

A Unified AI Framework for Early Threat Detection in Healthcare: Combining Anomaly Detection, Sequence Learning, and CTI Correlation

Kalyana Krishna Kondapalli

CEO, Mytecz, India.

To Cite this Article: Kalyana Krishna Kondapalli, "A Unified AI Framework for Early Threat Detection in Healthcare: Combining Anomaly Detection, Sequence Learning, and CTI Correlation", *Indian Journal of Computer Science and Technology*, Volume 05, Issue 02 (May-August 2026), PP: 123-128.



Copyright: ©2026 This is an open access journal, and articles are distributed under the terms of the [Creative Commons Attribution License](#); Which Permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract: The accelerating digitization of clinical environments—through cloud-hosted services, the Internet of Medical Things (IoMT), and interoperable electronic health records—has substantially expanded the attack surface of healthcare delivery organizations. Traditional, signature- and rule-driven defenses respond poorly to fast-evolving, polymorphic, and zero-day intrusions that target sensitive patient data and life-critical clinical workflows. This work proposes an Adaptive AI-Driven Threat Intelligence Framework engineered to deliver proactive cyber defense across heterogeneous healthcare information systems. The framework couples supervised classifiers, unsupervised anomaly detectors, and sequence-aware deep learning to perform continuous monitoring, behavioral profiling, and predictive threat detection. By correlating real-time telemetry from network flows, system and authentication logs, IoMT device events, and curated external Cyber Threat Intelligence (CTI) feeds, it surfaces novel and evasive attack patterns that fixed-rule systems miss. An adaptive learning loop retrains detection models and refreshes security policies as new evidence arrives, while explainable-AI components surface human-readable rationale for each alert. Experimental evaluation in a simulated hospital environment shows measurable gains in detection accuracy, sharp reductions in mean time to detect (MTTD) and mean time to respond (MTTR), and stronger alignment with healthcare data-protection mandates such as HIPAA and GDPR.

Key Words: Artificial Intelligence; Threat Intelligence; Healthcare Cybersecurity; Proactive Defense; IoMT Security; Explainable AI; Anomaly Detection.

I. INTRODUCTION

Modern healthcare delivery is increasingly mediated by digital platforms. Electronic health record (EHR) systems, telemedicine portals, cloud-hosted analytics, and IoMT devices are now embedded in routine clinical workflows, reshaping both patient outcomes and operational throughput. While this transformation has unlocked clear gains in care coordination and efficiency, it has simultaneously produced a far broader and more complex attack surface. Adversaries have responded with targeted ransomware operations, large-scale data exfiltration, credential-phishing campaigns, and advanced persistent threats (APTs) aimed specifically at hospitals, payers, and clinical research entities, where downtime translates directly into patient risk.

Legacy defensive measures—static rule sets in firewalls, signature-based intrusion detection systems (IDS), and identity-bound access controls—were designed for slower, more predictable threat environments. They struggle with stealthy, polymorphic, and previously-unseen attacks, producing late detections, elevated false-positive rates, and degraded overall cyber resilience. In environments where clinical operations cannot tolerate prolonged disruption, this defensive lag is unacceptable.

Closing this gap requires a fundamentally more proactive and self-adjusting approach. To that end, this paper presents an Adaptive AI-Driven Threat Intelligence Framework that targets the layered, dynamic security challenges characteristic of healthcare information systems. The framework integrates artificial intelligence (AI), machine learning (ML), and deep learning to monitor, anticipate, and respond to threats in near real time. Fig. 1 illustrates the high-level architecture, comprising data ingestion, predictive analytics, behavioral profiling, and automated triage components. The overarching aim is earlier threat identification, automated containment, and continuous alignment with regulatory obligations governing patient information.

A. Research Objectives

The principal objectives of this study are: (i) to design and operationalize an adaptive, AI-based framework for proactive cybersecurity in dynamic healthcare environments; (ii) to strengthen detection of zero-day and emergent threats through the combined use of machine learning, behavioral analytics, and live threat intelligence; and (iii) to compress incident response times and improve overall cyber resilience while maintaining strict adherence to healthcare data-protection regulations.

