

# AI-Powered Smart Healthcare Monitoring System Using Federated Learning: A Privacy-Preserving Approach for Early Disease Detection

Anmol Budhewar<sup>1</sup>, Soham Kute<sup>2</sup>, Sahil Aher<sup>3</sup>, Jivan Kawde<sup>4</sup>, Vishal Yadav<sup>5</sup>

<sup>1</sup>Assistant Professor, Department of Computer Engineering

<sup>2,3,4,5</sup>Student, Department of Computer Engineering

Sandip Institute of Technology and Research Centre, Nashik, Maharashtra, India

E-mail: anmolsbudhewar@gmail.com, sohamkute72@gmail.com, ahersahil2005@gmail.com, jivankawde7@gmail.com, ydvishal89@gmail.com

\*\*\*

**Abstract** - Healthcare artificial intelligence is still commonly built by pooling patient data from every hospital on a single server and training a model there. This approach delivers accuracy but creates a serious privacy risk, and it conflicts with regulations such as HIPAA, GDPR and India's Digital Personal Data Protection Act, which is a major reason why hospitals are reluctant to release raw records [1], [2]. This paper proposes a system that avoids centralization instead of working within it. Wearable devices at each hospital node continuously collect physiological signals (heart rate, ECG, SpO<sub>2</sub> and blood pressure), and a Long Short-Term Memory (LSTM) model [3] is trained on these data locally and on-site. Only an encrypted set of model weights leaves the hospital; a central server combines the weights using Federated Averaging (FedAvg) [4] and returns an improved global model. Over 20 communication rounds, the global model reaches 97.6% accuracy and a 97.2% F1-score for early disease prediction over five categories, which is within about half a percentage point of a centralized model trained with full access to the same pooled data. Clinicians interact with the system through a Streamlit dashboard that shows prediction confidence, patient trends and alerts without exposing another hospital's raw records.

**Key Words:** federated learning, LSTM, healthcare monitoring, disease prediction, FedAvg, wearable IoT, privacy-preserving AI.

## 1. INTRODUCTION

Healthcare has changed considerably over the last decade, largely because of the volume of data it now produces. Hospitals, diagnostic laboratories, wearables and smart medical equipment generate health data continuously: smartwatches track heart rate, oximeters log blood oxygen, glucose monitors record sugar levels, and hospital systems maintain electronic health records (EHRs). This scale of data did not exist ten years ago, and it has made continuous monitoring possible in place of waiting for a patient's next clinic visit [5], [6]. The corresponding difficulty is that the data must be stored, secured and analysed properly, which is proving harder than collecting it [7].

Deep learning is one of the more promising tools for making sense of such data. LeCun, Bengio and Hinton [8] explained why these models generalize as well as they do, and healthcare has become one of the fields in which that generalization pays off. Rajkomar et al. [9] and Topol [10] describe how machine learning is beginning to change clinical practice and not only research publications. Models trained on medical records can flag early signs of heart disease, diabetes, respiratory problems and other chronic conditions, sometimes detecting patterns that a clinician would not notice during a routine check [11].

The way in which most of these models are trained, however, is the main obstacle. The standard approach collects patient data from several hospitals on one central server and trains there. This is the simplest way to build a model, but medical records are among the most sensitive data that exist (identities, diagnoses, laboratory results, prescriptions and scans), and storing all of them in one place means that a single breach becomes catastrophic instead of contained [2].

Regulators have responded accordingly. HIPAA, GDPR and India's Digital Personal Data Protection Act all place strict limits on how patient data may be stored, moved and shared. As a result, many hospitals decline to hand over raw datasets to outside researchers or partner institutions, even where pooling the data could improve diagnostic accuracy for every participant [1].

