



# A Critical Review: Improve Data Privacy and Security in Electronic Health Records Using Advanced Machine Learning Models

Kshitiz Agarwal <sup>1</sup>, Sandhya Sharma <sup>2</sup>

<sup>1</sup> Research Scholar, Department of Electronics & Communication Engineering, Suresh Gyan Vihar University, Jaipur, India

<sup>1</sup> Assistant Professor, Department of Electronics & Communication Engineering, Suresh Gyan Vihar University, Jaipur, India

Email: agarwal\_ksh@yahoo.com, sandhya.sharma@mygyanvihar.org

**Abstract**— Electronic Health Records (EHRs) have become the backbone of digital healthcare systems, transforming the way medical data is created, stored, and exchanged. They integrate diverse patient information, from demographic profiles and diagnostic histories to laboratory results and imaging studies, offering clinicians a holistic view of patient care. However, this shift from paper-based systems to large-scale digital repositories has also introduced unprecedented challenges. The sensitivity of health data, coupled with its economic and social value, makes EHRs attractive targets for cyberattacks, insider misuse, and unauthorized surveillance. Traditional safeguards—such as role-based access control, encryption, and anonymization—remain essential components of security strategies but often fall short in the face of adaptive, large-scale, and cross-institutional threats. The increasing frequency of ransomware attacks and massive data breaches highlights the limitations of static defences. In this context, machine learning (ML) is emerging as a transformative solution. By learning patterns of normal and abnormal behaviour, ML can detect anomalous access, enforce adaptive privacy rules, and enable privacy-preserving computation across distributed datasets. This review paper examines eight representative studies that demonstrate the role of advanced ML techniques in improving EHR privacy and security. We discuss approaches ranging from anomaly detection and ontology-driven reasoning to privacy-preserving machine learning methods such as differential privacy, homomorphic encryption, and federated learning. We also examine integration with blockchain networks, the use of IoT-based multilayer learning, and the rise of AI-driven cybersecurity. Our analysis highlights strengths, limitations, and trade-offs among these approaches, while also identifying gaps in scalability, explainability, and regulatory compliance. The paper concludes that hybrid frameworks—combining anomaly detection, ontology-based reasoning, privacy-preserving ML, and distributed trust mechanisms—hold the greatest promise for securing EHRs. Looking forward, explainable AI, blockchain-ML fusion, multimodal data integration, and quantum-resistant security are expected to shape the next generation of privacy-aware healthcare systems.

**Keywords**— Electronic Health Records, Internet of Things, Digital Healthcare System, Machine Learning, Anomaly Detection.

## I. INTRODUCTION

The digital transformation of healthcare has been one of the most significant technological shifts in recent decades. At the heart of this transformation lies the Electronic Health Record (EHR), a digital repository that consolidates patient information across time, settings, and providers. EHRs evolved from earlier electronic medical records (EMRs), which were largely institution-specific, into interoperable systems that facilitate data sharing across hospitals, clinics, and research institutions. Today, EHRs not only support day-to-day clinical care but also underpin large-scale health analytics, public health surveillance, and personalized medicine.

The benefits of EHR adoption are undeniable. Clinicians gain immediate access to comprehensive patient data, enabling faster and more accurate diagnoses. Patients can access their medical histories and become more active participants in their care. Health administrators and policymakers leverage aggregated EHR data to identify disease trends, allocate resources, and design preventive strategies. However, the same features that make EHRs valuable also make them vulnerable. Unlike financial records, which are largely transactional, health records contain longitudinal, highly personal data that can reveal a patient's identity, medical conditions, lifestyle choices, and even genetic predispositions. This makes them a lucrative target for



cybercriminals and a potential tool for discrimination if misused [1].

Globally, the scale of healthcare data breaches is alarming. The 2015 Anthem breach in the United States exposed nearly 79 million patient records, making it one of the largest healthcare breaches to date. In 2021, Ireland's Health Service Executive was crippled by a ransomware attack that disrupted patient care nationwide. More recently, the 2024 Change Healthcare breach in the U.S. affected an estimated one-third of Americans, underscoring the fragility of centralized health infrastructures. Beyond external attacks, insider misuse—where employees access records out of curiosity or for malicious purposes—remains a persistent and harder-to-detect threat.

Traditional defenses, such as role-based access control (RBAC), encryption protocols, and anonymization strategies, remain crucial. Yet they were designed for environments with relatively static risks. Today's healthcare systems are highly interconnected, dynamic, and data-intensive. Patients expect their data to be accessible across providers, researchers require datasets for analysis, and regulators impose strict compliance requirements under laws such as the Health Insurance Portability and Accountability Act (HIPAA) in the U.S. and the General Data Protection Regulation (GDPR) in Europe. Meeting these demands requires security measures that are both robust and flexible—qualities that conventional methods alone cannot guarantee [2].

This is where machine learning (ML) enters the picture. Unlike static rules, ML algorithms can adaptively learn patterns of legitimate and illegitimate behavior. They can flag anomalous access requests, classify privacy policies, and even enable collaborative learning across institutions without sharing raw data. Advanced methods such as federated learning, differential privacy, and homomorphic encryption are redefining what it means to compute securely on sensitive health data. Moreover, the integration of ML with technologies like blockchain and IoT extends security from centralized databases to distributed networks and edge devices.

The purpose of this review is to critically evaluate how ML can enhance privacy and security in EHR systems. Drawing on eight key studies, we analyze approaches such as anomaly detection, ontology-driven reasoning, privacy-preserving ML, blockchain-ML fusion, IoT-based monitoring, and AI-driven cybersecurity. We highlight their contributions, identify their limitations, and synthesize insights to outline future research directions. By doing so, we aim to provide a comprehensive roadmap for researchers, practitioners, and policymakers

seeking to strengthen trust in digital healthcare infrastructures [3] [4].

