IJIRSET
International Journal of Innovative Research in Science | Engineering | Technology

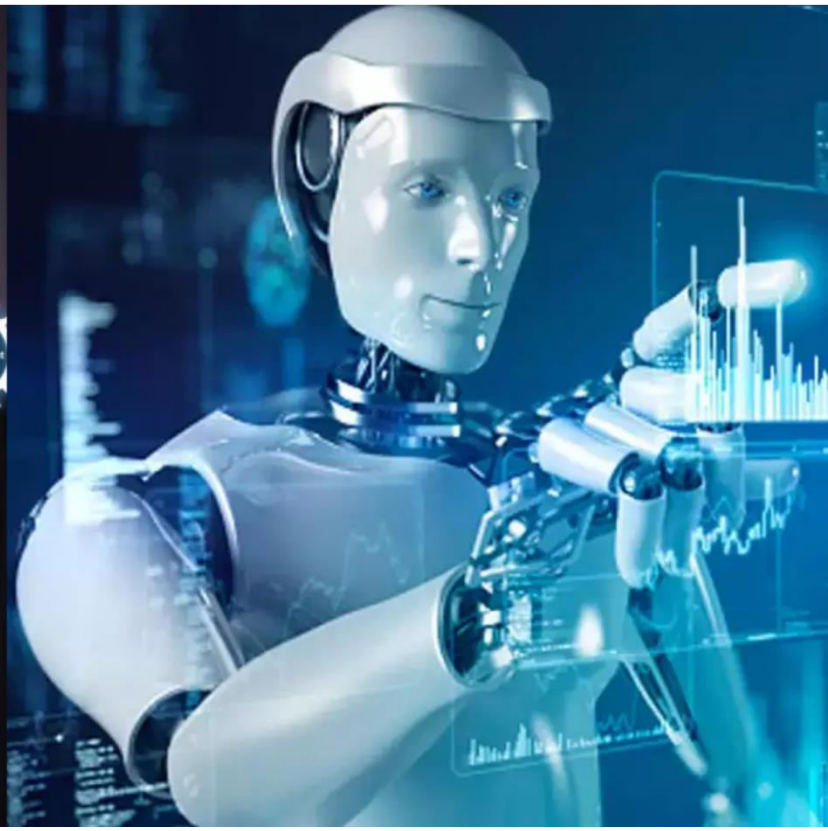
ISSN(O): 2319-8753
ISSN(P): 2347-6710



# International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

*(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)*



**Impact Factor: 9.935**

**Volume 15, Issue 10, October 2026**

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510089|

# AI-Powered Smart Security Framework for Unauthorized Access Detection and Data Privacy Preservation in Healthcare Systems

**Morampudi Divya Sri, Dr. S. Suhasini**

Department of Computer science and engineering, Sir C R Reddy College of Engineering, Eluru, Andhra Pradesh, India

Department of Artificial Intelligence, VR Siddhartha Engineering College, Kanuru, Andhra Pradesh, India

**ABSTRACT:** Electronic health records (EHRs) are high-value targets for external attackers and for insiders who misuse legitimate privileges. Existing detectors either ignore the privacy of the logs used for training, or protect privacy without detecting access misuse. This paper proposes a framework that (i) detects unauthorized access with an autoencoder trained only on normal access sessions, (ii) trains the detector across hospitals with federated averaging and client-level differential privacy so that raw logs never leave the hospital, and (iii) records alerts in a tamper-evident, hash-chained audit ledger. On simulated access-session data from 8 hospitals (5 runs), the private federated detector reached an F1-score of 0.851 ± 0.015 and an AUC of 0.955 at noise multiplier $\sigma = 1$ ($\varepsilon \approx 50$ at $\delta = 10^{-5}$ under the stated assumptions), compared with 0.896 for a centralized autoencoder trained on pooled logs and 0.872 for Isolation Forest. Detection stayed within the run-to-run variation of the clipped, noise-free federated model up to $\sigma = 2$ ($\varepsilon \approx 20$) and degraded sharply at $\sigma = 4$ ($\varepsilon \approx 8.8$, F1 0.659). The ledger detected every tampering attempt in our tests. The results are limited to simulated data.

**KEYWORDS:** Healthcare security; unauthorized access detection; anomaly detection; federated learning; differential privacy; BiLSTM autoencoder; blockchain; electronic health records.

