ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1548

AI-Driven Automated Incident Response in Healthcare Cybersecurity: A Systematic Review of SOAR Frameworks, IoMT Security, and Emerging Trends

Ghanakshari Khandre[1]; Shripad Bhide[2]

[1]Post Graduate Student; [2]Assistant Professor

[1,2]Department of Computer Engineering, P.E.S. Modern College of Engineering Pune, India

Publication Date: 2026/06/01

Abstract: Modern healthcare infrastructures face an escalating spectrum of cyber threats, driven in large part by their growing reliance on tightly integrated digital ecosystems — spanning electronic health records, cloud platforms, and networked clinical devices. This exposure is compounded by the rapid proliferation of the Internet of Medical Things (IoMT), a domain characterized by resource-constrained endpoints, inconsistent patch cycles, and legacy communication protocols ill-suited to contemporary security demands.

Concurrently, artificial intelligence and machine learning have emerged as promising instruments for advancing cyber threat detection. Yet a persistent and consequential gap remains: the transition from detection to effective, automated response. Though conceptually intertwined, these two functions present markedly different operational realities. Current literature disproportionately favors detection-centric approaches, leaving automated incident response comparatively underexplored.

This paper offers a systematic review of AI-driven automated incident response in the context of healthcare cybersecurity. It critically examines the deployment of Security Orchestratioshrin, Automation, and Response (SOAR) frameworks, delineates security challenges endemic to IoMT environments, and surveys emerging intelligent defense paradigms. Following a PRISMA-guided methodology, relevant literature was drawn from IEEE Xplore, SpringerLink, PubMed, Google Scholar, and arXiv, encompassing publications from 2019 to 2026. A rigorous multi-stage screening of 1,499 initial records, governed by predefined inclusion criteria, yielded 20 studies for substantive analysis.

The findings reveal a consistent pattern: while AI-based detection models attain strong performance benchmarks, the operationalization of automated response within clinical settings remains nascent. SOAR adoption continues to mature slowly, and prevailing approaches frequently fall short in delivering real-time mitigation and recovery. Compounding these limitations are dependence on narrow benchmark datasets, insufficient model interpretability for clinical contexts, and inadequate incorporation of privacy-preserving methodologies such as federated learning.

Taken together, this review makes a compelling case for end-to-end security architectures that transcend detection alone. Next-generation systems must embed SOAR capabilities across the full incident response lifecycle, positioning healthcare organizations to adopt cybersecurity solutions that are not only technically robust but practically deployable at scale.

How to Cite: Ghanakshari Khandre; Shripad Bhide (2026) AI-Driven Automated Incident Response in Healthcare Cybersecurity: A Systematic Review of SOAR Frameworks, IoMT Security, and Emerging Trends. *International Journal of Innovative Science and Research Technology*, 11(5), 2598-2607. https://doi.org/10.38124/ijisrt/26may1548

I. INTRODUCTION

The healthcare sector has experienced an accelerated shift toward digital infrastructure, with interconnected technologies — including electronic health records (EHRs), cloud-based platforms, and the Internet of Medical Things (IoMT) — now forming the operational backbone of modern clinical environments. While these developments have brought measurable gains in care efficiency and service delivery, they have simultaneously expanded the attack

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1548

surface available to malicious actors. Healthcare organizations have increasingly become high-value targets for cyberattacks, ransomware in particular, owing to the sensitivity of medical data and the mission-critical nature of clinical workflows. Consequently, cybersecurity in healthcare has evolved well beyond conventional data protection, emerging as a direct determinant of patient safety and operational continuity.

A significant driver of this expanding threat landscape is the pervasive integration of IoMT devices — ranging from infusion pumps and wearable monitors to diagnostic imaging systems — many of which function under severe computational constraints and limited native security capabilities. These endpoints frequently operate on outdated firmware and are architecturally incompatible with standard endpoint protection mechanisms. Compounding this, healthcare security operations continue to depend heavily on manual incident response workflows that are both labor-intensive and contingent on specialist expertise. Security analysts routinely contend with overwhelming alert volumes, fostering alert fatigue and introducing dangerous latency into response cycles. The resulting vulnerability is stark: threats may be identified in time, yet not acted upon swiftly enough to avert meaningful harm.

To address these structural deficiencies, the research community has increasingly turned to Artificial Intelligence (AI) and Machine Learning (ML)-based intrusion detection systems, which have demonstrated considerable efficacy in recognizing anomalous behavior and identifying threats across healthcare networks. In parallel, Security Orchestration, Automation, and Response (SOAR) platforms have been developed to streamline incident response by integrating disparate security tools and executing structured, predefined playbooks. Individually, both paradigms offer substantial promise; however, their coordinated application within healthcare settings — and specifically within IoMT ecosystems — remains markedly underdeveloped.

This observation points to a deeper, more systemic problem in the current research landscape. Detection capabilities have advanced at a notably faster pace than response capabilities, leaving a critical operational gap. The majority of existing studies address the identification of attacks without extending their scope to encompass response-oriented mechanisms such as containment, remediation, or system recovery. Healthcare environments further complicate the adoption of automated frameworks through a distinct set of constraints — regulatory compliance obligations, pronounced system heterogeneity, and the imperative to preserve patient safety throughout any automated intervention. These factors collectively impede the translation of theoretical automation models into clinically viable deployment.

