

AI-Driven Fraud Detection and Cybersecurity Framework for Protecting U.S. Healthcare Systems Against Emerging Cyber Threats

Partha Chakraborty

School of Business, International American University, Los Angeles, CA 90010, USA

<https://orcid.org/0009-0006-3203-8902>

Email: parthachk64@gmail.com

Abstract

The rapid digitalization of the U.S. healthcare sector has created a large and increasingly valuable attack surface that cybercriminals continue to exploit. This paper examines two major threats facing healthcare organizations: financial fraud involving fraudulent claims and billing, and cybersecurity attacks such as ransomware and data breaches. It proposes an integrated, AI-driven framework for addressing both threats. Drawing on peer-reviewed literature, the paper first examines the changing healthcare threat landscape, including the sharp increase in hacking-related breaches and the operational and patient-safety consequences of ransomware. It then reviews artificial intelligence and machine learning approaches to fraud detection, including supervised and unsupervised classifiers, deep-learning anomaly detection, and natural language and network analytics. These capabilities are then considered within established cybersecurity frameworks, particularly the NIST Cybersecurity Framework and HIPAA. The paper's primary contribution is a layered framework that connects AI-based fraud detection with the Identify, Protect, Detect, Respond, and Recover functions of cybersecurity risk management while incorporating privacy-preserving techniques and human oversight. The paper concludes by discussing ongoing challenges involving model interpretability, data privacy, class imbalance, and benchmarking gaps, as well as future opportunities such as real-time prevention, federated learning, and explainable AI.

Keywords: healthcare fraud detection, cybersecurity framework, artificial intelligence, machine learning, NIST Cybersecurity Framework, HIPAA, ransomware, data breaches

1. Introduction

The United States healthcare system has become one of the most digitally dependent sectors of the economy. Electronic health records (EHRs), Internet of Medical Things devices, networked diagnostic equipment, and cloud-based systems have transformed clinical workflows and improved the coordination of patient care. At the same time, this digital transformation has introduced significant cybersecurity risks, including ransomware, data breaches, and unauthorized access to sensitive health information. As these technologies become more interconnected, they create additional opportunities for malicious actors to gain access to healthcare systems (George et al.). Healthcare organizations may benefit from EHRs, networked medical devices, and cloud-based platforms, but these same technologies can create new entry points for attackers. The consequences extend beyond compromised data because cyberattacks can interrupt healthcare services, threaten patient safety, and undermine public trust in healthcare institutions (Ahmed et al.).

Alongside cybersecurity threats, healthcare organizations face a persistent and equally costly problem: financial fraud. Healthcare insurance fraud is a major concern in the United States, with an estimated \$300 billion lost annually. Although machine learning has been investigated as a fraud-detection tool for more than a decade, significant challenges remain, including class imbalance, differences among datasets, and difficulties with model interpretability (Curtis et al.). Financial fraud not only creates substantial monetary losses but also undermines the integrity and trustworthiness of healthcare systems. The increasing complexity of fraudulent activities has made conventional detection approaches, such as manual audits and rule-based systems, less effective (Oladokun et al.). Fraud, waste, and abuse in medical claims therefore represent a continuing drain on resources that could otherwise support patient care.

Cybersecurity and financial fraud are often treated as separate problems and managed by different teams, technologies, and budgets. However, the two threats share an important foundation: healthcare data and the digital infrastructure used to process it. Both require continuous monitoring of large volumes of diverse information, both depend on identifying abnormal behavior in near real time, and both operate within a regulatory environment that can struggle to keep pace with rapidly changing threats. Regulations such as the Health Insurance Portability and Accountability Act (HIPAA) establish important requirements for protecting patient information, but compliance requirements alone may not fully address the changing nature of cyber threats (Ahmed et al.).

This paper argues that an integrated, AI-driven framework can address both threats by incorporating machine-learning-based fraud detection into a structured cybersecurity risk-management approach. Section 2 examines the evolving cyber threat landscape in U.S. healthcare. Section 3 reviews artificial intelligence and machine learning techniques used for fraud

detection. Section 4 examines cybersecurity frameworks applicable to healthcare. Section 5 proposes an integrated framework that combines these capabilities. Section 6 discusses challenges and ethical considerations, Section 7 examines future directions, and Section 8 concludes the paper.

