

Federated Learning for Privacy-Preserving Personalized Cardiovascular Monitoring with Wearable Sensors

Edwin R. Jorgensen

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

Abstract

Wearable sensors now support continuous cardiovascular monitoring at population scale, providing longitudinal photoplethysmographic and electrocardiographic data that can inform early detection of arrhythmia, hypertension-related instability, and other cardiovascular conditions. However, centralized machine learning approaches require transferring sensitive physiological data to cloud environments, creating privacy, security, governance, and scalability problems. Federated learning offers a system-level alternative in which model updates are computed on distributed devices and only aggregated parameters are shared. This paper examines the architectural, infrastructural, and socio-technical dimensions of federated learning for privacy-preserving personalized cardiovascular monitoring with wearable sensors. It analyzes data quality challenges arising from motion artifacts and sensor heterogeneity, the design of secure aggregation and personalization mechanisms, and the trade-offs among privacy, utility, fairness, and energy consumption. The discussion further addresses robustness, clinical safety, regulatory compliance, and lifecycle management. The paper argues that federated cardiovascular monitoring should be understood not merely as a machine learning method but as a distributed health information infrastructure requiring coordinated governance, adaptive signal processing, equitable cohort representation, and sustainable edge operation. It concludes by identifying future directions for cross-institutional validation, energy-aware personalization, and regulatory alignment.

Keywords

federated learning; wearable sensors; cardiovascular monitoring; privacy preservation; personalization; edge computing; health data governance.

1. Introduction

Wearable cardiovascular monitoring has shifted from episodic clinical measurements toward continuous, longitudinal observation of photoplethysmographic, electrocardiographic, and accelerometric signals. This shift creates opportunities for early detection of arrhythmias, heart failure decompensation, and stress-related cardiovascular instability. At the same time, it generates sensitive physiological data at very high frequency across heterogeneous consumer devices. Traditional cloud-centric machine learning pipelines aggregate these signals in centralized data lakes, where model training can uncover population-level patterns but also concentrates privacy risk, increases communication burdens, and assumes stable connectivity. Federated learning has emerged as an architectural alternative that distributes model training across devices while keeping raw physiological data localized [1].

Federated learning is not a single optimization technique but a family of distributed protocols with distinct coordination, communication, and trust assumptions. The core idea is that

participating clients compute model updates on local data and transmit only these updates, or compressed versions thereof, to an aggregation server. This approach can reduce the movement of raw health data and create new possibilities for collaborative model development across institutions and device fleets. However, the system-level challenges are substantial: statistical heterogeneity across clients, intermittent availability, uneven data quality, adversarial actors, and regulatory obligations differ from conventional centralized learning environments [2]. In medical and wellness settings, federated learning has been proposed as a route toward multicenter collaboration that respects data governance boundaries and enables access to diverse patient populations [3].

Privacy preservation in such systems must be treated as a composite property rather than a binary condition. Federated learning alone does not guarantee confidentiality because model updates can encode information about the training data. Formal privacy frameworks, such as differential privacy, provide rigorous bounds on the influence of individual records [4]. These mechanisms interact with accuracy, personalization, and fairness in ways that require careful architectural reasoning. This paper addresses those interactions from a systems perspective, emphasizing structural trade-offs, deployment constraints, governance, robustness, and policy implications rather than a single algorithmic contribution.

2. Wearable Cardiovascular Monitoring and Data Quality Challenges

Continuous cardiovascular monitoring using wearable devices depends fundamentally on the quality of the underlying biosignals. Photoplethysmography is widely used because it can be acquired with low-cost optical sensors, yet it is highly susceptible to motion artifact, ambient light, and sensor displacement. In ambulatory environments, movement can obscure the pulsatile component and produce misleading heart rate estimates. Signal conditioning therefore becomes a prerequisite for reliable model training. Adaptive notch filtration methods have been developed to recover photoplethysmographic signals during physical activity by suppressing artifact frequencies while preserving the physiological waveform [5]. Such preprocessing is especially important in federated settings because low-quality updates from corrupted signals can degrade the global model for all participants.

Photoplethysmography has a long history in physiological measurement and remains an attractive modality for noninvasive cardiovascular assessment, but its interpretation requires attention to vascular tone, tissue perfusion, and sensor placement [6]. Wearable sensor platforms vary widely in sampling rate, optical wavelength, contact pressure, battery capacity, and onboard processing capability [7]. These differences create heterogeneous data distributions even before demographic and clinical variation is considered. A federated infrastructure must therefore incorporate device metadata, sensor calibration, and quality assurance flags as first-class signals rather than assuming that all local datasets are statistically exchangeable.

