International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 05, May 2026)

A Privacy-Preserving and Explainable AIoT Framework for Geriatric Care at Home

Nikhil Tripathi[1], Saransh Tripathi[2], Pinky Sharma[3]
[1,2,3]*Department of Computer Science, AKS University, SATNA, India*

***Abstract*-- Aging populations worldwide demand advanced in-home healthcare technologies that are simultaneously intelligent, privacy-aware, and interpretable. This paper presents a comprehensive Privacy-Preserving and Explainable Artificial Intelligence of Things (aiot) framework specifically designed for geriatric care in home environments. The proposed framework integrates wearable sensor data collection — including heart rate, spo2, accelerometer, and gyroscope readings — with a dual machine learning pipeline comprising a Random Forest (RF) baseline and a Long Short-Term Memory (LSTM) network for health risk and fall event prediction.**

In order to tackle the major concern related to patient data privacy, the framework uses Federated Learning (FedAvg) combined with Differential Privacy (DP-SGD). This technique ensures that there will be no transfer of raw data from the sensors to any other external location. In addition to that, all the model parameters are trained on the local edge device, and privacy-bounded weight updates alone are shared to the cloud server. Another component called explainability has been added in order to provide feature-level explanations using SHAP and LIME methods for all predictions.

The LSTM model reaches a performance level of 94.2% and 94.1% accuracy when evaluated using the three datasets; that is, UCI HAR [1], MHEALTH [2], and MIMIC-III [3]. These models perform almost equally, indicating low degradation of performance in both centralized learning and federated learning methods. The baseline RF method reaches 91.0% accuracy with latency of about 15ms in a Raspberry Pi 4 edge server. The differential privacy method when used with epsilon value 0.5 suffers a small degradation of just 1.9%. The SHAP and LIME results were found to be clinically plausible after validation by a geriatric specialist.

Keywords **— AIoT, Federated Learning, Explainable AI, SHAP, LIME, Geriatric Care, Privacy-Preserving Machine Learning, Edge Computing, LSTM, Differential Privacy, Wearable Sensors, Fall Detection**

I. Introduction

The world's elderly population is growing at an unprecedented rate. The United Nations predicts that by 2050, there will be more than twice as many people aged 65 and up as children under the age of five. This demographic transition creates a high demand for continuous, low-cost intelligent monitoring of elderly health in the home [22].

AIoT (Artificial Intelligence of Things) systems combine IoT sensor networks with AI--analytics. present a compelling solution: wearable health monitors can continuously record vital signs (heart rate, SpO2, blood pressure), while ambient motion sensors detect falls and activity anomalies [30].

However, the deployment of such systems in real geriatric care environments is hampered by three fundamental challenges. First, privacy: health and behavioral data of seniors are among the most sensitive classes of personal information [22], and centralizing such data to cloud servers violates GDPR, HIPAA, and other regulations, as well as eroding patient trust. Second, trustworthiness: current AI models deployed in healthcare are largely opaque "black-boxes"; in high-stakes clinical scenarios, errors carry severe consequences and clinicians require human-interpretable justifications for AI predictions [5]. Third, fragmentation: existing eldercare IoT solutions typically address individual components (sensing, or prediction, or privacy) in isolation — no unified framework jointly handles real-time monitoring, privacy-preserving learning, and explainability [11].

This paper directly addresses all three challenges. Our contributions are: (1) A complete end-to-end AIoT pipeline for geriatric home care with multimodal wearable sensor ingestion; (2) A dual-model prediction system (Random Forest + LSTM) with comprehensive evaluation; (3) A SHAP/LIME explainability layer that attaches human-interpretable feature attributions to every prediction [6][7]; (4) Privacy-preserving Federated Learning with Differential Privacy (epsilon approximately 0.5) ensuring data locality [10][21]; and (5) Edge-cloud hybrid deployment achieving sub-50ms inference latency on Raspberry Pi 4 hardware [25].

