AI-Based MITRE ATT&CK Detection System: A Feasibility Study

Dimitris Koutras
Focal Point-sprl
Waterloo, BELGIUM
0000-0002-9154-8340

Michalis Karamousadakis
Plaixus P.C.
Athens, Greece
0000-0002-9411-2139

Giannis Konstantinidis
Plaixus P.C.
Athens, Greece
0009-0006-9852-6941

Christos Grigoriadis
Focal Point-sprl
Waterloo, BELGIUM
0000-0002-3192-667X

Vangelis Malamas
Department of Informatics
University of Piraeus
Piraeus, Greece
0000-0001-9238-6796

Panayiotis Kotzanikolaou
Department of Informatics
University of Piraeus
Piraeus, Greece
0000-0002-8771-9020

***Abstract*—The rise in cyber threats necessitates automated detection systems that can effectively identify and respond to hostile techniques. This paper presents a feasibility assessment of adopting Large Language Models (LLMs) to enhance cybersecurity operations within the MITRE ATT&CK framework. We research how AI can automate Kusto Query Language (KQL) development to better cyber threat detection in Microsoft Sentinel. We start with prompt engineering to improve AI-generated queries, then compare LLMs to determine the top models. Through successive breakthroughs, we progressed from a naïve prompting method to an advanced Chain of Thought (CoT) prompting technique, enabling AI models to give more contextually accurate and structured KQL queries. We extensively tested both open-source and closed-source models, evaluating their performance using two separate accuracy scoring formulae. Our results demonstrate that CoT significantly enhances the precision of AI-generated queries, while ChatGPT-4o-mini surpasses other models in generating structured KQL queries. Our technology leverages real-time MITRE ATT&CK Intelligence and Microsoft Sentinel log analysis for automated threat identification and response in order to minimize human effort and enhance productivity. Our approach applies AI to automate cybersecurity tasks, whereas most other research on LLM-assisted security analytics remains theoretical and thus fills an important gap between theory and practice.**

***Index Terms*—MITRE ATT&CK, AI-driven cybersecurity, LLMs, prompt engineering, threat detection, KQL queries**

I. Introduction

This feasibility study describes the initial steps in adapting Large Language Models [1] such as GPT-4 [2] within the domain of detection engineering. The AI-driven MITRE ATT&CK Detection system was designed to help cybersecurity professionals automate the detection of cyber threats. It will use advanced natural language processing techniques, integrating several external data sources such as Microsoft Sentinel and the MITRE ATT&CK framework, to automatically develop KQL queries that detect certain attack techniques. These were several stages: from crafting a fine prompt to testing the latest language models for which one better will provide proper query generation. The following report details the most significant steps taken during its development: in the stage of engineering of the prompts, on the stage of selecting models, and in the further features based on those developments.

A. Motivation - Contribution

Rapid development of cyber threats requires more effective and automated mechanisms for threat detection. Traditional cybersecurity products rely on rule-based detection, requiring human definition, which naturally involves developments that are very time-consuming and prone to human error. In this respect, LLMs provide a promise for automating the generation of KQL queries based on up-to-date intelligence about threats in real time. However, the integration of LLMs into cybersecurity workflows also comes with several challenges related to prompt engineering, model selection, and system scalability. The work presented here seeks to address these challenges by developing an AI-based MITRE ATT&CK detection system for improving the efficiency and precision of cyber threat detection.

This study describes an AI-driven cyber threat detection system that uses Large Language Models (LLMs) to automate the development of the Kusto query language (KQL) within the MITRE ATT&CK framework, enhancing efficiency and precision in cybersecurity operations. We create a better prompt engineering method that uses Chain of Thought (CoT) prompting to improve context and accuracy in inquiry generation. A detailed study of cutting-edge LLMs is undertaken to find the most efficient models for cybersecurity objectives. Furthermore, we apply a dual-scoring approach for complete performance evaluation, assuring a dependable assessment of AI-generated inquiries. Our tech enhances security operations by automating and providing thorough query explanations, helping analysts optimize threat detection. This study shows that AI-driven automation can speed detection, reduce human labor, and improve decision making in threat response, making it a strong foundation for its use in cybersecurity.