From an infrastructure perspective, the reliability of cardiovascular monitoring depends on the entire acquisition chain, including firmware filtering, motion compensation, segmentation of cardiac cycles, and periodic recalibration. Failure at any layer can propagate into model updates and clinical interpretations. Federated architectures should support local data quality assessment, selective participation, and update rejection based on signal confidence. This design stance recognizes that data quality is not a fixed property of a dataset but an ongoing operational condition that must be governed across devices, manufacturers, and clinical contexts.

3. Architectural Foundations for Federated Cardiovascular Learning

The basic federated architecture for cardiovascular monitoring can be described as a hierarchical system in which wearable devices or nearby edge nodes maintain local model components and exchange updates with a central aggregation service. Secure aggregation protocols allow the server to combine encrypted updates without inspecting individual contributions, thereby reducing the risk of reconstructing a participant signal from gradient information [8]. This protection is relevant because cardiovascular waveforms are quasi-periodic and may contain identifying features. Secure aggregation, however, introduces communication overhead, cryptographic assumptions, and limitations on server-side quality control. These costs must be balanced against the sensitivity of the model and the threat model of the deployment environment.

Federated learning has been explored in medical imaging and distributed hospital data as a mechanism for multi-institutional collaboration without direct sharing of patient records [9]. Cardiovascular monitoring differs in several important respects. First, the data are longitudinal and streaming rather than static, requiring temporal models that can capture evolving physiological states. Second, wearable deployments are cross-device rather than cross-silo, meaning that the number of clients may be very large, device withdrawal is common, and local availability is uncertain. Third, labels for cardiac events are often sparse, noisy, or derived from user reports. These conditions demand asynchronous aggregation, staleness tolerance, and careful weighting of updates by data quality and recency.

Architectural design choices also determine how accountability is distributed. A central server can enforce versioning, audit logging, and access control, but it may also become a single point of failure or coercion. Alternatively, decentralized topologies can distribute aggregation across multiple intermediaries, reducing centralized control but increasing coordination complexity. For clinical-grade cardiovascular monitoring, hybrid architectures may be most appropriate: local fleets managed by care providers or device manufacturers, regional aggregation services, and global model coordination only when regulatory agreements permit.

4. Personalization Across Populations and Contexts

Cardiovascular physiology varies markedly across age, sex, medication use, fitness, sleep, and comorbid conditions. A single global model trained across many devices may perform well on average but poorly for older adults, patients with arrhythmias, or individuals with darker skin tones when optical sensors are poorly calibrated. Personalized federated learning has emerged as a response to this heterogeneity, aiming to retain the benefits of shared representations while adapting local models to individual or subgroup physiology [10]. This approach is particularly important for monitoring applications where baselines differ strongly across patients and where abnormal events are defined relative to personal history.

Personalization techniques vary in their distribution of adaptation across the client-server boundary. Some architectures maintain a shared core model for feature extraction and allow each participant to update only a local output layer, thereby limiting communication overhead and device computation. Others train local models under a regularization penalty that discourages divergence from a global aggregate, preserving population-level information while enabling individual adaptation. These choices have direct consequences for cardiovascular monitoring because patient-specific baselines for heart rate variability, ectopic burden, and circadian variation are more informative than absolute population thresholds. A system that adapts the decision boundary to individual physiology may detect deviations

earlier, but it also risks overfitting to short-term noise or transient sensor artifacts. Therefore, personalization should be coupled with temporal smoothing, confidence estimation, and periodic re-evaluation on held-out local data. Contextual factors such as sleep stage, physical activity, caffeine intake, and medication schedules can shift cardiovascular dynamics independently of disease state. Personalization mechanisms that condition predictions on coarse activity categories may improve specificity while avoiding detailed behavioral profiling that would undermine privacy. In wearable health settings, federated transfer learning approaches have been used to align feature representations across different device distributions before local adaptation, illustrating the practical value of separating population-level representation learning from individualized decision calibration [11].

Fairness in federated cardiovascular monitoring is closely connected to data representation and device access. If training cohorts are dominated by healthy, relatively young users with high-end wearables, the resulting models may underperform for older adults, patients with chronic diseases, or populations with lower-quality sensors. Federated systems can inadvertently reinforce these disparities by weighting updates according to data volume, connectivity frequency, or participation consistency. Governance frameworks should therefore require participation audits, subgroup performance reporting, and mechanisms to correct for recruitment imbalances. Personalization alone does not solve fairness; without explicit oversight, it can create accurate models for well-represented groups while leaving others with poorly calibrated predictions. This issue is not unique to federated learning, but its distributed nature makes auditability more complex because individual datasets are not centrally visible. Audit protocols that rely on secure multiparty computation or privacy-preserving statistics can provide some assurance without exposing raw physiological data.