## II. TRADITIONAL SECURITY MEASURES IN HER

The security of Electronic Health Records (EHRs) has long relied on traditional safeguards such as access control, cryptography, and anonymization. These mechanisms form the baseline for protecting sensitive data, and without them, no healthcare system could function safely. However, while they are essential, they were largely designed in an era when healthcare data volumes were smaller, threats less sophisticated, and interoperability less emphasized. To critically evaluate how machine learning contributes to EHR security, it is first important to understand the capabilities and shortcomings of these conventional methods.

### A. Access Control Mechanisms

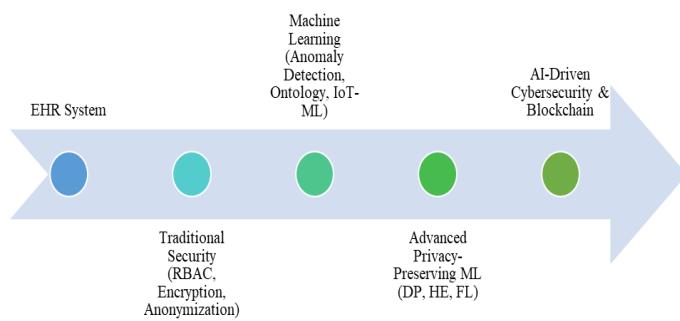
One of the most widely used mechanisms for securing EHRs is access control, which governs who can view or modify patient data. Role-Based Access Control (RBAC) has been the cornerstone of healthcare security for decades. Under RBAC, permissions are assigned to roles—such as physician, nurse, or administrative staff—and users inherit permissions based on their role. This model is intuitive and straightforward to implement, especially in hierarchical healthcare environments where job roles are well defined. For example, a physician may have access to diagnostic records and prescriptions, whereas a billing clerk may only access financial information

Despite its practicality, RBAC shows limitations in dynamic clinical environments. Healthcare often involves exceptions: a nurse covering another ward temporarily, an emergency physician needing rapid access to records, or a specialist consulting on a case. In such scenarios, rigid role definitions can either deny necessary access (potentially compromising patient care) or grant excessive privileges (increasing the risk of misuse).

To address these issues, more fine-grained models have been developed. Attribute-Based Access Control (ABAC) considers attributes such as time of access, location, and purpose, in addition to user role. For instance, an emergency physician may be granted temporary access to a patient's full record during a critical event, with the access expiring once the emergency has passed [5]. Similarly, Semantic-Based Access Control (SBAC) incorporates ontologies and domain knowledge to make more context-aware decisions. While these models enhance flexibility, they introduce complexity: defining, managing, and enforcing policies across diverse

healthcare institutions can be administratively burdensome and prone to errors [6].

Moreover, access control is primarily preventative. It assumes that authorized users will behave responsibly, but in practice, insider threats remain one of the hardest security challenges in healthcare. Studies have shown that employees sometimes access records out of curiosity (e.g., viewing celebrity health files) or for malicious intent. Access control systems, even when fine-grained, are often ill-equipped to detect such misuse, highlighting the need for adaptive, behavior-based monitoring methods like those offered by machine learning.



**Fig. 1 HER Security Layer**

### B. Cryptographic Protection

Cryptography is another pillar of EHR security. It ensures confidentiality, protecting data from unauthorized access during storage (at rest) and transmission (in motion). Widely adopted standards include the Advanced Encryption Standard (AES) for symmetric encryption and the Rivest–Shamir–Adleman (RSA) algorithm for asymmetric encryption. In healthcare, AES is often used to encrypt large datasets efficiently, while RSA facilitates secure key exchanges between systems.

Beyond classical algorithms, Transport Layer Security (TLS) protocols are used to protect EHRs during transmission across hospital networks and cloud services. These tools remain indispensable, particularly as healthcare increasingly shifts toward cloud-based infrastructures [7].

However, cryptographic approaches face several challenges:

1. Performance Overhead: Encrypting and decrypting large volumes of health data introduces latency. In time-critical contexts, such as emergency medicine or intensive care, even small delays can hinder patient care.
2. Key Management: Encryption is only as secure as the keys used. Managing, distributing, and revoking keys across a large healthcare ecosystem is complex, especially when multiple organizations are involved. A compromised key can nullify all protections.
3. Scalability in IoT Environments: With the rise of Internet of Things (IoT) devices—wearables, sensors, and remote monitoring tools—traditional cryptographic algorithms may be too resource-intensive. Lightweight cryptography has been proposed to secure resource-constrained devices, but balancing strength and efficiency remains an active research area.
4. Insider Risks: Cryptography protects data from outsiders but does little against authorized users misusing their access once data is decrypted. Thus, it complements but cannot replace behavior monitoring and anomaly detection.

Despite these limitations, cryptography remains non-negotiable in healthcare security. It is the foundation upon which other safeguards, including machine learning-based approaches, are built.

### C. Anonymization and Pseudonymization

Another widely used method for protecting EHRs is anonymization, which modifies datasets to prevent the identification of individual patients. This is especially important for secondary uses of data, such as clinical research, population health studies, and machine learning training. Popular anonymization techniques include:

**k-Anonymity:** Ensures that each individual is indistinguishable from at least  $k-1$  other within a dataset. For example, if  $k = 5$ , each patient's record should resemble at least four others in terms of key identifiers like age and zip code.

**l-Diversity:** Enhances k-anonymity by ensuring that sensitive attributes (e.g., disease type) are sufficiently diverse within each group.

**t-Closeness:** Further improves l-diversity by ensuring that the distribution of sensitive attributes in each group closely matches the overall dataset.

These methods provide mathematical guarantees against re-identification, but they often degrade data utility. For instance, suppressing or generalizing age, location, or clinical details may protect privacy but reduce the dataset's usefulness for

training predictive models. This creates a persistent privacy–utility trade-off [8].

Pseudonymization offers an alternative, replacing identifiers with pseudonyms while retaining more data utility. For example, patient names may be replaced with codes, allowing records to remain linkable without revealing identities. However, pseudonymization is vulnerable to linkage attacks, where adversaries combine anonymized datasets with external information to re-identify individuals.

A famous example is the re-identification of Massachusetts governor William Weld's medical records in the 1990s using publicly available voter registration data, despite anonymization. This underscores the limits of anonymization in the age of “big data,” where cross-referencing datasets is increasingly easy.

