Federated Learning–Enabled Privacy-Preserving PPG Foundation Models for Intelligent Healthcare Agent Systems

Massimo A. Hunt
Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.
massimo.work@unr.edu

Abstract

The convergence of photoplethysmography as a near-ubiquitous non-invasive sensing modality, large-scale health foundation models, and intelligent agent systems promises a transformative leap toward personalized, longitudinal, and scalable healthcare. However, the centralization of sensitive physiological data introduces profound privacy risks and increasingly stringent regulatory constraints. This paper presents a system-level exploration of federated learning–enabled privacy-preserving PPG foundation models designed to serve as sensing and representation backbones within intelligent healthcare agent architectures. We examine the complete pipeline, from on-device self-supervised pre-training of a shared PPG foundation model across heterogeneous wearable populations to its privacy-preserving aggregation and integration into multi-agent reasoning frameworks that deliver context-aware health monitoring, triage, and decision explanation. The analysis foregrounds structural trade-offs among communication efficiency, model utility, differential privacy guarantees, secure aggregation, and fairness across diverse demographic and device strata. Particular attention is given to the governance, sustainability, and deployment implications of such a system, including Byzantine resilience, regulatory compliance with the GDPR and HIPAA, and the challenges of continual adaptation on resource-constrained edge devices. Through conceptual synthesis of recent advances and critical system design perspectives, we argue that a privacy-by-design federated foundation model combined with modular agent orchestration can form the trustworthy infrastructural core of next-generation intelligent healthcare ecosystems. Open research directions concerning adversarial robustness of agent reasoning, cross-silo governance, and post-deployment lifecycle management are identified to inform future interdisciplinary efforts.

Keywords

Federated learning; PPG foundation model; Privacy-preserving AI; Intelligent healthcare agents; System governance; Robustness.

1. Introduction

The rapid proliferation of consumer and clinical wearable devices equipped with photoplethysmography sensors has created an unprecedented longitudinal data stream capturing hemodynamic variations, heart rate variability, oxygen saturation, and surrogate indicators of autonomic and cardiovascular state. In parallel, the rise of foundation models in machine learning, capable of learning high-capacity, generalizable representations from massive unlabeled data, has begun to reshape the biomedical signal processing landscape. A pre-trained PPG foundation model promises to unify tasks as diverse as arrhythmia detection,

blood pressure estimation, stress quantification, and sleep staging under a single adaptable representation, dramatically reducing the need for task-specific labelled data. Simultaneously, intelligent healthcare agent systems built around large language models are emerging as interactive, reasoning-capable mediators between raw physiological data and clinical decision support, offering patients and providers conversational insight, triage recommendations, and longitudinal trend analysis. Integrating these three technological vectors into a cohesive system could realize a paradigm in which wearable-derived PPG signals are transformed into actionable, privacy-respected health intelligence.

Yet this vision collides with fundamental tensions surrounding data privacy and sovereignty. PPG waveforms inherently contain identifiable morphological patterns and can reveal sensitive health conditions; centralizing such data for foundation model training not only violates the principle of data minimization enshrined in modern privacy regulations but also creates a high-value target for adversarial and inferential attacks. Federated learning has emerged as the principal architectural response, enabling collaborative model training across decentralized data silos without raw data exchange. Although federated learning has been studied extensively in general machine learning and in certain health informatics contexts, its intersection with the unique demands of PPG foundation models and their downstream deployment into intelligent agent systems remains underexplored from a systems perspective. The challenges are not purely algorithmic but involve deep structural trade-offs concerning communication budgets, statistical heterogeneity of sensor populations, the embedding of differential privacy and secure aggregation in resource-constrained wearables, fairness across demographic groups, and the governance of an evolving model that becomes the perceptual foundation of safety-critical agent behavior.

