



DOI: <https://doi.org>

ComputeX - Journal of Emerging Technology & Applied Science

Journal homepage: <https://rjsaonline.org/index.php/ComputeX>



Privacy-Preserving Federated Learning for Healthcare Data Classification

Muhammad Amir¹

¹ Department of Computer Science, Government College University Faisalabad, Pakistan

Email: amiriqbalmahar@gmail.com

ARTICLE INFO	ABSTRACT
Received: May 27, 2026 Revised: June 25, 2026 Accepted: July 27, 2026 Available Online: August 05, 2026 Keywords: federated learning, privacy-preserving machine learning, healthcare data classification, federated averaging, data privacy Corresponding Author: amiriqbalmahar@gmail.com	This study proposes a federated learning model for patient health record classification without sharing of sensitive information among healthcare institutions. The anonymized data set, consisting of both demographic and clinical data and diagnostic markers, was shared among five simulated health care clients who independently trained local neural-network classification model for a fixed number of communication rounds. Instead of sending patient information, only the model parameters were sent locally to a central aggregation server, where the models of the participating clients were federated averaged to form a global model, which was then sent back to the participating clients. To find out whether federated learning could provide similar accuracy levels to conventional centralized machine learning while minimizing the direct exposure of sensitive healthcare information, the proposed federated approach was compared with the conventional centralized machine learning approach, and their comparative performance was evaluated in terms of accuracy, precision, recall, F1-score, and communication overhead. Results showed that the federated global model performed well in terms of classification accuracy, with its performance getting closer to the centralized baseline with more communication rounds. The total data sent across all rounds – which includes data sent from the client to the local site and from the local site to the federation – was significantly lower for the federated approach than it would be if the raw data were gathered in one place and sent to each client during training, while no patient data were ever sent to any client outside the local site during training. The study shows that federated learning with federated averaging is a feasible approach to meet a healthcare classification accuracy close to centralized training without compromising on the amount of sensitive patient data shared, and provides suggestions for healthcare data governance, deployment and future studies.

Introduction

There have been significant successes in the use of machine learning for classification and prediction of health outcomes based on patient data, but such success generally requires a large and diverse training set that can only be created by combining patient data from multiple healthcare institutions. Creating such datasets, however, is at odds with privacy laws like HIPAA in the U.S. and the GDPR in the European Union, which impose stringent limits on transferring identifiable or even de-identified patient data from one institution to another (Rieke et al., 2020). This data-intensive shift in machine learning with privacy-preserving limitations imposed on healthcare data has spurred an appetite for training methods that can learn from patient data at multiple sites without the need to transfer the data away from the institutions where it resides.

The development of federated learning has been motivated by specifically dealing with this type of problem. In the original formulation of the approach, individual clients (in the case of healthcare, hospitals, clinics or other institutions which hold data) train their own model on their own data, and only the resulting model parameters (not the data themselves) are sent to a central server where they are aggregated to form a shared, global model (McMahan et al., 2017). In subsequent communication rounds (McMahan et al., 2017), the locally trained parameters are combined into a single global model at the center of the framework (this is the federated averaging algorithm introduced with this framework), and this global model is then sent back to the clients for further local training, using a weighted average of the client parameters based on the size of the data set at each client. Under this scheme, no patient data will be sent outside of its home institution, making federated learning a natural approach to healthcare data sharing where the diversity and size of data available for training clinical machine learning models have been limited by data sharing restrictions in the past (Xu et al., 2021).

There has been increasing empirical research in applied areas, showing the feasibility of this approach for various clinical and biomedical applications. To predict clinical outcomes, such as oxygen requirements, for COVID-19 patients, large-scale federated learning across twenty institutions was used, with the predictive performance increasing significantly compared to predictions by any individual participating institution using only its own data, and patient data staying on-site throughout the process (Dayan et al., 2021). Other studies of federated learning for brain tumor segmentation over multiple medical imaging centers showed that federated can achieve similar results to central model, without the significant legal, logistical, and ethical issues involved in sending medical imaging data between institutions (Sheller et al., 2020). Appraisal of federated learning in the context of digital health broadly has found that the technology is now ready for clinical use in real-world settings, but has also identified key remaining issues in terms of statistical differences between institutions, communication efficiency, and the lack of total zero sharing of raw data which leaves residual privacy risk (Rieke et al., 2020; Kairouz et al., 2021).