## I. INTRODUCTION

The digitization of healthcare has improved care quality and availability, but it has also exposed patient data to cyber threats such as credential theft, ransomware, insider snooping and attacks on connected medical devices. Health records are especially sensitive because, unlike payment data, they cannot be reissued once leaked. Regulations such as HIPAA and GDPR require institutions to protect patient information and to keep auditable records of who accessed what.
Traditional safeguards such as role-based access control (RBAC) decide who may access a record, but they cannot tell whether a permitted access is appropriate. A nurse who browses the records of a neighbor, or a stolen credential used at an unusual hour, passes RBAC checks. Machine learning on access logs can catch such behavior, but two problems remain. First, training a detector normally requires pooling sensitive audit logs from several hospitals in one place, which itself creates a privacy risk. Second, detection alerts are usually stored in ordinary databases that an insider with administrative rights could alter.

This paper addresses both problems in one framework. The main contributions are:

1. A session-level autoencoder detector for unauthorized EHR access that needs no labeled attack data.
2. A privacy-preserving training scheme that combines federated averaging with client-level differential privacy (clipped, noised updates with Rényi accounting), so raw logs stay local.
3. A tamper-evident, hash-chained audit layer for alerts, with a tamper-detection test.
4. An evaluation on simulated multi-hospital data covering detection quality, the privacy-utility trade-off and overhead, with the code provided.

The rest of the paper is organized as follows. Section II reviews related work, Section III presents the framework, Section IV describes the experimental setup, Section V reports results, Section VI analyzes security and privacy, and Section VII concludes.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510089|

## II. RELATED WORK

### A. AI-based unauthorized access and insider threat detection

Early work on EHR access monitoring used outlier detection on audit logs. Boddy et al. proposed a hybrid density-based outlier detection model for detecting unauthorized access in hospital networks through real-time audit-log analysis [1]. Yeng et al. compared several machine learning classifiers on simulated EHR logs, motivated by "break the glass" and self-authorized access scenarios [2]. Hurst et al. applied supervised learning to EHR access data and reported accuracy of about 0.99 with a support vector machine [3]. Tabassum et al. studied contextual anomalies in smart health and found that an Isolation Forest combined with an SVM performed best [4]. More recent studies explore hybrid meta-classifiers for access logs [5] and graph and Transformer models for national EHR audit logs [6]. A common limitation is that these detectors are trained centrally on pooled logs and treat privacy of the training data as out of scope.

### B. Privacy-preserving learning in healthcare

Federated learning (FL) trains a shared model without moving raw data [7]. Surveys of FL in healthcare identify data heterogeneity as a major challenge and highlight differential privacy, secure multiparty computation and homomorphic encryption as the main protection mechanisms [8]. Rahman et al. review the combination of FL and differential privacy for decentralized healthcare and discuss open research questions [9]. A systematic review of 80 studies reports a shift toward hybrid privacy-preserving FL architectures that combine several mechanisms [10]. Most of this work focuses on diagnosis and imaging rather than on access misuse.

### C. Integrated frameworks with intrusion detection and blockchain

For the Internet of Medical Things, Al-Hawawreh and Hossain proposed a privacy-aware attack detection framework using a differentially private contractive deep autoencoder [11]. Bezanjani et al. combined blockchain-enabled request encryption with BiLSTM-based request pattern recognition to detect unauthorized access attempts [12]. A recent framework integrates blockchain, homomorphic encryption, federated learning and an AI-based IDS for healthcare IoT [13], and a 2026 review of AI-powered IoMT intrusion detection assesses detection performance, explainability and privacy preservation together [14].

### D. Research gap

Existing access-log detectors lack built-in privacy preservation, while privacy-preserving FL studies rarely target access misuse. Few designs unify detection, privacy and auditability, and many integrated designs are resource intensive. This work targets that gap with a lightweight, unified framework.

## III. PROPOSED FRAMEWORK

The framework detects unauthorized access locally, trains the detector across hospitals without sharing raw logs, and records every alert in a tamper-evident ledger.

### A. Architecture overview

The framework has three layers deployed across N participating hospitals and one aggregation server:

1. **Local detection layer (at each hospital).** Access logs from the EHR system are summarized into session features and used to train a local autoencoder. At run time the model scores each session.
2. **Privacy-preserving collaboration layer.** Hospitals send only clipped model updates to the aggregation server. The server adds Gaussian noise to the aggregate, performs federated averaging and returns the global model.
3. **Audit layer.** Every alert (session summary, anomaly score, model version, timestamp) is hashed and appended to an append-only ledger shared by the hospitals.