Federated Learning (FL) addresses this problem by reversing the usual arrangement. Instead of moving the data to the model, it moves the model to the data: each hospital trains locally on its own patients and sends only the learned parameters, never the underlying records, to an aggregation server. The server merges the updates from all participants into a stronger global model and distributes it for the next round of local training, so the raw data never leaves the institution where it originated [12], [4], [13]. The approach is not only theoretical: FL models have reached accuracy close to that of centralized training on tasks such as ECG classification, diabetic retinopathy detection and cancer diagnosis, while offering privacy guarantees that centralized training cannot match [11], [12], [14].

Wearable Internet of Things (IoT) devices have meanwhile become common in both hospitals and homes, continuously streaming heart rate, ECG, SpO<sub>2</sub>, respiratory rate and activity data [15], [6]. Such streams allow abnormal readings to be detected well before they become emergencies. They also send more sensitive data to cloud servers, which reintroduces the privacy problem that FL is intended to solve [16].

The system presented here combines wearable sensor data, EHRs and an LSTM-based prediction model [3] within a Federated Learning framework. Each participating hospital trains locally, and FedAvg [4] builds a shared global model without centralizing patient data. A Streamlit dashboard allows doctors to view patient status, prediction confidence and alerts in real time. The aim is a system that is accurate enough to be clinically useful and private enough for a hospital to agree to use it.

The work makes three contributions. First, it applies Federated Averaging to a multi-institution LSTM pipeline covering cardiovascular, diabetic, respiratory and hypertensive conditions, using wearable and EHR data together and not either source alone [17], [18]. Second, it measures the accuracy trade-off between federated and centralized training round by round, instead of assuming that federated training eventually converges to the centralized result [4], [19]. Third, it keeps a human in the loop:

the dashboard is designed so that the clinician, and not the model, makes the final decision [10].

The remainder of the paper is organized as follows. Section 2 reviews related work, Section 3 describes the proposed framework and methods, Section 4 gives the mathematical formulation, Section 5 reports the experimental configuration, performance and discussion, and Section 6 presents the conclusions and future work.

**Table -1:** Summary of the reviewed literature

| Ref. | Study                  | Method                                 | Main contribution                                                                                    | Limitation                                              |
|------|------------------------|----------------------------------------|------------------------------------------------------------------------------------------------------|---------------------------------------------------------|
| [20] | Rieke et al. (2020)    | FL for digital health                  | Showed that FL is suitable for healthcare because hospitals cannot freely share raw patient data.    | Mostly conceptual; limited experiments.                 |
| [21] | Sheller et al. (2020)  | FL for brain tumor segmentation        | Achieved performance close to centralized learning without sharing medical data.                     | Focused only on medical imaging.                        |
| [22] | Kaissis et al. (2020)  | Privacy-preserving FL                  | Combined FL with privacy protection for medical data.                                                | Secure computation increases computational cost.        |
| [23] | Nguyen et al. (2021)   | FL and IoT survey                      | Explained the growing integration of FL and IoT, including healthcare.                               | Healthcare-specific details were limited.               |
| [18] | Ghosh et al. (2022)    | IoT-based smart healthcare             | Found that most existing healthcare monitoring systems were centralized.                             | Did not develop a federated solution.                   |
| [2]  | Abaoud et al. (2023)   | Privacy-enhanced FedAvg                | Improved protection against privacy attacks in healthcare FL.                                        | Additional privacy layer causes computational overhead. |
| [17] | Thrasher et al. (2023) | Multimodal FL                          | Demonstrated the feasibility of combining imaging, time-series and EHR data.                         | Data-fusion methods are still developing.               |
| [24] | Li et al. (2024)       | FL deployment review                   | Identified practical issues such as device synchronization, dropped rounds and uneven participation. | Recommendations were not experimentally validated.      |
| [15] | Rghioui et al. (2020)  | Wearable sensing with machine learning | Developed an end-to-end wearable monitoring system for diabetic patients.                            | Centralized approach; no privacy-preserving FL.         |

Three observations follow from the review. First, Federated Learning has emerged as a promising privacy-preserving approach for distributed healthcare systems [1], [25], while deep learning models such as LSTM have shown strong performance in disease prediction from physiological data [3], [9]. Second, existing studies concentrate mainly on prediction accuracy, and comparatively few integrate real-time IoT monitoring, privacy preservation and continuous health-risk assessment within a single deployable framework [18], [15]. Third, although wearable devices enable continuous patient monitoring, data privacy, communication efficiency, model personalization and real-time emergency alert generation remain insufficiently addressed [24], [26].