B. Contributions

The contributions of this work are summarized as follows. First, we introduce a threat intelligence pipeline tailored to the

operational and regulatory constraints of healthcare information systems. Second, we present an adaptive learning component capable of incrementally adjusting to drifting threat landscapes and clinical workflow changes. Third, we fuse predictive analytics with policy-driven automation so that defensive posture shifts from reactive containment to anticipatory mitigation. Finally, we incorporate explainability mechanisms that expose the drivers of each alert, enabling security analysts and clinical stakeholders to act on AI output with confidence and auditability.

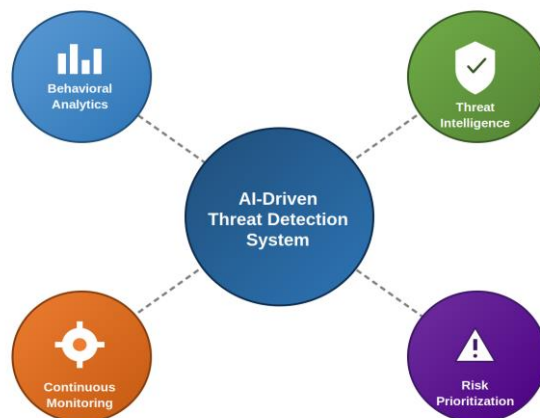


Fig. 1. AI-Driven Threat Intelligence Framework for Healthcare Cyber Defense.

II. LITERATURE REVIEW

The rapid digitalization of healthcare has reshaped both clinical practice and operational efficiency, but it has also expanded organizations' exposure to cyber risk. As providers adopted EHRs, telemedicine, and IoMT devices, security challenges grew in parallel. Early scholarship in healthcare cybersecurity focused largely on perimeter and signature-based controls—firewalls, encryption-at-rest, role-based access management, and rule-driven IDS [1], [2]. While effective against catalogued threats, these mechanisms have repeatedly fallen short against zero-day exploits, insider misuse, and coordinated attacks against clinical infrastructure [3]. The brittleness of static rule logic in interoperable, real-time clinical settings has been documented across multiple operational case studies [4].

Subsequent research turned toward machine learning (ML) and AI as candidate accelerators for healthcare defense. ML-driven IDSs—built on anomaly detection, supervised classification, and unsupervised clustering—have outperformed legacy approaches in flagging malicious traffic and atypical user behavior [5], [6]. Yet these systems frequently exhibit elevated false-positive rates, limited generalization to novel attack distributions, and difficulty integrating heterogeneous healthcare telemetry [7], [8]. Many also depend on labeled incident data, which remains scarce in healthcare due to privacy constraints, fragmented data ownership, and reluctance to disclose breach information [9].

More recent advances in deep learning and behavioral analytics have demonstrated stronger performance against sophisticated adversarial patterns and have helped suppress false alarms. Several proposals couple AI models with live threat intelligence feeds to keep pace with shifting attack tradecraft [10]. Nevertheless, recurring concerns persist around model interpretability, scalability under clinical loads, and verifiable regulatory alignment [11]. Most published work continues to address narrow slices of the problem— anomaly detection, threat correlation, or risk scoring in isolation—rather than offering an end-to-end, dynamic solution shaped to healthcare-specific operational realities [12].

A more recent line of work emphasizes AI-centric threat intelligence architectures that combine deep learning, behavioral analytics, and real-time external feed correlation to enable proactive defense [13]. These contributions show that adaptive, self-improving models can anticipate emerging attack patterns, shorten response timelines, and improve resilience to ransomware and APT activity [14]. However, open issues around explainability, scaling, and demonstrable compliance remain, motivating the design of domain-specialized adaptive frameworks for healthcare information systems [15].

