



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/242057/>

Version: Accepted Version

Proceedings Paper:

Fan, Y., Salama, A., Qazzaz, M.M.H. et al. (Accepted: 2026) Intent-Driven Agentic AI for Safe Energy-Efficient Open RAN Optimization via a 5G Network-in-a-Box Digital Twin. In: 2026 IEEE International Conference on Communications Workshops (ICC Workshops). 2026 IEEE International Conference on Communications (ICC), 24 May - 28 Jun 2026, Glasgow, Scotland, UK. IEEE. (In Press)

This is an author produced version of an conference paper accepted for publication in 2026 IEEE International Conference on Communications Workshops (ICC Workshops), made available via the University of Leeds Research Outputs Policy under the terms of the Creative Commons Attribution License (CC-BY), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Intent-Driven Agentic AI for Safe Energy-Efficient Open RAN Optimization via a 5G Network-in-a-Box Digital Twin

Yejing Fan¹, Abdelaziz Salama¹, Mohammed M. Hameed Qazzaz¹, Maryam Hafeez¹,
Syed Ali Raza Zaidi¹, Ahmed Murad², Mahdi Sharara², and Vara Prasad Talari³

Abstract—As networks evolve with virtualization and disaggregation, operators face increasing complexity, rising energy costs, heterogeneous devices, and highly dynamic demand. This work responds to the challenge of enabling networks operating with high levels of autonomy, understanding the directives of the operator in natural language, reasoning over scenarios, and acting intelligently to optimize themselves using Agentic AI. This paper presents an end-to-end demonstration architecture that couples a portable 5G network-in-a-box with an AI-enabled Open RAN (O-RAN) digital twin for intent-driven energy-efficient network optimization. The framework is deployed in a cloud-native configuration aligned with O-RAN principles, exploiting Amazon Web Services (AWS) scalable rApp hosting, elastic compute, and MLOps support, while the VIAVI AI Radio System Generator (AI-RSG) which provides a realistic digital twin for repeatable benchmarking and safe policy validation. The framework contains two tightly coupled agentic subsystems: (i) an intent and scenario generation agent that translates natural-language operator inputs into structured simulation schemas; and (ii) a traffic forecasting and energy optimization agent based on a Long Short-Term Memory (LSTM) model, which provides short-horizon per-RU load prediction and confidence-aware three-class classification to guide base station sleep decisions. A purpose-built energy-efficiency optimizer executes these control actions on a real OpenAirInterface gNB testbed with Universal Software Radio Peripheral (USRP) hardware and commercial user equipment (UE), where controlled channel impairments provide labeled ground-truth events for supervised learning. Experimental results demonstrate a 45.6% reduction in RAN power consumption while maintaining full UE connectivity, corresponding to 6.03 kWh/day energy savings (513 kg CO₂/year). The LSTM forecasting model achieves 96.3% training accuracy and 95.7% validation accuracy, with stable validation loss, indicating strong generalization across traffic regimes. These results demonstrate a practical pathway toward safe and autonomous energy management in next-generation 5G/6G RAN systems.

Index Terms—O-RAN, Agentic AI, Energy Efficiency, Digital Twin, Network-in-a-Box, OpenAirInterface

I. INTRODUCTION

The evolution of mobile networks toward virtualized, disaggregated, and programmable architectures increases operational complexity, energy cost, device heterogeneity, and

traffic variability. In this setting, there is growing interest in autonomous control frameworks that can interpret operator intent, reason over deployment scenarios, and optimize network behavior. Open Radio Access Network (O-RAN), through its disaggregated architecture and RAN Intelligent Controller (RIC) framework, supports both near-real-time and non-real-time optimization via xApps and rApps, respectively [1]. At the same time, the radio access network (RAN) remains the dominant contributor to end-to-end network energy consumption in dense deployments, making energy-aware RAN control a key target for intelligent automation [2].

