

# MobiLLM: An Agentic AI Framework for Closed-Loop Threat Mitigation in 6G Open RANs

Haohuang Wen<sup>\*23</sup>, Prakhar Sharma<sup>\*12</sup>, Vinod Yegneswaran<sup>12</sup>, Ashish Gehani<sup>12</sup>, Phillip Porras<sup>12</sup>, Zhiqiang Lin<sup>23</sup>

<sup>1</sup>SRI, <sup>2</sup>SE-RAN.ai, <sup>3</sup>The Ohio State University

{prakhar, haohuang, phil}@se-ran.ai, {vinod.yegneswaran, ashish.gehani}@sri.com, {lin.3021}@osu.edu

**Abstract**—The evolution toward 6G networks is being accelerated by the Open Radio Access Network (O-RAN) paradigm—an open, interoperable architecture that enables intelligent, modular applications across public telecom and private enterprise domains. While this openness creates unprecedented opportunities for innovation, it also expands the attack surface, demanding resilient, low-cost, and autonomous security solutions. Legacy defenses remain largely reactive, labor-intensive, and inadequate for the scale and complexity of next-generation systems. Current O-RAN applications focus mainly on network optimization or passive threat detection, with limited capability for closed-loop, automated response.

To address this critical gap, we present an agentic AI framework for fully automated, end-to-end threat mitigation in 6G O-RAN environments. MobiLLM orchestrates security workflows through a modular multi-agent system powered by Large Language Models (LLMs). The framework features a Threat Analysis Agent for real-time data triage, a Threat Classification Agent that uses Retrieval-Augmented Generation (RAG) to map anomalies to specific countermeasures, and a Threat Response Agent that safely operationalizes mitigation actions via O-RAN control interfaces. Grounded in trusted knowledge bases such as the MITRE FiGHT framework and 3GPP specifications, and equipped with robust safety guardrails, MobiLLM provides a blueprint for trustworthy AI-driven network security. Initial evaluations demonstrate that MobiLLM can effectively identify and orchestrate complex mitigation strategies, significantly reducing response latency and showcasing the feasibility of autonomous security operations in 6G.

## I. INTRODUCTION

The rise of sixth-generation (6G) wireless networks is tightly coupled with the Open Radio Access Network (O-RAN) paradigm[1]—a shift toward openness that is transforming the telecom network ecosystem. By enabling third-party developers to deliver modular xApps and rApps that interoperate across vendors, O-RAN unlocks rapid innovation and paves the way for 6G adoption in mission-critical domains such as remote healthcare, smart manufacturing, and autonomous transportation. Yet, this open, multi-vendor environment also introduces major challenges. The cost to manage such a complex network is expected to be extremely high, due to its sophisticated protocols and the need for deep operational expertise. Meanwhile, 6G networks will likely inherit known vulnerabilities from prior standards—including over-the-air

attacks from malicious devices, rogue base stations, and man-in-the-middle threats—while also experiencing new run-time faults like registration failures, software misconfigurations, and other unexpected behaviors [11], [10], [15], [4], [12], [6]. If not addressed properly, these persistent threats and operational instabilities will weaken the trust required to deploy 6G in security-critical domains.

The current research landscape reveals two parallel streams of work in O-RAN: passive security detection and non-security-focused automation. On one hand, significant research has concentrated on threat detection. This includes various xApps designed to identify network anomalies [22] and malicious activities such as the presence of rogue UEs and base stations [21], [24], [9], [25]. A critical limitation of these solutions, however, is their passive nature; they are designed to report threats but lack the capability to automatically respond or mitigate them. On the other hand, numerous closed-loop xApps and rApps have been developed for network automation and performance optimization tasks like network slicing and resource management [19], [13], [5], [28]. These systems, while demonstrating the potential of automation, were not created to address cybersecurity challenges. This leaves a crucial gap for a framework that can unify intelligent detection with proactive, closed-loop security response.

