Journal of Information Technology Management
ISSN-OnLine: 2980-7972
University of Tehran

Synthetic Immunity: A Generative AI Framework for Privacy-Aware Cyber Risk Management in Healthcare Information Systems

Raghad Tariq Al-Hassani *

*Corresponding author, Ministry of Higher Education and Scientific Research, Minister's Office, Baghdad, Iraq. E-mail: eng_raghadtariq@yahoo.com

Zainab Ali Abbood

Department of Computer Technology Engineering, Al-Mansour University College, Baghdad, Iraq. E-mail: zainab.a.abbood@muc.edu.iq

Adil M. Salman

Department of Cybersecurity Engineering, College of Engineering, Al-Iraqia Science University, Baghdad, Iraq. E-mail: adelmsk63@baghdadcollege.edu.iq

Haider D. Albonda

Department of Control and Systems Engineering, College of Artificial Intelligence Engineering, University of Technology, Baghdad, Iraq. E-mail: haider.d.albonda@uotechnology.edu.iq

Journal of Information Technology Management, 2026, Vol. 18, Special Issue, pp. 46-77.
Published by the University of Tehran, College of Management
doi: https://doi.org/10.22059/jitm.2026.109973
Article Type: Research Paper
© Authors
Received: August 24, 2026
Received in revised form: October 18, 2026
Accepted: November 27, 2026
Published online: December 01, 2026

Abstract

Effective management of healthcare information systems is crucial to the proper functioning of healthcare organizations, as they need to be secure to protect the privacy of patient data, maintain their critical operations, and aid in the secure and effective decision-making processes of the organization. But as threats become more sophisticated, strict privacy needs and the evolving nature of healthcare environments have highlighted the shortcomings of traditional ID methods. This study presents a framework for cybersecurity management based on generative artificial intelligence, called the Synthetic Immunity Framework (SIF), with Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), differential privacy, and an intrusion cognition engine that mimics the adaptive behavior of biological immune systems. The proposed framework produces realistic synthetic healthcare behavioral data, builds a dynamic synthetic immune memory, and conducts context-aware intrusion analysis while maintaining the confidentiality of sensitive patient information. Experimental evaluation was performed on privacy-preserving access logs of Electronic Health Record

(EHR) and data from events in a healthcare system. This framework is able to detect intrusion with an accuracy of 96%, which is greater than that of the traditional approach based on machine learning and deep learning. In addition, synthetic data generation did not require patient-sensitive data, and differential privacy achieved good detection performance across various privacy settings. The results show that the proposed approach enhances the capability of intrusion detection while enabling privacy-preserving management of cybersecurity, mitigating cyber risk, and facilitating decision-making in the context of healthcare information systems, which makes it a promising solution for future smart hospitals and Internet of Medical Things (IoMT) networks.

Keywords: Healthcare Information Systems, Cybersecurity Management, Synthetic Immunity Framework, Generative Artificial Intelligence, Differential Privacy, Intrusion Detec-tion

Introduction

The proliferation of Electronic Health Records (EHRs), Internet of Things (IoT), and the Internet of Medical Things (IoMT) has revolutionized the landscape of healthcare information systems, introducing new challenges and complexities in the management and security of healthcare data. Patient information is vital in a healthcare organization, and with the reliance on interconnected digital infrastructures, safeguarding patient data and ensuring uninterrupted healthcare services have become key organizational concerns. Patient information in healthcare is considered one of the most sensitive types of data to be stolen by hackers, due to the abundance of highly sensitive information such as patient records, diagnostics, treatment history, and real-time physiological data. This makes the possible ransomware attack, breach of unauthorized access to EHRs, insider threats, and attacks on interconnected medical devices a threat to patient safety and healthcare service availability, organizational resilience, and clinical decision-making processes within healthcare institutions (Alauthman et al., 2026; Mehmood et al., 2026).

ISO/IEC 27001 and the NIST Cybersecurity Framework (CSF) are traditional information security frameworks that offer extensive information security management guidance based on risk assessment, access control, continuous monitoring, and incident response. Although these frameworks are largely reactive and perimeter-based, they are not very effective at detecting adaptive, intelligent, and unknown cyber threats in highly dynamic healthcare environments (Malatji, 2023; Pascoe et al., 2024). Healthcare information systems are becoming more interconnected, data-driven, cloud-based, and intelligent; thus, organizations increasingly look for intelligent cybersecurity solutions to increase operational resilience, secure information assets, manage cybersecurity, and comply with regulatory and privacy policies.

These challenges can now be met with the promise of recent advancements in Artificial Intelligence (AI) and Generative Artificial Intelligence (Generative AI). Generative models like Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) can create realistic synthetic healthcare data while maintaining statistical properties of original

datasets without revealing sensitive patient information (Giuffrè & Shung, 2023; Yoon et al., 2020). These features enable secure data sharing, privacy-conscious machine learning, intelligent anomaly detection, and help address healthcare privacy compliance standards like HIPAA and GDPR. While adversarial learning is used to create highly realistic synthetic samples for GANs, structured latent representations are learned for healthcare analytics and privacy-preserving applications using VAEs. These technologies, taken together, hold significant promise for improving the intelligent detection of threats, ensuring secure data use, adaptive cybersecurity, and privacy-centric analytics in contemporary healthcare settings.

In addition to their technical features, intelligent cybersecurity solutions have emerged as a critical part of healthcare IT management, playing a vital role in shaping the resilience, compliance, continuity, and security of healthcare assets. To this end, modern healthcare organizations need a cybersecurity framework capable of detecting cyber threats, enabling informed managerial decisions, facilitating cyber risk management, and strategic governance of healthcare information systems. Therefore, the combination of privacy-preserving generative AI and adaptive intrusion cognition offers a viable solution to enhance cybersecurity efforts and IT management in increasingly complex healthcare settings. This approach inspires the design of the proposed Synthetic Immunity Framework (SIF), an adaptive generative intelligence and privacy-preserving framework for advanced cyber defense and effective management of healthcare information systems.

Research Motivation

While AI-powered cybersecurity tools have made significant strides in healthcare, safeguarding sensitive patient data while ensuring the system can accurately and adaptively detect intrusions remains a primary challenge. In the healthcare sector, there is the constant production of diverse data from EHR systems, IoMT devices, wearable sensors, and cloud-based healthcare services. Traditional intrusion detection systems (IDS) are largely signature-based or have pre-defined behavior and thus have limited ability to detect new attacks and retain patient privacy. These challenges drive the creation of a privacy-enhancing, adaptive, and intelligent cybersecurity framework that can learn from the dynamic behaviors in healthcare and proactively tackle new cyber risks.

Research Gap

Recent research results demonstrate the effectiveness of machine learning, deep learning, federated learning, and privacy-protection methods in healthcare cybersecurity, but there are significant research gaps to be addressed (Bishnu et al., 2026; Meena et al., 2026). Most existing solutions are dedicated to enhancing either the accuracy of intrusion detection, the protection of privacy, or the generation of synthetic data, separately. There are few studies that combine Generative AI, synthetic healthcare data generation, privacy-preserving learning, adaptive intrusion cognition, and contextual threat interpretation in a single

cybersecurity approach. In addition, existing systems are not equipped with biologically-inspired adaptive mechanisms that can build a synthetic immune memory that is constantly evolving to recognize new cyberattacks that haven't been encountered before. The aforementioned restrictions indicate the necessity of developing a comprehensive system that focuses on preserving privacy, learning adaptively, and understanding the intrusions in healthcare network systems.

To overcome these drawbacks, this paper introduces a Synthetic Immunity Framework (SIF) that is based on the adaptive features of the biological immune system. The proposed framework is not just an ordinary anomaly detection system; it also proposes a new notion, that of intrusion cognition, which, in addition to detecting suspicious activities, also attempts to interpret suspicious behaviors, comprehend the context in which they take place, and adjust in real time to new cyber threats. In this context, GANs and VAEs jointly learn to generate a synthetic immune memory that can identify abnormal ("non-self") healthcare behavior from both synthetic and real healthcare data, as well as normal ("self") healthcare behavior. Also, the framework incorporates differential privacy, encryption mechanisms, and secure data handling practices to guarantee the protection of sensitive patient information while also enabling smart healthcare cybersecurity (de Medeiros et al., 2025).

Research Objectives

In this context, the main goal of this study is to create a Hybrid GAN–VAE Learning-based private Synthetic Immunity Framework for healthcare network systems, which incorporates the elements of synthetic immune memory, differential privacy, and cognitive intrusion analysis into a single framework. In particular, the objective of this study is to:

a) Generate realistic and privacy-preserving synthetic healthcare datasets for secure model training.

b) Construct a synthetic immune memory capable of continuously learning evolving healthcare behaviors.

c) Improve the detection of known and previously unseen cyber threats through cognitive intrusion analysis.

d) Use differential privacy and secure data management systems to preserve sensitive patient information.

e) Enable real-time, intelligent threat monitoring and adaptive cybersecurity decision-making in healthcare settings.

Controlled healthcare simulations and synthetic datasets that mimic realistic EHR access logs and healthcare network activities are used to validate the framework conceptually.

Results from the experimental evaluation confirm the effectiveness of the proposed framework in terms of detection accuracy, precision, recall, and privacy preservation, suggesting that it can be implemented in future intelligent healthcare cybersecurity systems.

