Robust Pulse Waveform Analysis for Wearable Cardiovascular Health Assessment

Nilesh Chowdhury
Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
nilesh-chowdhury96@colostate.edu

Jeffrey Herrera
Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO, USA.
jeffreywork@missouri.edu

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

Abstract

Wearable cardiovascular monitoring has become a central component of longitudinal health assessment, with pulse waveform analysis offering a non-invasive window into cardiac timing, vascular dynamics, and autonomic regulation. However, the translation of wearable photoplethysmography and related pulse signals into trustworthy health information is constrained by motion artifacts, sensor variability, physiological heterogeneity, and inconsistencies in data quality. This paper presents a system-level examination of robust pulse waveform analysis for wearable cardiovascular health assessment. It argues that robustness cannot be achieved through algorithmic refinements alone, but must be understood as an emergent property of sensor design, signal quality governance, layered analytical architecture, and deployment infrastructure. The discussion addresses the physical and contextual sources of waveform degradation, the structural trade-offs among on-body processing, edge computing, and centralized cloud inference, and the need for quality-aware, activity-adaptive processing pipelines. Attention is also given to fairness, clinical governance, and health equity, including the ways in which optical sensing and machine learning can reproduce or amplify disparities. Sustainability and interoperability are examined as policy-relevant constraints on large-scale deployment. The paper concludes by outlining directions for adaptive multimodal learning, energy-efficient inference, real-world validation, and governance frameworks that support safe and equitable cardiovascular monitoring outside controlled clinical settings. Throughout, the emphasis is placed on structural design choices rather than isolated signal processing techniques, providing an interdisciplinary perspective for researchers, clinicians, and health system architects.

Keywords

pulse waveform analysis; wearable sensors; photoplethysmography; robustness; edge computing; fairness; clinical governance; health equity.

1. Introduction

Wearable cardiovascular health assessment has expanded rapidly because of advances in low-power optical sensing, embedded computation, and wireless connectivity.

Photoplethysmography has become a dominant modality in this space because it enables the continuous acquisition of pulse waveforms from peripheral vascular beds without requiring invasive instrumentation [1]. The resulting waveforms contain information about cardiac cycle timing, arterial compliance, peripheral resistance, and respiratory modulation. Yet the same signals are highly susceptible to motion, sensor displacement, ambient light, and interindividual differences in tissue properties. The central challenge is therefore not merely the extraction of heart rate from clean recordings, but the construction of an integrated system that can preserve analytical validity across heterogeneous conditions and populations [2].

The broader field has moved from fingertip measurements toward wrist-worn, earlobe, upper-arm, and even imaging-based configurations, each of which introduces distinct optical and anatomical trade-offs [3]. Consumer devices now collect vast quantities of longitudinal pulse data, but the raw information is often noisy and incomplete. Robustness in such systems should be viewed as a structural property that arises from the interaction of sensor physics, signal quality assessment, contextual inference, computational placement, and institutional oversight. This paper provides a system-level analysis of robust pulse waveform analysis for wearable cardiovascular health assessment. It examines measurement contexts, failure modes, analytical architecture, edge-cloud trade-offs, fairness, governance, sustainability, and future directions. In doing so, it prioritizes architectural reasoning over the presentation of specific equations or algorithmic pseudocode.

2. Pulse Waveform Acquisition and Measurement Contexts

Photoplethysmography operates by illuminating a region of skin with a light source and measuring the reflected or transmitted light with a photodetector [1]. The raw signal represents volumetric changes in the microvascular bed during each cardiac cycle. Its morphology is shaped not only by cardiac ejection but also by vascular tone, respiration, local temperature, sensor pressure, and coupling between the device and tissue. Fingertip sensors typically provide high signal amplitude because of the dense capillary network, but their use is incompatible with continuous daily wear. Wrist-worn sensors offer far greater usability and have become common in consumer devices, yet the wrist is a challenging anatomical site because of low perfusion and frequent movement [2]. Earlobe and upper-arm sensors present intermediate trade-offs in terms of comfort, stability, and waveform quality.

