

Federated Machine Learning Models for Privacy-Preserving Analysis of Distributed Healthcare Data

Dr. Benjamin Carter

Professor, California Institute of Privacy-Preserving Technologies, United States of America

Abstract

The increasing adoption of digital healthcare technologies has generated substantial volumes of patient information distributed across hospitals, clinical laboratories, research institutions, and healthcare information systems. These datasets provide opportunities for predictive modeling, disease-risk assessment, and evidence-based medical research. However, conventional centralized machine learning approaches often require transferring sensitive patient records into shared repositories, creating concerns related to confidentiality, institutional governance, information security, and restricted data access. Federated machine learning offers an alternative computational approach in which participating institutions contribute to model development while retaining patient-level records within their local environments. This hypothetical methodological study examines federated machine learning models for privacy-preserving analysis of distributed healthcare data. The proposed framework incorporates institution-level data preprocessing, local model training, parameter aggregation, privacy-enhancing mechanisms, and independent predictive evaluation. A hypothetical dataset distributed across four healthcare institutions is developed to illustrate differences in sample size and clinical outcome prevalence. A second analytical scenario compares assumed classification outcomes for local learning, basic federated learning, and federated learning incorporating additional privacy-protection mechanisms.

The illustrative calculations demonstrate how predictive evaluation can be conducted across distributed learning configurations using accuracy, sensitivity, and F1 score. Differences between the assumed configurations highlight the need to evaluate predictive utility independently of privacy protection. No formal privacy guarantee is attributed to the hypothetical configurations. The study emphasizes that retaining patient records locally may reduce the need for centralized data collection but does not automatically prevent information leakage through transmitted model updates. Appropriate privacy safeguards, secure aggregation, differential privacy, institutional coordination, and reliable evaluation remain essential considerations. The paper contributes a methodological foundation for future distributed healthcare analytics research and concludes that federated machine learning requires empirical validation before claims concerning predictive improvement or patient-data protection can be established.

Keywords: Federated Machine Learning, Distributed Healthcare Data, Privacy-Preserving Analytics, Electronic Health Records, Federated Optimization, Secure Aggregation, Differential Privacy, Clinical Prediction, Healthcare Informatics, Collaborative Learning.

1. Introduction

The increasing digitization of healthcare services has transformed the manner in which clinical information is collected, maintained, and analyzed. Hospitals and medical institutions routinely generate electronic health records containing diagnostic information, laboratory measurements, treatment histories, demographic characteristics, and clinical outcomes. These records provide opportunities for developing predictive models capable of identifying statistical relationships relevant to healthcare research. However, information collected across different institutions is frequently maintained within independent databases, limiting the availability of comprehensive datasets for conventional machine learning applications.

Centralized machine learning generally requires relevant information to be transferred into a shared computational environment before model development. Although this approach may simplify data integration, the movement and consolidation of sensitive healthcare records create substantial confidentiality and governance concerns. Patient information may contain attributes capable of identifying individuals directly or through combinations of apparently nonidentifying variables. Research by El Emam et al. (2011) examined re-identification risks associated with health information, emphasizing the importance of evaluating privacy safeguards within their specific technical and organizational circumstances.

Distributed learning provides an alternative approach by enabling participating institutions to perform computations without necessarily transferring complete patient-level datasets. Earlier research investigated privacy-preserving statistical analysis across separate data repositories and established methods for collaborative learning. Jiang et al. (2013) examined distributed logistic regression through WebGLORE, while Lu et al. (2015) investigated distributed survival analysis without patient-level data sharing. These contributions demonstrated that selected forms of statistical learning could be coordinated across healthcare institutions while reducing dependence on centralized patient-data repositories.

Federated machine learning extends distributed optimization through procedures in which local participants contribute model updates to a shared learning process. McMahan et al. (2016) developed an influential decentralized learning approach based on aggregated local model updates, while Konečný et al. (2016) investigated associated communication and optimization challenges. These methods offer a foundation for collaborative predictive modeling across independently maintained datasets. However, keeping raw information at its original location is not equivalent to guaranteeing confidentiality, because shared model parameters or gradients may still reveal sensitive information under certain conditions.