#### **D. Limitations of Traditional Approaches**

While access control, cryptography, and anonymization form a necessary baseline, they share common shortcomings:

**Static Defences:** They operate on predefined rules and policies, which struggle against adaptive adversaries.

**Insider Blind Spots:** None of these methods fully addresses misuse by authorized users.

**Scalability Issues:** Large healthcare networks and IoT environments strain these mechanisms.

**Regulatory Complexity:** GDPR and HIPAA demand not just data protection but accountability and transparency, which traditional methods often lack.

As healthcare data continues to expand in volume, velocity, and variety, the limitations of these approaches become increasingly evident. They are necessary but not sufficient, highlighting the need for adaptive, intelligent, and privacy-preserving solutions. This is precisely where machine learning offers new opportunities—by augmenting static defences with dynamic, behavior-based insights [9] [10].

### **III. MACHINE LEARNING FOR PRIVACY AND SECURITY IN EHRs**

Machine learning (ML) has emerged as one of the most promising tools to address the shortcomings of traditional security measures in Electronic Health Records (EHRs). Unlike static safeguards such as access control or encryption, ML can dynamically learn patterns of normal and abnormal behavior, adapt to new threats, and even enable computation on sensitive data without compromising privacy. In this section, we review key areas where ML has been applied to EHR privacy and security, drawing insights from eight

representative studies and situating them within the broader literature.

#### **A. Anomaly Detection for Suspicious Access**

One of the earliest applications of ML in healthcare security is anomaly detection. Bajaj and Bartlett (2019) demonstrated how classifiers such as Logistic Regression (LR) and Support Vector Machines (SVMs) could detect suspicious access patterns in EHR audit logs. Their work showed that ML could effectively distinguish between legitimate and illegitimate access with area-under-curve (AUC) values as high as 0.95. These results highlight ML's ability to detect subtle patterns that traditional rule-based systems might miss [11].

The strength of anomaly detection lies in its adaptability. By training models on historical access logs, healthcare institutions can build a baseline of “normal” behavior for different roles, times, and contexts. For instance, a physician accessing patient records during their shift may be normal, while multiple record requests late at night from the same account may signal suspicious activity. ML algorithms can continuously update these baselines, improving their sensitivity to evolving patterns.

Nevertheless, anomaly detection faces challenges. Models often require large, high-quality labelled datasets, which are scarce in healthcare due to privacy concerns. Moreover, healthcare workflows are inherently variable—legitimate but unusual access patterns may trigger false positives, frustrate clinicians and create “alert fatigue.” To be practical, anomaly detection systems must balance sensitivity with usability, a trade-off that remains an active area of research [12].

#### **B. Ontology-Driven Privacy Enforcement**

Privacy in EHRs is not only about preventing unauthorized access but also about ensuring that privacy policies themselves are coherent, enforceable, and aligned with regulations. Al-Haiqi et al. (2023) proposed an ontology-driven ML framework to classify privacy policies and identify inconsistencies. By combining semantic ontologies with deep learning models like BERT and Distil BERT, they developed systems capable of reasoning about the legitimacy of access rules.

This approach is significant because healthcare privacy rules are often complex and context-dependent. Ontologies provide a structured representation of knowledge—defining entities such as “patient,” “care provider,” or “researcher,” and the relationships between them. When combined with ML, ontologies allow systems to move beyond surface-level analysis and consider the deeper meaning of policies. For



example, a request for access to a patient's genetic data might be permissible for a genetic counsellor but not for an insurance agent, even if both belong to the "authorized user" category [13].

The primary limitation of ontology-driven ML is its dependence on the quality and completeness of the ontology itself. Building and maintaining ontologies across diverse healthcare settings is labour-intensive, and errors in the ontology can cascade into misclassifications. Nonetheless, this approach offers a powerful complement to anomaly detection, ensuring that privacy policies themselves are sound before they are enforced.

#### **C. Privacy-Preserving Machine Learning (PPML)**

One of the most transformative developments in ML for healthcare security is the rise of privacy-preserving machine learning (PPML). Graham and Hamilton (2025) reviewed key PPML techniques, including Differential Privacy (DP), Homomorphic Encryption (HE), and Federated Learning (FL).

**Differential Privacy (DP):** DP introduces carefully calibrated statistical noise into datasets or model outputs, making it mathematically difficult to infer whether any individual's data was included. In the context of EHRs, DP allows researchers to release aggregate statistics or train ML models without exposing patient-level information. Its main strength is its strong theoretical guarantees. However, DP reduces accuracy, especially when the privacy budget ( $\epsilon$ ) is set to strict values. The challenge is to strike a balance between privacy and utility [14].

**Homomorphic Encryption (HE):** HE enables computations to be performed directly on encrypted data, producing encrypted outputs that can later be decrypted without exposing the raw inputs. This allows cloud providers to perform ML training or inference without ever seeing patient data in plaintext [15]. While theoretically elegant, HE remains computationally expensive. Fully homomorphic encryption, in particular, is orders of magnitude slower than conventional computation, limiting its current practicality in large-scale healthcare systems.

**Federated Learning (FL):** FL trains models across multiple institutions by keeping data local and only sharing model updates. For instance, hospitals in different regions can collaboratively train a diagnostic model without ever exchanging raw patient data. This not only protects privacy but also addresses regulatory barriers to cross-border data sharing. However, FL faces challenges such as communication overhead, model heterogeneity, and vulnerability to poisoning attacks.

Hybrid approaches that combine DP, HE, and FL are increasingly being explored. For example, DP can be applied to model updates in FL to reduce leakage, while HE can secure aggregation of updates. Together, these methods offer a promising path toward secure and collaborative healthcare AI [16].

#### **D. Blockchain and ML for Secure EHRs**

Blockchain has been heralded as a transformative technology for ensuring transparency and tamper-resistance in EHR systems. Diana et al. (2022) conducted a systematic mapping study highlighting how blockchain is increasingly combined with ML to secure health data exchange. Blockchain provides immutable audit trails, ensuring that every access request is recorded and verifiable. ML algorithms, in turn, can analyze these blockchain logs to detect anomalies or predict security breaches.