Among available mechanisms, switching Radio Units (RUs) between active and sleep states during low-demand periods offers substantial energy-saving potential [3]. However, safe RU sleep control must account for not only current load, but also predicted demand, propagation conditions, and contextual factors such as weather and local events. Otherwise, energy savings may come at the cost of degraded Quality of Service (QoS). This motivates a closed-loop framework that combines forecasting, contextual reasoning, and rapid mitigation within the O-RAN control hierarchy.

Recent work has addressed parts of this problem, including mobility-aware power optimization in O-RAN [3], LSTM-based traffic prediction and agentic control [4], O-RAN resource allocation for federated learning [5], [6], adaptive model selection for efficient forecasting [7], and safety issues in multi-agent telecom systems [8]. However, these studies largely treat forecasting, optimization, and safety separately. In contrast, this work integrates intent translation, forecasting, contextual reasoning, digital twin validation, and real hardware execution into a single deployable framework.

Despite this progress, forecasting, optimization, and safety have largely been treated separately. Existing studies do not close the loop from human-readable operator intent, through a cloud-native O-RAN optimization pipeline, to real hardware validation with measurable energy savings. The main contribution of this work is to integrate intent translation, agentic reasoning, topology-aware optimization, digital twin validation, and real hardware execution within a single deployable framework.

To address this gap, we present an end-to-end demonstration architecture that combines a portable 5G network-in-a-box (NIB) with an AI-enabled O-RAN digital twin for intent-

¹School of Electronic and Electrical Engineering, University of Leeds, Leeds, UK. (e-mail: {Y.Fan1, A.M.Salama, ml14mmh, M.Hafeez, S.A.Zaidi}@leeds.ac.uk)

²VIAVI Solutions. (e-mail: Ahmed.Murad@viavisolutions.com, Mahdi.Sharara@viavisolutions.com)

³Amazon Web Services. (e-mail: vatalari@amazon.com)

driven energy-efficient network optimization. The framework follows cloud-native O-RAN principles: Amazon Web Services (AWS) provides scalable rApp hosting, elastic compute, and MLOps support, while the VIAVI AI Radio System Generator (AI-RSG) provides a realistic digital twin for repeatable benchmarking and safe policy validation. Within the proposed AgentRAN framework, intent processing, forecasting, and energy optimization are implemented as Non-RT RIC rApps, while anomaly detection and rapid mitigation are handled by a Near-RT RIC xApp, enabling autonomous yet safe RU ON/OFF control. This paper makes four main contributions:

- **End-to-end agentic architecture:** An intent and scenario agent compiles natural-language intent into a reproducible O-RAN scenario schema, while agentic control functions coordinate forecasting, anomaly detection, and energy optimization across RIC layers.
- **Decision-oriented forecasting and control:** A confidence-aware LSTM-based forecasting agent provides 10-minute-ahead per-RU load predictions and three-class traffic regime classification (Low, Moderate, High), supporting safe base station sleep and wake decisions.
- **Cloud-native O-RAN deployment:** The framework is deployed on AWS, providing scalable rApp hosting, elastic GPU/CPU resources for ML training and inference, MLOps orchestration, and secure API integration with the Non-RT RIC.
- **Dual validation in DT and hardware:** The complete stack is validated using the VIAVI AI-RSG digital twin for repeatable benchmarking and a portable NVIDIA DGX Spark with USRP B210 NIB platform for real UE and controlled-impairment validation.

II. SYSTEM ARCHITECTURE

Fig.1 presents a multi-agent-based framework that enables energy-efficient optimization in O-RAN by coordinating intelligent agents across the Near-RT RIC, Non-RT RIC, and external digital twin environments. Agentic control loops as explained below are built using AWS Strands Agents. These agents interact with the Non-RT RIC via the A1 interface for policy guidance and model updates, while control actions are enforced in the Near-RT RIC through the E2 interface. The Viavi digital twin platform (Teravm AI RSG) is leveraged to simulate RAN scenarios and validate optimization strategies before deployment, providing closed-loop feedback into the agent system. At the core of this orchestration lies the Model Context Protocol (MCP) server, which acts as a shared reasoning and coordination layer, exposing capabilities (rApps/xApps as “skills”), maintaining context across agents, and enabling structured interaction between LLM-driven agents and network functions.

