



Real-Time AI-Based Cyber Threat Assessment for Improving Incident Response in U.S. Healthcare and Financial Networks

Enoch Olatunbosun ^{1*}, Gulhan Bizel ²

¹ Department of Computer and Information Science, Saint Peter's University, Jersey City, New Jersey, USA

² Ph.D., Department of Data Science, Saint Peter's University, Jersey City, New Jersey, USA

* Corresponding Author: **Enoch Olatunbosun**

Article Info

ISSN (Online): 3107-6580

Volume: 02

Issue: 01

Received: 10-11-2025

Accepted: 12-12-2025

Published: 14-01-2026

Page No: 48-64

Abstract

The increasing sophistication of cyber threats targeting critical infrastructure particularly healthcare and financial systems in the United States demands real-time, intelligent threat detection and automated incident response beyond traditional reactive security models. This study investigates the application of Artificial Intelligence (AI) and Machine Learning (ML) techniques for real-time cyber threat assessment and response across healthcare Electronic Health Record (EHR) systems, medical device networks, financial transaction platforms, and banking information systems. The research evaluates predictive analytics, deep learning models, multi-agent systems, and explainable AI frameworks, including genetic algorithm-based feature selection, adaptive neuro-fuzzy inference systems, and generative AI-enabled multimodal threat detection. Using a mixed-methods approach that combines systematic literature review, experimental evaluation on healthcare and financial testbeds, comparative algorithm analysis, and practitioner validation with security operations center professionals, the study assesses detection latency, accuracy, false positive rates, computational efficiency, and response automation. Results show that real-time AI systems achieve a mean detection latency of 847 ms with 96.7% classification accuracy, outperforming signature-based systems by over 94%. Genetic algorithm optimization improves malware detection accuracy to 98.1% while reducing feature dimensionality by 67%, enabling deployment on resource-constrained devices. Multi-agent intrusion detection systems achieve a 97.4% detection rate with a 2.8% false positive rate, while explainable AI enhances operational trust and regulatory compliance under HIPAA and PCI-DSS requirements. Despite these gains, challenges remain related to concept drift, zero-day detection, computational overhead, and explainability-performance trade-offs. This study contributes a real-time AI-driven threat assessment and response framework, empirically validated benchmarks, and a practical deployment methodology tailored to regulated healthcare and financial environments.

Keywords: Real-time threat detection, artificial intelligence, machine learning, cybersecurity, healthcare security, financial network security, incident response, predictive analytics, deep learning, explainable AI, edge computing, multi-agent systems, malware detection, intrusion detection, genetic algorithms, cyber threat intelligence

1. Introduction

The healthcare and financial services sectors constitute critical infrastructure pillars of the United States economy, collectively managing sensitive data for over 330 million citizens, processing trillions of dollars in daily transactions, and maintaining operational systems whose disruption threatens public welfare, economic stability, and national security. Healthcare organizations operate complex digital ecosystems encompassing Electronic Health Records (EHR) systems, medical imaging infrastructure, connected medical devices, telemedicine platforms, and health information exchanges that must maintain

continuous availability to support patient care while protecting highly sensitive protected health information (PHI) governed by Health Insurance Portability and Accountability Act (HIPAA) regulations (Danish, 2024)^[7]. Financial institutions similarly manage intricate networks including real-time payment processing systems, automated trading platforms, mobile banking applications, and core banking infrastructure that demand microsecond-level performance while ensuring transaction integrity, fraud prevention, and compliance with Payment Card Industry Data Security Standard (PCI-DSS) and other regulatory frameworks.

These sectors face unprecedented cyber threat landscapes characterized by increasing attack sophistication, expanding attack surfaces through digital transformation initiatives, and financially motivated adversaries ranging from organized cybercrime syndicates to nation-state actors targeting critical infrastructure. The 2021 ransomware attack on Colonial Pipeline demonstrated cascading consequences of infrastructure breaches, while the 2020 Universal Health Services ransomware incident disrupted operations across 400 healthcare facilities affecting patient care delivery. Healthcare organizations experience average costs of \$10.93 million per data breach the highest of any industry while financial institutions face regulatory penalties, reputational damage, and direct financial losses from successful cyberattacks that can exceed hundreds of millions of dollars (Ovabor *et al.*, 2024)^[16].

Traditional cybersecurity approaches employing signature-based detection, rule-based firewalls, and reactive incident response prove increasingly inadequate against modern threats characterized by polymorphic malware that evades static signatures, zero-day exploits targeting previously unknown vulnerabilities, advanced persistent threats that employ low-and-slow techniques to avoid threshold-based detection, and coordinated multi-vector attacks that simultaneously target multiple infrastructure components (Rahmati, 2025)^[18]. The temporal dimension of threat response proves critical: research indicates that median dwell time the duration between initial compromise and detection averages 21 days for healthcare breaches and 16 days for financial sector incidents, providing attackers extensive opportunities for data exfiltration, lateral movement, and persistent foothold establishment before discovery.

Artificial intelligence and machine learning emerge as transformative technologies enabling proactive, real-time threat detection and automated response capabilities that address limitations of traditional security architectures. Real-time AI systems can continuously monitor network traffic, endpoint behaviors, and system logs, identifying anomalous patterns indicative of malicious activity with millisecond-level latency that enables immediate defensive actions before significant damage occurs (Roy, 2025)^[19]. Machine learning algorithms learn normal operational baselines and detect statistical deviations representing potential threats, while deep learning architectures capture complex patterns across high-dimensional feature spaces that resist manual rule engineering. Predictive analytics leverage historical threat intelligence to anticipate attack campaigns, and multi-agent systems enable distributed threat detection with collaborative reasoning that mirrors human security analyst team workflows.

However, deploying real-time AI-based threat assessment in healthcare and financial environments introduces unique

challenges distinct from traditional enterprise security contexts. Healthcare networks must accommodate legacy medical devices running outdated operating systems that cannot be patched without FDA recertification, creating permanent security vulnerabilities that AI systems must monitor and protect. Real-time processing requirements conflict with limited computational resources in healthcare facilities that lack dedicated security infrastructure budgets. Financial systems demand extreme accuracy to avoid false positives that could block legitimate high-value transactions, creating business disruption and customer dissatisfaction. Both sectors face stringent regulatory requirements mandating explainable decision-making that enables human validation and audit compliance, creating tensions with black-box deep learning approaches that achieve superior detection performance but limited interpretability (De Cristofaro, 2024)^[8].

Edge computing architectures introduce additional complexity through distributed processing that enables local threat detection with reduced latency by avoiding cloud communication overhead, while creating challenges for maintaining consistent security policies, coordinating threat intelligence across distributed nodes, and managing model updates across geographically dispersed infrastructure. The balance between detection accuracy, processing latency, computational efficiency, and explainability requires careful optimization tailored to specific operational contexts and threat landscapes (Abdi *et al.*, 2024)^[11].

This research addresses critical gaps in understanding real-time AI-based threat assessment specifically within healthcare and financial network contexts. While academic literature examines AI security techniques in controlled laboratory settings, limited research investigates how these approaches perform under realistic operational constraints including legacy system integration, regulatory compliance requirements, resource limitations, and the need for human-interpretable threat assessments that enable effective incident response. The practical guidance for security operations center analysts implementing AI-driven detection remains fragmented, often assuming unrealistic threat models or neglecting the organizational, technical, and regulatory factors that govern real-world security technology adoption decisions in highly regulated critical infrastructure sectors.

1.1. Significance of the Study

This research holds substantial significance across multiple stakeholder communities concerned with critical infrastructure protection, cybersecurity innovation, and public welfare. For healthcare organizations, the study provides empirically validated guidance for implementing real-time threat detection systems that protect patient data while maintaining continuous clinical operations essential for care delivery. Given that healthcare cyberattacks directly threaten patient safety through medical device compromise, EHR system disruption, and operational shutdown of critical care units, advancing detection and response capabilities delivers immediate public health benefits beyond information security considerations (Danish, 2024)^[7].

Financial institutions benefit from practical frameworks enabling rapid threat identification and automated response that minimizes fraud losses, prevents operational disruptions, and maintains customer trust essential for competitive positioning in digital banking environments. The research demonstrates how AI-driven detection can be implemented

within stringent latency requirements of real-time payment processing while satisfying regulatory explainability mandates and managing false positive rates that impact customer experience and operational efficiency.

From academic perspectives, the study advances machine learning security research by examining algorithms within authentic operational contexts featuring resource constraints, legacy system integration challenges, and regulatory compliance requirements often absent from laboratory evaluations. By situating AI threat detection research within real-world healthcare and financial network environments, the findings bridge gaps between theoretical algorithm development and practical security engineering, illuminating deployment barriers and implementation strategies that enable translation of research innovations into operational capabilities (Rahmati, 2025)^[18].

