Cybersecurity in Robotic Healthcare Systems: A Critical Thematic Review of Challenges, Solutions, and Research Directions

Hadi Dastan Elikhchi
University of Greater Manchester

Thaier Hamid, PhD
University of Greater Manchester

ABSTRACT

Nowadays, Robotic systems are increasingly deployed in the healthcare system to enable modern autonomous delivery, real-time diagnostics, and precision in surgery. However, these contemporary and intelligent machines, connected into complex systems which introduce new cybersecurity challenges that can impact operational availability, data integrity, also threaten patient safety. This review critically examines the recent robotic healthcare systems' current landscape of cybersecurity. By utilizing the CIA security model (Confidentiality, Integrity, Availability) and socio-technical systems theory. Therefore, we organize the literature into three thematic domains such as: behavioral anomaly detection using machine learning, access control, and secure communication. Each one of the domains is evaluated for its integration potential, limitations, real-world applicability, and theoretical foundations. This review reveals significant research gaps in deployment realism, regulatory alignment, and resilience. This study proposes a comprehensive research direction and agenda to lead future work toward layered and robust ethical cybersecurity frameworks for clinical robotic systems and environments.

General Terms

Secure Architecture, Cybersecurity, Intrusion Detection, Healthcare Technology, Clinical Safety, Machine Learning, Robotic Systems.

Keywords

Healthcare Robotic Systems, Role-Based Access Control, Human-Centered Security, Resilience, Behavioral Anomaly Detection, GRU, TLS 1.3, Secure Communication.

1. INTRODUCTION

Healthcare delivery is experiencing a revolutionary shift, spurred by the repeatedly increasing integration of robotic technology [22]. These technologies and systems are increasingly deployed to perform monitoring patient vitals, automate hospital logistics, and perform high-precision surgical procedures with minimal human intervention. However, such modern technologies are typically embedded and network-enabled with complex, entrusted systems like artificial intelligence (AI) to handle sensitive medical data, positioning them within the broader category of cyber-physical systems (CPS) and uncertainties. These advanced technologies embody a tight coupling between real-world physical processes and computational logic which creates new intersections between human health and digital infrastructure outcomes [21].

Unlike traditional breaches in robotic healthcare systems or information technology failures can have consequences far beyond data loss. In this case, any Cyberattacks may lead to outages during critical medical interventions, incorrect dosing of medication, or even unauthorized surgical maneuvers. This type of scenario elevates cybersecurity from a purely technical concern into a direct impact on medical ethics, operational continuity and a matter of clinical safety. To address and understand the current state of security in this evolving field, this study conducted a critical thematic review of cybersecurity literature by highlighting and focusing on robotic healthcare systems.

Furthermore, the analysis identifies three principal domains of concern in anomaly detection systems leveraging machine learning, access control architectures, and secure communication protocols. These domains signify the most rigorously examined vectors for safeguarding robotic technology which functions in clinical and hospital environments.

This study employs two conceptual frameworks to guide its analysis, such as: three essential pillars for protecting digital systems of confidentiality, integrity, and availability, which well-established CIA triad, which frames security objectives through the lenses. The second is the socio-technical systems model, which emphasizes that implementation cannot be isolated from human and technological design, and contextual and organizational factors as well. These two frameworks and models provide a multidimensional lens for evaluating the integration and applicability of cybersecurity mechanisms in real-world clinical environments as well.

Therefore, by critically applying literature and synthesizing academic research, the review aims to address, clarify the current weaknesses and strengths of cybersecurity strategies in robotic healthcare systems. This also leads to key outlines to propose a forward-looking agenda by highlighting research gaps to support the development of more operationally viable, ethically grounded, and resilient cybersecurity frameworks for clinical robotics systems.

2. METHODOLOGY

This paper employs critical thematic literature review synthesis approaches, surveys and evaluates cybersecurity issues, strategies within robotic healthcare environments. This review is grounded in a systematic literature search conducted across different academic databases, including PubMed, ScienceDirect, SpringerLink, ACM Digital Library, and IEEE Xplore. A defined set of keyword combinations, such as "GRU-based anomaly detection", "RBAC hospital access control," "TLS 1.3 in robots," and "robotic healthcare cybersecurity," was used to ensure a comprehensive retrieval of relevant literature across interdisciplinary domains. To ensure scholarly rigor, only peer-reviewed publications were selected from 2015 to 2025 written in English that were considered.