Fig. 2 illustrates the end-to-end operational architecture, organized into two logical planes: a simulation and scenario plane and a live RAN control plane, connected by a context-flag channel.




Fig. 1. AgentRAN multi-agent architecture across Non-RT and Near-RT RIC layers

Operationally, the framework can be viewed as an instance of AgentRAN, a cloud-native agentic AI architecture for autonomous energy optimization in O-RAN. The intent and forecasting functions reside in the Non-RT control domain as decision-oriented rApps, while a lightweight anomaly-detection xApp operates in the Near-RT RIC to monitor streaming key performance measurement (KPM) and QoS indicators, including PRB utilization, throughput degradation, UE drops, and abrupt load spikes. When degradation is detected, the xApp triggers re-optimization and, if necessary, reactivation of previously sleeping RUs. RU switching actions are then executed through O-RAN control interfaces toward the simulator and live test environment.

In the simulation and scenario plane (top, Fig. 2), the intent agent accepts a free-form natural-language description from the network operator (e.g., “Elland Road Stadium, Leeds. After the match, only 20 people showed up. There are four base stations (BS) serving these people.”) and executes a four-stage tool pipeline: a Nominatim geocoder resolves venue coordinates, an OpenWeatherMap API call retrieves current weather, a Tavily web search corroborates geographic context, and a retrieval-augmented generation (RAG) knowledge base supplies 3GPP deployment knowledge. The pipeline output populates a structured JSON schema, with fields annotated with LLM-generated reasoning traces. The schema seeds the digital twin, which renders a geographically accurate 3-D Mapbox visualization overlaid with real-time path-loss heatmaps, BS-UE link animations, and receive-power contours. In the present implementation, this role is provided by the VIAVI AI-RSG.

In the live RAN control plane (bottom, Fig. 2), the NIB hardware block comprises the USRP B210 radio, the DGX Spark compute node, over-the-air (OTA)-connected UEs, and the controlled channel-impairment interface (channelmod). The OAI gNB exports a live KPM stream to the Forecasting Agent, which applies adaptive LSTM model selection: the lightweight LW-LSTM handles stable Low and Moderate regimes, while the full FP-LSTM is activated in response to anomaly or event triggers. The agent produces a three-class regime label, along with per-class confidence scores, which are fed to the energy efficiency (EE) Optimizer xApp. The EE Optimizer issues BS sleep and wake commands and exposes a single-click operator WAKE ALL override for human-in-the-loop safety. A dashed context-flags arrow (Fig. 2, center)




Fig. 2. End-to-end workflow of the proposed AgentRAN framework

carries weather caution and event proximity flags from the scenario plane directly into the EE Optimizer’s sleep gating function, blocking unsafe sleep decisions whenever contextual risk is elevated.

III. INTENT AND SCENARIO AGENT

The intent and scenario agent translates natural-language operator requests into reproducible O-RAN scenarios for the digital twin. A large language model (LLM) reasoning module coordinates external tools for geocoding, weather data retrieval, web-based geographic context lookup, and RAG-based telecom knowledge extraction, subsequently compiling the results into a structured JSON schema.

The resulting schema defines parameters such as location, activity context, anticipated user density, deployment constraints, RU placement, mobility patterns, and temporal traffic behavior. This structure enables controlled and repeatable digital twin configurations. For operator oversight, each field is linked to a concise reasoning note and can be manually overridden if necessary.