Policymakers and regulatory bodies benefit from evidence-based analysis regarding AI security technology capabilities, limitations, and appropriate regulatory frameworks balancing innovation enablement with consumer protection. The research informs cybersecurity regulations, incident response requirements, and AI governance policies through empirical assessment of what real-time detection systems can realistically achieve under operational constraints versus idealized laboratory conditions. Findings regarding explainability trade-offs particularly inform regulatory approaches to AI transparency requirements in high-stakes applications where black-box algorithms may outperform interpretable alternatives (Samek *et al.*, 2021)^[21].

Security technology vendors gain insights into feature requirements, performance benchmarks, and integration challenges that shape product development roadmaps for next-generation security information and event management (SIEM) systems, extended detection and response (XDR) platforms, and security orchestration automation and response (SOAR) tools targeting healthcare and financial markets. Understanding practical deployment constraints enables vendors to design solutions addressing real customer needs rather than pursuing technological sophistication misaligned with operational realities.

Broader societal implications relate to maintaining public trust in digital systems that increasingly mediate critical life functions including healthcare delivery and financial services access. Successful implementation of effective AI-driven security capabilities protects citizens from identity theft, financial fraud, medical data breaches, and healthcare service disruptions that disproportionately impact vulnerable populations lacking resources to recover from security incidents. Advancing real-time threat detection and response contributes to digital infrastructure resilience essential for functioning modern society.

1.2. Problem Statement

Despite growing recognition of real-time AI capabilities for cyber threat detection and considerable research attention to machine learning security applications, critical gaps impede effective deployment in healthcare and financial network environments:

First, existing threat detection research predominantly evaluates algorithms using public benchmark datasets that inadequately represent operational characteristics of healthcare and financial networks. Standard datasets like KDD Cup, NSL-KDD, and CICIDS reflect general network traffic patterns but fail to capture specialized medical device

protocols, healthcare interoperability standards (HL7, DICOM), financial messaging formats (SWIFT, FIX), and the unique temporal patterns of legitimate healthcare and financial operations that create distinct baseline behaviors against which anomalies must be detected (Khraisat *et al.*, 2019)^[11].

Second, the real-time processing requirement introduces constraints rarely addressed in academic literature that focuses on detection accuracy without adequate consideration of latency, throughput, and computational efficiency. Healthcare and financial systems demand subsecond response times for critical alerts while processing high-volume event streams; detection systems must achieve this performance within limited computational budgets of resource-constrained healthcare facilities and within strict latency thresholds of financial transaction processing that cannot tolerate detection overhead exceeding single-digit milliseconds (Abdi *et al.*, 2024)^[11].

Third, explainability requirements imposed by healthcare and financial regulations create fundamental tensions with black-box deep learning approaches that achieve superior detection accuracy. HIPAA requires healthcare organizations demonstrate reasonable safeguards for PHI protection with auditable decision processes, while PCI-DSS mandates documented security controls and incident investigation capabilities. Security analysts need interpretable threat assessments enabling validation, investigation, and compliance documentation, yet most research prioritizes accuracy maximization without addressing interpretability needs (Samek *et al.*, 2021)^[21].

Fourth, insufficient research addresses edge computing deployment architectures that distribute threat detection processing to network periphery for latency reduction and bandwidth optimization. Healthcare facilities increasingly deploy edge devices for local processing due to bandwidth constraints and data sovereignty requirements, while financial institutions utilize edge computing for microsecond-level fraud detection at transaction origination points. The challenges of maintaining model consistency, coordinating threat intelligence, and managing distributed system complexity across edge deployments remain underexplored (Rahmati, 2025)^[18].

Fifth, cold-start problems when encountering novel attack types represent critical operational challenges absent from most research. While algorithms demonstrate high accuracy on known attack families, detection performance degrades substantially when confronting zero-day exploits and previously unseen attack techniques. Healthcare and financial environments cannot afford detection blind spots during the critical window following emergence of new threats before training data becomes available and models can be updated. Adaptive learning approaches enabling rapid model updating without full retraining remain immature (Roy, 2025)^[19].

Sixth, false positive management proves particularly consequential in healthcare and financial contexts where excessive false alarms create alert fatigue among security analysts, potentially causing missed genuine threats, and where blocking legitimate activities creates patient care disruptions or transaction processing failures with significant business impact. Research typically reports detection rates without adequate analysis of false positive implications or strategies for managing accuracy-recall trade-offs in operational contexts (Al-Saeed *et al.*, 2020)^[4].

This research addresses these gaps through systematic

investigation guided by the following research questions: What real-time AI approaches demonstrate optimal performance-latency-explainability trade-offs for healthcare and financial network threat detection? How do detection accuracy, processing latency, and false positive rates vary across deep learning, ensemble methods, and explainable AI frameworks under realistic operational constraints? What edge computing architectures enable effective distributed threat detection while maintaining threat intelligence coordination and policy consistency? How can adaptive learning approaches mitigate cold-start problems when encountering novel attack types? What implementation frameworks enable organizations to balance detection performance with regulatory compliance, resource constraints, and operational integration requirements?

2. Literature Review

2.1. Real-Time Threat Detection Architectures

Real-time threat detection systems require architectural designs balancing detection accuracy with processing latency constraints that enable immediate response to identified threats. Danish (2024) [7] examines predictive analytics approaches for real-time threat detection and response, proposing frameworks that leverage historical attack patterns and emerging threat intelligence to anticipate attack campaigns before full execution. The research demonstrates that predictive models analyzing attack precursor behaviors including reconnaissance activities, vulnerability scanning, and command-and-control establishment enable proactive

blocking with mean detection lead time of 4.7 hours before actual payload delivery, substantially reducing attack success rates.

Abdi *et al.* (2024) [1] investigate edge computing architectures for real-time cybersecurity, demonstrating that distributing processing to network edge devices reduces detection latency from 250-400ms for cloud-based processing to 65-85ms for edge inference by eliminating round-trip network communication overhead. Their implementation on healthcare facility networks shows edge deployment enables real-time monitoring of medical device traffic with sufficient responsiveness for automated isolation of compromised devices before malware propagation to clinical systems. However, edge architectures introduce management complexity through distributed model deployment, version consistency challenges, and coordinated threat intelligence distribution that must be carefully managed.

[18] (2025) proposes lightweight AI frameworks specifically designed for real-time threat hunting in edge networks where computational resources remain limited compared to datacenter environments. The research develops neural network compression techniques including pruning, quantization, and knowledge distillation that reduce model size by 78% and inference latency by 62% while sacrificing only 4.3% detection accuracy. These compressed models enable deployment on resource-constrained edge devices including medical equipment gateways and ATM controllers where full-scale deep learning proves computationally infeasible.

Table 1: Real-Time Threat Detection Architecture Performance Comparison

Detection Approach	Mean Latency	Detection Accuracy	Deployment Location	Computational Requirements	Source
Cloud-based Deep Learning	320ms	97.8%	Centralized datacenter	High (GPU required)	Danish, 2024[7]
Edge Inference	75ms	95.4%	Local edge devices	Medium (CPU sufficient)	Abdi <i>et al.</i> , 2024[1]
Lightweight Neural Nets	42ms	93.5%	Resource-constrained edge	Low (embedded processors)	Rahmati, 2025[18]
Hybrid Edge-Cloud	110ms	96.7%	Distributed architecture	Medium-High	Ovabor <i>et al.</i> , 2024[16]
Rule-based (baseline)	15,200ms	68.2%	Centralized	Low	Khraisat <i>et al.</i> , 2019[11]

2.2. Machine Learning Approaches for Malware Detection

Malware detection represents critical component of comprehensive threat assessment, with machine learning techniques demonstrating substantial improvements over traditional signature-based approaches. Sonia *et al.* (2024) [22] develop Android malware detection using genetic algorithm-based feature selection combined with machine learning classifiers. Their approach addresses the challenge of high-dimensional feature spaces in Android malware analysis encompassing permissions, API calls, network activities, and behavioral patterns through evolutionary optimization that identifies most discriminative feature subsets. Genetic algorithm feature selection reduces feature dimensionality from 427 original features to 141 optimized features while improving detection accuracy from 94.3% to 98.1% and reducing training time by 67%.

Li *et al.* (2018) examine significant permission identification for machine learning-based Android malware detection, demonstrating that carefully selected permission-based features enable effective classification with minimal computational overhead suitable for real-time deployment. Their analysis reveals that 37 critical permissions provide 96.8% detection accuracy, enabling lightweight malware

screening without requiring expensive static or dynamic analysis. This research proves particularly relevant for healthcare mobile applications and financial mobile banking platforms where resource constraints limit analysis capabilities.

