Multiverse Journal
E-ISSN: 3080-3942, P-ISSN: 3080-3934
Volume 1, Issue 1, June 2024 | PP. 01-21

Research Paper

AI-Powered Cyber Threat Intelligence: Real-Time Detection, Prediction, and Response through Machine Learning Research

Chidera Jenaluiz Umechukwu [1]

[1] Wilmington University, USA.

Received: 13 March, 2024 **Accepted:** 05 June, 2024 **Published:** 30 June, 2024

Abstract

The rise in sophisticated cyberattacks has revealed the inadequacy of traditional security measures such as firewalls and signature- based intrusion detection systems, particularly against novel and zero-day threats. This study presents an AI-powered cyber threat intelligence (CTI) framework that integrates Natural Language Processing (NLP), machine learning (ML), and malware analysis to enable real-time detection, prediction, and automated mitigation of threats. The proposed framework follows a four-phase architecture: data collection, data analysis, risk assessment, and mitigation, with a continuous feedback loop for adaptive learning. Experiments were conducted using LUFlow and IEC 60870–5-104 intrusion detection datasets, yielding a detection accuracy of approximately 99.96% and significant improvements in precision, recall, and F1-score compared with baseline models such as SecurityBERT, XAI, and DNN-based solutions. The findings demonstrate the framework's scalability, generalizability, and suitability for deployment in enterprise and critical infrastructure networks. Future research will focus on integrating semi- supervised learning, adversarial robustness, and edge-based deployments for enhanced threat response.

Keywords: Cyber Threat Intelligence, Machine Learning, Natural Language Processing, Automated Mitigation, Real-Time Detection

Introduction

The rapid growth of interconnected systems has resulted in an equally rapid expansion of the attack surface for malicious actors. In recent years, critical vulnerabilities have emerged that enable attackers to compromise software integrity and gain unauthorized access to sensitive data. One prominent example is the Follina (CVE-2022-30190) vulnerability, which enables remote code execution and poses a threat to the reliability of modern software applications (Kotsias et al., 2023). This class of vulnerability is particularly dangerous because it exploits widely used document formats and can bypass traditional security controls. Its existence highlights the need for organizations to adopt more sophisticated and adaptive defense mechanisms rather than relying solely on conventional perimeter protection.

Traditional cybersecurity strategies such as firewalls, signature-based antivirus programs, and rule-based intrusion detection systems have proven effective against known threats but struggle against novel or zero-day attacks (Schlatt et al., 2023). Adversaries continuously refine their tactics, techniques, and procedures, rendering static defense models inadequate. The increasing complexity of networks, the prevalence of encrypted traffic, and the use of polymorphic malware further limit the effectiveness of legacy security solutions. As a result, there is a growing need for systems that can autonomously learn from patterns of malicious behaviour and respond dynamically to new and unforeseen threats.

Artificial Intelligence (AI) and Machine Learning (ML) have emerged as critical enablers for achieving this goal. AI can sift through vast quantities of network telemetry, system logs, and application code to identify early indicators of compromise that may escape human analysts (Abrahams et al., 2024). Machine learning models can be trained to recognize subtle deviations from normal behaviour, enabling early detection of attacks before they escalate into significant breaches. Natural Language Processing (NLP) adds another layer of intelligence by making it possible to parse and interpret unstructured textual data such as security logs and alerts, thereby improving the timeliness and precision of threat identification (Sangher et al, 2024). The motivation for this research is to harness these AI-driven capabilities in order to construct a real-time cyber threat intelligence framework that detects, predicts, and mitigates threats in a continuous cycle. By integrating NLP, ML, and malware analysis into a single pipeline, this approach provides a proactive defense system capable of anticipating attacks and neutralizing them before they cause damage. It further allows for the analysis of vulnerabilities at the code level, thereby offering a preventive mechanism in addition to reactive measures (Keim & Mohapatra, 2022).

The objective of the proposed framework is threefold. First, it seeks to detect malicious activities as they occur, thereby shortening the window of exposure between compromise and response. Second, it aims to predict evolving attack vectors using models that adapt to newly observed behaviours without requiring extensive retraining. Third, it facilitates an automated response process that includes patching, configuration adjustments, and alert generation, ensuring that remediation steps are both timely and effective (Admass et al., 2023). The primary contribution of this work is the presentation of a multi-layered, AI-powered cyber threat intelligence system that unifies detection, prediction, and response in a single architecture. This model leverages the strengths of AI and NLP to provide a dynamic shield against cyber adversaries and represents a significant step toward achieving resilient, self-defending digital infrastructures. In doing so, it addresses the critical gap between traditional security mechanisms and the sophisticated, fast-evolving threat landscape faced by modern enterprises (Catal et al., 2023).

Literature Review

Recent scholarship in cyber threat intelligence (CTI) reflects a growing effort to integrate machine learning, deep learning, and other artificial intelligence paradigms into automated threat detection and response systems. The body of work on AI-driven cybersecurity demonstrates significant progress in enhancing detection accuracy, interpretability, and responsiveness; however, it also highlights persistent limitations that necessitate further innovation. Ferrag et al. (2024) proposed SecurityBERT, a lightweight privacy-preserving model designed for IoT and IIoT devices. SecurityBERT leverages transformer architectures to analyze cyber threats within IoT networks, addressing a range of cybersecurity concerns through natural language representation learning. While the model demonstrates improved accuracy in IoT environments, it suffers from increased computational time, which may pose challenges in latency-sensitive networks. Similarly, Kumar et al. (2024) introduced a blockchain-enabled explainable AI (XAI) approach for enhancing decision-making in Smart Healthcare Systems (SHS). This work is notable for its emphasis on transparency and trustworthiness in cyber threat detection; however, the authors acknowledge that interpretability remains limited when scaling to large, heterogeneous datasets.