The agent further derives context-sensitive control flags from the schema. For example, large nearby events may inhibit sleep decisions for affected base stations, while adverse weather conditions trigger caution flags throughout the scenario. These control flags, along with historical KPM data and traffic forecasts, are provided to the EE optimizer to prevent unsafe sleep actions. The flags are also displayed in the operator dashboard as part of the confidence assessment.

IV. TRAFFIC FORECASTING AND ENERGY SAVING AGENT

The Traffic Forecasting and Energy Saving Agent provides short-horizon, 10-minute-ahead per-RU load predictions to support proactive, safe EE optimization. The agent collects real-time and historical KPMs, applies an adaptive model selection strategy to choose between a lightweight LSTM and

a larger pre-trained LSTM depending on the detected traffic regime, and outputs a three-class regime label (Low / Moderate / High) with an associated confidence score for every active RU. This confidence-aware output is the primary input to the sleep gating logic in the EE Optimiser xApp.

A. KPM Feature Space and Prediction Horizon

At each reporting interval $\Delta t = 10s$, the agent collects a feature vector $x_t \in R^d$ per RU, comprising dimensions $d = 8$ KPM: downlink (DL) Physical Resource Block (PRB) use, uplink (UL) PRB use, cell DL throughput (Mbps), cell UL throughput (Mbps), mean SINR UE (dB), mean Received Power from the Reference Signal (RSRP, dBm), number of active UEs, and time-of-day encoding (normalized hour). A sliding window of $T = 60$ consecutive steps (10 minutes of history) is assembled into the input tensor $\mathbf{X} \in R^{T \times d}$, providing sufficient temporal context for both daytime traffic patterns and event-driven load spikes, consistent with the horizon established in our prior LSTM deployment [4]. The agent produces a prediction of the load state at the horizon of $\tau = 60$ steps (10 minutes) ahead.

B. LSTM Encoder Architecture

Both model variants share the same LSTM encoder skeleton. A two-layer LSTM with hidden dimensions $h_1 = 128$ and $h_2 = 64$ processes the input sequence $\mathbf{X}$ and emits a sequence of hidden states $\mathbf{H} = \{h_1, \dots, h_T\}$. The final hidden representation h_T is fed into a two-layer fully connected classifier with ReLU activations and dropout ($p = 0.2$), producing a three-class probability distribution via a softmax function:

$$\hat{\mathbf{y}}(t+\tau) = \text{softmax}(\mathbf{W}_2 \cdot \text{ReLU}(\mathbf{W}_1 \cdot h_T + \mathbf{b}_1) + \mathbf{b}_2) \quad (1)$$

where the three output classes correspond to the low (PRB utilization $\leq 30\%$), Moderate (30–70%), and High ($> 70\%$) regimes. The class boundaries are derived empirically from the NIB KPM distribution and are aligned with the PRB thresholds used by the EE Optimizer sleep-gating function (Section V). The model is trained with the categorical cross-entropy loss using the Adam optimizer (learning rate = 10^{-3} , batch size = 32, maximum 50 epochs, early stopping patience = 10).

C. Adaptive Model Selection

A central feature of the forecasting agent is its ability to dynamically select between two pre-trained LSTM variants based on the currently detected traffic regime, following the model-complexity philosophy of the NAS-based framework in [7]:

- **Lightweight LSTM (LW-LSTM):** Used when the traffic regime has been classified as Low or Moderate for at least three consecutive reporting windows, indicating a stable operating condition. This model has a reduced encoder ($h_1 = 64, h_2 = 32$, 38K parameters) and achieves $R \approx 0.93$ on stable-regime held-out data with 70–75% lower

computational cost than the full model, enabling energy-efficient inference on the DGX Spark edge compute node.