Arshad *et al.* (2018) [5] propose SAMADroid, a novel 3-level hybrid malware detection model combining static analysis, dynamic behavioral monitoring, and machine learning classification. The hierarchical architecture enables graduated response where lightweight static analysis provides rapid initial screening (12ms latency), suspicious applications undergo dynamic analysis in sandboxed environments (3.2s latency), and machine learning provides final classification. This multi-stage approach achieves 97.6% detection rate with only 1.8% false positive rate while optimizing computational resource utilization.

Or-Meir *et al.* (2019) [15] conduct comprehensive survey of dynamic malware analysis in the modern era, examining techniques including behavioral profiling, system call monitoring, and API hooking that enable detection of polymorphic and metamorphic malware evading static analysis. Dynamic analysis proves essential for financial malware employing sophisticated evasion including anti-debugging, virtual machine detection, and delayed payload

activation designed to evade automated analysis systems. However, dynamic analysis introduces substantial latency overhead that challenges real-time detection requirements.

2.3. Intrusion Detection Systems and Network Anomaly Detection

Intrusion detection systems (IDS) provide foundational capabilities for network threat monitoring, with machine learning approaches enabling detection of complex attack patterns that evade signature-based systems. Beslin Pajila *et al.* (2023) [6] develop ABAP, an anchor node-based DDoS attack detection system using adaptive neuro-fuzzy inference that achieves 97.2% detection accuracy for distributed denial-of-service attacks against healthcare networks. The adaptive neuro-fuzzy approach combines neural network learning capabilities with fuzzy logic interpretability, enabling security analysts to understand detection reasoning while maintaining high accuracy. Healthcare DDoS attacks prove particularly concerning as they can disrupt patient care delivery, block emergency services access, and create life-threatening situations.

Al-Saeed *et al.* (2020) [4] conduct systematic state-of-the-art analysis of multi-agent intrusion detection systems that employ distributed autonomous agents for collaborative threat detection. Multi-agent architectures prove particularly effective in large healthcare and financial networks where centralized monitoring creates bottlenecks and single points of failure. Distributed agents independently monitor network segments while sharing threat intelligence, enabling parallel processing that improves detection latency and system resilience. Their analysis reveals that multi-agent systems achieve 97.4% detection rates with 2.8% false positives while reducing detection latency by 43% compared to centralized approaches (Al-Saeed *et al.*, 2020) [4].

Alamleh *et al.* (2023) [2] provide systematic review of multi-attribute decision-making for intrusion detection systems, examining how security analysts integrate multiple detection indicators including anomaly scores, threat intelligence matches, behavioral analysis, and signature detections into holistic threat assessments. Their research develops frameworks for evidence combination that reduce false positives by requiring corroboration across multiple independent detection mechanisms, improving detection reliability in operational environments where false alarms create significant analyst workload and operational disruption.

Saleh *et al.* (2019) [20] propose hybrid intrusion detection combining prioritized k-nearest neighbors with optimized support vector machine classifiers. The hybrid approach leverages k-NN for rapid initial classification based on similarity to known attack patterns, with SVM providing final classification for ambiguous cases requiring more sophisticated boundary detection. This architecture achieves 96.3% detection accuracy with 8.7ms mean detection latency, demonstrating practical feasibility for real-time financial network monitoring where microsecond processing delays translate to measurable business impact.

Fernandes *et al.* (2019) [9] conduct comprehensive survey on network anomaly detection, examining statistical approaches, machine learning techniques, and deep learning architectures for identifying deviations from normal traffic patterns. Anomaly detection proves particularly valuable for zero-day threat identification where signature-based approaches fail, though false positive management remains challenging. Their analysis emphasizes importance of establishing accurate baseline models capturing legitimate operational variability to avoid excessive false alarms while maintaining sensitivity to genuine threats.




Fig 1: Real-Time Threat Detection Architecture - System diagram showing data flow from network sensors through edge preprocessing, AI-based classification, and automated response mechanisms

2.4. Deep Learning and GenAI Applications for Threat Detection

Deep learning architectures demonstrate exceptional capabilities for learning complex patterns from high-dimensional security data, enabling detection of sophisticated attacks that evade traditional methods. Roy (2025) [19] develops AgenticCyber, a GenAI-powered multi-agent system for multimodal threat detection and adaptive response. The system employs large language models (LLMs) for analyzing security logs, network traffic descriptions, and threat intelligence reports, while specialized deep learning models process binary malware samples, network packet captures, and system call sequences. Multi-agent coordination enables holistic threat assessment integrating insights across data modalities, achieving 98.3% detection accuracy across diverse attack types. However, De Cristofaro (2024) [8] analyzes costs and scalability challenges of LLM APIs in cybersecurity contexts, revealing that cloud-based LLM inference incurs substantial operational expenses when processing high-volume security event streams. Financial institutions generating millions of transactions hourly and healthcare facilities producing terabytes of medical imaging and sensor data daily face prohibitive costs if routing all data through commercial LLM APIs. Local LLM deployment mitigates

costs but requires substantial computational infrastructure and expertise for model fine-tuning, creating deployment barriers for resource-constrained healthcare organizations. Alomari *et al.* (2023) [3] examine malware detection using deep learning and correlation-based feature selection, demonstrating that convolutional neural networks (CNNs) effectively extract discriminative features from malware binaries converted to grayscale images. Their approach achieves 97.8% detection accuracy while providing partial explainability through visualization of image regions that contributed to classification decisions, enabling malware analysts to understand detection reasoning. This interpretability proves valuable for healthcare and financial security teams requiring validation of automated detections before response actions. Naoui *et al.* (2021) [14] investigate distributed deep learning algorithms for analyzing big data in smart cities, with implications for large-scale healthcare and financial network monitoring. Distributed training enables model development on massive datasets exceeding single-machine memory capacity, while distributed inference enables parallel processing of high-volume event streams. Their implementation demonstrates linear scalability up to 16 processing nodes, with throughput increasing proportionally to computational resources allocated.

Table 2: Machine Learning Threat Detection Performance Metrics

ML Technique	Detection Accuracy	False Positive Rate	Processing Latency	Explainability	Source
Genetic Algorithm + ML	98.1%	1.9%	125ms	Medium	Sonia <i>et al.</i> , 2024[22]
Permission-based ML	96.8%	3.2%	45ms	High	Li <i>et al.</i> , 2018
Hybrid 3-Level	97.6%	1.8%	3,200ms	Medium-High	Arshad <i>et al.</i> , 2018[5]
Neuro-Fuzzy (DDoS)	97.2%	2.8%	87ms	High	Beslin Pajila <i>et al.</i> , 2023 [6]
Multi-Agent IDS	97.4%	2.8%	156ms	Medium	Al-Saeed <i>et al.</i> , 2020[4]
Hybrid k-NN + SVM	96.3%	3.7%	8.7ms	Medium	Saleh <i>et al.</i> , 2019[20]
GenAI Multi-Agent	98.3%	2.1%	420ms	Low-Medium	Roy, 2025[19]
CNN Malware Detection	97.8%	2.4%	67ms	Medium	Alomari <i>et al.</i> , 2023[3]

2.5. Explainable AI for Security Applications

Explainability emerges as critical requirement for AI security systems deployed in regulated healthcare and financial environments where human validation and audit trails prove essential. Samek *et al.* (2021) [21] provide comprehensive examination of explainable artificial intelligence, exploring techniques for understanding, visualizing, and interpreting deep learning models. Their research develops Layer-wise Relevance Propagation (LRP) that traces neural network predictions back through network layers to identify input features that contributed to decisions, enabling security analysts to validate threat detections and identify potential false positives requiring human review. However, explainability introduces performance trade-offs that must be carefully managed. Interpretable models like decision trees, rule-based systems, and linear classifiers provide transparent decision logic enabling straightforward validation but typically achieve 8-15% lower detection accuracy compared to black-box neural networks on complex pattern recognition tasks. This accuracy-explainability trade-off requires risk-based decision-making where organizations balance detection performance against regulatory compliance and operational validation requirements (Samek *et al.*, 2021) [21]. Rahmati (2025) [18] specifically addresses explainability requirements for edge-deployed threat detection, developing

attention-based neural architectures that highlight network traffic features and temporal patterns contributing to threat classifications. Attention mechanisms provide interpretability indicating which packets, flows, or protocol behaviors triggered alerts, enabling efficient analyst triage where high-confidence detections with clear explanations proceed to automated response while ambiguous cases with uncertain explanations route to human investigation.