These gaps motivate the proposed AI-powered Smart Healthcare Monitoring System, which combines IoT-based data collection, LSTM-based health prediction, Federated Learning for privacy-preserving model training and a real-time monitoring dashboard to provide secure and accurate healthcare support [20], [23].

## 3. PROPOSED FRAMEWORK AND METHODS

### 3.1 System Architecture

Fig. 1 shows the overall architecture, which follows the federated learning paradigm [4], [13]. Each hospital forms a Local Edge Layer: wearables stream data to an Edge Client Device, which trains an LSTM model on-site. Only the encrypted weights from this training leave the hospital, travelling to the Central Cloud Aggregator, which runs FedAvg and produces an updated global model. The global model is then sent to every node for the next round. A Streamlit dashboard operates on top of this pipeline, and patient records, predictions and analyst actions are logged to a local SQLite/PostgreSQL database that never synchronizes raw data outside the hospital [20].

## 2. LITERATURE REVIEW

Representative studies published between 2020 and 2026 on AI-based healthcare monitoring, Federated Learning, IoT-enabled wearable devices, LSTM-based disease prediction, privacy-preserving machine learning and real-time health monitoring were collected from Google Scholar. The studies were analysed to identify recent advances and existing research gaps and to guide the design of the proposed system. Table 1 summarizes the representative works reviewed.




**Fig -1:** System architecture of the AI-powered smart healthcare monitoring framework.

### 3.2 Dataset and Pre-processing

For model development, a mix of public benchmark datasets was used: MIMIC-IV [27], PhysioNet [28], the UCI Heart Disease dataset [29] and WESAD [30]. These were split across simulated hospital nodes to approximate a realistic non-IID federated setting [19], and were complemented by wearable streams covering heart rate, ECG, SpO<sub>2</sub>, body temperature, blood pressure, respiratory rate and physical activity [15], [6]. Table 2 lists the feature set used at each node.


Before training, missing values are imputed statistically and duplicates are removed. An outlier check flags sensor readings that resemble hardware faults and not genuine physiological extremes. Categorical fields (gender and medical-history category) are one-hot encoded, and continuous features are scaled with Min-Max normalization computed separately at each local node, using only that node's own training split, so that no information about another hospital's data range is leaked through the scaling step [19]. The data at each node are split 70-15-15 into training, validation and test sets.

**Table -2:** Physiological features used per local node

| Feature           | Typical range      | Source         |
|-------------------|--------------------|----------------|
| Heart Rate (HR)   | 50–180 bpm         | Wearable       |
| ECG               | Waveform           | Wearable       |
| SpO <sub>2</sub>  | 80–100 %           | Wearable       |
| Body Temperature  | 35–40 °C           | Wearable       |
| Blood Pressure    | 90/60–180/120 mmHg | Wearable / EHR |
| Respiratory Rate  | 10–30 breaths/min  | Wearable       |
| Glucose Level     | 70–300 mg/dL       | EHR            |
| Physical Activity | Steps / accel.     | Wearable       |


### 3.3 Federated Learning Framework (FedAvg)

Every hospital node trains the same LSTM architecture on its own local data, which never leaves the institution. After a fixed number of local epochs, each node sends only its updated weights to the Central Cloud Aggregator, which averages them with weights proportional to the amount of local data used at each node; this is Federated Averaging [4]. Restricting communication to model updates also limits bandwidth requirements [26]. The resulting global model is redistributed to every node for the next round, as illustrated in Fig. 2. The experiments used 20 rounds, 40 local epochs per round and 8 simulated hospital nodes with deliberately uneven, non-IID patient populations [19], since this is closer to what a real deployment would encounter.