### B. Session features

Raw access events (user ID, role, patient ID, action type, timestamp, location) are grouped into one record per user session and summarized into 15 features scaled to [0, 1]: event count, distinct patients, fraction of accesses outside the user's department, off-hours fraction, export or print fraction, mean inter-event interval, distinct locations, failed logins, fraction of sensitive-record accesses, fraction of VIP-record accesses, weekend fraction, and a one-hot role (4 values). Fixed scaling constants are used so that no hospital has to share statistics. Direct identifiers are replaced by pseudonymous tokens before training.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510089|

**C. Autoencoder for anomaly detection**

The detector is a dense autoencoder with layers 15–16–6–16–15, tanh hidden units and a linear output layer (725 parameters). It is trained only on normal sessions to minimize the mean squared reconstruction error of a session x with d = 15 features, given in Eq. (1):

$$L(x) = (1/d)\ \Sigma_{j=1}^{d}\ (x_j - \hat{x}_j)^2 \ \ldots\ldots..(1)$$

The anomaly score of a session is its reconstruction error, s(x) = L(x). Each hospital k sets its own threshold $\tau_k$ as the 99th percentile ($P_{99}$) of the scores on its normal validation sessions $V_k$, as in Eq. (2):

$$\tau_k = P_{99}(\ \{\ s(x) : x \in V_k\ \}\ )\ \ldots\ldots\ldots(2)$$

A session is flagged when $s(x) > \tau_k$, so the expected false-positive rate is about 1%. Because the model learns normal behavior, it can flag attack patterns it has never seen, which supervised classifiers may miss. When event-level sequences are available, a sequence model such as a BiLSTM autoencoder can replace the dense autoencoder; we leave this for future work.

**D. Federated training with differential privacy**

Let hospital *k* hold local normal sessions $D_k$. In each communication round *r*, the server sends the global weights $w_r$ to all *N* hospitals. Hospital *k* trains locally for E epochs and computes its update, Eq. (3):

$$\Delta_k = w_k - w_r\ \ldots\ldots\ldots(3)$$

The update is clipped to norm C, as in Eq. (4):

$$\Delta_k \leftarrow \Delta_k \cdot \min(\ 1,\ C / \|\Delta_k\|_2\ )\ \ldots\ldots\ldots(4)$$

The server sums the clipped updates, adds Gaussian noise $\xi_r$ and averages, as in Eq. (5):

$$w_{r+1} = w_r + (1/N)\ (\ \Sigma_{k=1}^{N}\ \Delta_k + \xi_r\ ),\ \xi_r \sim N(0,\ \sigma^2 C^2 I)\ldots\ldots\ldots(5)$$

The noise is added to the aggregate, which assumes secure aggregation or a trusted aggregator. Privacy is defined at the hospital level: the guarantee covers the addition or removal of one hospital's data. Because every hospital participates in every round, each round is a Gaussian mechanism with sensitivity C and noise scale σC, whose Rényi divergence of order α is $\alpha/(2\sigma^2)$ [17]. After T rounds, the Rényi privacy loss ε(α) is given by Eq. (6):

$$\varepsilon(\alpha) = T\alpha\ /\ (2\sigma^2)\ldots\ldots(6)$$

It is converted to an (ε, δ) guarantee by Eq. (7), minimizing over $\alpha > 1$:

$$\varepsilon = \min\ [\ \varepsilon(\alpha) + \ln(1/\delta)\ /\ (\alpha - 1)\ ]\ \ldots\ldots\ldots\ldots(7)$$

**E. Audit layer**

When a session is flagged, the hospital appends a record (hash of the session summary, anomaly score, model version, timestamp) to an append-only ledger in which each block's hash covers the previous block's hash. Changing or removing an earlier record breaks the chain and is detected on verification. Only hashes and metadata are stored, never patient data. Our prototype is a local SHA-256 hash chain; a production deployment would run on a permissioned blockchain with several hospitals as validating nodes.

**F. Algorithm summary**

**Algorithm 1: Federated private autoencoder training and detection**

1. Initialize global model $w_0$, clipping norm C, noise multiplier σ and number of rounds T.
2. For r = 0 to T − 1: the server sends $w_r$ to all hospitals.
3. Each hospital trains the autoencoder locally on normal sessions, clips its update to norm C and returns it.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510089|

4. The server adds noise to the summed updates, averages them to obtain $w_{\{r+1\}}$, and the privacy accountant updates ε.
5. After training, each hospital computes its threshold $\tau_k$ on local validation data.
6. At run time, for each session compute the reconstruction error; if it exceeds $\tau_k$, raise an alert and append its hash to the ledger.