B. Related Work

We can find in the literal review several references concerning the MITRE ATT&CK framework in combination with the Kusto Query Language (KQL). Georgiadou et al. [3]extend the usefulness of MITRE ATT&CK beyond adversary emulation and red-teaming by relating it to aspects of organizational and individual culture influencing security vulnerabilities [4]. While their work stresses security culture and organizational elements in risk assessment, our study focuses on the technical automation of cyber threat detection using Large Language Models (LLMs). Unlike their qualitative risk rating technique, our system delivers a realistic implementation by automating KQL development to detect real-time adversary behavior. Our research also improves cybersecurity operations by applying advanced rapid engineering approaches. Huang et al. [5] presents a method that will improve the efficiency of Cyber Threat Intelligence analysis, which enables automatic extraction of MITRE ATT&CK tactics from unstructured threat reports using deep learning such as BERT and ontology-based fusion with the purpose of increasing classification accuracy to solve some problems of identifying TTPs manually. Unlike their approach (using text classification), our technology directly integrates with Microsoft Sentinel and employs advanced prompt engineering with Large Language Models to create actionable detection questions. Cao et al. [6] suggest Chain of Thought (CoT) Prompting as a security-focused improvement to Large Language Models (LLMs) to prevent jailbreak attacks. They improve model self-regulation without extra training using a five-stage CoT methodology. While their research deals with making LLMs more robust against malicious manipulation, our contribution employs CoT prompting in a cybersecurity detection context to improve the accuracy of automatic KQL query generation. While their security hardening strategy works by impeding malicious outputs, our effort reinforces reasoning capabilities to detect cyber risks on the fly. Szmurlo et al. [7] explore how chatbots in cybersecurity serve dual purposes as both defensive and offensive tools. Our research, rather than comparing the overall defensive and offensive roles of chatbots, builds a customized AI-powered detection system using LLMs to automate KQL query generation in improving threat intelligence and detection workflows. Cardenoso et al. [8] compare open-source and closed-source models to find the best LLM-based methods for translating natural language inquiries into structured query formats. Their study evaluates fine-tuning versus prompt engineering, highlighting that commercial tools can achieve strong performance without extensive training, provided that resource constraints, data privacy, and dataset availability are considered. Our research shows LLMs' real-world performance in security operations and incident handling by providing a practical, domain-specific implementation that automates threat identification and response, unlike their wide model comparison approach. Hossain et al. [9] introduce the automated event categorization system in SIEM to overcome the emerging problem of human event classification due to the increase in cyberattack incidents. The technology uses IBM QRadar's SIEM categorization architecture for the evaluation and classification of security events using machine learning. Kasri et al. [10] discussed the role of Large Language Models in cybersecurity: threat detection, vulnerability assessment, malware analysis, and automatic development of security policies. Their contribution also displays some beneficial LLM-driven advancement: large datasets analysis, phishing detection, and assistance for real-time decisions.

II. Methodology

Building an AI-based system for detecting MITRE ATT&CK requires a well-thought-through iterative process that will ensure the resulting KQL queries are accurate, effective, and of high quality. This section describes the basic steps in our process: crafting prompts, choosing models, and evaluating results. We improved our method step by step, starting with a basic setup and moving to more advanced techniques. Our goal was to help the system create better and more accurate questions for detecting hacking threats. Our method includes comparing different LLMs to make sure we pick the best one for real-life security tracking.

