

Explainable AI-Driven Federated Learning for Privacy-Preserving Healthcare Analytics and Disease Detection

Himani Adsar¹, Ramesh Shahabade² and Shudhodhan Bokefode^{3*}

¹Pune Institute of Computer Technology, Pune University, Pune, Maharashtra, India

²Terna Engineering College, Nerul, Navi Mumbai, Maharashtra, India

³MAEER's MIT College of Railway Engineering and Research, Barshi, Solapur, Maharashtra, India

E-mails: ¹himaniadsar7@gmail.com, ²rameshshahabade@ternaengg.ac.in and ³shudhodhan358@gmail.com

*Corresponding Author: shudhodhan358@gmail.com

Received: 14th June, 2026; Revised: 20th July 2026; Accepted: 12th September, 2026; Available Online: 15th September 2026

ABSTRACT

Modern healthcare environments generate massive volumes of sensitive data from hospitals, laboratories, and wearable devices, creating opportunities for improved diagnostics but also raising significant privacy and regulatory challenges. This study introduces an Explainable Federated Learning (FL) framework that enables multiple medical institutions to collaboratively train deep neural networks without sharing raw patient data. By exchanging only model parameters, the framework preserves data confidentiality while maintaining strong predictive performance. To ensure transparency, explainability modules—including SHAP, Grad-CAM, and attention maps—are integrated to help clinicians interpret the model's decision-making process. Recent advancements demonstrate that FL improves generalization and strengthens privacy using differential privacy and secure aggregation techniques. Our proposed system unifies these developments by combining distributed learning with an interpretable inference layer. Experiments on multi-hospital EHR and medical imaging datasets show a 6–12% improvement in diagnostic accuracy while achieving a privacy budget of $\epsilon < 2$, supporting trustworthy and regulation-compliant healthcare AI.

Keywords: Federated Learning, Explainable Artificial Intelligence (XAI), Medical Disease Prediction, Privacy Preservation, Interpretability, Healthcare Data Analytics, Secure Aggregation, Differential Privacy

How to cite this article: Adsar H, Shahabade R, Bokefode S, Explainable AI-Driven Federated Learning for Privacy-Preserving Healthcare Analytics and Disease Detection. Int J Drug Deliv Technol. 2026;16(83s): 242-250; DOI: 10.25258/ijddt.16.83s.26

Source of support: Nil.

Conflict of interest: None

1. INTRODUCTION

The rapid growth of electronic health records (EHRs), biomedical imaging, and wearable sensor data has fueled a new wave of data-driven innovation in healthcare analytics. Machine learning (ML) models built on such data have shown great promise in predicting disease risks, assisting clinical decision-making, and personalizing treatment strategies. However, traditional centralized training methods typically require aggregating all patient data onto a single server, raising serious privacy concerns and violating data protection laws such as HIPAA and GDPR.

To address these issues, Federated Learning (FL) has emerged as a transformative paradigm that enables collaborative model training across multiple healthcare institutions—without sharing raw data [1], [2], [4], [5], [7]. By keeping data local and exchanging only encrypted model parameters, FL provides a scalable and privacy-preserving solution that overcomes data silos and facilitates secure collaboration across hospitals [8], [13].

Recent reviews emphasize the growing adoption of FL across various healthcare applications—from medical imaging and genomics to wearable IoT-based health

monitoring [3], [4], [9]. For instance, Zhang et al. [2] proposed advanced methods to tackle data heterogeneity and communication delays in FL networks, while Li et al. [3] developed privacy-aware workflows aligned with hospital data governance standards. Abbas et al. [4] highlighted the integration of FL with smart healthcare ecosystems to enable real-time disease surveillance, and Eden et al. [7] explored governance models for the responsible deployment of FL in clinical environments.

Despite these achievements, a key limitation remains: most FL models operate as “black boxes”, offering limited transparency into how predictions are made. This opacity reduces clinicians' confidence in AI-assisted diagnoses, as medical decisions must be supported by interpretable and justifiable reasoning [6], [15], [17]. Explainable Artificial Intelligence (XAI) directly addresses this challenge by generating human-understandable explanations for complex models using methods such as SHAP, LIME, Grad-CAM, and attention mechanisms [16]–[18]. The fusion of XAI and FL—often referred to as Federated Explainable AI (Fed-XAI)—represents a promising direction for achieving both privacy and transparency in medical AI [8], [10], [11].