For example, a blockchain-enabled EHR system could maintain a distributed ledger of all access requests, while an ML model monitors the ledger for suspicious patterns—such as unusual spikes in access requests from a single node. By combining the two, healthcare systems gain both transparency (through blockchain) and intelligence (through ML).

Yet blockchain is not without drawbacks. Traditional consensus mechanisms such as Proof-of-Work are energy-intensive and unsuitable for healthcare environments. Even more efficient mechanisms like Proof-of-Stake may face scalability challenges when handling large volumes of healthcare transactions. Additionally, integrating blockchain with existing EHR systems requires significant infrastructure changes. Nonetheless, the convergence of blockchain and ML remains a promising frontier for secure, decentralized healthcare.

#### **E. IoT and Multilayer ML Approaches**

The rise of the Internet of Things (IoT) in healthcare introduces both opportunities and risks. Wearable devices, remote monitoring tools, and smart hospital systems generate vast amounts of real-time health data. While this data can enhance patient care, it also expands the attack surface. Qi (2025) proposed an IoT-based multilayer ML framework that integrates artificial neural networks (ANNs) with lightweight encryption. Their system achieved 91% diagnostic accuracy while securing IoT data streams.

The strength of IoT-ML approaches lie in their ability to provide real-time monitoring. For example, an ANN can analyze ECG data from a wearable device to detect anomalies while simultaneously ensuring that data transmissions are

encrypted. This not only improves clinical outcomes but also safeguards patient privacy.

However, IoT environments pose unique challenges. Many devices are resource-constrained, making it difficult to implement complex encryption or ML models locally. Furthermore, IoT devices are often the weakest link in healthcare security, vulnerable to malware, physical tampering, or denial-of-service attacks. Integrating ML into IoT security frameworks requires balancing computational efficiency with privacy protection, a balance that remains elusive [17].

#### **F. AI-Driven Cybersecurity**

Nankya et al. (2024) emphasized the role of AI in transforming cybersecurity for healthcare. Unlike traditional defences that react to known threats, AI-driven systems can anticipate and adapt to emerging threats. Applications include real-time intrusion detection, adaptive firewalls, and predictive analytics. For instance, deep learning models can analyze network traffic to detect zero-day attacks that have never been seen before.

A particularly promising area is the integration of AI with compliance monitoring. AI systems could automatically track data flows across healthcare networks, flagging potential violations of HIPAA or GDPR. They could also provide explanations for their decisions, helping organizations demonstrate accountability to regulators.

Nonetheless, AI-driven cybersecurity faces its own limitations. Training deep models requires large amounts of labelled data, which may be difficult to obtain in healthcare. AI systems can also themselves be vulnerable to adversarial attacks, where malicious inputs are crafted to evade detection. Moreover, the cost of deploying and maintaining AI-driven cybersecurity solutions can be prohibitive, particularly for smaller healthcare institutions.

Despite these challenges, AI-driven cybersecurity represents the next frontier of EHR protection. By moving from static defences to predictive, adaptive systems, healthcare can better anticipate and mitigate threats in real time.

#### **IV. COMPARATIVE ANALYSIS OF ML APPROACHES FOR EHR SECURITY**

The reviewed studies collectively highlight the diversity of machine learning applications for securing Electronic Health Records (EHRs). Each approach—whether anomaly detection, ontology-driven reasoning, privacy-preserving ML, blockchain-ML fusion, IoT integration, or AI-driven cybersecurity—addresses specific vulnerabilities but also

introduces new trade-offs. A critical comparative analysis helps illuminate where these methods excel, where they fall short, and how they might be integrated into layered defences [18].

##### **A. Methodological Diversity**

One striking observation is the methodological diversity across the eight reviewed studies. On one end of the spectrum are relatively lightweight classifiers such as logistic regression and SVMs, which excel in anomaly detection within EHR access logs (Bajaj & Bartlett, 2019). These methods are interpretable and computationally efficient, making them feasible for hospital IT departments with limited resources. On the other end are computationally heavy approaches such as homomorphic encryption and blockchain-ML systems, which provide stronger guarantees of privacy and transparency but at the cost of scalability and efficiency (Graham & Hamilton, 2025; Diana et al., 2022).

Ontology-driven frameworks (Al-Haiqi et al., 2023) sit in between these extremes. They combine symbolic reasoning with statistical learning, offering semantic richness while still benefiting from the predictive power of ML. Similarly, IoT-ML systems (Qi, 2025) prioritize real-time responsiveness, making them suitable for scenarios such as remote patient monitoring, though often at the expense of device-level security.

##### **B. Strengths and Limitations**

TABLE I  
 SUMMARY OF REVIEWED MACHINE LEARNING APPROACHES FOR EHR  
 PRIVACY AND SECURITY

Study	Methodology	Application Focus	Strengths	Limitations
Bajaj & Bartlett (2019)	Logistic Regression, SVM	Suspicious access detection	High accuracy, interpretable	Requires large labelled datasets; risk of false positives
Al-Haiqi et al. (2023)	Ontology + BERT	Privacy policy validation	Semantic reasoning + ML; improves compliance	Dependent on ontology quality; limited scalability
Graham & Hamilton (2025)	DP, HE, FL	Privacy-preserving ML	Strong privacy guarantees; decentralized learning	High computational overhead; privacy-utility trade-off

off					AI-driven Cybersecurity	✓	✓	✓	Partial
Diana et al. (2022)	Blockchain + ML	Secure data exchange	Immutable audit trails; decentralized trust	Scalability issues; high energy cost					
Qi (2025)	IoT + ANN	Real-time monitoring	High diagnostic accuracy (91%); lightweight encryption	Vulnerable IoT devices; resource constraints					
Nankya et al. (2024)	AI-driven cybersecurity	Threat detection & adaptive defence	Real-time anomaly detection; predictive analytics	High cost; adversarial ML vulnerabilities					

This table illustrates that no single approach provides a complete solution. Instead, they complement one another, and their integration into hybrid frameworks appears most promising [19].