1) First experimental prompting approach: First of all, formulating the prompt that the AI model would use to generate an accurate and contextually relevant KQL query was the first crucial step in the development of the system. Initially, we employed a preliminary experimental prompting strategy utilizing open-source models, predicated on the premise that for a specific MITRE ATT&CK method ID or technique name, the LLM would identify its corresponding technique name or technique ID. We assessed thirty-six (36) distinct open-source state-of-the-art large language models, with the findings presented in I.

As can be observed, the results from this approach were suboptimal, with the models achieving an accuracy of less than 10%. The models struggled to correctly associate the MITRE ATT&CK technique names with their corresponding IDs (and vice versa) due to a lack of contextual understanding. Key issues observed were the length of the prompt (very small, 1-2 sentences), lack of contextual information (e.g. lack of a detailed description of the MITRE ATT&CK techniques) and requesting numerical outputs (e.g. IDs) from the LLMs, a task that is well known that LLMs struggle on.

2) Chain of Thought prompting approach: We improved our approach since the previous one did not go that well. We used the Chain of Thought prompting technique to enhance the form of our prompts; in doing this, it helps the AI model think step by step before answering, improving at understanding and generating accurate KQL queries for cyber threat detection. Instead of just asking the model to provide a query in one step, we gave it instructions to explain its reasoning before creating the final result. This made a big difference in accuracy.

To get better results, we tested more advanced AI models, including closed-source models like GPT-4, Claude, and Qwen because they are generally more accurate than open-source models. These models (ChatGPT-4, ChatGPT-3.5, Llama 3.1,

TABLE I
First Experimental Prompting Approach Results

Model Name	ID Accuracy	Name Accuracy
Llama3-70B	0.6%	0.3%
WizardLM2	0.9%	0.0%
Mistral-7B	3.0%	0.0%
Gemma-7B	0.64%	0.0%
Llama2-7B	0.96%	0.0%
CodeGemma-7B	5.7%	0.0%
Llava-7B	1.3%	0.3%
Codellama-7B	8.7%	0.0%
Qwen-4B	0.0%	0.0%
Phi3	0.0%	0.0%
Llama2-Uncensored	0.3%	0.0%
DeepSeek-Coder-6.7B	0.0%	0.0%
Mistral-OpenOrca	1.6%	0.0%
Dolphin-Mistral	0.9%	0.3%
Phi	0.0%	0.0%
Orca-Mini-7B	0.3%	0.0%
Zephyr-7B	2.8%	0.0%
Wizard-Vicuna-Uncensored-7B	0.0%	0.0%
Vicuna-7B	0.0%	0.0%
TinyLlama	0.0%	0.0%
Dolphin-Llama3-8B	0.3%	0.0%
OpenHermes	1.9%	0.0%
Yi-6B	0.0%	0.0%
OpenChat	0.9%	0.0%
TinyDolphin	0.0%	0.0%
Stable-Code	0.9%	0.0%
Neural-Chat	0.6%	0.0%
Wizard-Math-7B	0.0%	0.0%
Starling-LM	0.64%	0.0%
DolphinCoder-7B	0.64%	0.0%
Nous-Hermes-7B	0.0%	0.3%
Orca2-7B	0.0%	0.0%
StableLM2-1.6B	0.0%	0.0%
SQLCoder-7B	0.0%	0.0%
Dolphin-Phi	0.0%	0.0%
CodeQwen-7B	4.5%	0.0%

CodeStral, and Qwen 2) are better at understanding complex cybersecurity information and handling structured reasoning.

We had to experiment during testing with the number of real-world examples that we should use in the prompt. We needed to know whether the model will do better at generating queries once it is provided with more examples of cybersecurity cases. The truth is, it helped with extra examples, but there was a point beyond which additional examples didn't really make much difference. We also modified a setting known as temperature that controls how much randomness the model adds to its answers. This resulted in a more predictable model and gave quite good accuracy with a lowered temperature. With higher temperature, there could be some occasions when mistakes could occur.