More recently, Large Language Models (LLMs) have emerged as a promising tool for O-RAN network management, such as network slicing and resource management [17], [26]. While these studies demonstrate the applicability of LLMs for network orchestration [23], [2], they exhibit several limitations that prevent them from filling the aforementioned security gap. First, none of these LLM-based frameworks address the crucial aspect of closed-loop security management. An effective security system requires more than simple anomaly detection; it needs a sophisticated analysis engine grounded in the latest threat intelligence and an intelligent response engine capable of enacting changes to the O-RAN network configuration. Second, most of these works rely on basic prompt engineering with general-purpose LLMs, a method prone to hallucinations and irrelevant output when dealing with the highly technical and standardized nature of cellular networks. This highlights the need for robust knowledge grounding to ensure accurate, task-specific outputs. Consequently, the challenge of creating a framework that can autonomously detect, analyze, and respond

The first two authors contributed equally to this work

remains unsolved.

In this paper, we present MobiLLM to address these gaps by introducing a knowledge-grounded, multi-agent framework dedicated to closed-loop security response in O-RAN environments. However, pre-trained models, despite their impressive general capabilities, lack the deep, domain-specific knowledge contained within the extensive and highly technical specifications that govern cellular networks. Hence entrusting security-related tasks to an LLM is inherently risky. Unlike human experts, LLMs can hallucinate—generating plausible but false information, which could lead to incorrect and potentially disruptive mitigation actions in a live network. Therefore, the design of MobiLLM focuses on implementing robust guardrails that ensure its trustworthiness and operational safety.

To achieve this, we introduce two key innovations: first, to leverage the LLM as a powerful natural language parser that maps real-time threat descriptions to the structured MITRE FiGHT framework [18], and second; employing the ReAct pattern [27] to translate the high-level mitigation strategy from FiGHT into a sequence of calls to a small, predefined set of safe, human-written APIs. This ensures that all interactions with the network are constrained to vetted, predictable operations, making the system’s behavior both effective and verifiable.

**Contributions.** Our main contributions include:

- Development of MobiLLM, the first multi-tiered agentic framework for cellular threat analysis, planning, and response<sup>1</sup>.
- Evaluation of MobiLLM with a suite of five real-world 5G threat scenarios. We release associated RRC protocol traces and initial incident reports for each threat.
- Discussion of guardrails to prevent misaligned or hallucinated LLM outputs from affecting changes in the network.

## II. SYSTEM DESIGN AND IMPLEMENTATION

We present the overall design of MobiLLM, a modular agentic AI framework created to manage the complex security needs of 6G networks. Applying existing LLM-based agentic frameworks directly to the high-stakes, dynamic 6G environment is not feasible, as mobile networks require highly specialized knowledge, an area where generic pre-trained LLMs are deficient. For instance, a recent GSMA report highlights that even state-of-the-art models like GPT-4 struggle to answer highly technical questions related to 3GPP specifications and telecommunication protocols [7].

To overcome this critical knowledge gap, MobiLLM is built as a multi-agent architecture that augments LLM capabilities with explicit domain awareness. The framework decomposes the complex security pipeline, i.e., from threat analysis to closed-loop mitigation, into three distinct stages. Each stage is handled by a specialized LLM-based agent that operates in sequence, creating a streamlined workflow. By assigning each agent a specific subtask with a precisely crafted prompt, our

design maximizes accuracy and efficiency, ensuring that the system’s reasoning and actions are grounded in the specific context of 6G network operations.

**Deployment Scenarios.** A potential deployment scenario for MobiLLM is within the O-RAN’s Service Management and Orchestration (SMO) infrastructure [20], where it can be instantiated as a modular rApp while not compromising the overall network latency in a non-critical network path. Within this architecture, MobiLLM is positioned to work synergistically with a wide range of threat and anomaly detection xApps and rApps deployed in the O-RAN control plane [21], [24], [9], [22], and perform event analysis, classification, as well as closed-loop mitigation. The major interactions are conducted through the standard O-RAN interfaces such as O1 and A1, to seamlessly work with other network components including xApps, rApps, O-CU, and O-DU. The following describes the detailed design of MobiLLM.