A system-level perspective must also consider wavelength selection, optical geometry, sampling rate, and energy consumption. Green light is frequently used in wrist-based devices because it is strongly absorbed by hemoglobin and is less affected by venous pulsations, although it can also be sensitive to melanin absorption [3]. Red and infrared wavelengths penetrate more deeply and are used in pulse oximetry and in some approaches to blood pressure estimation. The optical choices interact with skin pigmentation, tissue thickness, and ambient light conditions [4]. These physical factors constrain which cardiovascular features can be reliably extracted. Heart rate is relatively robust to modest waveform distortion, whereas pulse arrival time, reflected wave features, and other vascular stiffness indices require much higher signal quality. Robust architecture must therefore acknowledge that different physiological parameters have different quality thresholds and that sensor placement fundamentally shapes downstream analytical possibilities.

3. Signal Quality Degradation and Failure Modes

Motion artifact is the most pervasive threat to wearable pulse waveform analysis. Physical activity produces relative movement between the sensor and the skin, alters tissue

compression, and changes local blood volume in ways that are not synchronized with cardiac activity [5]. Periodic activities such as running can introduce strong narrowband interference that overlaps with plausible heart rate frequencies. Irregular gestures can produce transient spikes that corrupt waveform morphology and complicate feature extraction. Because motion artifacts often share spectral and temporal characteristics with the pulse signal, simple bandpass filtering is usually insufficient. Robust systems must treat physical activity as a context that changes the very signal model rather than as a nuisance to be removed after the fact.

Beyond motion, several other failure modes degrade measurement quality. Ambient light can saturate the photodetector or introduce harmonics, particularly when the sensor is not firmly coupled to the skin. Excessive sensor pressure reduces pulsatile perfusion, while insufficient pressure increases sensitivity to displacement. Low peripheral perfusion, cold extremities, and vasoconstriction reduce waveform amplitude and may lead to signal dropout. Skin pigmentation affects light absorption and can reduce effective signal strength when device optics are not designed and tested across diverse skin tones [4]. Autonomic changes caused by stress, medication, or illness further alter pulse morphology in ways that can be mistaken for cardiovascular deterioration. These varied failure modes indicate that wearable pulse analysis requires explicit uncertainty management. A segment that is sufficient for heart rate estimation may be entirely unsuitable for advanced cardiovascular assessment, and systems must be designed to recognize this distinction.

4. Robust Analytical Architecture for Wearable Cardiovascular Assessment

A resilient analytical pipeline must integrate multiple layers that operate from acquisition to decision support. The first layer involves adaptive signal conditioning, including baseline correction, interference suppression, and dynamic gain control. The second layer is quality-aware segmentation, in which short windows of pulse signal are evaluated before being passed to parameter extraction. The third layer is motion-tolerant estimation, combining time-domain, frequency-domain, and model-based strategies to recover cardiac timing from compromised signals. The fourth layer is contextual inference, in which accelerometer and gyroscope information is used to label activity states and adjust processing parameters accordingly [6]. This layered structure reduces the probability that a localized failure in one module propagates into a clinically misleading output.

Motion-tolerant heart rate estimation has been substantially improved by approaches that use accelerometer references to separate pulse-related components from motion-related components. Early frameworks demonstrated that wrist-type photoplethysmographic signals during intensive exercise could be processed through accelerometer-assisted spectral tracking and subtraction [6]. Later work refined motion reference selection, adaptive filtering, and tracking mechanisms to improve robustness in free-living conditions [7]. Deep learning has also been applied to learn direct mappings from raw or minimally processed photoplethysmography and accelerometer signals to heart rate estimates [8]. Such models can outperform conventional methods when training distributions align with deployment conditions, but they introduce new concerns about domain shift, dataset bias, interpretability, and computational cost.