Overall, by doing incremental reasoning, using superior AI models, and cleaning the training instructions better, much more precise and useful KQL queries could finally be developed. This approach gave finer intelligence of the attack to the system and made more structured queries find threats appropriately for Microsoft Sentinel.

III. Results and Findings

1) First scoring formula: The percentages of the LLMs accuracy are based on the success of KQL generation as provided through manual evaluation (the raw results can be found in the file named "query-results-2024-08-08"), categorized into four levels of correctness:

- Category 1 includes statements that are completely accurate and provide the right information, and it has a weight of 1.
- Category 2 includes estimates that are grammatically correct but only provide some of the right information. It has a weight of 0.70.
- Category 3 includes predictions that are grammatically right but do not provide any data. This has a weight of 0.3.
- Category 4 includes wrong predictions that are not only grammatically incorrect but also do not provide any true information. This category has a weight of 0.

To find the overall percentage for each model, add up the weighted counts for Categories 1, 2, 3, and 4 predictions. Then, divide that amount by the overall number of predictions for the model. This shows varying accuracy in predictions and makes sure that each model's performance is assessed based on its importance.

$$\text{Accuracy} = \left(\frac{W_1C_1 + W_2C_2 + W_3C_3 + W_4C_4}{\text{Total Number of Predictions Per Model}} \right) \times 100$$

Based on the previous scoring formula, Table II summarizes the accuracy percentages of the models.

TABLE II
Accuracy percentages using the first scoring formula

Model	Percentage
ChatGPT-3.5	55.71%
ChatGPT-4o-mini	51.79%
CodeStral (Mistral)	44.05%
Llama 3.1-70B	42.74%
ChatGPT-4o	42.62%
ChatGPT-4	39.64%
Llama 3.1-8B	39.17%
Qwen 2-72B	33.69%
Claude-3.5 Sonnet	25.00%
Qwen 2-110B	22.62%

2) Second scoring formula: Additionally, a simpler accuracy formula was also tested, which calculates the percentage of correct predictions using the following formula.

$$\text{Accuracy} = \left(\frac{\text{Number of Correct Predictions}}{\text{Total Number of Predictions Per Model}} \right) \times 100$$

The results using this simpler calculation method are described in Table III.

Both methods give a comprehensive view of model performance, with the weighted approach giving more granularity to evaluate partial correctness.

TABLE III
ACCURACY PERCENTAGES USING THE SECOND SCORING FORMULA

Model	Percentage
ChatGPT-3.5	54.76%
ChatGPT-4o-mini	50.00%
CodeStral (Mistral)	42.86%
ChatGPT-4	33.33%
Llama 3.1-70B	33.33%
Llama 3.1-8B	30.95%
ChatGPT-4o	28.57%
Qwen 2-72B	21.43%
Qwen 2-110B	19.05%
Claude-3.5 Sonnet	11.90%

3) The Next Step: A number of useful conclusions were drawn from the extensive testing: ChatGPT-4o-mini proved to be the best model for generating appropriate and relevant KQL queries. This model generated high quality questions for MITRE ATT&CK recognition, as it showed better performance in understanding challenging directions and examples given in the prompt. It also proved to be faster and cheaper than ChatGPT-3.5.

It was also found that the most accurate and precise searches were achieved by using six instances in the prompt. Multiple instances allowed the AI model to grasp a wide range of possibilities, improving its ability to generalise responses. In addition, different temperature settings on the studies yielded an optimal value of 0.5 as a temperature balance between producing deterministic, sharp results and enough flexibility to allow for variability in the protocols.

Once the best performing model was selected, the next step was to integrate the required APIs to ensure smooth operation. The system was integrated with the Microsoft Sentinel API for sample data collection and KQL searching, but up-to-date threat information was obtained from the MITRE ATT&CK TAXII server. To make the system more reliable, a backup solution has been developed using the local dataset in the absence of the TAXII server.