Main Contributions

The key innovation of this study is the seamless combination of Generative AI (GenAI), Synthetic Immune Memory, cognitive intrusion analysis, and privacy-preserving learning within a comprehensive healthcare cybersecurity ecosystem. The proposed Synthetic Immunity Framework integrates adaptive learning, generation of synthetic data, contextual threat interpretation, and privacy protection in compliance with regulations in one intelligent architecture, thereby delivering an integrated approach that goes beyond the current approaches, which are primarily focused on either intrusion detection or privacy preservation. The major scientific achievements of this work have been summarized as follows:

a) A novel SIF inspired by the adaptive characteristics of biological immune systems is proposed for privacy-aware healthcare cybersecurity.

b) A unified architecture integrating (GANs, VAEs), differential privacy, and cognitive intrusion analysis is developed to provide both intelligent threat detection and privacy preservation.

c) A SIM mechanism is introduced to continuously learn representative healthcare behavioral patterns and improve the detection of previously unseen and zero-day cyber threats.

d) Privacy-preserving learning is achieved through the integration of differential privacy and secure data handling mechanisms, enabling compliance with healthcare regulations such as HIPAA and GDPR without exposing sensitive patient information.

e) The experimental validation protocol is based on healthcare-related datasets, showing that the proposed framework provides better intrusion detection capabilities, high model generalization through the generation of synthetic data, and greater privacy preservation than traditional baseline frameworks.

The paper is organized as follows: Section 2 examines recent studies on the relationship between generative AI and healthcare security. In Section 3, the proposed Synthetic Immunity Framework is presented. The implementation, experimental analysis, and discussion of the results obtained is described in Section 4. Finally, Section 5 is dedicated to summarizing the paper and pointing towards some future research lines.

Literature Review

The only thing that has changed is that the security solutions of today's digital infrastructure are no longer static environments, but are more intelligent, privacy-centric, and adaptive to address the ever-evolving landscape of cyber threats (Aljumaiah et al., 2025). Current cybersecurity methodologies, like NIST Cybersecurity Framework and ISO/IEC 27001 (Malatji, 2023; Pascoe et al., 2024), have explicit methods for assessing risk, for managing access points, and for responding to incidents, but these remain reactive rather than proactive and are rule-based. In fact, they play a critical role in healthcare environments where timely decision-making is crucial, and data security is paramount. All these challenges have been the subject of a lot of studies in the last few years to incorporate Artificial Intelligence (AI) into cybersecurity systems. In particular, generative models and deep learning techniques have proven to be very effective in enhancing data privacy and anomaly detection.

Taxonomy of AI-Driven Cybersecurity Approaches

While all of these advancements have positively enhanced, there are still some fundamental challenges with existing AI-driven cybersecurity solutions:

a) Detection-Centric Design: All current methods are mainly based on detecting anomalies and lack any contextual knowledge or interpretation of cyber threats. This is limiting their capacity to aid intelligent decision-making (Alkasassbeh & Al-Haj Baddar, 2023).

b) Limited Cognitive Awareness: Current models don't use context-aware reasoning in making decisions about attacks, interpreting behavioral patterns, or adjusting when they know about the context, which is not useful in a complex environment such as a healthcare system (Ahmad et al., 2021).

c) Privacy-utility trade-off: Preserving privacy through the creation of synthetic data has a price in terms of utility, and ensuring a balance between utility and privacy is a tough challenge, particularly in regulated areas (Kaabachi et al., 2025).

d) Limited Proactive Capabilities: Most of the current systems and systems using artificial intelligence are dependent on historical data and are not good at identifying zero-day attacks and unknown intrusion patterns (Alkasassbeh & Al-Haj Baddar, 2023).

e) Non-Domain Specific Adaptation: The majority of the frameworks are general and do not consider the particular requirements of healthcare systems, including regulatory requirements (such as HIPAA), real-time reactivity, and patient data sensitivity.

Generative Models for Cybersecurity and Data Privacy

Generative Adversarial Networks (GANs) (Goodfellow et al., 2014; Yang et al., 2023) have recently emerged as a technique for generating synthetic data. In this framework, the generator learns to create synthetic data, while the discriminator learns to distinguish between real and generated data. More advanced models, such as StyleGAN2 and BigGAN, have demonstrated remarkable capabilities in data fidelity, making them valuable for cybersecurity simulations and data augmentation (Karras et al., 2020).

Similarly, Variational Autoencoders (Wei & Mahmood, 2020) provide a probabilistic representation and reconstruction of data, which are efficiently represented and controlled in the latent spaces. VAEs are well suited for applications where privacy is paramount, and direct access to sensitive data is not allowed.

The importance of synthetic data has emerged in recent times, particularly in the aspects of mitigating privacy concerns while maintaining the analytical value. However, GANs and VAEs have been shown to be capable of generating realistic, but non-identifiable data sets that can be used for machine learning applications, overcoming classic anonymization methods (Tucker et al., 2020).

Advanced generative architectures, such as autoregressive models (e.g., PixelCNN) and flow-based models (e.g., RealNVP and Glow), can model complex data distributions and simulate various types of attacks (van den Oord et al., 2016). These developments have paved the way for generative AI to play an important role in cybersecurity, particularly in intrusion detection and threat modeling.

Limitations of Existing AI-Based Cybersecurity Approaches

To systematically recognize these limitations, a comparison of the approaches and gaps found in the existing studies was provided in Table 1. The existing studies consider cybersecurity as a static detection problem and not as a dynamic cognitive process. In addition, however, few efforts have been made to create a single, configurable architecture that addresses all of the challenges of privacy, intelligence, and domain-specific restrictions, as well as takes advantage of generative AI.

Table 1. Critical analysis of existing approaches and research gaps

Category	Existing Approaches	Key Limitations	Research Opportunities	Refs.
ML-Based Detection	RF, SVM	Limited generalization and static learning	Integrate adaptive and context-aware learning	(Ahmad et al., 2021; Alkasassbeh & Al-Haj Baddar, 2023)
Generative AI	GANs, VAEs	Primarily used for data generation rather than system-level intelligence	Embed generative models within adaptive security frameworks	(Tucker et al., 2020; Yang et al., 2023)
Privacy Mechanisms	Differential privacy, encryption	Trade-off between privacy and model performance	Develop privacy-aware learning frameworks	(Abadi et al., 2016; Mosaiyebzadeh et al., 2024)
Domain Adaptation	General-purpose frameworks	Limited adaptation to healthcare-specific requirements	Design healthcare-aware cybersecurity architectures	(Khan et al., 2024; Mosaiyebzadeh et al., 2024)
Threat Detection	Reactive anomaly detection	Limited ability to detect unknown and zero-day attacks	Enable proactive and predictive threat modelling	(Alkasassbeh & Al-Haj Baddar, 2023; Gueriani et al., 2024)
DL-Based Models	LSTM, CNN	High data dependency and computational complexity	Optimize models for real-time healthcare environments	(Gueriani et al., 2024)

In order to tackle the above-mentioned drawbacks, this study presents a Synthetic Immunity Framework, which is inspired by the biological immune system, where the concept of Cybersecurity is redefined as an adaptive, self-evolving, and cognitive process. The proposed framework is novel in the following aspects compared to the existing approaches:

a) Intrusion Cognition: Moving beyond detection to enable interpretation and contextual understanding of cyber threats.

b) Synthetic Immune Memory: Utilizing generative AI to simulate past and potential attack patterns for proactive defense.

c) Privacy-Aware Learning: Combining synthetic data generation techniques with differential privacy to meet healthcare regulations.

d) Security and Privacy: Specialized security and privacy measures for healthcare networks, such as interoperability between IoMT devices and EHRs.

Recent research has focused on the use of artificial intelligence, privacy-preserving techniques, and intelligent intrusion detection methods to tackle the ever-evolving cybersecurity landscape in the healthcare sector. But most current approaches concentrate on a single or two of the aspects. To further emphasize the novelty of the proposed Synthetic Immunity Framework (SIF), the proposed approach is compared with representative state-of-the-art studies in terms of the AI techniques, applicability to healthcare, privacy preservation, synthetic data generation, cognitive intrusion analysis, and adaptive learning capabilities available in the proposed approach and the representative state-of-the-art studies in Table 2.