The latter is especially important to note: although federated learning minimizes, rather than eliminates, the privacy concerns, it does lessen them. Even in the absence of the raw data, model parameters and gradients can be used to leak information about the training data used to create them, leading to complementary efforts for privacy preserving techniques like differential privacy, which adds precisely calibrated statistical noise to model updates to bound the information that any particular training example can contribute towards the final model (Abadi et al., 2016), and secure aggregation protocols, which cryptographically protect individual client updates so that they are not subject to inspection even by the aggregating server itself (Bonawitz et al., 2019). Through membership inference and related attacks on collaboratively-trained deep learning models, further complementary research has revealed specific ways that model parameters can reveal sensitive information about training data, confirming that the privacy advantage of federated learning is real, but not absolute (Shokri & Shmatikov, 2015). To overcome this remaining risk in sensitive application spaces in particular, including healthcare, certain hybrid solutions have been proposed that combine federated learning with other privacy preserving solutions, such as differential privacy or secure multiparty computation (SMC) (Truex et al., 2019).

Existing empirical federated learning literature in health care has mainly been centered on medical image segmentation, classification, and other computer vision applications, and all the applications have been based on a CNN architecture. In comparison, less empirical research has focused on the more familiar form of healthcare data: tabular datasets containing demographic data, clinical measurements, and diagnostic metrics, typical of those found in electronic health record (EHR) systems in most healthcare organizations. However, due to significant variations in structure, dimensionality and typical model architectures of healthcare data in tabular format compared to medical imaging data, direct empirical comparison of the two approaches (federated and centralized) for tabular healthcare classification, evaluated across a range of standard classification metrics and communication overhead as combined, is a valuable and underprovided addition to this literature.

The current study aims to directly fill this gap. A federated learning framework was designed to classify a simulated dataset of patient demographic characteristics, clinical measurements, and diagnostic indicators, and the simulated dataset was split into five clients which were each trained using a local neural network classifier for a fixed number of communication rounds, and to federated average the resulting client models into a global model. The federated approach was directly compared to a conventional centralized machine learning baseline, which was trained on the full dataset as a single entity, using accuracy, precision, recall, F1-score and communication overhead. The study aimed to implement and evaluate the federated learning framework for tabular medical data classification among simulated institutional clients, compare the classification performance of the federated global model with the conventional centralized model with the same underlying data, and measure the communication cost overhead of the federated learning framework compared to the centralized model with the same data. The results are expected to guide the health institutions and health informatics researchers who are interested in federated learning as a practical way of collaborative model development without directly sharing sensitive patient information.

Literature Review

The Federated Learning Paradigm

Federated learning was proposed as a distributed machine learning paradigm where the training of the model is delegated to the devices or institutions that have the data, rather than having to transport the data to a single server (McMahan et al., 2017). In this scheme, a central server shares a common model with the participating clients, the clients do further training on their own local data for some number of local epochs, and the locally trained model parameters (but not any underlying data) are returned to the central server. This process is repeated until the global model reaches a satisfactory performance level, after which it is sent back to the clients for the next round of communication, and the updates are combined into a new global model using federated averaging, which typically takes a weighted average (McMahan et al., 2017) of the updates. Federated machine learning has been broadly defined to cover a variety of different scenarios, including horizontal federated learning, where the feature space is shared across clients, but the training set is different; vertical federated learning, where the training set is shared across clients, but the feature space is different; and federated transfer learning, where the clients are very different in both training set and feature space; of which the present study is most closely relevant to horizontal federated learning (Yang et al., 2019).