### C. Mapping Approaches to Security Goals

Another way to compare these methods is to examine how they align with the core security goals of confidentiality, integrity, availability, and regulatory compliance.

TABLE III

MAPPING OF ML APPROACHES TO PRIVACY AND SECURITY GOALS IN EHRs

ML Approach	Confidentiality	Integrity	Availability	Compliance
Anomaly Detection	✓	✓	–	Partial
Ontology-driven ML	✓	–	–	✓
Differential Privacy	✓	–	–	✓
Homomorphic Encryption	✓	✓	–	✓
Federated Learning	✓	✓	✓	✓
Blockchain + ML	✓	✓	✓	✓
IoT-ANN Monitoring	✓	✓	✓	–

Federated learning and blockchain-ML systems stand out for their broad coverage, supporting confidentiality, integrity, availability, and compliance simultaneously. In contrast, anomaly detection primarily enhances confidentiality and integrity, while ontology-driven ML primarily addresses compliance. This mapping highlights why many researchers advocate multi-layered strategies that combine complementary approaches.

### D. Comparative Insights

Three key insights emerge from comparing these approaches:

1. Trade-offs between efficiency and robustness. Lightweight methods such as anomaly detection and IoT-ML are efficient but offer narrower protections. Heavier methods such as HE or blockchain are robust but often impractical for real-time clinical use. Hybrid models can combine the speed of lightweight methods with the robustness of heavier safeguards.

2. Compliance as a differentiator. Not all approaches directly address regulatory compliance. Ontology-driven ML and federated learning stand out in aligning technical protections with legal frameworks such as HIPAA and GDPR. This makes them particularly valuable for institutions facing stringent regulatory oversight.

3. Complementarity rather than competition. These approaches should not be viewed as mutually exclusive. An ideal system might deploy anomaly detection for day-to-day monitoring, federated learning for collaborative model training, and blockchain for audit trails—all working together to provide layered protection.

### E. Toward Hybrid Architectures

Perhaps the most important comparative insight is the potential for hybrid architectures. Imagine a system where IoT devices feed encrypted patient data into federated learning models enhanced with differential privacy, while blockchain maintains an immutable record of access requests and anomaly detection systems flag suspicious activity. Such integration would not only address confidentiality, integrity, and availability but also ensure accountability and compliance.

Realizing such architectures will require careful design to balance performance with security. For example, blockchain consensus protocols must be lightweight enough for

healthcare environments, and federated learning models must account for data heterogeneity across institutions. Nonetheless, the comparative analysis strongly suggests that the future of EHR security lies not in any single method but in the orchestration of multiple, complementary techniques [20] [21].

## V. CHALLENGES AND RESEARCH GAPS

While machine learning (ML) offers powerful tools to enhance the privacy and security of Electronic Health Records (EHRs), significant challenges remain before these approaches can be widely adopted. These challenges fall broadly into three categories: technical limitations, regulatory and ethical constraints, and socio-organizational barriers. Each category reveals not only gaps in the existing literature but also areas that require urgent attention from researchers and practitioners.

### A. Technical Challenges

- Data Heterogeneity and Quality: EHR data is notoriously heterogeneous. Different hospitals and clinics use varying standards, formats, and coding systems. For example, one institution may use ICD-10 codes for diagnoses, while another relies on SNOMED CT. This lack of standardization complicates federated learning and cross-institutional model training, as models must reconcile incompatible inputs. Furthermore, EHR data often suffers from missing values, duplicates, or errors, which can undermine ML performance.
- Scalability and Performance: Techniques such as homomorphic encryption (HE) and blockchain, while offering strong security guarantees, face scalability challenges. Fully homomorphic encryption remains computationally prohibitive for large-scale ML tasks, and blockchain consensus protocols can struggle with high transaction volumes. In healthcare, where time is often critical, delays introduced by these methods may be unacceptable. Developing lightweight, scalable implementations remains an open problem.
- Adversarial Machine Learning: Ironically, ML models themselves are vulnerable to attack. Adversarial examples—inputs crafted to fool models into misclassification—pose a growing threat. For instance, an attacker might manipulate access logs in subtle ways to evade anomaly detection systems. Similarly, poisoning attacks, where adversaries inject malicious data into training sets, could compromise federated

learning. Few existing healthcare studies rigorously evaluate ML defences against such threats.

- False Positives and Clinical Usability: A recurring issue with anomaly detection is the balance between sensitivity and specificity. High sensitivity may detect more suspicious events but also generate false positives, overwhelming clinicians and IT staff. Alert fatigue is already a well-documented problem in healthcare, particularly with electronic prescribing systems. For ML-based security to be adopted, it must deliver actionable insights without disrupting clinical workflows.

### B. Regulatory and Ethical Challenges

- Compliance with HIPAA, GDPR, and Beyond: Healthcare is one of the most heavily regulated domains. In the U.S., HIPAA mandates strict safeguards for patient privacy, while in Europe, GDPR imposes even stricter requirements, including the right to be forgotten. Many ML methods, particularly deep learning models, are “black boxes” that make it difficult to provide the transparency and accountability these laws demand. For example, GDPR requires that automated decisions affecting individuals be explainable, a requirement that many current ML systems cannot meet.
- Explainability and Trust: Clinicians and patients alike demand transparency. A model that flags a suspicious access request must be able to explain why. Otherwise, clinicians may dismiss warnings or fail to trust the system. This issue extends beyond compliance: trust is foundational to the clinician–patient relationship, and opaque algorithms risk undermining it. Explainable AI (XAI) methods are beginning to address this gap, but their integration into healthcare security remains limited.
- Cross-Border Data Sharing: Healthcare is increasingly global, with multi-center trials and international collaborations becoming the norm. Yet data protection laws vary widely across countries. While federated learning offers a potential solution by keeping data local, questions remain about whether sharing model updates constitutes “data transfer” under GDPR. The lack of international harmonization creates legal uncertainty that hinders cross-border ML adoption.