- Automatically activated whenever any of the following anomalies triggers fires: (i) the current regime prediction transitions to High with *confidence* > 0.80 ; (ii) the scenario agent raises a weather caution flag $F_{weather}(t) = 1$; (iii) the event agent raises an event proximity flag $F_{event}(b, t) = 1$ for any active BS; or (iv) the PRB utilization gradient $\Delta p / \Delta t$ exceeds a threshold of 2%/s over the previous five steps. The FP-LSTM uses the full $h_1 = 128, h_2 = 64$ encoder and was pre-trained on the dataset, simulated urban mobility traces from [3], and event-day traffic from [4], achieving 95.7% validation accuracy on the test set.

The model selector $S(t)$ is implemented as a stateful finite automaton with hysteresis: once the FP-LSTM is activated, it remains active for a minimum hold time of $T_{hold} = 6$ reporting windows (1 minute) to prevent rapid oscillation at regime boundaries. This design is consistent with the Conservative/Moderate/Aggressive threshold framework of [9], with the hold time acting analogously to the Conservative stability margin.

D. Confidence-Aware Output and Sleep Gating

The softmax output $\hat{y}(t + \tau)$ provides a calibrated confidence signal for regime prediction. Based on this output, the EE Optimiser defines a sleep-gating function $G(b, t)$, which permits a sleep action for base station b only when the predicted confidence of the *Low* regime exceeds a predefined threshold $\theta = 0.75$, and no adverse contextual flags are present:

$$G(b, t) = 1 \quad (2)$$

$$\iff \hat{y}_{Low}(t + \tau) \geq \theta \wedge F_{weather}(t) = 0 \wedge F_{event}(b, t) = 0 \quad (3)$$

Otherwise, if $G(b, t) = 0$, the base station remains active and the associated decision rationale is recorded for operator inspection. This conservative gating mechanism reduces the risk of unsafe sleep transitions under uncertain operating conditions and prevents ambiguous regime predictions from degrading service continuity. For example, a prediction for which $\hat{y}_{Low}(t + \tau) = 0.72$, which remains below the sleep threshold.

V. ENERGY-EFFICIENT BASE STATION SLEEP OPTIMIZATION

Within AgentRAN, the EE optimizer acts as the decision-fusion layer, combining predicted traffic regime, contextual flags, and anomaly feedback before issuing RU sleep/wake actions. The EE optimizer xApp operates within the Near-RT RIC at a control granularity of 100 ms. It maintains a real-time representation of the network state, including the set of base stations $\mathcal{B}$, the set of UEs $\mathcal{U}$, and the time-varying radio and load conditions. At each decision epoch t , the optimizer determines a binary sleep-control policy for base stations. Let

$$s_m(t) \in \{0, 1\}, \quad \forall m \in \mathcal{B} \quad (4)$$

denote the decision variable, where $s_m(t) = 1$ indicates that base station m is transitioned to sleep mode, and $s_m(t) = 0$ indicates that it remains active [9].

Objective Function: The goal is to maximise the total energy saving achieved by switching selected base stations into sleep mode:

$$\max_{\{s_m(t)\}} \sum_{m \in \mathcal{B}} s_m(t) P_{sleep}(m) \quad (5)$$

where $P_{sleep}(m)$ denotes the power reduction obtained when base station m is switched from active to sleep mode [4].

QoS Constraint: To ensure service continuity, each UE must maintain a minimum throughput requirement:

$$\sum_{m \in \mathcal{B}} \alpha_{k,m}(t) R_{k,m}(t) \geq R_{min}, \quad \forall k \in \mathcal{U} \quad (6)$$

where $\alpha_{k,m}(t) \in \{0, 1\}$ denotes the UE-BS association indicator, and $R_{k,m}(t)$ is the achievable data rate between UE k and BS m .

Feasibility Condition: A candidate sleep decision is considered feasible only if each UE originally served by a sleeping BS can be reassigned to at least one neighbouring active BS with sufficient spare PRB capacity and a reference signal received power (RSRP) above -90 dBm.

Solution Approach: The optimization is solved using a greedy search over candidate subsets of base stations, evaluated in descending order of potential energy saving. At each step, feasibility is verified under the QoS constraint.