The principal gap that emerges from this body of work is the absence of a unified, adaptive, AI-driven framework for proactive healthcare cyber defense that simultaneously integrates supervised learning, unsupervised anomaly detection, deep behavioral analysis, real-time threat intelligence, and explainable AI in one cohesive system. Existing models often underweight the dynamic nature of clinical environments, the necessity of continuous online learning, and the operational constraint that defensive actions must not jeopardize service continuity or regulatory standing. Few studies further examine how predictive analytics and orchestrated automation can jointly intercept threats before they impact patient care or clinical workflows.

This paper addresses these gaps by proposing an Adaptive AI-Driven Threat Intelligence Framework that adopts a hybrid analytical approach—pairing supervised classification, unsupervised outlier detection, deep-learning-based behavioral analysis, and real-time threat intelligence correlation. In contrast to prior approaches, the framework is engineered to absorb new threat patterns continuously, automate response within safety boundaries, and align by design with healthcare data-protection mandates.

The combination of predictive analytics, behavioral analysis, and adaptive learning produces a scalable, explainable, and proactive cybersecurity layer purpose-built for healthcare information systems.

A. Problem Statement

The expanding adoption of EHRs, cloud-based clinical services, telemedicine, and IoMT endpoints has elevated the exposure of healthcare information systems to advanced cyber threats. Many deployed defenses remain reactive and fragmented, providing inadequate coverage against zero-day exploits, APTs, insider misuse, and rapidly mutating ransomware operations. Such mechanisms typically lack real-time situational awareness, continuous learning, and predictive insight, leading to delayed detections, inflated false-positive rates, and prolonged incident response cycles. The challenge is compounded by the diversity of healthcare data sources, stringent regulatory expectations, and the non-negotiable requirement of uninterrupted clinical service. Consequently, there is a pressing need for an integrated, adaptive, intelligence-driven framework capable of proactively detecting, analyzing, and mitigating emerging threats while also delivering scalability, explainability, and verifiable compliance—an objective whose realization is essential for protecting patient data, sustaining system availability, and reinforcing the cyber resilience of contemporary healthcare information systems.

III. METHODOLOGY

The proposed methodology realizes an Adaptive AI-Driven Threat Intelligence Framework that supports proactive cyber defense across dynamic healthcare information systems by combining multi-source data acquisition, feature engineering, hybrid AI modeling, threat scoring, and automated response. Stages are arranged in a staged pipeline to safeguard reliability, scalability, and regulatory compliance in clinical settings.

1) Data Collection and Threat Intelligence Ingestion

Security-relevant telemetry is continuously collected from a broad range of healthcare assets, including application logs, authentication and access events, IoMT device telemetry, endpoint records, network flow records (e.g., NetFlow), DNS and HTTP traces, and cloud audit trails. In parallel, external Cyber Threat Intelligence (CTI) feeds supply indicators of compromise (IoCs), tactics-techniques-procedures (TTPs), and vulnerability advisories. Every event is normalized into a shared schema (timestamp, asset identifier, user identifier, event type, severity, source, destination) so that cross-source correlation becomes tractable.

2) Preprocessing, Normalization, and Privacy-Aware Handling

Raw logs commonly contain missing fields, duplicates, and inconsistent formats. After deduplication and cleansing, structured features are encoded. To safeguard patient information, identifiers are tokenized or pseudonymized and only security-relevant attributes are retained, in line with privacy-by-design principles. Numeric features are standardized as:

$$x' = (x - \mu) / \sigma$$

where x denotes the original feature value, μ its mean, and σ its standard deviation.

3) Feature Engineering and Temporal Context

Features are constructed at the user, device, and network levels to surface healthcare-specific attack signals: failed-login frequency, off-hours access, bulk record export volume, lateral-movement indicators, IoMT command counts, and privilege-escalation markers. Sliding windows aggregate event streams to retain temporal dynamics:

$$f_t = \sum \varphi(e_\tau) \text{ for } \tau = t-W \text{ to } t$$

where $\varphi(e_\tau)$ projects event e_τ into the feature space, and f_t is the aggregated feature vector at time t .