**Fig -2:** Federated Averaging (FedAvg) workflow across hospital nodes.

### 3.4 LSTM Disease Prediction Model

Each local model is an LSTM network, the architecture proposed by Hochreiter and Schmidhuber [3] to overcome the vanishing-gradient problem that affects plain recurrent networks on longer sequences. It was chosen because physiological data are inherently a time series, and the gating mechanism retains long-range temporal patterns in a way that a plain feed-forward network cannot. The network takes 8 normalized physiological features as input and passes them through an LSTM layer with 128 units, dropout at 0.3, a second LSTM layer with 64 units, a further dropout layer and a dense layer with 32 ReLU units, followed by a softmax output over 5 classes: Normal, Cardiovascular Risk, Diabetes, Respiratory Disorder and Hypertension. Training uses the Adam optimizer [31] (learning rate 0.001,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ) with categorical cross-entropy loss, a batch size of 32, up to 40 local epochs per round, and early stopping with a patience of 5 on the local validation loss. The architecture is shown in Fig. 3.


**Fig -3:** LSTM neural network architecture with layer configuration.

### 3.5 Privacy and Security Mechanism

No raw patient record leaves a hospital node; the only data crossing the network are encrypted model weights, sent over TLS once per communication round. Because the central aggregator sees only model parameters and never identifiable patient data, the system is designed to be consistent with HIPAA, GDPR and India's Digital Personal Data Protection Act [1], [16]. This does not eliminate every privacy risk, since gradient-based attacks are a known issue in the FL literature [2], [16], [32], but it removes the largest one, namely a central store of raw medical records that forms a single attack target.

### 3.6 Human-in-the-Loop Dashboard

The dashboard is a Streamlit application designed for doctors to check patients' current readings, the model's predicted risk category, a confidence score and a recent trend chart in one place. When an abnormal reading occurs, an alert is displayed with enough context for the doctor to review it and decide whether to acknowledge, escalate or dismiss it, without consulting a separate system first. Every action is logged, so that nothing flagged by the model disappears silently. The doctor remains the decision-maker, and the role of the system is to provide the information needed to decide quickly [10], [9].

## 4. MATHEMATICAL FORMULATION

### 4.1 Federated Averaging

At round  $t$ , each participating node  $k$  starts from the current global weights  $w_t$ , trains locally, and produces updated weights  $w_{t+1}^k$ . The aggregator then computes the new global model as a weighted average over all  $K$  nodes [4]:

$$w_{t+1} = \sum_{k=1}^K (n_k/n) w_{t+1}^k \quad (1)$$

where  $n_k$  is the number of local training samples at node  $k$  and  $n = \sum n_k$  is the total over all participating nodes. Weighting by  $n_k$  prevents a small clinic with few patients from influencing the global model as strongly as a large hospital with a much larger dataset [4].

### 4.2 LSTM Gating Equations

For input  $x_t$  at time step  $t$  and previous hidden state  $h_{t-1}$ , the LSTM cell computes the forget gate  $f_t$ , the input gate  $i_t$  and the output gate  $o_t$  as follows [3]:

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \quad (2)$$

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \quad (3)$$

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \quad (4)$$

$$\tilde{C}_t = \tanh(W_c \cdot [h_{t-1}, x_t] + b_c) \quad (5)$$

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \quad (6)$$

$$h_t = o_t \odot \tanh(C_t) \quad (7)$$

where  $\sigma$  is the sigmoid function,  $\tilde{C}_t$  is the candidate cell state and  $\odot$  denotes elementwise multiplication. The forget and input gates together decide how much of the previous cell state is kept and how much is overwritten, which allows the network to track a physiological trend over a longer window than a plain recurrent network [3].

### 4.3 Min-Max Normalization

Every continuous feature is scaled independently at each local node using

$$X_{\text{scaled}} = (X - X_{\min}) / (X_{\max} - X_{\min}) \quad (8)$$

with  $X_{\min}$  and  $X_{\max}$  computed only from that node's own training split, so that nothing about another hospital's data range leaks through the scaling [19].