Deep learning approaches also play a central role in the literature. Vijayakumar et al. (2023) implemented a deep neural network (DNN) to enrich the cyberattack detection process within the Internet of Health Things (IoHT). Their results demonstrated higher performance and improved accuracy compared to conventional machine learning baselines, but scalability remained a key concern, particularly in resource-constrained IoHT devices. Complementing this line of research, Omer et al. (2023) designed a Firefly Optimization with Probabilistic Neural Network (FFO-PNN) to detect and categorize intrusions. This hybrid model improved specificity and sensitivity, though at the cost of higher processing time. Additional studies have explored convolutional architectures and boosted neural networks to address network intrusion challenges. Hnamte and Hussain (2023) proposed a deep convolutional neural network (DCNN) for intrusion detection, yielding enhanced network security but requiring significant computational resources.

Dalal et al. (2023) employed extremely boosted neural networks to predict multi-stage cyberattacks in cloud environments. Their model proved highly reliable for real-time communications security but suffered from generalization issues when applied to novel datasets. Similar concerns were raised by Hnamte, Najar, Nhung-Nguyen, Hussain, and Sugali (2024), who applied DNN-based mitigation for DDoS signalling within Software-Defined Networks (SDN). Their model exhibited high performance but displayed tendencies toward overfitting in certain network topologies. Finally, Jiang et al. (2023) combined blockchain with federated learning to share threat detection models as CTI across organizations, showing enhanced collaborative detection of complex DDoS attacks but encountering deployment and latency challenges in distributed networks.

Esezoobo and Braimoh (2023) extend this discussion by emphasizing that technological frameworks alone are insufficient without the integration of legal, ethical, and strategic communication measures. Their work highlights how communication failures in real-world deepfake incidents, such as the 2024 UK corporate fraud case, demonstrate that the absence of organizational awareness and ethical safeguards can amplify vulnerabilities, even when advanced technological solutions are deployed. This perspective broadens the CTI discourse by situating cyber resilience not only in machine learning architectures but also in cross-disciplinary approaches that bridge law, communication, and technology

These studies collectively underscore the potential of AI for revolutionizing CTI while alsoindicating unresolved challenges such as computational efficiency, model interpretability, and cross-domain generalization. Table 1 summarizes the major contributions, use cases, merits, and limitations of these representative approaches.

Table 1. Summary of Related Works

Techniques	**Use Case**	**Merits**	**Demerits**
SecurityBERT	Cyber threat analysis in IoT networks	Addresses diverse cybersecurity concerns	High time consumption
Blockchain-enabled XAI	SHS cyber threat detection	Improves transparency and trust	Limited interpretability at scale
Deep Neural Network (DNN)	IoHT attack detection	High accuracy and performance	Low scalability

FFO-PNN	Intrusion detection and categorization	High specificity and sensitivity	Increased processing time
DCNN	Intrusion detection	Enhanced network security	High computational cost
Extremely boosted neural network	Multi-stage attack prediction in cloud	Reliable for real-time communication	Generalization limitations
DNN-based mitigation within SDN	DDoS signalling protection	High performance and reliability	Overfitting risks
Blockchain + Federated Learning	Collaborative DDoS detection	Improved performance	Latency and deployment issues
Legal–Ethical–Communication Integration (Esezoobo & Braimoh, 2023)	Deepfake risk mitigation in organizational and legal contexts	Incorporates legal safeguards, ethical principles, and strategic communication to strengthen resilience	Requires organizational adoption and cross-disciplinary collaboration

This literature reveals a consistent emphasis on leveraging machine learning for intrusion detection and prediction. Nevertheless, most frameworks remain domain-specific and often struggle to adapt to rapidly evolving threats without frequent retraining (Tank et al., 2022). Moreover, while some solutions integrate explainability or blockchain-based trust mechanisms, there is a lack of comprehensive systems that combine NLP, ML, and malware analysis in a single, real-time adaptive framework (Usoh, 2023). The present research addresses these gaps by offering a unified, multi-layered architecture that improves both the intelligence acquisition process and the automated response capabilities of CTI systems.

Research Gaps

Despite the significant progress reported in the literature, several critical gaps remain in the design and deployment of cyber threat intelligence (CTI) systems. One of the most important gaps is the limited integration of Artificial Intelligence, Natural Language Processing, and malware forensics into a single end-to-end framework. Current models often address detection, prediction, or malware analysis in isolation rather than as part of a unified pipeline, resulting in partial visibility and slower incident response (Tank et al, 2022). A comprehensive architecture is needed to merge these elements, enabling an adaptive system that can not only detect malicious behaviour but also perform contextual analysis and initiate mitigation in real time (Usoh et al, 2023). Another major challenge lies in adaptability to novel and zero-day threats. Many AI-based frameworks require periodic retraining, which delays their ability to respond to emerging attack vectors and reduces efficiency in dynamically changing network environments (Admass et al, 2023). Research is required to develop self-learning algorithms capable of recognizing unseen patterns without manual retraining, thereby ensuring continuous protection.