4) Hybrid Detection Model (Supervised + Unsupervised + Deep Learning)

Because healthcare attacks span both catalogued and previously-unseen patterns, a hybrid modeling strategy is adopted:

- Supervised classifier (catalogued attacks): logistic regression, random forest, or XGBoost trained on labeled incidents. For logistic regression, the attack probability is modeled as $P(y = 1 | x) = \sigma(w^T x + b) = 1 / (1 + e^{-(w^T x + b)})$.
- Unsupervised anomaly detection (unknown / zero-day): One-Class SVM or Isolation Forest is used to identify outliers relative to nominal healthcare operational behavior.
- Sequence-aware deep model: an LSTM is trained on rolling event sequences with hidden state $h_t = \text{LSTM}(x_t, h_{t-1})$; the anomaly score for a sequence is computed from the prediction error $S_{seq} = \|x_{t+1} - \hat{x}_{t+1}\|^2$.

Training uses stratified splitting and cross-validation. Class imbalance—an unavoidable feature of attack data—is mitigated through weighted loss functions or oversampling. A representative binary cross-entropy loss is given by:

$$L = -(1/N) \sum [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)]$$

5) Threat Correlation and Risk Scoring

Outputs from each detector are fused into a single per-entity (user, device, or session) threat score using a weighted aggregation:

$$R = \alpha P_{sup} + \beta A_{unsup} + \gamma S_{seq} + \delta I_{cti}$$

where P_{sup} is the supervised attack probability, A_{unsup} is the unsupervised anomaly score, S_{seq} is the sequence-based

anomaly score, I_{cti} captures CTI alignment strength (IoC and TTP matching), and $\alpha + \beta + \gamma + \delta = 1$. Thresholds map R into discrete alert tiers: low, medium, high, and critical.

6) Concept-Drift Handling and Adaptive Learning

Healthcare environments evolve continuously as new devices, workflows, and traffic profiles are introduced, producing concept drift. The framework counters this through scheduled retraining and online recalibration triggered by drift signals such as feature-distribution divergence. Drift is monitored using a Kullback–Leibler divergence:

$$D_{KL}(P \parallel Q) = \sum P(i) \log [P(i) / Q(i)]$$

where P denotes the baseline distribution and Q denotes the current distribution.

7) Explainability and Clinically-Safe Alerting

To establish trust and audit readiness, explainable-AI tools such as SHAP and LIME are used to surface the principal features behind each alert (for example, off-hours access combined with abnormal record-query volume and a CTI match). Alert prioritization additionally accounts for clinical impact—events touching ICU-adjacent systems are escalated as critical.

8) Automated Mitigation and Response Orchestration

For high-risk events, the framework launches safety-aware response actions governed by policy: endpoint isolation, step-up authentication, malicious IP or domain blocking, token revocation, and IoMT-device quarantine. Execution proceeds via SOAR/SIEM integration, with human-in-the-loop authorization required for any action that could disrupt clinical availability.

9) Evaluation Metrics and Validation

Performance is assessed using both detection-quality and operational metrics:

$$Precision = TP / (TP + FP); Recall = TP / (TP + FN); F1 = 2 \cdot (Precision \cdot Recall) / (Precision + Recall)$$

Mean Time to Detect (MTTD) and Mean Time to Respond (MTTR) are also captured to quantify proactive capability under healthcare-realistic load. Together, these measures verify that the framework remains data-driven, adaptive, and proactive while also being deployable in real-world clinical environments where privacy, compliance, and service continuity are non-negotiable.

IV. RESULTS AND DISCUSSION

The proposed Adaptive AI-Driven Threat Intelligence Framework was evaluated in a controlled, healthcare-style testbed comprising simulated EHR access logs, IoMT device telemetry, network traffic captures, and security event records. The dataset reproduced both nominal operating conditions and injected adversarial scenarios, including ransomware activity, credential misuse, lateral movement, and data exfiltration. The framework was benchmarked against a conventional rule-based IDS and standalone machine-learning baselines, with detection accuracy, response efficiency, and robustness used as the primary axes of comparison.

