Table 4. Comparative evaluation of the proposed framework against state-of-the-art 5G base station sleep mode methodologies

Study	Methodology	Dataset Context & Training Requirements	Energy Savings	QoE Metric/Impact Profile	Computational Complexity & Budget	Deployment Feasibility & Constraints
[10]	Deep Reinforcement Learning (DRL) + Spatiotemporal Forecast	Telecom Italia Milan (Real); High offline training data overhead.	~45.0%	Latency-aware; prone to transient drops during convergence spikes.	Very High (Requires continuous reward shaping and online loop updates).	Low (High risk of exploratory downtime; requires heavy edge GPU infrastructure).
[13]	Multi-Agent Reinforcement Learning (MARL)	Real-world Swedish operator dataset; multi-cell simulation models (MIMO, user association).	~50.0%	Negligible degradation of Quality of Service (QoS); collaborative user offloading mitigates interference.	Very High (Joint control of multi-level sleep, antenna switching, and user scheduling).	Low (State-space scaling bottlenecks; heavy decentralized model update loops at edge).

[16]	LSTM + Lyapunov Optimization	Synthetic traffic traces; Minimal initialization data required.	~15.0%	Queue stability and bounded delay enforcement.	High (Requires solving optimization matrices at every single decision step).	Low (Mathematical formulation becomes intractable for dense urban deployments).
[17]	LSTM + Base Station Sleeping	Spatio-temporal dataset; Moderate training time needed.	Moderate	Implicit; lacks explicit multi-region threshold safety boundaries.	Moderate (Standard deep network forecasting execution).	Moderate (Lacks real-time safety mechanisms for edge cases).
[19]	Deep Learning (DL) + Reconfigurable Intelligent Surfaces (RIS)	Simulated multi-cell environment; Demands continuous hardware state inputs.	~49.6%	Delay-energy tradeoff optimized dynamically.	High (Simultaneous prediction and meta-surface element phase control).	Low (Dependent on expensive, specialized hardware deployment).
[20]	Federated LSTM + DRL	Real 5G operator dataset; High distributed communication overhead.	~40.0%	Moderate; bounded by federated aggregation steps.	High (Frequent decentralized parameter aggregations).	Low (Requires strict security protocols and high network bandwidth).
Proposed Framework	LSTM Neural Network + Threshold Control	Milan & Trentino (Real); Lean historical traffic training requirements.	43.9%–50.4%	Zero (0) violations; structured threshold safety buffers.	Low (Deterministic, single-model forward pass; no training loop).	High (Immediately compatible with legacy O-RAN edge processors).

1) Computational budget vs. runtime latency

At runtime, Multi-Agent Reinforcement Learning (MARL) frameworks such as Ref. [13] adopt a Multi-Agent Deep Deterministic Policy Gradient (MADDPG) model under Centralized Training with Decentralized Execution (CTDE). Within this setup, the system must continuously balance decision accuracy against online inference latency. The base station is an agent that is autonomous in translating local traffic demand, power state, and neighbor interference to discrete sleep-mode actions. By separating training from execution, online inference can only perform a localized evaluation of the actors with $O(1)$ time complexity, and the decision latency can be kept within sub-millisecond bounds to meet strict CPU budget constraints.

The Federated DRL models [20] provide substantial peak energy savings—yielding approximately 50% and 40% reductions, respectively—but also introduce significant state-space explosions, protracted training convergence loops, and the requirement to continuously iterate on decentralized parameters. Such overheads lead to processing times that are not suited for 5G NG-RAN, which is only on the order of 5G Next-Generation Radio Access Network (NG-RAN). In contrast to these, our framework only uses the forward pass of a lightweight LSTM. Its deterministic inference step takes less than a millisecond, which is the most realistic for hardware profiles with low processing power.

2) Training data requirements

Many state-of-the-art DRL designs (DeepBSC [10] or RIS-assisted designs [19]) require extensive environmental simulations and complex reward shaping to prevent exploratory divergence. Our architecture avoids this challenge by only looking at traffic volume statistics and not complex multi-cell signaling loops or active environment tracking.