6. Privacy, Security, and Governance Trade-offs

Federated learning reduces the movement of raw physiological data, but it does not by itself eliminate privacy risk. Model updates, particularly gradients from recurrent or attention-based models, can encode information about the training instances from which they were computed. In cardiovascular applications, these updates may reveal heart rate patterns, medication-response signatures, or the presence of rare arrhythmias. Secure aggregation protects individual updates from direct inspection by the coordination server and from other clients, yet it also limits the server's ability to identify malicious contributions, verify data quality, or detect model poisoning. These limitations require system designers to introduce complementary safeguards such as client attestation, update clipping, and audit traces that do not depend on raw data visibility. Heterogeneity and communication constraints remain central challenges for practical federated learning deployments, and their interaction with privacy mechanisms is especially acute in health monitoring [12].

Differential privacy provides a formal framework for limiting the influence of individual records on trained models by adding calibrated noise to model updates or aggregated parameters [13]. The strength of this guarantee is governed by a privacy budget, which must be allocated across training rounds, personalization steps, and model versions. Strong privacy guarantees generally reduce the sensitivity of the model to rare events and may disproportionately degrade performance for underrepresented cardiovascular conditions. This trade-off is particularly serious in health monitoring, where minority groups may have distinctive physiological presentations. Intermediate edge aggregation nodes can reduce the need for direct device-to-cloud communication and serve as points where local policies are enforced before data-derived signals are forwarded [14]. Federated differential privacy has

been explored for medical data as a way to combine decentralized training with formal confidentiality bounds, but the practical deployment of these methods requires careful tuning of noise scales and aggregation protocols [15]. Governance bodies must therefore decide what level of formal privacy is acceptable for different deployment modes, distinguishing between consumer wellness applications, clinical decision support, and post-market surveillance.

Regulatory frameworks such as the European General Data Protection Regulation establish principles of data minimization, purpose limitation, and accountability that align conceptually with federated learning, but they also impose requirements for transparency, individual rights, and international transfer safeguards that distributed systems must operationalize. A federated cardiovascular monitoring platform should maintain provenance records for model updates, document the legal basis for each data processing stage, and provide mechanisms for participants to withdraw consent or delete local models. These governance requirements are not merely legal formalities; they shape the architecture by requiring traceable aggregation, versioned model releases, and contractual agreements among device manufacturers, cloud providers, and clinical institutions. Without such agreements, federated systems may become fragmented, with inconsistent privacy protections and uneven clinical oversight across regions.

7. Robustness, Safety, and Clinical Integration

The robustness of federated cardiovascular monitoring depends on the ability of the system to function under heterogeneous sensors, missing clients, adversarial perturbations, and distribution shift. Wearable devices may change firmware, calibration, or sensor quality over time, producing covariate shifts that are not synchronized across the population. A federated model trained on a stable cohort may degrade when deployed to new devices or clinical subpopulations. Monitoring systems should therefore include operational safeguards such as confidence scoring, trend analysis, and escalation paths that do not rely solely on model output. Clinical integration also requires that outputs be interpretable by care teams and reconcile with established diagnostic pathways. The distributed nature of federated learning means that model behavior can vary across participating sites, and this variation must be understood before clinical deployment.

Model poisoning is a particular concern in large-scale federated deployments because a small number of compromised devices can distort the shared model if their updates are assigned disproportionate weight. Defensive mechanisms include robust aggregation rules that downweight outliers, update similarity checks, and limiting the influence of any single client. However, such mechanisms must be carefully balanced against legitimate heterogeneity: a patient with a rare arrhythmia may produce an update that appears as an outlier but is clinically important. Overly aggressive outlier rejection can suppress exactly the signals that medical monitoring seeks to detect. Federated systems should therefore maintain multiple model views, including population-level, subgroup-level, and personalized components, so that rare patterns are not erased by consensus. Fairness in machine learning requires attention to both statistical bias and systemic biases in data collection, and federated health systems inherit these challenges across device, institutional, and demographic dimensions [16].

Safety assessment in federated learning is complicated by the absence of a single central test set. Clinical validation requires prespecified protocols, external evaluation cohorts, and prospective monitoring of model changes. Continuous learning systems may drift after deployment, making it difficult to define a stable version that regulators can review. One approach is to separate model development, frozen release, and local personalization layers so that the core clinical model remains auditable while adaptation occurs only at the edge. This

separation helps align federated learning with medical device regulation, which typically assumes a locked algorithm with documented performance characteristics. Machine learning in health applications is known to be vulnerable to distribution shift, label noise, and hidden confounders, and these risks are amplified when data are generated continuously outside controlled clinical settings [17]. A clinically usable system must also provide fallback behavior when model confidence is low or when the device is unable to communicate with aggregation infrastructure. Safety engineering therefore extends beyond model accuracy to include clear indications for users, escalation to telehealth services, and integration with electronic health records where clinically appropriate.

