

Midocean University

**Bridging the Gap Between Traditional Detection and Intelligent
Response in Banking Cybersecurity**

Student Name: Abdullah Al-Faifi

Course: Research Methodology in Computer Science (AI800)

Instructor: Dr. Mohammed Anbar

Contents

Abstract.....	3
Chapter One:.....	4
1.1 Introduction	4
1.2 Background	4
1.3 Problem Statement	5
1.4 Research Objectives	5
1.5 Scope of the Study	5
Chapter Two:.....	6
2.1 Literature Review	6
2.2 Critical Summary of the Literature Review	7
Chapter Three:	7
3.1 Research Methodology	7
3.2 Research Design	7
Phase 1 – Systematic Literature Review:	7
Phase 2 – Semi-structured Interviews:	8
Phase 3 – Quantitative Survey:	8
Phase 4 – Framework Development and Validation:.....	8
3.3 Study Population and Sampling	8
3.4 Data Collection Instruments	9
3.5 Data Analysis Procedures.....	10
3.6 Proposed Framework.....	10
3.7 Ethical Considerations and Study Limitations.....	11
Expected Outcomes	12
Gantt Chart.....	12
References	13

Abstract

This research proposal aims to develop an intelligent, machine learning–based framework to enhance cyber incident detection and response within Security Operations Centers (SOCs) in the banking sector. Despite ongoing advancements in cybersecurity technologies, many SOC environments continue to rely heavily on rule-based systems and static signatures, limiting their ability to detect evolving and sophisticated threats while generating high volumes of false positives and alert fatigue.

To address these challenges, the study adopts a mixed-methods approach that combines a systematic literature review, semi-structured interviews with cybersecurity professionals, and a quantitative survey targeting SOC practitioners. The research seeks to design a practical and adaptable framework that integrates anomaly detection, alert prioritization, and decision support mechanisms to improve detection accuracy and operational efficiency.

The proposed framework also considers critical factors such as data privacy, regulatory compliance, and model explainability—key requirements in the highly sensitive banking environment. In addition, the study draws on recent advancements in machine learning, including techniques such as federated learning, transfer learning, and few-shot learning, to enhance adaptability and scalability across diverse data sources.

The expected contributions of this research are both theoretical and practical. It aims to provide a comprehensive and implementable framework that reduces false positives, accelerates incident response, and improves the overall effectiveness of SOC teams. Furthermore, the study will offer an actionable roadmap aligned with current regulatory standards and cybersecurity best practices in the banking sector.

Chapter One:

1.1 Introduction

Banking cybersecurity has become one of the most critical domains requiring advanced research and practical solutions due to the rapid growth in security data volumes, the increasing diversity of cyberattacks, and the continuous evolution of adversarial techniques. Although financial institutions rely on Security Operations Centers (SOCs) to monitor events, analyze logs, and detect incidents, these centers continue to face significant operational challenges, including alert fatigue, difficulty distinguishing between benign and malicious activities, and delayed response to complex, multi-stage attacks.

In this context, the integration of machine learning techniques into NOC/SOC environments represents a crucial research direction. Such integration enables a transition from traditional rule-based detection approaches toward intelligent, data-driven mechanisms based on pattern recognition, anomaly detection, and predictive analytics.

This study builds on established research principles by identifying existing gaps in current SOC practices, defining clear and measurable objectives, and adopting a structured research design that aligns data collection and analysis with the research problem. The aim is to develop more adaptive and efficient detection capabilities tailored to the specific needs of banking cybersecurity operations.

1.2 Background

This proposal intersects two complementary research streams. The first focuses on the application of machine learning techniques for intrusion detection in domains such as the medical Internet of Things (IoT) and other critical infrastructures. Prior studies have demonstrated the effectiveness of approaches such as fog computing, which improves detection speed by processing data at the network edge (Sharifi & Goli-Bidgoli, 2022), and feature selection techniques that enhance model accuracy while reducing computational overhead (Dey et al., 2022). Federated learning has also been widely explored for its ability to preserve data privacy through decentralized training (Yang et al., 2023), although challenges related to model synchronization and heterogeneous environments remain. In addition, few-shot and transfer learning approaches have shown promise in detecting novel attacks with limited labeled data, but their generalizability across domains requires careful adaptation (Ma et al., 2023; Xu et al., 2025).