Against this backdrop, this paper presents a systematic review of AI-driven automated incident response within healthcare cybersecurity, with focused attention on SOAR framework adoption, IoMT-specific security challenges, and emerging intelligent defense strategies. The study aims to examine how existing research approaches the integration of detection and response, surface key limitations in prevailing methodologies, and delineate the research gaps that must be addressed before effective, real-world automation in healthcare security operations becomes achievable. Through cross-study synthesis, this review seeks to offer a structured analytical foundation that can meaningfully inform the design of scalable, practical security solutions.

This paper is structured as follows.Section 2 establishes the conceptual background of healthcare cybersecurity, covering incident response models, SOAR frameworks, and IoMT-specific vulnerabilities. Section 3 details the systematic methodology governing study selection and analytical approach. Section 4 presents a comparative review of the selected literature. Section 5 offers an in-depth discussion of central findings and identified research gaps. Section 6 outlines prospective directions for future inquiry. Section 7 concludes the paper.

II. BACKGROUND

- *Healthcare Cybersecurity Evolution*

Over the past decade, cybersecurity in healthcare has undergone a fundamental transformation — evolving from a peripheral IT concern into a mission-critical dimension of clinical infrastructure. Early defensive postures were largely perimeter-oriented, relying on firewalls and signature-based detection to safeguard relatively isolated systems. The widespread integration of electronic health records, cloud platforms, and networked clinical technologies has since dismantled these boundaries. Contemporary healthcare environments operate as deeply interconnected ecosystems, with sensitive data flowing continuously across internal and external networks — dramatically broadening the exploitable attack surface.

Simultaneously, the character of cyber threats has shifted from opportunistic intrusions to coordinated, high-impact campaigns. Ransomware attacks, in particular, have proven capable of paralyzing hospital operations, postponing critical treatments, and placing patients in direct jeopardy. This trajectory reveals an uncomfortable truth: cybersecurity failures in healthcare are no longer measured solely in data loss but increasingly in clinical consequences. Despite the proliferation of sophisticated monitoring and detection tools, defensive strategies have remained largely reactive. Systems can identify anomalies with growing precision, yet consistently struggle to convert those identifications into swift, proportionate action.

This asymmetry reflects a structural deficiency in the current paradigm. Detection technologies have advanced at a pace that response mechanisms have not matched, leaving healthcare systems most exposed at the very moment an attack materializes. The gap between identifying a threat and acting upon it constitutes one of the central, unresolved challenges in healthcare cybersecurity research.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1548

➢ *Manual vs. Automated Incident Response*

Incident response remains a foundational pillar of cybersecurity operations, yet in healthcare settings it continues to be executed predominantly through manual, expertise-dependent processes. Security analysts are required to triage alerts, verify threats, and formulate mitigation strategies under significant time pressure. As alert volumes in modern healthcare networks grow, this approach becomes increasingly untenable. In practice, analysts face cognitive overload, resulting in delayed responses, inconsistent decision quality, and an elevated probability that critical threats go unaddressed.

Automation has been positioned as the logical remedy — enabling faster, more consistent handling of security events by executing predefined actions in response to detected conditions, without requiring human intervention at each step. In principle, this shifts organizations from a reactive posture to one approaching near real-time defense, a capability that is particularly vital in high-stakes clinical environments.

However, healthcare presents constraints that have no direct analogue in conventional enterprise IT. Patient safety must function as a non-negotiable variable in any automated decision. Actions that are routine in corporate networks — device isolation, process termination, traffic blocking — carry the potential to disrupt active clinical procedures when applied in medical contexts. Consequently, most current implementations of automation in healthcare stop well short of full autonomy, focusing instead on alert enrichment and prioritization rather than active response. This cautious but necessary conservatism reinforces the fundamental limitation: even as detection improves, response remains slow, fragmented, and reliant on human judgment — perpetuating the very gap that automation was intended to close.

➢ *SOAR Frameworks*

Security Orchestration, Automation, and Response platforms have been broadly recognized as a structural solution to the inefficiencies endemic to traditional incident response. By consolidating disparate security tools and orchestrating their interaction through automated, playbook-driven workflows, SOAR systems aim to standardize and accelerate the response process. In enterprise environments, these platforms have demonstrated clear value — reducing response times, minimizing operational overhead, and enabling coordinated action across detection systems, firewalls, and threat intelligence feeds.

Despite this track record, SOAR frameworks have not been meaningfully transposed into healthcare contexts. A central obstacle is the absence of domain-specific playbooks capable of accommodating the particular complexities of clinical environments — heterogeneous device ecosystems, specialized communication protocols, and stringent regulatory obligations. Existing SOAR implementations are largely premised on environments where automated containment can be executed without posing collateral risk, an assumption that is fundamentally untenable in clinical settings where devices may be actively supporting patient care.

Empirical evidence for SOAR deployment within healthcare is also conspicuously scarce. Most available studies present conceptual architectures or simulated evaluations, without validation in live hospital infrastructure. This absence of real-world testing undermines confidence in their clinical applicability and scalability. As a result, SOAR remains a technically sound but practically underutilized instrument in healthcare cybersecurity — one whose full potential is constrained by a failure to account for the unique demands of the domain it most urgently needs to serve.

➢ *IoMT Security*

The Internet of Medical Things has become an indispensable component of modern healthcare, underpinning continuous patient monitoring, remote diagnostics, and data-informed clinical decision-making. IoMT devices span a broad operational spectrum — from body-worn sensors to life-critical support systems — and are in constant communication across interconnected hospital networks. This pervasive connectivity, while enabling more responsive and efficient care, simultaneously introduces a security environment of considerable complexity and fragility.

