ICT9099CS01 - Mitigating Non-Determinism in LLM-Assisted Incident Response

Principal Investigator

Dr. Mohammad Siddiqui	
Polytechnic ID: 12011115	**Department/School**: ICT

Programme	ICT9099 Master of Science by Research in Computer Science
Research Area and Specialization:	Cybersecurity, Incident Response, Artificial Intelligence for Cybersecurity, Large Language Models (LLMs), Knowledge Graphs, Security Analytics.
Company Endorsement	Beyon Cyber
Industry mentor	Mr. Britto Irudhayasamy

Background of the problem

Large Language Models (LLMs) are increasingly employed in Security Operations Centers (SOCs) for alert analysis, incident investigation, and threat intelligence. Nevertheless, their probabilistic nature may result in inconsistent outputs, limited explainability, and diminished operational trust (Ismail et al., 2025; Loumachi et al., 2025).

Recent research has investigated Retrieval-Augmented Generation (RAG) and knowledge-based methods to enhance reliability. Ismail et al. (2025) demonstrated the application of a RAG-driven SOC copilot for security event response. Loumachi et al. (2025) found that integrating structured evidence retrieval with LLMs improves cyber incident investigations. Gao and Chang (2025) validated the effectiveness of knowledge-graph-based RAG for cybersecurity vulnerability analysis.

Despite these advancements, many SOCs still depend on alert titles or keyword matching when reusing investigation procedures, frequently neglecting critical contextual information. This project aims to address this gap by integrating cybersecurity knowledge graphs, contextual metadata, deterministic retrieval, and controlled LLM assistance to enhance the consistency, explainability, and effectiveness of incident investigations.

References
Gao, X., & Chang, X. (2025). https://doi.org/10.1109/ACCESS.2025.3627277

Ismail, R. K., Widyatama, F., Wibawa, I. M., Brata, Z. A., Ukasyah, Nelistiani, G. A., & Kim, H. (2025). https://doi.org/10.3390/s25030870

Loumachi, F. Y., Ghanem, M. C., & Ferrag, M. A. (2025). https://doi.org/10.3390/computers14020067.

List of publications related to the field

- Siddiqui, M. A., Al Masud, S. M. R., & Kamal, M. B. (2016). Towards Improving the Quality of Present MAC Protocols for LECIM Systems. International Journal of Advanced Computer Science and Applications (IJACSA), 7(6), 280–288. DOI: 10.14569/IJACSA.2016.070636. Available at: http://dx.doi.org/10.14569/IJACSA.2016.070636
  • Siddiqui, M. A., Ahmad, Q. S., & Khan, M. H. (2010). Link State based High Throughput Routing (HTR) Protocol for Wireless Mesh Networks. Egyptian Computer Science Journal (ECS Journal), 34(4), 50–59, May 2010. Available at: http://ecsjournal.org/JournalArticle.aspx?articleID=400
  • Siddiqui, M. A., Ahmad, Q. S., Khan, M. H., & Khan, R. A. (2010). Patient Health Monitoring using Wireless Body Area Networks. Egyptian Computer Science Journal (ECS Journal), 34(4), 75–81, May 2010. Available at: http://ecsjournal.org/JournalArticle.aspx?articleID=402
  • Ullah, S., Higgins, H., Siddiqui, M. A., & Kwak, K. S. (2008). A Study of Implanted and Wearable Body Sensor Networks. In Proceedings of the 2nd KES International Symposium on Agent and Multi-Agent Systems: Technologies and Applications (KES-AMSTA 2008), Lecture Notes in Artificial Intelligence (LNAI), Vol. 4953, Springer, pp. 464–473. Published 2008; Springer edition 2009. DOI: 10.1007/978-3-540-78582-8_47. Available at: https://link.springer.com/chapter/10.1007/978-3-540-78582-8_47