*Author for Correspondence: shudhodhan358@gmail.com

For example, Amato and Branco [5] introduced the SemFedXAI framework for interpretable decision-making in federated settings, while Alkhanbouli et al. [6] provided a comprehensive review of XAI methods for disease prediction, outlining future pathways toward interpretable federated architectures.

To further enhance privacy, researchers have explored combining FL with differential privacy and homomorphic encryption [10], [12], [20], [21]. Shukla et al. [10] demonstrated a privacy-enhanced FL approach for breast cancer detection, achieving strong diagnostic accuracy while preserving patient confidentiality. Similarly, Gupta et al. [21] proposed an encrypted federated ensemble model for breast cancer imaging. However, most existing approaches treat privacy and explainability as two separate challenges, lacking a unified solution that effectively ensures both.

This research bridges that gap by introducing the Explainable Federated Medical Pattern Analysis (EFMPA) framework. EFMPA integrates distributed model training, secure aggregation, and post-hoc interpretability within a single, cohesive architecture. It empowers clinicians to visualize and validate model decisions without exposing sensitive data. The framework's novelty lies in its combination of differential privacy-based aggregation and XAI-driven interpretability metrics, fostering accountability, transparency, and trust in federated medical AI systems.

The remainder of this paper is organized as follows: Section II reviews prior work on federated learning and explainable AI integration; Section III details the proposed EFMPA methodology; Section IV presents the experimental results and analysis; and Section V concludes with key insights and directions for future research in ethical, explainable, and privacy-preserving healthcare intelligence.

II. RELATED WORK

A. Federated Learning in Healthcare

Federated learning (FL) has emerged as a cornerstone of privacy-preserving data analytics across distributed medical environments. Unlike traditional centralized machine learning systems, FL enables hospitals and research institutions to collaboratively train predictive models without transferring sensitive patient data to a central repository.

Early frameworks demonstrated its strong potential for medical imaging diagnostics, multi-institutional EHR analysis, and wearable sensor data modeling. One of the earliest systematic reviews outlined federated architectures for clinical data integration, showing how cross-silo training enhances disease-pattern generalization. Subsequent implementations provided deployment guidelines for hospital networks, confirming that decentralized aggregation can achieve performance comparable to centralized models while ensuring compliance with healthcare privacy regulations.

The adoption of FL has since expanded to multiple domains, including pathology, ophthalmology, and intensive care analytics. Comprehensive surveys have explored communication efficiency, data heterogeneity management, and cross-domain scalability strategies. Studies on smart healthcare systems integrating Internet-of-Things (IoT) devices have shown that federated edge models can process real-time physiological signals locally, eliminating the need to transmit raw data to cloud servers. Governance research has further emphasized the importance of auditable pipelines and standardized consent mechanisms for safe cross-institutional collaboration.

Despite these advancements, several challenges persist. Heterogeneous data distributions across hospitals often lead to model divergence and bias, while limited bandwidth and client computation constraints slow convergence. Moreover, malicious clients can still extract information from gradient updates unless robust cryptographic protections are implemented. Consequently, current research focuses on hybrid solutions that embed privacy-enhancing mechanisms—such as secure aggregation, homomorphic encryption, and differential privacy—into the federated learning workflow.

B. Explainable Artificial Intelligence in Medical Applications

While federated architectures successfully protect data confidentiality, they often compromise interpretability—a crucial factor for clinical acceptance. Explainable Artificial Intelligence (XAI) seeks to make complex models more transparent by illustrating how particular features, signals, or image regions influence model predictions.

Recent literature classifies XAI methods into two main categories: post-hoc model-agnostic techniques (e.g., SHAP and LIME) and intrinsically interpretable models (e.g., attention-based networks or rule-based systems). Systematic reviews show a rapid increase in XAI applications across healthcare—spanning radiology, cardiology, genomics, and intensive care prognosis—demonstrating that interpretability improves clinician trust and enhances diagnostic reliability.

However, most existing XAI approaches are designed for centralized learning frameworks. When applied to federated environments, interpretability becomes more complex: each participating institution may have domain-specific data distributions and distinct local explanation contexts. This fragmentation leads to inconsistencies in explanations and limits the reproducibility of global model reasoning, hindering clinical validation across sites.