Other key features implemented include session management, feedback collection, dynamic query generation with explanations, on-demand log analysis and conversation management. In addition, these features allow the user to quickly identify and assess cyber threats, in addition to learning. The system is constantly being improved, with user comments incorporated to ensure flexibility and increased performance in real-world cybersecurity environments.

IV. THE LLM-BASED CHATBOT

This chapter describes the design and function of an LLM-based chatbot-a state-of-the-art AI system designed to improve cybersecurity by automating threat detection and analysis. It leverages LLMs for generating well-targeted KQL inquiries, thereby assisting security analysts in attack detection and response.

We start with a summary of the chatbot's main characteristics: query generation, real-time log analysis, and interaction with the cybersecurity framework. Next comes the system architecture-review of the chatbot architecture in general and all crucial parts included therein, such as user interface, AI model, threat intelligence, and SIEM systems like Microsoft Sentinel. Further comes the workflow, depicting steps involved from input by a user to execution of queries for finding out threats. Finally, the developed chatbot assessment is performed regarding the automation of cybersecurity operations and possible upgrades.

By the end of the chapter, we explain how the LLM-based chatbot represents an advanced cybersecurity solution connected with AI automation and practical threat detection.

1) Architecture: AI-driven MITRE ATT&CK Detection System amplifies the operation of cybersecurity by integrating Large Language Models, Microsoft Sentinel, and the servers of the MITRE ATT&CK TAXII to develop queries with appropriate Kusto Query Language for the detection of threats automatically. Key Highlights:

- **Integration of MITRE ATT&CK data**: Pulls the latest threat intelligence from the TAXII server and provides a way to have the data when offline.
- **Dynamic Query Generation**: To generate LLM-based KQL queries with MITRE ATT&CK techniques for high accuracy in detection.
- **Educational Insights**: Explain every line of generated KQL queries line by line for learning and threat analysis.
- **Real-Time Log Analysis**: Pulls the security logs from Microsoft Sentinel for real-time analysis in threat detection.
- **Feedback and Learning**: Capture feedback provided by users to re-train fine-tune the AI-generated questions for better results in the future.
- **Session Management**: Enables users to save, rename, and keep track of previous detection scenarios for organized investigations.
- **API Integration**: Provides seamless integration with external cybersecurity tools for smooth data exchange.
- **Customizable Inputs**: Supports predefined and user-defined log tables along with time ranges for flexible analysis.
- **Prompt Engineering**: Enhanced LLM prompting techniques ensure topmost quality and syntactically correct query generation.
- **User-Friendly Interface**: A web-based interface that allows users to interact with the system intuitively.

2) Workflow: The LLM-based MITRE ATT&CK Detection System has well-structured the process that enables users to better interact with an AI model for retrieving information regarding cybersecurity threats, thus creating proper KQL search for log analysis in Microsoft Sentinel. The work process is presented as a successive chain that supports good interaction among the parts of the system to effectively identify a cyber threat. The process starts when a person signs into the system using a website. The system has two main choices: talk to