Those studies were included if they highlighted cybersecurity measures in the context of robotic systems operating in hospital or clinical settings. Furthermore, ensure scholarly rigor by only selecting peer-reviewed publications from 2015 to 2025 written in English that were considered. Those studies were included if they highlighted cybersecurity measures in the context of robotic systems operating in hospital or clinical settings. However, those studies were excluded if they did not explicitly engage with cybersecurity frameworks or implementations, offered unvalidated theoretical models, or lacked a focus on healthcare systems.

Therefore, out of an initial pool exceeding 260 records, 65 publications were selected following the title: abstract, and full-text screening, like analyzed and classified into three different principal themes, such as: machine learning-based anomaly detection, access control mechanisms, and secure communication protocols.

Those classifications were informed by the conceptual alignment of each research study within the CIA security triad model, and were examined for integration, real-world applicability, and feasibility with clinical practices. This methodology was designed and structured evaluation of a highly interdisciplinary domain while preserving its pertinence to both healthcare systems and technical challenges.

3. SECURE COMMUNICATION IN ROBOTIC HEALTHCARE

3.1 Rationale for Secure Communication

Robotic technology platforms in hospitals rely heavily on data communication through network connections to both external (remote commands, cloud diagnostics) and internal (sensor-actuator loops). These types of data communications on networks often transmit control signals and sensitive patient data. If tampered with or intercepted, these data communications can be used to: Extract sensitive patient or procedural information, Hijack surgical sessions, and replay legitimate commands (causing denial-of-service or incorrect operations).

3.2 TLS 1.3 and Cryptographic Safeguards

Transport Layer Security (TLS) is known for standard cryptographic protocol which directly involves securing data in transit, like cyber-physical systems (robotic healthcare environments). With its ratification in 2018, TLS 1.3 introduced several performance and security improvements over its predecessors, and TLS 1.2. Additionally, it reduced the handshake latency through 1-RTT and 0-RTT mechanisms and eliminated obsolete cryptographic algorithms, such as ensuring forward secrecy and RSA key exchange using ephemeral Diffie-Hellman key exchanges as well.

Transport Layer Security (TLS 1.3) offers crucial latency improvements for the surgical robotics systems, within research by Naylor et al. [7] and Sun et al. [1] confirming its viability when combined with trust anchors or even hardware acceleration. In addition, its adoption in robotic middleware such as DDS, ROS 2 remains limited, hindered by performance and integration constraints. Furthermore, clinical deployment requires further testing under legacy conditions and a real-world network.

3.3 Challenges and Trade-offs

Notwithstanding the benefits of implementing TLS 1.3 for safeguarding communication in robotic healthcare systems, numerous problems remain in its execution. Latency is known as one of the main issues. Although TLS 1.3 provides faster handshakes than its predecessors, in high-stakes situations like surgical robots, even small delays of 20 to 30 milliseconds might cause real-time operations to be disrupted.

Furthermore, another issue lies in cryptographic key management. In most of the dynamic robotic systems networks, where nodes frequently leave or join, managing certificate revocation and the process of rotating encryption keys becomes increasingly complex. In addition to increasing administrative workloads, if not done perfectly, this increases the possibility of security breaches. Legacy compatibility is still a major obstacle, too.

In this case, most of the hospital systems continuously use outdated software and device environments which are incompatible with advanced and modern cryptographic libraries needed for TLS 1.3.

This type necessitates either system upgrades that are time-consuming and may be costly, or the development of backwards-compatible security solutions, which might potentially endanger robustness.

For robotic healthcare systems, TLS 1.3 is the most secure transport protocol; nonetheless. This is not always the most practical in latency-sensitive environments. At the same time, TLS 1.2 is older and continues to be extensively used in hospitals because of its wider device support and lower integration cost. In contrast, Datagram TLS (DTLS), which avoids connection-oriented handshakes and thus lowers latency, has been suggested as a lightweight substitute for Internet of Things-based medical equipment. IPSec offers integrity and strong confidentiality at the network layer but adds significant overhead that is inappropriate for robotic control in real time.

Comparative evaluations reveal that TLS 1.3 offers integrity and the strongest guarantee of confidentiality, but it can disrupt availability when facing handshake delays or denial-of-service. In surgical robotics systems, disruptions of any size can jeopardize patient safety, making availability equally as important as confidentiality. As a result, selecting a cryptographic protocol must strike a balance between theoretical resilience and operational limits in clinical settings.