Table 2. Comparison of Existing Healthcare Cybersecurity Approaches and the Proposed SIF

Study	Main Technique	Main Limitation	Research Gap
Mehmood et al. (2026)	AI-powered generative architecture for healthcare cybersecurity	Emphasizes scalable cybersecurity, privacy-aware synthetic-data generation, and anomaly detection without a biologically inspired adaptive immune-memory mechanism	Does not incorporate an evolving synthetic immune memory or explicitly formulate intrusion cognition
Mosaiyebzadeh et al. (2025)	Privacy-preserving federated-learning IDS incorporating differential privacy	Focuses on distributed, privacy-preserving intrusion detection without integrating generative learning or biologically inspired cognitive adaptation	Does not combine federated detection with synthetic immune memory, generative threat modelling, and contextual intrusion cognition
Gueriani et al. (2024)	CNN–LSTM-based intrusion detection for IoT environments	Depends primarily on previously observed training patterns and may have limited adaptability to evolving and previously unseen threats	Does not integrate generative learning, synthetic immune memory, or context-aware intrusion interpretation
Proposed SIF	Hybrid GAN–VAE architecture with synthetic immune memory	Requires comprehensive experimental validation of its privacy, utility, adaptability, computational efficiency, and zero-day detection capabilities	Integrates generative learning, adaptive intrusion cognition, privacy preservation, contextual threat interpretation, and evolving synthetic immune memory within a unified healthcare cybersecurity framework

As can be seen from Table 2, the existing efforts in healthcare cybersecurity focus mostly on single areas like intrusion detection, privacy protection, or synthetic data generation. Synthetic Immunity Framework, on the other hand, adopts a hybrid GAN–VAE learning approach and integrates all the above-mentioned capabilities within a unified framework, including differential privacy, synthetic immune memory, cognitive intrusion analysis, and adaptive learning. This comprehensive integration allows the framework to create realistic health care behavioral patterns, maintain patient information, and identify known and unknown cyber threats. This results in the proposed framework being more adaptive and privacy-preserving than the current state-of-the-art methods, which constitutes the main scientific contribution of this work.

Proposed Synthetic Immunity Framework

Sustainable Intelligent Framework (SIF) is an adaptive, privacy-aware, cognition-driven cybersecurity framework that has been proposed for securing healthcare network systems. This framework is inspired by the biological immune system, which learns normal ("self") and malicious ("non-self") behaviors in healthcare continuously by combining Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs). This way, the framework will evolve with time and evolve to detect new cyber threats as they emerge and create a dynamic Synthetic Immune Memory (SIM) for each threat.

The proposed framework is not a conventional intrusion detection system (IDS) that is based on the detection of predefined rules or anomalies, but one in which intrusion cognition

is introduced. The detection of irregular activity is not the only reason for using these, but the analysis of how they are functioning with respect to their operations, which enables more intelligent and proactive cybersecurity decision-making. In the healthcare sector, where the timely and accurate processing of information is crucial for maintaining patient confidentiality and uninterrupted clinical operations, this capability proves to be very beneficial.

To safeguard data confidentiality, the framework incorporates various methods for differential privacy, secure data handling, and generating synthetic data, ensuring that AI models can be trained on representative healthcare data without compromising sensitive patient information. Generative AI powered by privacy-preserving technologies strikes a balance between detection accuracy and regulatory compliance, meeting healthcare standards such as HIPAA and GDPR. The overall architecture of the proposed Synthetic Immunity Framework is depicted in Figure 1, with the generative models being connected to the privacy-preserving devices, the Synthetic Immune Memory and the Intrusion Cognition Engine.

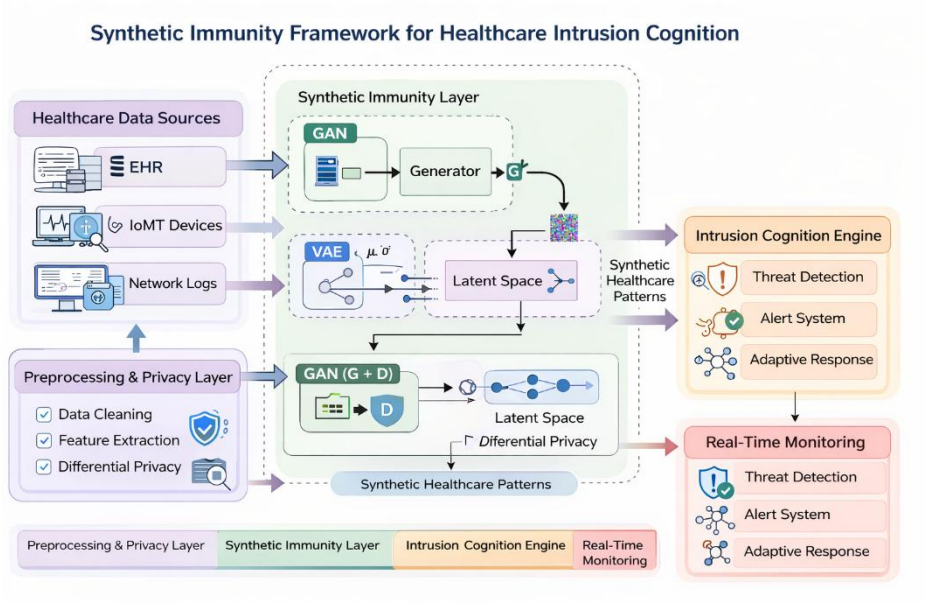


Figure 1. Synthetic Immunity Framework for Healthcare Intrusion Cognition

In the proposed approach, an adversarial learning method is used, in which Generator (G) generates realistic healthcare behavioral patterns from latent noise and context information, while the Discriminator (D) distinguishes between real and synthetic samples. In this iterative adversarial training process, both models gradually learn to better approximate the complicated distributions of healthcare data, as well as to produce statistically consistent synthetic samples. Generated samples undergo realism and context evaluation before being added to the training set. High-quality synthetic samples are retained; low-quality samples are discarded in order to enrich downstream machine learning and deep learning models. This

control mechanism improves the generalization of the model and avoids propagation of unrealistic data.

The generator and discriminator continuously interact, creating an adaptive learning cycle similar to the biological immune system. New behaviors for the healthcare sector and new attack patterns are added to the Synthetic Immune Memory, which continually updates context-aware intrusion cognition and allows proactive detection of new cyber threats. The full process for the proposed framework when deploying it in real-time is described by Algorithm 1, which outlines the learning process from data pre-processing to real-time deployment.

Algorithm 1. Synthetic Immunity Learning Process

```
Input:
   H_data → Healthcare data (EHR, IoMT, logs)
   ε → Privacy budget
   θG, θD → Generator and Discriminator parameters
Output:
   C_model → Intrusion Cognition Model
Begin
1:  Normalize(H_data)
2:  Extract features (user behavior, timestamps, device activity)
3:  H_priv ← Apply_Differential_Privacy (H_data, ε)
4:  Initialize Generator G (z; θG)
5:  Initialize Discriminator D (x; θD)
6: while not converged do
7:     z ← Random noise vector
8:     X_fake ← G(z)
9:     Train D using real data H_priv and fake data X_fake
10:    Update G to maximize fooling D
11: end while
12: X_syn ← Generate high-quality synthetic healthcare patterns
13: Train intrusion detection models (Random Forest, SVM, and LSTM) using:
       (H_priv + X_syn)
14: C_model ← Learn behavioral patterns (self vs non-self)
15: For each incoming healthcare event x:
16:    score ← Compute_ICS (x)
17:    If score > τ then
18:       Trigger alert and adaptive response
19:    End if
End
```

Algorithm 1 outlines the entire operation procedure of the proposed Synthetic Immunity Framework. The framework starts with data preprocessing and privacy preservation in the healthcare domain and then adversarially trains the GAN-based synthetic data generator and discriminator. The Intrusion Cognition Engine's knowledge base is built by creating synthetic healthcare patterns and combining these with privacy-preserved real healthcare data to create

https://jitm.ut.ac.ir/

the Synthetic Immune Memory (SIM). Lastly, the trained cognitive model conducts real-time intrusion cognition and adaptive response to support continuous learning from new healthcare cyber threats.

Framework Architecture Overview

The proposed Synthetic Immunity Framework (SIF) is a hierarchical solution that integrates generative AI, privacy-preserving technologies, and cognitive intrusion detection to provide a comprehensive cybersecurity solution for healthcare network systems. The design consists of 5 tightly coupled and interdependent layers that can deliver the required functions within the structure as part of the system pipeline.

The proposed architecture is not the traditional layered security as it allows for continuous interaction between layers, allowing on-the-fly adaptation and on-the-fly learning. The interconnection enables the system to respond to new threats in a way similar to how a biological immune system would respond. The 5 elements of this structure are:

a) Data Acquisition Layer: Collects data from diverse sources in the healthcare landscape, such as Electronic Health Records (EHR), IoMT devices, as well as network activity logs. This layer guarantees data availability and integrity in the downstream process.

b) Synthetic Immune Generation Layer: This layer is the backbone of the framework, which aims to generate synthetic healthcare behaviors with the help of GANs and VAEs. This layer builds a dynamic immune memory, which means that it can simulate things like normal and malicious activities.

c) Privacy Preserving Layer: Safeguards patients' sensitive data by implementing differential privacy, encryption, and secure data handling procedures. This layer ensures adherence to healthcare-related rules and regulations like HIPAA and the utility of the data for model training.

d) Intrusion Cognition Engine: Performs sophisticated behavioral analysis with machine learning and deep learning algorithms such as LSTM. This layer adds the notion of intrusion cognition -- not just detection, but understanding of threats in the context of their operation.

e) Real-Time Monitoring Layer: Constant monitoring of healthcare systems, which would enable the real-time detection of threats, generation of alerts, and adaptive response.

Note: These layers are not layers per se, but an intelligent closed-loop system. The data will be passed from the acquisition to the synthetic generation, where synthetic immune

patterns will be created, privacy-preserving mechanisms will be applied before passing them to the cognition engine for analysis. The results are fed into the monitoring layer, which can initiate adaptive responses and dynamically alter the system. It is a feedback-driven architecture, which enables the system to learn more from historical data and synthetic data, and to be more robust and resilient to new cyber threats. Table 3 shows the elements of the synthetic immunity framework.