### A. Threat Analysis Agent

The Threat Analysis Agent serves as the initial triage system. Its primary input is not just a simple prompt, but a detected threat or abnormal event, which could originate from network monitoring tools, an O-RAN xApp/rApp, or a natural language query from a human operator. Upon receiving an alert, the agent’s first move is beyond the surface-level description by actively inspecting the event and correlating it with related network data, querying sources such as performance metrics, RAN-level data, and logs. This process allows the agent to analyze the context, evaluate the potential risk, and distinguish a genuine security threat from a benign network fault or false positive. If a credible threat is identified, the agent generates a structured, machine-readable threat report containing the event summary, affected components, and an initial risk assessment. If no threat is detected, it logs the analysis and closes the event, ensuring that downstream agents do not expend unnecessary computational resources.

### B. Threat Classification Agent

The main goal of the Threat Classification Agent is to solve a critical problem with using LLMs for security: their lack of specialized knowledge and their tendency to hallucinate. Our novel approach doesn’t trust the LLM to invent solutions. Instead, we use it as an intelligent *translator* to ground its reasoning in a reliable, human-curated knowledge base. The key innovation is leveraging the MITRE FiGHT framework [18] as this trusted source. This agent bridges the gap between a potentially ambiguous threat alert and a concrete, verifiable action plan. By forcing the agent to map every threat to a standard FiGHT technique, we ensure its classifications are accurate, consistent, and based on industry best practices. This transforms the LLM from an unreliable narrator into a dependable reasoning engine. The agent follows a systematic, two-phase process to achieve this:

- **Phase-I: Offline Knowledge Base Preparation.** First, we crawl the entire public MITRE FiGHT framework to gather

<sup>1</sup>MobiLLM has been released at <https://github.com/5GSEC/MobiLLM>

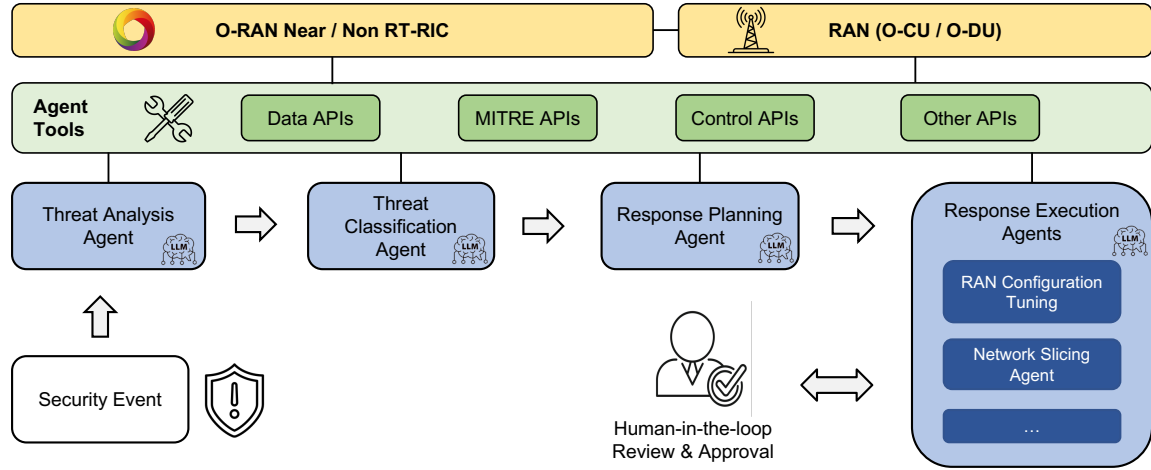


Figure 1: Overview of MobiLLM’s agentic architecture. The config tuning agent can be invoked autonomously, but the effects are only applied after human approval is gathered.

all attack techniques and their detailed descriptions. This text data is then processed by a text embedding model, which converts the description of each technique into a numerical vector. Finally, these vectors are stored and indexed in a specialized vector database, creating a searchable library of 6G threats.