C. Integrating Explainability with Federated Learning

The integration of Explainable Artificial Intelligence with Federated Learning—commonly known as Explainable Federated Learning (Fed-XAI)—has recently gained momentum as a means of achieving both privacy and interpretability in medical AI systems. Semantic frameworks such as SemFedXAI apply ontological

reasoning to harmonize local explanations across federated clients, demonstrating that shared semantics can significantly improve global interpretability.

Empirical studies further validate the potential of Fed-XAI. For example, differential privacy-enhanced models for breast cancer detection have achieved high diagnostic accuracy while providing interpretable feature attributions for clinicians. Similarly, hybrid explainable vision transformer architectures combine global aggregation with local saliency maps, and ensemble-based federated systems offer real-time interpretability for multimodal medical data.

Beyond classification tasks, fairness-aware explainable federated models for electronic health records (EHRs) demonstrate how interpretability can support bias detection and model auditing. Other studies have integrated blockchain and homomorphic encryption to enable verifiable and privacy-preserving model updates while maintaining transparency in local decision-making.

Despite these encouraging developments, several research gaps remain. Existing frameworks often struggle to balance privacy strength with explanation fidelity, while the communication overhead from transmitting explanation metadata can affect scalability. Furthermore, standardized evaluation metrics for assessing explanation quality in distributed environments are still underdeveloped.

III. METHODOLOGY

A. Overview of the Proposed Framework

The Explainable Federated Medical Pattern Analysis (EFMPA) framework combines the strengths of Federated Learning (FL), Differential Privacy (DP), and Explainable Artificial Intelligence (XAI) to enable interpretable and privacy-preserving disease prediction across multiple medical institutions. The framework is designed for deployment in multi-hospital or multi-clinic environments, allowing participants to collaboratively train a global model without transferring raw patient data.

Each participating healthcare institution maintains its own local dataset, comprising electronic health records (EHRs), laboratory test results, and medical imaging data. EFMPA operates through three main layers:

1. **Local Model Training Layer** – Each client independently trains a local model using its private dataset.
2. **Federated Aggregation Layer** – The trained model parameters, encrypted and noise-protected, are transmitted to a secure aggregator for global model updates.
3. **Explainability Layer** – Post-training, interpretability modules such as SHAP, Grad-CAM, and attention-based visualizations generate transparent explanations for clinical interpretation.

To ensure privacy during parameter exchange, EFMPA applies differential privacy noise injection and secure aggregation protocols, preventing any leakage of sensitive medical information through gradient updates or model weights. After each global communication round, the updated global model is redistributed to all clients for further training, creating a continuous, adaptive, and privacy-preserving learning cycle.

B. System Architecture

The EFMPA system architecture is organized into five logical modules, as conceptually illustrated in *Fig. 1*.

1. Data Acquisition and Preprocessing Module

Each healthcare institution collects multimodal data from EHR systems, wearable sensors, and imaging modalities. Data preprocessing is performed locally to ensure quality and uniformity—handling missing values, normalizing features, and removing noise.

For imaging data: contrast normalization and resizing are applied.

For EHR and structured data: tokenization and encoding ensure consistent feature representation across sites.

2. Local Model Training Module

Each client trains a model suited to its data type:

CNNs for imaging data, GRUs for sequential data, and Fully connected neural networks for structured tabular records.

Training uses stochastic gradient descent with batch normalization and dropout to reduce overfitting. Before transmission, differential privacy noise is added to model gradients or weights to protect patient information.

3. Federated Aggregation Module

At the central server, local model updates are aggregated using the Federated Averaging (FedAvg) algorithm. Weight updates are averaged proportionally to each client's dataset size and local model performance. The updated global model is redistributed to all clients, and the process repeats until convergence.

4. Privacy Preservation and Security Module

EFMPA enforces a dual-layer privacy protection strategy:

- **Differential Privacy (DP):** Gaussian noise is injected into gradients before transmission to ensure individual-level privacy.
- **Secure Aggregation & Homomorphic Encryption:** The server can only view aggregated model updates, not individual contributions.

This ensures protection against reconstruction or inference attacks, maintaining strong data confidentiality throughout the training process.

5. Explainable AI and Interpretation Module

After each training cycle, the XAI layer interprets model predictions through tailored explainability tools:

- SHAP and LIME quantify the most influential clinical features for structured data.
- Grad-CAM highlights disease-relevant regions in medical images.
- Attention visualizations identify critical time-series or physiological patterns in sequential data.

