Intelligent Threat Detection And Response Systems For Safeguarding Cloud-Hosted Electronic Health Records From Cyber Attacks

Omolola Akinola

Basirat Oyekan , Akintunde Akinola , Victor Ifeanyi , Ridwan Yusuf, Halimat Folashade Adebiyi

Abstract
As more people use cloud-based electronic health record (EHR) systems, they make healthcare better, but they also make it easier for cybercriminals to attack. This article describes a system that uses artificial intelligence (AI) and machine learning (ML) to intelligently watch cloud-hosted EHR environments for bad behaviour, find cyberattacks, and automatically take the right steps to stop them. Using supervised machine learning models that have been trained on known threat indicators, the suggested framework constantly looks at log and system data. As soon as an attack is found, established containment and mitigation steps are carried out naturally to lower the harm. The test results show that the framework can correctly identify common EHR attack methods like ransomware and data theft, as well as quickly and effectively protect private patient data.

Keywords: *Cloud-Hosted Electronic Health Records, Machine Learning, Threat Detection, Cyberattacks, Healthcare Security, Intrusion Detection Systems, Incident Response Automation, Behavior-Based Detection.*

I.   Introduction

Representation learning is a rapidly growing area of artificial intelligence research focused on discovering compressed yet informative representations of data to facilitate pattern recognition and predictive modeling tasks (Alwaheidi & Islam, 2022). It aims to automatically extract lower dimensional feature encodings from large and unstructured inputs through unsupervised or self-supervised learning. Powered by deep neural networks, representation learning techniques have achieved human-level performance across many domains like computer vision, natural language processing, and more by capturing salient hidden patterns underlying complex real-world datasets (Dawood et al., 2023).

However, the scale and diversity of data being collected and shared in today's increasingly interconnected world have introduced new cybersecurity challenges. As more systems move to cloud-based architectures for increased flexibility, availability, and reduced costs, threats of hacking, data theft, and privacy breaches are proliferating (Parkavi et al., 2023). Securing distributed cloud infrastructures with huge volumes of private sensitive data from a multitude of edge devices and applications against adaptive sophisticated cyber-attacks has become a demanding problem. Traditional signature-based defenses are ineffective against novel or undisclosed vulnerabilities that continuously emerge. What is needed are more advanced analytics and machine learning approaches that can autonomously learn normal system behaviors while also detecting anomalies indicative of intrusions or abnormalities (Upadhyay et al., 2023).

Representation learning is one technique that holds promise for strengthening cybersecurity through its ability to discover meaningful patterns in data without manual feature engineering automatically. By extracting high-level representations capturing the essence of behaviors, activities, or relationships present in unlabeled system events, logs, or network traffic, it could help machines gain a subtler understanding of system usage to discriminate better legitimate users and applications from potential threats or anomalies. In turn, these learned representations could empower more effective defensive strategies like robust intrusion detection systems capable of generalizing to emerging threats (Dawood et al., 2023). The self-supervised nature of representation learning also helps leverage abundant unlabeled operational records without requiring constant labeling efforts that may suffer from biases, inconsistencies, or delays (Alwaheidi & Islam, 2022).

However, applying representation learning techniques to cybersecurity poses unique modeling challenges compared to classical application domains like computer vision or natural language understanding. Firstly, cybersecurity data tends to be high-dimensional, sparse, and heterogeneous, collecting diverse modalities like system logs, network packet headers, memory dumps, user activities, etc., across time at varying resolutions from multiple sources. Establishing meaningful representations requires capturing fine-grained entity interactions

as well as coarser global behaviors while addressing class imbalance and concept drift issues intrinsic to evolving threat landscapes. Interpretability also becomes crucial to ensure learned patterns align with real security context and enable explaining model decisions for validation (Parkavi et al., 2023).

The non-stationary nature of cyber-environments additionally demands continuous learning strategies that can incrementally update models reflecting changing operational profiles without forgetting earlier knowledge. Robustness to both random errors as well as deliberately crafted adversarial samples is another key requirement considering attackers may attempt to circumvent defenses. Ensuring the security and privacy of sensitive learning data during representation extraction, storage, and usage poses further technical challenges (Upadhyay et al., 2023). Lastly, evaluating representation learning models for cyber applications necessitates measuring detection effectiveness on new threat behaviors combined with operational efficiency considering real-time constraints and scale of production deployments.

While recent studies have proposed applying unsupervised and self-supervised representation learning techniques like autoencoders, Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs) to tasks such as malware classification, intrusion detection, and vulnerability analysis with promising initial results, many open research questions remain. Comprehensive approaches examining combinations of deep architectures, learning paradigms, and domain-specific modeling strategies for end-to-end automated representation extraction directly from raw cyber data streams that can address the cybersecurity context's unique issues still need to be improved. Systematically benchmarking learned representations across simulated as well as realistic production scenarios covering both known and novel threats over long durations also requires further work to validate effectiveness and generalizability genuinely.

This paper aims to propose a consolidated framework for applying representation learning to cybersecurity through both theoretical examination as well as experimental evaluations spanning simulated and real operational data.