Signal quality assessment is essential for clinically meaningful wearable monitoring. Quality indices can identify segments with sufficient morphology for reliable feature extraction [9]. Lightweight quality classifiers can be embedded on the device itself, preventing corrupted data from being transmitted or used for downstream decisions. More detailed quality scoring

may be performed at the edge or in the cloud [10]. The combination of quality gating and activity context allows systems to operate in a tiered manner. During low-motion rest, advanced morphological features may be derived safely. During high-motion activity, the system may restrict itself to heart rate estimation with wider uncertainty bounds. This tiered inference model is important because it matches analytical ambition to demonstrated data quality rather than forcing every segment through the same diagnostic pathway.

Pulse waveform analysis has also been extended beyond heart rate to respiratory rate estimation, where amplitude modulation, baseline wander, and pulse interval variability provide indirect respiratory information [11]. Deep learning architectures originally developed for electrocardiographic rhythm classification have influenced cardiovascular analytics by demonstrating that high-dimensional temporal signals can be mapped to clinical labels [12]. Similar approaches have been adapted for ambulatory settings, although caution is warranted when transferring models trained in hospital environments to wearable photoplethysmography [13]. Cuffless blood pressure estimation has also attracted significant attention, with multistage deep models attempting to learn vascular state from pulse morphology [14]. Rapid extraction of vital signs under changing physical activity intensity remains a difficult problem, and recent work has emphasized real-time opto-physiological monitoring that adapts to varying activity levels [15]. These advances illustrate a broader shift from static signal processing toward adaptive, context-aware analytical systems.

5. Structural Trade-offs in On-Body, Edge, and Cloud Processing

The placement of analytical functions across wearable hardware, edge gateways, and cloud infrastructure is a decisive structural choice. On-body processing minimizes reliance on connectivity and can support immediate feedback, but it is limited by battery capacity, thermal constraints, and available memory. Edge processing on a smartphone or home hub offers additional computational resources while maintaining geographic proximity to the user. Cloud processing enables cohort-scale analysis, model retraining, and longitudinal pattern discovery, but it introduces latency, connectivity dependence, and privacy risk. A robust deployment strategy distributes tasks according to the minimum viable quality required at each layer. Basic heart rate and activity context can be computed on-device, while more complex arrhythmia screening may require edge or cloud resources.

The placement of computation also shapes model lifecycle management. Wearable sensor populations drift over time due to hardware revisions, firmware updates, seasonal physiological variation, and demographic shifts. Models trained centrally may degrade when deployed across different devices and user groups. Federated learning has been proposed as a governance-compatible strategy for updating models across distributed devices without centralizing raw waveform data [16]. This approach preserves privacy but introduces coordination challenges related to statistical heterogeneity, communication cost, and fairness across participating sites. Robustness at the system level therefore depends on how model updates are governed, how data provenance is recorded, and how performance degradation is detected after deployment. These organizational considerations are as important as the choice of signal processing or learning algorithm.

6. Fairness, Bias, and Clinical Governance

Wearable pulse waveform analysis intersects directly with health equity. Bias can emerge from sensor physics, training data composition, annotation procedures, and deployment contexts. Skin pigmentation can influence optical signal quality when devices are not tested

across diverse populations, potentially leading to systematic underestimation of cardiovascular features for some users [4]. Machine learning models trained predominantly on data from young, healthy, or light-skinned individuals may learn features that generalize poorly to older adults, individuals with darker skin, or patients with peripheral vascular disease. Fairness must therefore be designed into data collection, quality annotation, model training, and post-market evaluation ratherthan treated as an afterthought. Large-scale wearable datasets must include representative samples across skin tones, ages, body compositions, and cardiovascular disease states, and quality annotations must be generated under varied motion and environmental conditions. The consequences of failing to do so are evident in other health algorithm contexts, where biased risk prediction has systematically underestimated the health needs of marginalized groups despite similar underlying morbidity [17]. Optical pulse analysis carries the same hazard because melanin concentration and peripheral perfusion are not merely demographic variables but direct physical modifiers of signal generation. If model development proceeds without explicit stratification and subgroup performance testing, devices may meet aggregate accuracy targets while failing systematically for darker-skinned users, individuals with low cardiac output, or older adults with stiffened vasculature. Intersectional evaluation has shown that disaggregated performance reporting is necessary to expose failures that are hidden by overall averages [18]. A robust wearable cardiovascular system must therefore embed fairness as a measurement discipline, with prospective subgroup analyses, transparent quality thresholds, and recalibration procedures linked to physiological and demographic metadata.