The present study examines federated machine learning as a methodological approach for analyzing distributed healthcare data while considering privacy-related requirements. It develops a hypothetical institutional learning environment, identifies major computational components, and evaluates illustrative predictive outcomes through transparent calculations. Particular attention is directed toward differences in institutional datasets, the role of privacy-enhancing mechanisms, and the relationship between predictive performance and information protection. The study does not claim to have collected clinical records or implemented a federated learning system. Instead, it establishes a structured foundation for subsequent empirical investigation.

2. Objectives of the Study

The primary objective of this study is to examine federated machine learning as a methodological approach for collaborative analysis of distributed healthcare information. The research considers decentralized training, privacy-enhancing techniques, and predictive performance evaluation. It also identifies technical and institutional conditions that influence the reliability of privacy-conscious healthcare analytics.

2.1 Examination of Federated Learning Architecture

The first objective is to examine how independently maintained healthcare datasets can contribute to collaborative model development without requiring routine transfer of complete patient-level records. Particular attention is directed toward local computation, shared model parameters, and aggregation procedures. The study considers the conceptual differences between centralized learning and federated optimization. Understanding these components provides a foundation for evaluating whether distributed learning architectures are appropriate for healthcare institutions that maintain separate information systems and operate under different organizational requirements.

The objective also involves examining the influence of institutional differences on collaborative learning. Hospitals may serve populations with different demographic characteristics, disease prevalence, and clinical recording practices. These differences can influence local training outcomes and the behavior of a shared predictive model. The study therefore considers how unequal dataset sizes and heterogeneous observations affect model aggregation. This assessment supports the development of methodological procedures that recognize institutional variability instead of assuming that every participating healthcare organization contributes statistically equivalent information.

2.2 Assessment of Predictive Performance

The second objective is to examine how predictive outcomes can be evaluated across alternative healthcare learning configurations. The proposed analysis considers local-only learning, basic federated learning, and federated learning incorporating additional privacy safeguards. Hypothetical classification outcomes provide the basis for calculating accuracy, sensitivity, and F1 score. These calculations are illustrative and do not represent experimental comparisons. Their purpose is to establish a transparent evaluation method rather than demonstrate that a particular learning strategy provides superior clinical predictions.

Another purpose is to clarify the distinction between predictive utility and privacy protection. A model may produce satisfactory classification results while providing insufficient safeguards against disclosure through shared information. Conversely, privacy-enhancing procedures may influence model utility depending on their implementation and parameters. The study examines these considerations as separate evaluation requirements. This objective emphasizes that acceptable federated learning performance must be supported by both reliable predictive evidence and a clearly defined assessment of privacy-related risks.

2.3 Identification of Privacy and Implementation Challenges

The third objective is to identify challenges associated with privacy-conscious machine learning across healthcare institutions. These include information heterogeneity, unequal participation, communication constraints, model-update confidentiality, and data-governance requirements. The study considers how these conditions may affect collaborative model development and the interpretation of results.

Identifying relevant difficulties is important because a technically feasible distributed learning architecture may remain unsuitable when institutional policies, security expectations, or operational capabilities are not adequately addressed.

The investigation further aims to establish requirements for reproducible and responsible evaluation. Future studies should specify the participating institutions, dataset characteristics, model architecture, aggregation strategy, privacy mechanisms, and procedures for evaluating clinical predictions. Evidence concerning information security should remain distinct from evidence concerning prediction accuracy. This objective therefore supports a methodological framework in which federated learning is assessed through its statistical, computational, organizational, and privacy-related characteristics rather than through unsupported claims of comprehensive protection.

3. Review of Literature

3.1 Foundations of Privacy-Preserving Distributed Learning

Agrawal and Srikant (2000) examined privacy-preserving data mining and demonstrated that useful statistical analysis could be performed while limiting direct access to original records. Sweeney (2002) introduced the k-anonymity framework for protecting information in released datasets, while Vaidya and Clifton (2002) investigated collaborative data mining involving vertically partitioned information. These contributions established important foundations for understanding the relationship between analytical utility and confidentiality. They also demonstrated that different privacy techniques address distinct disclosure problems and computational settings.