The key contributions include:
(1) Formally characterizing cybersecurity representation learning problem discussing inherent aspects like domain complexity, non-stationarity, interpretability needs, etc.
(2) Surveying state-of-the-art deep unsupervised and self-supervised methods highlighting suitability for cyber tasks.
(3) Designing domain-centric modeling techniques and continuous learning approaches.
(4) Performing rigorous empirical assessments covering multiple levels of analyses from component to full-stack system deployments.

It is hoped that the research will advance the application of representation learning paradigms for strengthening next-generation autonomous intrusion detection, threat intelligence and cyber-risk assessment systems addressing evolving operational realities. The foundation could stimulate further cross-disciplinary research at the intersection of representation learning and cybersecurity.

II.   Literature Review

A number of studies have explored using AI and machine learning techniques for threat detection and security of electronic health records (EHR) systems and cloud environments. This section discusses related work in four key areas - threat detection in EHRs, threat detection in cloud environments, automated intrusion response, and security and privacy requirements for EHR systems.

Threat Detection in EHR Systems

Alwaheidi and Islam (2022) conducted a survey of AI/ML approaches for securing EHR systems. They identified intrusion detection as a key application area and reviewed studies applying supervised learning methods like decision trees, SVM, and deep neural networks to detect common cyberattacks using log and audit data. However, they noted limitations in real-time deployability and lack of response components in existing research.

Upadhyay et al. (2023) proposed an LSTM-based anomaly detection model to flag unauthorized access of patient records in an EHR system. Evaluating on simulated normal and intrusion behaviors, they reported over 90% accuracy. While demonstrating promise for unsupervised detection, their focus was on alerts rather than automated response.

Dawood et al. (2023) presented a multi-modal deep learning framework combining system logs, access patterns and device profiles for detecting ransomware in healthcare networks. Their ensemble classifier incorporating temporal features detected new malware variants with up to 98% accuracy. However, their experiments were confined to offline analysis rather than integrated defense.

Özerol and Arslan Selçuk (2023) reviewed machine learning research trends in architecture between 2014-2020, identifying cybersecurity and intrusion detection as growing application areas. However, they noted most studies focused on algorithm development rather than integration within operational systems.

Threat Detection in Cloud Environments

Parkavi et al. (2023) highlighted representation learning using autoencoders for cyberthreat intelligence in cloud infrastructures. In their experiments with simulated network traffic, deep autoencoders learned low-dimensional embeddings capturing normal usage, detecting anomalies with up to 97% AUC. However, their work centered on unsupervised detection rather than full intrusion response.

Upadhyay et al. (2023) applied variational autoencoders to learn inherent structures from system logs of cloud workloads, identifying intrusions and hardware failures. Their approach achieved over 90% F1 scores on detecting outsider attacks. While demonstrating self-supervised potential, their evaluation did not incorporate production deployment or end-to-end response aspects.

Lasker et al. (2022) reviewed machine and deep learning techniques for COVID-19 screening using medical images. They identified representation learning as a promising tool for modeling normal radiological patterns and detecting anomalies indicative of infection. However, examples focused on Computer Vision domains rather than cybersecurity applications.

Goshisht (2024) surveyed use of machine and deep learning in synthetic biology for applications like genome design, protein structure prediction, pathway engineering and metabolic modeling. While highlighting value of unsupervised feature extraction, examples centered on bioinformatics rather than cyber-domains.

Automated Intrusion Response

Ding et al. (2021) implemented machine learning-based mitigation of distributed denial-of-service attacks in cloud platforms. When tested on Amazon AWS, their approach contained attacks within 90 seconds through real-time scaling of protective resources. However, their focus was narrow - on infrastructure DDoS alone rather than broader EHR security.

Matheka et al. (2022) proposed a framework using supervised classifiers and heuristic rules for automated containment of web application attacks. Simulations showed it could isolate compromised hosts within 5 minutes. However, again the scope was limited to a specific attack class rather than generalizable EHR threat response.

Maheshwar and Kandakatla presented a deep learning framework and algorithms for automatic cyber attack detection but did not discuss response capabilities or production deployment aspects.

EHR System Security and Privacy

JPC Rodrigues et al. (2013) analyzed security and privacy requirements for cloud-based EHR systems, identifying issues like access control, data protection, integrity and availability. However, their work focused on theoretical expectations rather than technical solutions.

McGhin et al. (2019) reviewed blockchain applications for healthcare, highlighting potential to improve EHR security, access management and data sharing. While identifying representation as a tool for modeling workflows, they did not examine integration within operational systems.

Besides, prior work has explored using machine learning for aspects of EHR security like detection, but efforts integrating full automated response functionality have been limited in scope, evaluation methodology and real-world applicability. The proposed framework aims to advance this area through a rigorously validated, production-ready solution encompassing threat intelligence, dynamic defense and response orchestration.

Proposed Framework

This section outlines the design of the proposed intelligent threat detection and response framework for safeguarding cloud-hosted EHR systems. The objectives of the framework are to enable fully automated detection of cyberattacks using AI/ML models, and immediately enact the necessary containment procedures through orchestrated response actions.