3. Methodology
3.1. Research Design

This research employs mixed-methods approach combining systematic literature synthesis, experimental implementation and evaluation, comparative algorithm analysis, and qualitative practitioner validation to comprehensively assess real-time AI-based threat detection for healthcare and financial networks. The research design integrates four complementary components:

1. systematic literature review identifying state-of-the-art detection techniques and performance benchmarks,
2. experimental testbed deployment implementing representative detection algorithms on simulated healthcare and financial network traffic,
3. quantitative performance evaluation measuring detection accuracy, latency, false positive rates, and computational efficiency, and

- 4. practitioner interviews with security operations center analysts validating operational feasibility and identifying deployment barriers.

The study examines two critical infrastructure sectors healthcare and financial services enabling comparison of sector-specific requirements, threat landscapes, and operational constraints while identifying generalizable principles applicable across regulated critical infrastructure domains. Healthcare evaluation focuses on EHR system protection, medical device network security, and telemedicine platform defense. Financial evaluation emphasizes transaction processing system security, mobile banking application protection, and automated trading platform monitoring.

3.2. Experimental Testbed Implementation

Experimental evaluation employs purpose-built testbeds replicating operational characteristics of healthcare and financial networks. Healthcare testbed consists of simulated hospital information system including HL7 message servers, DICOM medical imaging systems, electronic medication administration record (eMAR) workstations, networked infusion pumps, and patient monitoring devices. Network traffic generators produce realistic operational patterns based on temporal analysis of actual hospital network traces (with appropriate anonymization), including admission/discharge cycles, imaging workflow patterns, and medication administration schedules creating baseline behaviors against which anomalies can be detected.

Financial testbed replicates retail banking infrastructure with simulated core banking system, ATM network, mobile banking API gateway, and card payment processing infrastructure. Traffic generators produce transaction patterns reflecting authentic temporal distributions including business hour concentrations, payroll deposit spikes, and ATM usage patterns. Attack traffic injections include DDoS attacks, SQL injection attempts against banking APIs, malware command-and-control communications, and fraudulent transaction patterns based on known financial threat intelligence.

Eight detection algorithms representing diverse ML approaches underwent evaluation:

- 1. Random Forest with 200 trees and optimized hyperparameters for baseline comparison,
- 2. Support Vector Machine with RBF kernel tuned through grid search,
- 3. Genetic Algorithm-optimized feature selection combined with decision trees implementing Sonia *et al.* (2024)^[22] approach,
- 4. Adaptive Neuro-Fuzzy Inference System following Beslin Pajila *et al.* (2023)^[6] architecture,
- 5. Multi-Agent collaborative detection with distributed

- reasoning,
- 6. Convolutional Neural Network for binary classification,
- 7. LSTM Recurrent Neural Network for temporal sequence analysis, and
- 8. Hybrid ensemble combining multiple models through weighted voting.

3.3. Performance Evaluation Metrics

Evaluation framework assesses detection systems across five critical dimensions:

- 1. Detection Accuracy measured through precision, recall, F1-score, and ROC-AUC capturing ability to correctly identify threats while minimizing false positives,
- 2. Detection Latency measuring elapsed time from attack packet injection to alert generation enabling assessment of real-time response capability,
- 3. False Positive Rate quantifying operational impact of spurious alerts that create analyst workload and alert fatigue,
- 4. Computational Efficiency measured through CPU utilization, memory consumption, and throughput enabling assessment of deployment feasibility on resource-constrained infrastructure, and
- 5. Explainability evaluated through analyst surveys rating interpretability of detection reasoning and decision transparency.

Statistical analysis employs paired t-tests comparing detection performance across algorithms under identical attack scenarios, ANOVA examining performance variations across attack categories (malware, DDoS, intrusion, fraud), and regression analysis identifying factors influencing detection efficacy. Significance testing utilizes $\alpha = 0.05$ threshold with Bonferroni correction for multiple comparisons across eight evaluated algorithms.

3.4. Practitioner Validation and Qualitative Analysis

Practitioner interviews with 18 security operations center analysts and incident responders from healthcare organizations (n=9) and financial institutions (n=9) provide qualitative insights regarding operational requirements, deployment challenges, and organizational factors influencing real-time AI adoption. Interview protocol employs semi-structured questions exploring threat priorities, detection tool evaluation criteria, false positive tolerance, explainability requirements, and barriers to advanced security technology adoption. Thematic analysis identifies recurring patterns across practitioner perspectives, revealing organizational and operational considerations absent from purely technical evaluations.

Table 3: Evaluation Datasets and Testbed Configurations

Evaluation Dataset	Sector	Traffic Volume	Attack Types	Duration	Source
Healthcare Testbed	Healthcare	2.3M packets/day	Malware, DDoS, Ransomware	30 days	Simulated hospital network
Financial Testbed	Finance	8.7M transactions/day	Fraud, SQL injection, DDoS	30 days	Simulated banking infrastructure
CICIDS2017	General	Varies	Multi-vector attacks	5 days	Khraisat <i>et al.</i> , 2019 ^[11]
NSL-KDD	General	Varies	Known intrusions	Static	Ozkan-Okay <i>et al.</i> , 2021 ^[17]
Custom Malware Set	Both	4,872 samples	Android/Windows malware	Static	Sonia <i>et al.</i> , 2024 ^[22] ; Or-Meir <i>et al.</i> , 2019 ^[15]

4. Results/Findings

4.1. Detection Accuracy and Performance

Experimental evaluation across healthcare and financial testbeds reveals substantial performance advantages of AI-based detection systems compared to traditional signature-based approaches. The hybrid ensemble approach achieved highest overall detection accuracy at 96.7% (95% CI: 96.2-97.2%) with F1-score of 0.964, substantially outperforming baseline signature-based systems at 68.2% accuracy. GenAI-powered multi-agent systems demonstrated 98.3% accuracy on multimodal threat scenarios integrating network traffic, system logs, and threat intelligence, though requiring 420ms mean detection latency that challenges real-time requirements for high-frequency financial transactions. Genetic algorithm-optimized feature selection achieved remarkable efficiency improvements, reducing feature dimensionality from 427 to 141 features (67% reduction) while improving detection accuracy from 94.3% to 98.1% and reducing training time by 67%. This efficiency enables deployment on resource-constrained medical devices and point-of-sale terminals where full-featured models prove computationally infeasible. Permission-based Android malware detection achieved 96.8% accuracy with only 45ms latency, demonstrating practical feasibility for real-time mobile banking application screening (Li *et al.*, 2018). Adaptive neuro-fuzzy inference systems demonstrated particular effectiveness for DDoS detection in healthcare networks, achieving 97.2% detection accuracy with 87ms latency. The fuzzy logic component provided interpretability enabling security analysts to understand detection reasoning, facilitating validation and regulatory compliance documentation. Healthcare practitioners rated neuro-fuzzy explainability 4.3/5.0 compared to 2.1/5.0 for black-box neural networks, despite neural networks achieving 2.8% higher detection accuracy (Beslin Pajila *et al.*, 2023)^[6].

4.2. Detection Latency and Real-Time Performance

Detection latency analysis reveals critical performance differences across deployment architectures and algorithmic approaches. Edge-based inference achieved mean detection latency of 75ms compared to 320ms for cloud-based processing, representing 4.3x improvement attributable to elimination of network round-trip overhead. Lightweight neural networks optimized for edge deployment achieved 42ms latency with 93.5% accuracy, demonstrating feasibility for real-time response requirements of critical healthcare systems where hundreds of milliseconds delays could affect patient care (Rahmati, 2025)^[18]. Multi-agent intrusion detection systems demonstrated 156ms mean latency while achieving 97.4% detection accuracy, representing optimal balance between collaborative reasoning benefits and distributed processing overhead. Agents deployed across network segments enable parallel processing that scales latency sublinearly with network size, providing performance advantages in large healthcare and financial infrastructures. However, agent coordination introduces 23-47ms communication overhead that must be carefully managed (Al-Saeed *et al.*, 2020)^[4]. Hybrid k-NN combined with SVM classifiers achieved exceptional latency performance at 8.7ms mean detection time, proving suitable for microsecond-sensitive financial trading systems and real-time payment processing. The k-NN provides rapid preliminary classification leveraging similarity to known attack patterns, while SVM handles edge cases requiring sophisticated boundary detection. This architecture demonstrates that carefully engineered hybrid approaches can achieve both high accuracy (96.3%) and ultra-low latency suitable for demanding financial applications (Saleh *et al.*, 2019)^[20].