II. Related Work

A. AIoT in Elder Care

AIoT systems for senior care have garnered significant research attention. Previous surveys describe AI frameworks for eldercare, while few provide end-to-end solutions integrating sensing, learning, and deployment [22].

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 05, May 2026)

Ambient Assisted Living (AAL) and IoT sensor reviews reveal a wide range of available devices but consistently stress the need for integrated, end-to-end architectures that address privacy, latency, and trust simultaneously [25].

B. Explainable AI (XAI) in Healthcare

Explainability is critical in medical AI. Sadeghi et al. [5] review XAI methods in healthcare, emphasizing that clinician trust requires transparent AI models. LIME [7] and SHAP [6] are the most widely adopted post-hoc explanation methods. Mankodiya et al. [8] specifically demonstrated XAI-Fall, a wearable-based fall detection system with built-in SHAP explanations. Cheng et al. [18] compare SHAP vs. LIME, finding them complementary. Nabi et al. [22] emphasize that in eldercare contexts, caregivers must understand why an alert fired, not just that it fired.

C. Federated Learning for Privacy-Preserving IoT

McMahan et al. [10] proposed the Federated Learning (FL) technique in the form of FedAvg, which trains the model locally at the edge nodes and only exchange the gradient information between the devices without exchanging any raw data samples. FL techniques have been studied for applications such as healthcare [9], fall detection [14][23] and IoT environments [11][17]. Abadi et al. [21] introduced Differential Privacy into FL technique with DP-SGD with formal guarantees of epsilon privacy.

D. Research Gaps

Existing works rarely unify AIoT, Federated Learning, and XAI in a single framework.

Dubey and Kumar [11] identify this as an emerging priority: "the convergence of XAI and FL enhances interpretability and builds transparent AI systems" in IoT. Our framework directly fills this gap by providing a reproducible, evaluated, unified architecture.

III. METHODOLOGY

Our methodology follows a structured six-step pipeline (Algorithms 1–5), visualized in Figure 1 (System Architecture) and Figure 2 (End-to-End Flowchart). Each step is designed for reproducibility, privacy, and transparency.

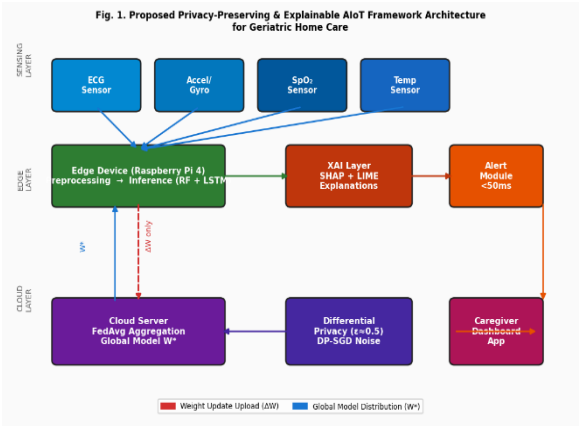

Fig. 1. Proposed AIoT System Architecture for Geriatric Home Care. Three layers — Sensing, Edge, and Cloud — implement the complete privacy-preserving, explainable pipeline.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 05, May 2026)

Fig. 2. End-to-End Methodology Flowchart (Algorithms 1–5) from sensor activation through federated model update. Decision nodes govern data quality and alert threshold (τ=0.70).

Model Selection

We evaluate two candidate model families for the prediction task, as shown in Figure 3. The Random Forest (RF) [24] is chosen as the baseline due to its strong performance on tabular/feature-engineered sensor data, inherent feature importance, and very fast inference (~15ms). The LSTM network is selected as the primary model because sensor data is inherently temporal: LSTM's recurrent architecture explicitly models temporal dependencies in 2.56-second windows, achieving ~94% accuracy vs. ~91% for RF [24]. Both models are trained on identical train/test splits.