This paper examines these interconnected dimensions through an integrative system-level lens. We deliberately avoid mathematical formalisms and algorithmic recipes, instead concentrating on architectural reasoning, cross-domain comparison, and governance-aware design. The following sections unpack the motivations and background, propose a federated learning architecture tailored to PPG foundation models, delve into privacy-preserving mechanisms, explore the integration with intelligent agents, and assess fairness, sustainability, and deployment realities. Throughout, we emphasize that building a trustworthy, large-scale, privacy-preserving PPG intelligence fabric is not solely a machine learning challenge but a socio-technical infrastructure problem demanding interdisciplinary collaboration among engineers, clinicians, regulators, and ethicists.

2. Background and Motivations

PPG is an optical sensing technique that measures volumetric changes in microvascular blood flow, typically via a light source and photodetector integrated into wearables such as smartwatches, finger clips, and earbuds. The resulting waveform, while lower in signal-to-noise ratio than electrocardiography, captures rich physiological correlates that make it attractive for continuous, unobtrusive monitoring. Foundation models for time-series data have recently leveraged self-supervised objectives such as masked signal modelling and contrastive predictive coding to extract universal representations analogous to those achieved by large language models in natural language processing. In the PPG domain, SIGMA-PPG exemplifies this trend by incorporating statistical-prior informed generative masking to build a robust encoder capable of handling signal corruption and inter-subject variability [3]. Concurrent advances such as PPG-MAE further demonstrate how masked autoencoder pretraining can yield representations that transfer effectively across downstream health tasks

[4]. These foundation models are poised to serve as shared feature extractors that decouple the low-level signal processing from high-level clinical inference, thereby enabling rapid adaptation to novel applications with minimal labelled data.

Federated learning provides the necessary privacy framework for training such models without consolidating sensitive PPG records. The canonical Federated Averaging algorithm [1] coordinates a population of client devices each holding local data; clients compute model updates on their private data and transmit only these updates, after which a central server aggregates them to improve a global model. Comprehensive surveys have catalogued the manifold challenges that arise in this setting, including non-identically distributed data across clients, systems heterogeneity, and communication bottlenecks [2]. In the PPG context, these challenges are amplified by the wide variability in sensor quality, sampling rates, ambient noise conditions, and the diverse physiological profiles of populations across geography, age, gender, and skin pigmentation. Consequently, a federated PPG foundation model must contend with a highly non-IID landscape where some clients generate clean, clinically-validated signals while others produce motion-corrupted, low-fidelity streams.

Privacy threats in federated learning extend beyond the mere withholding of raw data. Sophisticated gradient inversion and membership inference attacks can reconstruct sensitive inputs or determine whether a particular individual's data participated in training. Differential privacy [5] offers a rigorous probabilistic guarantee, typically implemented by clipping per-sample gradients and injecting calibrated noise during local training, ensuring that the contribution of any single individual is obscured. Secure aggregation protocols [6] complement differential privacy by cryptographically hiding individual updates from the central server, revealing only the aggregated result. When combined, these mechanisms form a defense-in-depth strategy highly relevant to health data governed by frameworks such as the GDPR. The broader vision of federated learning for digital health, articulated by Rieke et al. [7], envisions a world in which hospitals, research institutions, and consumer device ecosystems collaborate on model development while retaining full control of their data assets. Specialized systems like FedPPG have already begun to demonstrate the feasibility of federated self-supervised learning for PPG-based stress detection [8], but these efforts have not yet addressed the scale, representational depth, and system-wide implications of a full foundation model.

Furthermore, the reliability of any PPG-driven health system depends on the quality of the input signal. Established signal quality indices [9] must be integrated into both client-side preprocessing and the model training loop to filter out unreliable segments and to condition the representation on confidence estimates. In a federated setting, the distribution of signal quality across clients becomes a fairness dimension, if models inadvertently learn to rely on high-quality data from affluent demographics with premium devices while underperforming for underserved populations using older sensors. All these facets motivate a system architecture that is not only technically proficient but also socio-technically aware, prompting the detailed architectural discussion that follows.