In deployed settings, many IoMT devices contend with severe resource constraints: limited processing capacity, extended operational lifecycles that preclude frequent replacement, and restricted ability to receive security updates. These characteristics leave devices chronically exposed to known vulnerabilities, with no practical mechanism for timely remediation. The coexistence of diverse communication protocols and legacy system integrations further compounds security management, creating numerous potential entry points for adversarial exploitation and enabling lateral movement across hospital networks once an initial compromise occurs.

Research in this domain has concentrated predominantly on detection — applying machine learning and anomaly detection techniques to identify malicious behavior within IoMT traffic. While these methods consistently demonstrate strong performance in controlled evaluations, they rarely address what should follow detection. In operational terms, identifying a threat without a defined mechanism to contain or neutralize it yields limited security benefit, particularly in environments where delayed or inappropriate responses can translate directly into harm. The absence of integrated detection-and-response capabilities within IoMT security therefore stands as a critical barrier one that amplifies the broader detection–response gap that runs throughout this field.

III. METHODOLOGY

➢ *Search Strategy*

A structured, systematic search was conducted across five established academic databases: Google Scholar, IEEE Xplore, arXiv, SpringerLink, and PubMed. To

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1548

comprehensively capture the relevant literature, multiple keyword combinations were employed, including: *"SOAR in healthcare cybersecurity," "automated incident response in healthcare systems," "AI-based incident response for IoMT," "machine learning in IoMT security," "cybersecurity for medical devices with automation," "LLM-based incident response,"* and *"AI intrusion detection in healthcare."*

The search was bounded to publications appearing between 2019 and 2026, with retrieval limited to English-language sources. Peer-reviewed journal articles and conference proceedings were prioritized to ensure scholarly rigor and methodological quality.

➢ *Included and Excluded Criteria*

Table 1 Inclusion and Exclusion Criteria

Criteria	Included ✓	Excluded ✗
Publication Year	2019–2026	Before 2019
Source Type	Peer-reviewed journals & conferences	Blogs, reports, non-reviewed sources
Research Focus	AI/ML, IoMT, healthcare cybersecurity	Irrelevant domains
Scope	SOAR and automated response	Hardware-only or unrelated topics
Accessibility	Full-text available	Abstract-only papers

➢ *PRISMA Flow*

The study selection process adhered to the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) framework. Figure 1 depicts the full progression from initial database identification through final study inclusion.