Scholarly articles related to the project:
• Ismail, R. K., Widyatama, F., Wibawa, I. M., Brata, Z. A., Ukasyah, Nelistiani, G. A., & Kim, H. (2025). Enhancing Security Operations Center: Wazuh Security Event Response with Retrieval-Augmented-Generation-Driven Copilot. Sensors, 25(3), 870. DOI: 10.3390/s25030870.
• Loumachi, F. Y., Ghanem, M. C., & Ferrag, M. A. (2025). Advancing Cyber Incident Timeline Analysis Through Retrieval-Augmented Generation and Large Language Models. Computers, 14(2), 67. DOI: 10.3390/computers14020067.
• Gao, X., & Chang, X. (2025). A Retrieval-Augmented Generation Framework Based on a Knowledge Graph of Cybersecurity Vulnerabilities in Power Networks. IEEE Access, 13, 186693–186710. IEEE. DOI: 10.1109/ACCESS.2025.3627277.

Research Objectives and Expected Outcomes:

Aim

The aim is to develop a context-aware framework that reduces non-determinism in large language model (LLM)-assisted incident response by employing deterministic retrieval, cybersecurity knowledge graphs, and structured contextual reasoning.

Objectives:

• Analyze operational alert datasets, associated metadata, and investigation histories.
•Develop a hierarchical cybersecurity knowledge graph that represents domains, scenarios, objectives, and investigation procedures.
• Design a context-aware metadata framework to support alert classification and routing.
• Develop a deterministic retrieval engine that identifies relevant investigation paths using contextual similarity.
• Integrate a locally hosted large language model (LLM) to support investigation activities when deterministic rules are insufficient.
• Evaluate the proposed framework using operational security operations center (SOC) datasets and predefined performance metrics..

Expected Outcomes

- Development of a context-aware metadata framework to facilitate alert classification and routing.
- Construction of a hierarchical cybersecurity knowledge graph that represents domains, scenarios, objectives, and investigation procedures.
- Creation of a reusable database of logic templates derived from historical investigations.
- Implementation of a deterministic retrieval engine designed to identify appropriate investigation paths.
- Deployment of a locally hosted prototype, assisted by a large language model (LLM), to support incident response.
- Provision of quantitative evidence demonstrating enhanced investigation consistency and reduced analyst effort.
- Completion of a Master's dissertation and submission of at least one peer-reviewed publication.

Research Design/Methodology

An Applied Research approach using Design Science Research (DSR) will be adopted to design, develop, and evaluate a framework that combines cybersecurity knowledge graphs, contextual metadata, deterministic retrieval, and controlled LLM assistance.

Data Collection: Historical alerts, metadata, investigation records, and expert input will be analyzed to identify investigation patterns and workflows.

Framework Development: A cybersecurity knowledge graph, context-aware metadata model, and deterministic retrieval engine will be developed to select appropriate investigation procedures. A locally hosted LLM will provide support when deterministic methods are insufficient.

Evaluation: The framework will be tested using operational datasets and compared with traditional title-based approaches using accuracy, precision, recall, F1-score, investigation consistency, time reduction, and analyst satisfaction.

Ethics: All data will be anonymized and handled in accordance with research ethics and confidentiality requirements.

Significance and potential impact of research in Bahrain

The project aims to address the challenge of achieving reliable and explainable AI-assisted investigations in Security Operations Centers (SOCs). The proposed framework leverages contextual information to improve investigation consistency, reduce analyst workload, enhance knowledge reuse, and support dependable AI-driven incident response. These advancements are intended to contribute to Bahrain's cybersecurity and digital transformation objectives.

Alignment with Bahrain's national research strategy

Contribution to Bahrain Society: Enhances Bahrain's cybersecurity resilience and incident response capabilities.

Alignment with Industry: Improves SOC investigation quality, efficiency, and analyst productivity through AI-driven solutions.

National Priorities: Supports Bahrain's digital transformation, cybersecurity, AI, and knowledge economy goals.

Benefits for Industry: Improves investigation consistency, knowledge reuse, auditability, and reduces analyst workload through a deployable AI-assisted framework.