Power Consumption Model: The instantaneous power consumption of base station m is modeled as:

$$P_m(t) = P_{tx,m}(t) + P_{idle,m} \quad (7)$$

where $P_{tx,m}(t)$ represents load-dependent transmission power and $P_{idle,m}$ is the static idle power [5].

The transmission power is further expressed as:

$$P_{tx,m}(t) = P_{mode,m} \cdot \alpha_m(t) \quad (8)$$

where $\alpha_m(t) \in [0, 1]$ denotes the transmission activity factor, and $P_{mode,m} \in \{P_S, P_M, P_F\}$ corresponds to the sleep, medium, and full power modes, respectively [5].

Energy Efficiency Metric: The overall network energy efficiency is defined relative to the full-power baseline:

$$\eta_{ee}(t) = 1 - \frac{P_{total}(t)}{|\mathcal{B}| \cdot P_F} \quad (9)$$

where $P_{total}(t) = \sum_{m \in \mathcal{B}} P_m(t)$ is the total network power consumption [5].

Operational Metrics: The optimizer reports the following metrics to the operator dashboard: (i) the number of base stations eligible for sleep (*CAN SLEEP*), (ii) instantaneous power saving (kW), (iii) extrapolated daily energy saving (kWh/day), and (iv) the number of base stations currently in sleep mode (*APPLIED*).




Fig. 3. Portable 5G NIB platform used for practical validation

An estimated CO_2 reduction is computed based on a grid carbon intensity of 233 g/kWh. Additionally, a *WAKE ALL* override mechanism is provided to ensure human-in-the-loop control under uncertain or emergency conditions.

VI. PRACTICAL DEPLOYMENT: 5G NETWORK-IN-A-BOX TESTBED

To validate the complete pipeline under real OTA conditions, we developed a portable 5G NIB platform. Housed in a ruggedized carry case, the NIB integrates an NVIDIA DGX Spark, a packaged USRP radio front-end, a packaged UE, and a Raspberry Pi display into a self-contained demonstrator. As shown in Fig. 3, this design enables real-time execution of the OAI gNB stack, live KPI visualization, and hardware-in-the-loop validation of energy-efficient control actions, while remaining compact enough for rapid deployment in laboratory, field-trial, and demonstration environments.

A. Hardware Configuration

The NIB integrates an NVIDIA DGX Spark compute node, an Ettus USRP B210 radio front-end, a commercial 5G UE, and a local display in a portable enclosure. The system runs the OAI 5G NR gNB on band n78 with 40 MHz bandwidth and exports live KPMs for forecasting and control. Two simulated UEs are added to enrich the scheduling and impairment dataset used for classifier validation.

The practical deployment is complemented by a cloud-native orchestration layer and a digital-twin validation layer where AWS provides deployment-ready orchestration and life-cycle management, while AI-RSG provides safe, reproducible validation before actions are applied to the NIB hardware.

B. Controlled Impairment Injection

A key experimental advantage of the NIB testbed is the availability of a controllable channel impairment interface via the OAI channelmod telnet command. At any simulation step, the operator can inject a calibrated pathloss step (channelmod modify 0 ploss < dB_value >), inducing an immediate, precisely-timed change in the received power of the specified UE channel. This produces a causal signature in the SINR and BLER time series that serves as ground-truth event timing for supervised training of the traffic-regime classifier and for validating the anomaly trigger logic of the model selector.

C. Live KPM Collection and xApp Interface

The xApp subscribes to the OAI telnet KPM stream and collects per-UE metrics. These include RNTI, RSRP (dBm), CQI, MCS, BLER (DL/UL), PRBs allocated, and cell throughput at 10-second intervals. The BS-1 dashboard panel presents a live CQI and PRB Allocation table for nine active UEs. It also shows a PRB LOAD (%) time-series chart and a CELL THROUGHPUT (MBPS) trace. A FORECAST button triggers the active LSTM variant to generate a 10-minute-ahead regime prediction with confidence scores. The EE optimizer uses these predictions for prospective sleep planning. Representative gNB logs confirm that two UEs achieve SINR values of 23.0 and 22.5 dB. Downlink BLER is 0.000% and uplink BLER ranges from 1.8 to 6.9% at MCS 28 (256 QAM). This demonstrates stable, high-SNR operation consistent with the lab-based NIB environment.