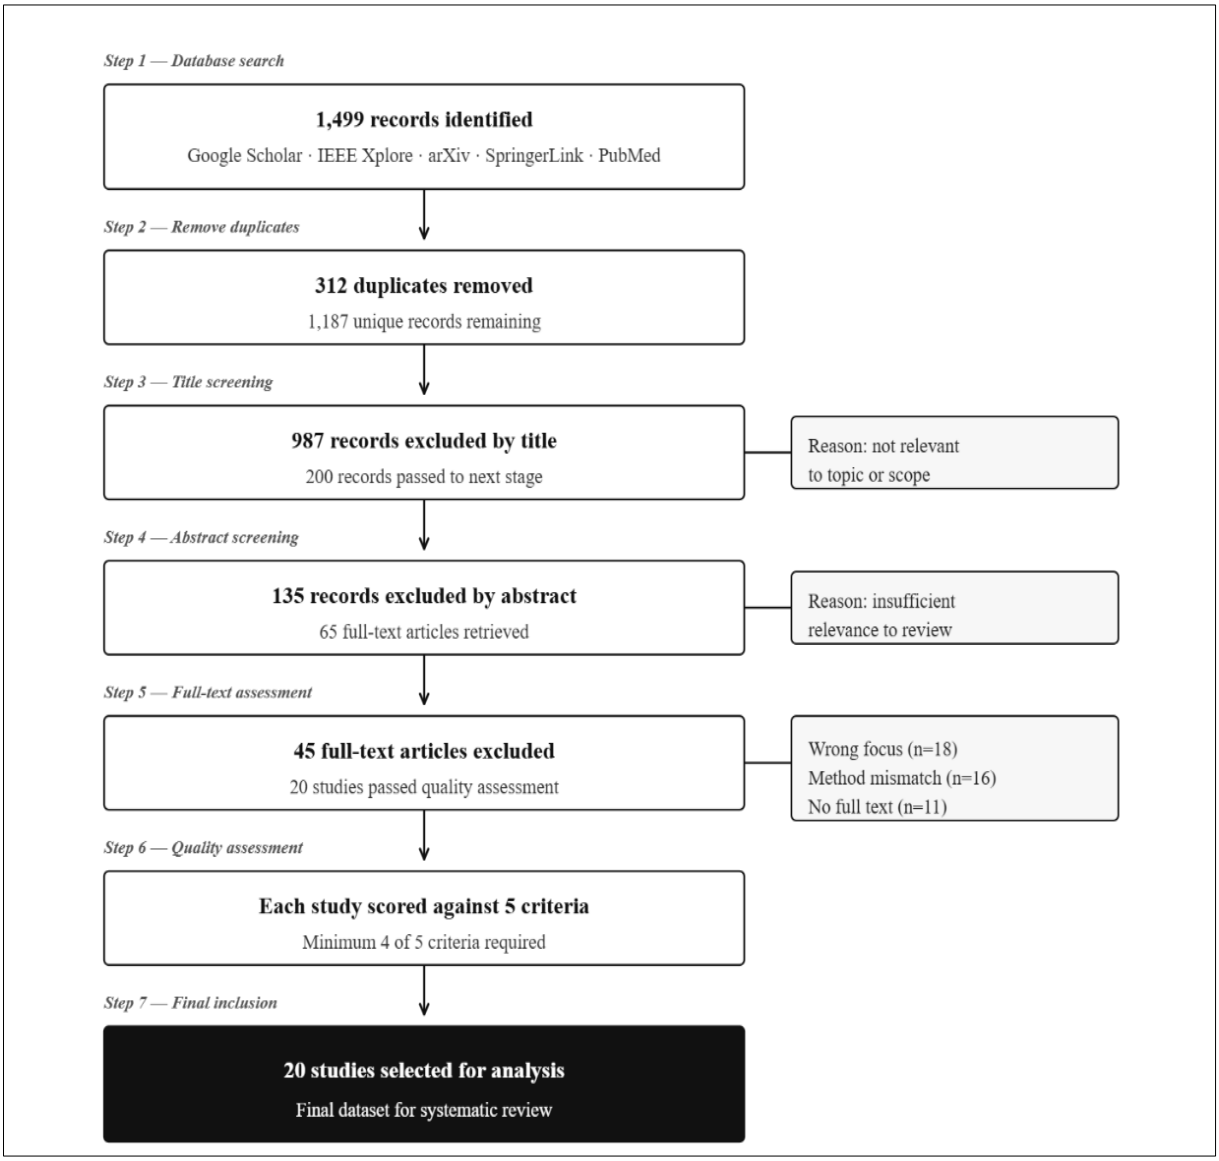

Fig 1 PRISMA Flow Diagram

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1548

➢ *Quality Assessment*

Each candidate study was evaluated against five quality indicators to confirm relevance and methodological soundness:

- Clearly defined research objectives
- Transparent and well-described methodology
- Quantitative or otherwise measurable results
- Explicit discussion of study limitations
- Direct relevance to AI, IoMT, SOAR, or healthcare cybersecurity

Only studies satisfying a minimum of four out of five criteria were retained for final inclusion.

➢ *Data Extraction Process*

A standardized extraction protocol was applied to each selected study, capturing the following dimensions:

- Publication details (year and source)
- AI/ML techniques employed
- Dataset characteristics
- Reported performance metrics (e.g., accuracy, precision)
- Presence or absence of SOAR integration
- Use of Explainable AI (XAI)
- Availability of automated response mechanisms
- Validation within healthcare environments
- Principal findings and reported limitations

IV. LITERATURE REVIEW

➢ *IoMT Security Studies*

The accelerating proliferation of IoMT devices has substantially enlarged the attack surface within healthcare environments, giving rise to layered and increasingly difficult-to-manage security challenges. The reviewed literature thoroughly examines vulnerabilities spanning device, network, and cloud tiers, demonstrating how architectural heterogeneity and the persistence of legacy systems sustain chronic security exposures [1], [2]. Resource-constrained endpoints operating without regular security updates are shown to remain vulnerable for prolonged periods, creating conditions ripe for exploitation.

Several contributions offer structured taxonomies of IoMT threats, categorizing attack types such as spoofing, denial-of-service, and man-in-the-middle intrusions, and systematically mapping these vulnerabilities across sensor and communication layers [3]. Other studies illustrate how compromised devices — infusion pumps and patient monitoring systems, in particular — can serve as footholds for lateral movement across hospital networks, enabling localized breaches to escalate into system-wide incidents [4].

Despite these contributions, IoMT security research remains anchored predominantly in threat characterization rather than threat management. Vulnerability identification and risk mapping, while valuable, do not translate automatically into actionable mitigation. In time-sensitive clinical environments, documenting a vulnerability without a viable containment mechanism affords little practical protection. The downstream consequences of this limitation are significant: compromised medical devices left uncontained can directly disrupt treatment workflows and amplify risks to patient safety.

Collectively, IoMT security studies establish a robust conceptual foundation for understanding the threat environment, but provide limited traction for real-time mitigation or automated incident handling — a gap that the following section explores through the lens of AI-driven detection.

➢ *AI/ML-Based Detection Approaches*

Artificial intelligence and machine learning have become central instruments in healthcare intrusion detection, particularly within IoMT contexts. The reviewed literature encompasses a broad range of methodological approaches — including Random Forest classifiers, gradient boosting techniques, and deep learning architectures — applied to the detection of anomalous patterns in network traffic and device behavior [5], [6]. Performance outcomes are consistently strong, with reported metrics frequently indicating high discriminative accuracy between benign and malicious activity.

Ensemble learning methods have been applied to healthcare-relevant datasets with notable classification performance and computational efficiency [7], while deep learning models — including convolutional neural networks and variational autoencoders — have been deployed to capture intricate traffic patterns and identify previously unseen attack variants [8]. Explainable AI techniques, such as feature attribution methods and local interpretability models, have also been incorporated to improve transparency in detection reasoning [9].

A critical and recurrent limitation, however, constrains the practical value of these advances. Virtually all AI/ML-based systems reviewed terminate at the detection boundary: they produce alerts or class labels, but provide no mechanism for automated containment, remediation, or system recovery. The operational burden consequently remains with human analysts, which undermines the practical utility of even the most technically sophisticated detection pipelines.