Table 3. Components of the Synthetic Immunity Framework

Layer	Core Function	AI Technique	Key Contribution	Healthcare Relevance
Data Acquisition	Collect and preprocess heterogeneous healthcare data	Data preprocessing, feature extraction	Enables structured input for AI models	Continuous patient and system monitoring
Synthetic Immunity Layer	Generate synthetic behavioral patterns (self & non-self)	GAN + VAE	Builds adaptive immune memory	Privacy-preserving model training
Privacy Preservation	Protect sensitive data and ensure secure processing	Differential Privacy, Encryption	Ensures data confidentiality and compliance	HIPAA/GDPR compliance
Intrusion Cognition Engine	Analyze and interpret abnormal behaviors	LSTM, ML/DL models	Introduces cognitive threat understanding	Detect unauthorized medical access
Real-Time Monitoring	Detect and respond to threats in real time	Stream processing, real-time analytics	Enables adaptive and low-latency response	Critical care system protection (ICU, IoMT)

Synthetic Immunity Modeling

The Synthetic Immunity Modeling layer is the main intelligence of the proposed framework, using generative AI to simulate the action of a biological immune system. The layer could provide a dynamic and adaptive synthetic immune memory to detect normal (self) and malicious (non-self) patterns in healthcare network environments. To achieve this, it brings in two complementary generative models: the Generative Adversarial Networks (GANs) for adversarial learning and attack simulation, and the Variational Autoencoders (VAEs) for learning latent representations and structured pattern encoding.

Generative Adversarial Learning

Within the proposed framework, GANs simulate the immune response process, where the generator acts as a synthetic pattern creator and the discriminator acts as an immune validator. The adversarial learning process is defined as:

$$min_G max_d V(D,G) = E_{\mathrm{x}\sim \mathrm{p}_{\{\mathrm{data}\}(\mathrm{x})}}\left[\log \mathrm{D}(\mathrm{x})\right] + E_{\mathrm{z}\sim \mathrm{pz}(\mathrm{z})}\left[\log\left(1-\mathrm{D}\left(\mathrm{G}(\mathrm{z})\right)\right)\right] \quad (1)$$

where:

- X represents real healthcare data samples (EHR, IoMT signals)
- Z is a random noise vector
- G(z) is the generated synthetic sample
- D(x) is the probability that the input data are real

The proposed Synthetic Immunity Framework includes a generator that creates synthetic healthcare behavioral patterns and a discriminator that will serve as an immune validator to detect legitimate healthcare activities from those that could be malicious. For the purposes of further explaining the relationship between some of the elements of this generative AI and the elements of the biological immune system, a mapping between the two is provided in Table 4.

Table 4. Synthetic Immunity Interpretation of GAN Components

Component	Synthetic Immunity Role	Functional Description
Generator (G)	Immune Stimulation	Produces synthetic healthcare behaviors representing both normal and anomalous patterns
Discriminator (D)	Immune Recognition	Validates authenticity by distinguishing between real and synthetic data
Adversarial Learning Loop	Immune Adaptation	Enables continuous system learning and adaptation to evolving threats

Conditional and Context-Aware Generation

To enhance realism and applicability in healthcare systems, the generator is extended to a conditional GAN framework:

$$G(z, y) \rightarrow X\,\text{synthetic} \quad (2)$$

where:

- G denotes the generator network.
- z is a latent random noise vector.
- y represents contextual healthcare information.
- Xsynthetic denotes the generated synthetic healthcare sample.

VAE-Based Latent Immune Memory

While GANs focus on adversarial realism, VAEs are used to construct a structured latent immune memory by encoding healthcare data into a continuous latent space.

$$L = Eq(z \mid x)[\log p(x \mid z)] - D_{KL}\,(q(z|x)||p(z)) \quad (3)$$

where:

- $q(z \mid x)$: encoder distribution
- $p(x \mid z)$: decoder distribution
- D_{KL}: divergence between latent and prior distributions

The purpose of this representation is to allow the system to store behavioral patterns of normal system behavior, capture relationships between healthcare events, and support reconstruction and comparison with anomalies.

Hybrid GAN–VAE Integration

This approach is based on GANs and VAEs, which complement each other in their strengths, and is suggested to be used together in a single model. To generate high-quality and realistic synthetic health care data through adversarial learning, GANs are employed, and to generalize and stabilize the representations, VAEs are employed to produce structured and interpretable latent representations.

This integration enables this framework to gain an acceptable compromise between realism of the data and robustness of the model. The probabilistic encoding in VAEs helps to achieve stable training and meaningful organization of the latent space, while the adversarial approach of GANs also guarantees the generation of diverse and high-fidelity samples. Thus, the hybrid model can boost the overall performance of the Synthetic Immunity Framework by enhancing the diversity of the data, ensuring a consistent structure, and avoiding common problems like mode collapse in the standalone GAN model. The objective function of the hybrid model is given by:

$$\mathrm{Ltotal} = \mathrm{LGAN} + \lambda \mathrm{LVAE} \quad (4)$$

Where:

- LGAN represents the adversarial learning loss.
- LVAE denotes the variational autoencoder reconstruction loss.
- λ is the balancing coefficient.

Synthetic Immune Memory Construction

The proposed framework uses the generative training period to create a Synthetic Immune Memory (SIM), which is a dynamic memory of learned healthcare behavior patterns. The memory comprises synthetic normal and anomalous data representations, generated by the hybrid GAN–VAE architecture. The Synthetic Immune Memory can formally be defined as:

$$SIM = \{xnormal, xattack\} \quad (5)$$

Where:

- xnormal represents synthetic samples corresponding to normal healthcare system behavior
- xattack represents synthetic intrusion patterns simulating malicious activities

The built SIM is highly dependent on the Intrusion Cognition Engine to give context-aware analysis. In particular, any events that are received from the health care system are constantly compared with the stored synthetic event patterns, and the system recognizes the deviations from normal behavior.

The proposed SIM is not a single snapshot where it is static, but rather it is a continuous process and will continually change to respond to new cyber threats and changing system dynamics. Here, the Synthetic Immune Memory is able to amplify the capability of the system to recognize deviations from normal operating patterns, which allows for detection of subtle and new anomalies and enables the system to interpret the changes in behavior in a contextual healthcare setting. This adds a lot to the strength and wisdom of the proposed design, extending its generativeness and capacity for cognitive intrusion analysis.

$$ICS = \alpha A + \beta S + \gamma C \quad (6)$$

Where:

- A: Anomaly score
- S: Similarity score with Synthetic Immune Memory
- C: Contextual healthcare behavior score

The overall confidence level of the proposed framework for classification of healthcare events is represented by Equation 6 and is defined as the Intrusion Cognition Score (ICS). The score is a composite of three complementary factors: anomaly score, similarity of the current event to stored Synthetic Immune Memory, and contextual healthcare behavior score. The weighting coefficients (α, β, and γ) indicate the relative weight of each component, and the constraint is $\alpha + \beta + \gamma = 1$ (normalization). If the computed ICS is above the predetermined decision threshold τ, the event is labelled an intrusion; otherwise, it is deemed normal system behavior.

$$\text{Decision} = \begin{cases} \text{Attack}, & \text{if } ICS > \tau \\ \text{Normal}, & \text{Otherwise} \end{cases} \quad (7)$$

- ICS represents the Intrusion Cognition Score computed using Equation (6).
- τ denotes the decision threshold used to distinguish malicious from normal healthcare behaviors.
- Attack indicates that the healthcare event is classified as an intrusion requiring an adaptive security response.
- Normal indicates that the healthcare event is considered legitimate and no security action is required.

The last stage in the proposed Synthetic Immunity Framework is captured in Equation 7. Once the Intrusion Cognition Score is calculated from Equation 6, it compares the computed score with the predefined threshold τ. Events with scores above the threshold are defined as cyber intrusions, while scores below the threshold are considered normal behaviors. This decision mechanism will allow for context-aware and adaptive intrusion detection yet keep the detection accuracy high in healthcare environments.

Relevance to Healthcare Systems

One of the most important problems in healthcare network environments is the lack of availability of data, the sensitivity of the data itself, and the dynamic nature of the system, which are addressed by the proposed Synthetic Immunity Modeling layer. The framework takes advantage of generative models to generate useful synthetic data that lessens the need to access real patient data, yet doesn't sacrifice data utility.

The privacy-enabling features ensure compliance with healthcare regulations, such as HIPAA and GDPR, and enable secure model training. Moreover, its adaptive characteristics permit ongoing learning from the changing data flows within healthcare systems, such as IoMT-based environments. Moreover, the synergistic use of synthetic immune memory and the cognitive intrusion analysis increases the reliability of decision-making, which is particularly important in medical applications. Thus, the suggested framework is appropriate to be implemented in smart hospitals, remote patient monitoring, and IoMT-based health care systems.

Results

Implementation Setup

The effectiveness of the proposed Synthetic Immunity Framework (SIF) has been evaluated through a comprehensive experimental work, which has been conducted in the context of a healthcare cybersecurity environment. The experiments were conducted on a set of healthcare-related datasets from access logs of Electronic Health Records (EHR) and activity

events in a healthcare system. The main goal of this assessment was to determine the framework's ability to effectively identify unusual healthcare activities without compromising sensitive patient data and to comply with healthcare privacy rules.