4. ACCESS CONTROL MECHANISMS IN ROBOTIC ENVIRONMENTS

4.1 Role-Based Access Control (RBAC)

RBAC is a well-established approach and mature model in which individuals are given roles with certain rights, such as administrator, technician, or surgeon. This is helpful in robotic systems for limiting authorized physicians to high-level control, such as: Preventing non-technical people from updating software and auditing or recording user activity.

According to Ferraiolo and Kuhn [2], RBAC's simplicity makes implementation easier, and it fits in nicely with hierarchical medical processes.

4.2 Dynamic and Context-Aware Models

Conventional RBAC fails to consider contextual elements like device state, time, or location. Recent works in RBAC involve several mechanisms such as emergency role escalation, geo-fencing and temporal constraints, Attribute-Based Access Control (ABAC), which integrates environmental variables into access determinations, where RBAC permissions are temporarily modified under specified operational or clinical circumstances. Chen et al. [10] have introduced hybrid RBAC-ABAC models for Internet of Things sensor inputs for robotic systems for nursing assistants, achieving a better balance between usability and security.

4.3 Threats and Limitations

A significant danger to access control in robotic healthcare systems is the potential for insider misuse, frequently resulting from compromised credentials. Role-Based Access Control (RBAC) models, although proficient in allocating permissions, are deficient in procedures for identifying unauthorized actions post-access approval. Furthermore, the growing implementation and adoption of context-aware RBAC engenders considerable policy complexity. At the same time, administrators must take control and manage dynamic rules which depend on environmental variables, which can lead to unintended privilege escalation and configuration conflicts.

Overall, enforcing context-aware regulations and policies in real time incurs additional computing complexity. This might cause delays in access validation during time-sensitive clinical processes, potentially impacting system reliability and affecting patient care.

RBAC is conceptually basic and matched with medical hierarchies. Its rigidity frequently contrasts with emergency situations in which access must be granted immediately. As an example, during an unexpected complication, a surgeon may need immediate access to robotic controls, which RBAC alone does not allow. ABAC models highlight this by considering contextual characteristics such as device status, location, and time. However, adaptation comes at the expense of policy complexity and administrative burden.

Real-world evaluations demonstrate that hybrid ABAC-RBAC models provide better performance and usability but are rarely adopted in health care systems or hospitals due to their maintenance difficulty and burden of ensuring policy correctness. Therefore, access control models are theoretically mature; their practical integration into robotic middleware such as DDS or ROS2 is still in its early stages and underdeveloped.

5. MACHINE LEARNING FOR INTRUSION AND BEHAVIOR ANOMALY DETECTION

5.1 ML in Cyber-Physical Systems

Traditional signature-based intrusion detection systems (IDS) in robotic systems fail to detect subtle or novel threats. Machine learning models, particularly Gated Recurrent Units (GRUs) and Recurrent Neural Networks (RNNs), can model temporal dependencies in enabling and behavior detection of: Deviations in actuator-sensor cycles, Network traffic anomalies, control command injection. Salim et al. [4] have trained a GRU model on normal robotic system task sequences, which identified attacks with 97% accuracy and low false positives as well.

5.2 Limitations and Deployment Concerns

While ML-based intrusion detection systems offer significant promises in enhancing robotic healthcare cybersecurity, numerous constraints must be overcome before broad use. A primary concern is elucidation.

In most current models, particularly deep neural networks, function as "black boxes," which makes it difficult for security and clinicians and auditors to defend their choices in safety-critical scenarios. Therefore, the lack of transparency hinders and undermines trust in healthcare settings. Furthermore, the paucity of high-quality, real-world labelled datasets for robotic cyber-physical systems hinders successful validation and training.

Unlike standard IT systems, robotic systems generate complex and context-sensitive data which are hard to annotate or even simulate. Additionally, adversarial ML represents an emerging threat as attackers can deliberately craft input sequences which cause the model to misclassify risky activity as harmless. Similar IoT-based healthcare studies reinforce these risks, highlighting privacy and intrusion vulnerabilities across connected medical devices [11]. It addresses the need for ongoing model adaptation and robust defense systems.

In response to these challenges and complexity, researchers have begun developing different methods and models, like hybrid models, which combine ML with rule-based logic, and are also exploring explainable AI (XAI) frameworks, for example, SHAP and LIME, to improve interpretability. These initiatives, while promising, must be confirmed in actual clinical settings.