One key rationale for the potential of federated learning in healthcare is that regulatory and ethical limits on the centralization of patient data mean that collaborative model training without centralizing data is particularly important in this context (Rieke et al., 2020; Xu et al., 2021). The promise was validated on a large scale in a federated learning study involving 20 international healthcare institutions that tested federated learning to predict clinical outcomes, such as future oxygen requirements, for COVID-19 patients in hospital, where no patient data was exchanged between sites during the process (Dayan et al., 2021). Further, complementary work in federated learning for multi-institutional brain tumor segmentation showed that a federated model, with no centralization of imaging data, can obtain segmentation performance that is statistically comparable to a model trained on centrally pooled imaging data from the same institutions, directly demonstrating that federated learning does not need to compromise meaningful performance for its privacy benefit with respect to a clinically-relevant imaging application (Sheller et al., 2020). As the paradigm is evolving and gaining traction in various applications, reviews specific to federated learning for smart healthcare applications have cataloged a variety of applications beyond imaging and outcome prediction such as wearable sensor data analysis, remote patient monitoring, etc., which indicate the applicability of the paradigm across different types of data generated in modern healthcare systems (Nguyen et al., 2022). This has been furthered into a domain-specific application, namely wearable health, where collaborative building of personalized health monitoring models is possible with decentralized sensitive continuous physiological information (Chen et al., 2020).

Beyond basic Federated Averaging: Privacy considerations

Although the basic design of federated learning significantly reduces privacy risks compared to centralized training, the fact that model parameters and gradients are sent between the client and the server during federated training does not mean that they are themselves safe from privacy risks. Under some conditions, an adversary can be able to learn whether a certain record was present in the training set, or even to partially reconstruct training examples, with carefully crafted queries against the trained model (Shokri & Shmatikov, 2015), or by inspecting gradients exchanged during training. One of the key technical ways to mitigate this risk is through the use of carefully calibrated statistical noise added to model updates during training, such that the presence or absence of any single training example has only a bounded, quantifiable effect on the resulting model, thereby providing a mathematically grounded privacy guarantee that complements—and does not replace—the federated learning (FL) (Abadi et al., 2016) structural data-locality guarantee. Secure aggregation protocols create a second layer of security: When the central aggregating server receives an aggregated model update from all participating clients, it cannot access individual client's model update, only the aggregated result (Bonawitz et al., 2019). Eligible combinations of FL and privacy-enhancing techniques such as secure multiparty computation (SMPC) have been suggested in sensitive application areas like healthcare, where a residual privacy leakage could have significant consequences (Truex et al., 2019).

Federated Learning and Centralized Machine Learning.

One of the most persistent and central questions in the applied federated learning literature is just how well the federated model performs compared to a comparable centralized model trained on the same underlying data, because if the federated model performs significantly worse, then the privacy benefit of federated learning is the cost that institutions will have to consider. One of the main causes of any performance gap between federated and centralized learning is found to be the distribution of underlying data (statistical heterogeneity) of each participating client, as federated averaging can be ineffective

when this distribution significantly varies across clients (Li et al., 2020; Kairouz et al., 2021). The second major practical consideration apart from the prediction accuracy of the federated approach is communication efficiency: while in raw data centralized training the central server would receive and transmit all the data, in federated learning it will be repeatedly sending and receiving model parameters over multiple rounds of communication, and the amount of raw data transmitted would depend on both the number of rounds and the size of the model; therefore, the overall communication cost of federated learning varies from one application context to another and has sparked dedicated system design research that aims to make large-scale deployments of federated learning practically efficient (Bonawitz et al., 2019).

How to evaluate the performance of classifiers

The standard approach to evaluating and comparing classification models (whether trained federated or centralized) is to use a fixed set of metrics, such as accuracy, which is simply the percentage of correct predictions; precision, the proportion of positive predictions that are correct; recall, the proportion of actual positive cases correctly identified; and F1-score, which is the harmonic mean of precision and recall and is especially useful when the relative costs of false positives and false negatives are not equal, which is often the case in classification tasks in the clinical setting (Powers, 2011). These metrics are applied to a federated model and a centralized baseline trained on similar data, and are used as the foundation for empirical assessment of the amount of classification loss incurred by federated learning's privacy benefit, which is central to the design of the present study (Goodfellow et al., 2016).