The proposed framework was implemented using Python 3.10 as the main implementation language. The architecture of the hybrid GAN–VAE, together with the LSTM-based intrusion cognition module, was implemented using TensorFlow 2.x and PyTorch, while conventional machine learning models such as Random Forest and Support Vector Machine (SVM) models were implemented using the Scikit-learn library. The integration of the PySyft framework allowed for incorporating differential privacy mechanisms during model training and evaluation, thereby achieving privacy-preserving learning.

The experimental process is divided into three steps. First, the hybrid GAN-VAE architecture is used to generate synthetic healthcare data to build a Synthetic Immune Memory (SIM) and augment the training dataset with privacy-preserving synthetic samples. Second, the real and synthetic healthcare data are mixed and used for training the intrusion detection models to learn both legitimate and malicious healthcare behaviors. Lastly, the trained models are tested with previously unseen healthcare data to measure their generalization capability, their anomaly detection performance, and their capability to be privacy-preserving. All experimental trials were conducted under the same experimental conditions to ensure reproducibility, scalability, and an even assessment process. Table 5 shows the environment, software tools, learning models used, privacy settings, and hardware used in the experimental evaluation, as well as the key training parameters used.

Table 5. Implementation Configuration for Healthcare Experiments

Component	Description
Programming Language	Python 3.10
Deep Learning Framework	TensorFlow 2.x / PyTorch
Machine Learning Models	Random Forest, SVM, LSTM
Generative Models	GAN, VAE
Privacy Technique	Differential Privacy (ε = 0.1)
Data Type	EHR Access Logs, Healthcare System Events
Dataset Size	8,000 Records
Training Strategy	80% Training / 20% Testing
Learning Rate	0.001
Batch Size	32
Number of Epochs	100
Optimizer	Adam
Latent Dimension	128
Hardware Configuration	Intel Core i7 Processor, 16 GB RAM
Evaluation Metrics	Accuracy, Precision, Recall, F1-Score

The values for the hyperparameters reported in Table 5 are typical of those used in recent cybersecurity and deep learning studies in healthcare, and were initially tuned in small trials to ensure stability of convergence and good detection performance. It's because of its

adaptive optimization capability and quick convergence qualities that the Adam optimizer was picked. The setting of a learning rate of 0.001 and a batch size of 32 resulted in a good compromise between the computational expense and model accuracy. The latent dimension was fixed at 128 to avoid having a too complex model and to keep the representative healthcare behavioral features. In addition, the differential privacy parameter ($\varepsilon = 0.1$) was chosen to ensure a high level of privacy protection without compromising the performance of anomaly detection. In all the experimental tests, the hyperparameters were kept constant to give an objective, consistent, and repeatable assessment of the system under different experimental conditions.

The proposed implementation configuration, detailed in Table 5, combines generative AI models, privacy-preserving approaches, and machine learning techniques into a cohesive healthcare cybersecurity system. The architecture sought to appropriately integrate speed, efficiency, scalability, and data privacy requirements for patients with interoperability and compatibility requirements for existing health care infrastructures.

Access to large-scale real-world medical data is extremely restricted because of patient information sensitivity and the strict regulations and standards for management of medical data. To assess the proposed framework and subject it to privacy regulations and ethical data handling practices, controlled simulations along with synthetic healthcare-like datasets were used. This approach will enable you to carry out a comprehensive evaluation of the experiment without any confidential patient data being shared.

For reliable evaluation of the experiments, each experiment was repeated a few times under the same experimental conditions. All of the performance metrics in the remainder of the document are the average results of all the experimental runs, thus minimizing the effect of random variation and yielding consistent, robust, and reproducible results.

Healthcare Case Study

To evaluate the overall effectiveness of the proposed Synthetic Immunity Framework (SIF), a healthcare-related case study was conducted using access logs of Electronic Health Records (EHR) and activity events of a healthcare system. This assessment aims to ascertain if the framework is applicable to be able to differentiate between legitimate healthcare operations and malicious ones without losing the privacy-preserving requirements.

The experimental data consists of approximately 8000 data samples with information about user access actions and events generated by the system. The dataset comprises normal and abnormal activities, including unauthorized access, privilege misuse, and other suspicious actions, to mimic potential attacks in healthcare settings. To ensure the confidentiality of data and comply with healthcare regulations, differential privacy ($\varepsilon = 0.1$) and encryption were used during preprocessing and training, while data masking methods were employed. Before

evaluating the results of anomaly detection from the proposed framework, it is essential to have a good understanding of the healthcare dataset. Table 6 presents the distribution of events used in the experiments, categorized as normal or abnormal.

Table 6. Healthcare Dataset Distribution

Data Type	Number of Records	Description
Normal Activities	7,970	Legitimate EHR access and healthcare system operations
Anomalous Events	30	Simulated unauthorized access and abnormal behavioral activities
Total	8,000	Combined healthcare dataset

As seen in Table 6, the data is very skewed, characteristic of realistic healthcare systems, in which only a small percentage of incidents involve security. The imbalance is a key problem for IDSs and will make the proposed Synthetic Immunity Framework a challenging test case. A confusion matrix analysis was performed for further evaluation of the detection capability of the proposed framework. The classification results of the actual class with the predicted class are plotted in Figure 2, and the results show the accuracy of the framework in predicting the actual class for the normal and malicious healthcare activities.

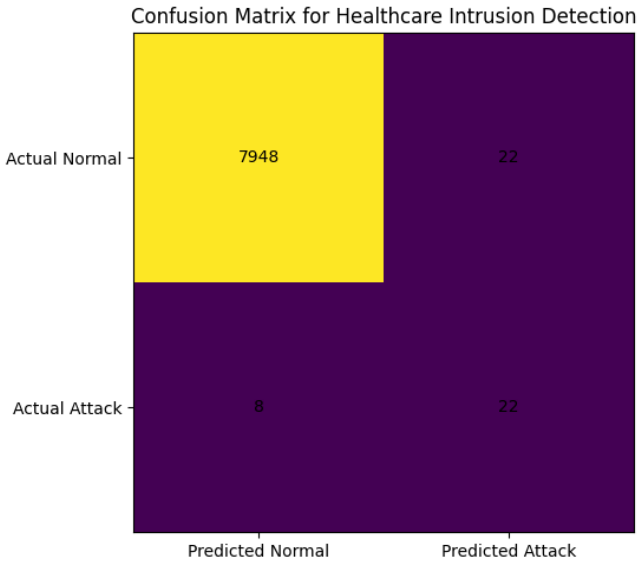


Figure 2. Confusion Matrix of Healthcare Intrusion Detection Using the Proposed Synthetic Immunity Framework.

The confusion matrix shows that the proposed framework is able to accurately separate normal healthcare activities from the malicious ones with high accuracy. The number of false negatives and false positives is very small, so the classification is good despite the very imbalanced data. The results confirm that the Synthetic Immunity approach can effectively detect rare and critical threats and reduce false alarms.

Furthermore, the flexibility and strength of the framework are demonstrated by the ability to recognize anomalous events accurately, based on a small sample of attack events. While cyberattacks in healthcare are few, they can cause devastating impacts to patient health and

privacy, as well as impact service availability. The findings are promising and suggest that similar arrangements can be implemented in a proactive and privacy-preserving manner to monitor cybersecurity in today's health-care systems.

Performance Evaluation

A detailed performance analysis was carried out to measure the efficiency of the proposed synthetic immunity framework (SIF) by adhering to the measures of standard classification, including accuracy, precision, recall, and F1 score. These metrics help to strike the right balance between the correct detection of anomalous healthcare activities and minimizing false alarms. In addition, the processing time was also measured to evaluate the suitability of the framework regarding real-time applications with regard to healthcare cybersecurity. The healthcare testing dataset defined in Section 4.2 was used for the evaluation. The reported results are the average, and several experimental runs were carried out to achieve consistency and reliability.

Table 7. Performance Evaluation in Healthcare Environment

Metric	Value
Accuracy	0.96
Precision	0.93
Recall	0.93
F1-Score	0.93
Processing Time	1.5 s

As it can be noticed in Table 7, the proposed framework has a high accuracy of 96%, that is, it has the ability to differentiate very well between the legitimate healthcare activities and the malicious ones. A balanced performance indicates the accuracy and precision of the threat identification process, with indicators of 93% precision and recall rate, which helps reduce the risk of false alarms and false negative threats. Therefore, the F1 score obtained is proof of the hardness and reliability of the model in highly imbalanced healthcare datasets.

Further, the processing time for the approach in the proposed design (1.5 seconds) was measured, which proves the computational efficiency of the proposed approach and enables it to be used for near real-time healthcare intrusion monitoring. The effectiveness of the Synthetic Immunity Framework is highlighted by its detection accuracy, ability to handle powerful cyber classification, and low processing cost, which are demonstrated and are effective in protecting the modern healthcare infrastructure system in the ever-changing cyber environment.