Comparative studies demonstrate that while GRU-based models excel at capturing temporal dependencies in robotic processes, they are challenging and resource-intensive to retrain in dynamic environments. At the same time, RNNs offer the same temporal modelling, but frequently suffer from vanishing gradient difficulties, reducing long-term accuracy. CNNs, though not sequence-based, have demonstrated effectiveness in classifying network intrusion patterns at a lower computing cost. Each strategy has some advantages, but none of them ensures consistent distribution across diverse healthcare settings.

Importantly, outcomes from laboratory simulations rarely translate into clinical practice. Models trained on synthetic datasets are very accurate (>95%), but when exposed to noisy hospital data, performance suffers dramatically. Furthermore, adversarial ML attacks expose a new vulnerability frontier in which attackers can delicately modify inputs to confuse detection systems.

Table 1. Evaluation of GRU, RNN, and CNN models for robotic intrusion detection

Algorithm	Lab	Hospital	Hospital
GRU	97%	75%	High
RNN	94%	70%	Medium
CNN	92%	73%	Low

6. RESILIENCE, ETHICS, AND REGULATORY CONSIDERATIONS IN ROBOTIC CYBERSECURITY

6.1 Resilience Engineering in Clinical Robotics

In robotic healthcare systems, resilience is the ability to continue critical operations in the face of system failures or cyber disturbances. While security efforts often focus on detection and prevention, the ability to operate safely and recover under compromised or degraded conditions remains undeveloped. Clinical robots must be capable of supporting manual intervention, maintaining critical operations, and isolating affected components when needed. Although Models like dynamic node quarantine in robotic swarms show potential, they are rarely used in hospital-grade systems. Recent work on resilient robotic swarms demonstrates adaptive strategies for containing failures, which could inspire clinical robotics resilience [17]. Additionally, recent middleware lacks native support for resilience features such as failover or reconfiguration, regardless of the existence of supporting technologies like SDN. To improve resilience, recovery methods must be integrated as essential design features rather than as after-the-fact precautions.

Due to the high financial cost, certification requirements, and lack of middleware support, hospitals rarely implement resilience engineering solutions. Many robotic systems are certified for modifications and static configurations to enable resilience characteristics can invalidate compliance.

6.2 Ethical Dimensions: Transparency, Explainability, and Trust

Cybersecurity in healthcare robotics systems poses serious ethical issues. Such decisions must be grounded, transparent in clinical and ethical reasoning, also open to post-event review. Opaque ML models, particularly those utilized in anomaly detection systems, have the potential to undermine clinician trust since their outputs are neither interpretable nor accountable [23]. According to Kesarwani et al., the lack of explainability is a significant obstacle to adoption. To guarantee that safety and trust remain key to deployment, ethical integration of cybersecurity in robots necessitates systems that are understandable, auditable, and designed with active physician participation.

The contrast between the opaque requirement for explainability and black-box ML models depicts a basic ethical dilemma. Clinicians frequently avoid adopting models that they cannot interpret, particularly when patient safety is at risk. Reasonable AI methods such as LIME and SHAP offer partial solutions but incur additional computing expenses. It is imperative to thoroughly examine this trade-off between efficiency and interpretability to win therapeutic trust.

6.3 Regulatory Compliance and Legal Frameworks

In healthcare, robotic systems must adhere to established legal limitations such as GDPR [6] and HIPAA, which require health data protection, breach responsibility and informed consent.

The FDA’s guidance on cybersecurity for networked medical devices also emphasizes continuous monitoring and patch management as essential safeguards [5]. At minimum, compliance with international security standards such as ISO/IEC 27001 [25] provides a structured baseline for assurance. However, these frameworks provide minimal advice for self-driving robotic systems, particularly those employing embedded artificial intelligence. The FDA's expanding cybersecurity requirements are a crucial step forward, but they remain vague in terms of actual enforcement in robotic agents. Privacy-preserving frameworks, such as those discussed by Hossain et al. [14], can complement encryption and access-control approaches in mitigating risks. Scholars are advocates for more robotics-specific legislation, highlighting the necessity of ongoing integrity assurance, built-in authentication, and real-time legal tracking methods.

Adoption is further restricted by regulatory framework conflicts. HIPAA allows exceptions for life-critical interventions, while GDPR emphasizes strict user consent and data minimization. Disputes between Cybersecurity guidelines tailored to individual devices have been introduced by the FDA; nevertheless, their implementation in robotic systems is unclear, also restricting adoption. This imbalance makes compliance more difficult and frequently deters hospitals from integrating modern cybersecurity.