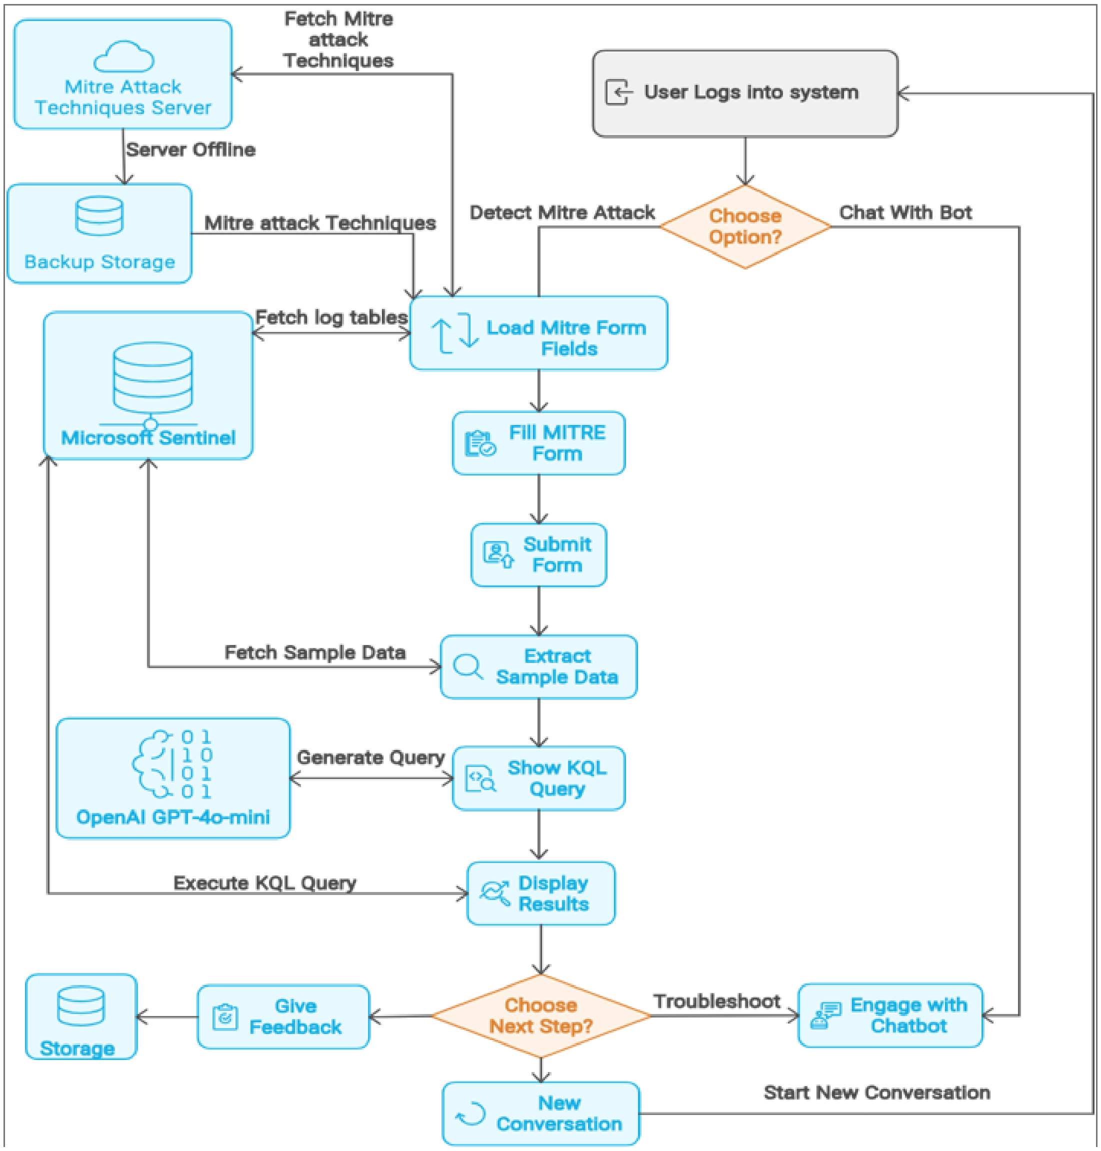

Fig. 1. Overall Architecture Approach

the AI, where users can ask the chatbot questions about cybersecurity, KQL queries, or specific MITRE ATT&CK methods, or identify a MITRE ATT&CK technique, where the system helps users follow a clear process to find certain computer threats.

If the user picks threat detection, the system initiates several actions to gather relevant data. It first gets the newest MITRE ATT&CK methods from the TAXII server to ensure the identification process uses the latest threat intelligence. If the TAXII server is not available, the system will use a local file so that it can keep working without interruption. Upon startup, the application communicates with Microsoft Sentinel to enumerate a list of available log tables. These contain event data that will be analyzed against known patterns indicative of an attack. Once identified, the log sources are summarized, and a form is presented to the end-user to fill out, specifying one of the detected MITRE ATT&CK methods to analyze with the appropriate log table and timestamp to investigate within.