2. The Evolving Cyber Threat Landscape in U.S. Healthcare

The scale of the cybersecurity threat facing healthcare organizations is evident in annual breach reporting. Analysis of healthcare breach data found that hacking attacks against healthcare providers increased by 320% in 2016, with 81% of records breached that year resulting specifically from hacking attacks (Abouelmehdi et al., 2018). The same reporting identified 325 large breaches involving protected health information and the compromise of 16,612,985 individual patient records. The largest single incident accounted for 3,620,000 compromised records, while 40% of large breach incidents involved unauthorized access or disclosure (Abouelmehdi et al., 2018). These findings demonstrate the need for healthcare providers to adopt more proactive and comprehensive strategies for protecting information assets and addressing the growing cyber threat (Abouelmehdi et al., 2018).

Ransomware has become one of the most visible and disruptive forms of cyberattack. It is a type of malware that encrypts data and holds it hostage until a ransom is paid, and it has been identified as one of the most prominent threats facing hospitals (Abouelmehdi et al., 2018). The operational effects of ransomware are increasingly well documented. A study connecting a database of hospital ransomware attacks with Medicare claims data found that ransomware attacks reduced hospital volume by 17–24% during the initial week of an attack, with recovery occurring within approximately three weeks. More concerning, among patients who were already hospitalized when an attack began, in-hospital mortality increased by 34–38% (Neprash et al.,). These findings demonstrate that cyberattacks are no longer simply an information-security or financial issue. They can directly affect patient safety and clinical outcomes.

The structural factors contributing to healthcare's vulnerability are also important. Healthcare organizations often operate with limited resources and outdated systems while facing shortages of cybersecurity expertise. Their security can be further weakened by third-party vendors and supply-chain partners that may provide attackers with additional pathways into healthcare environments (Ahmed et al.,). Repeated ransomware attacks have disrupted medical services, while millions of patient records remain vulnerable to identity theft and financial exploitation following breaches of EHR systems (Ahmed et al.,). The financial incentive for attackers is clear. Healthcare information is valuable on the black market, and the sector's dependence on digital networks combined with the value of its data has made healthcare an increasingly attractive target (Ahmed et al.,).

A comprehensive review of healthcare cybersecurity literature emphasizes that healthcare is a critical societal environment that attracts malicious activity and has experienced a significant number of damaging attacks. At the same time, advances in technology, computing systems, and wireless communications continue to change healthcare by creating both new capabilities and additional complexity (Martínez et al., 2022). The literature also maps healthcare threats against the MITRE ATT&CK knowledge-based framework, identifies security mechanisms designed to protect healthcare systems, and catalogs publicly available security-focused datasets that can support machine learning research in the medical domain (Martínez et al., 2022). This connection is important because it helps healthcare organizations translate general threat intelligence into specific defensive priorities while also identifying the data resources needed to develop AI-based security tools.

3. AI-Driven Fraud Detection in Healthcare

3.1 The Limits of Rule-Based Detection

Traditional healthcare fraud detection has relied heavily on manual reviews and rule-based systems. However, these approaches can generate large numbers of false positives while failing to identify more complex fraud patterns. As healthcare becomes increasingly digital, artificial intelligence provides new opportunities to improve fraud prevention and detection (Kolla,). Conventional approaches based on manual audits and predefined rules are becoming less effective as fraudulent activities become more sophisticated (Oladokun et al.,). The fundamental weakness of rule-based detection is that rules are designed around known patterns. Fraudsters, however, continuously change their methods. Effective fraud-detection systems therefore need to adapt as new patterns and behaviors emerge (Matloob et al.,).

3.2 Supervised and Unsupervised Machine Learning

Machine learning provides an alternative to static rules by allowing systems to identify patterns within large datasets. Research examining supervised and unsupervised techniques, including decision trees, neural networks, and anomaly-detection models, has shown that these methods can identify suspicious activity in insurance claims and reduce false positives (Dey et al.,). By analyzing large EHR and claims datasets, machine learning models can identify unusual patterns that may indicate fraud. Predictive analytics can also support proactive risk assessment, allowing insurers and healthcare organizations to address potential fraud before payments are made rather than discovering it afterward (Dey et al.,). More

broadly, AI can identify complex patterns and anomalies that conventional approaches may overlook, and machine learning models have demonstrated strong performance in predicting fraudulent claims and transactions (Oladokun et al.,).