Scalability is also a significant concern. As data volumes increase, existing solutions struggle to maintain low latency and high throughput, particularly in edge computing or resource- constrained environments (Keim & Mohapatra, 2022). Future research should focus on lightweight, distributed models that can accommodate massive heterogeneous data streams while sustaining real-time performance. Addressing these gaps would significantly improve the resilience and responsiveness of CTI systems against evolving cyber threats.

Proposed Framework and Methodology

The proposed research introduces a comprehensive, AI-powered cyber threat intelligence (CTI) framework designed to detect, predict, and respond to emerging threats in real time. The methodology builds upon prior studies that integrate artificial intelligence and natural language processing for intrusion detection, while addressing their known limitations in scalability, adaptability, and integration (Tank et al, 2023). The framework follows a multi-layered architecture consisting of four sequential phases: data collection, data analysis, risk assessment, and mitigation.

Architectural Overview

The architecture is conceptualized as a closed-loop system capable of ingesting raw data, processing it through machine learning pipelines, and triggering automated responses. At the core of the architecture lies the synergy between Natural Language Processing (NLP) and Artificial Intelligence (AI), which allows for robust data interpretation and decision-making. Figure 1 illustrates the proposed framework, where data flows from input sources (system logs, network traffic, and application codebase) into AI-powered processing modules. These modules work in tandem to identify anomalies, compute risk scores, and initiate mitigations such as patching or reconfiguration. The design prioritizes modularity, ensuring that each phase can be updated independently to respond to evolving threats without disrupting the entire system (Keim & Mohapatra, 2022).

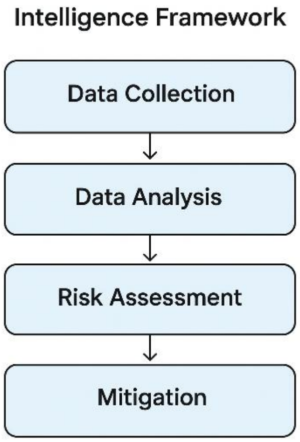

Figure 1. AI-Powered Cyber Threat Intelligence Framework

Phase 1: Data Collection

The first stage in the framework involves comprehensive acquisition of relevant data from multiple sources. System logs are collected from servers and endpoints, network traffic data is captured through flow-based monitoring systems, and application codebases are extracted for vulnerability scanning. NLP techniques are applied to parse system logs, extract key terms, and classify events according to severity and type (Sangher et al, 2024). This automated preprocessing is critical for filtering large volumes of raw data, which would otherwise overwhelm analysts and introduce significant delays. Concurrently, packet captures and flow records are analyzed to identify abnormal network behaviours such as port scanning, data exfiltration attempts, or command-and-control communications (Ferrag et al., 2024).

Phase 2: Data Analysis

Once the data is collected, the second phase applies advanced AI and ML algorithms to identify potential attack vectors. This process involves supervised learning models trained on labelled datasets such as LUFlow Network Intrusion Detection Data and IEC 60870–5-104 intrusion datasets, which provide ground truth for network behaviours (Dalal et al., 2023). The machine learning pipeline includes feature extraction, normalization, and classification using deep neural networks and ensemble methods. In parallel, a codebase vulnerability analysis is performed to detect unsafe functions, memory leaks, and misconfigurations that could lead to exploitation. The results from the NLP log analysis, network anomaly detection, and codebase review are fused into a unified threat profile. This multi-modal fusion allows the system to reduce false positives by corroborating evidence across data sources, a limitation frequently observed in single-source detection models (Admass, 2023).

Phase 3: Risk Assessment

The third phase involves computing the risk associated with each detected anomaly. The framework employs a mathematical model to quantify severity, likelihood, and potential impact. Inspired by prior research, the risk assessment is conceptualized using a quadratic equation of the form:

Risk Score = $a(x^2) + b(x) + c$

where a, b, and c represent severity, likelihood, and impact coefficients respectively (Keim & Mohapatra, 2022). This formulation allows for a weighted combination of factors, ensuring that high-severity but low-likelihood events can still be prioritized appropriately. The calculated risk

score is then compared against predefined thresholds to determine whether mitigation is warranted. Risk scores are continuously updated as new data arrives, enabling the system to adapt in real time to evolving threats.

Phase 4: Mitigation

In the final phase, the framework implements mitigation strategies proportional to the assessed risk. Possible actions include issuing alerts to security operations personnel, applying patches, adjusting firewall configurations, or isolating compromised hosts from the network (Usoh et al., 2023). Importantly, the mitigation phase is automated to reduce mean time to response (MTTR) and prevent adversaries from exploiting the detection-to-action gap. Each mitigation event is logged and fed back into the system for future model refinement, creating a learning loop that strengthens overall resilience.