Once submitted, it pulls the sample log data from Microsoft Sentinel, adding context to the AI-generated KQL query. With that, the AI model creates a detailed prompt including the selected MITRE ATT&CK method, sample log data, and syntax rules for the generation of KQL queries. The latter then feeds into an AI model, while the response comes as a KQL query that reflects the chosen threat, crafted by ChatGPT-4o-mini or other LLMs. After generation, the query will be presented to the end user at the web interface, line by line, to understand its structure and purpose.

After reviewing the query, the system immediately executes it in Microsoft Sentinel to retrieve log analysis results. These results indicate whether there is evidence of the selected MITRE ATT&CK method in the user's logs. At this point, the user has two main options: they can continue interacting with the AI chatbot to refine the query, adjust parameters, or troubleshoot issues, or they can start a new detection process by selecting a different MITRE ATT&CK method, log table, or time period and repeating the workflow.

The system is also designed with a feedback mechanism that allows users to rate the accuracy and relevance of the AI-generated KQL queries. This feedback will be stored and applied to constant improvements in the generation of future AI queries so that the system is effective and adaptive for cybersecurity professionals. This architecture is demonstrated in figure 1.

V. Conclusions

This paper presented an AI-powered MITRE ATT&CK detection system, illustrating the feasibility of employing Large Language Models (LLMs) to automate cyber threat identification. We enhanced AI-generated KQL search accuracy and cybersecurity efficiency using advanced prompt engineering and Chain of Thought (CoT) prompting. The system seamlessly integrates Microsoft Sentinel, MITRE ATT&CK TAXII servers, and OpenAI's API, providing a robust automation framework for query generation, log analysis, and threat intelligence processing. The novel feature of this work is the systematic improvement of LLM-generated queries for cybersecurity. Unlike prior research that focused on static analysis or security policy formulation, our system generates, performs, and explains KQL queries for real-time threat detection. Our self-improving AI assistant for security analysts detects threats and evolves with user feedback and session management. ChatGPT-4o-mini was the best model after testing for accuracy, efficiency, and cost. Our scalable, adaptive technique works in several SIEM settings and attack surfaces. This work indicates that LLMs may automate complex security procedures, saving manual effort and boosting detection precision when fine-tuned with cybersecurity-specific limits.

A. Future Work

Our solution represents a quantum leap in AI-driven cybersecurity automation; there are still areas for further improvement. Refining AI Models for Cybersecurity Tasks: Training domain-specific LLMs with security datasets may result in massive improvements to generate more accurate, attack-aware KQL searches. Mechanisms for Continuous Learning: With the real-time adaptive learning, this model will get better with new attack strategies, user feedback, and security intelligence updates to make the long-term efficiency achievable. Improved Threat Intelligence Expansion: Future updates may include increasing the threat intelligence feeds, adding dark web monitoring, industry-specific security databases, and anomaly detection algorithms to attain better detection accuracy in case of an attack. Metaverse [11]Security and Virtual Threat Detection: The growth of digital ecosystems means the dire need to protect virtual space, blockchain-based assets, and decentralized apps. Future work could include AI-driven security monitoring for metaverse environments, LLM-assisted cyber forensics, identity verification, and virtual attack simulations. Blockchain for Cybersecurity and Threat Intelligence Sharing: Blockchain [12] can be used to enhance data integrity, threat intelligence distribution, and tamper-proof security logs. Future research might be required on smart contract-driven security automation that will enable decentralized verification of security alarms to avoid manipulation of data and fabrication of SIEM logs.