Yu et al. (2008) investigated privacy-preserving Cox regression for distributed survival analysis, including applications involving clinical information held by separate institutions. Dwork et al. (2006) developed foundational techniques for calibrating statistical noise to privacy sensitivity, while Dwork and Roth (2014) provided a comprehensive theoretical account of differential privacy. These studies established that privacy protection requires explicitly defined mechanisms and assumptions. Their contributions remain relevant to federated learning because decentralized computation alone does not provide a mathematically specified guarantee against information disclosure.

3.2 Distributed Healthcare Analytics and Collaborative Models

Jiang et al. (2013) introduced WebGLORE, a distributed logistic regression service designed to support analysis across sensitive datasets. Their work illustrated how aggregated local statistics could contribute to shared model estimation. Ohno-Machado et al. (2014) described pSCANNER, a healthcare research network incorporating distributed analytical methods and institutional governance. Wiens et al. (2014) examined predictive modeling using information from multiple hospitals and highlighted the importance of accounting for institutional differences when developing models intended for local clinical environments.

Lu et al. (2015) developed WebDISCO, a distributed approach for Cox regression that avoids transferring patient-level records between participating sites. Their research demonstrated the technical feasibility of collaborative survival modeling through selected intermediate statistics. Collins et al. (2015) established reporting principles for clinical prediction models, emphasizing transparency concerning model development and evaluation. Together, these studies provide important foundations for healthcare-oriented federated learning, although distributed statistical analysis and modern iterative federated optimization should not be treated as identical computational procedures.

3.3 Federated Optimization and Privacy Mechanisms

Konečný et al. (2015) examined optimization across unevenly distributed datasets, identifying communication efficiency as an important constraint. McMahan et al. (2016) subsequently described collaborative model training through the aggregation of local updates while keeping original training data distributed. Konečný et al. (2016) investigated additional communication-efficient approaches to federated learning. These contributions established the methodological foundations of modern federated optimization, particularly the coordination of multiple local training processes and the challenges associated with heterogeneous participant data.

Abadi et al. (2016) examined differentially private training for deep learning, establishing an approach for managing privacy loss during model optimization. Bonawitz et al. (2016) investigated secure aggregation for protecting individual contributions during collaborative computation. Fredrikson et al. (2015) demonstrated privacy concerns associated with extracting information from predictive models. These studies collectively emphasize that privacy protection requires more than decentralized data storage. Effective safeguards must consider the information exposed through model outputs, intermediate updates, and the assumptions governing participating computational entities.

4. Research Methodology

4.1 Research Design and Federated Learning Framework

The present study adopts a hypothetical methodological research design supported by illustrative analytical calculations. The proposed framework considers four independent healthcare institutions maintaining locally stored clinical information. Each institution contributes to a conceptual collaborative model-training process without transmitting its complete patient-level dataset to a coordinating server. The framework comprises local preprocessing, model optimization, parameter aggregation, privacy assessment, and predictive validation. No hospital data were collected, and no federated learning software implementation was executed for this manuscript.

The conceptual learning procedure begins with a shared initial model distributed to participating institutions. Each site would optimize the model using locally available training observations and return an authorized update for aggregation. The coordinating process would combine eligible contributions to construct a revised shared model. For standard weighted model averaging, the conceptual update can be expressed as $w = \sum (n_k/N)w_k$, where n_k represents each institution's training sample size, N represents the combined sample size, and w_k represents its locally updated parameters.

4.2 Hypothetical Healthcare Dataset and Privacy Assumptions

The illustrative dataset comprises 1,000 hypothetical healthcare observations distributed among four institutions. The assigned institutional sample sizes are 400, 300, 200, and 100 records respectively. A binary clinical outcome is assumed, with 120 positive and 880 negative observations across the participating institutions. These quantities are hypothetical assumptions introduced to illustrate institutional data imbalance and differences in outcome prevalence. They do not represent actual patient records, clinical diagnoses, or information obtained from any healthcare organization.

