

Intelligent Wearable Systems for Remote Patient Monitoring Using Multimodal Biomedical Signal Analytics

Wayne Jimenez

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.

contactwayne@unr.edu

Mikko L. Coleman

Department of Computer Science, University of New Hampshire, Durham, NH, USA.

mikkolcoleman227@unh.edu

Abstract

The integration of wearable sensing, multimodal biomedical signal processing, and intelligent decision support is transforming remote patient monitoring from episodic measurement into continuous and context-aware physiological surveillance. This paper presents a system-level analysis of intelligent wearable systems for remote patient monitoring, emphasizing architectural trade-offs, multimodal data fusion, robustness, governance, equity, and sustainability. The discussion begins with the structural organization of body area networks, edge computing, and cloud-based clinical services, focusing on the need for adaptive allocation of processing across layers to balance latency, energy consumption, privacy, and diagnostic quality. It then examines core physiological modalities, including photoplethysmography, electrocardiography, accelerometry, respiratory effort, and electrodermal activity, and their integration through feature-level and decision-level analytics. A recurrent theme is that multimodal systems are more resilient to motion artifact and missing data than single-modal monitors, but only when signal quality estimation, synchronization, and context-aware inference are treated as first-class infrastructure. The paper further addresses governance and fairness challenges, including data minimization, federated learning, algorithmic bias, and regulatory compliance. It argues that sustainable clinical deployment requires not only signal processing accuracy but also attention to device lifecycle, energy consumption, interoperability, workflow integration, and socio-technical equity. Future directions highlight the need for auditable adaptive systems, standardized multimodal data models, and policy frameworks that ensure intelligent monitoring is safe, fair, and robust across heterogeneous populations and care settings.

Keywords

wearable systems, remote patient monitoring, multimodal biomedical signals, signal analytics, health data governance, fairness, sustainability.

1. Introduction

The growing prevalence of chronic cardiovascular, respiratory, and neurological conditions has intensified interest in continuous remote monitoring outside conventional clinical environments. Wearable sensor platforms now support longitudinal acquisition of physiological signals such as cardiac activity, peripheral blood volume changes, movement,

and skin conductance, enabling clinicians to observe trends that are obscured by episodic measurements [1], [2]. Consumer and research-grade devices have also diversified rapidly, with photoplethysmographic sensors, single-lead electrocardiograms, multi-axis accelerometers, and temperature sensors becoming available in small and low-power form factors [3]. In parallel, modern artificial intelligence methods have made it possible to learn clinically relevant patterns from large and heterogeneous physiological datasets obtained from distributed wearable and mobile health platforms [4]. Yet realizing clinical value from these data streams requires more than isolated signal processing; it demands a coherent systems architecture that spans sensing, communication, inference, and clinical integration. This paper therefore examines intelligent wearable monitoring as a socio-technical system rather than a collection of algorithms and sensors. The analysis integrates signal-level considerations with infrastructure governance, deployment sustainability, and equity, reflecting the complexity of current healthcare environments. Early research on wearable health monitoring identified the need for scalable architectures and contextual awareness [1], [2], while subsequent reviews emphasized the diversity of sensor modalities and the importance of designing for heterogeneous patient populations [3]. The central challenge is not the absence of physiological information, but the difficulty of transforming noisy, irregular, and multimodal streams into trustworthy clinical decisions without overburdening patients, clinicians, or network infrastructure.

Machine learning approaches have expanded the analytical capacity of wearable systems by enabling detection of subtle patterns across long observation windows [4]. However, the successful translation of these methods into remote monitoring has been limited by data quality, label scarcity, and domain shift across devices and populations. Signal processing remains a foundational layer because downstream models inherit the distortions introduced during acquisition and preprocessing [5]. Artifacts arising from motion, ambient light, and electrode displacement are particularly problematic in ambulatory settings, and their mitigation has motivated extensive research in adaptive filtering, time-frequency methods, and signal quality estimation [6]. Heart rate monitoring during physical activity, for example, illustrates the interplay between motion artifact and physiological signal extraction in wrist-type photoplethysmography [7]. Similarly, reliable wireless monitoring depends on automated assessment of signal quality before physiological parameters are computed [8]. These challenges call for a system-level perspective in which multimodal sensing, robust preprocessing, and context-aware analytics are treated as interdependent components of a larger clinical infrastructure.