These explanations are presented to clinicians through visual or textual dashboards, enabling them to understand and validate the model's reasoning in a transparent and trustworthy manner.

C. Workflow of EFMPA

The EFMPA workflow follows a federated iterative learning process comprising four main phases:

1. Initialization:

The central coordinator initializes a global model with random parameters and distributes it to all participating clients.

2. Local Training:

Each client trains the model locally for several epochs, computes gradients, and applies differential privacy noise before sharing updates.

3. Model Aggregation:

The server aggregates the encrypted updates using FedAvg, producing a refined global model.

4. Explainability and Evaluation:

Upon convergence, the XAI module computes interpretability metrics (e.g., feature importance, heatmaps, and attention weights). These results are displayed on a federated dashboard for clinician review and cross-client consistency analysis.

This process continues across multiple communication rounds until both accuracy and explainability metrics stabilize, achieving equilibrium between learning performance and interpretability.

D. Algorithmic Design

Input:

Local datasets $\mathcal{D}_i$ from multiple clients (e.g., hospitals, devices)

Global model parameters θ_0 (initial model)

Differential privacy noise level σ

Learning rate η

E. Evaluation Metrics

The EFMPA framework is assessed using six key performance indicators:

1. **Accuracy (ACC):** Proportion of correctly classified disease outcomes.

2. **F1-Score:** Harmonic mean of precision and recall, especially relevant for imbalanced datasets.
3. **AUC (Area under the ROC Curve):** Reflects the model's discriminative power.
4. **Interpretability Index:** A composite measure combining fidelity, stability, and alignment with human understanding.
5. **Privacy Budget (ϵ):** Quantifies privacy protection level under differential privacy.
6. **Communication Cost:** Represents data transmission overhead per training round.

By optimizing these metrics jointly, EFMPA achieves a balanced trade-off between diagnostic accuracy, interpretability, and data privacy.

F. Implementation Details

The EFMPA framework was implemented using TensorFlow Federated and PyTorch Lightning, integrated with SHAP and Captum libraries for model explainability. Experiments were performed on a distributed simulation of ten hospitals, each hosting partitions of benchmark datasets such as:

MIMIC-IV (EHR data), ChestX-ray14 (medical imaging), and UCI Chronic Kidney Disease (tabular data).

Model training employed the Adam optimizer with a learning rate of 0.001, a mini-batch size of 32, and 50 global communication rounds. Privacy noise levels were varied between $\epsilon = 1.0$ and $\epsilon = 2.5$ to study trade-offs between privacy and accuracy.

A federated visualization dashboard was developed to present:

Patient-level feature importance, Global feature distributions, and Cross-client explanation consistency.

This dashboard ensures clinically interpretable, transparent, and auditable insights into the federated model's decision-making process across participating healthcare institutions.

IV. EXPERIMENTAL RESULTS AND DISCUSSION

A. Experimental Setup

To evaluate the effectiveness of the proposed Explainable Federated Medical Pattern Analysis (EFMPA) framework, experiments were conducted on three well-known medical benchmark datasets that represent different types of healthcare data and diagnostic tasks:

MIMIC-IV – a structured electronic health record (EHR) dataset containing over 400,000 patient admissions, including demographic details, lab results, and treatment records.

ChestX-ray14 – a large-scale image dataset with more than 100,000 chest radiographs, annotated for 14 disease categories such as pneumonia, tuberculosis, and pulmonary fibrosis.

Chronic Kidney Disease (CKD) Dataset (UCI Repository) – a smaller, structured dataset with 400 patient records and 24 clinical features used for binary classification.

Each dataset was divided among ten simulated hospital clients to reflect real-world scenarios where data are distributed and non-identically distributed (non-IID). The federated learning (FL) environment was implemented using TensorFlow Federated, with secure aggregation protocols ensuring that model updates were exchanged safely between clients and the central server.

All experiments were carried out on a distributed setup equipped with an Intel i7 processor, 32 GB RAM, and an NVIDIA RTX 4070 GPU.

The performance of EFMPA was compared with four baseline configurations:

1. **Centralized Model:** Traditional model trained on all combined data without privacy constraints.
2. **Local Models:** Individual models trained at each client independently, without collaboration.
3. **Standard Federated Learning (FedAvg):** Classic FL setup without privacy or explainability features.
4. **Differentially Private FL (DP-FL):** FL incorporating Gaussian differential privacy but lacking interpretability mechanisms.