Privacy Compliance Evaluation

One of the most crucial requirements of the modern Healthcare Cybersecurity system is the security of sensitive Healthcare details. In order to address this problem, the proposed Synthetic Immunity Framework (SIF) has been developed, which is based on a multi-layered

approach to privacy preservation, namely differential privacy, encryption, and data masking. These are ways to comply with important healthcare policies like the Health Insurance Portability and Accountability Act (HIPAA) and the General Data Protection Regulation (GDPR) that ensure the analytics used for effective intrusion detection remain the same.

A series of experiments were conducted using different settings of differential privacy, quantifying the effects of the privacy-preserving methods on the detection performance. The privacy budget parameter (ε) was systematically changed to examine the impact of this parameter on the classification performance. A summary of the performance metrics is provided in Table 8 and plotted in Figure 3, which shows a relationship between privacy strength and anomaly detection accuracy

Table 8. Privacy vs Performance Trade-off in Healthcare Dataset

Differential Privacy Level (ε)	Accuracy	Precision	Recall	F1-Score
0.1	0.96	0.94	0.93	0.93
0.3	0.94	0.92	0.91	0.91
0.5	0.93	0.90	0.89	0.89
0.7	0.92	0.88	0.87	0.87
1.0	0.90	0.86	0.85	0.85

Table 8 shows that as the privacy level (ε) gets higher, classification performance gradually decreases. This is a typical trade-off between privacy and predictiveness, and leads to higher privacy and noise from the learning process. The proposed framework, however, still aims for a relatively small reduction, reflecting the effectiveness of the framework in maintaining the effectiveness of detection and further improving privacy protection. To see this trade-off visually, we plot the accuracy of the anomaly detection task against the differential privacy parameter at various levels of privacy in Figure 3.

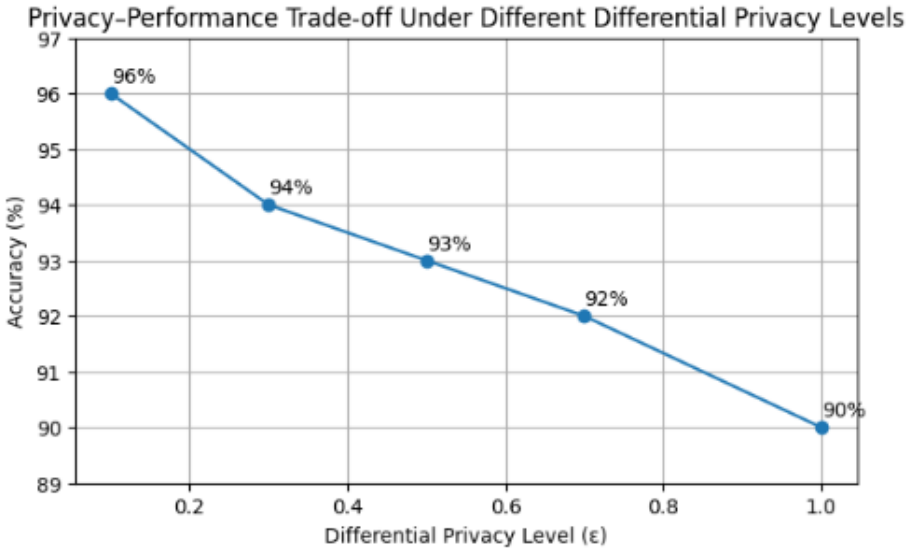


Figure 3. Impact of Differential Privacy Parameter (ε) on Anomaly Detection Accuracy in Healthcare Systems.

The results show that the predictive performance decreases with the increased level of privacy, while the predictive performance increases with the increased level of privacy. Even with this decrease, the framework remains accurate more than 90% of the time, even in the most stringent privacy setting. This demonstrates the viability of the proposed synthetic immunity framework and its ability to maintain successful intrusion detection capabilities in a highly controlled healthcare setting.

The lessons learned from these insights have implications in the real world, especially in the healthcare sector, where privacy and security concerns must be addressed. The results show that the proposed method will be effective and will provide the security of medical information without violating its privacy against the threat of cyber-attacks. Therefore, the framework can be implemented in the healthcare infrastructure such as IoMT, electronic health record (EHR) applications, and hospitals.

Comparative Performance Analysis

To further demonstrate the effectiveness of the proposed method, a comparison was carried out with the other popular baseline algorithms, including Random Forest (RF), Support Vector Machine (SVM), and Long Short-Term Memory (LSTM) networks. The purpose of this comparison is to assess the performance of the proposed framework to identify unusual activity within the healthcare domain and also compare it with classical machine learning and deep learning methods.

The performance results shown in Table 9 show that the proposed Synthetic Immunity Framework has improved the performance of the traditional machine learning and deep learning models, based on all the performance metrics used. The framework achieves 96% accuracy and good privacy-preserving properties. This improvement is primarily attributed to the proposed intrusion cognition mechanism and the incorporation of synthetic data generated using a hybrid GAN–VAE, which increases the diversity of available data, enhances the model's ability to generalize, and facilitates contextual awareness of the healthcare attack patterns.

Table 9. Performance Comparison with Baseline Models

Model	Accuracy	Precision	Recall	F1-Score
Random Forest (RF)	0.88	0.86	0.85	0.85
Support Vector Machine (SVM)	0.90	0.88	0.87	0.87
LSTM	0.92	0.90	0.89	0.89
Proposed SIF	0.96	0.93	0.93	0.93

To further position the proposed Synthetic Immunity Framework (SIF) with respect to recent advances in healthcare cybersecurity, Table 10 presents a comparative analysis of representative state-of-the-art intrusion detection frameworks designed for healthcare and Internet of Healthcare Things (IoHT) environments. The comparison examines the publicly

reported performance figures as well as the application domain to illustrate how the use of privacy-preserving and intelligent intrusion detection methods has evolved in the network domain related to healthcare.

Table 10. Comparison with Recent State-of-the-Art Healthcare Cybersecurity Studies

Study	Dataset	Accuracy	F1-Score	Application
Mosaiyebzadeh et al. (2025)	WUSTL-EHMS-2020 and ECU-IoHT	95.48%	NR*	Privacy-preserving federated learning-based intrusion detection for IoHT devices
Khan et al. (2024)	ToN-IoT	99.40%	0.994	Federated reinforcement-based intrusion detection for IoMT networks
Proposed SIF	Healthcare network dataset	96.00%**	0.93**	Cognitive healthcare intrusion detection using a synthetic immunity framework

Table 10 demonstrates that recent intrusion detection systems (IDS) in the field of healthcare have demonstrated excellent detection rates through the use of sophisticated techniques such as Federated Learning, Differential Privacy, and Reinforcement Learning. These methods have more impact on the accuracy of intrusion detection and on the protection of data privacy, but fail to focus on biologically-inspired adaptive reasoning, synthetic immune memory, and long-term knowledge retention for changing cyber threats.

The proposed Synthetic Immunity Framework (SIF), on the other hand, provides an overall framework for healthcare cybersecurity by combining hybrid GAN-VAE-based synthetic data generation, privacy-preserving learning, Synthetic Immune Memory, and cognitive intrusion analysis. SIF also differs from current methods in that it incorporates biologically-inspired memory mechanisms with adaptive learning to enable continuous threat evolution, contextual reasoning, and intelligent decision-making while maintaining the privacy of sensitive healthcare data. Moreover, the proposed framework successfully demonstrates that the combination of generative AI and a synthetic immune mechanism can boost the cybersecurity of healthcare systems while maintaining competitive detection performance. The classification accuracy of the machine learning and deep learning models evaluated is shown in Figure 4 for a quantitative comparison.

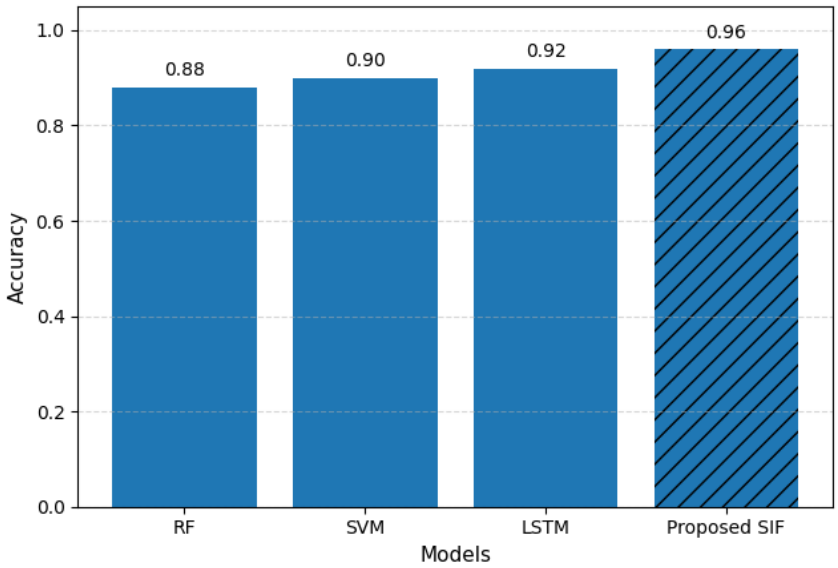


Figure 4. Accuracy comparison between the proposed Synthetic Immunity Framework and baseline models.