2. Architectural Foundations of Wearable Remote Monitoring

Wearable remote monitoring systems are typically organized as layered architectures spanning body-area sensing, edge or gateway processing, and cloud-based clinical services. The structural choices made at each layer influence not only technical performance but also clinical safety, patient autonomy, and system scalability. A centralized architecture in which raw waveforms are streamed continuously to a remote server can simplify algorithm deployment and data management, yet it imposes severe energy and bandwidth burdens on wearable nodes and may introduce unacceptable latency for real-time alarms. Conversely, a highly decentralized architecture that performs all inference locally on the wearable can protect privacy and reduce communication overhead, but it limits the complexity of models that can be executed within tight power and memory budgets. The appropriate balance depends on the clinical use case, the characteristics of the physiological signals, and the

reliability of the communication infrastructure. For example, fall detection or arrhythmia alarms may require local edge processing to guarantee low latency, whereas retrospective trend analysis or population health modeling can tolerate cloud latency but may require large aggregated datasets. Multi-sensor fusion complicates this layering because data from accelerometry, electrocardiography, and photoplethysmography must be synchronized and contextualized before higher-level inference is attempted [9]. A well-designed system treats the edge not merely as a data aggregator but as an intelligent intermediary that performs signal quality assessment, adaptive filtering, and early warning detection.

Fog and edge computing paradigms offer a middle path by distributing computational tasks according to their latency, privacy, and resource requirements [10]. In a remote monitoring deployment, a home gateway or smartphone can perform data cleaning, compression, and preliminary anomaly detection before transmitting clinically meaningful features to the cloud. This reduces both energy consumption on battery-constrained sensors and the exposure of raw physiological data to long-distance networks. However, edge-based inference introduces its own governance problems, including the need to update models across heterogeneous devices, monitor algorithmic drift, and ensure that locally computed decisions are auditable by clinicians. Federated learning has emerged as a promising approach to train models across distributed patient devices without centralizing raw biomedical signals [11]. By exchanging model updates instead of patient data, federated architectures can reduce privacy risks, but they require robust protection against model inversion and membership inference attacks, as well as mechanisms to ensure statistical representativeness across participating sites. The increasing convergence of artificial intelligence and clinical medicine highlights the potential of such distributed learning systems, while also underscoring the importance of external validation and transparent reporting [12]. These architectural considerations are not merely technical; they shape how patients and providers perceive trust, accountability, and control over health information.

3. Multimodal Signal Acquisition and Preprocessing

Photoplethysmography has become a foundational modality for wearable monitoring because it enables estimation of heart rate, oxygen saturation, and vascular dynamics using low-cost optical sensors, yet its signal quality is highly sensitive to motion artifact and skin pigmentation [13]. The green and red wavelengths used in many wrist-worn devices interact differently with tissue, and the resulting photoplethysmographic waveform can be corrupted by periodic limb movement, ambient light fluctuations, and sensor displacement. Robust extraction of physiological parameters from photoplethysmographic signals therefore requires adaptive preprocessing that separates motion-induced disturbances from genuine hemodynamic variations. A range of signal processing strategies has been proposed to address this challenge, including adaptive notch filtering techniques that recover photoplethysmographic components during physical activity without requiring additional motion reference signals [14]. Such methods are important because they reduce dependence on accelerometer arrays that may themselves be noisy or misaligned. The effectiveness of these preprocessing steps is typically evaluated using heart rate estimation error during treadmill protocols, where motion artifacts are particularly severe.