3. Federated Learning Architecture for PPG Foundation Models

Designing a federated system for PPG foundation models requires orchestrating edge-side training, aggregation strategy, and model partitioning to balance representation quality, communication cost, and privacy protection. The architecture we envisage comprises a population of wearable devices, each embedding a local instance of the foundation model encoder, and a coordinating server that may be hosted in a cloud or a more privacy-sensitive

edge data center within a hospital trust. Training proceeds in rounds: the server broadcasts the current global model to a sampled subset of clients, each client performs several local epochs of self-supervised learning on its private PPG stream, and the resulting model updates are communicated back. Aggregation at the server, typically via Federated Averaging, produces a new global model that encapsulates the collective knowledge of the participating population without ever observing the raw data.

The central structural tension in this design arises from the opposing pulls of communication efficiency and model convergence. Large foundation models, which can easily exceed a hundred million parameters, impose a heavy communication footprint that is often prohibitive for battery-powered wearables connected via intermittent Bluetooth or Wi-Fi. Techniques such as model compression, gradient sparsification, and quantization are indispensable, yet they introduce additional noise that can degrade the learned representations. The number of local epochs executed before each update is a critical knob: more local training reduces communication frequency but exacerbates client drift, especially under non-IID data distributions, potentially leading to a global model that fails to generalize across the entire device fleet. Hierarchical federated architectures that introduce edge aggregation servers, situated for example at the level of a hospital ward or a telecom base station, can alleviate the core-network burden and enable more responsive model personalization by grouping clients with similar data characteristics.

Fairness considerations must be woven into the architecture, not treated as an afterthought. When data distributions are imbalanced across demographic groups, standard federated averaging can produce a model that performs well for the majority while systematically disadvantaging minority populations. Li et al. [10] have introduced fairness-aware aggregation objectives that re-weight client contributions to mitigate such disparities. Incorporating these objectives into the PPG domain requires careful operationalization of fairness, as group membership may be defined by factors such as device type, skin tone, or clinical condition, many of which the federated protocol is not permitted to access directly due to privacy constraints. Integrating explainability methods, such as Shapley value-based feature attribution [11], into the local or server-side auditing process provides a pathway to detect and remedy representation biases without compromising the privacy membrane.

The architecture must also support continual learning. Wearable firmware, sensor characteristics, and population health profiles evolve over time. The global model should therefore be updated incrementally, with safeguards against catastrophic forgetting and concept drift. A promising direction is the adoption of modular adapter layers that remain client-specific while the bulk of the foundation model is shared; this strikes a balance between global knowledge consolidation and local personalization. The system should further incorporate signal quality indices [9] at the client side to gate training examples, preventing the model from overfitting to noisy or artifactual segments and preserving the robustness that is essential for downstream agent systems. Having established this federated backbone, the next critical layer is the rigorous embedding of privacy-preserving mechanisms that ensure the architecture fulfills its data protection promises.

4. Privacy-Preserving Mechanisms and Systemic Implications

Privacy in a federated PPG foundation model is not a monolithic guarantee but a multi-layered property that must be engineered, governed, and continuously verified. While the core federation protocol already prevents raw data from leaving the device, model updates themselves can leak information, necessitating formal privacy amplification. Differential

privacy, instantiated through DP-SGD [5], provides a measurable privacy budget by clipping per-example gradients and adding calibrated Gaussian noise during local training. When local DP is applied before aggregation, the server receives anonymized update statistics, and even the aggregated global model enjoys a degree of privacy amplification through the composition of many noisy contributions. Secure aggregation [6] fortifies this further by ensuring that the server learns only the aggregated update, never any individual model delta, even if the server is compromised. Together, this combination establishes a strong defense-in-depth posture.

However, the system-level repercussions of these protections are profound. The noise injected by DP-SGD reduces the effective signal of each local update, slowing convergence and potentially eroding the representational fidelity of the foundation model, particularly for rare physiological patterns such as infrequent arrhythmias that are inherently sparse in client populations. This privacy-utility trade-off interacts with fairness: if a fixed global privacy budget is applied uniformly, minority subpopulations whose data are already scarce may suffer disproportionate utility loss. Adaptive privacy allocation, where clients with more sensitive data or those contributing to rare classes employ a tighter privacy budget while others use a more relaxed one, is an active area of system design that requires careful orchestration to avoid introducing new biases.