## IV. EXPERIMENTAL SETUP

All experiments use simulated multi-hospital data and a NumPy and scikit-learn implementation, with 5 independent runs.

### A. Data

No real hospital audit logs were available, so access sessions were simulated. The generator creates one record per user session for four roles (doctor, nurse, administrator, billing), with role-specific distributions (beta, lognormal and Poisson models) for the 15 features. Each of N = 8 hospitals has its own role mix, activity volume (0.8–1.3×) and off-hours baseline, which makes the data non-IID. Each hospital has 2,000 normal sessions for training, 500 normal sessions for threshold validation, and a test set of 900 normal and 100 anomalous sessions, giving 8,000 test sessions per run with 10% anomalies.

Four anomaly types are injected: snooping on out-of-department, VIP and sensitive records (30%); stolen credentials with off-hours, multi-location and failed-login signatures (30%); bulk export (25%); and subtle behavioral drift (15%). All detectors are trained on normal sessions only, except the supervised reference, which also sees labeled anomalies (10% of its training set). Because the generator is our own, absolute scores say little about real hospitals; the comparisons between methods on identical data are the meaningful part.

### B. Baselines

- Isolation Forest (200 trees) [18]
- One-Class SVM (RBF kernel, $\nu = 0.05$, trained on 3,000 sessions) [19]
- Random Forest (200 trees, trained with labels; an upper reference, not a like-for-like competitor) [20]
- Centralized autoencoder trained on pooled normal sessions (no privacy)
- Federated autoencoder with plain averaging (no clipping, no noise)
- Federated autoencoder with clipping only (no noise)

### C. Metrics

Precision, recall, F1-score, ROC-AUC and false-positive rate (FPR) for detection, computed on all test sessions with per-hospital thresholds. For AUC, each score is divided by its hospital's threshold so that scores are comparable across hospitals. For privacy: ε at $\delta = 10^{-5}$. For efficiency: model size, time per round, inference time and ledger append time.

### D. Implementation details

Autoencoder 15–16–6–16–15 with tanh hidden units; plain SGD with learning rate 0.05 and batch size 64; T = 40 rounds with E = 2 local epochs; all 8 hospitals take part in every round. Clipping norm C = 0.05, chosen near the median update norm (0.057) observed in an unclipped run, so clipping binds on a large share of updates; this data-dependent choice is not covered by the reported ε. Noise multipliers $\sigma \in \{0.5, 1, 2, 4\}$, $\delta = 10^{-5}$. Thresholds are the 99th percentile of each hospital's validation errors. The centralized autoencoder is trained for 40 epochs on pooled data. Seeds 0–4; results are mean ± standard deviation over the 5 runs. Code: run_experiments.py (single CPU core).

## V. RESULTS AND DISCUSSION

The private federated detector gives up about 0.045 F1 against centralized training at $\sigma = 1$ ($\varepsilon \approx 50$) and holds its accuracy up to $\sigma = 2$ ($\varepsilon \approx 20$), but it does not beat the non-private baselines. All values are mean ± standard deviation over 5 runs on simulated data.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510089|

### A. Detection performance

Table 1 compares the detectors on 8,000 test sessions per run, 10% of them anomalous.

| Model | Precision | Recall | F1 | AUC | FPR |
|---|---|---|---|---|---|
| Isolation Forest | 0.895 ± 0.020 | 0.850 ± 0.031 | 0.872 ± 0.021 | 0.992 ± 0.002 | 0.011 ± 0.002 |
| One-Class SVM | 0.897 ± 0.012 | 0.885 ± 0.001 | 0.891 ± 0.006 | 0.978 ± 0.002 | 0.011 ± 0.001 |
| Centralized autoencoder | 0.904 ± 0.008 | 0.888 ± 0.015 | 0.896 ± 0.011 | 0.988 ± 0.004 | 0.011 ± 0.001 |
| Federated autoencoder (plain averaging) | 0.892 ± 0.010 | 0.888 ± 0.009 | 0.890 ± 0.010 | 0.986 ± 0.004 | 0.012 ± 0.001 |
| **Federated autoencoder + DP ($\sigma = 1$)** | 0.873 ± 0.012 | 0.831 ± 0.022 | 0.851 ± 0.015 | 0.955 ± 0.006 | 0.013 ± 0.001 |
| Random Forest (supervised, uses labels) | 0.998 ± 0.001 | 0.992 ± 0.001 | 0.995 ± 0.001 | 1.000 ± 0.000 | 0.000 ± 0.000 |

The supervised Random Forest is best, but it needs labeled attacks, which hospitals rarely have. Among detectors trained on normal data only, the centralized autoencoder (F1 0.896), one-class SVM (0.891) and plain federated autoencoder (0.890) are similar. Isolation Forest has a lower F1 (0.872) but the highest AUC of the unsupervised methods (0.992). The private federated detector trails the centralized autoencoder by about 0.045 F1. FPR is 1.1–1.3% for every unsupervised method, as expected from the 99th-percentile thresholds.