- **Phase-II: Run-time Threat Classification.** The agent receives a structured threat report from the Threat Analysis Agent. It uses a Retrieval-Augmented Generation (RAG) process. The natural language description from the report is used as a query to search the vector database. The search retrieves the top matching MITRE FiGHT technique(s) based on semantic similarity. The agent then outputs a formal classification using the FiGHT technique ID and, most importantly, extracts the mitigation guidance associated with that technique, providing a clear plan for the next agent in the chain.

### C. Response Planning and Execution Agent

The final agent in the MobiLLM pipeline is responsible for translating a mitigation plan into safe, real-world action. The central motivation behind its design is to solve the most critical challenge in autonomous security: how to grant an LLM control over a live network without exposing the system to the risks of hallucination or unsafe command execution. Our novel contribution is a two-layer architecture that separates high-level planning from low-level execution. We use the LLM as a sophisticated planner that reasons about what to do, but we strictly limit how it does it. The LLM’s creative capabilities are used to dynamically construct an action plan by mapping the mitigation guidance from MITRE FiGHT [18] to a set of available network control APIs. However, the final execution is delegated to specialized *Action Agents*, which operate within rigid, pre-defined workflows. An example of such includes a *Config Tuning Agent* that specifies a workflow to update the RAN configuration and perform a reboot to let

it take effect. This design provides the flexibility of LLM-based planning while enforcing the safety of deterministic, human-vetted code, ensuring that every action taken is both intelligent and trustworthy. The agent executes its task through the following step-by-step process:

- **Action Plan Generation.** The agent receives the mitigation goal from the previous stage. It then analyzes its library of available, human-written safe APIs and generates a step-by-step plan to achieve the goal. If a viable plan cannot be constructed using only the provided tools, it proceeds directly to Step 3.
- **Execution via Safe Workflows.** If a valid plan is created, it is dispatched to the appropriate specialized Action Agents. Each of these agents follows its own pre-defined, hard-coded workflow to execute its part of the plan. This critical safety guardrail ensures the LLM can only provide parameters while the rigid script controls the execution logic (the how), preventing any unsafe or unintended operations.
- **Revert on Failure.** If no valid plan could be created in the first step, the agent does not attempt to improvise. Instead, it flags the issue and presents the original mitigation guidance as a high-level recommendation to a human operator, ensuring that a human is always in the loop for complex or unsupported threats.

### D. Implementation Details

Following the aforementioned architecture of MobiLLM, we now describe the implementation details.

**Knowledge Grounding.** We base the design of MobiLLM’s agentic architecture on two recent LLM advancements: ReAct agents [27], and retrieval-augmented generation (RAG). The former is a few-shot prompting technique that has demonstrated good results in reasoning and action-based (tool calling) tasks. The latter is a technique to provide inference-time context to the LLM, and helps grounding the LLM response

in existing verified knowledge bases and curb issues like hallucinations, meandering output and ill-formed responses. We employ numerous sources as RAG knowledge bases viz. MITRE FiGHT [18], 3GPP standard specifications, O-RAN documents and curated academic research papers documenting relevant 5G/LTE attacks and their mitigations. The main ones that interface with respective agents are described below:

- **Network Data APIs:** These enable the threat analysis agent to gain an in-depth insight about the current network status at various levels of granularity. (e.g., fetching network traffic, logs, and detected events)
- **MITRE APIs:** The MITRE FIGHT framework [18] documents known 5G / 6G tactics, techniques and mitigations. This knowledge base of adversary behaviors is recorded as a database of popular techniques, their operational modalities, target attack surfaces and proposed mitigations. For any identified threat, this database is one of the most direct resources for crafting a response. These APIs enable the threat classification agent to search the MITRE FIGHT database for relevant tactics and techniques (and associated procedures and mitigations) given a certain security incident and threat analysis.
- **Control APIs:** These functions enable the RAN configuration tuning agent to perform an action in a human-supervised control loop. An agent can fetch the active configuration of a CU, suggest modifications, and, upon receiving human approval, apply the updated configuration and reboot the RAN to enact the changes. The scope of controllable parameters can be readily expanded by incorporating additional parameters from the O-RAN E2SM-RC (RAN Control) standard [20].