Beyond the core DP and aggregation protocols, the integration of the PPG foundation model into intelligent healthcare agent systems raises adversarial robustness concerns. Agents that ingest PPG-derived embeddings to generate clinical recommendations are susceptible to adversarial perturbations, both at the signal level and through crafted textualprompts designed to mislead reasoning or circumvent safety filters. Because the agent's clinical logic relies on the fidelity of the embeddings produced by the perceptive foundation model, an adversary who perturbs the PPG signal with carefully crafted noise may induce erroneous feature representations that cascade into dangerous recommendation errors. At the language interface, adversarially constructed prompts could manipulate the agent into disregarding critical contextual information or divulging inferential outputs contrary to privacy constraints. The intersection of these vulnerabilities demands that the privacy-preserving architecture be extended with robust training and monitoring capabilities, a requirement that recent research on security enhancement for adversarial robust large language model agents in medical decision-making tasks has brought into sharp focus [12]. While that line of work primarily targets the text-based reasoning layer, its implications for multi-modal agent systems that integrate perceptual streams are immediate: the entire inference pipeline must be hardened, and the foundation model must be trained not only for accuracy under natural distribution but also under adversarial signal transformations that respect the physical constraints of PPG sensors. Consequently, the federated training protocol should incorporate adversarial augmentation and certified robustness objectives that are themselves privacy-compliant, a delicate multi-objective optimization that warrants further system-level investigation.

5. Integration with Intelligent Healthcare Agent Systems

Once a privacy-preserving PPG foundation model is trained and operational, the next architectural dimension concerns its integration into intelligent healthcare agents that turn raw representations into explainable, personalized, and actionable health intelligence. In our envisioned system, the agent layer consists of a modular set of reasoning components orchestrated around a large language model core. The foundation model serves as a dedicated perception module, consuming raw PPG time-series inputs from the user's wearable and converting them into a compact, privacy-processed embedding vector. This embedding is then

contextualized with other available data modalities, such as electronic health record extracts, self-reported symptoms, environmental sensor readings, and activity logs, within the agent's prompt construction and reasoning workflow. The agent may be deployed in a tiered fashion: a lightweight on-device version processes embeddings to provide immediate, low-latency wellness feedback, while a more powerful cloud-side instance handles deeper diagnostic reasoning, long-term trend analysis, and coordination with clinical decision support systems under strict data governance agreements.

The coupling of the foundation model and the agent architecture forces a reexamination of the conventional boundaries between sensing, learning, and reasoning. Because the PPG foundation model was trained under a federated privacy budget, its internal representations are already shaped by noise calibration, making them inherently somewhat blurred with respect to fine-grained individual details. This property can be leveraged to design a "privacy-preserving input layer" for the agent: the embedding itself carries a quantifiable disclosure risk, and the agent's responses can be further post-processed or constrained by a privacy policy layer that limits the granularity of shared insights. For example, the agent might be instructed never to output exact heart rate variability numbers if the associated embedding retains too much identifiable information, but rather to communicate risk categories or trends. Such a design harmonizes with data minimization principles and aligns with the accountability requirements of the GDPR and HIPAA, where automated decision-making must be accompanied by meaningful explanations and the ability for individuals to contest decisions.

The agent system further introduces a spatial and temporal orchestration layer. In a practical scenario, an individual might interact with the agent through a smartphone application that receives model updates from a hospital-provided PPG foundation model, while also connecting to a pharmacist-side agent for medication adherence monitoring and a cardiologist-side agent for arrhythmia surveillance. Each agent instance may access the same PPG embedding but apply different inference heads and reasoning policies tailored to its domain. Federated learning can extend to these downstream modules as well, allowing each clinical specialty to fine-tune a lightweight prediction head on top of the frozen representation without sharing the underlying raw data. This layered federation promotes modular verification: the safety properties of each agent can be audited independently, while the shared representation ensures cross-domain consistency. The agent architecture must also implement robust context management to track dialogue states, maintain continuity of care across multiple sessions, and gracefully handle ambiguous or conflicting user prompts without overriding professional clinical protocols. Research on fitting intelligent decision support into critical clinical workflows emphasizes that agents should operate as augmentative, not autonomous, entities, always deferring to licensed professionals when confidence bounds are exceeded or when high-stakes clinical decisions are involved [20].

