International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 6, Nov-Dec-2026, ISSN (Online): 2395-566X

Healthcare Monitoring Using Artificial Intelligence, Federated Learning and Blockchain Technology

Suryawanshi Mohit Pravin, Prof. Sakshi S.S., Prof. Nalawade T.S., Prof. Maniyar A.A
Department of Artificial Intelligence and Machine Learning
Sahyadri Valley College of Engineering and Technology, Rajuri, Pune, Maharashtra, India

Abstract—Healthcare organizations generate large volumes of sensitive information, yet collaboration across institutions is often restricted by privacy, governance, and security requirements. This paper presents a conceptual review of an integrated healthcare framework that combines Artificial Intelligence (AI), Federated Learning (FL), and Blockchain. In the framework, participating institutions retain patient records locally and train local models, while model updates are exchanged for collaborative learning. A permissioned blockchain layer is considered for recording selected model-training and access-control events, supporting traceability and tamper-evident auditing. Privacy-enhancing mechanisms discussed in the underlying seminar material include differential privacy, homomorphic encryption, secure communication, and access control. The paper explains the system workflow, examines its security and privacy properties, and compares it conceptually with centralized AI and federated learning without blockchain. It also discusses practical concerns such as non-identically distributed hospital data, communication overhead, malicious updates, scalability, interoperability, and regulatory governance. Numerical performance values cited from the referenced base paper are identified as previously reported results and are not presented as experiments conducted by the author. The analysis indicates that the combination may support collaborative healthcare AI, but clinical deployment requires independent validation, careful threat modelling, and institutional oversight.

Keywords—Artificial Intelligence, Federated Learning, Blockchain, Healthcare Data Privacy, Smart Contracts, Privacy-Preserving Machine Learning

I. INTRODUCTION

Digital healthcare systems increasingly rely on electronic health records, diagnostic images, laboratory reports, and data collected by monitoring devices. These sources can support clinical decision-making and research, but they also contain sensitive personal information. When several hospitals attempt to develop a common analytical model, transferring all records to a central repository can create privacy, governance, and security concerns. Differences in institutional policies and data formats may further restrict collaboration.

Artificial Intelligence can identify patterns in clinical data and assist with prediction or classification tasks. However, conventional centralized model training generally requires data to be collected or accessed at a common location. Federated Learning offers an alternative: each participating institution trains a model using its own records and shares model parameters or updates rather than the underlying patient data [1], [2]. This reduces the need for raw-data transfer, although it does not by itself eliminate privacy risks.

Blockchain can complement this process by maintaining a shared, tamper-evident record of selected transactions. In a healthcare collaboration setting, the ledger may record model-update submissions, validation outcomes, access requests, and consent-related events. Smart contracts can enforce predefined rules, subject to correct design and governance. The purpose of this paper is to examine how AI, FL, and blockchain can be combined into a privacy-aware healthcare collaboration framework.

Healthcare monitoring systems increasingly combine data from electronic health records, laboratory systems, diagnostic equipment, and Internet of Things (IoT) devices. These sources are useful for different tasks: electronic records provide clinical history, diagnostic images support image-based assessment, and wearable sensors can provide repeated measurements such as heart rate, temperature, and oxygen saturation. An AI model can learn patterns from these data, but its reliability depends on the quality and representativeness of the information used for training. A model developed at one hospital may not perform equally well at another hospital because patient populations, clinical procedures, and equipment can differ.

The proposed project therefore considers collaboration between multiple healthcare institutions rather than relying on a single data repository. Each institution remains responsible for its own patient records and local governance requirements. The shared component is the learning process: institutions contribute model updates under an agreed protocol, while the system maintains records of relevant participation and verification events. This separation between local data custody and collaborative model development is central to the patient-centric objective of the framework.