8. Energy, Scale, and Sustainable Deployment

Wearable cardiovascular sensors operate under tight energy and thermal constraints. On-device training can increase battery drain, reduce sensor sampling time, and generate heat, which may be unacceptable for continuous wear. The computational cost of local gradient computation depends on model size, update frequency, and the incorporation of adaptive signal processing steps such as motion artifact removal. A sustainable federated deployment must therefore treat energy as a first-class design constraint, selecting lightweight model architectures, scheduling training during charging periods, and adapting model complexity to battery state and sensor context. Energy-aware federated learning has become an active area of research because wearable devices cannot sustain frequent gradient computation without degrading their primary sensing function. Compression of model updates through quantization, sparsification, and low-rank approximation reduces communication energy but introduces noise that may interact with differential privacy mechanisms. The interaction is nontrivial: compression that preserves most signal for common patterns can destroy the rare patterns that matter for cardiovascular anomaly detection. A sustainable system must therefore evaluate the combined effect of compression, privacy noise, and local model adaptation on clinical utility rather than optimizing each component in isolation.

Scalability further complicates energy management because the number of participating wearables may range from thousands to millions, while the communication infrastructure must tolerate intermittent connectivity, variable latency, and heterogeneous network conditions. Asynchronous aggregation allows the server to incorporate updates as they arrive rather than waiting for stragglers, but it also introduces temporal bias because devices that synchronize frequently may come from different demographic or behavioral groups than devices that connect rarely. This bias can distort the global model and undermine fairness. Hierarchical edge aggregation reduces the load on the central server by combining updates from nearby devices before forwarding aggregated summaries [18]. Edge nodes can enforce local data quality checks, apply regional privacy policies, and adapt to network partitions, but they also introduce new trust boundaries and require additional security measures. The overall architecture must therefore be understood as a socio-technical infrastructure in which energy, connectivity, trust, and clinical accountability are jointly optimized across wearable, edge, and cloud tiers.

9. Policy Implications and Cross-Domain Comparisons

Federated cardiovascular monitoring inherits policy tensions from broader digital health, data protection, and medical device regulation. Consent mechanisms designed for one-time data collection are insufficient for continuous, decentralized model training that may persist across software updates and device ownership changes. Policies must address the longitudinal nature

of participation, including the right to delete local contributions, the portability of personalized models, and the conditions under which model updates may be retained by aggregators. Such policies are better developed through multi-stakeholder governance involving device manufacturers, clinical societies, patient advocates, and regulatory agencies. Cross-domain comparisons with the financial sector, where federated learning has been used for fraud detection across institutions, illustrate both the potential and the limitations of consensus-based learning under regulatory constraints [19]. In banking, data are tabular, labels are relatively clear, and institutions share a strong common interest in detecting fraud. In cardiovascular monitoring, signals are continuous, labels are often probabilistic, and participants include individuals, clinics, and consumer device companies with asymmetric incentives. These differences mean that governance models from other sectors cannot be transferred directly but can inform mechanisms for audit, accountability, and dispute resolution.

Telecommunications provides another instructive comparison because federated learning has been applied to on-device keyboard prediction and network optimization at very large scale. In those domains, models can tolerate moderate personalization loss because the primary objective is average utility across many users. In health monitoring, average utility is insufficient; clinicians require predictable behavior for vulnerable subgroups, and errors may carry direct physiological consequences. This asymmetry argues for a differentiated regulatory posture: consumer wellness features may operate under lighter oversight, while algorithms intended to inform clinical decisions should be subject to prospective validation, post-market surveillance, and transparent performance reporting across demographic strata. Policy instruments such as algorithm impact assessments, which have been proposed for high-risk artificial intelligence systems [20], could be adapted to federated health deployments by requiring periodic evaluation of subgroup accuracy, privacy budget consumption, and energy sustainability.