6. Structural Trade-offs in Fairness, Robustness, and Continual Adaptation

Deploying a federated PPG foundation model within a heterogeneous and evolving population exposes deep structural trade-offs that connect fairness, robustness, and the ability to adapt over time. Fairness in this context is multi-dimensional: it includes parity of predictive performance across demographic groups delineated by skin tone, age, and comorbid conditions, as well as equity in the quality of service provided by the agent system itself. The root of many fairness concerns lies in the training data distribution, which in a federated setting is inherently determined by the device fleet composition and user opt-in

patterns. Previous studies have established that commercial pulse oximeters and PPG sensors exhibit decreased signal-to-noise ratio in individuals with darker skin pigmentation [17], and age-related morphological changes in the pulse waveform further differentiate subgroups [24]. If a standard federated averaging protocol is employed without corrective measures, the global model will gravitate toward the characteristics of the majority client subgroup, reinforcing existing health disparities.

Addressing this requires going beyond simple re-weighting heuristics. The architecture must incorporate statistically grounded fairness constraints into the aggregation objective, as explored in federated optimization frameworks that allow clients with disadvantaged data distributions to exert greater influence on the global model shape [10]. Yet imposing such constraints inevitably interacts with the privacy mechanism and communication budget. Stronger representation for minority subgroups often necessitates that those clients contribute more frequent or less noised updates, which in turn increases their privacy risk and widens the overall privacy-utility frontier for the system. A possible middle ground involves stratified secure aggregation, where clients are grouped by self-declared or inference-avoidable characteristics, and separate sub-models are maintained for distinct population strata, later combined via a meta-aggregation step. This approach, however, raises its own governance questions about categorizing individuals and must be implemented with transparency and ethical oversight. Beyond model training fairness, the agent layer must be monitored for differential treatment in its interactive behavior, ensuring that it does not respond with systematically less informative guidance to certain user groups due to latent biases learned from the representation.

Robustness concerns span multiple strata of the system. At the foundation model level, the PPG encoder must be resistant not only to naturally occurring motion artifacts and sensor noise, but also to adversarial perturbations that could be introduced by a malicious on-device application or a compromised device. Byzantine-robust aggregation rules that filter out anomalous updates from adversarial or faulty clients have been shown to protect federated learning against targeted model poisoning [21]. Incorporating such rules into the federated PPG training loop is non-trivial because the heterogeneity of legitimate clients can make it difficult to distinguish a genuinely atypical but benign update from a malicious one without auxiliary quality signals. Signal quality indices [9] can serve as a preliminary filter, but they themselves can be spoofed. Multi-layered defenses combining robust statistics, update clipping, and anomaly detection on embedding geometry are essential. At the agent level, robustness must be extended to prompt-level adversarial inputs and to data fusion risks. An adversary might exploit the agent's ability to combine PPG embeddings with free-text symptom reports by inserting contradictory or manipulative information, a scenario that underscores the relevance of the adversarial robust agent design principles articulated in recent medical agent security literature [12]. Systematic red-teaming and continuous monitoring must be institutionalized as core operational practices.

Continual adaptation links fairness and robustness through the temporality of deployment. The global PPG foundation model, once deployed, cannot remain static, because sensor hardware evolves, population demographics shift, and new disease phenotypes emerge. Federated continual learning strategies that allow the global model to absorb novel information without catastrophically forgetting previously acquired competencies are essential for sustained clinical utility [13] [23]. Yet every model update cycle reopens the privacy budget and may inadvertently shift the model's fairness profile, if, for instance, newly

onboarded devices are predominantly from a different socioeconomic stratum than the original cohort. Consequently, the governance framework must mandate periodic privacy audits, fairness impact assessments, and adversarial robustness recertification as part of the model lifecycle. The technical tension between model plasticity needed for adaptation and the stability required for regulatory certification is a defining systems challenge that extends well beyond algorithm design into legal and organizational policy.