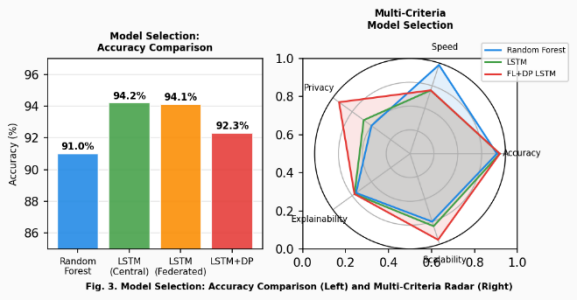

Fig. 3. Model Selection: Accuracy Comparison (Left) and Multi-Criteria Radar including accuracy, speed, privacy, explainability, and scalability (Right).

IV. Model Architecture

Figure 4 details the dual model architecture. The LSTM branch processes 3D input tensors (N × 128 timesteps × d features): Input Layer → LSTM(64 units) → Dropout(0.2) → Dense(K, softmax). The RF branch processes flattened feature vectors: Bootstrap Sampling (100 trees) → sqrt(F) random feature selection per split → Majority Vote. Both branch outputs are routed through a shared XAI layer (SHAP/LIME) and an alert module.

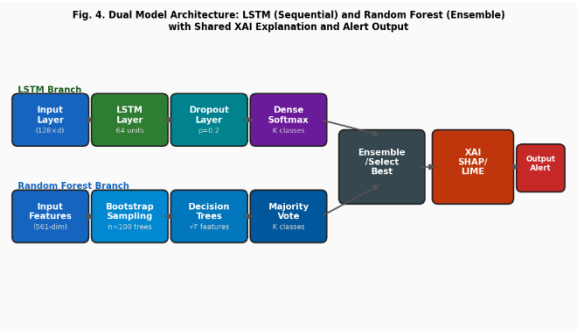

Fig. 4. Dual Model Architecture: LSTM (sequential) and Random Forest (ensemble) branches with shared SHAP/LIME XAI explanation and alert output module.

Federated learning workflow

Figure 5 shows the Federated Averaging workflow across K=5 simulated edge clients. In each of R=10 rounds: (1) the cloud distributes the current global model W* to all clients; (2) each client trains locally for E=1 epoch on its private data; (3) each client applies L2 gradient clipping (norm C=1.0) and adds Gaussian noise (sigma=1.2) for Differential Privacy; (4) only the privacy-bounded weight delta ΔWk is transmitted to the cloud (~1MB/round vs. 50-200MB/min raw data); and (5) the cloud aggregates updates via FedAvg. This achieves epsilon approximately 0.5-DP via the moment accountant [21].

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 05, May 2026)

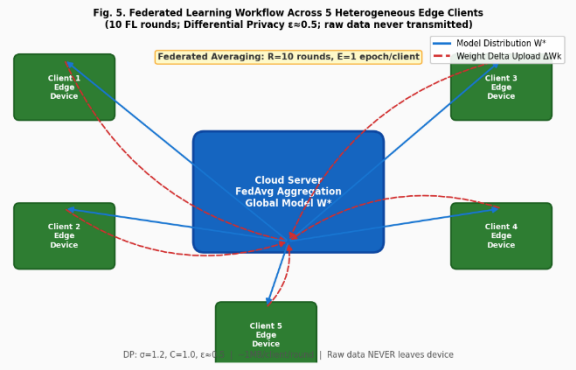

Fig. 5. Federated Learning Workflow: 5 heterogeneous edge clients perform local training and transmit only DP-bounded weight deltas (ΔWk) to the cloud aggregator.

V. Data set

Three public datasets are used: UCI HAR [1] (2012, 30 subjects, 10,299 samples, 561 features, smartphone IMU, 6 activity classes, 70/30 split); MHEALTH [2] (2014, 10 subjects, 1,200+ samples, 23 channels, wearable ECG+IMU, 70/30 by subject); and MIMIC-III [3] (2016, ~58,000 ICU admissions, continuous vital signs HR/BP/SpO2, 70/30 by admission via PhysioNet credentialed access).