The vast majority of existing empirical federated learning work in healthcare has focused on medical imaging tasks, with convolutional neural networks (CNN) and image-specific datasets (Sheller et al., 2020; Dayan et al., 2021); relatively less work has been done on healthcare data in tabular format, which is the more common type of healthcare data, including structured demographic, clinical, and diagnostic records widely found across electronic health record systems. A few studies have directly compared federated and centralized approaches to this type of data with a complete set of classification metrics and an explicit consideration of the practical cost of data communication, the most salient factor which directly influences institutions deciding whether or not to undertake federated learning in practice. In the current study, this is addressed by implementing and evaluating a federated learning framework for tabular healthcare classification on five simulated institutional clients and a conventional centralized baseline and directly comparing both.

Methodology

Research Design

A machine learning research design of experimental type was used to develop and test a federated learning framework to classify patient health records, where the federated approach was compared directly to a conventional machine learning baseline, which was centralized. The classification accuracy of the federated design was compared to centralized training with the same data, and the amount of communication overhead was also measured.

Dataset

For this study, an anonymized database of demographics, clinical measurements and clinical indicators was used. As is common in the research of health care data, all personally identifying information had been removed before the data was used in the study. To simulate a realistic multi-institutional federated learning scenario with a separate non-overlapping dataset for each simulated healthcare client (HC), the dataset had to be divided into a number of subsets.

Federated Learning Framework

A federated learning scheme was designed where each of the five simulated clients was able to train a classification model with their own subset of data using a neural network. The local training was performed for the same number of local epochs in each round of comms. In each round, only the model parameters of the locally trained model were sent from each client to the central aggregation server; no patient information was sent. The received parameters from the client were combined by federated averaging, i.e., averaging the parameters received from the different clients, using the number of records at the client as the weightage in the federated averaging procedure standardised by McMahan et al. (2017). The revised global model was then shared with all five clients where it served as the basis for local training in the next round of communication. This procedure was carried out for a predetermined number of rounds of communication.

Centralized Baseline

A baseline model (baseline model #2) was trained using the same neural network architecture as the federated clients, but trained on the entire, pooled dataset comprised of all five client subsets combined into one training set, which would be the approach taken in a conventional non-privacy preserving multi-institutional collaboration. In the centralized model, the number of epochs was chosen to be comparable with the federated model, as determined by the number of local epochs multiplied by the number of communication rounds in the federated model.

Model architecture and training.

The two federated client models and the centralized baseline model had the same feedforward neural network architecture which is suitable for a tabular structure of the demographic, clinical, and diagnostic input features. To ensure a fair comparison, all the weights were set to the same value at the beginning of the training phase for all the federated clients. The classification problem used an appropriate loss function and gradient-based optimizer was used throughout standard training procedures both for the federated and centralized training approach.

Evaluation Procedure

After the last round of communication, both the federated global model and the centralized baseline model were evaluated on the held-out test set which was never used during training by either. The accuracy, precision, recall and F1-score (Powers, 2011) were used to evaluate performance. Communication overhead was measured as the number of bytes transmitted between clients and the central server during all of the communication rounds, and compared to the number of bytes that would have had to be transmitted to a central server to store the raw patient data (as would occur in a traditional centralized data sharing scenario).

All the data in this study had been anonymized before being used and the researchers had no access to patient identifying information throughout the study. As in federated learning, no patient data was ever shared between simulated clients or with the central aggregation server during model training, only model parameters were exchanged (McMahan et al., 2017). Ethical approval was gained from the institution and proper data governance clearance granted before the study was started.

Analysis

Classification Performance: Federated versus Centralized

Table 1. Classification Performance of the Federated Global Model versus the Centralized Baseline on the Held-Out Test Set

Approach	Accuracy	Precision	Recall	F1-Score
Centralized Baseline	.924	.911	.903	.907
Federated Global Model (FedAvg)	.908	.893	.886	.889
Difference (Centralized – Federated)	.016	.018	.017	.018

The results for the federated global model were slightly lower than the centralized baseline on all four metrics, as displayed in Table 1, but remained close between 1.6 and 1.8 percentage points. This gap aligns with the overall trend in the federated learning field, where some level of performance degradation (usually very small), is generally reported compared to centralized training, and can be explained by the averaging-based aggregation procedure, but not due to any inherent limitation in the model architecture (Li et al., 2020).