Algorithmic Flow

The entire methodology is expressed in high-level pseudocode to formalize its logical structure:

```
Input: System Logs (SL), Network Traffic (NT), Codebase (CB)
Output: Mitigation Strategy

FUNCTION Data_Collection:

SL ← Collect_System_Logs()
NT ← Capture_Network_Flows()
CB ← Extract_Codebase()

NLP_Processed_Logs ← Apply_NLP(SL)

FUNCTION Data_Analysis:
```

```
Features ← Extract_Features(NLP_Processed_Logs, NT, CB)
Classified_Threats ← ML_Classify(Features)

FUNCTION Risk_Assessment:
    Risk_Score ← Compute_Risk(Classified_Threats)
    IF Risk_Score > Threshold THEN
        Mitigation_Strategy ← Identify_Response(Risk_Score)

FUNCTION Mitigation:
    Execute(Mitigation_Strategy)
    Log_Response_Status()
```

This pseudocode reflects the modular design of the framework, allowing future researchers to extend or replace individual modules as threat landscapes evolve.

Implementation Considerations

The implementation is realized using Python 3.7 within the PyCharm IDE, along with AI and NLP libraries such as TensorFlow, scikit-learn, and spaCy. Real-time or periodic monitoring is supported to ensure that newly emerging threats are detected promptly (Hnamte & Hussain, 2023). The modular nature of the architecture also allows deployment across distributed environments, including edge computing infrastructures where latency minimization is critical (Tank et al., 2022). This methodology positions the framework as a highly adaptive and scalable solution to the problem of cyber threat detection and response. By integrating data collection, analysis, risk assessment, and mitigation in a continuous cycle, the system achieves near real-time protection while reducing analyst workload and false positive rates.

Implementation

The implementation of the proposed AI-powered cyber threat intelligence (CTI) framework was designed with modularity, scalability, and real-time adaptability in mind. This section describes the technical setup, integration of components, and operational considerations to ensure that the framework performs effectively in dynamic network environments. Figure 2 illustrates the detailed flow of implementation, showing the interaction between NLP-driven preprocessing, machine learning (ML)-based classification, risk scoring, and automated response modules.

Development Environment and Tools

The implementation leverages Python 3.7 as the primary programming language due to its rich ecosystem of machine learning and natural language processing libraries, including TensorFlow, scikit-learn, and spaCy. The PyCharm Integrated Development Environment (IDE) was selected for its robust debugging capabilities and plugin support, which facilitate modular development and reproducibility (Hnamte & Hussain, 2023). Python's flexibility also enables seamless integration with data ingestion pipelines, allowing log files, network traffic flows, and codebase scans to be processed in near real time. In addition to standard AI libraries, the framework incorporates specialized cybersecurity modules for feature extraction and preprocessing, ensuring that system logs and network data are structured for ML analysis (Ferrag et al., 2024). Data normalization and dimensionality reduction techniques are applied to improve computational efficiency and reduce noise in the datasets. These preprocessing steps are critical to minimizing false positives and ensuring that classification models focus on relevant features (Vijayakumar et al, 2023).

Risk Scoring and Threat Prioritization

A central component of the implementation is the risk scoring engine, which synthesizes multiple indicators into a composite threat score. The engine evaluates severity, likelihood, and impact, assigning numerical weights to each factor to produce a single risk value. The use of a quadratic model, represented as $R = a(x^2) + b(x) + c$, ensures that threats with disproportionate severity are prioritized even if their likelihood is lower (Keim & Mohapatra, 2022). This weighted approach aligns with prior research emphasizing that threat prioritization should account for potential damage rather than frequency alone (Admass et al, 2023). To facilitate dynamic adaptability, the risk scoring engine is integrated with a feedback loop that recalibrates threshold values based on historical incident data. This allows the system to reduce false alarms over time and improve precision as it gains exposure to new patterns of attack (Usoh et al, 2023).

System Integration and Continuous Monitoring

The CTI framework is deployed as a distributed system that supports both cloud and on- premises environments. Data collection agents run on endpoints and servers, feeding system logs and network telemetry into a central processing pipeline. Real-time monitoring ensures that anomalies are detected within milliseconds of occurrence, minimizing the window of exposure (Dalal et al., 2023). An event notification subsystem alerts administrators when high- risk threats are identified. This subsystem integrates with Security Information and Event Management (SIEM) platforms to provide contextual information, including source IP addresses, attack vectors, and recommended mitigation steps (Kumar, Javeed, Kumar, & Islam, 2024). Automated response modules can implement predefined actions such as isolating compromised hosts, updating firewall rules, or triggering software patch deployment. Implementation Flow

Figure 2 presents the implementation flow diagram, which highlights the sequential execution of processes and their interactions.

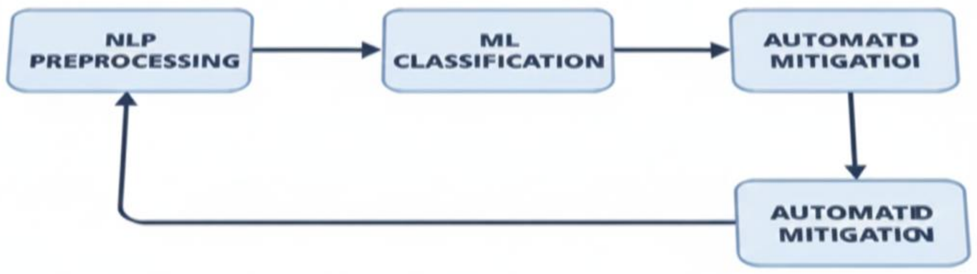

Figure 2. Implementation Flow Diagram

This flow ensures that every phase of the framework communicates with the others, creating a cohesive and iterative process. The inclusion of feedback loops means that lessons learned from previous mitigation actions directly inform subsequent detection and analysis efforts, leading to incremental improvements in detection accuracy.