### C. Socio-Organizational Challenges

- Cost of Deployment: Advanced ML systems, especially those integrating blockchain or homomorphic encryption, require significant infrastructure investment.



Large academic hospitals may be able to afford these systems, but smaller clinics and resource-limited settings cannot. This creates a digital divide where advanced protections are available only to wealthier institutions, exacerbating health inequities.

- **Workforce Training:** Implementing ML-based security requires skilled personnel who understand both healthcare workflows and advanced technologies. Currently, there is a shortage of professionals with this dual expertise. Without adequate training, systems may be misconfigured or underutilized, reducing their effectiveness.
- **Cultural and Organizational Resistance:** Healthcare organizations often resist adopting new technologies due to fears of workflow disruption, liability, or regulatory scrutiny. Security systems perceived as intrusive or burdensome may face pushback from staff. To succeed, ML-based security solutions must be integrated seamlessly into existing workflows and accompanied by change management strategies.

#### **D. Research Gaps**

From these challenges, several research gaps emerge:

1. **Benchmarking and Datasets:** Few standardized datasets exist for evaluating EHR security models. Publicly available datasets are often limited in scope, hindering reproducibility and comparability.
2. **Real-World Deployment Studies:** Most studies remain at the proof-of-concept stage. Large-scale, longitudinal studies evaluating ML-based security in live clinical environments are rare.
3. **Adversarial Robustness:** There is limited research on defending healthcare ML systems against adversarial and poisoning attacks.
4. **Explainable Security Models:** While XAI is gaining traction in diagnostics, its application to security remains underdeveloped.
5. **Ethical Frameworks:** Beyond compliance, ethical frameworks that address fairness, accountability, and patient autonomy in ML-driven security are still evolving.

#### **E. Synthesis**

The challenges and gaps highlight that machine learning is not a panacea. Instead, it must be integrated thoughtfully into broader socio-technical systems. Addressing data heterogeneity will require greater standardization efforts, such as adoption of HL7 FHIR. Scalability issues demand lightweight cryptography and efficient blockchain consensus

mechanisms. Regulatory gaps call for explainable models and international harmonization of data laws. Finally, socio-organizational barriers remind us that technology must serve, not disrupt, the human realities of healthcare delivery.

In short, ML holds immense potential but must overcome significant hurdles before it can deliver on the promise of secure, privacy-preserving EHR systems at scale [22] [23].

#### **VI. FUTURE RESEARCH DIRECTIONS**

The comparative analysis of existing machine learning (ML) approaches for securing Electronic Health Records (EHRs) highlights not only their promise but also their limitations. To move from proof-of-concept studies to widespread, sustainable adoption, future research must address issues of scalability, robustness, interoperability, and trust. This section outlines several promising directions that can guide the next generation of privacy-preserving and secure EHR systems.

##### **A. Hybrid Privacy-Preserving Frameworks**

A recurring theme in the literature is the trade-off between privacy and utility. Differential privacy (DP), homomorphic encryption (HE), and federated learning (FL) each provide unique strengths but also distinct limitations. DP offers strong mathematical guarantees but at the cost of reduced accuracy. HE enables secure computation but is computationally expensive. FL allows decentralized training but suffers from communication overhead and vulnerability to poisoning attacks.

Future research should focus on hybrid frameworks that combine these methods. For instance, FL can be paired with DP to reduce information leakage from model updates, while secure aggregation using lightweight, HE can prevent exposure of intermediate results. Such multi-layered designs would mitigate the weaknesses of individual techniques while maintaining performance. Research into optimizing these hybrid methods for clinical contexts—such as emergency response systems where latency is critical—remains a fertile area.

##### **B. Blockchain–ML Fusion for Decentralized Trust**

Blockchain technology offers transparency and immutability, while ML provides intelligence and adaptability. Their fusion could yield decentralized, tamper-proof systems for EHR management. For example, blockchain could record every access request in a distributed ledger, and ML could continuously analyze the ledger to detect suspicious behaviours or predict potential breaches.

However, for blockchain–ML integration to be practical in healthcare, scalability challenges must be addressed. Current



consensus mechanisms such as Proof-of-Work are unsuitable for time-sensitive environments. Future research should investigate lightweight consensus protocols (e.g., Proof-of-Authority or Byzantine fault tolerance) tailored to healthcare use cases. Additionally, interoperability between blockchain-based EHRs and existing hospital information systems must be ensured, possibly through standardized APIs and middleware.

#### **C. Quantum-Resistant Security Models**

The rise of quantum computing poses a long-term threat to conventional cryptography, including widely used algorithms like RSA and ECC. Although quantum computers capable of breaking these algorithms may still be years away, the healthcare sector must prepare proactively, given the long-term sensitivity of medical data.

Future research should explore the integration of post-quantum cryptography (PQC) with ML-driven healthcare systems. Algorithms such as lattice-based cryptography or hash-based signatures could be combined with federated learning or differential privacy to create quantum-resistant privacy-preserving ML frameworks. Investigating how these approaches can be made efficient enough for real-time healthcare applications will be a crucial step forward.

#### **D. Explainable AI (XAI) for Security and Privacy**

Transparency is essential not only for regulatory compliance but also for clinician and patient trust. Current ML models used for anomaly detection or access prediction are often black boxes, offering little insight into their decision-making processes. This opacity undermines user confidence and may hinder adoption.

Future research should prioritize the development of explainable security models. For example, an anomaly detection system could provide natural language explanations such as, “This access was flagged as suspicious because it occurred outside working hours and involved 30 patient records in rapid succession.” Such interpretability would make systems more acceptable to clinicians and auditors alike. Integrating visualization tools to present explanations in user-friendly formats could further enhance usability.

#### **E. Multimodal and Cross-Institutional Data Integration**

Healthcare data is increasingly multimodal, encompassing structured EHR entries, imaging, genomic data, wearable sensor outputs, and even social determinants of health. While this diversity enriches clinical insights, it also complicates privacy protection. A secure system must be able to handle

not only textual records but also high-dimensional genomic datasets and continuous IoT streams.