The second research stream is situated within the banking sector, where institutions must process large-scale, real-time security data under strict regulatory and operational constraints. Existing studies highlight the benefits of transfer learning in reducing training costs (Vu et al., 2020) and emphasize the growing importance of explainable artificial intelligence (XAI) to ensure transparency, trust, and regulatory compliance (Reddy, 2022). However, rapid advancements in AI have also exposed the need for more robust evaluation frameworks in safety-critical environments (Zhang & Lu, 2021).

Despite these advancements, the literature indicates a clear gap: the lack of integrated and operationally viable frameworks that adapt advanced machine learning techniques to the specific requirements of banking SOC environments. In particular, challenges related to privacy, explainability, real-time processing, and seamless operational integration remain insufficiently addressed. This research aims to bridge this gap by proposing a comprehensive framework that combines intelligent detection with practical deployment considerations tailored to the banking sector.

1.3 Problem Statement

Many banking Security Operations Centers (SOCs) continue to rely on rule-based and signature-driven systems. While these approaches are effective in detecting known attack patterns, they are inherently limited in identifying emerging threats, anomalous behavior, and complex multi-stage attacks. Consequently, organizations face high false-positive rates, delayed incident response, analyst fatigue, and an increased risk of overlooking critical security events.

This challenge is particularly severe in the banking sector due to the sensitivity of financial data, strict regulatory requirements, and the significant financial and reputational impact of detection failures. At the same time, the integration of machine learning into SOC environments introduces additional challenges, including model explainability, trust in automated decisions, and compatibility with existing operational workflows.

Therefore, the core research problem lies in the lack of a comprehensive and practical framework that effectively integrates intelligent anomaly detection, alert prioritization, and decision support mechanisms, while ensuring privacy, regulatory compliance, and operational feasibility within banking SOC environments.

1.4 Research Objectives

- To analyze the limitations of traditional rule-based and signature-based detection methods within banking Security Operations Centers (SOCs), particularly in terms of false-positive rates and incident response delays.
- To identify and assess the most relevant features and behavioral patterns for developing machine learning–based anomaly detection mechanisms across network, host, and user activity data.
- To design and develop an intelligent framework that integrates data collection, preprocessing, alert prioritization, and initial incident response processes.
- To evaluate the applicability and effectiveness of advanced machine learning techniques—such as feature selection, federated learning, transfer learning, and few-shot learning—within the banking cybersecurity context.
- To develop a practical implementation roadmap for the proposed framework, ensuring alignment with regulatory requirements, cybersecurity governance standards, and seamless integration with existing SOC systems.

1.5 Scope of the Study

This research focuses on Security Operations Centers (SOCs) within the banking sector, with particular emphasis on monitoring and analysis environments that ingest diverse data sources, including logs from network devices, servers, endpoints, and security platforms. The study aims to design a research-oriented and practical framework rather than develop a fully functional commercial product or deploy a production-ready system.

The scope of the research is limited to the application of machine learning techniques in the stages of cyber incident detection, alert prioritization, and decision support. It does not extend to digital forensics, incident recovery, or post-incident management processes.

In terms of evaluation, the proposed framework will be assessed through a combination of literature analysis, expert feedback, and comparison with existing approaches in professional practice. To ensure compliance with privacy and regulatory constraints, the study does not utilize real or sensitive banking data; instead, it relies on simulated scenarios, documented use cases, and practitioner insights.

These limitations are acknowledged as necessary constraints to maintain feasibility, ethical compliance, and relevance within the context of academic research.

Chapter Two:

2.1 Literature Review

Recent studies indicate that machine learning has become a central component in enhancing intrusion detection capabilities within sensitive and high-demand environments. For example, fog computing approaches enable faster detection by processing data closer to its source, thereby reducing response latency (Sharifi & Goli-Bidgoli, 2022). Feature selection techniques have also been shown to improve model accuracy while reducing computational overhead, particularly in environments with heterogeneous and high-volume data (Dey et al., 2022).