Electrocardiography provides a complementary electrical view of cardiac activity and is often used as a reference for validating photoplethysmography-derived heart rate. However, single-lead wearables are susceptible to baseline wander and electrode motion artifacts, and the morphology of the recorded waveform may be insufficient for diagnosing complex

arrhythmias. Combining photoplethysmography with electrocardiography can improve confidence in physiological parameter estimation, especially when one modality is temporarily degraded. For instance, breathing rate can be estimated from both electrocardiographic and photoplethysmographic signals through respiratory modulation of amplitude, frequency, and pulse transit time [15]. The multimodal approach allows cross-validation and continuity of monitoring even when a single signal becomes unreliable. Signal synchronization is a prerequisite for such fusion, as time delays between wearable nodes can distort estimates of pulse arrival time or heart rate variability. Preprocessing pipelines must therefore include rigorous timestamp alignment, resampling, and quality flagging before any attempt at data fusion. Without these safeguards, multimodal analytics can amplify rather than reduce uncertainty.

In addition to cardiovascular signals, accelerometry provides motion context, electrodermal activity reflects sympathetic arousal, and temperature sensors capture thermal regulation. These modalities are often sampled at different rates and have distinct noise profiles. The preprocessing of each stream must be tailored to its physical origin. For example, accelerometer signals benefit from low-pass filtering to isolate postural changes while retaining high-frequency components for step detection. Electrodermal activity consists of slow tonic and rapid phasic components that respond to thermoregulation and emotional stimuli, and their separation requires careful baseline estimation. The integration of such diverse signals into a common timeline is a systems challenge that goes beyond individual algorithm design. Contextual labeling, whether the patient is sleeping, exercising, or resting, can be inferred from accelerometry and used to adapt thresholds for other physiological parameters. This context-aware preprocessing reduces false alarms and improves the interpretability of downstream analytics.

4. Data Fusion and Intelligent Analytics

Multimodal data fusion in wearable monitoring can be organized conceptually at different levels, including raw data fusion, feature-level fusion, and decision-level fusion. Raw data fusion is attractive because it preserves the richest information, but it requires precise synchronization and can be computationally expensive. Feature-level fusion extracts salient characteristics from each modality, such as heart rate from electrocardiography, pulse amplitude from photoplethysmography, and activity count from accelerometry, then combines these features for classification or regression. Decision-level fusion allows separate classifiers or estimators to operate on individual modalities and then combines their outputs using voting, Bayesian methods, or learned meta-classifiers. The appropriate fusion level depends on the clinical question, the degree of redundancy among modalities, and the available computational resources. In many remote monitoring scenarios, feature-level fusion offers a favorable trade-off between interpretability and performance, because clinicians can inspect the contributing features and understand which physiological channels influenced a given alert.

Machine learning methods applied to fused multimodal features can detect deterioration trends that are not evident in any single signal. For example, subtle increases in resting heart rate combined with reduced physical activity and fragmented sleep may signal an impending exacerbation of chronic obstructive pulmonary disease or heart failure. The clinical value of these analytics hinges on the quality of the training data, the representativeness of the patient cohort, and the robustness of the model to missing modalities. Wearable data are frequently incomplete due to charging gaps, sensor detachment, or communication failures, and models

that assume complete input vectors may degrade significantly in real deployments. Imputation methods can fill missing values, but they must be used cautiously because overconfident imputation can mask uncertainty and lead to false reassurance or alarm fatigue. Probabilistic outputs and uncertainty quantification are therefore important components of intelligent wearable analytics, enabling clinicians to interpret model confidence in the context of incomplete or noisy observations. Emerging communication technologies such as 5G can support lower latency and more reliable streaming of multimodal data, but they also introduce new infrastructure and energy trade-offs that must be evaluated within the broader system architecture [16].