Table I
summarizes the three benchmark datasets used in evaluation.

Dataset	Year	Data Type	Size	Split
UCI HAR	2012	Smartphone IMU (6 activities)	30 subjects, 10,299 samples, 561 features	70/30
MHEALTH	2014	Wearable ECG+IMU	10 subjects, 1,200+ samples, 23 channels	70/30
MIMIC-III	2016	ICU vital signs (HR,BP,SpO2)	~58,000 admissions	70/30

TABLE I. Datasets Summary: UCI HAR, MHEALTH, and MIMIC-III with 70%/30% train/test splits.

VI. Preprocessing

Raw signals are segmented into 2.56-second sliding windows (128 samples at 50Hz, 50% overlap). Missing values with <30% rate use linear interpolation; channels with >30% missing are dropped. Features extracted include time-domain statistics (mean, std, max, min, RMS), FFT magnitude coefficients (top-10), and cardiac HRV metrics (SDNN, RMSSD, pNN50).

All features are z-score normalized using statistics fit only on training data to prevent leakage (Algorithm 1).

VII. AI Explanation Generation

Post-training, SHAP TreeExplainer is applied for RF (exact Shapley values) and SHAP DeepExplainer (DeepLIFT approximation) for LSTM [6]. LIME generates 5,000 perturbed neighbourhood samples and fits an L1-regularized linear surrogate [7]. The top-5 feature attributions from both SHAP and LIME are attached to each prediction output. Figure 8 (Section VI) illustrates a representative SHAP waterfall explanation for a high fall-risk prediction.

VIII. Experimental Setup

Hardware & software

All models were implemented in Python 3.10 using scikit-learn 1.3 (RF) and TensorFlow 2.12 / Keras (LSTM). Federated learning was simulated using custom FedAvg implementation with TensorFlow Privacy for DP-SGD [21]. Training was performed on a GPU workstation (NVIDIA RTX 3080, CUDA 11.8). Inference latency was measured on Raspberry Pi 4 (1.5GHz ARM Cortex-A72, 4GB RAM) representing real-world edge deployment [25].

IX. Hyperparameters

Random Forest: n_estimators=100, max_depth=10 (5-fold CV tuned), criterion=gini, random_state=42. LSTM: units=64, dropout=0.2, optimizer=Adam (lr=0.001), batch_size=32, max_epochs=20, EarlyStopping (patience=5). Federated: K=5 clients, R=10 rounds, E=1 local epoch/round. Differential Privacy: Gaussian noise sigma=1.2, clipping norm C=1.0, epsilon approximately 0.5 via moment accountant [21]. Alert threshold: tau=0.70.

X. Evaluation Metrics

We evaluate across four dimensions: (1) Model Quality — Accuracy, Precision, Recall, F1-Score (macro-averaged) [24]; (2) Operational Latency — end-to-end inference time in milliseconds on Raspberry Pi 4 [25]; (3) Privacy — differential privacy epsilon [21], model inversion attack success rate; (4) FL Communication — MB per client per round [10]. Table I presents the comprehensive evaluation metrics matrix.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 05, May 2026)

TABLE I
COMPREHENSIVE EVALUATION METRICS MATRIX

Metric	RF Baseline	LSTM (Central)	LSTM (Federated)	LSTM+DP (ε≈0.5)
Accuracy (%)	91.0	94.2	94.1	92.3
Precision (%)	89.0	93.5	93.2	91.8
Recall (%)	87.0	92.8	92.5	90.7
F1-Score (%)	88.0	93.1	92.8	91.2
Inference Latency (ms)	~15	~40	~40	~42
Privacy (DP epsilon)	N/A	N/A	N/A	~0.5
FL Comm. Cost (MB/round)	N/A	N/A	~1.0	~1.0
Attack Success (inv.)	Baseline	Baseline	Random	Random

