

Multimodal Artificial Intelligence for Early Detection of Cardiometabolic Risk Using Wearable Data

Raohe Long

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

long96r@unr.edu

Ivan Adams

School of Computing, Clemson University, Clemson, SC, USA.

ivanmail@clemson.edu

Dylan L. Lowe

Department of Computer Science, Binghamton University, Binghamton, NY, USA.

lowe.dylan@binghamton.edu

Rowan Harrison

Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO, USA.

harrisonrowan@missouri.edu

Abstract

Cardiometabolic risk arises from interacting physiological, behavioral, and environmental processes that frequently precede overt cardiovascular disease and type 2 diabetes. Wearable devices now generate continuous, multimodal data streams that can capture early perturbations in cardiovascular, metabolic, respiratory, and activity-related function. Artificial intelligence offers a promising pathway for transforming these heterogeneous data into clinically meaningful early warning signals. However, translating multimodal learning into reliable cardiometabolic risk detection requires addressing deeper system-level challenges beyond model accuracy. This paper presents a systems-oriented examination of multimodal artificial intelligence for early detection of cardiometabolic risk using wearable data. It analyzes the structural characteristics of wearable data streams, the architectural trade-offs involved in multimodal fusion, and the infrastructure requirements for longitudinal data integration. The discussion further considers robustness under real-world deployment, algorithmic fairness, privacy preservation, regulatory governance, and long-term system sustainability. Throughout, the paper argues that the central scientific and practical challenge is not solely the design of predictive models, but the construction of an integrated sociotechnical infrastructure in which multimodal models can operate safely, equitably, and sustainably. The analysis draws on recent advances in wearable sensing, deep learning, federated learning, clinical decision support, and algorithmic ethics to identify structural conditions under which early cardiometabolic risk detection may become a dependable component of preventive health systems.

Keywords

wearable sensors; multimodal artificial intelligence; cardiometabolic risk; digital biomarkers; federated learning; algorithmic fairness.

1. Introduction

Cardiometabolic disease continues to impose a substantial and growing burden on health systems worldwide, driven by the combined effects of hypertension, dyslipidemia, obesity, insulin resistance, and related social and environmental determinants [1], [2]. Early detection of cardiometabolic risk offers a critical opportunity to shift care from reactive treatment toward preventive intervention. Yet traditional risk assessment often depends on sporadic clinical measurements, self-reported behaviors, and population-level risk equations that may fail to capture individual temporal dynamics. Wearable devices have begun to change this picture by enabling continuous or near-continuous measurement of physiological signals outside clinical settings. These devices generate longitudinal data on heart rate, heart rate variability, physical activity, sleep, respiration, skin temperature, and increasingly metabolic indicators. The resulting data streams provide a foundation for detecting subtle changes that may precede clinically apparent disease.

Recent advances in artificial intelligence have expanded the analytical possibilities for wearable data. Deep learning methods can learn hierarchical representations from large physiological datasets, while multimodal modeling can integrate diverse signals that individually carry limited predictive value [3], [4], [5]. In cardiometabolic risk assessment, this integration is particularly important because no single wearable signal reliably captures the full complexity of metabolic dysregulation, autonomic imbalance, vascular dysfunction, and behavioral risk. A multimodal system can combine patterns across cardiac, activity, sleep, and respiratory data to produce a more coherent picture of physiological deterioration. However, realizing this potential requires careful attention to how such systems are designed, deployed, maintained, and governed across heterogeneous populations and care contexts.

This paper addresses the problem from a systems perspective. Rather than focusing narrowly on the performance of a specific predictive model, it examines the structural trade-offs that influence whether multimodal artificial intelligence can become a trustworthy component of early cardiometabolic risk detection. The paper analyzes the nature of wearable data streams, the architecture of multimodal learning systems, infrastructure requirements, generalization and robustness challenges, fairness and governance concerns, and long-term sustainability. In doing so, it argues that meaningful progress depends on aligning technical design with clinical workflows, regulatory realities, privacy constraints, and equity considerations.