The performance of these methods is increasingly supported by empirical research. One proposed framework combines machine learning and deep learning to analyze health insurance data and identify fraud at the actor level. Rather than simply identifying a suspicious claim, the framework distinguishes among patients, services, and providers such as physicians, hospitals, and pharmacies. An anomaly transformer is then used to evaluate the behavior of each identified element (Matloob et al.,). In that study, detected fraud was distributed across patient-level fraud at 50%, service-versus-doctor fraud at 12%, service-versus-patient fraud at 13%, and physician-level fraud at 25%. The proposed framework achieved 97% accuracy, with results validated using insurance data from local hospital employees and confirmed by a domain expert (Matloob et al.,). Actor-level attribution is particularly useful because it allows organizations to determine which specific elements are involved and take appropriate action when fraud is confirmed (Matloob et al.,).

3.3 Deep Learning, NLP, and Network Analytics

Advanced AI systems can combine information from previously isolated data sources, transform raw information into meaningful features, and use specialized training approaches to improve detection accuracy while reducing false positives (Kolla,). Machine learning, natural language processing (NLP), and network analytics are among the technologies driving this development (Kolla,). Together, these technologies support a shift from retrospective fraud management toward more proactive detection, allowing suspicious patterns to be identified in near real time before payment occurs (Kolla,). This movement from post-payment recovery to pre-payment prevention may be one of the most significant operational changes AI brings to healthcare fraud management.

A survey of machine learning approaches to healthcare fraud detection reviewed 22 supervised, unsupervised, and semi-supervised classification techniques published between December 2017 and October. The study focused on techniques that were novel within the healthcare fraud-detection field and categorized supervised approaches into deep learning, graph-based, and meta-learning methods (Curtis et al.,). The researchers found that deep learning has become a popular and effective approach for supervised learning. However, important gaps remain, particularly in transfer learning and incremental learning. The study also identified a major benchmarking problem, with few or no studies establishing consistent benchmarks across diverse techniques (Curtis et al.,). Addressing these gaps could improve the effectiveness, transparency, and adaptability of healthcare fraud-detection systems.

4. Cybersecurity Frameworks for Protecting Healthcare Systems

4.1 HIPAA and the Regulatory Baseline

In the United States, HIPAA provides a foundational standard for protecting sensitive patient information and is one of the primary regulatory frameworks governing healthcare cybersecurity (Layode et al.,). However, HIPAA is only one component of a broader healthcare cybersecurity environment. Multiple frameworks and regulations contribute to an organization's overall security posture (Layode et al.,). A major limitation of relying primarily on compliance is that HIPAA requirements, while mandatory, may not fully address the rapidly changing nature of cyber threats. As a result, an organization can be technically compliant while remaining operationally vulnerable (Ahmed et al.,).

4.2 The NIST Cybersecurity Framework

The National Institute of Standards and Technology Cybersecurity Framework (NIST CSF) provides guidance that organizations can use to assess and improve their ability to prevent, detect, and respond to cyberattacks. It is widely used by healthcare organizations as a foundation for developing cybersecurity policies and practices (Layode et al.,).

The CSF is organized around five core functions. Identify focuses on developing an understanding of organizational cybersecurity risks involving systems, assets, data, and capabilities. Protect focuses on implementing safeguards that support the delivery of critical services. Detect involves identifying the occurrence of cybersecurity events. Respond focuses on taking action after a cybersecurity event has been detected. Recover addresses resilience and the restoration of capabilities or services affected by a cybersecurity event (Aljumaiah et al.,).

Each function is further divided into categories that provide more specific outcomes. Within Protect, important categories include identity management and access control, awareness and training, data security, information protection processes and procedures, maintenance, and protective technology (Aljumaiah et al.,). Detect includes anomalies and events, security continuous monitoring, and detection processes. Respond includes response planning, communications, analysis, mitigation, and improvements. Recover includes recovery planning, improvements, and communications (Aljumaiah et al.,). This structure makes the NIST CSF practical for healthcare because individual categories can be connected to specific technical controls and organizational processes.

4.3 Toward Sector-Specific Guidance

Healthcare combines sensitive information, safety-critical operations, and complex legacy infrastructure, making sector-specific adaptations of general cybersecurity frameworks particularly important. Existing research identifies the security mechanisms used to protect healthcare systems while also cataloging publicly available security datasets that can support machine learning research in healthcare (Martínez et al., 2022). The connection between framework guidance and available data is important because an effective framework needs data to support AI-based detection, while data without appropriate governance cannot be responsibly deployed.