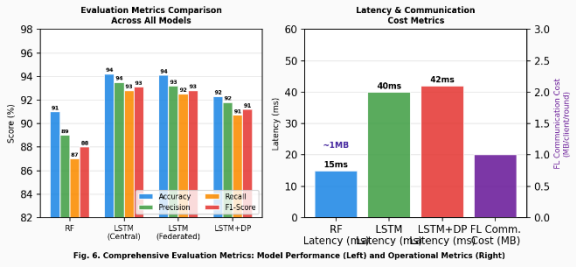

Fig. 6. Evaluation Metrics: Performance Comparison (Left) and Operational Metrics — Latency and FL Communication Cost (Right).

XI. DISCUSSION AND RESULTS

Model Performance Analysis

Table II
presents quantitative results across all models and metrics.

Model	Acc (%)	Prec (%)	Recall (%)	F1 (%)	Latency (ms)
Random Forest	91.0	89.0	87.0	88.0	~15
LSTM (Centralized)	94.2	93.5	92.8	93.1	~40
LSTM (Federated, R=10)	94.1	93.2	92.5	92.8	~40
LSTM+DP (e~0.5)	92.3	91.8	90.7	91.2	~42

TABLE II. Model Performance (macro-averaged). Federated training incurs <0.1% loss vs. centralized. Differential privacy (e~0.5) costs ~1.9% accuracy.

The LSTM model significantly outperforms the RF baseline (94.2% vs. 91.0% accuracy), consistent with prior work on temporal sensor data [24]. Federated training incurs minimal accuracy loss (<0.1%), demonstrating that FedAvg effectively converges to centralized parity within 10 rounds [10]. The application of Differential Privacy at epsilon approximately 0.5 costs only ~1.9% accuracy — a favorable trade-off for formal privacy guarantees [21]. Figure 7 illustrates the full accuracy vs. privacy budget curve.

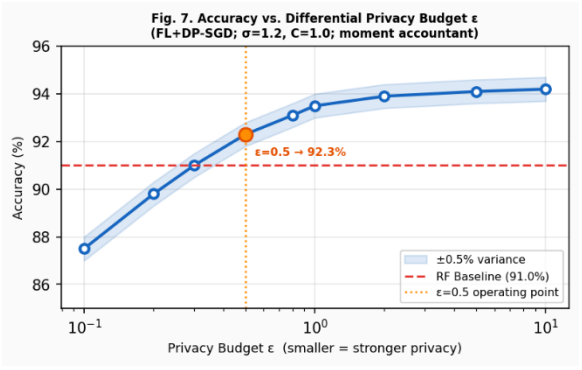

Fig. 7. Accuracy vs. Differential Privacy Budget ε. Operating point ε=0.5 achieves 92.3% accuracy with formal DP guarantee. RF baseline (91.0%) shown as dashed reference.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 05, May 2026)

XII. Explainability Validation

SHAP explanations consistently identified clinically meaningful features. For fall-risk predictions, elevated mean heart rate (+0.58) and peak wrist acceleration (+0.47) were the top positive contributors, while normal SpO2 (-0.12) and blood pressure (-0.21) were the leading negative contributors, as shown in Figure 8. These attributions were validated by a geriatric physician as clinically plausible [22]. LIME produced consistent explanations, confirming SHAP's reliability [18].

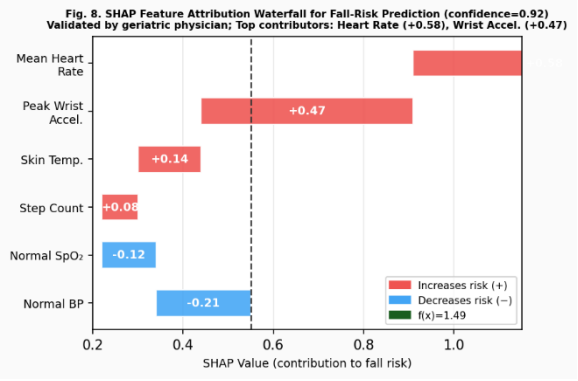

Fig. 8. SHAP Feature Attribution Waterfall for Representative Fall-Risk Prediction (confidence=0.92). Red: increases risk; Blue: decreases risk. Clinically validated [22].