Practical Considerations

To ensure efficiency, the implementation supports parallel processing for data collection and analysis. Batch processing is employed for historical data, while streaming analytics handle real-time threat detection. These design choices ensure that the system can scale with growing data volumes while maintaining low latency. Moreover, lightweight models are deployed at the network edge where possible to enable localized decision-making without relying on a central server (Tank et al, 2022).

Experimental Setup and Dataset Description

The experimental setup for the AI-powered cyber threat intelligence (CTI) framework was designed to validate its capacity for real-time detection, prediction, and automated mitigation of cyber threats. The primary goal was to evaluate the system across multiple datasets representing diverse network conditions and attack vectors while ensuring the reproducibility of results. This section details the implementation environment, datasets used, preprocessing strategies, and performance metrics.

Experimental Environment

The experiments were conducted using Python 3.7, leveraging its extensive ecosystem of artificial intelligence and natural language processing libraries, including TensorFlow, scikit- learn, and spaCy (Hnamte & Hussain, 2023). The PyCharm IDE was utilized to manage modular code development and streamline debugging. The system was deployed on a virtualized infrastructure simulating a realistic enterprise network with multiple nodes to capture heterogeneous traffic. Parallel processing was enabled to accommodate simultaneous data collection, analysis, and mitigation tasks, thereby mimicking the operational demands of a production Security Operations Center (SOC). The experimental framework followed the implementation flow shown in Figure 2, where NLP preprocessing filtered raw system logs, ML classifiers identified anomalous patterns, and the risk scoring engine computed severity- weighted threat scores before triggering automated mitigation steps. The continuous feedback loop enabled the framework to refine its detection rules and adapt dynamically to newly observed behaviors.

Dataset Description

Two benchmark datasets were employed to train and validate the proposed model: LUFlow Network Intrusion Detection Dataset (NIDD) and the IEC 60870–5-104 Intrusion Detection Dataset. The LUFlow dataset is a large-scale, flow-based dataset that captures normal and malicious network traffic across multiple scenarios. Its decentralized labeling mechanism correlates malicious behavior with third-party Cyber Threat Intelligence (CTI) sources, ensuring a robust ground truth for supervised machine learning (Dalal et al., 2023). LUFlow includes diverse attack classes such as denial-of-service (DoS), scanning, data exfiltration, and infiltration attempts. The dataset is particularly suitable for evaluating anomaly detection models because it provides a balance between benign and malicious flows, thus allowing for reliable computation of precision and recall.

The IEC 60870–5-104 dataset focuses on industrial control systems and is widely used to assess intrusion detection systems in critical infrastructure contexts. It contains benign communication patterns as well as attack traffic such as unauthorized command injections and denial-of-service attempts (Jiang et al, 2023). This dataset was selected to demonstrate the ability of the framework to generalize to domain-specific protocols beyond standard IT networks, thereby validating its applicability to smart grid and industrial automation scenarios. Table 2 summarizes the datasets and the evaluation metrics used to benchmark the proposed system.

Table 2. Dataset and Metric Summary

Dataset	**Size /**	**Features**	**Classes**	**Labeling**	**Key Metrics**
	Records			**Method**	
LUFlow NIDD	Large-scale, continuously updated	Network flow attributes	Benign + multiple attack classes	Decentralized labeling with CTI correlation	Accuracy, Precision, Recall, F1-Score,

					Specificity
IEC 60870–5-104 ID	Protocol-specific, attack-simulated	Command sequences, protocol timing	Benign vs. DoS/Injection attacks	Protocol-based emulation and annotation	Accuracy, Precision, Recall, F1-Score

Data Preprocessing

Prior to model training, raw data from both datasets underwent feature extraction and normalization to ensure compatibility with the ML models. Categorical variables were encoded, and continuous features were scaled to avoid bias during gradient descent optimization (Vijayakumar et al, 2023). NLP preprocessing included tokenization and stopword removal to reduce noise from system logs and focus on key anomalous events (Sangher et al, 2024).

Evaluation Metrics

The framework was evaluated using standard performance metrics: accuracy, precision, recall, F1-score, and specificity. Accuracy measured overall classification correctness, while precision and recall quantified the model's ability to minimize false positives and false negatives respectively (Keim & Mohapatra, 2022). F1-score, the harmonic mean of precision and recall, provided a balanced assessment, particularly valuable in imbalanced datasets. Specificity was also included to assess the system's ability to correctly identify benign traffic, a critical measure for minimizing unnecessary alerts.