References

[1] Y. Chen, M. Cui, D. Wang, Y. Cao, P. Yang, B. Jiang, Z. Lu, and B. Liu, "A survey of large language models for cyber threat detection," *Computers A Security*, vol. 145, p. 104016, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167404824003213

[2] M. Charfeddine, H. M. Kammoun, B. Hamdaoui, and M. Guizani, "Chatgpt's security risks and benefits: Offensive and defensive use-cases, mitigation measures, and future implications," *IEEE Access*, vol. 12, pp. 30 263–30 310, 2024.

[3] A. Georgiadou, S. Mouzakitis, and D. Askounis, "Assessing mitre attack risk using a cyber-security culture framework," *Sensors*, vol. 21, no. 9, 2021. [Online]. Available: https://www.mdpi.com/1424-8220/21/9/3267

[4] D. Koutras, P. Kotzanikolaou, E. Paklatzis, C. Grigoriadis, and C. Douligeris, "A framework for automating environmental vulnerability analysis of network services," *ITU Journal on Future and Evolving Technologies*, vol. 5, 2024.

[5] Y.-T. Huang, R. Vaitheeshwari, M.-C. Chen, Y.-D. Lin, R.-H. Hwang, P.-C. Lin, Y.-C. Lai, E. H.-K. Wu, C.-H. Chen, Z.-J. Liao, and C.-K. Chen, "Mitretrieval: Retrieving mitre techniques from unstructured threat reports by fusion of deep learning and ontology," *IEEE Transactions on Network and Service Management*, vol. 21, no. 4, pp. 4871–4887, 2024.

[6] Y. Cao, N. Gu, X. Shen, D. Yang, and X. Zhang, "Defending large language models against jailbreak attacks through chain of thought prompting," in *2024 International Conference on Networking and Network Applications (NaNA)*, 2024, pp. 125–130.

[7] H. Szmurlo and Z. Akhtar, "Digital sentinels and antagonists: The dual nature of chatbots in cybersecurity," *Information*, vol. 15, no. 8, 2024. [Online]. Available: https://www.mdpi.com/2078-2489/15/8/443

[8] F. C. Fernández, R. L. S. Garcia, and W. Caarls, "Comparison of llm models and strategies for structured query construction from natural language queries," in *2024 IEEE Latin American Conference on Computational Intelligence (LA-CCI)*, 2024, pp. 1–6.

[9] S. M. M. Hossain, R. Couturier, J. Rusk, and K. B. Kent, "Automatic event categorizer for siem," in *Proceedings of the 31st Annual International Conference on Computer Science and Software Engineering*, ser. CASCON '21. USA: IBM Corp., 2021, p. 104–112.

[10] W. Kasri, Y. Himeur, H. A. Alkhazaleh, S. Tarapiah, S. Atalla, W. Mansoor, and H. Al-Ahmad, "From vulnerability to defense: The role of large language models in enhancing cybersecurity," *Computation*, vol. 13, no. 2, 2025. [Online]. Available: https://www.mdpi.com/2079-3197/13/2/30

[11] V. Malamas, D. Koutras, T. K. Dasaklis, V. Vassilakopouls, and P. Kotzanikolaou, "Blockchain revolution in the metaverse: Challenges, applications and future directions," in *2024 International Conference on Artificial Intelligence, Metaverse and Cybersecurity (ICAMAC)*. IEEE, 2024, pp. 1–6.

[12] V. Malamas, D. Koutras, and P. Kotzanikolaou, "Uninterrupted trust: Continuous authentication in blockchain-enhanced supply chains," in *2023 8th South-East Europe Design Automation, Computer Engineering, Computer Networks and Social Media Conference (SEEDA-CECNSM)*. IEEE, 2023, pp. 1–6.

This work has been partially funded by the 'Harmonizing People, Processes, and Technology for Robust Cybersecurity' (CYberSynchrony) project, which has received funding from the European Union's Digital Europe Programme (DEP) under grant agreement No. 101158555.