A further consideration is the temporal dimension. Chronic disease monitoring often benefits from long-term trend analysis rather than instantaneous classification. Recurrent neural networks, temporal convolutional networks, and attention-based models can capture temporal dependencies across days or weeks, but they require careful handling of irregular sampling intervals and personalized baselines. The use of population-level models may miss individual deviations that are clinically meaningful, while personalized models trained on individual patient data may be vulnerable to overfitting when data are limited. Hybrid approaches that combine population pretraining with local adaptation offer a promising direction, but they introduce additional governance and validation burdens. The integration of multimodal wearable analytics into electronic health records requires standardized representations of derived features, alarms, and model metadata. Without such standardization, learned models may remain isolated research artifacts rather than becoming reliable components of clinical decision support.

5. Robustness and Context Adaptation

Robustness in wearable remote monitoring refers to the ability of the system to maintain acceptable performance across varying sensor placement, device hardware, patient behavior, and environmental conditions. Motion artifact is one of the most persistent threats to photoplethysmography and electrocardiography quality, and its characteristics differ between controlled laboratory settings and free-living conditions. Adaptive filtering methods that rely on reference accelerometry or motion-induced frequency tracking have improved performance, but many such methods are optimized for specific activities and may fail during transitions between walking, running, and postural changes [14]. Signal quality indices are essential for robustness because they allow the system to reject or down-weight corrupted segments rather than propagating errors into diagnostic algorithms [8]. However, quality thresholds are often tuned on healthy volunteers and may not generalize to patients with low perfusion, arrhythmias, or movement disorders. Dynamic quality assessment that adapts to the patient's baseline and activity context can improve performance, but it requires continuous calibration and validation. The system must also handle missing data gracefully, transitioning among modalities or switching to lower-frequency monitoring when high-quality streaming is not feasible. This graceful degradation is a key architectural principle for remote patient monitoring, because temporary loss of connectivity or sensor contact should not produce spurious alarms or silent failures.

Context adaptation extends robustness by using activity, location, and circadian information to adjust monitoring intensity and alarm thresholds. For example, a resting heart rate threshold that is appropriate during sleep may be too sensitive during daytime physical activity. Wearable systems can infer activity state from accelerometry and automatically adjust alarm logic, but this adaptation must be transparent to clinicians to avoid hidden rules

that undermine trust. Alert fatigue is a well-documented problem in continuous monitoring, and context adaptation can reduce false alarms by suppressing alerts that are likely explained by benign activity. However, excessive suppression may delay detection of true deterioration, especially in patients with atypical activity patterns. Designing adaptive thresholds that balance sensitivity and specificity across heterogeneous populations requires extensive validation and ongoing audit.

6. Governance, Fairness, and Policy Implications

The deployment of intelligent wearable systems in healthcare is governed not only by technical performance but also by data protection, algorithmic accountability, and equity. Wearable devices generate continuous, high-dimensional, and potentially identifiable physiological data, raising significant privacy concerns. Data minimization, purpose limitation, and user control are core principles that must be embedded into system architecture rather than treated as afterthoughts. Federated learning and on-device processing can reduce the volume of raw data leaving the patient's device, but they do not eliminate all privacy risks, and model updates may still leak information about individual participants [11]. Strong encryption, access controls, and audit logs are essential, but technical safeguards alone are insufficient without clear governance frameworks that define data ownership, consent, and secondary use.

Algorithmic fairness is a critical concern because wearable systems are often trained and validated on populations that do not reflect the diversity of actual patients. Photoplethysmographic signals, for example, can be affected by skin melanin concentration, peripheral perfusion, and sensor contact pressure, leading to systematic performance differences across demographic groups [13]. If such differences are not explicitly measured and addressed, risk prediction models may inadvertently encode bias and exacerbate health disparities. The broader literature on algorithmic fairness in healthcare has shown that biased proxies and unequal data representation can lead to substantial disparities in care allocation [17]. Responsible machine learning frameworks emphasize the need for transparent reporting, external validation, and ongoing monitoring for differential performance across subgroups [18]. Wearable systems should therefore include fairness audits as part of the development and post-deployment lifecycle.