Clinical governance for wearable pulse waveform analysis extends beyond model accuracy. Machine learning has transformed medical decision support, but its clinical integration requires carefully managed validation, monitoring, and updating, especially when data distributions shift across time and care settings [19]. Regulatory frameworks increasingly treat adaptive algorithms and wearable-derived biomarkers as dynamic medical devices, but the boundary between wellness products and diagnostic tools remains contested. Manufacturers can avoid formal medical device scrutiny by making only wellness claims, yet downstream use by clinicians may still occur. Robust governance therefore requires explicit statements about intended use, evidence quality, and the limitations of photoplethysmography-derived markers. Health systems that incorporate wearable pulse data into clinical workflows need institutional protocols for data provenance, model versioning, and periodic performance audits. Without these mechanisms, pulse waveform analytics can acquire a false legitimacy that is difficult to reverse once embedded in care pathways.

7. Sustainability, Interoperability, and Infrastructure Policy

Sustainability in wearable cardiovascular health assessment involves more than battery life. It also includes the energy cost of continuous sensing, the environmental footprint of device manufacturing and replacement, and the computational burden of cloud-based model inference. Federated learning offers one structural response to the tension between personalization and centralized data aggregation, as model updates can be computed on user devices or edge servers without transferring raw waveform data to central repositories [20]. This decentralized approach can reduce network traffic, improve privacy, and distribute training costs across existing hardware. However, federated learning introduces its own sustainability concerns, including communication overhead during model synchronization and the risk of statistical drift when participating devices differ in sensor quality or user

demographics. A robust system must balance the benefits of decentralized adaptation with the operational complexity of coordinating updates across heterogeneous device fleets.

Interoperability is equally important for sustainable deployment. Pulse waveform analytics cannot become a trusted component of cardiovascular care if data formats, quality annotations, and derived parameters cannot be exchanged across platforms. The digital medicine community has articulated the need for measurement metadata, sensor specifications, and contextual labels to accompany physiological time series, so that downstream consumers can assess whether a given estimate is fit for a particular purpose [21]. Electronic health record integration requires mapping wearable data to clinical terminologies while preserving uncertainty information. Semantic interoperability must also account for differences in sampling rates, waveform preprocessing, and activity classification. Without interoperable quality signals, clinicians may unknowingly compare heart rate variability derived from a green-light wrist sensor with values obtained from a red-light fingertip device, leading to spurious clinical conclusions. Robust infrastructure policy should therefore promote open data standards, metadata transparency, and vendor-neutral validation protocols.

8. Real-World Validation and Deployment Governance

The transition from laboratory prototypes to population-scale deployment introduces validation challenges that cannot be resolved through controlled benchtop testing alone. Wearable devices operate across a vast range of environmental temperatures, ambient light levels, skin hydration states, and activity patterns. Real-world evidence from observational cohorts has demonstrated both the promise and the fragility of wearable cardiovascular monitoring when applied outside supervised studies [22]. Longitudinal deployments require active monitoring of data completeness, signal quality drift, and device degradation. A pulse waveform platform that performs well in a short-term validation study may fail when users wear the device inconsistently or when firmware updates change sensor gain. Robust deployment therefore demands continuous post-market surveillance, with automated detection of performance anomalies and mechanisms for rolling back problematic model updates.