The research challenges identified in the literature also demonstrate that healthcare cybersecurity remains an evolving field rather than a solved problem (Martínez et al., 2022). Frameworks must therefore provide enough structure to support consistent security practices while remaining flexible enough to accommodate changing technologies and emerging threats.

5. An Integrated AI-Driven Fraud Detection and Cybersecurity Framework

The central argument of this paper is that healthcare organizations should integrate fraud detection and cybersecurity defense within a single, layered architecture aligned with the NIST CSF functions. Such an approach combines the pattern-recognition capabilities of AI-based fraud analytics with the structured risk-management approach provided by the NIST framework. The proposed framework consists of four primary layers.

Layer 1: Governance, Identity, and Data Security

The foundation of the framework is organizational understanding and protection. Consistent with the Identify and Protect functions of the NIST CSF, this layer establishes asset inventories, identity management and access controls, security awareness and training, data security, and protective technologies (Aljumaiah et al.,). Because HIPAA establishes important requirements for protecting sensitive patient information (Layode et al.,), this layer provides a point at which regulatory requirements can be incorporated into operational practices. However, the objective should extend beyond compliance because emerging threats require organizations to continuously adapt their defenses (Ahmed et al.,).

Vendor and supply-chain risk management should also be considered a required component of this layer. Third-party vendors and supply-chain partners can create pathways through which attackers enter healthcare environments (Ahmed et al.,). Therefore, managing these relationships should be integrated into the organization's overall security strategy rather than treated as an optional activity.

Layer 2: AI-Driven Detection Engine

The second layer contains the machine learning models responsible for continuous monitoring and anomaly detection. It combines supervised and unsupervised approaches, including decision trees, neural networks, and anomaly-detection models trained on large EHR and claims datasets (Dey et al.,). Deep-learning anomaly transformers can also be used to attribute suspicious behavior to specific actors, including patients, providers, and services (Matloob et al.,).

NLP and network analytics can extend detection into unstructured clinical documentation and relationships among providers, patients, services, and claims that may not be apparent in traditional tabular datasets (Kolla,). Because these technologies can combine previously siloed information and transform raw data into meaningful features (Kolla,), the detection engine should operate across both financial fraud and cybersecurity. For example, the same anomaly-detection infrastructure that identifies an unusual claim could also identify an abnormal authentication event or unusual network behavior.

Layer 3: Response and Human Oversight


Detection by itself is not enough. The framework must also support response planning, communication, analysis, mitigation, and continuous improvement (Aljumaiah et al.,). An important design principle is maintaining a balance between AI capabilities and human expertise through structured feedback mechanisms (Kolla,).

AI systems can generate alerts and risk scores, but human investigators should validate significant findings. This approach is consistent with fraud-detection frameworks in which domain experts validate identified cases (Matloob et al.,). Investigator feedback can then be incorporated into future model training and improvement. Actor-level attribution can help organizations determine which parties or processes require attention (Matloob et al.,), while near-real-time detection can allow organizations to intervene before fraudulent payments are completed rather than attempting to recover losses afterward (Kolla,).

Layer 4: Recovery and Resilience

The final layer focuses on recovery planning, improvement, communication, and the restoration of affected capabilities and services (Aljumaiah et al.,). Given the demonstrated effects of ransomware on hospital volume and patient mortality, recovery planning in healthcare should be viewed as a patient-safety responsibility rather than simply an IT continuity function (Neprash et al.,).

The framework can be represented as a continuous cycle:



The main value of this integration comes from the shared detection foundation. Financial fraud and cyber intrusions can both appear as abnormal patterns within healthcare data and digital infrastructure. A unified AI detection engine can therefore support both objectives, while the NIST CSF provides the governance structure needed to make the system auditable, compliant, and resilient.

Challenges, Limitations, and Ethical Considerations

Despite its potential, AI-driven fraud detection and cybersecurity face several persistent challenges. Machine-learning-based fraud-detection systems must address data privacy, model interpretability, and regulatory compliance (Dey et al.,). Interpretability is especially important because AI systems may influence decisions that affect a provider's livelihood or a patient's access to care. As a result, healthcare organizations need to understand not only what a model predicts but also how those predictions are being generated (Oladokun et al.,).