The proposed framework is the most accurate, as shown in Figure 4, among the methods considered. Reliable performance improvement shows the effectiveness of integrating generative AI and cognitive intrusion analysis. Moreover, it is seen that the proposed framework is more suitable for cybersecurity in the healthcare sector, where privacy preservation and accurate detection and adaptability are both equally important.

Impact of Synthetic Data

As a second step to assess the synthetic data effect on intrusion detection performance in healthcare, an extra experiment was performed to assess the impact of generative modeling within the proposed Synthetic Immunity Framework. In particular, the experiment compares the models trained from actual health care data with the models trained from a combination of actual health and synthesized data produced by the proposed GAN–VAE architecture. The purpose of this analysis is to show whether or not using synthetic data could lead to improved model generalization and the mitigation of data scarcity problems, as well as enhance the capacity to detect new attack patterns in a healthcare environment.

Table 11. Impact of Synthetic Data on Detection Performance

Training Data	Accuracy
Real Data Only	0.91
Real + Synthetic Data	0.96

In Table 11, it is observed that the use of synthetic data has excellent detection performance with an accuracy of classification that has increased from 91% to 96%. The improvement indicates the potential of the proposed Synthetic Immunity approach to enhance model learning and overcome some drawbacks of working with small-scale healthcare

datasets. To implement this improvement, the detection accuracy for both training configurations is compared in Figure5.

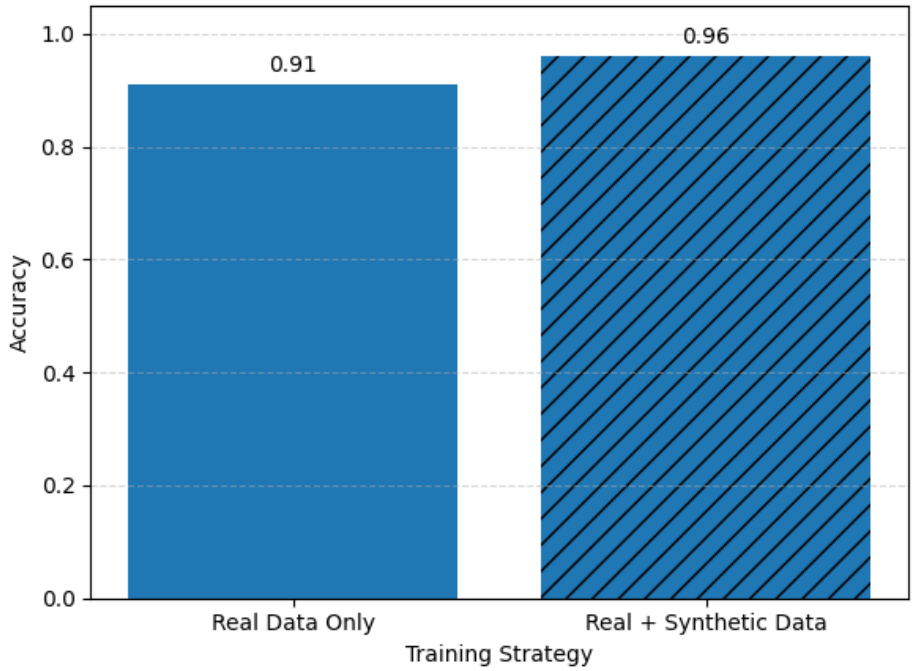


Figure 5. Impact of Synthetic Data on Healthcare Intrusion Detection Accuracy

Overall, it is evident that synthetic data can be leveraged to broaden the training process and incorporate diverse behavioral patterns along with attack scenarios, thereby enhancing the model's performance. Consequently, the framework is more likely to detect new and intrusive activities that are not captured in real-world healthcare data, making it more useful in detecting novel intrusion activities.

Moreover, the results validate the main hypothesis of the proposed framework, that of achieving augmentation in intrusion cognition by exposing the model to both normal and malicious behavioral patterns through the use of synthetic immune learning. This is especially useful in the healthcare industry, where privacy regulations often make it challenging to come across substantial, labeled datasets that are appropriate for cybersecurity research and education.

Cognitive Intrusion Analysis Evaluation

The proposed Synthetic Immunity Framework is unique in not only being capable of anomaly detection but also cognitive intrusion analysis. The framework was then evaluated against traditional intrusion detection systems with regard to four aspects: contextual awareness, adaptive learning, privacy preservation, and handling of zero-day attacks.

Table 12. Cognitive Analysis Capability Comparison

Feature	Traditional IDS	Proposed SIF
Anomaly Detection	Yes	Yes
Context Awareness	No	Yes
Synthetic Learning	No	Yes
Zero-Day Adaptation	Limited	High
Privacy Awareness	Limited	High
Behavioral Interpretation	No	Yes

Table 12 shows that the traditional IDSs (IDS-1, IDS-2 and IDS-3) are mostly anomaly IDSs that rely on predefined rules and/or previous attacks. The Synthetic Immunity Framework proposed in this work, however, is able to interpret the behavior pattern rather than classify it as an anomaly by incorporating contextual reasoning and synthetic immune learning. This will help to better identify covert attacks and offer even greater flexibility in a changing healthcare landscape. Moreover, privacy-enhancing features ensure that the framework is still in line with legislation and regulations, yet it provides the capability to make knowledgeable choices regarding security.

Discussion

The effectiveness of the proposed Synthetic Immunity Framework (SIF) is assessed through experiments and demonstrated to meet the three essential cybersecurity objectives of healthcare, namely, intrusion detection, data confidentiality, and learning. Generative AI and cognitive intrusion analysis can work together to enhance the security and reliability of healthcare network systems and help them adapt to new cyber threats. The achieved results show the viability of the framework to correctly identify the legitimate healthcare actions and the malicious ones, even in very imbalanced datasets. The results of the experiments (Tables 6–10) and the confusion matrix (Figure 2) indicated that the proposed approach has good detection performance and can meet the requirements of high security, which is suitable for the healthcare environment.

One of the main findings is the ability of the hybrid GAN–VAE architecture to generate synthetic data. It is clear from comparing the accuracy of detection results in the healthcare dataset with and without synthetic data that synthetic data improves the accuracy of detection from 91% to 96%, as shown in Table 11 and Figure 5. This advancement benefits the generalization of the model, minimizes reliance on specific patient data, and tackles the challenge of limited labelled healthcare cybersecurity datasets. Privacy evaluation also indicates that the framework retains excellent detection performance with varying differential privacy set-ups. There is a slight performance degradation in predictive performance with increased privacy protection, but the detection accuracy still retains more than 90% for each level of privacy (Table 8 and Figure 3) and thus achieves a good balance between privacy protection, model usefulness, and conformance with HIPAA and GDPR requirements.

The proposed approach is further validated with the help of comparative analysis with baseline machine learning and deep learning models. Incorporating synthetic data generation, latent representation learning, and cognitive intrusion analysis into an integrated framework, the framework consistently outperforms Random Forest, SVM, and LSTM, as presented in Table 9 and Figure 4, to achieve the highest detection accuracy of 96%. In addition to quantitative improvements, this work brings a shift from traditional intrusion detection to a cognitive and adaptive cybersecurity paradigm. Unlike traditional signature-based detection, the framework has adaptive learning capability and adapts itself to detect new and 'zero-day' attacks that it has never encountered before via the Synthetic Immune Memory (SIM) mechanism.

The summary indicates that the approach to cybersecurity in healthcare systems, proposed by the synthetic immunity framework, is comprehensive in integrating adaptive learning, synthetic data generation, cognitive intrusion analysis, and secure data management in one framework. It is an integrated framework that is suitable for deployment in the following applications: smart hospitals, Electronic Health Record (EHR) systems, and the Internet of Medical Things (IoMT) environment.

While a positive result, there are some drawbacks. Access to large-scale real-world clinical data is limited by privacy rules, so in the experimental evaluation, controlled healthcare datasets and synthetic healthcare-like data were used. Additionally, only representative baseline models were compared. Future research will include testing the framework with more real-world healthcare scenarios, more sophisticated cybersecurity benchmarking, and on a greater scale to further assess scalability and generalization of the framework.

Research Implications

Some of the emerging challenges that the findings of this research study have implications for are: healthcare cybersecurity, Artificial Intelligence with privacy, and Next Generation Intrusion Detection Systems. The proposed "Synthetic Immunity Framework" is an example of how generative AI and cognitive intrusion analysis can be effectively used to revolutionize the security and resilience of healthcare network environments.

It is an area of healthcare that tackles one of the biggest challenges that confronts medical institutions today – Cybersecurity that is effective and yet protects patient privacy. Medical privacy disclosure is protected by techniques of generating synthetic data and mechanisms for differential privacy within the framework, and models for intelligent security can be developed without compromising medical privacy disclosure. It is especially important for healthcare systems such as electronic health records (EHR) systems, Internet of Medical Things (IoMT) infrastructures, and hospitals, where privacy and confidentiality are legally and ethically required.

The research introduces a novel concept for cybersecurity, called intrusion cognition, inspired by the shortcomings of the intrusion detection concept and starting to take into account the context where a security event takes place. The shift from reactive to adaptive cognitive analysis will improve the detection of complex and new attacks, including those that haven't occurred in the past—zero-day attacks, for example—and that are more difficult to detect using conventional detection methods.