Resource-constrained edge deployment has further implications for governance. Model compression and TinyML techniques enable cardiovascular feature extraction on low-power microcontrollers, reducing latency and dependence on cloud connectivity [23]. However, aggressive quantization and pruning can alter model behavior in subtle ways, especially for underrepresented waveform morphologies. Deployment governance must include equivalence testing between full-precision and compressed models across demographic strata and activity classes. Regulatory review should not treat a compressed model as identical to its larger counterpart unless empirical equivalence has been demonstrated. The combination of continuous monitoring, model updates, and distributed inference requires versioned model registries, clear provenance for training data, and robust fallback behavior when confidence is low.

9. Future Directions for Robust Wearable Cardiovascular Intelligence

Future wearable cardiovascular systems will need to move beyond monomodal photoplethysmography toward adaptive multimodal fusion that preserves interpretability. Pulse waveform morphology has long been recognized as a rich but fragile source of cardiovascular information, and its analysis must be conditioned on sensor placement, tissue properties, and hemodynamic state [24]. Combining photoplethysmography with accelerometry, electrodermal activity, skin temperature, and occasional electrocardiogram

patches can provide cross-signal validation during ambiguous motion events. Such fusion should not simply concatenate inputs, but should model the causal structure of interference and physiological coupling. The system can then learn when pulse morphology is trustworthy and when it must rely on less granular parameters. This represents a shift from static feature extraction to context-dependent inference, where the analytical pathway is chosen according to real-time quality and reliability.

Explainability is another critical frontier. Clinicians are unlikely to trust a system that reports a high-risk cardiovascular event without providing interpretable evidence from the waveform and its context. Explanations that are disconnected from clinical workflow are of limited value, and researchers have emphasized that machine learning explanations must be contextualized for the specific decisions clinicians face [25]. For wearable pulse analysis, explainability might involve highlighting which waveform segments contributed to a confidence estimate, identifying whether motion artifact was adequately suppressed, and presenting quality indicators alongside risk scores. This aligns technical transparency with governance demands for auditability. The next generation of robust wearable cardiovascular intelligence should be designed as a socio-technical system in which machine learning, sensor engineering, clinical knowledge, and regulatory oversight co-evolve.

10. Conclusion

Robust pulse waveform analysis for wearable cardiovascular health assessment cannot be reduced to a single signal processing algorithm or a particular deep learning architecture. It is an emergent property of a layered system that spans optical sensor design, quality-aware signal conditioning, activity-adaptive inference, computational placement, and institutional governance. Motion artifacts, poor perfusion, skin tone, sensor displacement, and population heterogeneity are not peripheral concerns but central structural constraints that shape what can be reliably measured and interpreted. A robust system therefore adopts a tiered analytical strategy, restricting advanced cardiovascular assessment to segments of demonstrated quality while maintaining continuous basic monitoring under degraded conditions.

The fairness and equity dimensions of wearable pulse waveform analysis demand deliberate attention. Sensor physics and machine learning can reproduce systemic disparities unless data collection, performance evaluation, and post-market surveillance are explicitly designed to detect and correct them. Governance structures must keep pace with the adaptive nature of these technologies, ensuring that model updates, compressed inference engines, and federated learning protocols do not silently undermine validated performance. Interoperability and sustainability further shape the long-term viability of wearable cardiovascular health assessment, requiring open standards, metadata transparency, and energy-conscious infrastructure. The path forward lies in treating pulse waveform analysis as an integrated socio-technical challenge in which clinical utility, computational robustness, and health equity are pursued simultaneously.

References

1. Allen, J. (2007). Photoplethysmography and its application in clinical physiological measurement. Physiological Measurement, 28(3), R1–R39.

2. Majumder, S., Mondal, T., & Deen, M. J. (2017). Wearable sensors for remote health monitoring. Sensors, 17(1), 130.

3. Castaneda, D., Esparza, A., Ghamari, M., Soltanpur, C., & Nazeran, H. (2018). A review on wearable photoplethysmography sensors and their potential future applications in health care. International Journal of Biosensors & Bioelectronics, 4(4), 195–202.

4. Bent, B., Goldstein, B. A., Kibbe, W. A., & Dunn, J. P. (2020). Investigating sources of inaccuracy in wearable optical heart rate sensors. NPJ Digital Medicine, 3, 18.