© 2026 IJSRET

II. RELATED WORK

Research on federated learning in medicine has demonstrated the potential for multi-institutional model development without pooling patient-level records. Rieke et al. discussed the role of FL in digital health and highlighted opportunities as well as implementation challenges [1]. Sheller et al. reported a multi-institutional medical-imaging application in which institutions collaborated without directly sharing patient data [2]. Dayan et al. demonstrated federated learning for predicting clinical outcomes in patients with COVID-19, illustrating the relevance of distributed learning in a real healthcare setting [3].

The literature also considers blockchain as a means of strengthening integrity, accountability, and trust in distributed systems. A blockchain-based healthcare FL approach has been explored by Chang et al. [4]. Other works cited in the supplied seminar material examine privacy-preserving healthcare data exchange, hierarchical federated learning, and blockchain-supported IoT healthcare [5]–[8]. These studies indicate that the combination of FL and blockchain is promising, but the choice of consensus mechanism, privacy protection, communication design, and governance remains application-dependent.

The reviewed material also identifies a key distinction: FL limits the movement of raw patient records, but model updates can still expose information or be manipulated. Therefore, additional safeguards and a clearly defined threat model are necessary. Blockchain can improve auditability, but it cannot guarantee that submitted data are clinically correct or that every participant behaves honestly.

A key idea in prior healthcare research is that federated learning can enable model development across institutions that cannot directly exchange patient-level data. Studies involving medical imaging and clinical outcome prediction illustrate the potential of this approach. However, federated learning should not be interpreted as a complete privacy solution. Model parameters and gradients may contain information about local training examples, and the system must also address malicious or unreliable participants. Blockchain-supported coordination is consequently considered as an additional trust and audit layer, not as a replacement for privacy engineering.

III. PROPOSED FRAMEWORK

The framework considered in this paper consists of healthcare institutions, local AI training environments, a federated aggregation process, a permissioned blockchain network, and privacy and governance controls. Each hospital stores its electronic health records, medical images, or wearable-device information within its own controlled environment. The local training component uses the institution's permitted data to update a shared model.

After local training, the institution prepares a model update for transmission. Depending on the deployment design, the update may be protected using secure communication, differential privacy, encryption, or secure aggregation. A permissioned blockchain can record the update identifier, participant identity, timestamp, validation status, and related audit metadata. Large medical records and model artefacts should generally remain in approved off-chain storage; the ledger can store references or cryptographic hashes rather than the raw clinical content.

A federated aggregation component combines accepted updates and produces a new global model. The updated model is then distributed to participating institutions for another training round. A hierarchical arrangement may be used when many hospitals are organized into regional groups. Smart contracts may automate predefined checks, but clinical authorization, identity verification, consent management, and legal accountability must remain part of the wider governance process.

The architecture can be understood as five cooperating components. First, the healthcare data layer stores electronic records, medical images, and monitoring-device readings within each participating institution. Second, the local AI layer performs data preparation and model training. Third, the federated coordination layer manages training rounds and aggregates accepted updates. Fourth, the blockchain layer maintains permissioned identities and tamper-evident records of selected events. Fifth, the application layer presents validated model outputs to authorized healthcare users. These components are logically separated so that the blockchain does not need to store large medical files or raw patient records.

Before local training, an institution should check data quality, remove duplicate or invalid records where appropriate, standardize permitted features, and apply the organization's approved de-identification or pseudonymization process.

The exact preprocessing procedure depends on the clinical task and data type. For example, medical images require image-specific normalization, whereas time-series readings from wearable sensors require handling of missing values and irregular sampling intervals. The same feature definitions and label meanings should be documented across institutions to reduce avoidable differences in training data.

© 2026 IJSRET

IV. METHODOLOGY AND WORKFLOW

The workflow is described as a conceptual system process rather than as a newly executed clinical experiment. It contains the following stages:

- Local data preparation: Each participating hospital validates and preprocesses its permitted records within its own infrastructure.
- Global model distribution: An initial model is shared with the participating institutions.
- Local model training: Each institution trains the model using local data and computes an update.
- Update protection and submission: The update is transmitted through an authenticated channel, with suitable privacy safeguards applied according to the threat model.
- Blockchain recording and validation: The permissioned ledger records relevant metadata and the outcome of predefined validation rules.
- Federated aggregation: Accepted model updates are combined to create the next global model.
- Model redistribution: The updated model is returned to participating institutions for the next training round.