3) Operational risk and quality of experience (QoE)

To better understand how to minimize the QoE violation with a global MAE of 95.36 Mbps, we first need to see how the forecasting error distribution and the absolute traffic load are related. The ability to maintain service quality in the presence of large error values depends on two key characteristics of the network environment: The cellular connection between Milan and Trentino is very regular and cyclical in the diurnal cycle typical of dense cities. There is

no chaotic, high-amplitude traffic and no sudden capacity deficit in the sleep hours. In addition, our locally averaged error is dominated by the large predicted error in the daytime (~100 Mbps) and evening time. Thus, the control loop keeps the base station active and is not prone to incorrect shutdowns during peak periods. However, on the low-demand nights (01:00–05:00) when sleep is being done in the network (in which the sleep cycle is actively operated), the absolute traffic is well below 50 Mbps, and in these low-capacity regimes, the localized forecasting error is less than 5 Mbps. In these low-capacity regions, the model is very accurate and works well when the network remains close to the activation threshold and the system is able to control the small capacity drop and there are no QoE violations in our experiment.

Telecom operators value reliability for services over energy savings in real environments. Reinforcement learning agents are often plagued by operational risk due to their trial-and-error exploration phase that can cause unpredictable base station failures. In combination with a mathematically bounded threshold control law (100–150 Mbps) for LSTM predictions, our framework creates a safety buffer that guarantees zero QoE violations during the evaluation period.

On the other hand, our framework relies only on a lightweight LSTM configuration. As shown by our hardware profiling data in Section IV-D, our total inference loop achieves a deterministic mean execution latency of 0.235 ms ($P_{99} = 0.277$ ms). As a concrete proof-of-concept, our framework is able to overcome the processing bottlenecks of Lyapunov matrix drift tracking and DRL exploratory convergence loops in order to satisfy live 5G timing requirements.

To make sure that our proposed framework is scientifically sound compared with the state-of-the-art baselines in Table 4, two limitations should be considered. First, it is impossible to run the same DRL [10], MARL [13], and meta-surface RIS [19] at the same time on a computer and to make sure that they are all the same. Such complex reinforcement learning models are based on their own training environments, multi-cell simulation algorithms, Radio Resource Management (RRM) abstraction layers, and hyperparameter reward functions that cannot be replicated publicly in the literature. Therefore, the result obtained by comparing the results with data samples such as the Telecom Italia macro-urban dataset is the standard for the comparison

of the two. Second, our framework is completely different from the state-of-the-art frameworks in the sense that it addresses the simplicity-stability model in network automation. State-of-the-art learning models in this case are based on maximizing the energy savings by running exploration loops without being constrained by a constraint. We have chosen not to track the online optimization matrix at high complexity (e.g., Lyapunov drift calculation [16]) and instead adopted a mathematically regularized predictive model with a threshold boundary law (T_{thr}) to achieve similar energy savings (45.8%–50.4%) but with a deterministic latency (0.235 ms) and no exploration downtime. The scientific contribution is that we can achieve Pareto-optimal energy-delay efficiency in a safety-conscious way.

C. Limitations and Future Framework Extensions

The evaluation is subject to the following limitations. First, it is performed in a single cell (offline simulation) setting and not considering intercell interference and sleeping schedule between neighboring base stations. Second, the traffic threshold T_{thr} is a fixed value throughout the evaluation and adaptive threshold optimization to adjust for long-term traffic regime change would be beneficial. Third, the energy model utilized is based on 3GPP macro cell parameters and likely cannot be generalized to small cell or Open Radio Access Network (O-RAN). Fourth, the QoE is only evaluated for violation count and estimated latency overhead. The training data from the past might not be indicative of a non-stationary or anomalous traffic scenario. Future work will focus on multi-cell cooperative sleep scheduling with load balancing between cells and on adaptive threshold optimization online using lightweight reinforcement learning; hardware-in-the-loop validation on an O-RAN testbed; and use of recent graph-based spatiotemporal prediction models [22, 23] as prediction front-end for the threshold controller.