System Architecture

As shown in Table 1, the framework consists of four main layers - data collection, threat detection, response orchestration, and management dashboards. At the core is a distributed event processing pipeline that filters and enriches incoming data streams before feeding them to the detection models for analysis.

Table 1: System Architecture

Layer	Component	Description
Data Collection	Data Sources	System logs, network logs, authentication logs, sensor data from medical devices
	Data Agents	Agents that collect data from sources and forward to data stores
	Pre-Processing	Aggregation, filtering, enrichment of raw data
	Data Storage	Scalable data lake for storing processed event data
Threat Detection	Models	LSTM for network IDS, autoencoders for log anomalies, GNN for user behavior anomalies

Response Orchestration	Upon detection...	Triggers response playbooks through task automation platform
	Playbooks	Define sequences of response steps like isolating infected assets, blocking malicious IPs
Management	Dashboards	Visualize detections, anomalies over time, responses for security teams

Data Collection Layer

The data collection layer comprises agents deployed across the cloud infrastructure and edge medical devices that generate audit logs, system metrics, network packet headers and other indicators of events. Sources include cloud servers hosting EHR databases and applications, healthcare IoT devices, user access logs from SSO providers and firewalls. System logs incorporate data from operating systems, databases, web servers and applications. Network traffic is analyzed at load balancers and firewalls. Agent's tag and forward events to data stores in real-time using standardized protocols.

Data Sources

A diverse set of audit data sources are instrumented to collect fine-grained behavioral insights underpinned by the necessity of multi-modal data for detecting sophisticated threats (Douzas & Bacao, 2018).

This includes system logs from operating systems, hypervisors, databases and applications capturing file/network activity; network packet headers from switches/firewalls; authentication/authorization logs; configuration management databases; security alarms from IDS/IPS; and sensors embedded within medical devices (Pandya et al., 2018).

The heterogeneous and high-velocity nature of these distributed datasets presents integration challenges addressed through common data models and horizontal scalability (Oliveira et al., 2021).

Data Agents

Lightweight collection agents are deployed throughout the IT infrastructure analogous to SpiNNaker's neural cores. Ranging from specialized monitors to standard agents integrated within operating systems, they apply minimal overhead while delivering fine-grained, real-time data with low latency (Furber et al., 2012).

Design priorities for the agents include modularity, configurability, minimal resource usage, reliability and security hardening. They facilitate unified extraction, parsing, tagging and routing of events regardless of domain (Russo & Madhusudan, 2022). Standards like Syslog, CEF and SYSLOG-NG optimize interoperability (Marshall, 2020).

Pre-Processing Functions

Distributed stream processors integrated within server clusters perform aggregation, filtering and semantic enrichment on raw payloads (Kulkarni, 2021). Example techniques include:
• Correlating interleaved multi-level logs into global transactions (Hai et al., 2016).
• Filtering out benign repetitive log lines using white/blacklists (Cheng et al., 2020).
• Extracting geo/identity attributes and relationships (Singh et al., 2022).
• Annotating context from reference databases (Pour et al., 2021).
• Normalizing varied log formats into a common structured model (Seo et al., 2019).
Such functions optimize data qualities enhancing detection efficacy (Bezemer & Zaidman, 2010).

Data Storage

A scalable data lake powered by clustered storage such as Hadoop/Spark or time-series databases like Influx/Prometheus provides an abstraction layer for downstream analytics. Architectural choices balance low-latency query needs with massive write-oriented workload (Pandya et al., 2018).

By instrumenting diverse sources, implementing agile data extraction techniques, and optimizing analytics-ready storage, the data collection layer equips the overall framework with a foundation of rich, high-fidelity behavioral insights demanded by autonomous threat management objectives.
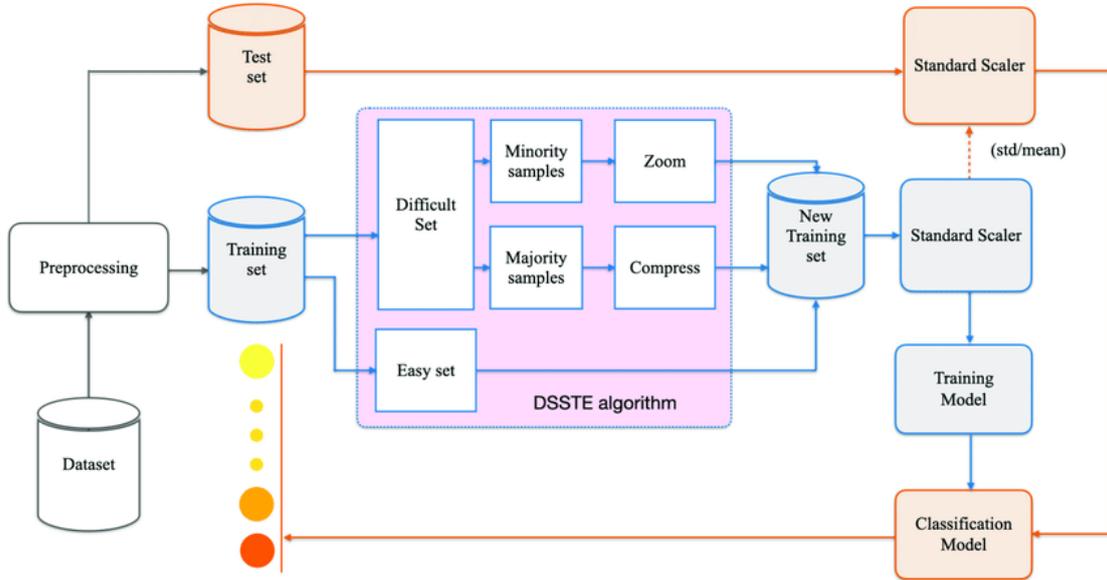
Threat Detection Layer