A commonly used aggregation method is Federated Averaging (FedAvg). If institution i has n_i local training examples and the total number of examples is N, a weighted global update can be expressed as:

$$W(t+1) = \Sigma_i (n_i / N) w_i(t+1) \qquad (1)$$

Here, $w_i(t+1)$ represents the locally trained model parameters received from institution i, while $w(t+1)$ is the aggregated global model. The equation illustrates weighted aggregation; the exact aggregation and validation policy must be selected according to the model architecture and deployment requirements.

The aggregation stage requires a policy for accepting, rejecting, or down-weighting updates. A basic FedAvg procedure weights each local model according to the number of training examples, but sample count alone may not represent data quality or clinical diversity. A practical system may therefore combine weighted aggregation with validation checks, update-norm limits, anomaly detection, and monitoring of model performance across participating sites. These controls should be tested carefully because overly strict filtering may also reject legitimate updates from smaller hospitals or underrepresented patient groups.

The framework described in the seminar material also considers hierarchical federated learning. In a hierarchical arrangement, local institutions may first coordinate within a regional group before regional updates are combined at a higher level. This structure may reduce direct communication between every hospital and a central coordinator, and it can reflect regional healthcare administration. Its actual benefit depends on the number of institutions, network conditions, aggregation policy, and governance model; it should be evaluated rather than assumed.

V. SECURITY AND PRIVACY CONSIDERATIONS

Keeping raw patient records at their originating institutions can reduce direct exposure caused by central data collection. Nevertheless, gradients and model parameters may reveal information about training examples under some conditions. Differential privacy can limit the influence of individual records by introducing calibrated noise, while homomorphic encryption may permit selected computations on encrypted values. These techniques have computational and utility trade-offs and should be selected based on the sensitivity of the task.

Blockchain contributes a tamper-evident history of recorded events and can improve traceability between participating organizations. Permissioned membership and smart-contract rules can restrict who is allowed to submit or validate transactions. Cryptographic hashes can help detect changes to off-chain files when the verification process is correctly implemented. However, immutability does not establish the truth of a record, and blockchain does not automatically prevent model poisoning, identity compromise, or unauthorized disclosure.

A practical design should therefore include identity and key management, role-based access control, secure communication, update screening, monitoring for anomalous contributions, backup and recovery procedures, and a process for handling consent and data-retention requirements. Security evaluation should consider both honest-but-curious participants and active adversaries.

Differential privacy can be used to limit the influence of an individual record on a released model update by introducing calibrated noise. The privacy level depends on the noise mechanism, clipping strategy, privacy budget, and number of training rounds. Homomorphic encryption may allow selected computations to be performed on encrypted values, although it can introduce substantial computational cost. Secure multi-party computation can distribute a computation among parties so that no single participant needs to reveal its private input. These techniques address different threat models and may be combined selectively;

© 2026 IJSRET

they are not interchangeable and do not automatically guarantee regulatory compliance.

The blockchain component is best designed as a permissioned network in which participating hospitals and approved organizations have known identities. Smart contracts can enforce agreed rules for participant registration, update submission, validation status, and access requests. The ledger can store hashes, timestamps, model-version identifiers, consent-event references, and transaction outcomes. Sensitive medical records and large model files should remain in appropriately protected off-chain repositories, with the ledger storing only the minimum metadata needed for verification and audit. This reduces unnecessary exposure and avoids treating an immutable ledger as a general-purpose medical database.

VI. COMPARATIVE ANALYSIS

The following qualitative comparison summarizes the main architectural differences described in the seminar report. It is not a measured benchmark.