#### 4.4 Evaluation Metrics

The usual confusion-matrix metrics are used throughout [9]. Accuracy is the overall proportion of correct predictions. Precision indicates how many of the predicted-positive alerts were genuine, and recall indicates how many of the true cases the model detected:

$$\text{Precision} = \text{TP} / (\text{TP} + \text{FP}) \quad (9)$$

$$\text{Recall} = \text{TP} / (\text{TP} + \text{FN}) \quad (10)$$

$$\text{F1} = 2 \cdot \text{Precision} \cdot \text{Recall} / (\text{Precision} + \text{Recall}) \quad (11)$$

The F1-score is the harmonic mean of precision and recall. The ROC-AUC is also reported for each disease class, since class imbalance arises naturally in this setting: at any given monitoring window, most patients are healthy [7].

## 5. DISCUSSION

### 5.1 Experimental Configuration

The federated pipeline was evaluated on public benchmark data distributed across simulated hospital nodes, as described in Section 3 [27], [28], [29], [30]. The nodes were given deliberately uneven, non-IID patient populations so that the setting resembles a real multi-institution deployment [19]. Table 3 lists the complete configuration used for the reported experiments.

**Table -3:** Experimental configuration

| Parameter                | Value                                           |
|--------------------------|-------------------------------------------------|
| Simulated hospital nodes | 8 (non-IID)                                     |
| Communication rounds     | 20                                              |
| Local epochs per round   | Up to 40                                        |
| Aggregation              | FedAvg, weighted by $n_k$                       |
| Input features           | 8 (Min-Max scaled per node)                     |
| Architecture             | LSTM 128, LSTM 64, Dense 32 (ReLU), Softmax 5   |
| Dropout                  | 0.3                                             |
| Optimizer                | Adam (lr 0.001, $\beta_1$ 0.9, $\beta_2$ 0.999) |
| Loss                     | Categorical cross-entropy                       |
| Batch size               | 32                                              |
| Early stopping           | Patience 5 (local validation loss)              |
| Data split per node      | 70-15-15 (train / val / test)                   |

**Table -5:** Comparison of centralized and federated training for healthcare monitoring

| Aspect                | Centralized learning                       | Federated learning (proposed)                            | Ref.       |
|-----------------------|--------------------------------------------|----------------------------------------------------------|------------|
| Location of raw data  | Pooled on one central server               | Remains at each hospital node                            | [7], [13]  |
| Breach exposure       | A single breach exposes all records        | No central store of raw records; only weights are shared | [2], [16]  |
| Regulatory fit        | Hospitals reluctant to release raw records | Only parameters leave the institution                    | [1]        |
| Communication         | One-time data transfer                     | Weights exchanged every round (20 rounds)                | [26]       |
| Data distribution     | Single pooled distribution                 | Non-IID across nodes                                     | [19]       |
| Residual privacy risk | Central attack target                      | Gradient-based attacks remain possible                   | [32], [16] |
| Accuracy              | Reference upper bound                      | Within about 0.5 pt of centralized (this work)           | [21], [11] |

Federated training does carry costs. Weights must be exchanged in every round, which adds communication overhead [26], and uneven participation, dropped rounds and device synchronization are known practical issues in deployment [24]. Weighting each node's contribution by its local sample count, as in Eq. (1), limits the influence of small clinics, but non-IID data remain a recognized source of difficulty for federated optimization [19]. Adding stronger privacy layers such as secure aggregation or differential privacy would further protect the weights but increases computational cost [22], [2].