7. Deployment, Sustainability, and Governance

Translating a federated PPG agent system from a proof-of-concept into a sustainable, large-scale healthcare infrastructure demands attention to a range of practical deployment considerations. The energy and computational footprint of on-device training is a primary sustainability concern. Wearables and smartphones possess limited battery and thermal budgets; performing multiple local epochs of a large transformer-based foundation model can be prohibitive. Several mitigation strategies can be integrated. Hierarchical federated learning architectures that offload part of the computation to nearby edge gateways, such as home hubs or hospital-edge servers, can reduce the client-side burden while preserving data locality [14]. Split learning, in which the foundation model is partitioned between a client-side compressor and an edge-side encoder, offers another path to reduce on-device workload, though it introduces latency dependencies and requires careful analysis of the information leakage from communicated activations. Additionally, parameter-efficient fine-tuning techniques can confine adaptation to a few trainable parameters, dramatically lowering per-round computational cost and communication volume. The sustainability calculus must also account for the server-side energy consumption of secure aggregation protocols and the carbon footprint of the overall training lifecycle, aligning with broader calls for green AI [22].

Governance of a federated PPG foundation model and the associated agent ecosystem requires a multi-stakeholder framework that addresses ownership, accountability, and auditability. The global model is a communal artifact co-produced by many data contributors; questions of data dividends, intellectual property, and liability in the event of a harmful agent error are unresolved. From a regulatory perspective, the system as a whole constitutes a medical device software pipeline under frameworks such as the FDA's SaMD (Software as a Medical Device) guidance and the European Medical Device Regulation. Since model updates occur continuously, a locked pre-market approval model is insufficient; instead, a lifecycle regulatory approach that includes predefined post-market surveillance, change control protocols, and real-world performance monitoring is required. Privacy compliance further demands that the system operationalize data subject rights under GDPR, including the right to deletion, which in a federated context translates into the ability to unlearn an individual's contribution from the global model. Machine unlearning algorithms remain nascent but are critical for legal compliance, adding yet another layer of complexity to the aggregation and update pipeline [18].

The interplay between technical design and governance is perhaps most visible in the allocation of differential privacy budgets across clients and over time. A centralized governance body, such as a data trust or a consortium of healthcare providers, could define a global privacy loss parameter, but implementing this equitably across a diverse device fleet requires technical policy engines that adapt the privacy noise level based on device capability, data sensitivity, and user consent preferences. Cryptographic key management for secure aggregation over large, dynamically changing cohorts also demands robust public key infrastructure and key rotation policies that do not become single points of failure. The

integration of intelligent agents adds further governance dimensions: agent outputs must be traceable to the provenance of the foundation model embeddings and the training data that shaped them, enabling post-hoc accountability when decisions are contested. Model cards, data nutritional labels, and agent behavior testing suites are emerging as templates for transparent documentation that can satisfy both regulatory scrutiny and user trust [19]. Ultimately, the sustainability of the entire ecosystem depends on building institutional trust through demonstrable, independently audited privacy and fairness guarantees rather than relying solely on technical assurances.

8. Conclusion

This paper has presented a system-level analysis of federated learning–enabled privacy-preserving PPG foundation models as the perceptive core of intelligent healthcare agent systems. Starting from the motivations of sensitive physiological data and the need for large-scale, generalizable representations, we walked through the federated architecture, privacy-preserving mechanisms, agent integration, and the structural trade-offs that govern fairness, robustness, and continual deployment. We argued that protecting privacy in such a system is not achieved by any single technique but by an orchestrated combination of federated topology, differential privacy, secure aggregation, and policy-aware agent design, each of which introduces tensions with utility, communication efficiency, and equitable performance. The embedding of a DP-trained PPG foundation model into a multi-agent reasoning environment promises to deliver longitudinal, non-intrusive health intelligence that respects individual sovereignty, but only if the system is accompanied by rigorous governance, lifecycle monitoring, and adversarial robustness measures that are currently incomplete.