B. Evaluation Criteria

Model performance was assessed using a mix of predictive, interpretability, and privacy-related metrics to provide a balanced evaluation:

Accuracy (ACC): The proportion of correctly classified cases.

Precision and Recall: Metrics that assess the model's ability to detect minority disease cases accurately.

F1-Score: The harmonic mean of precision and recall, capturing performance balance between false positives and false negatives.

AUC (Area under ROC Curve): Evaluates the model's ability to discriminate between disease and non-disease cases.

Interpretability Index (II): A composite measure combining explanation fidelity, stability across rounds, and human-alignment (clinician interpretability).

Privacy Budget (ϵ): Controls the trade-off between privacy and accuracy in differential privacy.

Communication Cost (CC): Represents the total bandwidth consumed per global aggregation round.

C. Quantitative Results

Table I summarizes the comparative performance of EFMPA against the baseline methods across the three datasets. The results demonstrate that EFMPA consistently achieves higher accuracy, stronger interpretability, and better privacy preservation than other configurations, proving its effectiveness for real-world medical AI deployment.

Table 1: Comparative Performance of EFMPA against Baseline Methods

Method	Accuracy (%)	F1-Score	AUC	Interpretability Index	Privacy (ϵ)
Local Models	79.2	0.76	0.81	–	∞
Centralized Model	90.8	0.88	0.92	0.42	∞
Standard FedAvg	88.1	0.85	0.89	0.44	∞
DP-FL ($\epsilon = 3.0$)	86.5	0.82	0.87	0.43	3.0
Proposed EFMPA ($\epsilon = 2.0$)	92.4	0.90	0.94	0.78	2.0

The proposed EFMPA framework consistently outperformed the baseline methods across all datasets. Compared with traditional FedAvg, EFMPA improved classification accuracy by an average of 4.3 %, AUC by 0.05, and interpretability index by 77 %, while maintaining a stricter privacy budget of $\epsilon = 2.0$.

The addition of differential privacy and secure aggregation caused minimal performance degradation (< 2 %) compared with the non-private centralized model, demonstrating that privacy can be enforced without substantial accuracy loss.

D. Visualization and Interpretability Analysis

D. Interpretability Analysis

Interpretability in the EFMPA framework was evaluated using several post-hoc explainability techniques integrated directly into the federated workflow. These tools provided clinicians with clear, data-driven insights into how predictions were made across different modalities and datasets.

SHAP Feature Importance (EHR and CKD Datasets):

For the MIMIC-IV and CKD datasets, SHAP (SHapley Additive exPlanations) was used to identify the most influential features driving disease predictions. The SHAP summary plots revealed that serum creatinine, hemoglobin

level, and blood pressure were the dominant contributors to CKD detection, while patient age and sodium concentration served as secondary predictors. Importantly, the variance in the top five feature rankings across all clients was below 5%, suggesting strong interpretability consistency among hospitals in the federated setup.

Grad-CAM Visual Explanations (ChestX-ray14):

For medical imaging data, EFMPA employed Grad-CAM (Gradient-weighted Class Activation Mapping) to visualize model attention on chest X-rays. The resulting heatmaps consistently highlighted pathologically relevant lung regions linked to pneumonia and pulmonary fibrosis. When compared to conventional (non-federated) CNN models, EFMPA's Grad-CAM explanations achieved a notably higher clinician plausibility score (8.6/10), indicating improved diagnostic transparency and greater trust among medical experts.

Attention-based Explanations (MIMIC-IV):

In sequential EHR analysis, the attention visualization module highlighted temporal dependencies and feature relevance over time. Features such as heart rate and oxygen saturation received the highest attention scores in patients with respiratory distress, enabling clinicians to trace the evolution of disease severity and identify critical progression patterns. This capability enhances both interpretability and clinical decision support within temporal data contexts.

E. Privacy and Efficiency Evaluation

Privacy protection and computational efficiency were thoroughly assessed to validate EFMPA's real-world feasibility.

Privacy Evaluation:

To ensure confidentiality, EFMPA was tested against membership inference and gradient inversion attacks. Results showed a 92% reduction in data-leakage rates compared with standard, non-private FL models. Under Gaussian differential privacy with a noise parameter of $\epsilon = 2$, the system maintained less than 2% accuracy loss, demonstrating a strong balance between privacy and model performance. This result confirms that EFMPA effectively safeguards sensitive patient data without compromising diagnostic reliability.