5. Ram, M. R., Madhav, K. V., Krishna, E. H., Komalla, N. R., & Reddy, K. A. (2012). A novel approach for motion artifact reduction in PPG signals based on AS-LMS adaptive filter. IEEE Transactions on Instrumentation and Measurement, 61(5), 1445–1457.

6. Zhang, Z., Pi, Z., & Liu, B. (2015). TROIKA: A general framework for heart rate monitoring using wrist-type photoplethysmographic signals during intensive physical exercise. IEEE Transactions on Biomedical Engineering, 62(2), 522–531.

7. Zhang, Z. (2015). Photoplethysmography-based heart rate monitoring in physical activities via joint sparse spectrum reconstruction. IEEE Transactions on Biomedical Engineering, 62(8), 1902–1910.

8. Reiss, A., Indlekofer, I., Schmidt, P., & Van Laerhoven, K. (2019). Deep PPG: Large-scale heart rate estimation with convolutional neural networks. Sensors, 19(14), 3079.

9. Orphanidou, C., Bonnici, T., Charlton, P., Clifton, D., Vallance, D., & Tarassenko, L. (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. Naeini, E. K., Azimi, I., Rahmani, A. M., Liljeberg, P., & Dutt, N. (2020). A real-time PPG quality assessment approach for healthcare Internet-of-Things. Procedia Computer Science, 175, 551–558.

11. Charlton, P. H., Birrenkott, D. A., Bonnici, T., Pimentel, M. A. F., Johnson, A. E. W., Alastruey, J., Tarassenko, L., Watkinson, P. J., Beale, R., & Clifton, D. A. (2018). Breathing rate estimation from the electrocardiogram and photoplethysmogram: A review. IEEE Reviews in Biomedical Engineering, 11, 2–20.

12. Hannun, A. Y., Rajpurkar, P., Haghpanahi, M., Tison, G. H., Bourn, C., Turakhia, M. P., & Ng, A. Y. (2019). Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. Nature Medicine, 25(1), 65–69.

13. Tison, G. H., Sanchez, J. M., Ballinger, B., Singh, A., Olgin, J. E., Pletcher, M. J., Vittinghoff, E., Lee, E. S., Fan, S. M., Gladstone, R. A., Mikell, C., Sohoni, N., Hsieh, J., & Marcus, G. M. (2018). Passive detection of atrial fibrillation using a commercially available smartwatch. JAMA Cardiology, 3(5), 409–416.

14. El Hajj, C., & Kyriacou, P. A. (2021). Cuffless and continuous blood pressure estimation from PPG signals using machine learning: A review. Biosensors, 11(12), 470.

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. 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., 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, 119.

17. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447–453.

18. Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. Proceedings of Machine Learning Research, 81, 1–15.

19. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. New England Journal of Medicine, 380(14), 1347–1358.

20. 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.

21. Coravos, A., Goldsack, J. C., Karlin, D. R., Nebeker, C., Perakslis, E., Zimmerman, N., & Erb, M. K. (2019). Digital medicine: A primer on measurement. Digital Biomarkers, 3(2), 31–71.

22. Dunn, J., Runge, R., & Snyder, M. (2018). Wearables and the medical revolution. Personalized Medicine, 15(5), 429–448.

23. Banbury, C., Zhou, C., Fedorov, I., Matas, R., Thakker, U., Gope, D., Janapa Reddi, V. J., Mattina, M., & Whatmough, P. (2021). Micronets: Neural network architectures for deploying TinyML applications on commodity microcontrollers. Proceedings of Machine Learning and Systems, 3, 517–532.

24. Elgendi, M. (2012). On the analysis of fingertip photoplethysmogram signals. Current Cardiology Reviews, 8(1), 14–25.

25. Tonekaboni, S., Joshi, S., McCradden, M. D., & Goldenberg, A. (2019). What clinicians want: Contextualizing explainable machine learning for clinical end use. Proceedings of Machine Learning Research, 106, 359–380.