Global health equity is another policy dimension that intersects with federated infrastructure. Low-resource settings may have limited access to high-quality wearables, reliable connectivity, or cloud computing, yet they often bear a disproportionate burden of cardiovascular disease. Federated architectures that rely on continuous connectivity and powerful edge devices may exacerbate rather than reduce these disparities. Alternatively, lightweight models that operate on low-cost sensors with intermittent synchronization could extend monitoring to underserved populations, provided that aggregation protocols are designed to accommodate irregular participation and that personalization does not depend on extensive labeled data. Cross-domain lessons from agricultural federated sensing and environmental monitoring, where devices are often solar-powered and connectivity is sparse, suggest that asynchronous, delay-tolerant update protocols can be robust under extreme infrastructural constraints. These lessons should inform the design of cardiovascular monitoring systems intended for global deployment.

The future evolution of federated cardiovascular monitoring will likely involve tighter integration with clinical pathways, including automated triage, remote patient management, and closed-loop interventions. Such integration raises new questions about liability when decisions are distributed across devices, edge nodes, and cloud models. Traditional medical device regulation assumes a clear manufacturer and a fixed algorithm, whereas federated systems are continuously evolving through contributions from many participants. Policy frameworks must therefore distinguish between the core model, which may be locked and

validated, and personalized adaptation layers, which may be updated more frequently but must remain within documented safety bounds. International harmonization of these frameworks will be slow, but interim solutions include contractual data protection agreements, technical audit logs, and independent model certification bodies that can evaluate federated systems without accessing raw participant data.

10. Conclusion

Federated learning presents a promising architectural direction for privacy-preserving personalized cardiovascular monitoring with wearable sensors, but its success depends on resolving deep structural tensions among privacy, personalization, robustness, energy, and fairness. The distributed nature of federated training aligns with the need to avoid centralizing sensitive physiological data, yet model updates themselves can leak information and require secure aggregation, differential privacy, and careful governance. Cardiovascular monitoring introduces distinctive challenges because biosignals are noisy, longitudinal, and highly heterogeneous across individuals and devices. Adaptive signal processing methods are essential to recover usable photoplethysmographic signals from motion-corrupted recordings, but they must be integrated into a federated infrastructure that treats data quality as an operational property rather than a preprocessing afterthought. Personalization is necessary for clinically meaningful monitoring, but it must be balanced against fairness and the risk of overfitting to short-term artifacts. Robustness against poisoning and distribution shift requires aggregation rules that respect the clinical importance of rare patterns. Energy and connectivity constraints demand lightweight models, hierarchical aggregation, and asynchronous protocols that can accommodate global disparities in device capability and network access. Policy frameworks must evolve beyond static consent and single-manufacturer device regulation toward ongoing, multi-stakeholder governance that accounts for continuous learning, subgroup performance, and the distributed allocation of responsibility. By attending to these system-level dimensions, researchers and practitioners can move federated cardiovascular monitoring from a technical prototype toward a trustworthy socio-technical infrastructure that improves early detection and personalized care without sacrificing privacy or equity.

References

1. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 54, 1273–1282.
2. 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.
3. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... & Cardoso, M. J. (2020). The future of digital health with federated learning. *npj Digital Medicine*, 3, 119.
4. Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference* (pp. 265–284). Springer.
5. Zheng, X., Dwyer, V. M., Barrett, L. A., Derakhshani, M., & Hu, S. (2022). Adaptive notch-filtration to effectively recover photoplethysmographic signals during physical activity. *Biomedical Signal Processing and Control*, 72, 103303.
6. Allen, J. (2007). Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3), R1–R39.

7. Majumder, S., Mondal, T., & Deen, M. J. (2017). Wearable sensors for remote health monitoring. *Sensors*, 17(1), 130.
8. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., ... & 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., ... & Bakas, S. (2020). Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 10, 12598.
10. Smith, V., Chiang, C.-K., Sanjabi, M., & Talwalkar, A. (2017). Federated multi-task learning. *Advances in Neural Information Processing Systems*, 30, 4424–4434.
11. 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.
12. 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.
13. 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.
14. Lim, W. Y. B., Luong, N. C., Hoang, D. T., Jiao, Y., Liang, Y.-C., Yang, Q., ... & Miao, C. (2020). Federated learning in mobile edge networks: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 22(3), 2031–2063.
15. Choudhury, O., Gkoulalas-Divanis, A., Salonidis, T., Sylla, I., Park, Y., Hsu, G., & Das, A. (2020). Differential privacy-enabled federated learning for sensitive health data. *arXiv preprint arXiv:1910.02578*.
16. 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), 1–35.
17. Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine*, 17, 195.
18. Wang, S., Tuor, T., Salonidis, T., Leung, K. K., Makaya, C., He, T., & Chan, K. (2019). Adaptive federated learning in resource constrained edge computing systems. *IEEE Journal on Selected Areas in Communications*, 37(6), 1205–1221.
19. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1–19.
20. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 1–21.