**Agent Tools and Prompts.** The threat analysis, threat classification, security planning and action agent interface with the knowledge bases, RAN control system, configuration tuning subroutines and system reboot functions to dynamically analyze and craft a response to an ongoing threat. The agents are free to call tools in any fashion (single-turn, multi-turn, multi-step, etc.). We use React-style few-shot prompts that encourage the model to think, reason, and act in an iterative loop. Max iterations for that loop is five.

**Guardrails:** To address LLM issues viz. hallucinations, irrelevant outputs, and misinterpretation of contextually similar inputs, we incorporate multi-layered guardrails designed to ensure robustness, correctness, and operational safety:

- **Prompt-Level Guardrails:** Carefully crafted few-shot prompts are used to constrain the LLM’s output to a predefined structure.
- **Rule-Based Guardrails:** Deterministic output checkers and JSON sanitizers are applied to filter out hallucinated or malformed outputs and enforce schema compliance.
- **Human-in-the-Loop Validation:** Given the high-stakes nature of integration with systems like 6G networks, human feedback is employed to review and validate LLM-generated

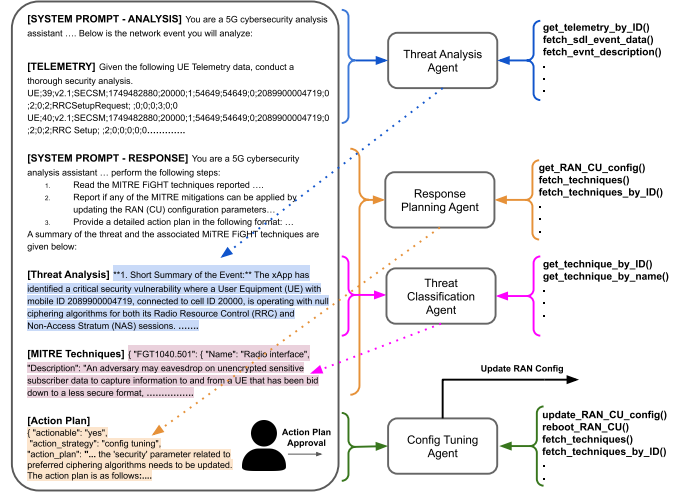


Figure 2: This figure illustrates the interaction flow of MobiLLM, where specialized agents sequentially process incident telemetry to classify a threat and generate a response plan. The resulting configuration change is then presented for mandatory human approval before being executed by the appropriate tool.

mitigation scripts. This step ensures that unintended or unsafe logic is not propagated into production environments.

### III. EVALUATION

Since MobiLLM represents the first agentic architecture capable of analyzing, classifying, and autonomously responding to threats within 6G O-RAN networks, defining a robust evaluation methodology posed a significant challenge. In the absence of a standardized dataset or established evaluation protocol for this domain, we curated a focused set of five representative attacks that 6G networks inherit from prior standards (e.g., LTE and 5G NR): Blind Denial-of-Service (DoS), Base Transceiver Station (BTS) resource depletion, Downlink IMSI exposure, Uplink IMSI exposure, and Null cipher exploitation [11], [6], [10], which were reproduced end-to-end within a testbed using open-sourced software including OpenAirInterface and OpenRAN Software Community RIC.

#### A. Evaluation Setup

We construct a 5-case evaluation suite representative of the aforementioned 6G attacks. For each threat we record (i) A threat scenario describing the threat and expected observables; (ii) Curated telemetry (RRC and NAS traffic telemetry) corresponding to the attacks; and (iii) The MITRE FiGHT [18] tactics classification associated with each threat and their corresponding mitigations. We use the scenario description and the curated telemetry as context with the initial prompt to the threat analysis agent. To evaluate the threat classification performance, we ask a few cellular network experts to generate ground truth results for all the targeted attacks.

We evaluate MobiLLM using the following metrics: (1) MITRE FIGHT technique retrieval accuracies (2) Agent tool invocation validity: whether agent tools were recommended in

the correct scenario (3) Overall remediation validity, which are validated by cellular network experts. We additionally report end-to-end latency per framework invocation.