Although the proposed LSTM-based threshold control framework illustrates strong energy efficiency and service quality preservation, there are some limitations:

- Cross-paper evaluation context: As mentioned in the comparison, an explicit comparison with state-of-the-art reinforcement learning methods has its limitations. These external models have been evaluated under different experimental conditions, network topologies, and traffic profiles, so this comparison is of a contextual benchmark rather than a definitive measure of the superior model.
- Single cell focus: We are mainly considering single-cell base station operations in this study, focusing on local traffic traces from the Milan and Trentino dataset. It does not model inter-cell interference, macro-micro coordination, or cooperative sleep scheduling in a dense multi-cell cluster dynamically.
- QoE metric depth: The Quality of Experience (QoE) validation relies on a threshold-crossing event model. It is effective for tracking service drops, but it does not capture even the most complex network dynamics specifically packet loss in real time, end-to-end user latency fluctuations, or sudden throughput fluctuations during state transitions.
- Lack of physical hardware validation: Our evaluation is based on offline algorithmic simulations driven by real-world historical data traces. While our runtime analysis shows that our execution timing is in line with 5G slot

constraints, we have not yet tested our framework on a physical hardware testbed (for instance, software-defined radio networks or Open Radio Access Network [O-RAN] architectures). Future work will be concerned with validating real-time synchronization overheads and hardware-induced latencies. We will also change our 3GPP macro-cell energy model to account for O-RAN functional splits.

VI. CONCLUSION

This research demonstrates that LSTM neural networks for traffic forecasting and deterministic threshold-based control law perform a very effective and easily implemented network to optimize 5G base station energy efficiency. Compared to the real-world Telecom Italia Milan and Trentino data sets, the proposed framework consistently realizes a net energy savings of 43.9%–50.4% compared with the default reactive level (32.0% in the real world) by a margin of more than 43%. Moreover, by selecting an optimal performance of 100 Mbps to 150 Mbps, we achieve these power savings while maintaining a service level of QoE neutrality and without a violation. The computational cost of this framework is still low, and the substantial processing time, convergence time, and state-space explosion in multi-agent and deep reinforcement learning models are not a major problem that we need to overcome. By breaking the initial model overfitting through a well-designed combination of early stopping and L2 regularization, our work provides a general time-series forecasting engine that can run reliably in real-world RAN infrastructure. To build off of these initial insights, the future research efforts will focus on scaling this architecture to multi-cell cooperative deployment scenarios, building online adaptive threshold optimization strategies through real-time feedback loops, and validating the whole framework on an O-RAN hardware testbed to enable scale commercial orchestration.

CONFLICT OF INTEREST

The author declares no conflict of interest.

REFERENCES

- [1] G. Barlacchi *et al.*, “A multi-source dataset of urban life in the city of Milan and the Province of Trentino,” *Scientific Data*, vol. 2, no. 1, 150055, 2015.
- [2] Telecom Italia. (2014). Telecom Italia big data challenge. [Online]. Available: <http://theodi.fbk.eu/openbigdata>
- [3] OpenCellID Community. (2023). OpenCellID: The world’s largest open database of cell towers. [Online]. Available: <https://www.opencellid.org>
- [4] D. Kotz and T. Henderson, “Crawdad: A community resource for archiving wireless data at Dartmouth,” *IEEE Pervasive Computing*, vol. 4, no. 4, pp. 12–14, 2005.
- [5] F. Xu, Y. Li, H. Wang *et al.*, “Understanding mobile traffic patterns of large scale cellular towers in urban environment,” *IEEE/ACM Transactions on Networking*, vol. 25, no. 2, pp. 1147–1161, 2016.
- [6] OpenAirInterface Software Alliance. (2026). 5G open RAN traffic traces and network measurements. [Online]. Available: <https://www.openairinterface.org>
- [7] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [8] M. T. Isyapar, S. Yildirim, and U. Uyan, “Machine learning based mobile traffic forecasting for large scale 5G networks,” in *Proc. 2023 14th International Conference on Electrical and Electronics Engineering (ELECO)*, 2023, pp. 1–5.
- [9] S. Tian, X. Zhang, Y. Zhang *et al.*, “Spatio-temporal position prediction model for mobile users based on LSTM,” in *Proc. 2019*

