# Federated Learning for Privacy-Preserving Wearable Health Monitoring

Richard Koskinen
Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
richardk@buffalo.edu

Mason Berry
Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO, USA.
mberry@missouri.edu

**Abstract**

Wearable health monitoring has transitioned from episodic clinical observation to continuous, high-resolution physiological data collection, creating new opportunities for early detection, chronic disease management, and personalized intervention. At the same time, the centralized aggregation of such sensitive data introduces substantial privacy, security, and governance risks. Federated learning has emerged as a system-level alternative that enables distributed model training without transferring raw physiological records to a central repository. This paper presents an interdisciplinary examination of federated learning for privacy-preserving wearable health monitoring, emphasizing structural trade-offs rather than algorithmic details. The analysis addresses the architecture of wearable sensing pipelines, the role of edge aggregation, the limits of federated learning as a privacy mechanism, and the complementary use of differential privacy and secure aggregation. It further considers data heterogeneity across devices, patient populations, and behavioral contexts, together with challenges related to robustness, fairness, personalization, and long-term sustainability. The discussion incorporates healthcare deployment experiences, edge computing constraints, and governance frameworks that shape clinical and consumer adoption. While federated learning can reduce the necessity of raw data sharing, it does not eliminate privacy risk, nor does it automatically ensure clinical validity or operational resilience. A layered systems perspective is therefore required, spanning sensing hardware, communication infrastructure, model coordination, institutional policy, and regulatory compliance. The paper concludes by identifying future directions for resilient, equitable, and sustainable federated health monitoring infrastructures.

**Keywords**

federated learning, wearable health monitoring, privacy preservation, edge computing, system architecture, health data governance.

## 1. Introduction

The expansion of wearable biosensors has profoundly altered the temporal and contextual granularity of physiological data collection. Continuous monitoring now extends across electrocardiography, photoplethysmography, accelerometry, skin temperature, electrodermal activity, and oxygen saturation, producing longitudinal records that are far richer than those generated in clinic-based spot checks. These data support early warning systems, ambulatory monitoring of chronic conditions, and individualized behavior change interventions. However, the same properties that make wearable data clinically valuable also make them difficult to

govern. Physiological signals can reveal not only disease states but also behavioral patterns, emotional responses, and aspects of identity, while location traces and activity logs add further re-identification risk. Federated learning has been proposed as a way to reconcile analytical utility with privacy protection by keeping raw data on user devices and exchanging model updates rather than records [1]. This approach has attracted sustained interest in digital health because it aligns with the emerging principle that sensitive health data should remain as close as possible to the point of collection while still contributing to population-level learning [2].

Understanding federated learning for wearable health monitoring requires moving beyond the narrow view that privacy is achieved simply by avoiding data centralization. Federated systems introduce their own challenges, including communication bottlenecks, statistical heterogeneity, client availability, update poisoning, and uneven participation incentives [3]. Large-scale deployments outside health, such as mobile keyboard prediction, have demonstrated that federated learning can operate across millions of devices, but they also reveal the engineering effort required to manage unreliable clients and heterogeneous data distributions [4]. Wearable health monitoring adds further constraints related to energy budgets, intermittent connectivity, regulatory oversight, and the need for clinically interpretable outputs. These constraints transform federated learning from a model training strategy into a socio-technical infrastructure problem in which sensing, edge computation, coordination protocols, institutional governance, and clinical accountability must be designed together [5]. This paper therefore adopts a systems perspective, emphasizing architecture, trade-offs, robustness, fairness, deployment sustainability, and policy implications rather than focusing on optimization or model internals.

## 2. Privacy Threats and Architectural Context in Wearable Health Monitoring

Wearable health monitoring environments are characterized by distributed data generation across heterogeneous devices and continuous background sensing. This architecture creates multiple points of privacy vulnerability, including device compromise, inference during transmission, central database exposure, and unwanted secondary use by service providers. Federated learning addresses only one part of this threat surface: it reduces the need to store complete physiological records in a central server. However, model updates themselves can leak information about individual training data, and a centralized aggregator may still infer sensitive attributes from repeated updates. Formal privacy mechanisms such as differentially private stochastic gradient descent introduce calibrated noise during training to limit disclosure risk [6]. Differential privacy provides a quantifiable framework for reasoning about membership inference and attribute leakage, with foundational work articulating its guarantees, composition properties, and limitations in statistical analysis [7]. In wearable health systems, differential privacy must be balanced against clinical utility, because excessive noise can degrade the subtle signal patterns that are essential for detecting arrhythmias, sleep disturbances, or metabolic changes.