Threat	Top-3	Top-1	CCR
BTS-Attack-1	1	0.6	0.8
BTS-Attack-2	1	1	0.4
BTS-Attack-3	1	0.6	0.8
Blind-DoS-1	1	1	1
Blind-DoS-2	1	1	1
Blind-DoS-3	1	1	0.6
Downlink-DoS	1	0.4	0.4
Downlink-IMSI	0.8	0	0
Null-Cipher-Integrity	0.8	1	1
Uplink-IMSI	0.8	0.4	0.4
Mean	0.94	0.72	0.64

Table I: Per-sample MITRE FiGHT Technique retrieval performance (Top-3 and Top-1 accuracy), and Per-threat tool correct calling ratio (CCR) across five MobiLLM runs

### B. Quantitative Results

Our per-attack evaluation, summarized in Table I, reveals that MobiLLM achieves high accuracy in threat classification. Specifically, the correct MITRE FiGHT mitigation technique was present in the Top-3 retrieved results 94% of the time. However, this success in classification does not fully translate to the action phase. As shown in Table I, the accuracy for correctly invoking the final mitigation action is considerably lower, at an average of 64%. To investigate the root causes of this performance gap, we analyze two representative case studies in the following section.

### C. Overall Remediation Validity

To illustrate MobiLLM’s remediation capabilities and limitations, we analyze two representative case studies: one successful mitigation and one failure.

**Null Cipher-Integrity Attack.** A Null Cipher-Integrity attack occurs when a misconfigured base station forces a UE to disable confidentiality (EEA0) or integrity protection (EIA0) during its security setup, exposing user data. In our evaluation, MobiLLM executed a successful, human-supervised remediation workflow, as in Figure 2. First, the Threat Classification Agent correctly mapped the anomaly to the corresponding MITRE FiGHT technique, ‘FGT1600.501’ (Disabling Encryption Over Radio Interface). Based on this classification, the Response Planning Agent used the *get\_ue\_description* tool to gather context and correctly drafted a plan to remove the insecure ‘nea0’ and ‘nia0’ algorithms from the CU’s configuration. Before execution, the agentic workflow paused for mandatory human approval. Once granted, the Config Tuning Agent successfully applied the changes using the *get\_ran\_cu\_config* and *update\_ran\_cu\_config* tools, resolving the underlying vulnerability.

**Downlink IMSI Extraction Attack.** A Downlink IMSI Extraction attack involves a man-in-the-middle attacker who intercepts and modifies a pre-authentication RRC message, forcing the UE to reveal its permanent identity (IMSI) in plain

text before security is established [6]. While MobiLLM again succeeded in classifying the attack with the correct MITRE FiGHT techniques, it failed during response planning. The agent mistakenly concluded that the threat could be mitigated by enforcing integrity protection on the connection. This plan was fundamentally flawed because the attack exploits a timing vulnerability that occurs before the authentication stage where such protections are activated. The root cause of this error lies in the limitations of the general-purpose LLM (Gemini 2.5 Flash) used in our evaluation, which lacks deep, specialized knowledge of the 3GPP state machine and specifications. This failure highlights the necessity of fine-tuning LLMs with domain-specific awareness to handle sophisticated cellular threats, a direction we explore in the discussion section.

## IV. DISCUSSION AND FUTURE WORK

### A. General purpose vs. specialized agents

MobiLLM is a promising direction towards fully automated threat incidence analysis and response. We perform all our experiments with Gemini-2.5-Flash as the core deduction engine for both analysis generation and function calling, and the distinction between agents is made through what tools they have at their disposal (geared towards a specific task). While knowledge bases and tools help ground the LLM response, further fine tuning is imperative for state-of-the-art performance for the following reasons:

- API-based off-the-shelf models are ephemeral and costly. Training bespoke LoRA [8] adapters enables light-weight SLMs [3] to achieve better performance in agentic tasks. We plan to bootstrap the current capabilities of our Agentic architecture to create a supervised fine-tuning dataset; to train small but specialized language models specifically for crafting a threat response in the context of O-RAN 5G.
- Currently, we supply network facts like CU configurations and parameters at inference time, which makes it slow and costly. Through supervised fine tuning, this information can be baked into the model parameters. We also plan to expand the use case to more than 5 attacks that we currently support, with tools designed to address a wide range of threats.
- Existing work has been successful in using telecom-specific knowledge to fine-tune an LLM to enhance domain awareness and understanding by fine-tuning on the vast corpus of 3GPP and O-RAN specifications [29]. In the security context, we can create a true network security expert with deeper protocol understanding and more accurate tool-use capabilities than any general-purpose LLM.