### B. Privacy-utility trade-off

Table 2 shows the effect of clipping and noise, with $\varepsilon$ computed for $T = 40$ rounds at $\delta = 10^{-5}$.

| Setting | $\sigma$ | $\varepsilon$ | F1 | Recall | AUC |
|---|---|---|---|---|---|
| Plain averaging (no clipping, no noise) | 0 | $\infty$ | 0.890 ± 0.010 | 0.888 ± 0.009 | 0.986 ± 0.004 |
| Clipping only | 0 | $\infty$ | 0.861 ± 0.015 | 0.842 ± 0.031 | 0.959 ± 0.008 |
| Clipping + noise | 0.5 | 140.7 | 0.864 ± 0.014 | 0.845 ± 0.028 | 0.958 ± 0.007 |
| Clipping + noise | 1 | 50.4 | 0.851 ± 0.015 | 0.831 ± 0.022 | 0.955 ± 0.006 |
| Clipping + noise | 2 | 20.2 | 0.847 ± 0.023 | 0.818 ± 0.029 | 0.952 ± 0.008 |
| Clipping + noise | 4 | 8.8 | 0.659 ± 0.076 | 0.552 ± 0.098 | 0.913 ± 0.018 |

Clipping at $C = 0.05$ lowers F1 from 0.890 to 0.861, because it caps the size of every hospital's update. Adding noise up to $\sigma = 2$ ($\varepsilon \approx 20$) changes F1 by less than one standard deviation (0.847 against 0.861). At $\sigma = 4$ ($\varepsilon \approx 8.8$) recall falls to 0.552 and F1 to 0.659. The $\varepsilon$ values are large: they bound the privacy loss of a whole hospital with only 8 participants and 40 rounds, so they describe modest protection, not a strong guarantee. More hospitals, fewer rounds or a better clipping-to-noise ratio should give a better trade-off and need to be tested.

### C. Overhead

| Component | Measure | Value |
|---|---|---|
| Model size | parameters; per hospital, per direction, per round | 725; about 2.9 KB (float32) |
| Federated round (8 hospitals, run one after another on one CPU core) | seconds | about 0.05 |
| Centralized training (40 epochs) | seconds | 0.92 ± 0.03 |
| Inference | ms per session | 0.032 |
| Ledger append (local hash chain, no consensus) | ms per alert | 0.0008 |
| Ledger tamper test | tampered chains detected | 5 of 5 runs |

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510089|

The model is small enough that communication and computation are negligible at this scale. The ledger figure is for a local hash chain without a consensus protocol, so it says nothing about the latency of a real permissioned blockchain.

**D. Ablation**

Table 2 serves as the ablation. Removing noise shows the cost of the privacy mechanism, removing clipping as well shows the cost of bounding updates, and the centralized autoencoder in Table 1 shows the cost of keeping data local. The audit layer does not affect detection, so it was not ablated.

## VI. SECURITY AND PRIVACY ANALYSIS

**Privacy of raw data.** Raw logs never leave the hospital. Only clipped updates, with noise added at the aggregator, are shared. The formal guarantee is a hospital-level ($\varepsilon$, $\delta$) bound that holds under three assumptions: secure aggregation or a trusted aggregator, participation of all hospitals in every round, and add-or-remove adjacency of one hospital. We did not test empirical attacks such as gradient inversion or membership inference.

**Insider tampering.** In all 5 runs, modifying one record in a 2,000-record ledger was detected on verification. A single hash chain held by one operator does not stop that operator from rewriting the whole chain, so a deployment needs several hospitals as validating nodes of a permissioned blockchain.

**Model poisoning.** Clipping bounds the influence of any single hospital's update, but we did not evaluate poisoning attacks. Robust aggregation (for example median or trimmed mean) is future work.