Beyond differential privacy, secure aggregation protocols reduce the risk that a coordinator can inspect individual client updates during model averaging. Practical secure aggregation enables multiple devices to contribute encrypted updates that are combined in such a way that the coordinator receives only the aggregated result [8]. This is particularly important in federated health networks where a university hospital, a technology firm, or a cloud provider may operate the aggregation service. When secure aggregation is combined with local training, it offers a stronger architectural guarantee against honest-but-curious aggregation

servers. Yet even secure aggregation does not address data quality, label noise, or malicious clients. Federated learning in medicine has been demonstrated most visibly through institutional-scale partnerships, where hospitals retain patient data while jointly training models across sites [9]. These examples show that federation can align with institutional governance and data minimization requirements, but they also reveal implementation barriers including institutional review protocols, heterogeneous electronic health record systems, and uneven computing capabilities across clinical sites.

Wearable health monitoring differs from hospital-centric federation in several important ways. Devices are numerous, computationally constrained, battery powered, and often connected through unstable wireless networks. The edge layer provides intermediate compute and storage, but it also creates new trust boundaries because edge nodes may be owned by network operators or smart home service providers. Recent analyses of federated learning in mobile edge networks highlight the interdependence between model training, bandwidth availability, and network latency [11], while surveys of smart healthcare report a rapidly growing but fragmented design space spanning communication efficiency, personalization, and security [12]. The data themselves are also more variable than institutional clinical data. Personal sensing research has shown that signals collected in daily life are shaped by movement, emotional state, social context, and environmental noise, making them simultaneously valuable and difficult to standardize [13]. This variability increases the importance of large and diverse datasets for learning robust representations, a concern mirrored in broader machine learning research showing that data scale can be as important as model design for achieving high performance [14]. Wearable health systems therefore require an architecture that preserves data locality while still enabling sufficient statistical diversity for generalization.

## 3. System Architecture and Deployment Considerations

A federated wearable health monitoring infrastructure typically consists of sensing devices, local model training on user devices or personal gateways, an aggregation service, and downstream clinical applications. This layered structure creates a series of design choices that affect privacy, energy use, latency, and clinical relevance. The sensing layer itself has improved significantly in recent years. For example, rapid vital sign extraction from opto-physiological monitoring under varying physical activity intensity supports more reliable real-time analytics at the edge [15]. Such advances are important because federated models trained on noisy or poorly synchronized signals cannot yield trustworthy health predictions, regardless of how well the privacy architecture is designed. The placement of inference and training is also critical. Training on a smartwatch may preserve data locality but limits model complexity and consumes battery, while offloading to a nearby phone or home gateway improves computational capacity but changes the trust boundary. Hybrid deployments in which lightweight feature extraction occurs on the device and heavier model updates occur on a personal edge node can balance these constraints, but they require careful synchronization and lifecycle management.

Security and robustness concerns become acute in federated health monitoring because a small number of compromised devices or fabricated updates can distort the aggregated model. Existing work on poisoning attacks in healthcare machine learning demonstrates that medical applications present attractive targets for adversarial manipulation, with consequences that range from misdiagnosis to insurance fraud [16]. Federated backdoor attacks further show that malicious clients can embed specific trigger patterns into a global model while

maintaining acceptable performance on normal inputs [17]. These risks are not theoretical; they are structural outcomes of distributed training systems in which local data are not directly inspectable by the coordinator. Defensive strategies include robust aggregation rules, anomaly detection over update vectors, and periodic audits of client contributions. However, such defenses must be implemented without compromising the privacy objectives of the system, because inspecting updates too closely may reintroduce disclosure risks. The result is a continuous tension between transparency for validation and opacity for privacy.