A. Detection Performance Analysis

Table I reports comparative detection results. The findings indicate that the proposed hybrid AI framework consistently outperforms both conventional and single-model baselines. By fusing supervised classification, unsupervised anomaly detection, and sequence-aware learning, the framework reliably catches both catalogued and previously-unseen attacks. Notably, recall is markedly higher, meaning a larger share of true attacks is captured—an essential property in healthcare, where missed detections can cascade into patient-safety consequences.

Method	Acc. (%)	Prec. (%)	Rec. (%)	F1 (%)
Rule-Based IDS	82.4	79.1	76.8	77.9
Standalone ML Model	88.7	86.5	84.2	85.3
Deep Learning Model	91.3	89.6	88.4	89.0
Proposed AI-Driven Framework	95.8	94.2	93.6	93.9

TABLE I. DETECTION PERFORMANCE COMPARISON

Key observations from Table I are: (i) the rule-based IDS records the weakest results across all metrics, with accuracy of 82.4%, recall of 76.8%, and F1 of 77.9%; (ii) the standalone ML baseline improves accuracy to 88.7%, recall to 84.2%, and F1 to 85.3%; (iii) the deep-learning baseline pushes accuracy to 91.3%, recall to 88.4%, and F1 to 89.0%; and (iv) the proposed framework outperforms all baselines, reaching 95.8% accuracy, 94.2% precision, 93.6% recall, and 93.9% F1.

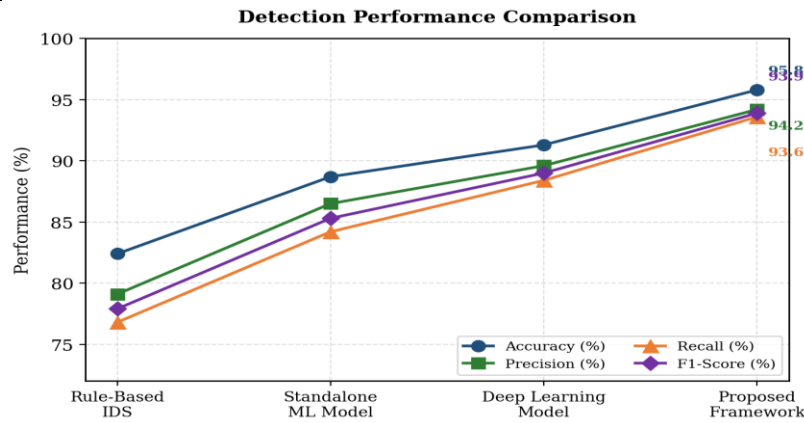


Fig. 2. Detection performance comparison across rule-based, ML, deep-learning, and proposed frameworks.

B. Threat Response Efficiency

Beyond detection quality, operational efficiency is equally consequential in healthcare cybersecurity. Table II summarizes response-time gains. By combining real-time correlation, threat scoring, and orchestrated automation, the proposed framework substantially shortens both Mean Time to Detect (MTTD) and Mean Time to Respond (MTTR). These reductions are decisive for limiting service disruption and protecting patient information once an intrusion attempt is in motion.

Approach	MTTD (minutes)	MTTR (minutes)
Rule-Based Security	28.5	64.2
Conventional SIEM + ML	18.7	41.6
Proposed Framework	7.9	19.4