- IEEE 25th International Conference on Parallel and Distributed Systems (ICPADS), 2019, pp. 967–970.
- [10] Q. Wu, X. Chen, Z. Zhou *et al.*, “Deep reinforcement learning with spatio-temporal traffic forecasting for data-driven base station sleep control,” *IEEE/ACM Transactions on Networking*, vol. 29, no. 2, pp. 935–948, 2021.
 - [11] S. Maghsudi and E. Hossain, “Multi-armed bandits with application to 5G small cells,” *IEEE Wireless Communications*, vol. 23, no. 3, pp. 64–73, 2016.
 - [12] A. Mughees, M. Tahir, M. A. Sheikh, and A. Ahad, “Towards energy efficient 5G networks using machine learning: Taxonomy, research challenges, and future research directions,” *IEEE Access*, vol. 8, pp. 187498–187522, 2020.
 - [13] T. Cai, “Mobile traffic classification and multi-cell base station control for energy-efficient 5G networks,” M.S. thesis, School of Electrical Engineering and Computer Science, KTH, Stockholm, Sweden, 2023.
 - [14] M. Na, S. Cho, F. Solat *et al.*, “Energy-efficient hybrid federated and centralized learning for edge-based wireless traffic prediction in aerial networks,” *IEEE Access*, vol. 12, pp. 130983–130994, 2024.
 - [15] D. Kulkarni, M. Venkatesan, and A. V. Kulkarni, “Deep learning traffic prediction and resource management for 5G RAN slicing,” *Journal of the Institution Engineers (India): Series B*, vol. 106, no. 2, pp. 593–606, 2025.
 - [16] G. Jang, N. Kim, T. Ha *et al.*, “Base station switching and sleep mode optimization with LSTM-based user prediction,” *IEEE Access*, vol. 8, pp. 222711–222723, 2020. doi: 10.1109/ACCESS.2020.3044242
 - [17] Y. Zhu and S. Wang, “Joint traffic prediction and base station sleeping for energy saving in cellular networks,” in *Proc. ICC 2021-IEEE International Conference on Communications*, 2021, pp. 1–6.
 - [18] 3GPP TS 38.300. (2020). 5G; NR; overall description; stage-2, release 16. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/138300_138399/138300/16.02.00_60/ts_138300v160200p.pdf
 - [19] S. Sun, C. Huang, G. Chen *et al.*, “Deep learning-based traffic-aware base station sleep mode and cell zooming strategy in RIS-aided multi-cell networks,” *IEEE Transactions on Cognitive Communications and Networking*, vol. 11, no. 4, pp. 2171–2184, 2025. doi: 10.1109/TCCN.2025.3529983
 - [20] H. Park and S. H. Yoon, “Deep reinforcement learning for base station switching scheme with federated LSTM-based traffic predictions,” *ETRI Journal*, vol. 46, no. 3, pp. 379–391, 2024. doi: 10.4218/etrij.2023-0065
 - [21] A. K. Jha *et al.*, “Optimizing 5G sustainable networks performance in urban areas using machine learning techniques,” *International Journal of Communication Systems*, vol. 38, no. 14, e70226, 2025. doi: 10.1002/dac.70226
 - [22] D. Vengaimarban and D. Rajiniginath, “Base station sleeping strategy in heterogeneous cellular networks using transformer swarm evolutionary adaptive memory gate convolutional LSTM model for cellular traffic prediction,” *Expert Systems with Applications*, 132037, 2026. doi: 10.1016/j.eswa.2026.132037
 - [23] N. Rajule, M. Venkatesan, R. Menon *et al.*, “Energy-efficient adaptive sleep mode control in cellular networks using grasshopper optimization and traffic-aware QoS management,” *Cluster Computing*, vol. 29, no. 3, 2026. doi: 10.1007/s10586-026-06003-1

Copyright © 2026 by the authors. This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited ([CC BY 4.0](https://creativecommons.org/licenses/by/4.0/)).