VII. SIMULATION RESULTS AND ANALYSIS

We evaluate the framework using three groups of KPIs: *energy-efficiency KPIs* (power reduction, energy saving, number of sleeping RUs, and RU reactivation response), *network KPIs* (throughput, PRB allocation efficiency, load balancing, QoS preservation, and UE stability during switching), and *AI/ML KPIs* (test accuracy, F1 score, and confidence-aware reliability of the forecasting model). Evaluation is conducted in the VIAVI AI-RSG digital twin under realistic RF and traffic dynamics, where RU ON/OFF decisions are computed at non-real-time timescales and protected by near-real-time anomaly monitoring for rapid mitigation. The same design supports generic energy-efficiency optimization across urban, suburban, and rural topologies without scenario-specific tuning.

A. LSTM Training Convergence

Table I and Fig. 4 show stable convergence of the FP-LSTM on the NIB dataset. Validation accuracy reaches 0.955 within the first few epochs and stabilizes at 0.957, while validation loss remains near 0.108, indicating good generalization without significant overfitting. Training accuracy continues to improve to 0.963 at epoch 25, with only a small train-validation gap. These results confirm that the model captures the dominant temporal structure of the KPM sequences over the 10-minute prediction horizon.

TABLE I
FORECASTING AGENT TRAINING SUMMARY (FP-LSTM, NIB DATASET)

Metric	Epoch 1	Epoch 5	Epoch 10	Epoch 15	Epoch 25 (Final)
Train Accuracy	0.943	0.956	0.958	0.959	0.963
Val Accuracy	0.954	0.957	0.956	0.957	0.957
Train Loss	0.142	0.110	0.106	0.103	0.097
Val Loss	0.110	0.108	0.107	0.107	0.108




Fig. 4. Forecasting Model Training and Validation Performance

B. Three-Class Classification Performance

Beyond the NIB training set, the forecasting pipeline is evaluated on a larger digital twin dataset comprising 185 realistic network scenarios representing urban, suburban, and rural deployments. Each scenario consists of five two-hour periods with 30-second timesteps, yielding 435,480 samples for supervised learning. The LSTM model uses 12-timestep sequences across 11 multimodal features, including RU operational metrics, temporal signals, topological context, and environmental conditions.

TABLE II
LSTM FORECASTING ANALYSIS PERFORMANCE

Performance Metric	Value
Test Accuracy	95.91%
F1 Score (Macro)	0.852
F1 Score (Weighted)	0.959
Dataset Size	435,480 samples
Training Epochs	50

TABLE III
LSTM PER-CLASS CLASSIFICATION PERFORMANCE

Load Class	Precision	Recall	F1-Score	Support
Low (0–33%)	97.61%	98.38%	0.980	75,039
Moderate (33–67%)	85.25%	81.58%	0.834	10,941
High (67–100%)	79.01%	70.16%	0.743	1,116

Table II indicates that the forecasting model achieves a test accuracy of 95.91%, with weighted and macro F1-scores of 0.959 and 0.852, respectively. Table III demonstrates strong per-class performance, with the low-load class achieving an F1-score of 0.980. This outcome is notable because accurate identification of low-load periods enables safe opportunities for RU sleep control. The Moderate and High classes exhibit lower performance, reflecting increased volatility and class imbalance in transitional and high-demand regimes. However, this limitation is mitigated by the safety-first design, which

provides near-real-time anomaly detection and rapid RU reactivation to maintain QoS when live conditions diverge from forecasts.