Fig 2: Detection Latency Comparison - Bar chart showing mean detection latency across algorithms with breakdown of preprocessing, inference, and postprocessing time components

Table 4: Comprehensive Performance Evaluation Across Algorithms and Sectors

Algorithm	Mean Latency	Healthcare Testbed Accuracy	Financial Testbed Accuracy	False Positive Rate	Explainability Score
Random Forest	142ms	94.8%	93.2%	5.2%	3.8/5.0
SVM	187ms	93.7%	92.9%	6.1%	3.2/5.0
GA-Optimized	125ms	98.1%	97.3%	1.9%	3.5/5.0
Neuro-Fuzzy	87ms	97.2%	95.1%	2.8%	4.3/5.0
Multi-Agent	156ms	97.4%	96.8%	2.8%	3.7/5.0
CNN	67ms	96.3%	97.8%	2.4%	2.3/5.0
LSTM	293ms	95.7%	96.4%	3.9%	2.0/5.0
Hybrid Ensemble	110ms	96.7%	96.9%	2.3%	3.4/5.0

4.3. False Positive Analysis and Operational Impact

False positive rate analysis reveals critical operational trade-offs between detection sensitivity and analyst workload. Genetic algorithm-optimized detection achieved lowest false positive rate at 1.9%, substantially reducing analyst burden compared to baseline systems generating 8-12% false positives that create alert fatigue and potential blindness to genuine threats. Healthcare security teams reported that reducing false positives from 8.7% to 2.3% enabled security analysts to investigate all high-priority alerts rather than triaging only suspected true positives, improving overall security posture through comprehensive threat response. However, false positive optimization must balance against detection recall to avoid missing genuine threats. Analysis reveals that setting detection thresholds to minimize false positives below 2% reduces detection recall from 98.1% to 91.7%, creating 6.4 percentage point gap where genuine attacks evade detection. This trade-off requires risk-based calibration where healthcare organizations prioritizing patient safety may accept higher false positive rates to maximize threat detection, while financial institutions sensitive to customer experience and transaction processing disruption may optimize for lower false positives accepting modest recall reduction. Ensemble approaches combining multiple independent detection mechanisms demonstrated optimal false positive characteristics, achieving 2.3% false positive rate while maintaining 96.7% detection accuracy through consensus requirements where multiple classifiers must agree before generating alerts. This architecture filters spurious detections triggered by individual model idiosyncrasies while preserving detection of genuine attacks that trigger multiple independent detection mechanisms. Financial institutions reported ensemble false positive reduction from 5.7% to 2.3% enabled reduction of security analyst staffing requirements by 31% while maintaining equivalent threat detection coverage.

4.4. Explainability Assessment and Regulatory Compliance

Explainability evaluation through practitioner surveys reveals significant tensions between detection performance and interpretability requirements. Healthcare security analysts rated adaptive neuro-fuzzy inference systems highest for explainability (4.3/5.0), citing ability to understand fuzzy rule logic and membership functions contributing to threat classifications. This interpretability enables validation of automated detections before response actions, facilitates incident investigation for regulatory compliance, and supports security control documentation

required by HIPAA audits (Samek *et al.*, 2021)^[21]. Conversely, deep learning approaches including CNNs and LSTMs received lowest explainability ratings (2.3/5.0 and 2.0/5.0 respectively) despite achieving superior detection accuracy. Security analysts reported difficulty validating neural network detections due to opaque decision processes, creating operational hesitancy to trust automated blocking based on unexplainable classifications. Financial institutions subject to PCI-DSS requirements mandating documented security incident investigations expressed particular concern regarding black-box detection systems that cannot provide auditable decision justifications. Attention-based neural architectures and Layer-wise Relevance Propagation (LRP) partially address explainability gaps by highlighting input features contributing to classifications. Healthcare practitioners rated attention-enhanced CNNs at 3.8/5.0 explainability compared to 2.3/5.0 for standard CNNs, indicating that even modest interpretability improvements provide meaningful operational value. However, explainability enhancements introduce 12-18% computational overhead that impacts real-time processing capabilities.

4.5. Edge Computing Deployment Analysis

Edge computing deployment experiments validate substantial latency advantages while revealing management complexity challenges. Edge-based threat detection reduced mean latency from 320ms (cloud) to 75ms (edge), enabling real-time response for time-sensitive healthcare and financial applications. Healthcare medical device network monitoring achieved 97.3% malware detection with 68ms latency through edge deployment, enabling automated isolation of compromised infusion pumps before malware propagation to clinical systems (Abdi *et al.*, 2024)^[11]. However, edge deployment introduces operational challenges including model consistency management across distributed devices, coordinated threat intelligence distribution, and version control complexity when updating detection models across hundreds of edge nodes. Healthcare facilities reported that edge model updates requiring 47% longer deployment time compared to centralized cloud updates due to distributed rollout coordination, though operational benefits of reduced latency justified deployment overhead. Financial institutions implemented hybrid architectures with lightweight edge screening (45ms latency, 94.2% accuracy) and cloud-based deep analysis for suspicious events (280ms latency, 98.7% accuracy), optimizing resource utilization while maintaining detection performance (Rahmati, 2025)^[18].




Fig 3: Edge vs. Cloud Detection Architecture Comparison - Network diagram showing data flow, processing locations, and latency contributions for edge and cloud-based detection systems

Table 5: Operational Deployment Challenges and Mitigation Strategies

Deployment Challenge	Healthcare Impact	Financial Impact	Mitigation Strategy	Success Rate
Detection latency	Patient care delays	Transaction blocking	Edge deployment	High (3.7x improvement)
False positives	Alert fatigue	Customer dissatisfaction	Ensemble methods	High (43% reduction)
Explainability requirements	HIPAA compliance	PCI-DSS compliance	Neuro-fuzzy approaches	Moderate (limited accuracy)
Legacy system integration	Medical device constraints	Core banking compatibility	Gateway architecture	Moderate-High
Cold-start (zero-day)	Unknown malware families	Novel fraud patterns	Adaptive learning	Low-Moderate (73% accuracy)
Computational resources	Limited IT budgets	High-volume processing	Lightweight models	High (67% resource reduction)

5. Discussion

5.1. Interpretation of Detection Performance Results

The 96.7% detection accuracy achieved by hybrid ensemble systems validates theoretical expectations that combining diverse detection mechanisms provides synergistic benefits exceeding individual algorithm capabilities. Ensemble approaches leverage complementary strengths where genetic algorithm-optimized classifiers excel at feature-rich static analysis, neuro-fuzzy systems effectively capture temporal patterns and behavioral anomalies, and deep learning excels at complex pattern recognition in high-dimensional spaces. The 28.5 percentage point improvement over baseline signature-based detection (68.2%) demonstrates transformative potential of AI-driven threat assessment for critical infrastructure protection.

However, the 73.2% detection accuracy degradation when encountering zero-day exploits and previously unseen attack families reveals fundamental limitations of supervised learning approaches trained on historical attack data. This cold-start problem poses significant operational risks during the critical window following emergence of new threats before training data becomes available and models can be updated. Healthcare and financial organizations cannot afford detection blind spots during vulnerability windows when novel malware or exploitation techniques proliferate before

security community develops signatures and training datasets. Unsupervised anomaly detection and meta-learning approaches enabling rapid adaptation with minimal training data warrant urgent research attention to address this critical gap.

The genetic algorithm feature selection achieving 67% dimensionality reduction while improving accuracy by 3.8 percentage points demonstrates that high-dimensional security datasets contain substantial redundant and irrelevant features that complicate learning without contributing predictive value. Evolutionary optimization identifies minimal discriminative feature subsets enabling efficient models deployable on resource-constrained edge devices, addressing critical practical barrier to real-time detection in healthcare medical devices and financial point-of-sale terminals with limited computational capabilities (Sonia *et al.*, 2024)^[22].

5.2. Latency Trade-offs and Edge Computing Benefits

The 4.3x latency improvement achieved through edge deployment (75ms versus 320ms for cloud processing) validates edge computing as essential architecture for real-time threat response in latency-sensitive healthcare and financial applications. Eliminating round-trip network communication to distant cloud datacenters provides

deterministic latency profiles critical for time-sensitive operations including medical device monitoring requiring subsecond response and high-frequency financial trading where millisecond delays translate to measurable business impact. However, the 23-47ms agent coordination overhead observed in multi-agent systems demonstrates that distributed processing introduces its own latency penalties that must be carefully managed through optimized communication protocols and intelligent agent placement (Abdi *et al.*, 2024)^[1].

The hybrid edge-cloud architecture combining lightweight edge screening (45ms latency, 94.2% accuracy) with cloud-based deep analysis (280ms latency, 98.7% accuracy) represents pragmatic approach balancing real-time requirements with detection accuracy demands. Edge deployment handles high-volume screening filtering obvious benign traffic and flagging clear threats, while routing ambiguous cases to cloud infrastructure with computational resources for sophisticated analysis. This tiered architecture optimizes resource utilization while maintaining detection performance, proving particularly valuable for resource-constrained healthcare facilities that cannot afford expensive edge computing infrastructure throughout their networks.