Regulatory oversight for wearable remote monitoring remains fragmented. Many devices are marketed as wellness products rather than medical devices, which can limit regulatory scrutiny while still influencing clinical decisions. The distinction between wellness and medical use depends on intended purpose, but the boundaries are increasingly blurred when patient-generated data are integrated into clinical workflows. Digital biomarkers derived from wearable sensors require rigorous analytical and clinical validation before they can support diagnosis or treatment decisions. Regulatory frameworks such as the European Medical Device Regulation and the United States Food and Drug Administration's digital health policies provide pathways for software as a medical device, but the rapid pace of algorithmic change poses challenges for static approval processes. Adaptive algorithms that update over time may require new approaches to post-market surveillance, including continuous performance monitoring and predefined update protocols. Policy must also address interoperability, data portability, and patient access to their own longitudinal records, which are foundational for trust and equity.

7. Sustainability and Deployment Considerations

Sustainable deployment of wearable remote monitoring extends beyond technical feasibility to include energy consumption, device lifecycle, clinical workflow integration, and economic viability. Wearable sensors and communication modules consume energy continuously, and frequent recharging burdens patients and reduces adherence. Energy-efficient signal acquisition can be achieved through adaptive sampling, duty cycling, and local preprocessing that reduces raw data transmission [10]. However, aggressive energy optimization can degrade signal quality or delay detection of clinically relevant events. System designers must consider the trade-off between battery life and monitoring continuity, particularly for populations with limited access to reliable charging infrastructure. Battery degradation over time also affects sensor stability and signal quality, requiring recalibration or replacement schedules that may be difficult to enforce in remote settings.

Infrastructure sustainability is closely linked to interoperability and modularity. Wearable devices from different manufacturers often use proprietary data formats and cloud services, creating fragmentation that hinders integration with electronic health records and secondary research. Open data standards and application programming interfaces can reduce lock-in and enable plug-and-play replacement of components, but they also introduce security and validation challenges. The Internet of Things and emerging communication technologies such as 5G can support more reliable and lower-latency transmission of multimodal data, but they may also increase energy demands and require new network infrastructure that is unevenly distributed across regions [16]. Equitable access to connectivity is therefore a determinant of health monitoring effectiveness, and deployment strategies must account for rural, low-income, and historically underserved communities.

Clinical workflow integration is a socio-technical challenge that is often underestimated. A remote monitoring system that generates a continuous stream of alerts may overwhelm clinicians unless it is designed to filter, prioritize, and present alarms in a clinically meaningful manner. Alert fatigue is a well-documented hazard in continuous telemetry and remote care settings, and intelligent systems must incorporate escalation pathways, contextual metadata, and user-specific preferences to reduce cognitive burden. The interpretability of model outputs is therefore not a secondary feature but a core requirement for clinical trust. Models that produce risk scores without explanation may be rejected in practice, particularly when they conflict with established clinical judgment. Explanations grounded in the contributing physiological modalities can help clinicians assess whether an alert reflects artifact, context, or genuine deterioration. The need for explainability must be balanced with the risk of oversimplification, and designers should avoid presenting post hoc saliency as causal evidence. A transparent workflow may present the temporal evolution of raw signal features alongside model confidence, allowing clinicians to inspect the evidence chain [20].