The proposed framework assumes that direct identifiers and original patient records remain within each institution's controlled information environment. Nevertheless, ordinary federated model updates may contain information about local training data and must not automatically be considered anonymous. The methodology therefore distinguishes basic federated learning from additional

protections such as secure aggregation and differentially private optimization. No actual cryptographic protocol, privacy budget, or protection mechanism was implemented. Consequently, the illustrative configurations do not establish any formal confidentiality or security guarantee.

4.3 Analytical Procedure and Evaluation Measures

The analytical procedure considers three hypothetical predictive configurations: a local-only model, a basic federated model, and a federated model with additional privacy safeguards. Their classification outcomes are assigned for illustrative comparison using a separate hypothetical evaluation set containing 400 observations. The evaluation includes 80 assumed positive cases and 320 negative cases. These assigned outcomes are not generated through model training, and their numerical differences do not provide empirical evidence of advantages or disadvantages associated with federated learning.

Predictive evaluation is conducted using accuracy, sensitivity, and F1 score. Accuracy represents the proportion of correctly classified observations, while sensitivity measures the proportion of positive outcomes identified. The F1 score combines precision and sensitivity into a single measure.

The calculations follow standard confusion-matrix definitions:

$$\text{Accuracy} = (TP + TN) / (TP + TN + FP + FN)$$

$$\text{Sensitivity} = TP / (TP + FN)$$

$$\text{F1 Score} = 2TP / (2TP + FP + FN)$$

A complete empirical study would additionally require clinical calibration, subgroup evaluation, external validation, communication analysis, and explicit privacy-risk assessment.

5. Analysis

5.1 Distribution of Hypothetical Healthcare Data Across Institutions

The first hypothetical analysis examines the distribution of clinical observations among four independently operating healthcare institutions. The conceptual training environment contains 1,000 records, with individual institutional contributions ranging from 100 to 400 observations. Positive outcome counts also differ across the institutions, demonstrating that equal outcome prevalence cannot be assumed. These differences illustrate the importance of examining local dataset composition before collaborative model development. The assigned values are not derived from genuine electronic health records or clinical databases.

Institutional data heterogeneity may influence how individual model updates contribute to a shared predictor. An institution with a larger dataset may receive greater aggregation weight under sample-size-based averaging, while smaller institutions may contain clinically important patterns insufficiently represented elsewhere. The hypothetical distribution therefore demonstrates why both sample quantity and outcome prevalence require consideration. However, the table does not establish the direction or magnitude of any resulting prediction bias, because no model was trained on the assigned observations.

Table 1: Hypothetical Distribution of Healthcare Data Across Four Institutions

Healthcare Institution	Total Records	Positive Outcomes	Negative Outcomes	Positive Prevalence (%)
Institution A	400	40	360	10.00
Institution B	300	45	255	15.00
Institution C	200	10	190	5.00
Institution D	100	25	75	25.00
Total	1,000	120	880	12.00

Note: All institutional identities, record counts, and outcomes are illustrative. No actual healthcare institution or patient dataset was involved.

5.2 Comparative Predictive Performance Across Learning Configurations

The second hypothetical analysis examines classification outcomes under three alternative learning configurations. A separate evaluation set containing 400 observations is assumed, including 80 positive cases and 320 negative cases. The local-only configuration is assigned 56 true positives and 288 true negatives, producing accuracy of 86%. The basic federated configuration is assigned 64 true positives and 284 true negatives, producing accuracy of 87%. These values illustrate confusion-matrix calculations rather than observed predictive improvements.

The privacy-enhanced federated configuration is assigned 60 true positives and 280 true negatives, producing hypothetical accuracy of 85%. Its sensitivity is 75%, compared with 70% and 80% for the local-only and basic federated configurations respectively. These differences are constructed assumptions and must not be interpreted as evidence that privacy safeguards necessarily reduce model performance. No differential privacy parameter, secure aggregation implementation, or actual model optimization was evaluated. Predictive utility and information security therefore remain separate, unverified properties of the proposed framework.