A further limitation concerns the experimental conditions under which these models are evaluated. The predominant reliance on static benchmark datasets fails to reflect the dynamic, heterogeneous conditions of live healthcare environments, raising legitimate questions about generalizability. Additionally, the prevalence of binary classification frameworks overlooks the need for severity stratification and contextual reasoning — capabilities that are essential prerequisites for any automated response system. In operational terms, the absence of integrated mitigation means that accurately detected threats may nonetheless go unaddressed within a clinically relevant timeframe, reinforcing the disconnect between analytical capability and operational impact.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1548

➢ *SOAR Frameworks and Automated Response Systems*

SOAR platforms were developed to address the systemic inefficiencies of manual incident response by enabling coordinated, automated security operations. Through the integration of multiple tools and the execution of structured playbook-driven workflows, these systems support consistent and rapid responses to detected threats [10]. Commercial implementations — including Splunk SOAR, IBM QRadar SOAR, and Microsoft Sentinel — have established the viability of response automation within enterprise cybersecurity contexts.

Recent research has explored meaningful enhancements to SOAR capabilities, including adaptive playbook generation, ML-driven response recommendation systems, and graph-based workflow modeling [11], [12]. Some approaches utilize historical incident data to inform dynamic strategy selection, while probabilistic models have been applied to adjust response actions under conditions of uncertainty.

Despite this progress, SOAR adoption within healthcare remains limited and largely theoretical. The majority of reviewed implementations are either conceptual in nature or validated against generic enterprise environments, with minimal accommodation for the distinctive constraints of clinical settings. The absence of healthcare-specific playbooks is a particularly significant gap: response actions appropriate in enterprise IT contexts may pose unacceptable risks in medical environments where automated interventions must be evaluated against patient safety implications.

The scarcity of empirical validation in real hospital deployments further limits confidence in these systems' practical reliability and scalability. Without clinical-environment testing, meaningful assessment of safety, interoperability, and performance under real-world conditions remains impossible. SOAR therefore represents a technically mature but clinically underadapted solution — one whose transformative potential in healthcare cybersecurity has yet to be substantively realized.

➢ *Emerging Trends: LLMs, Zero Trust, and Blockchain*

A set of emerging technologies is attracting growing attention for their potential to enhance cybersecurity within healthcare systems. Large language models (LLMs), in particular, have demonstrated capacity for processing and contextualizing threat intelligence, with recent work indicating their ability to support incident analysis and generate response recommendations by drawing on external knowledge sources [13].

In parallel, Zero Trust architectures have been proposed as a structural security model suited to healthcare environments, enforcing continuous authentication and granular access control to limit the risk of unauthorized access and lateral propagation [14]. Blockchain-based approaches have similarly been explored for their capacity to support tamper-evident audit trails and secure data exchange — particularly pertinent where sensitive medical records are involved [15].

These technologies introduce genuinely promising capabilities, but their integration into operational incident response remains exploratory. LLM outputs are predominantly advisory — capable of recommending actions but not of executing them — and face practical barriers including reliability concerns, hallucination risks, and the complexity of interfacing with operational security systems. Zero Trust and blockchain mechanisms likewise strengthen the broader security posture without directly addressing automated mitigation or real-time incident handling. Their synthesis into a unified, response-capable framework constitutes an open and consequential research challenge.

➢ *Critical Analysis of Existing Literature*

A cross-domain examination of the reviewed literature reveals a consistent and consequential limitation: the near-universal absence of integrated, automated incident response capabilities in healthcare cybersecurity research. While meaningful progress has been achieved in detection techniques, vulnerability characterization, and system design, these contributions seldom extend to encompass the full incident management lifecycle.

The most pervasive issue is the field's disproportionate emphasis on detection accuracy as the primary — and often sole — measure of system performance. While high accuracy is routinely reported, it does not inherently translate into improved security outcomes when response mechanisms are absent or delayed. The practical effect is a research landscape in which threats are identified but not reliably neutralized, leaving underlying vulnerabilities intact.

Equally significant is the failure to embed healthcare-specific considerations into proposed solutions. Many frameworks are designed for generic IT environments and do not account for the clinical realities of medical systems — patient safety obligations, regulatory constraints, and device heterogeneity. This undermines their direct applicability in real-world healthcare contexts. The limited availability of representative datasets compounds these issues, as does the predominant reliance on simulated or controlled experimental conditions. Systems that perform well in constrained test environments may not generalize to the complexity and dynamism of live hospital networks. Taken together, these limitations point to a clear imbalance: detection capabilities have reached relative maturity, while automated, contextually aware mitigation mechanisms remain substantively underdeveloped. Closing this gap demands a deliberate shift in research orientation — toward integrated frameworks that encompass detection, decision-making, and real-time response, designed specifically for healthcare deployment.