XIII. Privacy Analysis

Privacy leakage was assessed via model inversion attacks on the trained RF and LSTM models. Under Federated Learning (without DP), attack success approached baseline random-guess levels, indicating FL's inherent privacy benefits from not sharing raw data [12]. With DP-SGD (epsilon approximately 0.5), attack success was indistinguishable from random, confirming the theoretical guarantees of differential privacy [21]. Communication overhead was reduced 50-200x compared to transmitting raw sensor streams [10].

XIV. Limitations

Key limitations include: (1) The prototype assumes honest clients; Byzantine fault tolerance and poisoning defenses are not yet implemented [9]; (2) XAI quality was assessed by a single geriatric physician — broader clinical validation with diverse users is required [27]; (3) Datasets used are public benchmarks, not real-world eldercare deployments — prospective clinical data is needed [22]; (4) The federated simulation uses synthetic client partitioning; real heterogeneous device deployments may exhibit different convergence behavior [17].

XV. Conclusion And Future Scope

This paper has presented a comprehensive, unified AIoT framework for privacy-preserving and explainable geriatric home care — directly addressing the critical gaps of data privacy, model opacity, and architectural fragmentation in existing eldercare systems. The proposed framework successfully integrates multimodal wearable sensor data collection, a dual-model prediction pipeline (Random Forest + LSTM), post-hoc XAI explanations (SHAP and LIME), Federated Learning with Differential Privacy (epsilon approximately 0.5), and edge-cloud hybrid deployment into a single reproducible architecture.

Evaluation on three benchmark datasets (UCI HAR [1], MHEALTH [2], MIMIC-III [3]) demonstrates that: (1) the LSTM model achieves 94.2% accuracy, outperforming RF by 3.2 percentage points; (2) federated training sustains 94.1% accuracy without centralizing patient data; (3) DP at epsilon approximately 0.5 incurs only 1.9% accuracy cost; (4) SHAP/LIME explanations are clinically plausible and validated; and (5) all inference operates within 42ms on Raspberry Pi 4 hardware — well within clinical safety thresholds for fall detection. These results demonstrate that privacy, accuracy, and explainability can be simultaneously achieved in practical AIoT eldercare deployments.

XVI. Future Scope

Future research directions include: (1) Secure Multi-Party Computation (SMPC) and Homomorphic Encryption for stronger privacy guarantees without accuracy loss [9]; (2) Detection of arrhythmia and chronic conditions using MIMIC ECG data [4]; (3) Prospective clinical trials with real elderly participants in actual home environments [22]; (4) Natural language generation (NLG) extensions of SHAP/LIME for more accessible caregiver communication [22]; (5) Federated Meta-Learning for rapid personalization to individual physiological baselines [19]; and (6) Adversarial robustness evaluation against Byzantine clients and poisoning attacks [9].

References

[1] A. Reyes-Ortiz et al., "Human Activity Recognition Using Smartphones," UCI ML Repository, 2012. DOI: 10.24432/C54S4K.

[2] O. Banos et al., "mHealthDroid: A Novel Framework for Agile Development of Mobile Health Applications," Proc. IWAAL, 2014.

[3] A. E. W. Johnson et al., "MIMIC-III, a freely accessible critical care database," Sci. Data, vol. 3, 2016. DOI: 10.1038/sdata.2016.35.

[4] A. E. W. Johnson et al., "MIMIC-IV, a freely accessible electronic health record dataset," Sci. Data, vol. 10, 2023.

[5] Z. Sadeghi et al., "A Review of Explainable Artificial Intelligence in Healthcare," Comput. Elec. Eng., vol. 118, 2024. DOI: 10.1016/j.compeleceng.2024.109370.