Centralized AI typically collects or accesses records centrally and depends on central security controls. In blockchain-supported federated learning, raw records remain at participating institutions, local models are aggregated, and selected transactions can be recorded on a permissioned ledger. This reduces raw-data movement but adds identity, consensus, smart-contract, and communication overhead.

A useful comparison can be made across three approaches. In centralized AI, data from different sources are collected or accessed by a central platform, which can simplify training but increases the importance of central security, data-transfer controls, and governance. In conventional federated learning, raw records remain at participating sites and model updates are exchanged, reducing the need for direct record pooling but leaving questions about update privacy, participant trust, and auditability. In blockchain-supported federated learning, a permissioned ledger and smart contracts add traceability and shared verification, but also introduce consensus, storage, governance, and transaction overhead. The third approach is therefore not automatically superior in every deployment; its suitability depends on the collaboration requirements and threat model. Potential applications include remote patient monitoring, medical image analysis, disease-risk estimation, clinical outcome Prediction, electronic health record collaboration, and multi-hospital research. For remote monitoring, local devices or hospital gateways may collect vital-sign measurements and forward permitted data to the institution's AI environment. For medical imaging, each hospital can train on locally held images and contribute protected model updates. For clinical prediction, the participating institutions must agree on outcome definitions, observation windows, and evaluation measures. In every case, model outputs should support qualified healthcare professionals rather than replace clinical judgment.

VII. DISCUSSION AND RESEARCH CHALLENGES

The integrated approach can support collaboration when institutions cannot freely pool their patient records. Its value depends on the quality of local data, the reliability of participating organizations, and the design of the aggregation and audit processes. A permissioned ledger may make selected activities easier to trace, while federated training can make use of distributed datasets without routinely moving raw records.

Several issues remain unresolved. Hospital datasets may be non-identically distributed because patient populations, clinical practices, equipment, and coding systems differ. This can affect convergence and model generalization. Repeated communication between institutions may introduce latency and bandwidth costs. Blockchain consensus and ledger storage can add further processing overhead, particularly for smaller hospitals and edge devices.

Malicious or compromised participants may submit poisoned updates, while model updates may expose information if privacy controls are insufficient. Interoperability with existing electronic health record systems is also difficult because institutions may use different standards and interfaces. Finally, compliance cannot be inferred from the use of FL or blockchain alone; each deployment requires legal review, data governance, consent management, and clinical validation.

Ethical and governance considerations are equally important. Patients should receive clear information about the intended use of their data, and institutions should define lawful processing, consent or other applicable legal bases, retention periods, and procedures for access or correction. Model performance should be assessed across relevant patient groups and institutions because an overall accuracy score can hide unequal error rates. Audit records may help establish which organization submitted an update and when, but they cannot prove that a prediction is clinically correct, unbiased, or appropriate for a particular patient. Human oversight and a process for investigating harmful outputs remain necessary.

© 2026 IJSRET

VIII. REPORTED RESULTS AND LIMITATIONS OF THE EVIDENCE

The supplied reference paper reports several performance figures, including a 7–12% model-accuracy improvement, 94.6% protection effectiveness against model-inversion attacks, and a 98% figure for protection against adversarial manipulation. It also reports blockchain update latency values between 1.2 and 2.0 seconds for networks ranging from 10 to 100 participating hospitals. These values are reproduced here only as claims reported by that source; they are not results generated or independently verified in the present paper.

The available source material does not provide enough reproducibility detail in the seminar report to independently validate those figures. A rigorous evaluation would need to specify datasets, model architecture, data partitioning, baseline configuration, privacy parameters, threat model, hardware, consensus settings, and statistical procedures. Claims about regulatory readiness or clinical suitability should likewise be supported by documented compliance assessment and real-world validation.

A complete evaluation of the framework should measure more than predictive accuracy. Relevant measures include sensitivity, specificity, precision, recall, F1-score, area under the ROC curve, calibration, and performance variation across hospitals. System-level measures may include communication volume per training round, model convergence time, blockchain transaction throughput, confirmation latency, storage requirements, energy use, and the computational cost of privacy mechanisms. Security evaluation should specify the attacker capabilities and test scenarios such as model inversion, membership inference, poisoning, and unauthorized update submission. Results should be compared with centralized and non-blockchain federated baselines using the same datasets and evaluation protocol.