### B. Accessible and robust O-RAN security

We plan to integrate MobiLLM into real-world 6G deployments. As initiatives like the U.S. government’s Open Centralized Unit Distributed Unit (OCUDU) project accelerate the creation of open 6G platforms, MobiLLM can be embedded as a core security layer to not only enhance network resilience and trust from the outset, but also cut down network operator

training cost. We plan to develop more advanced security guardrails to protect MobiLLM from adversarial attacks like backdoor attacks [16] and prompt injection attacks [14]. This involves creating stricter, verifiable rules for its API interactions to defend against threats like prompt injection and ensure its operational integrity in mission-critical environments.

## V. CONCLUSION

MobiLLM is an agentic AI framework designed to provide fully automated, end-to-end security for 6G O-RAN environments. Our approach bridges the critical gap between passive threat detection and active, closed-loop response. The framework's design presents a novel blueprint for building trustworthy AI-driven network defenses by integrating RAG with trusted knowledge bases like MITRE FiGHT and incorporation of robust safety guardrails. Our evaluations confirm that MobiLLM can effectively identify threats from five types of cellular attacks reproduced in a testbed, orchestrate complex mitigation strategies, and significantly reduce response latency, demonstrating the feasibility and potential of autonomous security operations in next-generation networks.

## ACKNOWLEDGEMENTS

This research was supported by a Small Business Innovation Research (SBIR) Phase I award N6893625C0023 from the Naval Air Warfare Center (NAWC), the NSF convergence accelerator program under award ITE-2326882, and the Google Cloud Research Credits Program.