2. Cardiometabolic Risk and Wearable Data Streams

Cardiometabolic risk encompasses a constellation of interrelated conditions that increase the likelihood of cardiovascular events and metabolic disease. These conditions include elevated blood pressure, abnormal lipid profiles, central adiposity, impaired glucose tolerance, and chronic low-grade inflammation. Because these processes develop over years, they are often present long before a clinical diagnosis. Wearable sensors can capture physiological changes associated with such processes, including reduced heart rate variability, altered circadian rhythms, increased resting heart rate, disrupted sleep architecture, and reduced physical activity. Optical sensing approaches can extract vital signs under varying physical activity intensity, enabling continuous physiological monitoring in ambulatory conditions [6]. Such capability is important because cardiometabolic decompensation may be episodic or context-dependent, making isolated clinic measurements insufficient.

The informational value of wearable data depends on its temporal resolution, duration, and contextual richness. A single accelerometer count is relatively uninformative, but months of

continuous activity data can reveal patterns of sedentary behavior, exercise recovery, and circadian stability. Likewise, photoplethysmographic signals can support estimates of heart rate and pulse arrival characteristics, while electrodermal activity, temperature, and respiratory rate contribute complementary information about autonomic and thermoregulatory function [7], [8]. The medical relevance of these signals has been increasingly recognized, although their clinical integration remains uneven. Wearable data are characterized by high missingness, motion artifacts, sensor drift, and inter-individual variability, all of which pose significant analytical challenges. Successful early detection systems must therefore treat data quality and contextual uncertainty as first-order design problems rather than as preprocessing inconveniences.

The clinical deployment of wearables has expanded beyond wellness applications toward medical monitoring and risk stratification [9]. For cardiometabolic risk, this shift requires alignment with clinical evidence and preventive guidelines that define actionable thresholds and intervention pathways [10]. Wearable-derived biomarkers may not directly correspond to established laboratory measures, such as fasting glucose or low-density lipoprotein cholesterol. Consequently, an early detection system must formulate risk signals that are interpretable to clinicians and actionable within existing preventive care frameworks. A multimodal system can help bridge this gap by combining surrogate physiological signals with behavioral and contextual data, thereby producing estimates that are more robust than any single surrogate marker alone.

3. Multimodal Learning Architectures and System Trade-offs

Multimodal artificial intelligence for wearable data can be organized around three broad architectural strategies: early fusion, late fusion, and hierarchical or hybrid fusion. Early fusion combines raw or lightly processed signals before learning, allowing the model to discover cross-modal interactions. This approach can capture subtle dependencies, but it may be sensitive to missing modalities and heterogeneous sampling rates. Late fusion learns separate representations for each modality and combines them at a decision stage, which offers greater modularity but may miss early interactions. Hybrid architectures attempt to balance these trade-offs by learning modality-specific representations while also modeling higher-order correlations. None of these strategies is universally superior; the appropriate design depends on the stability of the sensors, the clinical use case, and the available computational infrastructure.

The structural trade-offs extend beyond model architecture. Multimodal systems require careful synchronization of data streams collected at different frequencies. Accelerometry may be sampled at tens of hertz, while glucose or temperature measurements may be recorded every few minutes. Sleep stages may be aggregated over hours, whereas heart rate variability features may require specific time windows. Aligning these streams into a coherent learning input requires imputation, temporal aggregation, and feature engineering that can introduce bias. If not carefully managed, these choices can amplify noise or obscure clinically relevant events. Furthermore, multimodal models often require substantially more training data and computational resources than unimodal models, raising sustainability concerns in resource-constrained health systems.