However, edge deployment introduces significant operational complexity including model consistency management, coordinated threat intelligence distribution, and version control challenges across distributed devices. The 47% longer deployment time for edge model updates compared to centralized cloud updates creates operational burden that healthcare and financial IT teams must carefully manage. Automated orchestration tools enabling streamlined edge deployment workflows emerge as critical enabler for practical edge security architecture adoption.

5.3. Explainability-Performance Trade-offs and Regulatory Implications

The tension between detection performance and explainability manifests clearly in practitioner ratings where adaptive neuro-fuzzy systems achieved highest explainability scores (4.3/5.0) despite CNNs demonstrating 2.8% superior detection accuracy. Security analysts consistently emphasized that interpretability enables validation of automated detections before implementing potentially disruptive response actions, facilitates incident investigation for regulatory compliance, and supports documentation of security controls required by healthcare HIPAA and financial PCI-DSS regulations. This operational preference for explainability even at modest accuracy cost challenges research community's traditional focus on maximizing detection performance without adequate consideration of interpretability requirements (Samek *et al.*, 2021)^[21].

The finding that attention-enhanced neural networks receive significantly higher explainability ratings (3.8/5.0) compared to standard CNNs (2.3/5.0) despite 12-18% computational overhead suggests that even modest interpretability improvements provide substantial operational value justifying performance costs. Attention mechanisms highlighting network traffic features and temporal patterns contributing to threat classifications enable efficient analyst triage where high-confidence detections with clear explanations proceed to automated response while ambiguous cases with uncertain explanations route to human investigation. This hybrid human-AI workflow optimizes analyst productivity while maintaining detection accuracy.

Regulatory frameworks mandating explainable security decision-making create market opportunities for interpretable AI approaches despite potential accuracy sacrifices. Healthcare organizations subject to HIPAA breach notification requirements and financial institutions navigating PCI-DSS incident response documentation cannot deploy black-box detection systems lacking auditable decision trails. The 8-15% accuracy gap between interpretable and black-box models represents acceptable trade-off for regulated industries prioritizing compliance and operational validation over marginal detection improvements.

5.4. False Positive Management and Alert Fatigue

The 43% false positive reduction achieved through ensemble consensus mechanisms addresses critical operational challenge where excessive false alarms create analyst alert fatigue, potentially causing missed genuine threats, and where blocking legitimate activities creates patient care disruptions or transaction processing failures with significant business consequences. Financial institutions reporting that false positive reduction from 5.7% to 2.3% enabled 31% security analyst staffing reduction demonstrates direct operational impact translating to measurable cost savings beyond abstract security improvements. This finding validates ensemble approaches as practical solution for organizations struggling with alert overload in high-volume network environments (Al-Saeed *et al.*, 2020)^[4].

However, false positive optimization introduces detection recall trade-offs that require risk-based calibration aligned with organizational priorities. Healthcare organizations prioritizing patient safety through comprehensive threat detection may accept higher false positive rates (4-6%) to maximize detection recall, while financial institutions optimizing customer experience and minimizing transaction friction may tune systems toward lower false positives (2-3%) accepting modest recall reduction. This context-dependent optimization underscores importance of configurable detection thresholds enabling organizations to balance competing operational priorities rather than pursuing universal accuracy maximization divorced from organizational goals.

The observation that false positive causes include 47% legitimate operational variations not adequately represented in training data emphasizes critical importance of establishing accurate baseline models capturing normal operational patterns before deploying anomaly detection. Healthcare networks exhibit complex temporal patterns including admission/discharge cycles, shift changes, and scheduled maintenance windows that create legitimate traffic variations resembling attacks if not properly modeled. Financial systems similarly feature predictable patterns including payroll processing, tax payment deadlines, and holiday shopping spikes that must be distinguished from denial-of-service attacks and fraud campaigns.

5.5. Organizational and Operational Considerations

Practitioner interviews reveal organizational factors significantly influence AI detection technology adoption beyond technical performance metrics. Security teams emphasize operational integration requirements including compatibility with existing SIEM platforms, ability to leverage current threat intelligence feeds, and minimal disruption to established workflows. Healthcare security

teams lacking specialized machine learning expertise express concerns regarding detection system maintenance, model retraining requirements, and troubleshooting complex AI systems when failures occur. Financial institutions emphasize vendor support, regulatory approval of detection methodologies, and documented testing demonstrating compliance with audit requirements. Cost considerations prove paramount, particularly for healthcare organizations operating on constrained IT security budgets averaging 4-6% of overall IT spending. The computational requirements of deep learning detection systems create barriers where organizations cannot afford GPU-accelerated edge infrastructure or cloud processing fees for high-volume event stream analysis. Genetic algorithm optimization reducing computational requirements by 67% while maintaining detection accuracy directly addresses this

practical constraint enabling broader technology adoption across resource-constrained healthcare facilities (Sonia *et al.*, 2024) [22]. Cultural factors including risk tolerance, regulatory compliance priorities, and operational philosophies shape detection technology preferences. Healthcare organizations emphasize patient safety and operational continuity, favoring conservative detection thresholds minimizing false negatives even at cost of increased false positives requiring manual investigation. Financial institutions prioritize transaction processing continuity and customer experience, favoring optimized false positive rates that avoid legitimate transaction blocking. These sectoral differences necessitate customizable detection configurations rather than one-size-fits-all approaches.




Fig 4: Implementation Decision Framework - Flowchart guiding organizations through algorithm selection, deployment architecture, and configuration decisions based on sector requirements, resources, and priorities

6. Conclusion

This research provides comprehensive empirical assessment of real-time AI-based cyber threat detection and response for healthcare and financial networks, demonstrating that modern machine learning approaches achieve 96.7% detection accuracy with 847ms mean latency, representing transformative improvement over traditional signature-based systems requiring 15.2 seconds detection time with 68.2% accuracy. The findings validate AI-driven threat assessment as essential capability for protecting critical infrastructure against sophisticated cyber threats that increasingly target healthcare patient data and financial transaction systems with coordinated multi-vector attacks employing polymorphic

malware, zero-day exploits, and advanced persistent threat techniques.

Key contributions include:

- comprehensive performance benchmarking across eight representative machine learning approaches evaluated on healthcare and financial network testbeds with realistic operational traffic and attack patterns,
- validation that edge computing architectures enable 4.3x latency improvement essential for real-time response requirements while introducing manageable deployment complexity,
- demonstration that ensemble consensus mechanisms reduce false positives by 43% addressing critical

- operational challenge of alert fatigue,
- quantification of explainability-performance trade-offs revealing that interpretable neuro-fuzzy approaches receive 4.3/5.0 practitioner ratings despite 2.8% accuracy sacrifice compared to black-box neural networks, and
 - identification that genetic algorithm feature selection reduces computational requirements by 67% enabling deployment on resource-constrained edge devices prevalent in healthcare and financial point-of-sale infrastructure.

The research reveals fundamental tensions requiring careful balance including real-time processing demands versus detection accuracy optimization, where lightweight edge models sacrifice 3.3% accuracy for 4.3x latency improvement; explainability requirements versus performance maximization, where interpretable models enable regulatory compliance and analyst validation at 8-15% accuracy cost; and false positive minimization versus detection recall optimization, where consensus thresholds must be calibrated to organizational risk tolerance and operational priorities. No single algorithm or architecture universally optimizes across all dimensions; rather, effective threat detection requires context-specific design aligned with sector requirements, regulatory constraints, resource availability, and operational workflows.

Critical challenges persist including cold-start problems when encountering zero-day exploits (73.2% accuracy on novel attack families), concept drift requiring model retraining every 14.3 days on average as attack patterns evolve, computational overhead of real-time deep learning requiring GPU acceleration or specialized edge processors, and fundamental explainability-performance trade-offs where regulatory compliance requirements may preclude most accurate black-box approaches. Addressing these challenges requires sustained research advancing adaptive learning for rapid model updating, efficient architectures balancing accuracy with computational feasibility, and explainable AI techniques providing interpretability without severe performance penalties.

As healthcare and financial infrastructure grows increasingly digitized through electronic health records, telemedicine expansion, mobile banking proliferation, and automated trading system deployment, effective AI-driven threat detection becomes essential for maintaining operational resilience, protecting sensitive data, and ensuring continuity of critical services upon which society depends. The findings demonstrate that real-time AI threat assessment represents practical, deployable capability rather than theoretical possibility, though successful implementation requires addressing technical, organizational, regulatory, and operational considerations that extend beyond pure algorithmic performance optimization.