7. COMPARATIVE ANALYSIS OF CYBERSECURITY STRATEGIES

In the domain of healthcare robotics systems, Cybersecurity cannot be seen as a standalone technological issue; rather, it necessitates a holistic approach that simultaneously addresses regulatory compliance, real-time responsiveness, operational efficiency, and confidentiality. To this end, the comparative framework presented below evaluates existing cybersecurity strategies across four different critical dimensions: (1) demands of real-world healthcare settings, and clinical readiness, or the extent to which each approach aligns with the constraints; (2) robustness, the transparency of system behavior and incorporating resilience to failures; (3) practical deplorability, infrastructural demands of each method and reflecting the computational, and (4) security coverage, as conceptualized by the CIA triad (Confidentiality, Integrity, Availability). Instead of favoring one approach over another, this synthesis aims to draw attention to their complementary advantages and provide guidance for integrating them into a unified, multi-layered security architecture.

The CIA triad is not equally covered when strategies are compared. Encryption protocols such as TLS 1.3 focus heavily on integrity and confidentiality but provide minimal assistance with availability in the event of a denial-of-service attack. Access control mechanisms reinforce availability and integrity, restrict users, yet they are susceptible to credential compromise. Although machine learning anomaly detection algorithms improve availability by seeing threats instantly, they frequently lack transparency, which erodes confidence.

This unequal focus shows that no one approach can fully address every aspect of CIA. The dearth of availability-focused research is especially troubling because robotic systems need to continue operating without interruption throughout crucial processes. Thus, for complete protection, layered designs including anomaly detection, access control, and encryption are crucial.

Table 2. Summary of cybersecurity strategies in robotic healthcare

Strategy	CIA Focus	Efficiency	Readiness
TLS 1.3	Confidentiality, Integrity	Fast, medium setup	Medium–High
RBAC / ABAC	Integrity, Availability	Lightweight. scalable	Low–Medium
GRU Based IDS	Integrity, Availability	Costly slow retrain	High
Hybrid Detection + RBAC	All three (via layers)	Complex. moderate latency	High
Resilience Engineering (SDN)	Availability	Heavy. dynamic control	High

8. RESEARCH GAPS AND AGENDA FOR FUTURE WORK

The swift implementation of autonomous robotic systems in healthcare environments has revealed numerous fundamental gaps in cybersecurity research and practice. While technological advancements such as machine learning-based intrusion detection, access control, and encryption protocols have individually improved system robustness, these initiatives frequently operate in isolation and do not establish a unified defense strategy. To realize trustworthy, safe, and resilient robotic systems in clinical environments, a more forward-looking and integrative research agenda must be adopted.

8.1 Toward Integrated, Layered Architectures

Current cybersecurity researchers in health robotics often threaten anomaly detection, access control, and encryption as isolated components. While Machine learning-based anomaly detection can detect and identify unusual behaviors, but typically lacks integration with upstream controls, limiting its ability to prompt timely mitigation. Similarly, RBAC enforces static permissions but lacks the contextual adaptability required in dynamic clinical environments. Furthermore, TLS 1.3 secures data in transit; it cannot prevent misuse by compromised but authorized users.

Additionally, highlighting this siloed structure requires an interoperable response that functions autonomously yet coordinates dynamically, behavioral monitoring, access control, and a layered architecture in which each security layer communicates [20]. For example, behavioral anomalies must trigger real-time adjustments to user privileges and encryption keys. Such cross-layer coordination is necessary to embrace and build adaptive security and resilient frameworks for robotic systems in clinical settings.

8.2 The Need for Adaptive Intelligence in Cyber Defense

Another noteworthy area of the research gap lies in the static nature of most deployed machine learning models in robotic intrusion detection systems (IDS). In general, these models are typically trained offline using historical data, which makes them more vulnerable to adversarial evolution and concept drift. In general, these models are typically trained offline using historical data, which makes them more vulnerable to adversarial evolution and concept drift.

In dynamic hospital settings where workflows change and use patterns change, static models deteriorate in reliability and accuracy. A future-ready research direction involves federated learning techniques and developing online which enable robotic systems to continuously update their threat models while maintaining low-latency inference and maintaining privacy. Furthermore, context-aware feedback loops and adaptive confidence thresholds can make such systems more acceptable in high-stakes clinical settings and reduce false alarms.