Technical challenges are also significant. Healthcare fraud-detection systems continue to face class imbalance because fraudulent claims represent a small proportion of legitimate claims. Dataset diversity and model interpretability remain additional concerns. The lack of consistent benchmarks across different techniques makes it difficult to compare model performance and determine which approaches are most effective (Curtis et al.,). Research gaps also remain in transfer learning and incremental learning, both of which are important because fraud tactics and cyber threats change over time (Curtis et al.,).

Privacy presents another major concern, particularly when AI models require access to sensitive patient information across multiple organizations. Federated learning provides one potential solution by allowing institutions to train models collaboratively without directly sharing their underlying data. However, federated learning introduces its own privacy and security risks because information exchanged during the training process can potentially reveal characteristics of institutional data (Rieke et al., 2020).

Federated learning therefore does not eliminate privacy concerns. Like other machine learning approaches, it involves some level of risk. Privacy-preserving techniques can provide stronger protection than many commercially available models, but these protections may involve performance trade-offs and could affect model accuracy (Rieke et al., 2020). The framework must therefore distinguish between trusted and less-trusted collaborations. In a consortium where participating organizations are bound by enforceable agreements, some malicious motivations may be reduced. Larger systems, however, may make such agreements more difficult to establish and enforce (Rieke et al., 2020). Any integrated framework must therefore explicitly consider the trust relationships among participants and determine an appropriate balance between privacy protection and model performance.

Future Directions

The future development of AI-based healthcare fraud detection and cybersecurity points toward several important areas. Emerging approaches include real-time prevention, predictive risk analytics, cross-payer collaboration, blockchain integration, and quantum-inspired detection capabilities (Kolla,). These developments could shift healthcare fraud management from primarily being viewed as a cost center toward becoming a strategic capability that supports the long-term sustainability of healthcare organizations.

Cross-payer collaboration is particularly important because fraudulent activities can span multiple insurers, providers, and healthcare organizations. A single organization may only see part of a larger fraud pattern. Sharing information and developing collaborative detection models could therefore provide a more complete view of suspicious behavior.

Explainable AI and privacy-preserving machine learning are also identified as important future directions for AI-driven fraud detection in the U.S. healthcare sector (Dey et al.). Federated learning may play an important role in collaborative model development because it can allow organizations to work together without directly sharing sensitive patient information (Rieke et al., 2020). Addressing the benchmarking gap identified in the literature would also improve the field. Shared datasets and standardized evaluation protocols could make it easier to compare different approaches and increase confidence in reported results (Curtis et al.).

Future systems should also more closely connect cybersecurity telemetry with fraud analytics. Instead of treating indicators of compromise and indicators of fraudulent billing as separate signals, organizations could analyze both within a common detection pipeline. This approach could help identify relationships that would otherwise remain hidden within organizational and technical silos.

Conclusion

The U.S. healthcare sector faces a threat environment in which cyber intrusions and financial fraud are both costly, increasingly sophisticated, and closely connected to the digital transformation that has improved patient care. Hacking attacks against healthcare providers have increased substantially, ransomware remains a major threat to hospitals, and the consequences of these attacks can include measurable increases in patient mortality (Abouelmehdi et al., 2018; Neprash et al.). At the same time, healthcare insurance fraud is estimated to cost the U.S. system approximately \$300 billion annually, while traditional manual and rule-based detection methods struggle to keep pace with increasingly sophisticated fraud schemes (Curtis et al.; Oladokun et al.).

This paper argues that AI-driven fraud detection and cybersecurity should not be treated as separate initiatives. By integrating supervised, unsupervised, and deep-learning techniques with NLP and network analytics and aligning them with the Identify, Protect, Detect, Respond, and Recover functions of the NIST Cybersecurity Framework, healthcare organizations can develop a more unified and adaptive defense strategy. Such an approach can support HIPAA compliance while also providing greater flexibility as threats continue to evolve.

However, realizing this potential requires organizations to address significant challenges involving model interpretability, data privacy, class imbalance, and inconsistent benchmarking. Privacy-preserving approaches such as federated learning may help organizations collaborate while reducing the need to share sensitive patient data, but these approaches must still balance privacy, security, and model performance. Ultimately, combining AI capabilities with human oversight and strong governance provides the most practical path toward protecting both the financial integrity and operational resilience of U.S. healthcare organizations.