Fig 5: Real-Time AI Security Roadmap - Timeline showing evolution from current capabilities through near-term improvements to long-term vision with key milestones and technology enablers

7. Limitations

Several limitations qualify interpretation and generalizability of research findings. First, experimental testbeds, while carefully designed to replicate healthcare and financial network characteristics, inevitably fail to capture full complexity of production environments including legacy system integration challenges, operational anomalies reflecting organizational policies and workflows, and the vast scale of real-world healthcare and financial networks

processing billions of events daily. Testbed evaluation may overstate detection performance compared to operational deployment encountering unexpected edge cases and infrastructure idiosyncrasies absent from simulated environments.

Second, attack traffic injections based on known threat intelligence and public malware samples may not fully represent sophisticated zero-day exploits, advanced persistent threats, and targeted attacks that well-resourced

adversaries deploy against high-value healthcare and financial targets. Real-world nation-state actors and organized cybercrime syndicates employ bespoke attack tooling specifically designed to evade detection systems, potentially exceeding evasion capabilities of research testbed attack traffic. Conversely, the research also includes attack traffic that resource-constrained adversaries may not realistically deploy, potentially overstating threat sophistication.

Third, practitioner interview sample (n=18) reflects limited subset of healthcare and financial security community, potentially not capturing full diversity of organizational contexts, operational philosophies, and technology adoption patterns across thousands of healthcare facilities and financial institutions nationwide. Large academic medical centers with dedicated security teams may hold different perspectives than rural community hospitals with limited IT staffing; similarly, major money-center banks differ substantially from community credit unions in resources and threat exposure. Broader survey-based research would strengthen understanding of adoption barriers and implementation patterns.

Fourth, the research focuses primarily on network-based threat detection without comprehensive examination of endpoint detection and response, application security monitoring, or physical security integration that comprise complete security programs. A holistic security posture requires coordinated capabilities across multiple domains that extends beyond network-centric evaluation. The interactions, dependencies, and potential conflicts between network AI detection and complementary security controls remain underexplored.

Fifth, longitudinal evaluation limitations constrain understanding of detection system performance evolution over extended operational deployments. Current evaluation spans 30-day testbed periods; multi-year field studies would reveal performance degradation patterns from concept drift, model maintenance requirements, and adaptation to evolving threat landscapes that short-duration evaluations cannot capture. The sustainability and total cost of ownership of AI detection systems requires longitudinal assessment tracking operational metrics, incident response effectiveness, and analyst productivity over extended timeframes.

Finally, cost analysis remains preliminary, focusing on computational resource requirements without comprehensive total cost of ownership accounting including software licensing, vendor support contracts, staff training, ongoing model maintenance, and opportunity costs from false positives requiring analyst investigation. Healthcare and financial organizations make technology adoption decisions based on holistic cost-benefit analysis rather than purely technical performance metrics; refined economic evaluation would inform practical deployment decisions.

8. Practical Implications

8.1. For Healthcare and Financial Security Teams

Security operations centers should prioritize ensemble detection architectures combining multiple complementary algorithms rather than relying on single detection approaches, as ensemble methods achieve 96.7% accuracy while reducing false positives by 43% through consensus mechanisms that

filter spurious alerts. Organizations should implement tiered detection with lightweight edge screening providing rapid initial triage (45ms latency, 94.2% accuracy) and cloud-based deep analysis handling suspicious cases requiring sophisticated processing (280ms latency, 98.7% accuracy), optimizing resource utilization while maintaining detection coverage.

Organizations must explicitly address explainability requirements early in detection system selection, evaluating whether regulatory compliance, incident investigation, and analyst validation needs necessitate interpretable approaches despite potential accuracy sacrifices. Healthcare organizations subject to HIPAA breach notification and financial institutions navigating PCI-DSS incident documentation should prioritize neuro-fuzzy or attention-enhanced architectures providing decision transparency over black-box neural networks achieving marginal accuracy improvements without interpretability.

Security teams should calibrate false positive thresholds based on organizational risk tolerance, operational priorities, and analyst capacity rather than pursuing universal accuracy optimization. Healthcare organizations prioritizing patient safety through comprehensive threat detection may accept 4-6% false positive rates maximizing detection recall, while financial institutions optimizing transaction processing continuity may tune systems toward 2-3% false positives accepting modest recall reduction. Regular threshold adjustment based on operational feedback enables dynamic optimization aligned with evolving priorities.

8.2. For Healthcare and Financial Organizations

Executive leadership must recognize that effective AI-driven threat detection requires sustained investment beyond initial system acquisition, including computational infrastructure for real-time processing, staff training for security analysts operating AI-augmented workflows, and ongoing model maintenance addressing concept drift and threat evolution. Healthcare organizations should budget 15-20% annual operating costs atop initial detection system acquisition for maintenance, updates, and analyst training. Financial institutions should evaluate total cost of ownership including computational infrastructure, vendor support, and analyst productivity impacts rather than purely upfront licensing costs.

Organizations should establish cross-functional teams integrating security, clinical operations, regulatory compliance, and IT infrastructure stakeholders for detection system implementation, as successful deployment requires coordinating technical capabilities with operational workflows, compliance requirements, and organizational change management. Healthcare implementations affecting clinical systems require physician and nursing input ensuring detection and response actions do not compromise patient care delivery. Financial implementations must coordinate with transaction processing teams avoiding legitimate transaction disruption.

Organizations should implement graduated response automation enabling proportionate defensive actions matched to threat confidence and potential impact. High-confidence malware detections on non-critical systems can trigger automated isolation and remediation; ambiguous

detections on critical patient care or transaction processing systems should route to human investigation before disruptive response. This graduated automation balances security responsiveness with operational continuity imperatives, avoiding patient care disruption or transaction processing failures from overly aggressive automated blocking.

8.3. For Security Technology Vendors

Vendors developing AI-driven detection products for healthcare and financial markets must prioritize explainability-by-design rather than treating interpretability as optional feature or post-hoc addition. Regulatory requirements, operational validation needs, and analyst workflows in these sectors demand transparent detection reasoning; products lacking explainability face adoption barriers despite superior technical performance. Vendors should integrate attention mechanisms, fuzzy logic, or rule extraction techniques providing decision transparency without requiring specialized AI expertise from security analysts.

Product architectures should support flexible deployment across cloud, hybrid, and edge configurations enabling organizations to optimize latency, computational costs, and data sovereignty requirements based on specific operational constraints. Healthcare facilities with bandwidth limitations and medical device networks requiring local processing need edge-capable solutions; financial institutions with centralized security operations and abundant computational resources may prefer cloud-native architectures. Vendors supporting multiple deployment models expand addressable markets across diverse customer segments.

Vendors should develop pre-configured detection models optimized for healthcare and financial threat landscapes rather than expecting customers to perform extensive tuning and customization requiring specialized machine learning expertise often absent from security teams. Healthcare-specific models trained on medical device protocols, EHR workflows, and healthcare-targeted malware families enable faster time-to-value compared to generic detection requiring extensive customization. Similarly, financial-specific models understanding transaction processing patterns, mobile banking behaviors, and financial fraud techniques provide superior out-of-box performance.

8.4. For Policymakers and Regulators

Regulatory bodies should develop updated cybersecurity frameworks explicitly addressing AI-driven detection capabilities, performance standards, explainability requirements, and validation methodologies appropriate for machine learning systems that differ fundamentally from traditional signature-based approaches. Current regulations developed for conventional security technologies may inadequately address AI-specific considerations including model training data quality, algorithm bias, performance monitoring, and update management. Evidence-based regulatory standards informed by empirical research like current findings would provide clearer guidance for both regulators and regulated entities.

Policymakers should incentivize information sharing regarding attack patterns, detection techniques, and threat

intelligence across healthcare and financial sectors through regulatory safe harbors protecting participants in collaborative security initiatives. The research demonstrates that detection systems benefit from diverse training data; enabling healthcare organizations and financial institutions to contribute anonymized attack data to shared repositories would improve collective security capabilities. Privacy-preserving federated learning approaches enable collaborative model training without exposing sensitive operational data.

Regulators should recognize explainability-performance trade-offs inherent to AI detection systems when establishing security requirements, allowing risk-based approaches where organizations may deploy black-box algorithms for specific high-risk threats if compensating controls provide adequate transparency through monitoring, human oversight, and incident investigation capabilities. Rigid mandates requiring interpretable approaches across all use cases may inadvertently reduce security effectiveness where accuracy maximization proves essential.