Table 2: Hypothetical Classification Performance of Three Learning Configurations

Learning Configuration	TP	FP	FN	TN	Accuracy (%)	Sensitivity (%)	F1 Score (%)
Local-Only Learning	56	32	24	288	86.00	70.00	66.67
Basic Federated Learning	64	36	16	284	87.00	80.00	71.11
Privacy-Enhanced Federated Learning	60	40	20	280	85.00	75.00	66.67

Note: Each configuration contains 80 positive and 320 negative evaluation observations. All outcomes are illustrative and do not constitute empirical model results or measured privacy protection.

6. Discussion

6.1 Federated Learning and Distributed Healthcare Collaboration

The hypothetical analysis provides a methodological illustration of how healthcare institutions with unequal datasets may participate in collaborative predictive modeling. Federated learning allows selected computations to be conducted locally while model-related information is exchanged for aggregation. This arrangement may reduce dependence on centralized patient-data repositories. Previous research involving WebGLORE and WebDISCO demonstrated the technical potential of distributed clinical modeling, supporting the broader rationale for collaborative analysis across independently maintained healthcare information systems.

Nevertheless, decentralization does not automatically establish effective collaboration or reliable clinical predictions. Participating institutions may differ in patient characteristics, diagnostic practices, and available explanatory variables. Such differences can influence the behavior of a shared model and limit its applicability at individual sites. The illustrative institutional distribution emphasizes these considerations without quantifying their practical effects. Actual federated implementations would require consistent data definitions, appropriate local validation, and transparent assessment of whether collaborative learning benefits each participating institution.

6.2 Predictive Utility and Privacy Protection

The illustrative performance comparison demonstrates why predictive evaluation should remain separate from privacy assessment. Accuracy, sensitivity, and F1 score describe how a classifier distinguishes assumed positive and negative outcomes. None of these measures establishes whether transmitted model information could reveal sensitive patient characteristics. Differential privacy and secure aggregation address particular disclosure concerns, but their effectiveness depends on explicitly specified algorithms, parameters, and threat models. Consequently, an apparently accurate model cannot automatically be described as privacy-preserving.

The methodological principles developed by Dwork and Roth (2014), Abadi et al. (2016), and Bonawitz et al. (2016) provide important foundations for strengthening privacy-related safeguards. Differential privacy requires an appropriate privacy-loss accounting procedure, while secure aggregation limits access to individual submitted contributions under specified assumptions. These techniques address different risks and should not be treated as interchangeable. Their combined use may be valuable, but empirical investigation and formal analysis are required to establish the protections achieved by a particular implementation.

6.3 Clinical Reliability and Institutional Differences

Healthcare prediction models require evaluation beyond overall classification accuracy because errors may have different consequences across clinical applications. A false-negative outcome can be particularly important when the target represents a condition requiring timely assessment, although the practical consequences depend on the specific task. Sensitivity, specificity, calibration, and external validation therefore remain relevant. Steyerberg et al. (2010) and Collins et al. (2015) emphasized methodological principles for evaluating and reporting clinical prediction models.

The differences in institutional sample size and outcome prevalence illustrated in Table 1 also raise questions concerning model generalizability. A shared predictor may perform differently across institutions because local patient populations or recording procedures differ. Wiens et al. (2014)

demonstrated the importance of considering hospital-specific characteristics when transferring predictive knowledge. Accordingly, future federated healthcare experiments should assess site-level performance rather than assuming that satisfactory combined accuracy demonstrates equivalent reliability across all participating organizations.

7. Challenges

7.1 Healthcare Data Heterogeneity and Interoperability

Distributed healthcare datasets frequently differ in measurement practices, clinical terminology, outcome definitions, coding conventions, and completeness. These differences can complicate the construction of shared learning objectives across institutions. A variable recorded consistently at one hospital may be unavailable or represented differently elsewhere. Such inconsistencies may influence local optimization and reduce the reliability of aggregated predictions. Federated learning does not eliminate these difficulties merely by keeping original records within their source institutions.