The reference material mentions adaptive differential privacy, homomorphic encryption, adversarial-attack detection, smart contracts, and Delegated Proof-of-Stake (DPoS) as possible elements of the integrated design. These techniques should be selected according to the deployment environment. For example, DPoS may reduce confirmation delays compared with energy-intensive proof-of-work designs, but it introduces questions about validator selection, concentration of authority, and governance. Similarly, encryption may improve confidentiality while increasing processing time. A practical implementation should document these trade-offs and justify each mechanism against explicit security and performance requirements.

Future Scope

Future work may investigate lightweight federated models for resource-constrained hospitals, communication-efficient aggregation, and scalable permissioned blockchain configurations. Further study is also needed on robust aggregation against poisoned updates, privacy-budget selection, secure multi-party computation, and interoperability with electronic health record standards. Evaluation should use clearly documented datasets and realistic multi-institutional settings, followed by security testing and clinical review before any operational deployment.

Implementation planning should also consider interoperability with existing hospital information systems. Healthcare institutions may use different electronic record platforms, coding systems, imaging formats, and identity-management services. Common data models, standardized interfaces, and documented metadata are important for meaningful collaboration. The system should support secure authentication, role-based authorization, key rotation, revocation of compromised credentials, backup and recovery, and incident response. Cross-border collaboration adds further complexity because data-protection obligations and institutional responsibilities may differ by jurisdiction.

IX. CONCLUSION

This paper has reviewed a conceptual healthcare collaboration framework integrating Artificial Intelligence, Federated Learning, and Blockchain. Federated Learning can support joint model development while reducing the need to transfer raw patient records, and blockchain can provide a shared audit trail for selected transactions and model activities. Privacy-enhancing techniques, identity controls, and secure aggregation are important because federated updates are not inherently private or trustworthy. The approach offers a potential direction for privacy-aware healthcare AI, but its effectiveness depends on careful system design, independent empirical evaluation, interoperability, regulatory governance, and clinical oversight.

Acknowledgment

The author gratefully acknowledges the guidance and academic support provided by the Department of Artificial Intelligence and Machine Learning, Sahyadri Valley College of Engineering and Technology, Rajuri, Pune.

REFERENCES

1. N. Rieke et al., “The future of digital health with federated learning,” npj Digital Medicine, vol. 3, article 119, 2020, doi: 10.1038/s41746-020-00323-1.

© 2026 IJSRET

2. M. J. Sheller et al., “Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data,” Scientific Reports, vol. 10, article 12598, 2020, doi: 10.1038/s41598-020-69250-1.
3. I. Dayan et al., “Federated learning for predicting clinical outcomes in patients with COVID-19,” Nature Medicine, vol. 27, pp. 1735–1743, 2021, doi: 10.1038/s41591-021-01506-3.
4. Y. Chang, C. Fang, and W. Sun, “A blockchain-based federated learning method for smart healthcare,” Computational Intelligence and Neuroscience, 2021, doi: 10.1155/2021/4376418.
5. D. Commey, S. Hounsinou, and G. V. Crosby, “Securing health data on the blockchain: A differential privacy and federated learning framework,” arXiv:2405.11580, 2024.
6. N. Waheed et al., “FedBlockHealth: A synergistic approach to privacy and security in IoT-enabled healthcare through federated learning and blockchain,” arXiv:2304.07668, 2023.
7. V. Stephanie et al., “Trustworthy privacy-preserving hierarchical ensemble and federated learning in Healthcare 4.0 with blockchain,” arXiv:2305.09209, 2023.
8. W. Li et al., “Privacy-preserving federated brain tumour segmentation,” in Proceedings of Machine Learning for Health, pp. 133–145, 2020.

© 2026 IJSRET