V. COMPARATIVE ANALYSIS OF SELECTED STUDIES

Table 3 presents a structured comparison of all 20 reviewed papers across key analytical dimensions, including the technique employed, dataset used, accuracy reported, SOAR integration, automated response capability, healthcare validation status, and primary limitation identified.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1548

Table 3 Summary Comparison of Reviewed Papers (P1–P20)

ID	Year	Technique	Dataset	Acc.	SOAR	Auto Resp.	HC Val.	Key Limitation
P1	2022	ECC security	Simulation	—	No	No	Partial	Encryption focus only
P2	2021	Honeypot + SOAR	Simulated traffic	—	Yes	Partial	No	Not adapted to healthcare
P3	2023	Intent-based SOAR	Conceptual	—	Yes	No	No	No real implementation
P4	2022	IoMT taxonomy	Survey data	—	No	No	Yes	No response mechanisms
P5	2023	Infusion pump SLR	Case study	—	No	No	Yes	No automated mitigation
P6	2024	CatBoost + XAI	CICIoMT2024	>99%	No	No	Partial	Detection only
P7	2022	Blockchain + FL	Conceptual	—	No	No	No	Not validated
P8	2021	Autonomous framework	Simulated	—	Partial	Partial	No	Unrealistic assumptions
P9	2023	SecBERT (NLP)	Threat intel data	—	No	No	No	No response integration
P10	2025	VAE + Random Forest	WUSTL-EHMS	~99%	No	No	Partial	Limited dataset
P11	2024	XGBoost (bi-layer)	SDN-IoMT	99.6%	No	No	Partial	No response pipeline
P12	2025	PSO + PNN	WUSTL-EHMS	~98%	No	No	Partial	Scalability issues
P13	2023	Bayesian graphs	Simulated	—	Yes	No	No	No latency validation
P14	2024	NCF playbook rec.	Enterprise alerts	Prec >0.8	Yes	Partial	No	Not healthcare-specific
P15	2023	Epidemiological AI	Incident data	—	No	No	No	Not healthcare-focused
P16	2024	XAI + Random Forest	IoMT dataset	~99%	No	No	Partial	No response justification
P17	2024	GTPDA model	BoT-IoT	99.5%	No	No	Partial	Detection-focused
P18	2025	LLM + RAG	CTI datasets	High	No	No	No	No execution layer
P19	2024	Systematic review	Literature	—	No	No	Yes	Confirms detection gap
P20	2023	SoK analysis	Literature	—	No	No	No	Fragmented research

➢ *Key Observations*

The comparative analysis surfaces several consistent patterns of note. Between 80 and 90 percent of reviewed studies concentrate exclusively on detection, with no automated response capability present. SOAR integration appears in only four of the twenty papers, and none of these

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1548

is validated within a real healthcare environment. Effective automated response is absent across the entire reviewed corpus. Healthcare validation is frequently partial or limited to simulated conditions, and benchmark datasets predominate over real hospital network data.

These observations collectively and unambiguously confirm the detection–response gap as the field's defining limitation, forming the analytical basis for the discussion that follows.

VI. DISCUSSION AND GAP ANALYSIS

The evidence synthesized across the reviewed studies reveals a clear and consistent pattern: while AI and machine learning have materially advanced threat detection in healthcare cybersecurity, the integration of these capabilities with automated incident response remains substantially underdeveloped. This imbalance is not incidental — it reflects deep structural limitations in how the field has framed the problem and defined success.

Detection performance has reached a relative ceiling of maturity. Multiple studies report accuracy exceeding 95% using Random Forest, XGBoost, and deep learning architectures. Yet this technical achievement is of limited operational value in the absence of corresponding response mechanisms. In live clinical environments, detecting an attack without the capacity to act on that detection does not meaningfully reduce exposure. The preponderance of accuracy-centric evaluation criteria over deployment-oriented metrics indicates a research culture that has optimized for measurable precision at the cost of practical impact.

A complementary observation concerns SOAR frameworks. Although a small number of studies engage with SOAR-based architectures or playbook automation, these contributions are largely conceptual or evaluated in generic enterprise settings. Mature commercial platforms — Splunk SOAR, IBM QRadar SOAR — demonstrate the technical feasibility of automated response in corporate IT; their adaptation to IoMT-specific constraints, however, remains largely unaddressed. The implication is clear: existing orchestration technologies are not yet calibrated to the operational realities of healthcare, and no current research has successfully bridged this divide.

- *SOAR Integration and Response Automation*

The most consequential gap identified is the near-total absence of end-to-end automated incident response pipelines in healthcare. SOAR platforms are mature in general IT sectors but remain conceptually and practically foreign to healthcare IoMT contexts. A critical missing element is the availability of standardized, vendor-interoperable playbooks capable of connecting AI-generated threat intelligence with concrete containment actions — such as dynamic network segmentation or targeted device isolation — within a clinically safe operational envelope.

The security challenges specific to IoMT add a further dimension of complexity that current research consistently underserves. Studies regularly identify device-level vulnerabilities and document risks such as lateral movement and unauthorized access, yet stop short of proposing structured response strategies. This omission is particularly consequential given that IoMT devices operate under strict safety constraints: automated containment of a device actively supporting patient care demands a level of clinical contextual awareness that generic security automation does not possess. The absence of clearly defined, clinically informed response policies represents one of the most significant implementation barriers currently facing the field.

- *Explainability and Regulatory Accountability*

A secondary but equally important gap concerns the limited adoption of Explainable AI in the context of automated response. In regulated clinical environments, high-performing black-box models frequently fail to satisfy the transparency requirements imposed by frameworks such as HIPAA and GDPR. While recent studies acknowledge explainability as a desirable property, few integrate tools such as SHAP or LIME in ways that would provide the interpretability necessary to satisfy clinical accountability standards. Without human-interpretable justifications for automated security actions affecting life-critical systems, meaningful institutional trust — and therefore adoption — will remain out of reach.