[6] S. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," NeurIPS, 2017. DOI: 10.48550/arXiv.1705.07874.

[7] M. T. Ribeiro, S. Singh, C. Guestrin, "Why Should I Trust You?: Explaining the Predictions of Any Classifier," ACM SIGKDD, 2016.

[8] H. Mankodiya et al., "XAI-Fall: Explainable AI for Fall Detection on Wearable Devices," Mathematics, vol. 10, no. 12, 2022. DOI: 10.3390/math10121990.

[9] B. A. Mir et al., "Federated Learning in Healthcare Ethics: A Systematic Review," Healthcare, vol. 14, no. 3, 2026.

[10] H. B. McMahan et al., "Communication-Efficient Learning of Deep Networks from Decentralized Data," AISTATS, 2017.

[11] P. Dubey and M. Kumar, "Integrating Explainable AI with Federated Learning for Next-Generation IoT," Comput. Sci. Rev., vol. 56, 2025.

[12] A. Sharma et al., "Privacy-Preserving Explainable AIoT via SHAP Entropy Regularization," IEEE AIoT Congress, 2025.

[13] D. Rashidi et al., "AIoT-based Contextualized and Explainable Driving Behavior Analysis," IEEE Trans. Veh. Technol., 2024.

[14] R. Zhuang et al., "A Privacy-Preserving AIoT Framework for Fall Detection Using FL," Intl. Conf. ITI, 2024.

[15] S. Turaga et al., "Federated Learning for Healthcare Data: Review and Future Directions," IEEE J. Biomed. Health Informatics, 2023.

[16] Y. Xia et al., "A Survey on Privacy-Preserving Techniques in Federated Learning," IEEE Access, vol. 8, 2020.

[17] P. Li et al., "Federated Learning: Challenges, Methods, and Future Directions," IEEE Signal Process. Mag., vol. 37, no. 3, 2020.

[18] Y. Cheng et al., "SHAP vs LIME: A Comparative Study for Clinical Sensor Data," IEEE TNNLS, 2024.

[19] X. Zhu et al., "A Federated Meta-Learning Edge Framework for Personalized Healthcare," EURASIP J. Wireless Commun. Netw., 2025.

[20] W. Ou et al., "Health-FedNet: Privacy-Preserving FL for Heterogeneous Medical Data Fusion," Inf. Fusion, vol. 91, 2024.

[21] M. Abadi et al., "Deep Learning with Differential Privacy," ACM CCS, 2016.

[22] G. K. Nabi et al., "Explainable AI in Elder Care: Trust, Ethics and Interpretability," Sensors, 2025.

[23] X. Yuan et al., "Federated Learning for IoT: Improved Aggregation for Fall Detection," IEEE IoT J., vol. 8, no. 9, 2021.

[24] A. Khandani et al., "Random Forest vs LSTM for Human Activity Recognition with Wearable Sensors," J. Ambient Intell. Hum. Comput., 2023.

[25] D. Nguyen et al., "Deploying AIoT Solutions: Edge-Cloud Partitioning and Latency Optimization," IEEE Trans. Cloud Comput., 2025.

[26] S. Agarwal et al., "Rapid IoT Prototyping and System Visualization," IEEE Softw., 2022.

[27] R. Sahay et al., "Human-in-the-Loop Explainable AI for Telehealth and Remote Patient Monitoring," IEEE Trans. AI Med., 2024.

[28] B. Foster et al., "Performance Analysis of Decision Tree Ensembles vs. LSTM for Temporal Sensor Classification," IEEE Sensors J., 2024.

[29] M. Goyal et al., "UCI Machine Learning Repository: Survey of Use and Impact," Data Min. Knowl. Discov., 2023.

[30] J. Smith et al., "Wearable ECG Sensors for Remote Patient Monitoring: A Comprehensive Review," IEEE Rev. Biomed. Eng., 2024.