Efficiency Evaluation:

The communication overhead introduced by explainability metadata was analyzed per training round. Although integrating explainable AI increased transmission volume by about 11%, this overhead was considered acceptable given the substantial gains in model transparency and clinician trust. Furthermore, applying model compression techniques such as quantization and sparsification effectively reduced bandwidth consumption, keeping overall system efficiency within practical limits.

F. Comparative Discussion

A detailed comparison across models and metrics revealed several key observations:

1. Performance Stability:

EFMPA achieved stable convergence across diverse and heterogeneous client datasets. The standard deviation of global accuracy across clients was limited to $\pm 1.7\%$, confirming that the model maintained fairness and consistent learning performance across sites.

2. Interpretability–Privacy Synergy:

Contrary to common assumptions that privacy constraints may hinder interpretability, EFMPA demonstrated that the two can complement each other. By integrating explainability within privacy-preserving federated learning, EFMPA improved model debugging and reliability, enabling better feature attribution even under strict differential privacy settings.

3. Clinical Applicability:

The interpretability outputs generated by the EFMPA framework were validated by domain experts for clinical relevance and consistency. The features identified by the explainable AI modules—such as serum creatinine for CKD, lung consolidation for pneumonia, and oxygen saturation trends for respiratory failure—aligned well with established clinical knowledge, strengthening confidence in model-driven insights.

4. Scalability and Generalization:

EFMPA demonstrated smooth scalability across ten simulated clients, and the framework can be readily extended to real-world hospital networks with minimal modification. Tests on unseen datasets confirmed strong generalization performance, indicating that federated collaboration enables effective knowledge transfer while maintaining data sovereignty within institutional silos.

G. Result Visualization




Figure 1: Illustrates a comparative bar chart of classification accuracy across the evaluated models. The proposed EFMPA framework achieved the highest performance, combining accuracy, interpretability, and privacy preservation.




Figure 2: Showing the comparative F1-Score and Interpretability Index for the five configurations — Local Models, Centralized Model, FedAvg, DP-FL, and EFMPA.

H. Discussion of Key Findings

The experimental findings clearly demonstrate that embedding explainability within federated learning frameworks significantly enhances both technical performance and clinical reliability. The interpretable feedback generated by the EFMPA system empowers clinicians to validate and trust AI-driven predictions, effectively narrowing the long-standing gap between opaque “black-box” models and real-world medical decision-making.

Equally important, EFMPA’s privacy-preserving mechanisms ensure full compliance with data governance and regulatory frameworks, including GDPR and HIPAA, without compromising model accuracy or utility. This balance between interpretability and privacy represents a crucial advancement for deploying AI responsibly in healthcare ecosystems.

In contrast to conventional methods that address explainability and privacy as separate challenges, EFMPA provides a unified, modular architecture that enables seamless collaboration among multiple hospitals or research centers. The dual-layer privacy protection—combining secure aggregation with differential privacy—effectively prevents sensitive data exposure, while the

explainability module enhances auditability, accountability, and transparency throughout the learning process.

Overall, the results underscore a pivotal insight: the future of healthcare AI must integrate federated learning, explainable AI, and privacy-preserving technologies in tandem. Such convergence not only promotes ethical and trustworthy clinical deployment but also lays the foundation for scalable, regulation-compliant, and human-centered medical intelligence systems.

V. CONCLUSION AND FUTURE SCOPE

A. Conclusion

This study introduced the Explainable Federated Medical Pattern Analysis (EFMPA) framework—an integrated approach that combines Federated Learning (FL), Differential Privacy (DP), and Explainable Artificial Intelligence (XAI) to achieve both privacy preservation and interpretability in medical disease pattern detection. Unlike conventional centralized AI models that rely on data aggregation, EFMPA enables collaborative model training across multiple healthcare institutions while ensuring that raw patient information remains securely within each site.

The framework employs secure aggregation and differential privacy mechanisms to protect sensitive model parameters, while post-hoc interpretability tools such as SHAP, LIME, and Grad-CAM generate transparent visual and feature-based explanations of predictions. Experimental evaluations on MIMIC-IV, ChestX-ray14, and Chronic Kidney Disease (CKD) datasets confirmed EFMPA's superior performance, achieving an average accuracy of 92.4%, F1-score of 0.90, and an interpretability index of 0.78, surpassing both conventional federated and centralized baselines.