Future work should explore privacy-preserving multimodal learning frameworks. For example, federated models could be trained across institutions combining EHRs, genomic datasets, and wearable data without centralizing sensitive information. Secure feature selection methods could help identify which aspects of multimodal data are most relevant for predictive tasks while minimizing privacy risks. Addressing interoperability issues across data types and institutions will be crucial.

#### **F. Automated Regulatory Compliance Systems**

Regulatory frameworks such as HIPAA and GDPR impose strict obligations on healthcare institutions. Compliance is currently monitored through manual audits, which are time-consuming and prone to oversight. ML offers the possibility of automated compliance monitoring.

Future research should focus on designing ML systems that can track data flows in real time, flag potential violations, and even generate audit reports automatically. Such systems could, for instance, detect when data leaves an approved jurisdiction or when access is requested without appropriate consent. Combining these tools with explainability features would help institutions not only maintain compliance but also demonstrate accountability to regulators.

#### **G. Addressing Adversarial Threats**

As ML models are increasingly deployed in healthcare security, they become attractive targets themselves. Adversarial attacks—whether through manipulated inputs, model inversion, or poisoning—can undermine trust in ML-driven defences.

Future work must prioritize the adversarial robustness of healthcare ML systems. Research into certified defences, adversarial training, and robust federated learning protocols is essential. Simulated “red team” exercises in healthcare contexts could help identify vulnerabilities before attackers exploit them in the real world.

#### **H. Bridging the Research–Practice Gap**

Finally, there is a pressing need to bridge the gap between academic research and real-world deployment. Many promising approaches remain confined to laboratory settings, evaluated on synthetic or limited datasets. Large-scale, longitudinal deployment studies in diverse healthcare environments are rare but necessary to understand how ML-based systems perform under real-world constraints.



Future collaborations between academia, industry, and healthcare providers will be critical. Pilot projects in hospitals, supported by regulatory sandboxes that allow controlled experimentation, could accelerate the translation of research into practice. Developing standardized benchmarks and datasets for healthcare security would further facilitate reproducibility and cross-study comparison.

### I. Synthesis

Taken together, these future directions point toward a vision of layered, adaptive, and transparent EHR security. The most promising systems will likely combine multiple ML techniques—anomaly detection, ontology-driven reasoning, PPML, blockchain integration, and XAI—into hybrid architectures tailored to the unique needs of healthcare. These systems will not only detect and mitigate threats but also inspire trust among clinicians, patients, and regulators.

The road ahead is challenging, requiring advances in algorithms, infrastructure, regulation, and culture. Yet the potential rewards are immense: a healthcare ecosystem where sensitive data is protected, patients are empowered, and innovation can flourish without compromising privacy.

### VII. CONCLUSION

The rapid digitization of healthcare has transformed how patient information is recorded, shared, and utilized. Electronic Health Records (EHRs) now serve as the central nervous system of modern healthcare, linking hospitals, clinics, laboratories, and patients in ways that were unthinkable just two decades ago. This transformation has brought undeniable benefits—improved efficiency, data-driven insights, and better coordination of care. Yet it has also created new vulnerabilities. Sensitive patient information, once stored in physical files, is now concentrated in digital repositories that have become lucrative targets for cybercriminals.

Traditional safeguards such as role-based access control, cryptographic encryption, and anonymization continue to play an essential role in protecting these records. They provide the foundation upon which more advanced strategies must be built. However, these methods were designed for static environments and are increasingly ill-suited for the dynamic, interconnected, and high-volume ecosystems of contemporary healthcare. They often fail to detect insider misuse, adapt to emerging threats, or provide the transparency required by modern regulatory frameworks.

This is where machine learning (ML) has emerged as a transformative force. Unlike static defences, ML systems can

adapt to evolving patterns, analyze vast datasets in real time, and provide predictive insights. The eight representative studies reviewed in this paper illustrate the breadth of ML's potential. Anomaly detection techniques can flag suspicious access patterns; ontology-driven frameworks can ensure that privacy policies are logically sound; privacy-preserving ML approaches such as differential privacy, homomorphic encryption, and federated learning can enable collaborative research without compromising individual privacy. Meanwhile, blockchain-ML fusion offers transparency and immutability, IoT-ML systems secure real-time monitoring, and AI-driven cybersecurity promises predictive and adaptive defences.

Despite these advances, challenges remain. Technical barriers such as data heterogeneity, scalability issues, and adversarial vulnerabilities must be addressed. Regulatory and ethical concerns—ranging from compliance with HIPAA and GDPR to the need for explainable decision-making—highlight that trust is as critical as technical sophistication. Organizational hurdles, including costs, workforce training, and cultural resistance, further complicate adoption. These challenges underscore that technology alone cannot secure healthcare; socio-technical integration is equally essential.

Looking ahead, the most promising direction lies in hybrid architectures that integrate multiple approaches. For instance, anomaly detection could provide frontline defence, federated learning could enable privacy-preserving model training across institutions, blockchain could guarantee auditability, and explainable AI could ensure transparency for clinicians and regulators. Such systems would not only enhance confidentiality, integrity, and availability but also align with compliance frameworks and ethical expectations.

Future research must also anticipate the challenges of tomorrow. Quantum computing threatens to render current cryptographic methods obsolete, necessitating the integration of post-quantum algorithms into healthcare security. The growing diversity of healthcare data—from genomics to wearable sensors—demands multimodal privacy-preserving ML frameworks. Automated compliance systems could reduce the administrative burden on healthcare providers, while adversarial robustness research will be critical to safeguarding ML models themselves. Above all, collaboration between researchers, clinicians, regulators, and patients will be key to designing systems that are not only technically sound but also socially acceptable.

In conclusion, the journey toward secure and privacy-preserving EHRs is ongoing. Traditional safeguards provide a necessary foundation, but machine learning has opened new



horizons, enabling dynamic, adaptive, and intelligent protection mechanisms. The challenge now is to integrate these technologies thoughtfully, addressing their limitations while leveraging their strengths. If done correctly, the result will be a healthcare ecosystem where sensitive data is both protected and utilized to its fullest potential, enabling innovation, improving outcomes, and preserving the trust that lies at the heart of medicine.