Several limitations should be noted. The hospital nodes are simulated from public datasets and not drawn from live clinical deployments, so the reported accuracy should be confirmed on real multi-institution data. The system reduces but does not remove privacy risk, because gradient-based attacks are still possible [32], [16]. Finally, the dashboard supports and does not replace the clinician, who makes the final decision on every alert [10].

| Parameter    | Value                                      |
|--------------|--------------------------------------------|
| Transmission | Encrypted weights over TLS, once per round |

### 5.2 Reported Performance

After 20 communication rounds the global model reaches an accuracy of 97.6% and an F1-score of 97.2% on early disease prediction over the five classes (Normal, Cardiovascular Risk, Diabetes, Respiratory Disorder and Hypertension). This is within about half a percentage point of a centralized model trained with full access to the same pooled data, which serves as the upper-bound reference [4], [21]. These values are summarized in Table 4. The small remaining gap is consistent with earlier reports that FL approaches the accuracy of centralized training on medical tasks while keeping the data at its source [11], [12], [14].

**Table -4:** Reported performance of the federated global model

| Metric             | Federated (round 20) | Reference                |
|--------------------|----------------------|--------------------------|
| Accuracy           | 97.6%                | Centralized, pooled data |
| F1-score           | 97.2%                | Centralized, pooled data |
| Gap to centralized | $\approx 0.5$ pt     | Centralized, pooled data |
| Prediction classes | 5                    | -                        |
| Rounds to report   | 20                   | -                        |

### 5.3 Federated versus Centralized Training

The practical value of the system lies in what is avoided as well as in the accuracy achieved. Table 5 compares centralized training with the proposed federated design on the properties that matter for hospital adoption. In centralized training, all records are pooled and a single breach exposes every patient [2]; in the federated design only model parameters are exchanged, so there is no central store of raw records and the arrangement is easier to reconcile with HIPAA, GDPR and India's Digital Personal Data Protection Act [1], [16].

## 6. CONCLUSIONS

The proposed work combines IoT and Federated Learning to provide secure and privacy-preserving healthcare monitoring. It allows healthcare data to remain on local devices while a machine-learning model is trained collaboratively. This approach reduces the privacy risks associated with centralized data storage and provides a scalable solution for intelligent healthcare monitoring [1], [12]. The system can help improve data security, patient privacy and monitoring efficiency.

Future work can focus on real-time monitoring, multimodal healthcare data [17], stronger privacy techniques such as secure aggregation [33] and differential privacy [34], and large-scale deployment [24].