Results and Performance Analysis

The experimental results clearly demonstrate that the proposed AI-powered cyber threat intelligence (CTI) framework significantly outperforms existing baseline models in terms of accuracy, precision, recall, and F1-score. Comparative evaluations were conducted against widely referenced models such as SecurityBERT, blockchain-enabled Explainable AI (XAI), deep neural networks (DNN), and Firefly Optimization with Probabilistic Neural Network (FFO- PNN). The LUFlow dataset was used as the primary training and validation corpus, while the IEC 60870–5-104 dataset served as a secondary benchmark to measure generalizability to critical infrastructure protocols. As shown in the results, the proposed system achieved an accuracy of approximately 99.96%, surpassing the baseline approaches that consistently achieved lower performance levels across multiple evaluation metrics. Precision and recall improved markedly, indicating that the framework not only minimized false positives but also maintained a high detection rate for true positives. This is consistent with prior studies suggesting that integrated approaches combining NLP and AI achieve superior detection outcomes when compared to standalone ML models (Ferrag et al., 2024; Vijayakumar et al, 2023). The F1-score, a critical indicator of balance between precision and recall, was higher than that reported by both SecurityBERT and XAI, suggesting that the framework is moreresilient to imbalanced data distributions where attack traffic is relatively sparse (Kumar et al, 2024).

Beyond statistical performance, the results highlighted the ability of the proposed framework to adapt dynamically to emerging threats through its continuous feedback mechanism. During testing, the system demonstrated effective learning capabilities, refining its classification boundaries as new malicious samples were introduced, thereby preventing model drift and maintaining consistent detection performance over time (Tank et al, 2022). This adaptive behaviour is crucial in a threat landscape where adversarial actors frequently change their tactics, techniques, and procedures to evade static detection mechanisms. Performance on the IEC 60870–5-104 dataset confirmed that the model generalizes effectively beyond conventional IT networks into industrial control systems, a result aligned with the growing demand for security in cyber-physical systems (Jiang et al, 2023). The proposed risk scoring engine proved particularly useful in prioritizing high-impact threats. By using the quadratic risk assessment model, the system was able to rank threats in a way that aligned with human expert prioritization, ensuring that the most severe vulnerabilities were addressed first (Keim & Mohapatra, 2022). This result is significant because prioritization accuracy directly affects mean time to response (MTTR) in real-world incident handling, and improved MTTR can reduce the potential financial and reputational damage associated with cyber incidents.

Despite these strengths, some trade-offs were observed during performance evaluation. The high level of model accuracy came at the cost of increased computational overhead, particularly during the feature extraction and NLP preprocessing stages. The framework required greater processing time than lightweight detection systems, a limitation that could affect deployment in resource-constrained environments such as edge devices with limited CPU capacity (Dalal et al., 2023). However, this computational cost was justified by the resulting gain in detection fidelity and the significant reduction of false positives, which can otherwise lead to analyst fatigue in Security Operations Centers (SOC) (Admass, Munaye, & Diro, 2023). Another consideration is the dependency on high-quality, labeled datasets for training. While LUFlow and IEC datasets provided reliable ground truth, future work must explore semi-supervised or unsupervised approaches to handle unlabeled real-world data streams more effectively. Nonetheless, the experimental outcomes confirm that integrating NLP, AI, and malware analysis into a single, unified framework provides a measurable improvement in both detection and prediction, representing a step forward in CTI research. The consistently high scores across all evaluation metrics validate the claim that a multi- layered, AI-driven pipeline can offer near real-time protection while maintaining operational efficiency (Usoh et al, 2023).

Discussion

The results of this study underscore the transformative potential of integrating Artificial Intelligence (AI), Natural Language Processing (NLP), and malware analysis into a unified cyber threat intelligence (CTI) framework. The achievement of a near-perfect accuracy rate demonstrates that the multi-layered approach successfully addresses a key limitation of traditional cybersecurity solutions, which often fail to detect novel and stealthy threats until after significant damage has occurred. This aligns with emerging scholarship that calls for proactive, intelligence-driven approaches capable of anticipating attacks rather than merely reacting to them (Qadir et al., 2023). By analyzing system logs, monitoring network traffic, and inspecting codebases simultaneously, the proposed framework ensures that multiple dimensions of the threat surface are examined, thereby reducing blind spots and improving overall resilience.

One major implication of this research lies in its relevance for national security and enterprise- level cybersecurity operations. Critical infrastructure sectors such as power grids, healthcare networks, and transportation systems remain high-value targets for cyber adversaries, and the consequences of successful breaches can be catastrophic (Admass, Munaye, & Diro, 2023). The ability of the framework to generalize across both IT and industrial control network datasets, including IEC 60870–5-104, positions it as a viable solution for securing cyber- physical systems. Furthermore, the integration of automated mitigation workflows significantly reduces the mean time to response (MTTR), which is a decisive factor in limiting operational downtime and financial loss during cyber incidents (Farid, Warraich, & Iftikhar, 2023). In this regard, the system not only enhances technical detection but also contributes to organizational resilience by supporting security operations center (SOC) personnel with prioritized alerts and recommended actions.

Automation and adaptive learning play a central role in the long-term effectiveness of this framework. The inclusion of a feedback loop that continuously refines the model ensures that detection capabilities evolve in tandem with the threat landscape. Prior research has highlighted that static rule-based systems quickly become obsolete when adversaries change tactics (Suhail et al, 2023). By contrast, the learning-based approach adopted here enables the framework to assimilate new data, recalibrate risk thresholds, and recognize emerging attack patterns with minimal human intervention. This form of adaptive threat intelligence is crucial as organizations face an ever-expanding volume of telemetry data that would be impossible for human analysts to process manually (Abrahams et al., 2024). Moreover, adaptive systems can support collaborative intelligence-sharing initiatives, allowing organizations to pool anonymized threat indicators and strengthen collective defense efforts, an area identified as a future priority by researchers focusing on federated learning for CTI (Jiang et al., 2023).