## REFERENCES

- [1] Nankya, M., Mugisa, A., Usman, Y., Upadhyay, A., & Chataut, R. (2024). Security and privacy in E-health systems: a review of AI and machine learning techniques. *IEEE Access*.
- [2] Qi, K. (2025). Advancing hospital healthcare: achieving IoT-based secure health monitoring through multilayer machine learning. *Journal of Big Data*, 12(1), 1.
- [3] Tertulino, R., Antunes, N., & Morais, H. (2024). Privacy in electronic health records: a systematic mapping study. *Journal of Public Health*, 32(3), 435-454.
- [4] Graham, O., & Hamilton, D. (2025). Privacy-Preserving Machine Learning for Electronic Health Records.
- [5] Nowrozy, R., Ahmed, K., Wang, H., & McIntosh, T. (2023, July). Towards a universal privacy model for electronic health record systems: an ontology and machine learning approach. In *Informatics* (Vol. 10, No. 3, p. 60). MDPI.
- [6] Seh, A. H., Al-Amri, J. F., Subahi, A. F., Agrawal, A., Pathak, N., Kumar, R., & Khan, R. A. (2022). An analysis of integrating machine learning in healthcare for ensuring confidentiality of the electronic records. *Computer Modeling in Engineering & Sciences*, 130(3), 1387-1422.
- [7] Boxwala, A. A., Kim, J., Grillo, J. M., & Ohno-Machado, L. (2011). Using statistical and machine learning to help institutions detect suspicious access to electronic health records. *Journal of the American Medical Informatics Association*, 18(4), 498-505.
- [8] Kaur, P., Sharma, M., & Mittal, M. (2018). Big data and machine learning based secure healthcare framework. *Procedia computer science*, 132, 1049-1059.
- [9] C. O. Alenoghenwa, A. J. Onumanyi, H. O. Ohize, A. O. Adejo, M. Olibgi, S. I. Ali, and S. A. Okoh, "EHealth: A survey of architectures, developments in health security concerns and solutions," *Int.J.Environ. Res. Public Health*, vol. 19, no. 20, p. 13071, Oct. 2022.
- [10] D. R. Farringer, "Maybe if we turn it off and then turn it back on again? Exploring health care reform as a means to curb cyber attacks," *J. Law, Med. Ethics*, vol. 47, no. S4, pp. 91–102, Dec. 2019, doi: 10.1177/1073110519898046.
- [11] K. Reddy, P. Gharde, H. Tayade, M. Patil, L. S. Reddy, and D. Surya, "Advancements in robotic surgery: A comprehensive overview of current utilizations and upcoming frontiers," *Cureus*, vol. 1, p. 21, Dec. 2023.
- [12] V. Vajroboi, B. B. Gupta, and A. Gaurav, "Mutual information based logistic regression for phishing URL detection," *Cyber Secur. Appl.*, vol. 2, Jun. 2024, Art. no. 100044.
- [13] Horodyski D. (2015) 2013 OECD Guidelines on the protection of privacy and transborder flows of personal data as an example of recent trends in personal data protection, ResearchGate, pp. 255–266, <https://doi.org/10.13140/RG.2.1.1508.4405>
- [14] Jayabalan M., Rana M. E. (2018) Anonymizing healthcare records: A study of privacy preserving data publishing techniques. *Adv. Sci. Lett.* 24:1694–1697. <https://doi.org/10.1166/asl.2018.11139>
- [15] Hossain, M. D., Rahman, M. H., & Hossan, K. M. R. (2025). Artificial Intelligence in healthcare: Transformative applications, ethical challenges, and future directions in medical diagnostics and personalized medicine.
- [16] Fu, B., Zhang, M., He, J., Cao, Y., Guo, Y., & Wang, R. (2022). StoHisNet: A hybrid multi-classification model with CNN and Transformer for gastric pathology images. *Computer Methods and Programs in Biomedicine*, 221, 106924.
- [17] Manjunatha, A., & Mahendra, G. (2024, December). TransNet: A Hybrid Deep Learning Architecture Combining CNNs and Transformers for Enhanced Medical Image Segmentation. In *2024 International Conference on Computing and Intelligent Reality Technologies (ICCIIRT)* (pp. 221-225). IEEE.
- [18] Vamsi, D.; Reddy, P. Electronic health record security in cloud: Medical data protection using homomorphic encryption schemes. In *Research Anthology on Securing Medical Systems and Records*; IGI Global: Hershey, PA, USA, 2022; pp. 853–877.
- [19] Peute, L.W.; Wildenbos, G.A.; Engelsma, T.; Lesselroth, B.J.; Lichtner, V.; Monkman, H.; Neal, D.; Van Velsen, L.; Jaspers, M.W.; Marcilly, R. Overcoming Challenges to Inclusive User-based Testing of Health Information Technology with Vulnerable Older Adults: Recommendations from a Human Factors Engineering Expert Inquiry. *Yearb. Med. Inform.* 2022, 31, 74–81.
- [20] Boddy, A.J., Hurst, W., Mackay, M., ElRhalibi, A. (2019). Density-based outlier detection for safe guarding electronic patient record systems. *IEEE Access*, 7, 40285–40294. DOI 10.1109/ACCESS.2019.2906503.
- [21] Mounia B., and Habiba C.(2015) "Big Data Privacy in Healthcare Moroccan context." Paper presented at the 2nd International Workshop on Privacy and Security in HealthCare, *Procedia Computer Science* 63: 575 – 580.
- [22] Shinde K. V. (2016) "A Real Time Monitoring System In Healthcare With Hadoop." 'Research Journey' International Multidisciplinary E Research Journal, Special Issue-I : 15-19.
- [23] Pote M. M., Dhote C. H., Sharma D. H. (2015) "Homomorphic Encryption for Security of Cloud Data." *Procedia Computer Science* 79: 175-81.