Federated learning offers a promising solution for privacy-preserving model training by enabling decentralized data processing, which is highly relevant for the banking sector; however, challenges related to model synchronization and environmental heterogeneity remain significant (Yang et al., 2023). Similarly, few-shot and transfer learning approaches address the scarcity of labeled data and reduce training costs, although they require careful adaptation to ensure cross-domain generalization and stable performance (Ma et al., 2023; Xu et al., 2025; Vu et al., 2020). In addition, hyperparameter optimization has been shown to improve the balance between detection accuracy and processing speed, albeit with increased computational demands (Türkoğlu et al., 2022).

Beyond performance improvements, recent research emphasizes the importance of explainability and trust in machine learning models. The lack of interpretability can limit adoption in critical domains such as banking, where transparency and accountability are essential (Reddy, 2022). Furthermore, the need for robust evaluation frameworks that account for operational complexity, regulatory constraints, and ethical considerations has been widely highlighted (Zhang & Lu, 2021).

Taken together, these studies reveal a clear shift toward more adaptive, privacy-aware, and efficient detection mechanisms. However, they also expose a significant research gap: existing approaches tend to address isolated technical challenges—such as accuracy, speed, or privacy—without providing an integrated and operationally viable framework. This limitation is particularly critical in banking SOC environments, where effective solutions must simultaneously address technical performance, regulatory compliance, and real-world operational requirements. Accordingly, this research aims to bridge this gap by proposing a comprehensive framework that unifies these dimensions into a practical and implementable solution.

2.2 Critical Summary of the Literature Review

The reviewed studies consistently demonstrate the effectiveness of machine learning in improving intrusion detection capabilities. However, they vary significantly in terms of application contexts, data sources, and evaluation methodologies, which limits the generalizability of their findings across domains.

Despite strong advancements in detection accuracy, key operational aspects—such as data privacy, model explainability, and integration with existing SOC workflows—remain insufficiently addressed. These factors are particularly critical in highly regulated environments such as the banking sector.

Furthermore, while research in sensitive domains such as the Internet of Medical Things (IoMT) provides valuable insights, direct transfer of these approaches to banking environments is not straightforward. Differences in regulatory requirements, risk profiles, and data governance constraints necessitate careful adaptation.

Overall, the literature reveals a significant gap: the absence of a comprehensive, end-to-end operational framework that effectively integrates intelligent detection mechanisms with alert prioritization and incident response within banking SOC environments.

Chapter Three:

3.1 Research Methodology

This study adopts a mixed-methods research design consisting of four main phases: a systematic literature review, semi-structured interviews with cybersecurity professionals, a quantitative survey of SOC practitioners, and the development and validation of the proposed framework.

The use of a mixed-methods approach is justified by the complex and multi-dimensional nature of the research problem. It enables an in-depth exploration of operational challenges through qualitative insights, while also providing quantitative evidence to support the generalization of findings across practitioner perspectives.

This methodology ensures a comprehensive understanding of both technical and operational aspects of SOC environments and supports the development of a practical, evidence-based framework aligned with real-world cybersecurity practices.

3.2 Research Design

The research is structured into four interconnected phases, ensuring a comprehensive and systematic investigation of the research problem.

Phase 1 – Systematic Literature Review:

A comprehensive review will be conducted following PRISMA guidelines to identify relevant peer-reviewed articles, conference papers, and industry reports published between 2015 and 2025. Inclusion criteria will focus on studies related to intrusion detection, anomaly detection, and security incident response in critical infrastructures. Studies lacking empirical validation will be excluded. A structured review matrix will be developed to summarize each study's objectives, methodologies, key findings, and limitations.

Phase 2 – Semi-structured Interviews:

A purposive sample of 10–12 banking SOC professionals, including analysts, engineers, and managers, will be recruited through professional networks and industry groups. Participants will have a minimum of three years of experience in SOC operations. The interviews will explore operational challenges, perceptions of machine learning techniques, and requirements for an integrated framework. Interview questions will be pilot-tested with domain experts to ensure clarity and relevance. All interviews will be conducted with informed consent, recorded, and transcribed for thematic analysis.

Phase 3 – Quantitative Survey:

Based on insights from the interview phase, a structured questionnaire will be developed using a five-point Likert scale to evaluate the perceived importance and feasibility of the proposed framework components. The survey instrument will be pilot-tested to ensure reliability, and internal consistency will be assessed using Cronbach's alpha. A stratified sampling approach will target 50–80 SOC practitioners across multiple banking institutions, ensuring representation across roles and regions. Responses will be collected through an online platform and anonymized.