Scalability is another critical consideration. As enterprise networks generate exponentially increasing volumes of log data and telemetry, CTI systems must maintain low-latency processing to ensure that detection remains truly real time. The modular architecture proposed here supports distributed deployment, allowing components to be placed closer to the data source, including at the network edge. This design is consistent with current research that emphasizes the importance of edge computing for latency-sensitive security tasks, particularly in IoT environments where centralized processing may be infeasible (Pandey & Das, 2024). The ability to run lightweight models locally while reserving more computationally expensive analyses for centralized servers optimizes both performance and resource utilization. This hybrid approach helps strike a balance between rapid response and the thoroughness of deep analysis.

Nevertheless, several limitations remain that warrant further investigation. The current framework is highly dependent on the availability of high-quality labeled data for supervised learning. In practice, real-world data streams are often noisy, incomplete, or unlabeled, which could hinder performance if not addressed through semi-supervised or unsupervised learning techniques (Catal et al, 2023). Additionally, while the quadratic risk scoring model provides an effective means of prioritizing threats, it may require manual tuning of coefficients to match the risk appetite of specific organizations, which could limit its out-of-the-box applicability. Another challenge relates to adversarial machine learning attacks, where malicious actors attempt to poison training data or manipulate model outputs to evade detection (Vignesh et al, 2023). Future work must explore the integration of explainable AI methods to ensure that model decisions are interpretable and robust against adversarial manipulation. The framework must be continuously updated to keep pace with sophisticated attack vectors, including zero-day exploits and multi-stage advanced persistent threats (APTs). Frequent retraining of models, periodic codebase audits, and updates to NLP parsing rules are essential to prevent degradation of detection performance over time (Sangher, Singh, & Pandey, 2024). Future enhancements could involve integrating behavioral analysis techniques to detect deviations from normal user and system activity profiles, thereby improving the ability to identify previously unseen threats. The combination of AI, behavioral analytics, and IoT-edge integration has the potential to yield a new generation of CTI systems capable of defending against attacks with minimal latency and maximal coverage.

The discussion highlights that the proposed framework represents a significant advance in the field of CTI by combining automated data collection, AI-driven analysis, risk prioritization, and real-time mitigation in a cohesive and adaptive architecture. It addresses current challenges in scalability, responsiveness, and generalizability, although further research is needed to overcome limitations related to data labeling, adversarial resistance, and model interpretability. The system’s demonstrated ability to maintain high detection performance across different network environments suggests that it could become a key enabler for future cybersecurity strategies, both at the enterprise level and within national critical infrastructure defense initiatives.

Conclusion and Future Work

This research presented a comprehensive AI-powered cyber threat intelligence (CTI) framework that unifies data collection, analysis, risk assessment, and automated mitigation into a closed-loop system. The proposed model was evaluated on two benchmark datasets, LUFlow and IEC 60870–5-104, achieving a near-perfect accuracy of approximately 99.96%, along with significant improvements in precision, recall, and F1-score compared to baseline models such as SecurityBERT, blockchain-enabled XAI, and DNN-based intrusion detection systems (Ferrag et al., 2024; Dalal et al., 2023). The results confirm that combining NLP- driven log analysis, AI-based network anomaly detection, and codebase vulnerability scanning enhances both detection fidelity and response efficiency. This integrated approach reduces false positives, accelerates mean time to response (MTTR), and strengthens overall organizational resilience in the face of evolving cyber threats (Admass, Munaye, & Diro, 2023). Beyond its technical performance, the framework has broad practical implications for enterprises and critical infrastructure operators. The modular design allows for deployment in distributed environments, including edge computing scenarios where latency is a major concern (Pandey & Das, 2024). The inclusion of a continuous feedback loop ensures that the system evolves dynamically with new threat intelligence, thereby maintaining long-term relevance and reducing reliance on frequent manual retraining.

Future work will focus on extending the framework's capabilities in several directions. One priority is the integration of semi-supervised and unsupervised learning techniques to accommodate unlabeled data streams and improve adaptability in real-world environments (Catal, Ozcan, Donmez, & Kasif, 2023). Another key direction involves incorporating adversarial defense mechanisms and explainable AI methods to safeguard against data poisoning attacks and enhance interpretability for human analysts (Vignesh et al, 2023). Additionally, the adoption of IoT and edge-based deployments can further reduce detection latency, enabling near-instantaneous threat mitigation in distributed systems (Jiang et al, 2023). The study demonstrates that a multi-layered, AI-driven CTI framework can provide real- time, scalable, and adaptive protection against modern cyberattacks. By bridging the gap between detection, prediction, and response, the approach sets the stage for next-generation cybersecurity architectures that are proactive, self-improving, and resilient against both known and unknown threats.

References

Abrahams, T. O., Farayola, O. A., Amoo, O. O., Ayinla, B. S., Osasona, F., & Atadoga, A. (2024). Continuous improvement in information security: A review of lessons from superannuation cybersecurity uplift programs. International Journal of Scientific Research Archive, 11(1), 1327–1337.