8.3 Open, Reproducible Evaluation Benchmarks

An ongoing cybersecurity constraint literature for robotics is the absence of reproducible benchmarks and shared datasets. Moreover, studies of cyber threats in eHealth underline how these gaps undermine readiness for real-world deployments [18]. Without a representative testbed, standardized metrics, it is nearly impossible to conduct longitudinal analysis, validate claims, or compare approaches.

Simulation environments which do exist are often undermining collaborative progress, limited in scope or proprietary. To mitigate this, future work must consider this by prioritizing the creation of annotated datasets, domain-specific, it includes system logs from robotics, attack signatures, and realistic traffic patterns of networks operating in clinical contexts. Additionally, open-source simulation toolkits which mimic multi-robot hospital workflows should become a commonplace fixture in experimental design.

8.4 Embedding Human-Centered Design in Security Policies

While there are numerous technical solutions, the human elements remain underexplored. Robotic systems must interact seamlessly with hospital IT staff, patients, and clinicians. Security engineering principles emphasize that system design must balance robustness with usability to avoid such failures [19]. If a security mechanism causes issues, delays, or lacks interpretability, generates excessive false positives, it is likely to be circumvented, disabled, or ignored. In this case, security is not an option; it is vital. Future research should integrate methods from human-computer interaction (HCI), workflow engineering, cognitive psychology to determine which security systems are structured and designed with clinical realities in mind. It includes studying how clinical systems mentally model robotic behavior, how override options are exercised, and alerts are perceived during security-related events [24].

8.5 Regulation-Aware System Engineering

Robotic systems are handling personal data and health information must conform to stringent legal frameworks such as emerging robotic ethics policies, GDPR, and HIPAA. However, regulatory compliance is frequently seen as a post-development issue, which leads to ad hoc adaptations that are potentially inefficient and noncompliant.

Therefore, future systems must embody data minimization strategies, user consent modules, transparent decision-making, incorporate real-time audit trails, and be compliant-by-design.

Additionally, collaboration with regulatory bodies and legal experts during system architecture planning should become standard practice and not only an exception.

A recurring deficiency in the literature is the preference for simulation-based assessments over actual hospital implementations. Although few frameworks have been tested in clinical contexts where noise, legacy systems, and workflow difficulties predominate, academic research frequently show great performance under controlled datasets. This theory-practice divide explains why cybersecurity frameworks have been sluggish to go from research to practical healthcare robotics. Establishing cooperative testbeds with hospitals, researchers, and regulators is necessary to close this gap.

Below is a proposed architectural model (schematic) illustrating how encryption, resilience mechanisms, anomaly detection, and access control interact dynamically within an autonomous robotic system in a hospital setting.

Figure 1 illustrates such a term as the Layered Cybersecurity Framework, conceptual model. It emphasizes the control across security modules and the flow of data reinforcing defense-in-depth principles.