Phase 4 – Framework Development and Validation:

Findings from the literature review, interviews, and survey will be triangulated to design the proposed framework. The framework will then be evaluated by a panel of domain experts to assess its feasibility, relevance, and alignment with regulatory and operational requirements. Feedback from this evaluation will be used to refine the framework and develop a practical implementation roadmap.

3.3 Study Population and Sampling

The target population of this study consists of cybersecurity analysts and engineers, SOC team leaders, and governance and compliance specialists within the banking sector.

For the qualitative phase, a purposive sampling strategy will be employed to recruit 10–12 participants with a minimum of three years of professional experience in SOC or banking cybersecurity environments. This approach ensures the inclusion of knowledgeable participants capable of providing in-depth insights into operational challenges and requirements.

For the quantitative phase, a stratified convenience sampling approach will be used to recruit approximately 50–80 practitioners through professional associations, industry forums, and cybersecurity networks. Efforts will be made to ensure representation across different organizational sizes, roles, and geographic regions to enhance the generalizability of the findings.

Participation in both phases will be voluntary, and informed consent will be obtained from all participants in accordance with ethical research standards.

3.4 Data Collection Instruments

The study employs a set of structured research instruments and analytical tools to ensure methodological rigor and reliability.

Literature Review Protocol:

A formal review protocol will be developed to define search strategies, including databases such as IEEE Xplore, ACM Digital Library, and Scopus. It will specify inclusion and exclusion criteria, as well as quality assessment metrics to evaluate the relevance and rigor of selected studies. The protocol will be reviewed to ensure consistency and transparency.

Interview Guide:

A semi-structured interview guide will be designed to explore key themes, including alert fatigue, detection of unknown threats, explainability requirements, regulatory compliance, and system integration challenges. The guide will be validated by domain experts and pilot-tested to refine question clarity and structure.

Survey Questionnaire:

The survey instrument will be developed based on insights from the interview phase. It will include items measured using a five-point Likert scale to assess perceptions of key framework components, such as data aggregation, feature selection, anomaly detection, interpretability, and governance. Demographic questions will also be included. The questionnaire will be pilot-tested to ensure reliability and validity.

Analytical Tools and Techniques:

Qualitative data from interviews will be analyzed using thematic coding supported by NVivo software. Quantitative data will be analyzed using SPSS or Python libraries (e.g., pandas, scipy). Reliability will be assessed using Cronbach's alpha, while descriptive statistics and exploratory factor analysis will be employed to identify patterns and relationships within the data.

3.5 Data Analysis Procedures

The study employs a combination of qualitative and quantitative data analysis techniques to ensure comprehensive and reliable findings.

Thematic Analysis:

Interview transcripts will be analyzed using a structured thematic analysis approach involving open, axial, and selective coding. Two independent coders will analyze the data to ensure objectivity, and discrepancies will be resolved through discussion. Inter-coder reliability will be assessed using Cohen's kappa. The resulting themes will be organized into a conceptual model representing key operational challenges and system requirements.

Descriptive and Inferential Statistics:

Survey data will be analyzed using descriptive statistics, including means, standard deviations, and ranking of framework components. The internal consistency of survey scales will be evaluated using Cronbach's alpha. Where appropriate, inferential statistical tests—such as Mann–Whitney U and Kruskal–Wallis—will be applied to examine differences between participant groups.

Triangulation:

Findings from the literature review, interviews, and survey will be systematically integrated using triangulation. This process will identify converging evidence, highlight discrepancies, and strengthen the validity and reliability of the study. The triangulated results will directly inform the design and refinement of the proposed framework.

3.6 Proposed Framework

The proposed framework is structured into five interconnected layers, each addressing a critical component of cybersecurity operations within banking SOC environments.

1. Data Aggregation Layer:

This layer is responsible for collecting and consolidating security data from multiple sources, including network devices, servers, endpoints, and third-party security platforms. It ensures data availability and consistency for downstream processing.

2. Preprocessing and Feature Selection Layer:

In this layer, raw data is cleansed, normalized, and transformed into structured formats suitable for analysis. Advanced feature selection techniques, including heuristic and meta-heuristic approaches, are applied to enhance model performance and reduce computational complexity.