Personalization adds another architectural challenge. A global model trained across many users may perform poorly for individuals with rare conditions, unusual physiological baselines, or distinct behavioral patterns. Federated transfer learning and personalized federated frameworks for wearable healthcare attempt to preserve shared representations while adapting to local user-specific characteristics [19]. In-home health monitoring systems such as FedHome extend this idea by using cloud-edge collaboration to create personalized models while limiting data exposure [20]. These approaches suggest that effective wearable health monitoring will not rely on a single monolithic global model. Instead, the infrastructure will need to support a hierarchy of models, including generalized population-level models, cohort-specific models, and user-specific adaptation layers. Each layer raises distinct evaluation questions: a population model may be evaluated on external cohorts, whereas a personalized model must be evaluated on longitudinal within-person stability and responsiveness to behavioral change.

Equity and fairness considerations are equally pressing. Wearable devices are distributed unevenly across socioeconomic, geographic, and demographic groups, and sensor performance can vary by skin tone, body composition, and activity type. Federated aggregation can exacerbate these disparities if the coordinator weights clients equally while some groups are underrepresented or if device heterogeneity creates systematic differences in gradient quality. Reviews of bias and fairness in machine learning identify numerous mechanisms through which data collection, labeling, and model design can produce discriminatory outcomes [18]. In federated health monitoring, fairness must be assessed not only in terms of aggregate accuracy but also in terms of equalized error rates across populations and consistent clinical decision support. This requires careful stratification of client populations, participatory design with diverse user communities, and auditing frameworks that persist across model versions. Technical aggregation alone cannot resolve structural inequities in device access and data representation.

## 4. Governance, Sustainability, and Policy Implications

The deployment of federated learning in wearable health monitoring takes place within complex regulatory and institutional environments. In many jurisdictions, physiological data collected by consumer wearables do not automatically receive the same legal protections as data generated in clinical care, even though their sensitivity may be equivalent. Federated architectures can help organizations meet data minimization obligations under health privacy frameworks, but they also require clear policies for consent, withdrawal, model provenance, and secondary use. Because federated systems obscure raw data from central servers, regulators and auditors may face new difficulties in verifying data quality, detecting unlawful discrimination, or reconstructing model decisions. Therefore, federated learning should be accompanied by governance mechanisms such as federated data registries, model cards, audit trails, and privacy impact assessments. These mechanisms convert privacy architecture into institutional accountability.

Sustainability is another underappreciated dimension. Wearable devices have limited hardware life cycles, and federated training consumes battery, bandwidth, and cloud computing resources. Continuous monitoring can reduce device utility if training interferes with user-facing features or accelerates hardware degradation. From an energy systems perspective, federated learning may not always be more efficient than centralized training, particularly when many communication rounds are required to reach convergence. Model compression, sparse updates, asynchronous aggregation, and intermittent participation can reduce resource demands, but they introduce additional complexity in convergence monitoring and clinical validation. Sustainable deployment also requires organizational continuity: a federated health monitoring service must maintain secure identity management, participant recruitment, hardware compatibility, and model versioning over years of operation. Without such lifecycle planning, federated health projects risk becoming short-term research prototypes that cannot support clinical or public health functions.

The choice between consumer wearable ecosystems and clinically regulated monitoring systems also shapes federation design. Consumer systems tend to prioritize scale, ease of use, and continuous engagement, while clinical systems emphasize validation, traceability, and integration with electronic health records. Federated models trained in consumer environments may not transfer directly to clinical use without recalibration and regulatory review. Conversely, models developed in supervised clinical studies may fail under the noisy, uncontrolled conditions of daily life. Cross-domain governance is therefore needed to bridge these worlds, establishing pathways for responsible translation from research prototypes to regulated health devices. Standardization efforts in federated interfaces, metadata exchange, and outcome definitions would support such translation, as would clearer guidance from health regulators on the evidentiary standards for continuously updated distributed models.

Finally, trust is a foundational governance issue. Users may be asked to consent to federated participation without fully understanding how their local data contribute to model improvement or how their privacy is protected. If federated learning is presented as a purely technical solution to privacy, it may create a false sense of security and undermine informed consent. Effective communication must acknowledge residual risks, including inference from model updates, device compromise, and adversarial participation. At the same time, institutional actors must demonstrate that federated systems produce tangible benefits for users, such as personalized feedback, better clinical risk stratification, or improved device functionality. The social acceptance of federated health monitoring will depend not only on technical safeguards but also on transparent governance, equitable benefit sharing, and meaningful user control over data contribution.