**Limitations.** The data is simulated by our own generator, so the results show relative behavior of the methods and not performance on real hospitals. Session features discard event order. The thresholds are fixed 99th percentiles. Only 8 hospitals were used, which keeps $\varepsilon$ high. The clipping norm was set from a non-private run and that choice is not accounted for in $\varepsilon$. Results are not broken down by anomaly type. A sequence model (BiLSTM autoencoder) and a real blockchain deployment were not evaluated. Reconstruction error alone does not explain an alert, so feature-level error attribution should be added for analysts.

## VII. CONCLUSION AND FUTURE WORK

This paper presented a framework that detects unauthorized EHR access with an autoencoder trained on normal sessions, trains it across hospitals with federated averaging and client-level differential privacy, and logs alerts in a hash-chained ledger. On simulated data from 8 hospitals, the private federated detector reached an F1-score of $0.851 \pm 0.015$ at $\sigma = 1$ ($\varepsilon \approx 50$), compared with 0.896 for centralized training and 0.872 for Isolation Forest. Detection held up to $\sigma = 2$ ($\varepsilon \approx 20$) and fell at $\sigma = 4$ ($\varepsilon \approx 8.8$, F1 0.659). The cost of keeping data local is therefore modest here, but the formal $\varepsilon$ values remain large. Future work includes validation on real audit logs, sequence models such as a BiLSTM autoencoder, a breakdown by anomaly type, a better privacy-utility trade-off with more hospitals, robust aggregation against poisoned updates, and deployment on a real permissioned blockchain.

## REFERENCES

[1] A. Boddy et al., "A hybrid density-based outlier detection model for privacy in electronic patient record system," conference paper, 2019.
[2] P. Yeng et al., "Comparative analysis of machine learning methods for analyzing security practice in electronic health records' logs," in Proc. IEEE Int. Conf. Big Data, 2020.
[3] W. Hurst, B. Tekinerdogan, T. Alskaif, A. Boddy, and N. Shone, "Securing electronic health records against insider-threats: A supervised machine learning approach," Smart Health, vol. 26, p. 100354, 2022.
[4] M. Tabassum, S. Mahmood, A. Bukhari, B. Alshemaimri, and A. Daud, "Anomaly-based threat detection in smart health using machine learning," 2024.
[5] "Energy efficient meta classifier model for log anomaly detection in healthcare," Int. J. Intelligent Computing and Cybernetics, vol. 19, no. 1. [Add authors.]
[6] M. Meneghello, "Security of the Italian Electronic Health Record: a multi-modal anomaly detection model," thesis, Univ. of Padova, 2025.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510089|

[7] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, 2017.
[8] C. Zhang and W. Li, "Survey of federated learning in healthcare: Applications, challenges, and future directions," Chinese J. Engineering, vol. 47, no. 9, pp. 1825–1840, 2025.
[9] M. M. Rahman et al., "Balancing privacy and performance: Federated learning with differential privacy for real-time, resilient healthcare AI," 2025.
[10] A. Alankamony and N. Nels, "Systematic review of privacy preservation in federated learning for secured healthcare applications," Information (MDPI), doi:10.3390/info17070647.
[11] M. Al-Hawawreh and M. S. Hossain, privacy-aware attack detection framework for IoMT using a differentially private contractive deep autoencoder. [Complete citation from the source paper.]
[12] B. Rezaei Bezanjani, S. H. Ghafouri, and R. Gholamrezaei, "Fusion of machine learning and blockchain-based privacy-preserving approach for health care data in the Internet of Things," arXiv:2510.19026.
[13] "Enhancing data security in healthcare IoT systems using emerging technologies," Genetics and Molecular Research. [Add authors and year.]
[14] J. Islam, D. Datta, and F. Akhter, "Trustworthy AI-powered intrusion detection for the Internet of Medical Things (IoMT): A review," Sensors, 2026, doi:10.3390/s26165182.
[15] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," Foundations and Trends in Theoretical Computer Science, vol. 9, no. 3–4, 2014.
[16] M. Abadi et al., "Deep learning with differential privacy," in Proc. ACM CCS, 2016.
[17] I. Mironov, "Rényi differential privacy," in Proc. IEEE Computer Security Foundations Symposium (CSF), 2017.
[18] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in Proc. IEEE ICDM, 2008.
[19] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson, "Estimating the support of a high-dimensional distribution," Neural Computation, vol. 13, no. 7, 2001.
[20] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, 2001.



INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INNO SPACE SJIF Scientific Journal Impact Factor

doi cross ref

निस्केयर NISCAIR

# INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

## IN SCIENCE | ENGINEERING | TECHNOLOGY

9940 572 462 6381 907 438 ijirset@gmail.com

www.ijirset.com

Scan to save the contact details