## REFERENCES

- [1] O-ran alliance. <https://www.o-ran.org/>, May 2024.
- [2] L. Bao, S. Yun, J. Lee, and T. Q. S. Quek. Llm-hric: Llm-empowered hierarchical ran intelligent control for o-ran. 2025.
- [3] P. Belcak, G. Heinrich, S. Diao, Y. Fu, X. Dong, S. Muralidharan, Y. C. Lin, and P. Molchanov. Small language models are the future of agentic ai. 2025.
- [4] E. Bitsikas and C. Pöpper. You have been warned: Abusing 5g's warning and emergency systems. In *Proceedings of the 38th Annual Computer Security Applications Conference*, pages 561–575, 2022.
- [5] S. D'Oro, L. Bonati, M. Polese, and T. Melodia. Orchestran: Network automation through orchestrated intelligence in the open ran. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*, pages 270–279. IEEE, 2022.
- [6] S. Erni, M. Kotuliak, P. Leu, M. Röschlin, and S. Capkun. Adaptover: adaptive overshadowing attacks in cellular networks. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*, pages 743–755, 2022.
- [7] GSMA. Gsma open-telco llm benchmarks. [huggingface/online](https://huggingface.com/gsma-open-telco-llm-benchmarks), February 2025.
- [8] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen. Lora: Low-rank adaptation of large language models. 2021.
- [9] J.-H. Huang, S.-M. Cheng, R. Kaliski, and C.-F. Hung. Developing xapps for rogue base station detection in sdr-enabled o-ran. In *IEEE INFOCOM 2023-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 1–6. IEEE, 2023.
- [10] S. R. Hussain, M. Echeverria, I. Karim, O. Chowdhury, and E. Bertino. 5greasoner: A property-directed security and privacy analysis framework for 5g cellular network protocol. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 669–684, 2019.
- [11] H. Kim, J. Lee, E. Lee, and Y. Kim. Touching the untouchables: Dynamic security analysis of the lte control plane. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1153–1168. IEEE, 2019.
- [12] M. Kotuliak, S. Erni, P. Leu, M. Roeschlin, and S. Čapkun. Ltrack: Stealthy tracking of mobile phones in lte. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 1291–1306, 2022.
- [13] A. Lacava, M. Polese, R. Sivaraj, R. Soundrarajan, B. S. Bhati, T. Singh, T. Zugno, F. Cuomo, and T. Melodia. Programmable and customized intelligence for traffic steering in 5g networks using open ran architectures. *IEEE Transactions on Mobile Computing*, 2023.
- [14] D. Lee and M. Tiwari. Prompt infection: Llm-to-llm prompt injection within multi-agent systems. 2024.
- [15] G. Lee, J. Lee, J. Lee, Y. Im, M. Hollingsworth, E. Wustrow, D. Grunwald, and S. Ha. This is your president speaking: Spoofing alerts in 4g lte networks. In *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*, pages 404–416, 2019.
- [16] Y. Li, H. Huang, Y. Zhao, X. Ma, and J. Sun. Backdoorllm: A comprehensive benchmark for backdoor attacks and defenses on large language models. 2025.
- [17] F. Lotfi, H. Rajoli, and F. Afghah. Oran-guide: Rag-driven prompt learning for llm-augmented reinforcement learning in o-ran network slicing. 2025.
- [18] MITRE. Mitre — fight. [mitre/online](https://mitre.org/fight).
- [19] M. Polese, L. Bonati, S. D'Oro, S. Basagni, and T. Melodia. Coloran: Developing machine learning-based xapps for open ran closed-loop control on programmable experimental platforms. *IEEE Transactions on Mobile Computing*, 22(10):5787–5800, 2022.
- [20] M. Polese, L. Bonati, S. D'Oro, S. Basagni, and T. Melodia. Understanding o-ran: Architecture, interfaces, algorithms, security, and research challenges. *arXiv preprint arXiv:2202.01032*, 2022.
- [21] A. Scalingi, S. D'Oro, F. Restuccia, T. Melodia, D. Giustiniano, et al. Det-ran: Data-driven cross-layer real-time attack detection in 5g open rans. In *IEEE International Conference on Computer Communications*, pages 1–10, 2024.
- [22] C. Sun, U. Pawar, M. Khoja, X. Foukas, M. K. Marina, and B. Radunovic. Spotlight: Accurate, explainable and efficient anomaly detection for open ran. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, pages 923–937, 2024.
- [23] Y. Tang, U. C. Srinivasan, B. J. Scott, O. Umealor, D. Kevogo, and W. Guo. End-to-end edge ai service provisioning framework in 6g oran. 2025.
- [24] H. Wen, P. Porras, V. Yegneswaran, A. Gehani, and Z. Lin. 5g-specter: An o-ran compliant layer-3 cellular attack detection service. In *Proceedings of the 31st Annual Network and Distributed System Security Symposium (NDSS'24)*, San Diego, CA, February 2024.
- [25] H. Wen, P. Sharma, V. Yegneswaran, P. Porras, A. Gehani, and Z. Lin. 6g-xsec: Explainable edge security for emerging openran architectures. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*, pages 77–85, 2024.
- [26] X. Wu, J. Farooq, Y. Wang, and J. Chen. Llm-xapp: A large language model empowered radio resource management xapp for 5g o-ran. *NDSS symposium*, 2025.
- [27] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafraan, K. Narasimhan, and Y. Cao. React: Synergizing reasoning and acting in language models. 2023.
- [28] S.-p. Yeh, S. Bhattacharya, R. Sharma, and H. Moustafa. Deep learning for intelligent and automated network slicing in 5g open ran (oran) deployment. *IEEE Open Journal of the Communications Society*, 5:64–70, 2023.
- [29] H. Zou, Q. Zhao, Y. Tian, L. Bariah, F. Bader, T. Lestable, and M. Debbah. Telecomgpt: A framework to build telecom-specific large language models. *IEEE Transactions on Machine Learning in Communications and Networking*, 2025.