TABLE II. RESPONSE TIME COMPARISON

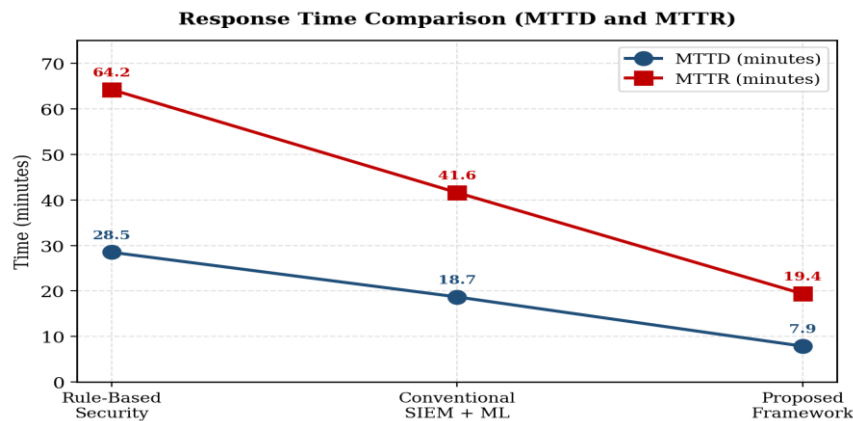


Fig. 3. MTTD and MTTR comparison across rule-based, conventional SIEM + ML, and proposed frameworks.

C. Comprehensive Comparison

Table III consolidates the comparison along multiple operational dimensions, including false-positive rate, scalability, adaptability to new threats, and regulatory alignment. The proposed framework records a very low false-positive rate alongside very high scalability, adaptability, and compliance-readiness ratings—properties that are jointly difficult for rule-based or single-model approaches to deliver.

Method	Acc. (%)	FPR	MTTD/MTTR (min)	S/A/C
Rule-Based IDS	82.4	High	28.5 / 64.2	L/L/L
Standalone ML	88.7	Med.	18.7 / 41.6	M/M/M
Deep Learning	91.3	Low	12.4 / 28.7	H/H/H
Proposed	95.8	V. Low	7.9 / 19.4	VH/VH/VH

Table III. Comparison With Existing Methods

D. Discussion

The experimental evidence indicates that the proposed framework strengthens proactive cyber defense by uniting adaptive learning, threat intelligence correlation, and explainable decision-making in a single pipeline. The combination of higher detection rates and shorter response times suggests that the design is well-matched to the layered, fast-changing threat landscape facing healthcare information systems. Adaptive learning further reduces operational dependence on manual rule updates—a longstanding limitation of conventional defenses. Taken together, the results support the position that an AI- and intelligence-centric architecture is more appropriate for safeguarding modern healthcare environments while preserving availability, regulatory standing, and patient safety.

V. CONCLUSION AND FUTURE SCOPE

This paper introduced an Adaptive AI-Driven Threat Intelligence Framework targeted at proactive cyber defense in dynamic healthcare information systems. By unifying supervised learning, unsupervised anomaly detection, deep-learning behavioral analysis, and real-time threat intelligence correlation, the framework directly addresses the limitations of reactive, fragmented defenses commonly deployed in clinical environments. Experimental results demonstrate substantial gains in detection accuracy, marked reductions in false positives, and meaningful improvements in mean time to detect and respond. Together, these outcomes confirm the framework's capacity to protect sensitive healthcare data, sustain system availability, and resist sophisticated adversarial campaigns—including ransomware and APT activity—while continuously adapting to drifting threat trends and preserving operational continuity and regulatory compliance.

Future work will pursue several extensions. First, federated learning is a natural next step, allowing healthcare institutions to share threat intelligence collaboratively without exposing patient data. Second, deeper integration of explainable-AI methods will further improve transparency in security decision-making and audit reporting. Third, large-scale clinical pilot deployments will be undertaken to validate the framework under realistic operational loads. Finally, blockchain-backed audit trails and automated compliance monitoring offer promising paths to strengthen accountability and trust within healthcare cybersecurity programs.