Beyond quantitative metrics, EFMPA enhances clinical interpretability by delivering meaningful explanations—feature attributions for EHR data, saliency maps for imaging, and attention plots for temporal sequences. These capabilities enable clinicians to verify and trust AI-generated predictions, supporting transparent and responsible medical decision-making. The privacy mechanisms integrated into EFMPA reduced data-leakage risks by over 90%, demonstrating that strong privacy protection can coexist with high diagnostic performance.

In summary, EFMPA represents a unified, modular architecture that simultaneously addresses the three key challenges of modern medical AI—data privacy, model interpretability, and cross-institutional collaboration. The framework lays the foundation for scalable, ethical, and regulation-compliant healthcare analytics capable of functioning under strict privacy and governance constraints.

B. Limitations

Despite its promising results, the EFMPA framework has several limitations that warrant further attention:

Communication Overhead:

The inclusion of explainability metadata increases communication costs by approximately 10–12%, which may strain bandwidth in resource-limited or edge environments.

Computational Complexity:

Generating interpretability reports for high-resolution medical images adds computational overhead, extending total training time per federated round.

Non-IID Data Handling:

Extreme heterogeneity across hospital datasets can affect convergence stability. More adaptive aggregation algorithms such as FedProx, FedDyn, or Scaffold could improve robustness.

Limited Real-World Validation:

Current evaluations rely on public datasets. Broader validation across actual hospital networks with real-time data exchange is necessary to assess clinical readiness and operational scalability.

Addressing these challenges will strengthen EFMPA's practicality and prepare it for large-scale deployment in real-world healthcare ecosystems.

C. Future Scope

The integration of explainability with privacy-preserving federated intelligence opens several promising directions for future research:

Multimodal Data Fusion:

Expanding EFMPA to handle multimodal medical data—including EHRs, genomic sequences, and medical imaging—can enhance disease-pattern recognition and enable personalized healthcare.

Edge AI and Federated Reinforcement Learning:

Implementing EFMPA on edge devices and incorporating reinforcement learning can support adaptive, real-time monitoring across wearable sensors and mobile health platforms.

Causal and Symbolic Explainability:

Integrating causal inference and knowledge-graph reasoning could move explainability beyond post-hoc visualization toward understanding causal relationships in patient health dynamics.

Fairness and Bias Mitigation:

Introducing federated fairness-aware optimization will ensure equitable model performance across demographic subgroups, addressing ethical concerns and reducing healthcare disparities.

Quantum-Safe Federated Security:

Future work can explore quantum-resistant encryption to protect federated communication against next-generation cyber threats.

Regulatory Compliance and Governance Frameworks:

Developing standardized governance protocols aligned with HIPAA, GDPR, and emerging AI regulations will be essential for clinical integration and global adoption.

Through these extensions, EFMPA can evolve into a comprehensive, trustworthy, and globally deployable healthcare AI ecosystem, prioritizing patient trust, transparency, and ethical intelligence in clinical practice.

REFERENCES

- [1] Z. L. Teo, L. Jin, S. Li, D. S. W. Ting, et al., "Federated machine learning in healthcare: A systematic review on clinical applications and technical architecture," *Cell Rep. Med.*, vol. 5, no. 2, p. 101179, Feb. 2024. doi:10.1016/j.xcrm.2024.101179.
- [2] F. Zhang, D. Kreuter, Y. K. Chen, et al., "Recent methodological advances in federated learning for healthcare," *Patterns*, vol. 5, no. 6, p. 101006, Jun. 2024. doi:10.1016/j.patter.2024.101006.
- [3] M. Li, A. A. H. Barakat, et al., "Implementing federated learning in healthcare," *Comput. Biol. Med.*, vol. 150, p. 106563, Apr. 2025. doi:10.1016/j.combiomed.2024.106563.
- [4] S. R. Abbas, Z. Abbas, A. Zahir, and S. W. Lee, "Federated learning in smart healthcare: A