Combined with the AI-RSG evaluation, these results suggest that the proposed framework generalizes effectively across diverse topologies and traffic patterns while maintaining a closed-loop, safety-aware energy optimization workflow.

VIII. CONCLUSION

This paper presented an end-to-end AgentRAN architecture that integrates intent-driven scenario generation, adaptive LSTM-based forecasting, Near-RT anomaly detection, and context-sensitive energy-efficiency optimization with a real 5G NIB testbed. The forecasting pipeline achieved 95.7% validation accuracy on the NIB dataset and 95.91% test accuracy on a larger AI-RSG dataset, while the full framework delivered a 45.6% RAN power reduction (6.03 kWh/day, 513 kg CO₂/year) in a four-BS urban scenario with zero outage events. These results show the practicality of a closed-loop O-RAN control pipeline for safe autonomous RU energy management. Future work will extend the framework to multi-agent cooperative optimization with spatial reasoning and temporal context modeling.

ACKNOWLEDGMENT

This research was supported by UK Research and Innovation (UKRI) through the EPSRC under two grants: the Technology Missions Fund project CHEDDAR (EP/Y037421/1), and Award UKRI851, focused on strategic decision-making and cooperation among AI agents in telecom safety and governance. This study does not involve human subjects or sensitive data, and raises no ethical or policy concerns.

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.
- [2] A. I. Abubakar, O. Onireti, Y. Sambo, L. Zhang, G. K. Ragesh, and M. Ali Imran, "Energy efficiency of open radio access network: A survey," in *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, 2023, pp. 1–7.
- [3] M. M. H. Qazzaz, A. Salama, M. Hafeez, and S. A. R. Zaidi, "Oran-map: A hybrid approach to mobility-aware power optimisation in open radio access networks (oran)," in *IEEE INFOCOM 2025 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2025, pp. 1–6.
- [4] A. Salama, Z. Nezami, M. M. H. Qazzaz, M. Hafeez, and S. A. Raza Zaidi, "Edge agentic ai framework for autonomous network optimisation in o-ran," in *2025 IEEE 36th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2025, pp. 1–7.
- [5] A. Salama, M. M. H. Qazzaz, S. D. Ali Shah, M. Hafeez, S. A. R. Zaidi, and H. Ahmadi, "Ecofl: Resource allocation for energy-efficient federated learning in multi-rat oran networks," in *2025 12th International Conference on Wireless Networks and Mobile Communications (WINCOM)*, 2025, pp. 1–7.
- [6] A. Salama, M. M. H. Qazzaz, S. D. A. Shah, M. Hafeez, and S. A. Zaidi, "Fedora: Resource allocation for federated learning in oran using radio intelligent controllers," 2025. [Online]. Available: <https://arxiv.org/abs/2505.19211>

- [7] A. Salama, M. M. H. Qazzaz, Z. Nezami, M. Hafeez, and S. A. R. Zaidi, "Poster: Agentic ai meets neural architecture search: Proactive traffic prediction for ai-ran," in *Proceedings of the 2nd ACM Workshop on Open and AI RAN*, ser. OpenRan '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 59–61. [Online]. Available: <https://doi.org/10.1145/3737900.3770175>
- [8] Z. Nezami, S. Bano, A. Salama, M. Hafeez, and S. A. R. Zaidi, "Safety and risk pathways in cooperative generative multi-agent systems: A telecom perspective," in *Proceedings of the 2nd ACM Workshop on Open and AI RAN*, ser. OpenRan '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 50–55. [Online]. Available: <https://doi.org/10.1145/3737900.3770174>
- [9] M. M. H. Qazzaz, A. Salama, M. Hafeez, and S. A. R. Zaidi, "Joint optimisation of load balancing and energy efficiency for o-ran deployments," 2025. [Online]. Available: <https://arxiv.org/abs/2509.09343>