The complexity of this socio-technical stack highlights the importance of interdisciplinary collaboration. Engineers must work with clinicians to define clinically relevant robustness criteria, with legal scholars to encode privacy policies into model architectures, and with ethicists to design fairness interventions that do not inadvertently essentialize sensitive attributes. Future work must advance practical unlearning in federated settings, formalize the composition of differential privacy across training and inference, and develop certified robustness bounds for multi-modal agent recommendations. Addressing these challenges can pave the way toward a digital health infrastructure that is at once powerful and respectful, capable of turning the continuous stream of PPG data into a trusted companion for lifelong wellness without ever betraying the trust of the individuals it serves.

References

1. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In Artificial Intelligence and Statistics (pp. 1273–1282). PMLR.

2. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210.

3. Guo, Z., Chen, T., Jiao, Y., Pan, Y., Hu, X., & Ferrario, M. (2026). SIGMA-PPG: Statistical-prior Informed Generative Masking Architecture for PPG Foundation Model. arXiv preprint arXiv:2601.21031.

4. Zhou, Y., Liu, J., Wang, X., & Li, Z. (2023). PPG-MAE: Self-supervised learning of photoplethysmography signals with masked autoencoders. IEEE Journal of Biomedical and Health Informatics, 27(9), 4512–4523.

5. 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).

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

7. Rieke, N., Hancox, J., Li, W., Milletarì, F., Roth, H. R., Albarqouni, S., ... & Cardoso, M. J. (2020). The future of digital health with federated learning. NPJ Digital Medicine, 3(1), 119.

8. Deng, Y., Liu, S., & Chen, Y. (2023). FedPPG: A federated learning framework for privacy-preserving PPG-based stress monitoring. IEEE Access, 11, 45326–45338.

9. Orphanidou, C., Bonnici, T., & Charlton, P. (2015). Signal-quality indices for the electrocardiogram and photoplethysmogram: Derivation and applications to wireless monitoring. IEEE Journal of Biomedical and Health Informatics, 19(3), 832–838.

10. Li, T., Sanjabi, M., & Smith, V. (2020). Fair resource allocation in federated learning. In International Conference on Learning Representations.

11. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems (pp. 4765–4774).

12. Hu, S. (2026). Research on Security Enhancement Methods for Adversarial Robust Large Language Model Intelligent Agents for Medical Decision-Making Tasks. arXiv preprint arXiv:2605.08257.

13. Yoon, J., Yang, E., Lee, J., & Hwang, S. J. (2021). Federated continual learning with weighted inter-client transfer. In International Conference on Machine Learning (pp. 12073–12086). PMLR.

14. Liu, L., Zhang, J., Song, S., & Letaief, K. B. (2020). Client-edge-cloud hierarchical federated learning. In IEEE International Conference on Communications (pp. 1–6).

15. Finlayson, S. G., Bowers, J. D., Ito, J., Zittrain, J. L., Beam, A. L., & Kohane, I. S. (2019). Adversarial attacks on medical machine learning. Science, 363(6433), 1287–1289.

16. Kairouz, P., Oh, S., & Viswanath, P. (2015). The composition theorem for differential privacy. In International Conference on Machine Learning (pp. 1376–1385). PMLR.

17. Fine, J., & Hester, R. (2021). The effect of skin pigmentation on the accuracy of pulse oximeters. British Journal of Anaesthesia, 127(2), e49–e51.

18. Veale, M., & Binns, R. (2017). Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data. Big Data & Society, 4(2), 2053951717743530.

19. Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... & Liang, P. (2021). On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258.

20. Yang, Q., Steinfeld, A., & Zimmerman, J. (2019). Unremarkable AI: Fitting intelligent decision support into critical, clinical decision-making processes. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (pp. 1–11).

21. Blanchard, P., El Mhamdi, E. M., Guerraoui, R., & Stainer, J. (2017). Machine learning with adversaries: Byzantine tolerant gradient descent. In Advances in Neural Information Processing Systems (pp. 119–129).

22. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (pp. 3645–3650).

23. Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. Neural Networks, 113, 54–71.

24. Allen, J., & Murray, A. (2003). Age-related changes in the characteristics of the photoplethysmographic pulse shape at various body sites. Physiological Measurement, 24(2), 297–307.

25. Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. In 2017 IEEE Symposium on Security and Privacy (pp. 3–18).