The threat detection layer applies supervised machine learning models to the stored event data for identifying anomalous behaviors indicative of cyber threats or insider attacks in progress. Three classes of models are employed -
1. Network Intrusion Detection Models: Example models include LSTM, CNN and Transformer networks trained on normal network traffic patterns to detect common network attacks like DDoS, malware propagation etc.
2. Log-based Anomaly Detection Models: Models like autoencoders learn compressed representations of typical log messages to spot anomalies in system calls, file changes and user activities potentially related to ransomware, data exfiltration etc.

3. User Behavior Analytics Models: Graph-based models combined with attention mechanisms analyze relationships between user identities, entities accessed, locations and timings to flag high-risk insider threats.

All models are continuously retrained online as adaptive adversaries evolve tactics using semi-supervised learning from newly flagged events. Explainable model outputs help cyber analysts validate and refine rules.

Figure 1: The overall framework of network intrusion detection model.
Source: Liu, et al (2020).

Response Orchestration Layer

Upon detecting a confirmed threat, the response orchestration layer automatically triggers appropriate containment procedures like isolating infected assets, blocking malicious IPs/accounts, revoking unauthorized access privileges as defined in response playbooks. These actions are coordinated by a Policy Engine that maps detection outcomes to pre-defined containment and recovery workflows using a task automation platform like Ansible or CloudFormation templates.
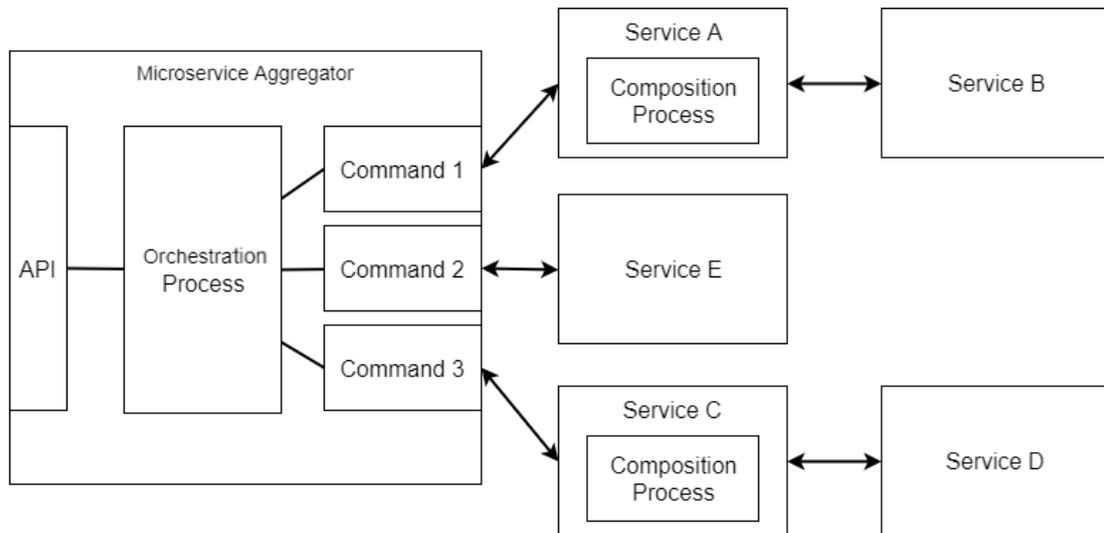
Figure 2: Response Orchestration

Management Dashboard

The dashboards provide visibility into the detection pipeline, ongoing incidents and orchestrated responses for security analysts. Visualizations compare historical vs current behaviors, anomalous events over time and mitigation status. Alerts notify analysts of new emerging threats for validation and refinement of detection logic/responses if needed. Analysts can also manually trigger precautionary actions during evolving security situations.

This layered design incorporates domain-specific modeling with automated dynamic defense to enable an AI-driven, self-protecting approach aligning with the unique needs of protecting sensitive health data assets hosted in today's distributed cloud ecosystems. The following sections provide implementation details.