The sustainability of wearable remote monitoring also depends on patient usability and engagement. Devices that are burdensome to wear, difficult to charge, or incompatible with daily activities will produce incomplete datasets and low adherence. Deployment programs should involve patients in the design of monitoring schedules, feedback mechanisms, and escalation protocols. The assumption that digital health interventions inherently reduce disparities is not supported by evidence; poorly designed systems can widen existing gaps by assuming high digital literacy or reliable connectivity [21]. Federated learning can contribute to privacy-preserving multi-institution model development, but participation in federated networks requires institutional capacity and governance alignment that may not be equally distributed [22]. Economic sustainability includes not only device cost but also the cost of

clinician time, data storage, algorithmic auditing, and technical support. Remote patient monitoring programs that neglect these hidden costs may appear promising in pilot studies but fail to scale. Evidence from randomized controlled trials suggests that remote patient monitoring can improve outcomes in specific clinical contexts, but effect sizes vary and depend heavily on workflow integration and patient selection [23].

8. Future Directions and Open Challenges

The future evolution of intelligent wearable systems will be shaped by the convergence of multimodal sensing with context-aware artificial intelligence, privacy-preserving learning, and automated personalization. One open challenge is the development of standardized multimodal data representations that can accommodate heterogeneous sampling rates, missing data, and device-specific characteristics without losing clinically relevant information. Another challenge is the alignment of adaptive algorithms with regulatory frameworks that were designed for static software. Continuous learning systems that update their parameters over time may require new forms of post-market surveillance and transparent versioning to ensure that performance drift is detected and corrected before patient harm occurs. The environmental footprint of large-scale artificial intelligence in healthcare also deserves attention, because the computational demands of cloud-based training and inference can be substantial, and sustainability must be evaluated across the entire lifecycle of data collection, model training, and device operation [24]. Cross-border governance is particularly complex for wearable monitoring because data may flow across jurisdictions with different privacy laws and clinical accountability standards. International harmonization of digital health regulations, including mutual recognition of software as a medical device and shared approaches to cybersecurity, would facilitate equitable deployment [25].

9. Conclusion

Intelligent wearable systems for remote patient monitoring represent a promising but complex convergence of multimodal signal processing, data infrastructure, machine learning, and healthcare governance. This paper has argued that the clinical value of such systems depends less on isolated algorithmic breakthroughs than on coherent architectural design that balances local and remote processing, protects privacy, and adapts robustly to real-world variability. Photoplethysmography, electrocardiography, accelerometry, and other modalities each contribute complementary information, and their fusion can improve resilience to artifact and missing data when synchronization and quality estimation are treated as foundational components. Adaptive notch-filtration techniques demonstrate that signal-level robustness remains indispensable, particularly during physical activity [14]. Yet technical accuracy alone is insufficient; systems must also address algorithmic fairness, explainability, regulatory alignment, and sustainable deployment. The future development of wearable monitoring will require interdisciplinary collaboration across engineering, clinical medicine, data science, law, and public health, as well as sustained attention to the patient populations that stand to benefit most. By integrating these perspectives, the next generation of intelligent wearable systems can become trustworthy, equitable, and enduring components of healthcare.

References

1. Majumder, S., Mondal, T., & Deen, M. J. (2017). Wearable sensors for remote health monitoring. *Sensors*, 17(1), 130. <https://doi.org/10.3390/s17010130>