References

1. Li, J., Huang, K.-Y., Jin, J., & Shi, J. (2008). A survey on statistical methods for health care fraud detection. *Health Care Management Science*, 11, 275–287.
2. DOI: <https://doi.org/10.1007/s10729-007-9045-4>
3. Relevance: Foundational methods for detecting healthcare insurance fraud and abnormal claims.
4. Joudaki, H., et al. (2015). Using data mining to detect health care fraud and abuse: A review of literature. *Global Journal of Health Science*, 7(1), 194–202.
5. DOI: <https://doi.org/10.5539/gjhs.v7n1p194>
6. Obodoekwe, N., & van der Haar, D. T. (2019). A comparison of machine learning methods applicable to healthcare claims fraud detection. In *Information Technology and Systems—Proceedings of ICITS 2019* (pp. 548–557).
7. DOI: https://doi.org/10.1007/978-3-030-11890-7_53
8. Relevance: Comparing machine learning approaches for healthcare claims fraud.
9. Zhang, C., Xiao, X., & Wu, C. (2020). Medical fraud and abuse detection system based on machine learning. *International Journal of Environmental Research and Public Health*, 17(19), 7265.
10. DOI: <https://doi.org/10.3390/ijerph17197265>
11. Coventry, L., & Branley, D. (2018). Cybersecurity in healthcare: A narrative review of trends, threats and ways forward. *Maturitas*, 113, 48–52.
12. DOI: <https://doi.org/10.1016/j.maturitas.2018.04.008>
13. Kruse, C. S., Smith, B., Vanderlinden, H., & Nealand, A. (2017). Security techniques for the electronic health records. *Journal of Medical Systems*, 41, Article 127.
14. DOI: <https://doi.org/10.1007/s10916-017-0778-4>

15. Jalali, M. S., & Kaiser, J. P. (2018). Cybersecurity in hospitals: A systematic, organizational perspective. *Journal of Medical Internet Research*, 20(5), e10059.
16. DOI: <https://doi.org/10.2196/10059>
17. Argaw, S. T., Bempong, N.-E., Eshaya-Chauvin, B., & Flahault, A. (2019). The state of research on cyberattacks against hospitals and available best practice recommendations: A scoping review. *BMC Medical Informatics and Decision Making*, 19, Article 10.
18. DOI: <https://doi.org/10.1186/s12911-018-0724-5>
19. Szymanski, R., & Levy, C. (2018). Health IT, hacking, and cybersecurity: National trends in data breaches of protected health information. *JAMIA Open*, 1(1), 15–19.
20. DOI: <https://doi.org/10.1093/jamiaopen/ooy019>
21. Choi, S. J., & Johnson, M. E. (2021). The relationship between cybersecurity ratings and the risk of hospital data breaches. *Journal of the American Medical Informatics Association*, 28(10), 2085–2092.
22. DOI: <https://doi.org/10.1093/jamia/ocab142>
23. He, Y., Aliyu, A., Evans, M., & Luo, C. (2021). Health care cybersecurity challenges and solutions under the climate of COVID-19: Scoping review. *Journal of Medical Internet Research*, 23(4), e21747.
24. DOI: <https://doi.org/10.2196/21747>
25. Argaw, S. T., et al. (2020). Cybersecurity of hospitals: Discussing the challenges and working towards mitigating the risks. *BMC Medical Informatics and Decision Making*, 20, Article 146.
26. DOI: <https://doi.org/10.1186/s12911-020-01161-7>
27. Lee, I., & McDonough, C. P. (2022 issue; published online December 22, 2021). An analysis of data breaches in the U.S. healthcare industry: Diversity, trends, and risk profiling. *Information Security Journal: A Global Perspective*, 31(3), 346–358.
28. DOI: <https://doi.org/10.1080/19393555.2021.2017522>
29. Relevance: U.S. healthcare breach trends and risk profiling. Note: Although it appeared in a 2022 journal issue, it was published online in 2021. If you require the final issue publication date to be before 2022, exclude it.
30. Rieke, N., et al. (2020). The future of digital health with federated learning. *npj Digital Medicine*, 3, Article 119.
31. DOI: <https://doi.org/10.1038/s41746-020-00323-1>
32. Bhuyan, S. S., et al. (2020). Transforming healthcare cybersecurity from reactive to proactive: Current status and future recommendations. *Journal of Medical Systems*, 44, Article 98.
33. DOI: <https://doi.org/10.1007/s10916-019-1507-y>