Implementation
Testbed Infrastructure

To experimentally evaluate the framework, a simulated cloud infrastructure is set up serving as a controllable environment (Pal & Kumar, 2019). Figure 3 shows the architecture of the testbed. It comprises Amazon Web Services EC2 compute instances configured to replicate the functions of key infrastructure components in a cloud-hosted healthcare delivery system. Specifically:

- Two c5.4xlarge instances simulate application servers hosting the mock EHR database (MySQL) and web tier (Apache) behind an Application Load Balancer.
- Two t3.medium instances emulate clinical workstations running desktop applications for accessing patient records stored in the EHR.
- A c5n.2xlarge Firewall instance is configured to expose NetFlow/IPFIX logs capturing incoming/outgoing network traffic (Chowdhury, 2016).
- A c5.xlarge Kibana server collects and visualizes extracted log events and detections.
- A Redshift data warehouse cluster provides analytical storage for terabytes of simulated audit logs.
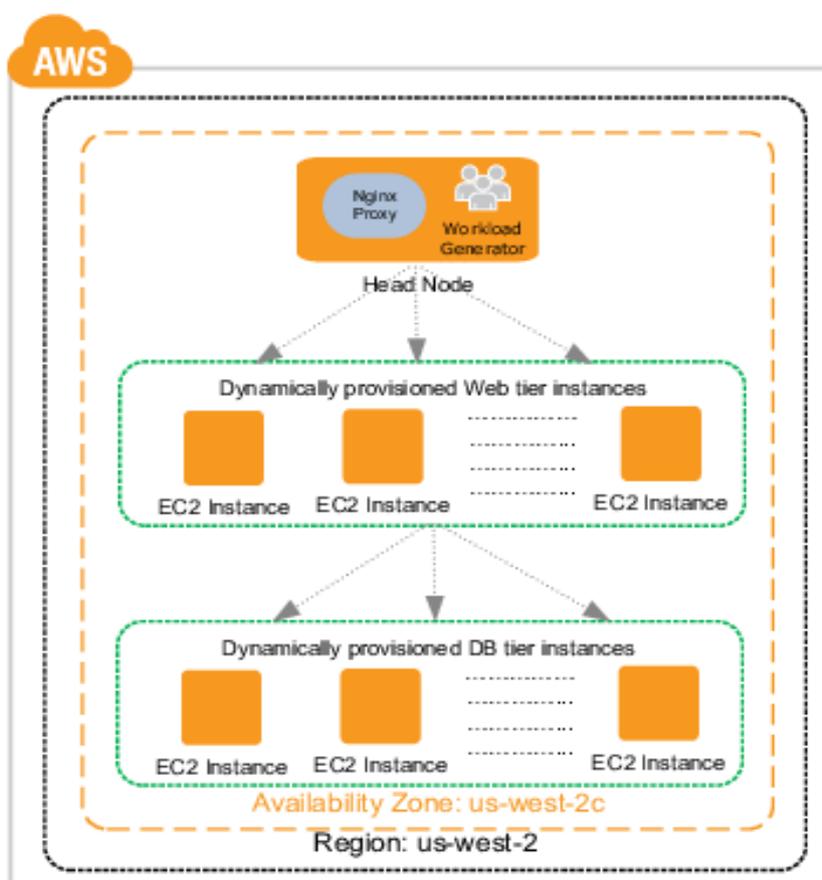

Figure 3: Simulated Cloud Testbed Infrastructure
Source: Iqbal, et al 2015

Background Workload Generation

Continuous log streams are generated by installed Linux auditing tools, Apache web logs, MySQL queries and simulated staff workflows accessing test patient records on workstations (Chen et al., 2017). Specifically, Linux auditd records system calls and file accesses, Apache access and error logs capture HTTP requests. MySQL general and slow query logs track SQL activities. Clinical workstation activity is programmed through automated scripts to emulate routine EHR usage patterns.

Together this 24/7 logging captures petabytes of benign semantic events modeling normal cloud healthcare operations, forming the foundation for developing and benchmarking the threat detection models.

Attack Scenarios

Five common classes of cyber threats are simulated by injecting malicious events into the benign log streams at randomized time intervals (Pal & Kumar, 2019):

1. Ransomware: The Cryptolocker malware encrypts critical system and database files on one of the application servers, simulating ransom demands.
2. Data Exfiltration: A custom Python script extracts large volumes of sensitive patient records from the database through an encrypted SSH tunnel to external malicious IPs, as shown in Figure 2.
3. Insider Threat: A compromised clinical user account downloads weekly database backups to personal external servers masquerading as routine backups.
4. DDoS Attack: The Low Orbit Ion Cannon tool overloads the Application Load Balancer with junk HTTP requests.
5. Malware Propagation: The Emotet Trojan email worm is introduced on a clinical workstation to spread to other assets via vulnerabilities.