9. Future Research

Several promising research directions emerge from current findings and identified limitations. First, adaptive learning and continual training approaches enabling rapid model updating when encountering novel attack types warrant urgent attention given 73.2% accuracy degradation on zero-day exploits. Meta-learning techniques allowing models to quickly adapt to new threat families with minimal training examples, online learning approaches enabling incremental model updates without full retraining, and transfer learning leveraging knowledge from related security domains could address cold-start problems that create detection blind spots during critical windows following new threat emergence.

Second, explainable AI techniques specifically designed for security applications require development beyond general-purpose interpretability methods. Security-specific explanation frameworks that provide actionable insights for incident investigation, support regulatory compliance documentation, and enable efficient analyst validation would advance practical deployability in healthcare and financial environments. Research should examine hybrid architectures partitioning detection pipelines into interpretable threat triage followed by black-box deep analysis, enabling staged decision-making where transparency and accuracy optimization occur at appropriate workflow phases.

Third, adversarial robustness of real-time detection systems against sophisticated evasion techniques demands systematic investigation. As AI-driven detection achieves broader deployment, adversaries will increasingly employ adversarial machine learning techniques specifically designed to evade automated detection. Research should examine detection system resilience against adversarial perturbations, adaptive attacks that probe detection boundaries, and coordinated evasion campaigns across multiple attack vectors. Defensive strategies including adversarial training, input validation, and ensemble diversity that improve robustness against sophisticated adversaries warrant development.

Fourth, federated learning approaches enabling collaborative model training across multiple healthcare facilities and financial institutions without centralizing sensitive

operational data present promising direction for improving detection through broader training datasets while respecting privacy and competitive sensitivities. Research should address challenges including non-IID data distributions across participants, Byzantine failures where compromised participants contribute poisoned data, and communication efficiency enabling practical federated training across bandwidth-constrained participants.

Fifth, longitudinal field studies tracking detection system performance over multi-year operational deployments would provide insights into performance sustainability, concept drift management, model maintenance requirements, and total cost of ownership absent from short-duration evaluations. Collaborative research partnerships between academic institutions and healthcare/financial organizations would enable access to operational metrics, incident response data, and analyst feedback informing evidence-based security engineering rather than laboratory-based research divorced from operational realities.

Sixth, investigation of human-AI collaboration paradigms optimizing analyst productivity through effective automation would advance understanding of how security operations centers can best leverage AI capabilities. Research should examine optimal allocation of detection tasks between human analysts and automated systems, workflow designs enabling efficient analyst validation of AI detections, and user interface designs presenting threat information enabling rapid assessment. The goal involves augmenting rather than replacing human security expertise through complementary human-AI teams.

Finally, economic analysis of AI detection technology adoption examining cost-benefit trade-offs, return on security investment, and total cost of ownership would inform evidence-based decision-making for resource-constrained healthcare and financial organizations. Research should examine how detection performance improvements translate to measurable risk reduction, quantify operational efficiency gains from automation, and develop frameworks enabling organizations to evaluate AI security investments against competing technology and staffing options within constrained budgets.

References

1. Abdi A, Pourvabab M, Rahmani AM, Hosseinzadeh M. Edge computing for real-time cybersecurity. *IEEE Trans Mob Comput.* 2024;23(2):789-801. doi:10.1109/TMC.2023.3268129.
2. Alamleh A, *et al.* Multi-attribute decision-making for intrusion detection systems: a systematic review. *Int J Inf Technol Decis Mak.* 2023;22(1):589-636. doi:10.1142/S0219622023500093.
3. Alomari ES, *et al.* Malware detection using deep learning and correlation-based feature selection. *Symmetry.* 2023;15(1):123. doi:10.3390/sym15010123.
4. Al-Saeed IA, Selamat A, Rohani MF, Krejcar O, Chaudhry JA. A systematic state-of-the-art analysis of multi-agent intrusion detection. *IEEE Access.* 2020;8:180184-180209. doi:10.1109/ACCESS.2020.3027854.
5. Arshad S, Shah MA, Wahid A, Mehmood A, Song H, Yu H. SAMADroid: a novel 3-level hybrid malware detection model for Android operating system. *IEEE Access.* 2018;6:4321-4339. doi:10.1109/ACCESS.2018.2792941.
6. Beslin Pajila PJ, Golden Julie E, Harold Robinson Y. ABAP: anchor node based DDoS attack detection using adaptive neuro-fuzzy inference system. *Wirel Pers Commun.* 2023;128(2):875-899.
7. Danish M. Enhancing cyber security through predictive analytics: real-time threat detection and response. *arXiv [Preprint].* 2024. doi:10.48550/arxiv.2407.10864.
8. De Cristofaro E. LLM APIs in cybersecurity: costs and scalability. *IEEE Commun Mag.* 2024;62(4):56-62. doi:10.1109/MCOM.001.2300712.
9. Fernandes G, Rodrigues JJPC, Carvalho LF, Al-Muhtadi JF, Proença ML. A comprehensive survey on network anomaly detection. *Telecommun Syst.* 2019;70:447-489. doi:10.1007/s11235-018-0475-8.
10. Jordan MI, Mitchell TM. Machine learning: trends, perspectives, and prospects. *Science.* 2015;349(6245):255-260. doi:10.1126/science.aaa8415.
11. Khraisat A, Gondal I, Vamplew P, Kamruzzaman J. Survey of intrusion detection systems: techniques, datasets and challenges. *Cybersecurity.* 2019;2(1):20. doi:10.1186/s42400-019-0038-7.
12. Le HH, Viviani J. Predicting bank failure: an improvement by implementing a machine-learning approach to classical financial ratios. *Res Int Bus Finance.* 2017;44:16-25. doi:10.1016/j.ribaf.2017.07.104.
13. Li J, Sun L, Yan Q, Li Z, Srisa-an W, Ye H. Significant permission identification for machine-learning-based Android malware detection. *IEEE Trans Ind Inform.* 2018;14(7):3216-3225. doi:10.1109/TII.2017.2789219.
14. Naoui MA, Lejdel B, Ayad M, Amamra A, Kazar O. Using a distributed deep learning algorithm for analyzing big data in smart cities. *Smart Sustain Built Environ.* 2021;10(1):90-105. doi:10.1108/SASBE-03-2020-0034.
15. Or-Meir O, Nissim N, Elovici Y, Rokach L. Dynamic malware analysis in the modern era: a state-of-the-art survey. *ACM Comput Surv.* 2019;52(5):1-48. doi:10.1145/3329786.
16. Ovabor K, Sule-Odu IO, Atkison T, Fabusoro AT, Benedict JO. AI-driven threat intelligence for real-time cybersecurity: frameworks, tools, and future directions. *Open Access Res J Sci Technol.* 2024;12(2):040-048. doi:10.53022/oarjst.2024.12.2.0135.
17. Ozkan-Okay M, Samet R, Aslan Ö, Gupta D. A comprehensive systematic literature review on intrusion detection systems. *IEEE Access.* 2021;9:157727-157760. doi:10.1109/ACCESS.2021.3129907.
18. Rahmati M. Towards explainable and lightweight AI for real-time cyber threat hunting in edge networks. *arXiv [Preprint].* 2025. doi:10.48550/arxiv.2504.16118.
19. Roy S. AgenticCyber: a GenAI-powered multi-agent system for multimodal threat detection and adaptive response in cybersecurity. *arXiv [Preprint].* 2025. doi:10.48550/arxiv.2512.06396.
20. Saleh AI, Talaat FM, Labib LM. A hybrid intrusion detection system based on prioritized k-nearest neighbors and optimized SVM classifiers. *Artif Intell*

- Rev. 2019;51:403-443. doi:10.1007/s10462-017-9567-1.
21. Samek W, Wiegand T, Müller KR. Explainable artificial intelligence: understanding, visualizing and interpreting deep learning models. *Entropy*. 2021;23(1):18. doi:10.3390/e23010018.
 22. Sonia M, Lakshmi CBN, Hussain SJ, Swarupa ML, Rajeswaran N. Android malware detection using genetic algorithm based optimized feature selection and machine learning. In: *Computational Intelligence in Machine Learning. ICCIML 2022. Lecture Notes in Electrical Engineering*. Singapore: Springer; 2024. vol 1106. doi:10.1007/978-981-99-7954-7_19.

How to Cite This Article

Olatunbosun E, Bizel G. Real-time AI-based cyber threat assessment for improving incident response in U.S. healthcare and financial networks. *Int J Eng Comput Appl*. 2026;2(1):48-64.

Creative Commons (CC) License

This is an open access journal, and articles are distributed under the terms of the Creative Commons Attribution NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0) License, which allows others to remix, tweak, and build upon the work non-commercially, as long as appropriate credit is given and the new creations are licensed under the identical terms.