Another challenge involves statistical heterogeneity in patient populations and disease prevalence. Institutions may serve communities with different demographic distributions, healthcare needs, and patterns of clinical utilization. These differences can create unequal relationships between predictors and target outcomes. A shared model may therefore favor patterns associated with larger institutions or more frequently represented patient groups. Appropriate evaluation requires examination of institution-level performance, careful handling of class imbalance, and consideration of whether global and local predictive objectives remain compatible.

7.2 Privacy Leakage and Information Security

Federated learning reduces the need to transfer raw patient records, but intermediate model updates can still contain information associated with local training observations. Potential disclosure concerns depend on the information transmitted, the learning algorithm, and the capabilities of parties observing the computation. Research on model inversion and privacy-preserving learning demonstrates the importance of examining information exposure beyond conventional dataset sharing. Consequently, decentralized architecture should be treated as one component of privacy-conscious system design rather than a complete confidentiality guarantee.

Security protection introduces additional implementation challenges because different safeguards address different risks. Secure aggregation can limit visibility into individual contributions, while differential privacy can provide a mathematically defined bound on selected forms of information disclosure when correctly implemented. Encryption protects information under specified cryptographic assumptions but does not independently establish statistical privacy. Healthcare institutions must therefore define appropriate threat models and evaluate whether their chosen mechanisms provide sufficient protection for the intended collaborative analysis.

7.3 Communication Efficiency and Institutional Governance

Federated training requires coordination among computational participants that may have different network capacities, hardware resources, and operational schedules. Repeated transmission of model updates can increase communication costs, particularly for complex predictive architectures. Unequal institutional participation may also affect model convergence and representativeness. Research by Konečný et al. (2016) emphasizes the importance of communication-efficient learning procedures.

Practical healthcare applications must additionally consider whether participating institutions can support reliable computation without interrupting clinical information systems.

Institutional governance presents challenges that cannot be resolved through optimization algorithms alone. Collaborative healthcare research requires clearly defined responsibilities concerning data access, system security, model updates, audit records, and permitted analytical purposes. Participating organizations may follow different internal policies and legal requirements. Accordingly, implementation should be supported by appropriate agreements and organizational review. Technical privacy mechanisms and institutional governance should function together to support responsible distributed learning rather than being treated as substitutes for one another.

8. Conclusion

Federated machine learning provides a methodological approach for collaborative analysis of healthcare information maintained across independent institutions. By performing selected training operations locally and coordinating model updates, distributed learning may reduce reliance on centralized repositories of sensitive patient records. The present hypothetical study examined this approach through a structured methodological framework addressing collaborative optimization, institutional data differences, and privacy-related safeguards. Its conclusions concern analytical principles and do not represent a completed clinical machine learning experiment.

The illustrative analysis demonstrated the importance of considering institutional sample distributions and evaluating predictive outcomes through clearly defined classification measures. Differences among the hypothetical learning configurations showed how accuracy, sensitivity, and F1 score can be calculated and compared. However, the assumed numerical outcomes cannot establish improvements attributable to federated learning or performance effects caused by privacy mechanisms. Such conclusions would require implemented algorithms, representative healthcare datasets, independent validation, and appropriately specified security assessments.

The study concludes that reliable privacy-conscious healthcare analytics requires coordination among predictive modeling, information security, clinical evaluation, and institutional governance. Retaining patient records locally represents a useful architectural consideration but does not independently guarantee confidentiality. Future research should investigate heterogeneous clinical datasets, secure aggregation, differential privacy, communication efficiency, and institution-level model performance. Reproducible empirical studies are necessary to determine when federated machine learning provides practical advantages while maintaining appropriate safeguards for sensitive healthcare information.

References

1. Abbas, S. R., Abbas, Z., Zahir, A., & Lee, S. W. (2024). Federated learning in smart healthcare: A comprehensive review on privacy, security, and predictive analytics with IoT integration. *Healthcare*, 12(24), Article 2587. <https://doi.org/10.3390/healthcare12242587>
2. Ali, M., Naeem, F., Tariq, M., & Kaddoum, G. (2023). Federated learning for privacy preservation in smart healthcare systems: A comprehensive survey. *IEEE Journal of Biomedical and Health Informatics*, 27(2), 778–789. <https://doi.org/10.1109/JBHI.2022.3181823>