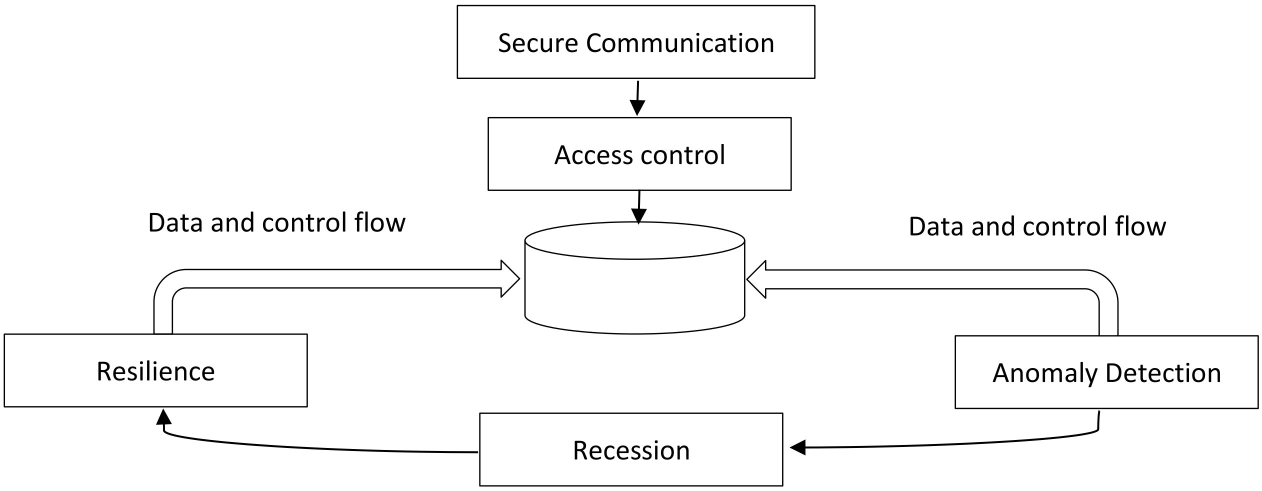

Fig 1: Layered Cybersecurity Architecture for Healthcare Robotics

9. CONCLUSION

Robotic systems provide tremendous possibilities for healthcare, but their ethical and safe deployment hinges on robust cybersecurity. As these platforms become more responsible for sensitive clinical functions, network-integrated, and increasingly autonomous, the same time become potential targets for sophisticated cyber threats.

This review has critically examined three core domains of robotic cybersecurity: behavioral anomaly detection, access control, and secure communication by addressing both unresolved vulnerabilities and technical advances. Importantly, the research goes beyond technical solutions to address regulatory compliance, ethical transparency, and resilience engineering.

To promote safe clinical robotics systems, the field must adopt human-centered governance, adaptive intelligence, layered defense strategy that tightly integrates cryptographic strength, and multidisciplinary. Only through this convergence can we ensure that healthcare robotics enhances care without compromising accountability, trust, or safety.

10. REFERENCES

[1] Sun, Y., et al. (2020). "TLS 1.3 Deployment in IoT." IEEE Internet of Things Journal.

[2] Ferraiolo, D., and Kuhn, D.R. (1992). "Role-Based Access Control." NIST.

[3] Khan, M., et al. (2019). "ML-Based Anomaly Detection in Cyber-Physical Systems." Computers & Security.

[4] Salim, R., et al. (2022). "GRU-Based IDS for ROS2." IEEE Transactions on Industrial Informatics.

[5] U.S. FDA. (2018). "Cybersecurity for Networked Medical Devices."

[6] European Union. (2016). "General Data Protection Regulation (GDPR)."

[7] Naylor, D., et al. (2019). "Performance of TLS 1.3 on Embedded Devices." SIGCOMM.

[8] Ahmed, M., et al. (2016). "Survey of Network Anomaly Detection." J. Network & Comp. Apps.

[9] Zhang, X., et al. (2020). "Fine-Grained Access Control in Healthcare Robotics." IEEE Access.

[10] Chen, T., et al. (2021). "Context-Aware RBAC Using IoT Sensors." J. Biomed. Informatics.

[11] Habibzadeh, H., et al. (2020). "IoT Security in Smart Healthcare Systems." Sensors.

[12] Kesarwani, A., et al. (2018). "Explainable AI for Healthcare." Nature Machine Intelligence.

[13] Rahman, F., et al. (2021). "Secure ROS2 Architecture." IEEE Embedded Systems Letters.

[14] Hossain, M., et al. (2018). "Privacy-Preserving Frameworks in Medical IoT." FGCS.

[15] Shafiq, M., et al. (2017). "Real-Time IDS with RNN for CPS." Information Sciences.

[16] Ristic, I. (2019). Bulletproof TLS and PKI. Feisty Duck.

[17] Boulton, M., & Krishnamachari, B. (2023). "Resilient Robotic Swarms." IEEE Transactions on Robotics.

[18] Casola, V., et al. (2021). "Cyber Threats in eHealth." IEEE Internet of Things Journal.

[19] Anderson, R. (2008). Security Engineering: A Guide to Building Dependable Distributed Systems.

[20] McGraw, G. (2006). Software Security: Building Security In. Addison-Wesley.

[21] Lee, E.A., and Seshia, S.A. (2017). Introduction to Embedded Systems: A Cyber-Physical Systems Approach.

[22] Yang, G.Z., et al. (2018). "Medical Robotics: Current Status and Future Trends." IEEE EMBS Magazine.

[23] Wright, A., and Sittig, D. (2008). "Clinical Decision Support and Malpractice Risk." JAMA.

[24] Van der Aalst, W.M.P. (2011). Process Mining: Discovery, Conformance and Enhancement of Business Processes.

[25] ISO/IEC 27001. (2013). "Information Security Management Systems.