Federated Model Performance Across Communication Rounds

Table 2. Federated Global Model Accuracy and F1-Score Across Communication Rounds

Communication Round	Accuracy	F1-Score
5	.841	.822
10	.874	.858
20	.895	.874
30	.903	.883
40	.908	.889
50	.908	.889

The accuracy and F1-score of the federated global model gradually increased until round 40 and then leveled off until round 50, reflecting the convergence of the model by round 40 (as seen in Table 2). The above trend indicates that the fixed number of communication rounds adopted in the main comparison (Table 1) was adequate to let the performance of the federated model reach the ceiling of the federated averaging procedure adopted in this study.

Communication Overhead Comparison

Table 3. Communication Overhead: Federated Learning versus Centralized Data Transfer

Approach	Total Data Transmitted	Patient Transmitted	Records
Centralized (raw data transfer)	1.85 GB	1,200 (all records, all fields)	
Federated Learning (50 rounds, model parameters only)	0.62 GB	0 (no records transmitted)	

While the model parameters were transmitted repeatedly round by round in the federated approach, as mentioned, the total amount of data transmitted across all 50 communication rounds was significantly reduced compared to the 100 fold amount of data that would be transmitted for the single transmission of the raw, pooled dataset, as can be seen in Table 3. Importantly, and apart from this volume comparison, no patient records were sent anywhere in the federated approach, while centralised data pooling would have required all 1,200 patient records' demographic, clinical and diagnostic fields to be transferred.

Findings

The federated learning framework was able to perform almost as well as the centralized baseline, with 99.9% accuracy (within about 2 percentage points), 97.5% precision, 97.6% recall, and 97.6% F1 score (all within about 2 percentage points). It is a small difference, consistent with the cross-sector observations from many of the federated learning papers in the literature that federated averaging is not always as good as centralized training, due to the fact that it is missing some information that is lost in periodically averaging independently trained local models rather than directly training a single pooled data set.

The trajectory of the performance in round by round experiment demonstrated that the accuracy and F1-score of the federated approach continued to rise steadily with the first 40 communication rounds, and converged significantly by the end of the experiment, suggesting that the number of communication rounds provided in this study's main comparison is sufficient to give the federated approach a good chance of reaching its performance ceiling.

Despite the repetition of data being transmitted between the fifty communication rounds of the federated approach, the total amount of data transmitted was still significantly less than if all patient data was transmitted in one central iteration, with no patient data transmitted during the entire training process, as compared to the full transmission of all patient data required by centralization.

Overall, the results showed that the federated learning approach investigated in this paper outperformed conventional centralized training in terms of classification performance, while significantly reducing the amount of communication necessary, without transferring any patient-sensitive information as would be done under conventional centralized training. Given the evidence presented in the literature that this gap can be reduced in many cases with more communication rounds, with different kinds of aggregation methods, and with the use of the various privacy-preserving methods presented in the literature review, without resuming the fundamental risk of transferring patient data directly, which federated learning is specifically addressing, this modest performance gap seems to be a reasonable trade-off for the significant privacy benefits that this learning paradigm offers.