2. Patel, S., Park, H., Bonato, P., Chan, L., & Rodgers, M. (2012). A review of wearable sensors and systems with application in rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, 9, 21. <https://doi.org/10.1186/1743-0003-9-21>
3. Piwek, L., Ellis, D. A., Andrews, S., & Joinson, A. (2016). The rise of consumer health wearables: Promises and barriers. *PLoS Medicine*, 13(2), e1001953. <https://doi.org/10.1371/journal.pmed.1001953>
4. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., & Dean, J. (2019). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24–29. <https://doi.org/10.1038/s41591-018-0316-z>
5. Goldberger, A. L., Amaral, L. A. N., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C.-K., & Stanley, H. E. (2000). PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. *Circulation*, 101(23), e215–e220. <https://doi.org/10.1161/01.CIR.101.23.e215>
6. 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. <https://doi.org/10.1109/TIM.2011.2175832>
7. 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. <https://doi.org/10.1109/TBME.2014.2359372>
8. 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. <https://doi.org/10.1109/JBHI.2014.2338351>
9. Mukhopadhyay, S. C. (2015). Wearable sensors for human activity monitoring: A review. *IEEE Sensors Journal*, 15(3), 1321–1330. <https://doi.org/10.1109/JSEN.2014.2370945>
10. Rahmani, A. M., Gia, T. N., Negash, B., Anzanpour, A., Azimi, I., Jiang, M., & Liljeberg, P. (2018). Exploiting smart e-Health gateways at the edge of healthcare Internet-of-Things: A fog computing approach. *Future Generation Computer Systems*, 78, 641–658. <https://doi.org/10.1016/j.future.2017.02.014>
11. Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R. R., & Bakas, S. (2020). Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 10, 12598. <https://doi.org/10.1038/s41598-020-69250-1>
12. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56. <https://doi.org/10.1038/s41591-018-0300-7>
13. Allen, J. (2007). Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3), R1–R39. <https://doi.org/10.1088/0967-3334/28/3/R01>

14. 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.
15. Charlton, P. H., Bonnici, T., Tarassenko, L., Clifton, D. A., Beale, R., & Watkinson, P. J. (2016). An assessment of algorithms to estimate respiratory rate from the electrocardiogram and photoplethysmogram. *Physiological Measurement*, 37(4), 610–626. <https://doi.org/10.1088/0967-3334/37/4/610>
16. Ahad, A., Tahir, M., & Yau, K.-L. A. (2019). 5G-based smart healthcare network: Architecture, taxonomy, challenges and future research directions. *IEEE Access*, 7, 950–963. <https://doi.org/10.1109/ACCESS.2019.2907635>
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. <https://doi.org/10.1126/science.aax2342>
18. Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V. X., Doshi-Velez, F., Jung, K., Heller, K., Kale, D., Saeed, M., Ossorio, P. N., Thadaney-Israni, S., & Goldenberg, A. (2019). Do no harm: A roadmap for responsible machine learning for health care. *Nature Medicine*, 25(9), 1337–1340. <https://doi.org/10.1038/s41591-019-0548-6>
19. Dunn, J., Runge, R., & Snyder, M. (2018). Wearables and the medical revolution. *Personalized Medicine*, 15(5), 429–448. <https://doi.org/10.2217/pme-2018-0044>
20. Tonekaboni, S., Joshi, S., McCradden, M. D., & Goldenberg, A. (2019). What clinicians want: Contextualizing explainable machine learning for clinical end use. *arXiv preprint arXiv:1905.05134*.
21. Veinot, T. C., Mitchell, H., & Ancker, J. S. (2018). Good intentions are not enough: How informatics interventions can worsen inequality. *Journal of the American Medical Informatics Association*, 25(8), 1080–1088. <https://doi.org/10.1093/jamia/ocy052>
22. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. *NPJ Digital Medicine*, 3, 119. <https://doi.org/10.1038/s41746-020-00323-1>
23. Noah, B., Keller, M. S., Mosadeghi, S., Stein, M., Johl, S., Delshad, S., Tashjian, V. C., Lew, D., Kwan, J. T., Jusufagic, A., & Spiegel, B. M. R. (2018). Impact of remote patient monitoring on clinical outcomes: An updated meta-analysis of randomized controlled trials. *NPJ Digital Medicine*, 1, 20172. <https://doi.org/10.1038/s41746-017-0002-4>
24. Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S. D., Tegmark, M., & Fuso Nerini, F. (2020). The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11, 233. <https://doi.org/10.1038/s41467-019-14108-y>
25. Gerke, S., Stern, A. D., & Minssen, T. (2020). Germany's digital health reforms in the COVID-19 era: Lessons and opportunities for other countries. *NPJ Digital Medicine*, 3, 94. <https://doi.org/10.1038/s41746-020-0306-7>