## REFERENCES

1. Dasaradharami Reddy, K., Gadekallu, T.R.: A comprehensive survey on federated learning techniques for healthcare informatics. *Comput. Intell. Neurosci.* 2023, Art. 8393990 (2023)
2. Abaoud, M., Almuqrin, M.A., Khan, M.F.: Advancing federated learning through novel mechanism for privacy preservation in healthcare applications. *IEEE Access* 11 (2023) 83562–83579
3. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* 9(8) (1997) 1735–1780
4. McMahan, B., Moore, E., Ramage, D., Hampson, S., Agüera y Arcas, B.: Communication-efficient learning of deep networks from decentralized data. In: *Proc. 20th Int. Conf. Artif. Intell. Statist. (AISTATS)* (2017)
5. Baig, M.M., Gholamhosseini, H.: Smart health monitoring systems: An overview of design and modeling. *J. Med. Syst.* 37(2), Art. 9898 (2013)
6. Baker, S.B., Xiang, W., Atkinson, I.: Internet of Things for smart healthcare: Technologies, challenges, and opportunities. *IEEE Access* 5 (2017) 26521–26544
7. Taha, Z.K., Yaw, C.T., Koh, S.P., Tiong, S.K., Kadirgama, K., Benedict, F., Tan, J.-D., Balasubramaniam, Y.: A survey of federated learning from data perspective in the healthcare domain: Challenges, methods, and future directions. *IEEE Access* 11 (2023) 45711–45735
8. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* 521 (2015) 436–444
9. Rajkomar, A., Dean, J., Kohane, I.: Machine learning in medicine. *N. Engl. J. Med.* 380 (2019) 1347–1358
10. Topol, E.J.: High-performance medicine: The convergence of human and artificial intelligence. *Nat. Med.* 25 (2019) 44–56
11. Sohan, M., Basalamah, A.: A systematic review on federated learning in medical image analysis. *IEEE Access* 11 (2023) 28628–28644
12. Shiranthika, C., Saeedi, P., Bajić, I.V.: Decentralized learning in healthcare: A review. *IEEE Access* (2023)
13. Yang, Q., Liu, Y., Chen, T., Tong, Y.: Federated machine learning: Concept and applications. *ACM Trans. Intell. Syst. Technol.* 10(2) (2019) 1–19
14. Zhang, F., et al.: Recent methodological advances in federated learning for healthcare. *arXiv preprint arXiv:2310.02874* (2023)
15. Rghioui, A., Lloret, J., Sendra, S., Oumnad, A.: A smart architecture for diabetic patient monitoring using machine learning. *Electronics* 9(9), Art. 1843 (2020)
16. Gu, X., Sabrina, F., Fan, Z., Sohail, S.: A review of privacy enhancement methods for federated learning in healthcare systems. *Int. J. Environ. Res. Public Health* 20(15), Art. 6539 (2023)
17. Thrasher, J., et al.: Multimodal federated learning in healthcare: A review. *arXiv preprint arXiv:2310.09650* (2023)
18. Ghosh, S., Chakraborty, A., Saha, S.: IoT-based smart healthcare monitoring systems: A literature review. *Health Technol.* 12 (2022) 1–19
19. Li, T., Sahu, A.K., Talwalkar, A., Smith, V.: Federated learning: Challenges, methods, and future directions. *IEEE Signal Process. Mag.* 37(3) (2020) 50–60
20. Rieke, N., et al.: The future of digital health with federated learning. *npj Digit. Med.* 3, Art. 119 (2020)
21. Sheller, M.J., et al.: Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. *Sci. Rep.* 10, Art. 12598 (2020)
22. Kaissis, G.A., Makowski, M.R., Rückert, D., Braren, R.F.: Secure, privacy-preserving and federated machine learning in medical imaging. *Nat. Mach. Intell.* 2 (2020) 305–311
23. Nguyen, D.C., et al.: Federated learning for Internet of Things: A comprehensive survey. *IEEE Commun. Surv. Tutor.* 23(3) (2021) 1622–1658
24. Li, M., Xu, P., Hu, J., Tang, Z., Yang, G.: From challenges and pitfalls to recommendations and opportunities: Implementing federated learning in healthcare. *arXiv preprint arXiv:2409.09727* (2024)
25. Zekiye, A., Özkasap, Ö.: Decentralized healthcare systems with federated learning and blockchain. *arXiv preprint arXiv:2306.17188* (2023)
26. Konečný, J., McMahan, H.B., Yu, F.X., Richtárik, P., Suresh, A.T., Bacon, D.: Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492* (2016)
27. Johnson, A.E.W., et al.: MIMIC-IV, a freely accessible electronic health record dataset. *Sci. Data* 10, Art. 1 (2023)
28. Goldberger, A.L., et al.: PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. *Circulation* 101(23) (2000) e215–e220
29. Janosi, A., Steinbrunn, W., Pfisterer, M., Detrano, R.: Heart Disease. *UCI Machine Learning Repository* (1988)
30. Schmidt, P., Reiss, A., Duerichen, R., Marberger, C., Van Laerhoven, K.: Introducing WESAD, a multimodal dataset for wearable stress and affect detection. In: *Proc. 20th ACM Int. Conf. Multimodal Interaction (ICMI)* (2018) 400–408
31. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *Proc. Int. Conf. Learn. Represent. (ICLR)* (2015)
32. Shokri, R., Shmatikov, V.: Privacy-preserving deep learning. In: *Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS)* (2015)
33. Bonawitz, K., et al.: Practical secure aggregation for privacy-preserving machine learning. In: *Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS)* (2017)
34. Dwork, C., Roth, A.: The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* 9(3–4) (2014) 211–407