REFERENCES

1. M. Dekker and L. Alevizos, "A threat-intelligence driven methodology to incorporate uncertainty in cyber risk analysis and enhance decision-making," *Security and Privacy*, vol. 7, no. 1, pp. 1–18, 2024, doi: 10.1002/spy2.333.
2. S. Hansen and A. J. Baroody, "Beyond the boundaries of care: Electronic health records and the changing practices of healthcare," *Information and Organization*, vol. 33, no. 3, p. 100477, 2023.
3. Z. Almahmoud, P. D. Yoo, O. Alhusein, I. Farhat, and E. Damiani, "A holistic and proactive approach to forecasting cyber threats," *Scientific Reports*, vol. 13, no. 1, pp. 1–15, 2023, doi: 10.1038/s41598-023-35198-1.
4. A. R. Iossifova and S. Meyer-Goldstein, "Impact of standards adoption on healthcare transaction performance: The case of HIPAA," *International Journal of Production Economics*, vol. 141, no. 1, pp. 277–285, 2013.
5. J. Xing and Z. Zhang, "Hierarchical network security measurement and optimal proactive defense in cloud computing environments," *Security and Communication Networks*, vol. 2022, pp. 1–12, 2022, doi: 10.1155/2022/6783223.
6. M. Tahmasebi, "Beyond defense: Proactive approaches to disaster recovery and threat intelligence in modern enterprises," *Journal of Information Security*, vol. 15, no. 2, pp. 106–133, 2024, doi: 10.4236/jis.2024.152008.
7. Y. Creado and V. Ramteke, "Active cyber defence strategies and techniques for banks and financial institutions," *Journal of Financial Crime*, vol. 27, no. 3, pp. 771–780, 2020, doi: 10.1108/JFC-01-2020-0008.
8. N. U. I. Hossain, S. Rahman, and S. A. Liza, "Cyber-susiliency index: A comprehensive resiliency-sustainability-cybersecurity index for healthcare supply chain networks," *Decision Analytics Journal*, vol. 9, p. 100319, 2023.
9. G. Apruzzese, P. Laskov, E. Montes De Oca, W. Mallouli, L. Brdalo Rapa, A. V. Grammatopoulos, and F. Di Franco, "The role of machine learning in cybersecurity," *Digital Threats: Research and Practice*, vol. 4, no. 1, pp. 1–38, 2023, doi: 10.1145/3545574.
10. Y. Zheng, Z. Li, X. Xu, and Q. Zhao, "Dynamic defenses in cyber security: Techniques, methods, and challenges," *Digital Communications and Networks*, vol. 8, no. 4, pp. 422–435, 2022, doi: 10.1016/j.dcan.2021.07.006.
11. H. I. Kure, S. Islam, M. Ghazanfar, A. Raza, and M. Pasha, "Asset criticality and risk prediction for an effective cybersecurity risk management of cyber-physical systems," *Neural Computing and Applications*, vol. 34, no. 1, pp. 493–514, 2022, doi: 10.1007/s00521-021-06400-0.
12. A. Yeboah-Ofori, S. Islam, S. W. Lee, Z. U. Shamszaman, K. Muhammad, M. Altaf, and M. S. Al-Rakhami, "Cyber threat predictive analytics for improving cyber supply chain security," *IEEE Access*, vol. 9, pp. 94318–94337, 2021, doi: 10.1109/ACCESS.2021.3087109.
13. Y. Kim, I. Lee, H. Kwon, K. Lee, and J. Yoon, "Ban: Predicting APT attack based on Bayesian network with MITRE ATT&CK framework," *IEEE Access*, vol. 11, pp. 91949–91968, 2023, doi: 10.1109/ACCESS.2023.3306593.
14. N. Thapa, Z. Liu, A. Shaver, A. Esterline, B. Gokaraju, and K. Roy, "Secure cyber defense: An analysis of network intrusion-based dataset CCD-IDSv1 with machine learning and deep learning models," *Electronics*, vol. 10, no. 15, pp. 1–13, 2021, doi: 10.3390/electronics10151747.
15. M. Javaid, A. Haleem, R. P. Singh, and R. Suman, "Towards insighting cybersecurity for healthcare domains: A comprehensive review of recent practices and trends," *Cyber Security and Applications*, vol. 1, p. 100016, 2023.