3. Intelligent Detection Layer:

This layer integrates machine learning techniques for anomaly detection and classification. Methods such as autoencoders and Isolation Forest are employed to identify abnormal patterns, while alert prioritization mechanisms help reduce noise. Privacy-preserving approaches, such as federated learning or transfer learning, may be utilized to train models without exposing sensitive data.

4. Interpretation and Operational Layer:

To support decision-making, this layer provides explainable outputs, including feature importance scores and confidence levels. It also generates actionable recommendations to assist analysts in understanding alerts and initiating appropriate response actions.

5. Governance and Compliance Layer:

This layer ensures that all processes adhere to regulatory and organizational requirements. It manages data privacy, access control, logging, and auditability by enforcing policies related to data retention, security governance, and documentation of decision-making processes.

3.7 Ethical Considerations and Study Limitations

This study will adhere to established ethical research standards and will obtain approval from the relevant Institutional Review Board (IRB) prior to data collection. Informed consent will be obtained from all participants, who will be clearly informed of the study's purpose, their rights, and their ability to withdraw at any time without consequence.

To ensure confidentiality and data protection, no sensitive banking data will be collected or accessed. All interview and survey responses will be anonymized and securely stored, with access restricted to the research team. Only aggregated findings will be reported to prevent the identification of individual participants or organizations.

Despite these safeguards, the study acknowledges several limitations. The relatively small qualitative sample size may limit the generalizability of findings. Additionally, the use of convenience and voluntary sampling may introduce self-selection bias, while reliance on self-reported data may be subject to recall bias. Furthermore, the proposed framework will be evaluated primarily through expert feedback and analytical validation rather than full-scale implementation in a live banking environment, which may limit the assessment of real-world performance.

Expected Outcomes

This research is expected to make both theoretical and practical contributions to the field of cybersecurity in banking environments.

- It provides a clear scientific characterization of the gap between traditional rule-based detection methods and intelligent, machine learning–driven response mechanisms within banking Security Operations Centers (SOCs).
- It identifies the key technical and organizational requirements necessary for the effective and secure adoption of machine learning, including considerations related to privacy, explainability, and regulatory compliance.
- It proposes a comprehensive and integrated framework that balances detection accuracy, operational efficiency, explainability, data privacy, and compliance with industry standards.
- It offers practical recommendations to support banking institutions in improving detection capabilities, reducing false positives, and enhancing decision-making processes within SOC environments.
- It establishes a foundation for future research by enabling subsequent empirical studies focused on implementation, validation, and performance evaluation of the proposed framework in real-world settings.

Gantt Chart

Task	May 2026	June 2026	July 2026	August 2026	September 2026	October 2026
Literature Review and Source Collection						
Research Tools Development and Validation						
Conducting Qualitative Interviews						
Survey Design and Quantitative Data Collection						
Data Analysis and Framework Development						
Final Writing and Revision						

References

- Sharifi, M., & Goli-Bidgoli, A. (2022). FogLearn++: A fog-based machine learning framework for intrusion detection in IoT environments. *Computers & Electrical Engineering*, 103, 108374. <https://doi.org/10.1016/j.compeleceng.2022.108374>
- Türkoğlu, M., Hanbay, D., & Sengur, A. (2022). An effective approach for intrusion detection using hyperparameter optimisation in machine learning models. *Applied Soft Computing*, 113, 107866. <https://doi.org/10.1016/j.asoc.2021.107866>
- Xu, C., Zhan, Y., Chen, G., Wang, Z., Liu, S., & Hu, W. (2025). Elevated few-shot network intrusion detection via self-attention mechanisms and iterative refinement. *PLOS ONE*, 20 (1), e0317713. <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0317713>
- Xu, C., Li, D., Liu, Z., Yang, J., Shen, Q., & Tong, N. (2025). Few-shot network intrusion detection method based on multi-domain fusion and cross-attention. *PLOS ONE*, 20 (7), e0327161. <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0327161>
- Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2023). Federated learning: Privacy-preserving collaborative machine learning. *ACM Computing Surveys*, 55 (1), Article 4. <https://doi.org/10.1145/3625558>
- Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial Information Integration*, 23, 100224. <https://doi.org/10.1016/j.jii.2021.100224>