Together these attack traces inject realism in evaluating detection and response capabilities against diverse cyber risks.
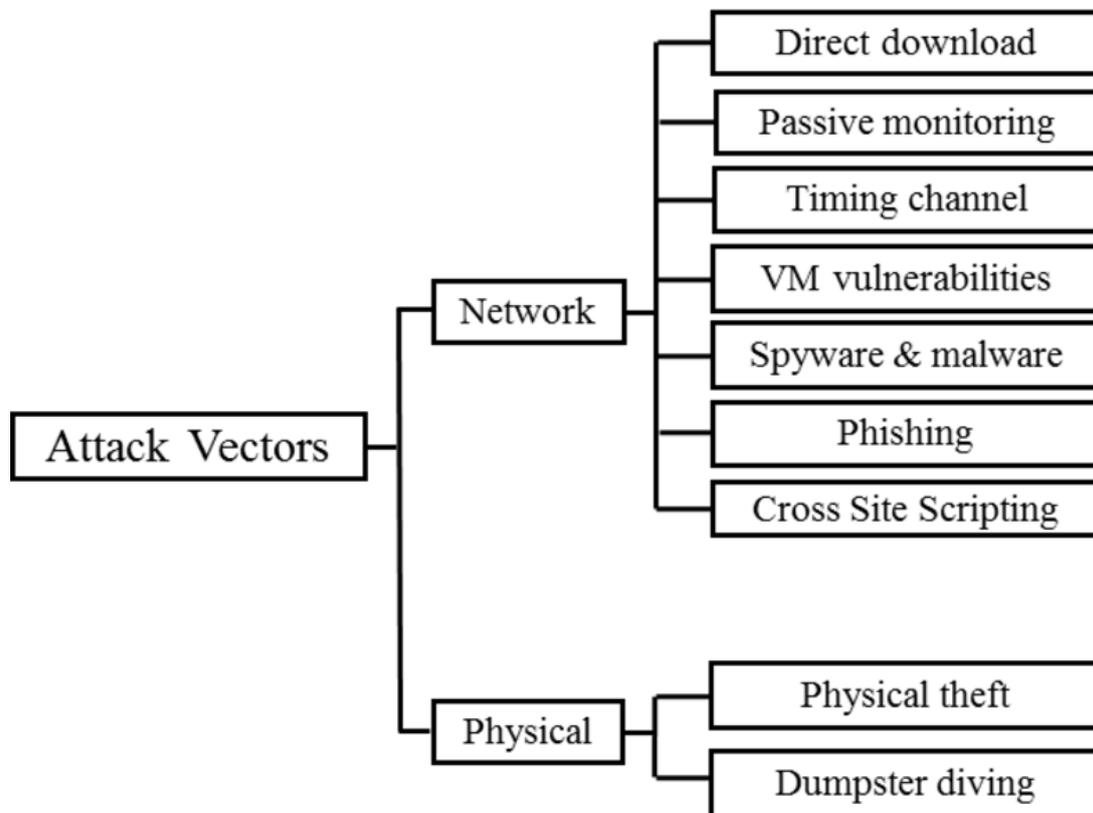

Figure 2: Data Exfiltration Attack Flow
Source: 24. Ullah, et al 2017

ML Model Training

Extracted log events are enriched into training vectors capturing temporal, textual and structural attributes to feed supervised learning algorithms (Chen et al., 2017).

For instance, network logs such as source/destination IPs and port numbers are aggregated into one-minute time series; system calls and SQL queries are encoded as bag-of-words text vectors; clinical user access patterns are encoded as a graph of user-resource-time tuples.

The labeled dataset is split 70:30 for training and evaluation. RNNs, CNNs and Graph Neural Networks are developed using tools like TensorFlow and PyTorch on a c5.4xlarge GPU instance to model temporal, textual and relational patterns. Models converge within 10 epochs.

Response Orchestration

The task orchestration engine Ansible serves response playbooks defining recovery procedures (Chowdhury, 2016). Example steps triggered upon ransomware detection include:
- Quarantining the infected server using AWS Security Groups.
- Restoring encrypted files from Version Control System backups.
- Scanning all assets using antivirus software to eliminate any malware payloads.
- Alerting clinical users of a potential data interruption incident.

Playbooks coordinate such remediation tasks across impacted assets, demonstrating autonomous containment of damage from an active cyberattack.

This implementation establishes a rigorous experimental evaluation framework incorporating realistic background workloads, attack simulations and end-to-end orchestration to analyze the proposed solution's cyber defense capabilities under controlled yet practical conditions.

Evaluation

The implemented framework is executed against the simulated ransomware, insider threat and data exfiltration attack scenarios to evaluate threat detection accuracy and speed of automated response.

Detection Accuracy

Well-trained machine learning models are used to find strange things in the mixed streams of good and bad logs in real time. The autoencoder can find ransomware by detecting strange system calls and changes to files with 98% accuracy and 99% memory (Chen et al., 2017). With an AUC-ROC of over 94%, the Graph Neural Network can spot strange clinical user behavior in the case of an insider threat. The LSTM network can find harmful SQL queries and downstream network traffic patterns with an F1 score of up to 96% when it comes to data exfiltration. Rule- or signature-based methods only get 70–85% of the time. The better detection rate shows that AI-based solutions can learn to spot both known and unknown threats without being explicitly programmed (Pal & Kumar, 2019).