## 5. Conclusion

Federated learning offers a compelling system-level response to the privacy challenges of wearable health monitoring by enabling distributed model training without centralizing raw physiological records. However, its adoption cannot be reduced to an algorithm replacement. The structural demands of wearable sensing, edge computation, secure aggregation, differential privacy, model personalization, and robust coordination require an integrated design perspective that spans devices, networks, institutions, and regulatory regimes. Privacy is not a property achieved once by design; it is an ongoing socio-technical accomplishment that depends on adversarial assumptions, lifecycle management, and governance practices. At the same time, federated architectures introduce new risks related to data heterogeneity, malicious clients, fairness degradation, and operational fragility. Future systems must

therefore be evaluated not only by predictive accuracy but also by their resilience under population drift, their energy sustainability, their capacity to support clinical accountability, and their equitable performance across diverse user communities. Advances in signal extraction, edge computing, and personalized federated methods will continue to improve feasibility, but meaningful progress will require interdisciplinary collaboration among engineers, clinicians, regulators, and user communities. The long-term value of federated learning for wearable health monitoring lies not merely in preserving privacy but in constructing trustworthy infrastructures that sustain both analytical insight and social legitimacy.

**References**


1. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 1273–1282.

2. Rieke, N., Hancox, J., Li, W., Milletarì, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M. J., 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.

3. 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., Rouayheb, S. E., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., Gruteser, M., Harchaoui, Z., He, C., He, L., Huo, Z., Hutchinson, B., Hsu, J., Jaggi, M., Javidi, T., Joshi, G., Khodak, M., Konečný, J., Korolova, A., Koushanfar, F., Koyejo, S., Lepoint, T., Liu, Y., Mittal, P., Mohri, M., Nock, R., Özgür, A., Pagh, R., Raykova, M., Qi, H., Ramage, D., Raskar, R., Song, D., Song, W., Stich, S. U., Sun, Z., Suresh, A. T., Tramèr, F., Vepakomma, P., Wang, J., Xiong, L., Xu, Z., Yang, Q., Yu, F. X., Yu, H., & Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210.

4. Hard, A., Rao, K., Mathews, R., Ramaswamy, S., Beaufays, F., Augenstein, S., Eichner, H., Kiddon, C., & Ramage, D. (2018). Federated learning for mobile keyboard prediction. arXiv. https://arxiv.org/abs/1811.03604

5. 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.

6. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, 308–318.

7. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407.

8. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 1175–1191.

9. 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.

10. 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.

11. Lim, W. Y. B., Luong, N. C., Hoang, D. T., Jiao, Y., Liang, Y.-C., Yang, Q., Niyato, D., & Miao, C. (2020). Federated learning in mobile edge networks: A comprehensive survey. IEEE Communications Surveys & Tutorials, 22(3), 2031–2063.

12. Nguyen, D. C., Pham, Q.-V., Pathirana, P. N., Ding, M., Seneviratne, A., Lin, Z., Dobre, O. A., & Hwang, W.-J. (2022). Federated learning for smart healthcare: A survey. ACM Computing Surveys, 55(3), Article 60.

13. Mohr, D. C., Zhang, M., & Schueller, S. M. (2017). Personal sensing: Understanding mental health using ubiquitous sensors and machine learning. Annual Review of Clinical Psychology, 13, 23–47.

14. Sun, C., Shrivastava, A., Singh, S., & Gupta, A. (2017). Revisiting unreasonable effectiveness of data in deep learning era. Proceedings of the IEEE International Conference on Computer Vision, 843–852.

15. Zheng, X., Dwyer, V. M., Barrett, L. A., Derakhshani, M., & Hu, S. (2023). Rapid vital sign extraction for real-time opto-physiological monitoring at varying physical activity intensity levels. IEEE Journal of Biomedical and Health Informatics, 27(7), 3107-3118.

16. Mozaffari-Kermani, M., Sur-Kolay, S., Raghunathan, A., & Jha, N. K. (2015). Systematic poisoning attacks on and defenses for machine learning in healthcare. IEEE Journal of Biomedical and Health Informatics, 19(6), 1893–1905.

17. Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020). How to backdoor federated learning. Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics, 2938–2948.

18. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. ACM Computing Surveys, 54(6), Article 115.

19. 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.

20. Wu, Q., Chen, X., Zhou, Z., & Zhang, J. (2020). FedHome: Cloud-edge based personalized federated learning for in-home health monitoring. IEEE Transactions on Mobile Computing, 21(8), 2818–2832.