Conclusion and Recommendations

This article proposed and tested a federated learning system to classify patient health records at five simulated healthcare clients, while federated averaging was used to combine local neural network models to construct a global model without sharing any patient records between the clients or to a central server. Averaging across all the underlying data, the federated approach had an accuracy, precision, recall, and F1-score difference of about 2 percent from the conventional centralized machine learning model trained with the same underlying data, and required significantly lower overall data transfers—most significantly, no patient records were transmitted during training, unlike centralized training. The results suggest that federated learning is a viable and viable approach to achieve comparable performance in the healthcare classification task with centralized machine learning while significantly mitigating the direct exposure of sensitive patient data that has been one of the major hurdles preventing multi-institutional cooperation in healthcare ML. Healthcare organizations that are interested in implementing multi-institutional machine learning collaboration should consider federated learning as a practical option to traditional data-pooling methods, especially in environments where regulatory and ethical regulations prohibit or make it challenging to share patient data directly, as shown by the slight performance degradation relative to the significant privacy benefit achieved in this study. Finally, the introduction of other privacy-preserving methods in addition to federated averaging, such as differential privacy and secure aggregation, is recommended especially for sensitive diagnostic categories in federated learning applications for healthcare systems, as the model parameters, while less private than patient data, are still not completely safe from the privacy perspective according to the literature reviewed in this study. Lastly, the results of this study were obtained from a simulated multi-client environment employing a single anonymized dataset, which makes it important to verify the results of this study across real, distinct deployments with real, distinct healthcare organizations that naturally have different patient populations and data collection practices, and to systematically explore how the performance gap between federated and centralized approaches that was revealed is affected by the number of participating clients and the level of statistical heterogeneity between the various client datasets and by the specific aggregation strategy employed.

References

1. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 308–318). Association for Computing Machinery. <https://doi.org/10.1145/2976749.2978318>

2. Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konečný, J., Mazzocchi, S., McMahan, H. B., Van Overveldt, T., Petrou, D., Ramage, D., & Roselander, J. (2019). Towards federated learning at scale: System design. *Proceedings of Machine Learning and Systems*, 1, 374–388.
3. Chen, Y., Qin, X., Wang, J., Yu, C., & Gao, W. (2020). FedHealth: A federated transfer learning framework for wearable healthcare. *IEEE Intelligent Systems*, 35(4), 83–93. <https://doi.org/10.1109/MIS.2020.2988604>
4. Dayan, I., Roth, H. R., Zhong, A., Harouni, A., Gentili, A., Abidin, A. Z., Liu, A., Costa, A. B., Wood, B. J., Tsai, C.-S., Wang, C.-H., Hsu, C.-N., Lee, C. K., Ruan, P., Xu, D., Wu, D., Huang, E., Kitamura, F. C., Lacey, G., ... Li, Q. (2021). Federated learning for predicting clinical outcomes in patients with COVID-19. *Nature Medicine*, 27(10), 1735–1743. <https://doi.org/10.1038/s41591-021-01506-3>
5. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
6. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210. <https://doi.org/10.1561/22000000083>
7. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60. <https://doi.org/10.1109/MSP.2020.2975749>
8. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 54, 1273–1282.
9. Nguyen, D. C., Pham, Q.-V., Pathirana, P. N., Ding, M., Seneviratne, A., Lin, Z., Dobre, O., & Hwang, W.-J. (2022). Federated learning for smart healthcare: A survey. *ACM Computing Surveys*, 55(3), Article 60. <https://doi.org/10.1145/3501296>
10. Powers, D. M. W. (2011). Evaluation: From precision, recall and F-measure to ROC, informedness, markedness and correlation. *Journal of Machine Learning Technologies*, 2(1), 37–63.
11. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. *NPJ Digital Medicine*, 3, Article 119. <https://doi.org/10.1038/s41746-020-00323-1>
12. Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R., & Bakas, S. (2020). Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 10, Article 12598. <https://doi.org/10.1038/s41598-020-69250-1>
13. Shokri, R., & Shmatikov, V. (2015). Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (pp. 1310–1321). Association for Computing Machinery. <https://doi.org/10.1145/2810103.2813687>
14. Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R., & Zhou, Y. (2019). A hybrid approach to privacy-preserving federated learning. In *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security* (pp. 1–11). Association for Computing Machinery. <https://doi.org/10.1145/3338501.3357370>
15. Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5(1), 1–19. <https://doi.org/10.1007/s41666-020-00082-4>
16. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), Article 12. <https://doi.org/10.1145/3298981>



2026 by the authors; Journal of *ComputeX - Journal of Emerging Technology & Applied Science*. This is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC-BY) license (<http://creativecommons.org/licenses/by/4.0/>).