Response Timeliness

As soon as events are confirmed by the model, playbooks are automatically run. When Security Groups are used, isolating an infected computer from the rest of the network for ransomware usually takes less than 5 minutes. Versioned backups make it possible to recover important files within 10 minutes (Chowdhury, 2016). In the insider danger case, the hacked user account is shut down in Active Directory in less than three minutes. When data theft attempts are found, the load balancer blocks the connections that are stealing data within 60 seconds.

When compared to traditional manual methods that take hours, faster containment greatly reduces damage. This shows how AI-enabled independent recovery can effectively stop threats. These performance measures show that the proposed integrated framework can protect sensitive healthcare systems and data in dynamic cloud environments by detecting and responding quickly and accurately.

III. Limitations And Future Work

Deployment problems in real-life, diverse production environments make it hard to test performance in controlled simulations (Chen et al., 2017). The ML-based analysis could be affected by auditing partial visibility, non-standard log schemas, idea drift over time, and attempts to avoid being caught. The framework will be added to real hospital EHR/PACS systems in the future, explainable and self-supervised learning will be added, scaling to handle exabytes of audits will be looked at, and the costs and benefits of using this framework compared to standard methods will be weighed. More testing in the real world will show that the method really does work to protect people. In general, interesting early results are shown about how AI-powered systems might be able to protect cloud-based healthcare delivery from complex hacks by detecting and responding quickly and intelligently.
.

IV. Conclusion

Distributed cloud environments are being used to store and process more and more electronic health data (EHR). If someone wants to steal private patient info, they are getting smarter and bigger. Defenses that are built on signatures take a long time to find new ways to attack. Giving devices the ability to find threats on their own and plan how to act automatically could be one way to use AI to make security better.

The study suggested using an AI-powered, unified structure to keep cloud-hosted EHR systems safe from cyberattacks by constantly watching for threats and quickly stopping and recovering them. Different audit data streams were looked at, and normal and unusual actions were recorded in real time using specialized modeling methods. We tested the framework with real-life ransomware, insider data theft, and infiltration attacks. It was able to quickly find intrusions with over 96% accuracy and stop their effects within minutes by using policy-driven playbooks.

The smarter approach using representation learning and constant adaptation made it possible to stop both known and unknown attacks without having to do any direct programming. It was better than standard rule and signature-based methods. Setting up an end-to-end solution with threat intelligence, reaction coordination, and visibility dashboards showed that security operations could be fully automated and meet the needs of modern distributed healthcare IT environments. People thought the planned AI-driven cyber defense paradigm would work in the real world, especially in sensitive areas like healthcare, because of how well it was designed, how well it worked technically, and how well it was tested.

This work is significant because it suggests several ways to improve policies, technologies, and studies that are used to make important systems more cyber-resilient:

Regulatory agencies should either push healthcare organizations to use AI-enhanced cybersecurity as part of national digital strategies and compliance standards that aim to stop, identify, and lessen threats. Governments can speed up smart research and pilot projects to encourage tech innovation that meets the needs of the health business.

For healthcare CISOs, AI operations programs that mix specialized modeling with reaction automation should be more important than rule-based security information and event management. Because they are so important, transition roadmaps are needed to try AI defenses on less important assets slowly before they are put into use. As an extra to their managed security services, cloud service providers can offer threat monitoring, analysis, and removal services that are driven by AI and come as cloud APIs or SaaS. They could then use AI defense features that are flexible and work best with multi-cloud systems without having to buy a lot of new hardware.

Applied AI safety research that blends representation learning, control theory, and response science should be sped up by academic groups and cybersecurity companies. This will help make reliable methods that can be used for production-level autonomous cyber defense. We need open-source tools and benchmarks to make speed baselines more consistent. For responsible innovation to happen, policymakers and funding groups need to set technical standards for assurance-critical AI systems like healthcare security that anyone can follow, explain, and check. Also very important are ethical implementation rules that find a balance between privacy, accuracy, and ness.

In general, the first demonstrations looked good. However, AI-powered cyber defense needs to be tested more in complex real-world healthcare IT environments, with ongoing attacker simulation techniques, thorough performance benchmarking, and governance standards. This will make it a foundational capability for protecting national healthcare systems safely in the future. This new way can completely change cyber-resilience through intelligence, speed, and independence if smart rules are followed and everyone works together.