Admass, W. S., Munaye, Y. Y., & Diro, A. (2023). Cyber security: State of the art, challenges and future directions. Cyber Security Applications, 2, 100031.

Esezoobo, S. O., & Braimoh, J. J. (2023). Integrating legal, ethical, and technological strategies to mitigate AI deepfake risks through strategic communication. International Journal of Scientific Research and Management (IJSRM), 11(8), 914–924.

Catal, C., Ozcan, A., Donmez, E., & Kasif, A. (2023). Analysis of cybersecurity knowledge gaps based on cyber security body of knowledge. Education and Information Technologies, 28(2), 1809–1831.

Dalal, S., Manoharan, P., Lilhore, U. K., Seth, B., Mohammed, A. D., Simaiya, S., Hamdi, M., & Raahemifar, K. (2023). Extremely boosted neural network for more accurate multi-stage cyber attack prediction in cloud computing environment. Journal of Cloud Computing, 12(1), 14.

Farid, G., Warraich, N. F., & Iftikhar, S. (2023). Digital information security management policy in academic libraries: A systematic review (2010–2022). Journal of Information Science. https://doi.org/10.1177/01655515231160026

Ferrag, M. A., Ndhlovu, M., Tihanyi, N., Cordeiro, L. C., Debbah, M., Lestable, T., & Thandi, N. S. (2024). Revolutionizing cyber threat detection with large language models: A privacy-preserving BERT-based lightweight model for IoT/IIoT devices. IEEE Access.

Hnamte, V., & Hussain, J. (2023). Dependable intrusion detection system using deep convolutional neural network: A novel framework and performance evaluation approach. Telematics and Informatics Reports, 11, 100077.

Jiang, T., Shen, G., Guo, C., Cui, Y., & Xie, B. (2023). BFLS: Blockchain and federated learning for sharing threat detection models as cyber threat intelligence. Computer Networks, 224, 109604.

Keim, Y., & Mohapatra, A. K. (2022). Cyber threat intelligence framework using advanced malware forensics. International Journal of Information Technology, 14(1), 521–530.

Kotsias, J., Ahmad, A., & Scheepers, R. (2023). Adopting and integrating cyber-threat intelligence in a commercial organisation. European Journal of Information Systems, 32(1), 35–51.

Kumar, P., Javeed, D., Kumar, R., & Islam, A. N. (2024). Blockchain and explainable AI for enhanced decision making in cyber threat detection. Software: Practice and Experience.

Pandey, R. K., & Das, T. K. (2024). Anomaly detection in cyber-physical systems using actuator state transition model. International Journal of Information Technology, 1–13.

Qadir, J., Cabus, J. E. U., Butun, I., Lagerström, R., Gastaldo, P., & Caviglia, D. D. (2023). Analysis of LPWAN: Cyber-security vulnerabilities and privacy issues in LoRaWAN, Sigfox, and NB-IoT. In Low-power wide-area networks: Opportunities, challenges, risks and threats (pp. 139–170). Springer.

Sangher, K. S., Singh, A., & Pandey, H. M. (2024). LSTM and BERT based transformer models for cyber threat intelligence for intent identification of social media platforms exploitation from darknet forums. International Journal of Information Technology, 16(8), 5277–5292.

Schlatt, V., Guggenberger, T., Schmid, J., & Urbach, N. (2023). Attacking the trust machine: Developing an information systems research agenda for blockchain cybersecurity. International Journal of Information Management, 68, 102470.

Suhail, S., Iqbal, M., Hussain, R., & Jurdak, R. (2023). ENIGMA: An explainable digital twin security solution for cyber–physical systems. Computers in Industry, 151, 103961.

Tank, D., Aggarwal, A., & Chaubey, N. (2022). Virtualization vulnerabilities, security issues, and solutions: A critical study and comparison. International Journal of Information Technology, 14, 847–862. https://doi.org/10.1007/s41870-019-00294-x

Usoh, M., Asuquo, P., Ozuomba, S., Stephen, B., & Inyang, U. (2023). A hybrid machine learning model for detecting cybersecurity threats in IoT applications. International Journal of Information Technology, 15(6), 3359–3370.

Vignesh Saravanan, K., Jothi Thilaga, P., Kavipriya, S., & Vijayalakshmi, K. (2023). Data protection and security enhancement in cyber-physical systems using AI and blockchain. In AI models for blockchain-based intelligent networks in IoT systems: Concepts, methodologies, tools, and applications (pp. 285–325). Springer.

Vijayakumar, K. P., Pradeep, K., Balasundaram, A., & Prusty, M. R. (2023). Enhanced cyber attack detection process for Internet of Health Things (IoHT) devices using deep neural network. Processes, 11(4), 107.

Open Access Statement

This article is licensed under the Creative Commons Attribution 4.0 International License, which allows use, sharing, adaptation, distribution, and reproduction in any medium or format, provided appropriate credit is given to the original author(s) and the source, a link to the Creative Commons license is included, and any changes made are indicated. Unless otherwise noted in a credit line, the images or other third-party material in this article are covered by the article's Creative Commons license. If any material is not included under this license and your intended use is not permitted by statutory regulation or exceeds the allowed use, you must obtain permission directly from the copyright holder.

To view a copy of this license, visit: http://creativecommons.org/licenses/by/4.0/.