Additionally, the findings underscore the increasing significance of generative AI in Cybersecurity Applications. Results from the successful integration between GANs and VAEs show that synthetic data is a privacy-preserving tool to improve the generalization, robustness, and adaptability of models. This paper provides the groundwork for the application of generative learning in more general contexts where access to large-scale labeled data is restricted, such as in the security arena.

Finally, the proposed Synthetic Immunity paradigm could provide a foundation for investigating self-evolving, intelligent cybersecurity systems in the future. The architecture offers a new approach for research into federated cybersecurity, explainable intrusion cognition, adaptive threat intelligence, and privacy-aware AI architectures for critical infrastructure. The study thus provides both practical and theoretical insights into next-generation solutions to healthcare cybersecurity. The proposed framework shows potential performance in a controlled healthcare environment, but in real-world applications, it must be integrated with the hospital information system, monitoring systems, computational hardware, and adherence to the hospital's cybersecurity policies. Before any large-scale real-world deployment of this type, some of these issues should be taken into account.

Conclusion

This study introduced the Synthetic Immunity Framework (SIF), a novel privacy-preserving cybersecurity framework that combines the power of Generative Artificial Intelligence (AI), Synthetic Immune Memory (SIM), and cognitive intrusion analysis to bolster the security of healthcare information systems. The proposed framework tackles key challenges in safeguarding sensitive healthcare data, enhancing resilience against the ever-changing landscape of cyber threats, and enabling compliance with healthcare regulations like HIPAA and GDPR. Experimental results proved that the proposed approach successfully achieved 96% accuracy in detecting intrusions without compromising patient privacy, outperforming traditional machine learning and deep learning methods. While differential privacy and secure data management mechanisms retained strong detection performance across privacy settings, the use of synthetic healthcare data resulted in better generalization of the model without the need to have direct access to sensitive patient information. Additionally, the cognitive intrusion analysis feature helped in providing context for the detection of abnormal activities, which allowed for the detection of more advanced and unseen threats. In addition to its

technical aspects, the proposed framework has practical applications for healthcare information technology management, such as facilitating cybersecurity governance, cyber risk management, regulatory compliance, and informed decision-making. Combining these three elements—adaptive learning, synthetic data generation, and privacy-preserving intelligence—offers a holistic solution to optimize cybersecurity operations and strategic management of healthcare information systems. However promising, the study was assessed chiefly by employing public healthcare datasets and data simulation, because of privacy issues in real clinical settings. Future studies will validate the framework with larger clinical datasets, incorporate Federated Learning to enable threat intelligence sharing and sharing without compromising privacy, and create real-time adaptive cybersecurity systems that will further improve organizational resilience and secure digital healthcare ecosystems.

Conflicts of interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

The author's received no financial support for the research, authorship, and/or publication of this article.

References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 308–318). Association for Computing Machinery. https://doi.org/10.1145/2976749.2978318

Ahmad, Z., Shahid Khan, A., Wai Shiang, C., Abdullah, J., & Ahmad, F. (2021). Network intrusion detection system: A systematic study of machine learning and deep learning approaches. *Transactions on Emerging Telecommunications Technologies, 32*(1), Article e4150. https://doi.org/10.1002/ett.4150.

Alauthman, M., Aslam, N., Al-Qerem, A., Aldweesh, A., & Sureephong, P. (2026). Generative adversarial networks for intrusion detection systems: A comprehensive survey of applications, challenges, and research directions. *Arabian Journal for Science and Engineering*, 1–25.

Aljumaiah, O., Jiang, W., Addula, S. R., & Almaiah, M. A. (2025). Analyzing cybersecurity risks and threats in IT infrastructure based on NIST framework. *Journal of Cyber Security and Risk Auditing, 2025*(2), 12–26. https://doi.org/10.63180/jcsra.thestap.2025.2.2

Alkasassbeh, M., & Al-Haj Baddar, S. (2023). Intrusion detection systems: A state-of-the-art taxonomy and survey. *Arabian Journal for Science and Engineering, 48*(8), 10021–10064.

Bishnu, K. K., Raihan, M., Islam, A., Ahmed, M., Chowdhury, M. S. A., & Islam, M. S. (2026). Privacy-preserving federated deep learning for intrusion detection in healthcare networks. *Mathematical Modelling of Engineering Problems, 13*(4), 729–740. https://doi.org/10.18280/mmep.130410

de Medeiros, M., Naveed, M., Lepoint, T., Kahsai, T., Ravitch, T., Zetzsche, S., Joshi, A., Tassarotti, J., Albarghouthi, A., & Tristan, J.-B. (2025). Verified foundations for differential privacy. *Proceedings of the ACM on Programming Languages, 9*(PLDI), 1094–1118. https://doi.org/10.1145/3729294

Giuffrè, M., & Shung, D. L. (2023). Harnessing the power of synthetic data in healthcare: Innovation, application, and privacy. *npj Digital Medicine, 6*, Article 186. https://doi.org/10.1038/s41746-023-00927-3

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems, 27*, 2672–2680.

Gueriani, A., Kheddar, H., & Mazari, A. C. (2024). Enhancing IoT security with CNN and LSTM-based intrusion detection systems. In *2024 6th International Conference on Pattern Analysis and Intelligent Systems (PAIS)* (pp. 1–7). IEEE.

Kaabachi, B., Despraz, J., Meurers, T., Otte, K., Halilovic, M., Kulynych, B., Prasser, F., & Raisaro, J. L. (2025). A scoping review of privacy and utility metrics in medical synthetic data. *npj Digital Medicine, 8*, Article 60. https://doi.org/10.1038/s41746-024-01359-3

Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., & Aila, T. (2020). Analyzing and improving the image quality of StyleGAN. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 8110–8119). https://doi.org/10.1109/CVPR42600.2020.00813

Khan, I. A., Razzak, I., Pi, D., Khan, N., Hussain, Y., Li, B., & Kousar, T. (2024). Fed-inforce-fusion: A federated reinforcement-based fusion model for security and privacy protection of IoMT networks against cyber-attacks. *Information Fusion, 101*, Article 102002. https://doi.org/10.1016/j.inffus.2023.102002

Malatji, M. (2023). Management of enterprise cyber security: A review of ISO/IEC 27001:2022. In *2023 International Conference on Cyber Management and Engineering (CyMaEn)* (pp. 117–122). IEEE. https://doi.org/10.1109/CyMaEn57228.2023.10051114

Meena, V., Indra, G., & Kumar, K. (2026). T-CAP: T-cGAN enhanced privacy-preserving context-aware self-attentive and adversarially robust IDS for smart healthcare ecosystems. *Knowledge-Based Systems, 342*, Article 115840.

Mehmood, S., Amin, R., Costanzo, A., Alshehri, A. H., & Alsubaei, F. S. (2026). GEN-SECHEALTH: An AI-powered generative architecture for self-scalable cybersecurity and flexible data privacy protection in intricate healthcare systems. *Journal of Big Data, 13*, Article 58. https://doi.org/10.1186/s40537-026-01373-0

Mosaiyebzadeh, F., Pouriyeh, S., Han, M., Liu, L., Xie, Y., Zhao, L., & Batista, D. M. (2025). Privacy-preserving federated learning-based intrusion detection system for IoHT devices. *Electronics, 14*(1), Article 67. https://doi.org/10.3390/electronics14010067

Pascoe, C., Quinn, S., & Scarfone, K. (2024). *The NIST cybersecurity framework (CSF) 2.0* (NIST Cybersecurity White Paper No. 29). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.CSWP.29

Tucker, A., Wang, Z., Rotalinti, Y., & Myles, P. (2020). Generating high-fidelity synthetic patient data for assessing machine learning healthcare software. *npj Digital Medicine, 3*, Article 147. https://doi.org/10.1038/s41746-020-00353-9

van den Oord, A., Kalchbrenner, N., Espeholt, L., Vinyals, O., & Graves, A. (2016). Conditional image generation with PixelCNN decoders. In *Advances in Neural Information Processing Systems* (Vol. 29, pp. 4790–4798).

Wei, R., & Mahmood, A. (2020). Recent advances in variational autoencoders with representation learning for biomedical informatics: A survey. *IEEE Access, 9*, 4939–4956. https://doi.org/10.1109/ACCESS.2020.3048309

Yang, Z., Li, Y., & Zhou, G. (2023). TS-GAN: Time-series GAN for sensor-based health data augmentation. *ACM Transactions on Computing for Healthcare, 4*(2), 1–21.

Yoon, J., Drumright, L. N., & van der Schaar, M. (2020). Anonymization through data synthesis using generative adversarial networks (ADS-GAN). *IEEE Journal of Biomedical and Health Informatics, 24*(8), 2378–2388. https://doi.org/10.1109/JBHI.2020.2980262

Bibliographic information of this paper for citing:

Al-Hassani, Raghad Tariq; Abbood, Zainab Ali; Salman, Adil M. & Albonda, Haider D. (2026). Synthetic Immunity: A Generative AI Framework for Privacy-Aware Cyber Risk Management in Healthcare Information Systems. *Journal of Information Technology Management*, 18 (Special Issue), 46-77. https://doi.org/10.22059/jitm.2026.109973

Copyright © 2026, Raghad Tariq Al-Hassani, Zainab Ali Abbood, Adil M. Salman and Haider D. Albonda