This limitation extends to emerging technologies. LLM-based systems, for instance, can generate analytically useful incident response recommendations, but these outputs are advisory rather than executable. Explainability techniques, similarly, are applied to justify detection decisions without extension to the response phase. In regulated environments where both transparency and accountability are prerequisites for deployment, this disconnect poses a formidable adoption barrier.

- *Dataset Diversity and Real-World Applicability*

The analysis further reveals a dataset homogeneity problem with material consequences for model generalizability. Benchmark datasets such as WUSTL-EHMS, while valuable, do not capture the full diversity of modern IoMT attack surfaces — including evolving ransomware variants and advanced persistent threats. The prevalence of centralized training paradigms additionally conflicts with the privacy-preserving imperatives of healthcare institutions. Decentralized architectures — federated learning in particular — represent a natural solution, enabling cross-institutional model development without the exposure of sensitive patient data, yet their integration into incident response frameworks remains nascent.

From an implementation perspective, these dataset limitations compound broader barriers to real-world deployment: regulatory ambiguity, the pervasiveness of legacy infrastructure, and the operational constraints of live hospital environments. Most proposed systems are designed

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1548

and tested in conditions that incompletely reflect clinical reality, limiting their transferability to operational settings.

➢ *Operational Constraints and Clinical Continuity*

A dimension that is systematically underaddressed in the reviewed literature is the operational reality of hospital environments — particularly the prevalence of legacy medical devices running end-of-life systems that cannot accommodate modern security agents. Automated response strategies proposed in the literature rarely engage with the concept of a clinical harm threshold: the point at which a security intervention — isolating a ventilator, for instance, or interrupting an infusion pump — poses a greater risk to the patient than the cyber threat it is intended to contain. Future research must treat this threshold not as an edge case but as a central design constraint, developing safety-aware, low-latency architectures capable of maintaining clinical continuity alongside security integrity.

Notwithstanding these limitations, the reviewed evidence also illuminates a set of tractable opportunities. The field is ripe for integrated frameworks that unify detection, decision-making, and automated response within a single coherent architecture — one incorporating safety-aware logic, healthcare-specific SOAR playbooks, and explainable response models capable of earning regulatory and clinical trust.

VII. THEMATIC ANALYSIS

Cross-examination of the selected literature reveals a defining architectural tendency: the development of increasingly sophisticated AI-driven detection systems alongside a conspicuous failure to operationalize automated response. Ensemble and gradient boosting methods — XGBoost and LightGBM in particular — consistently achieve high detection accuracy, often exceeding 91%, across IoMT-specific threat categories including man-in-the-middle attacks and data spoofing. Meta-heuristic optimization techniques such as Particle Swarm Optimization have further contributed by reducing computational overhead on resource-constrained endpoints through efficient feature selection.

Yet a persistent thematic pattern cuts across this technical progress: cybersecurity is treated as fundamentally a classification problem. Detection systems are refined for precision, but their operational scope concludes at alert generation. What follows detection — the containment, mitigation, and recovery phase — receives negligible attention in the literature. Most frameworks default to manual intervention for threat containment, a dependency that is structurally incompatible with the near real-time demands of clinical environments, where the latency introduced by human-in-the-loop response directly imperils patient safety and data integrity. The detection-response chasm is not merely a technical gap; it is an architectural omission with real-world consequences.

VIII. FUTURE RESEARCH DIRECTIONS

The synthesis presented in this review points unambiguously toward the need for a paradigm shift — from reactive detection toward proactive, automated resilience. Four interconnected research priorities emerge as essential to bridging the identified gaps.

➢ *Development of Safety-Aware SOAR Playbooks*

Future work must move beyond generic orchestration toward automation that is clinically informed. This requires the development of standardized SOAR playbooks designed specifically for IoMT environments — playbooks that embed clinical harm threshold logic, enabling the response system to evaluate device criticality before executing containment. For a life-critical device such as a ventilator, this may mean defaulting to bandwidth throttling or micro-segmentation rather than full isolation. Research into vendor-agnostic abstraction layers will be equally important to ensure interoperability across the heterogeneous infrastructure typical of real hospital deployments.

➢ *Integration of Explainable AI for Clinical Trust*

Overcoming adoption barriers requires that explainability be treated as an architectural requirement rather than an optional feature. Future frameworks should integrate low-latency XAI techniques — such as integrated gradients or computationally lightweight SHAP variants — capable of providing real-time interpretable justification for automated security decisions. Equipping clinicians and SOC analysts with transparent reasoning for automated interventions is not merely a usability enhancement; it is a prerequisite for regulatory compliance and the institutional trust necessary to sustain AI-driven autonomy in clinical environments.

➢ *Privacy-Preserving Collaborative Defense via Federated Learning*

The dataset limitations identified throughout this review call for a structural shift toward decentralized learning architectures. Federated learning enables hospitals to participate in collaborative model training without exposing sensitive patient data — generating more robust, representative models capable of recognizing a broader range of threat variants, including advanced persistent threats and novel ransomware strains. Hierarchical federated learning offers additional promise, enabling efficient load distribution between edge IoMT gateways and cloud-based global aggregators, and warrants dedicated investigation.