3. Brisimi, T. S., Chen, R., Mela, T., Olshevsky, A., Paschalidis, I. C., & Shi, W. (2018). Federated learning of predictive models from federated electronic health records. *International Journal of Medical Informatics*, 112, 59–67. <https://doi.org/10.1016/j.ijmedinf.2018.01.007>
4. Crowson, M. G., Moukheiber, D., Robles Arévalo, A., Lam, B. D., Mantena, S., Rana, A., Goss, D., Bates, D. W., & Celi, L. A. (2022). A systematic review of federated learning applications for biomedical data. *PLOS Digital Health*, 1(5), Article e0000033. <https://doi.org/10.1371/journal.pdig.0000033>
5. Dang, T. K., Lan, X., Weng, J., & Feng, M. (2022). Federated learning for electronic health records. *ACM Transactions on Intelligent Systems and Technology*, 13(5), Article 72. <https://doi.org/10.1145/3514500>
6. Guendouzi, B. S., Ouchani, S., El Assaad, H., & El Zaher, M. (2023). A systematic review of federated learning: Challenges, aggregation methods, and development tools. *Journal of Network and Computer Applications*, 220, Article 103714. <https://doi.org/10.1016/j.jnca.2023.103714>
7. Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, 2, 305–311. <https://doi.org/10.1038/s42256-020-0186-1>
8. 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>
9. Li, Q., Wen, Z., Wu, Z., Hu, S., Wang, N., Li, Y., Liu, X., & He, B. (2023). A survey on federated learning systems: Vision, hype and reality for data privacy and protection. *IEEE Transactions on Knowledge and Data Engineering*, 35(4), 3347–3366. <https://doi.org/10.1109/TKDE.2021.3124599>
10. 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>
11. Li, W., Milletari, F., Xu, D., Rieke, N., Hancox, J., Zhu, W., Baust, M., Cheng, Y., Ourselin, S., Cardoso, M. J., & Feng, A. (2019). Privacy-preserving federated brain tumour segmentation. In H.-I. Suk, M. Liu, C. Lian, & P. Yan (Eds.), *Machine learning in medical imaging* (pp. 133–141). Springer. https://doi.org/10.1007/978-3-030-32692-0_16
12. Loftus, T. J., Ruppert, M. M., Shickel, B., Ozrazgat-Baslanti, T., Balch, J. A., Efron, P. A., Upchurch, G. R., Jr., Rashidi, P., Tignanelli, C., Bian, J., & Bihorac, A. (2022). Federated learning for preserving data privacy in collaborative healthcare research. *DIGITAL HEALTH*, 8. <https://doi.org/10.1177/20552076221134455>
13. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (Vol. 54, pp. 1273–1282). *Proceedings of Machine Learning Research*. <https://proceedings.mlr.press/v54/mcmahan17a.html>
14. Pati, S., Kumar, S., Varma, A., Edwards, B., Lu, C., Qu, L., Wang, J. J., Lakshminarayanan, A., Wang, S.-H., Sheller, M. J., Chang, K., Singh, P., Rubin, D. L., Kalpathy-Cramer, J., & Bakas, S. (2024). Privacy preservation for federated learning in health care. *Patterns*, 5(7), Article 100974. <https://doi.org/10.1016/j.patter.2024.100974>

15. 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>
16. Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R. 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>
17. Teo, Z. L., Jin, L., Li, S., Miao, D., Zhang, X., Ng, W. Y., Tan, T. F., Lee, D. M., Chua, K. J., Heng, J., Liu, Y., Goh, R. S. M., & Ting, D. S. W. (2024). Federated machine learning in healthcare: A systematic review on clinical applications and technical architecture. *Cell Reports Medicine*, 5(2), Article 101419. <https://doi.org/10.1016/j.xcrm.2024.101419>
18. Upreti, D., Yang, E., Kim, H., & Seo, C. (2024). A comprehensive survey on federated learning in the healthcare area: Concept and applications. *Computer Modeling in Engineering & Sciences*, 140(3), 2239–2274. <https://doi.org/10.32604/cmescs.2024.048932>
19. 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>
20. 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>