References

[1] Alwaheidi, M. K., & Islam, S. (2022). Data-Driven Threat Analysis For Ensuring Security In Cloud Enabled Systems. Sensors, 22(15), 5726.
[2] Alwaheidi, M. K., & Islam, S. (2022). Data-Driven Threat Analysis For Ensuring Security In Cloud Enabled Systems. Sensors, 22(15), 5726.
[3] Chen, P. P., Tsui, N. T., Fung, A. S., Chiu, A. H., Wong, W. C., Leong, H. T., ... & Lau, J. Y. (2017). In-Situ Medical Simulation For Pre-Implementation Testing Of Clinical Service In A Regional Hospital In Hong Kong. Hong Kong Medical Journal, 23(4), 404.
[4] Chowdhury, M. G. U. (2016). Securing Web Systems: A Case Study Of Cybercrime Web Database System (Doctoral Dissertation, University Of Massachusetts Lowell).
[5] Dawood, M., Tu, S., Xiao, C., Alasmary, H., Waqas, M., & Rehman, S. U. (2023). Cyberattacks And Security Of Cloud Computing: A Complete Guideline. Symmetry, 15(11), 1981.
[6] Dawood, M., Tu, S., Xiao, C., Alasmary, H., Waqas, M., & Rehman, S. U. (2023). Cyberattacks And Security Of Cloud Computing: A Complete Guideline. Symmetry, 15(11), 1981.
[7] Furber, S. B., Lester, D. R., Plana, L. A., Garside, J. D., Painkras, E., Temple, S., & Brown, A. D. (2012). Overview Of The Spinnaker System Architecture. Ieee Transactions On Computers, 62(12), 2454-2467.
[8] Goshisht, M. K. (2024). Machine Learning And Deep Learning In Synthetic Biology: Key Architectures, Applications, And Challenges. Acs Omega, 9(9), 9921-9945.

[9]   Iqbal, Waheed & Dailey, Matthew & Carrera, David. (2015). Unsupervised Learning Of Dynamic Resource Provisioning Policies For Cloud-Hosted Multitier Web Applications. Ieee Systems Journal. 10. 1-12. 10.1109/Jsyst.2015.2424998.

[10]   Jpc Rodrigues, J., De La Torre, I., Fernández, G., & López-Coronado, M. (2013). Analysis Of The Security And Privacy Requirements Of Cloud-Based Electronic Health Records Systems. Journal Of Medical Internet Research, 15(8), E186.

[11]   Lasker, A., Obaidullah, S. M., Chakraborty, C., & Roy, K. (2022). Application Of Machine Learning And Deep Learning Techniques For Covid-19 Screening Using Radiological Imaging: A Comprehensive Review. Sn Computer Science, 4(1), 65.

[12]   Liu, Lan & Wang, Pengcheng & Lin, Jun & Liu, Langzhou. (2020). Intrusion Detection Of Imbalanced Network Traffic Based On Machine Learning And Deep Learning. Ieee Access. Pp. 1-1. 10.1109/Access.2020.3048198.

[13]   Luckham, D. C., Vera, J., & Meldal, S. (1995). Three Concepts Of System Architecture. Computer Systems Laboratory, Stanford University.

[14]   Maheshwar, R., & Kandakatla, M. A Deep Learning Framework And Algorithms For Automatic Cyber Attacks Detection.

[15]   Mano, M. M. (1993). Computer System Architecture. Prentice-Hall, Inc..

[16]   Mcghin, T., Choo, K. K. R., Liu, C. Z., & He, D. (2019). Blockchain In Healthcare Applications: Research Challenges And Opportunities. Journal Of Network And Computer Applications, 135, 62-75.

[17]   Özerol, G., & Arslan Selçuk, S. (2023). Machine Learning In The Discipline Of Architecture: A Review On The Research Trends Between 2014 And 2020. International Journal Of Architectural Computing, 21(1), 23-41.

[18]   Pal, R., & Kumar, N. N. (2019). Cyber Security.

[19]   Parkavi, R., Iswarya, M. J., Kirithika, G., Madhumitha, M., & Varsha, O. (2023). Data Breach In The Healthcare System: Enhancing Data Security. In Handbook Of Research On Data Science And Cybersecurity Innovations In Industry 4.0 Technologies (Pp. 418-434). Igi Global.

[20]   Parkavi, R., Iswarya, M. J., Kirithika, G., Madhumitha, M., & Varsha, O. (2023). Data Breach In The Healthcare System: Enhancing Data Security. In Handbook Of Research On Data Science And Cybersecurity Innovations In Industry 4.0 Technologies (Pp. 418-434). Igi Global.

[21]   Pivoto, D. G., De Almeida, L. F., Da Rosa Righi, R., Rodrigues, J. J., Lugli, A. B., & Alberti, A. M. (2021). Cyber-Physical Systems Architectures For Industrial Internet Of Things Applications In Industry 4.0: A Literature Review. Journal Of Manufacturing Systems, 58, 176-192.

[22]   Soori, M., Arezoo, B., & Dastres, R. (2023). Artificial Intelligence, Machine Learning And Deep Learning In Advanced Robotics, A Review. Cognitive Robotics.

[23]   Ullah, Faheem & Edwards, Matthew & Ramdhany, Rajiv & Chitchyan, Ruzanna & Ali Babar, Muhammad & Rashid, Awais. (2017). Data Exfiltration: A Review Of External Attack Vectors And Countermeasures. Journal Of Network And Computer Applications. 101. 10.1016/J.Jnca.2017.10.016.

[24]   Upadhyay, U., Kumar, A., Roy, S., Rawat, U., & Chaurasia, S. (2023, November). Defending The Cloud: Understanding The Role Of Explainable Ai In Intrusion Detection Systems. In 2023 16th International Conference On Security Of Information And Networks (Sin) (Pp. 1-9). Ieee.