- comprehensive review on privacy, security, and predictive analytics with IoT integration,” *Healthcare*, vol. 12, no. 24, art. 2587, Dec. 2024. doi:10.3390/healthcare12242587.
- [5] A. Amato and D. Branco, “SemFedXAI: A semantic framework for explainable federated learning in healthcare,” *Information*, vol. 16, no. 6, art. 435, May 2025. doi:10.3390/info16060435.
- [6] R. Alkhanbouli, H. M. A. Almadhaani, F. Alhosani, and M. C. E. Simsekler, “The role of explainable artificial intelligence in disease prediction: A systematic literature review and future research directions,” *BMC Med. Inform. Decis. Mak.*, vol. 25, no. 1, art. 110, Mar. 2025. doi:10.1186/s12911-025-02944-6.
- [7] R. Eden, A. Smith, and colleagues, “A scoping review of the governance of federated learning in healthcare,” *NPJ Digit. Med.*, vol. 8, art. 427, Jul. 2025. doi:10.1038/s41746-025-01836-3.
- [8] “Interplay between Federated Learning and Explainable AI: A scoping review,” arXiv:2411.05874, Apr. 2025. arXiv
- [9] A. Rieke, M. Hancox, A. Li, et al., “The future of digital health with federated learning,” *NPJ Digit. Med.*, vol. 3, art. 119, Nov. 2020 — cited as foundational background for FL deployments. doi:10.1038/s41746-020-00323-1.
- [10] S. Shukla, S. Rajkumar, A. Sinha, M. Esha, K. Elango, and V. Sampath, “Federated learning with differential privacy for breast cancer diagnosis: Enabling secure data sharing and model integrity,” *Sci. Rep.*, vol. 15, art. 13061, Apr. 2025. doi:10.1038/s41598-025-95858-2.
- [11] A. M. Al-Hejri, et al., “A hybrid explainable federated-based vision transformer framework for breast cancer prediction via risk factors,” *Sci. Rep.*, vol. 15, May 2025. doi:10.1038/s41598-025-96527-0.
- [12] S. Pati and colleagues, “Privacy preservation for federated learning in health care: Methods and open challenges,” *Comput. Secur.*, vol. 140, p. 103801, 2024. doi:10.1016/j.cose.2024.103801.
- [13] V. Pais, “Healthcare federated learning: A survey of applications and cross-silo implementations,” Taylor & Francis, 2025. (book/chapter)
- [14] H. Nasajpour, et al., “Federated learning in public health: A systematic review,” *Healthcare*, vol. 13, art. 2760, 2025. doi:10.3390/healthcare13212760.
- [15] “Federated Learning and Explainable AI in Healthcare,” in *Explainable AI in Healthcare* (book chapter), 2023. doi:10.1201/9781003333425-14.
- [16] M. Saarela and colleagues, “Recent applications of explainable AI (XAI) in healthcare: A systematic review,” *Appl. Sci.*, 2024. (open access survey)
- [17] Z. Sadeghi, “A review of Explainable Artificial Intelligence in healthcare,” *Comput. Electr. Eng.*, 2024.
- [18] J. Thrasher, “Multimodal federated learning in healthcare: A review,” arXiv:2310.09650, Oct. 2023.
- [19] R. Poulain, et al., “Improving fairness in AI models on electronic health records: The case for federated learning methods,” arXiv:2305.11386, 2023.
- [20] “FED-EHR: A privacy-preserving federated learning framework for EHR analytics,” *Electronics (MDPI)*, vol. 14, no. 16, art. 3261, 2025. doi:10.3390/electronics14163261.
- [21] C. Gupta, et al., “Applying YOLOv6 as an ensemble federated learning method for breast cancer detection with homomorphic encryption,” *Sci. Rep.*, 2025. doi:10.1038/s41598-024-80187-7.
- [22] N. Tahir, et al., “A federated learning-based model for predicting mortality using distributed EHR,” *J. Med. Internet Res.*, 2025. (online)
- [23] S. Sharma and N. Marchang, “A review on client-server attacks and defenses in federated learning,” *Comput. Secur.*, vol. 140, p. 103801, 2024. doi:10.1016/j.cose.2024.103801.
- [24] E. Daniel, et al., “FLEM-XAI: Federated learning based real time ensemble model with XAI for medical diagnosis,” *Front. Comput. Sci.*, 2025.
- [25] S. R. Abbas, “Federated Learning in Smart Healthcare: Applications, challenges and future directions,” review, 2024 (MDPI Healthcare — comprehensive survey). doi:10.3390/healthcare12242587