➢ *Real-Time Resilience in Legacy Environments*

A substantial portion of the healthcare ecosystem operates on legacy infrastructure that cannot support modern security tooling. Future research should explore edge-based virtual patching and AI-driven transparent proxy architectures as mechanisms for extending automated incident response to end-of-life medical devices without requiring modifications to device software. Lightweight, hardware-accelerated security models will be critical to ensuring these protective layers operate within latency bounds that do not interfere with real-time patient monitoring.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1548

IX. CONCLUSION

This systematic review has critically examined the current state of AI-driven security architectures within the healthcare IoMT ecosystem, synthesizing evidence from 20 high-quality studies published between 2019 and 2026. The overarching finding is unambiguous: while threat detection has reached a commendable level of technical maturity — with ensemble methods and deep generative models regularly achieving high accuracy — the operationalization of automated response within healthcare environments remains profoundly underdeveloped. A persistent detection-response chasm separates the ability to identify anomalies from the capacity to execute safe, autonomous containment

The review's findings make clear that next-generation security frameworks must move decisively beyond detection precision. Automated response in clinical environments requires safety-aware architectures that respect the clinical harm threshold — ensuring that security interventions do not inadvertently compromise patient care. Addressing the dataset homogeneity problem and the interpretability limitations of complex AI models, through federated learning and explainable AI respectively, is equally essential to enabling the institutional adoption that the field currently lacks.

Ultimately, the convergence of automated intelligence with clinical operational requirements is not simply a technical ambition — it is a foundational requirement for protecting patient safety in an increasingly vulnerable digital healthcare landscape. Future work should prioritize healthcare-specific SOAR playbook development, decentralized and privacy-preserving detection architectures, and XAI-enabled response justification mechanisms that together close the gap between detection insight and operational action.

REFERENCES

[1]. S. Kumari, M. Gaikwad, and S. A. Chavan, "Secure IoT-edge architecture with data-driven AI techniques for early detection of cyber threats in healthcare," Discover Internet of Things, vol. 5, no. 1, p. 14, 2025.
[2]. S. H. Almotiri, "AI driven IoMT security framework for advanced malware and ransomware detection in SDN," Journal of Cloud Computing, vol. 14, no. 1, 2025.
[3]. Ramya, Sudhakaran, Sivagnanam, and S. Krishnan, "Advanced intrusion detection technique (AIDT) for secure communication among devices in internet of medical things (IoMT)," Scientific Reports, vol. 14, 2024.
[4]. S. Kaur and Gupta, "Explainable AI assisted IoMT security in future 6G networks," Future Internet, vol. 17, no. 5, p. 226, 2025.
[5]. M. Mohale and O. Obagbuwa, "A systematic review on the integration of explainable artificial intelligence in intrusion detection systems," Frontiers in Artificial Intelligence, vol. 8, 2025.
[6]. R. Jain and A. Singh, "WUSTL-EHMS-2020: A new dataset for healthcare cybersecurity research," IEEE Internet of Things Journal, vol. 7, no. 9, pp. 8801-8812, 2020.
[7]. A. Tellache et al., "Advancing autonomous incident response: Leveraging LLMs and cyber threat intelligence," arXiv preprint arXiv:2508.10677, 2025.
[8]. R. Shinde et al., "Securing AI-based healthcare systems using blockchain technology: A state-of-the-art systematic literature review and future research directions," Artificial Intelligence Review, 2022.
[9]. R. Kremer et al., "IC-SECURE: Intelligent system for assisting security experts in generating playbooks for automated incident response," in Proc. IEEE Conference on Communications and Network Security (CNS), 2023.
[10]. S. Mushtaq, M. Mohsin, and M. M. Mushtaq, "A systematic literature review on the implementation and challenges of zero trust architecture across domains," Sensors, vol. 25, no. 19, p. 6118, 2025.
[11]. M. Yacoubi, O. Moussaoui, and C. Drocourt, "Enhancing IoMT security with explainable machine learning: A case study on the CICIOMT2024 dataset," Information, vol. 16, no. 2, p. 133, 2025.
[12]. J. Paulraj et al., "Autonomous AI-based cybersecurity framework for critical infrastructure: Real-time threat mitigation," PeerJ Computer Science, vol. 11, p. e2414, 2025.
[13]. S. Deb et al., "Securing the internet of medical things (IoMT): Real-world attack taxonomy and practical security measures," Journal of Computer Virology and Hacking Techniques, 2025.
[14]. R. Yener, M. Hassan, and M. Bashir, "Threats and security strategies for IoMT infusion pumps," Healthcare, vol. 10, no. 6, p. 1110, 2022.
[15]. Z. Huang et al., "Toward an intent-based and ontology-driven autonomic security response in security orchestration automation and response," Information, vol. 16, no. 12, p. 1036, 2025.
[16]. U. Bartwal et al., "Security orchestration, automation and response engine for deployment of behavioural honeypots," International Journal of Information Security, 2024.
[17]. Y. Chang, H. Liu, C. Lu, and N. Zhang, "SoK: Security and privacy risks of healthcare AI," JAMA Network Open, vol. 8, no. 5, 2025.
[18]. S. Abraham et al., "AI incident monitoring through a public health lens," Journal of Cyber Policy, vol. 4, no. 3, 2025.
[19]. FDA, "Cybersecurity in medical devices: Quality management system considerations and content of premarket submissions," U.S. Food and Drug Administration Guidance Document, Feb. 2026.
[20]. H. Wang and J. Liu, "Deep learning for real-time anomaly detection in IoMT gateways," IEEE Journal of Biomedical and Health Informatics, vol. 27, no. 5, 2023.