Another important system-level consideration is the potential for multimodal models to learn population-level associations that are not causally related to cardiometabolic risk. Wearable data reflect behavior, environment, and device usage patterns in addition to physiology. A model trained on such data may inadvertently learn device-specific or demographic-specific

signals rather than generalizable biomarkers. This challenge has been widely observed in clinical machine learning, where models can encode existing disparities in care and access [11]. Multimodal fusion may either mitigate or exacerbate such problems depending on how data are collected, labeled, and evaluated. A well-designed system should therefore include explicit evaluation across subgroups and settings rather than relying on aggregate performance metrics.

The architectural choices also interact with privacy and data governance. Centralized training on large repositories of wearable data offers computational convenience but raises concerns about data security and patient consent. Federated learning has emerged as an alternative in which models are trained across distributed devices or institutions without transferring raw data [12]. This strategy aligns well with wearable monitoring because data naturally reside on personal devices or within regional health systems. Federated digital health systems can support collaborative learning while reducing direct data exposure [13], [14]. However, federated approaches introduce additional complexity in model aggregation, communication efficiency, and protection against inference attacks. Privacy-preserving machine learning is therefore not simply a legal requirement but a systems design constraint that influences model architecture, data flow, and deployment feasibility [15].

4. Data Infrastructure and Integration Challenges

A robust early detection system for cardiometabolic risk depends on more than model development. It requires a data infrastructure that can ingest, clean, harmonize, and steward longitudinal wearable data across devices, platforms, and clinical contexts. Many wearable ecosystems are fragmented, with proprietary data formats and application programming interfaces that complicate integration. Data from different manufacturers may use different units, sampling rates, and quality flags. Building a reliable multimodal system therefore requires substantial investment in data standardization, metadata management, and provenance tracking. Without such infrastructure, model performance may degrade when devices are updated or when patients switch between platforms.

The integration of wearable data with electronic health records introduces further complexity. Cardiometabolic risk assessment ideally combines device-derived signals with laboratory results, medication histories, diagnoses, and social determinants of health. However, health record data are often irregular, incomplete, and semantically heterogeneous. Linking wearable data to clinical outcomes requires careful identity management and temporal alignment. In some settings, data linkage may be constrained by privacy regulations and institutional policies. These challenges highlight the need for governance frameworks that support responsible data sharing while maintaining patient trust [16].

Explainability and interpretability are also infrastructure concerns because clinical adoption depends on whether care teams can understand and act on model outputs. A continuous risk score derived from multimodal wearable data may be useful, but it must be accompanied by explanations that identify the main contributing signals and their temporal patterns. Different stakeholders may require different explanations: a cardiologist may want physiological detail, while a primary care nurse may prefer a simple risk trend. Explainability frameworks for healthcare artificial intelligence emphasize the need to connect model outputs to clinical reasoning [17], [18], although current approaches to post hoc explanation have well-documented limitations [19]. The design of explanation systems should therefore be treated as an integral part of the sociotechnical infrastructure rather than as a cosmetic layer added after model training.

The translation of multimodal models into useful clinical tools also requires alignment with clinical decision support systems. The National Academy of Medicine has highlighted the importance of evaluating artificial intelligence not only for predictive accuracy but also for workflow integration, usability, and impact on patient outcomes [20]. Wearable-based early detection systems must fit into existing preventive care pathways, which may include routine primary care visits, remote monitoring programs, or community health initiatives. Data infrastructure must support timely alerts without overwhelming clinicians, and it must enable longitudinal review rather than isolated point predictions. The sustainability of such systems depends on their ability to integrate with electronic health record workflows and to generate evidence of clinical value.

5. Generalization, Robustness and Deployment Realities

Generalization is a persistent challenge for wearable-based artificial intelligence. Models trained on data from a particular device, population, or environment may fail when deployed elsewhere. Sensor characteristics differ across manufacturers and generations, and physiological responses vary by age, sex, medication use, and comorbid conditions. Physical activity levels also influence signal quality; motion artifacts can corrupt optical measurements and reduce the reliability of derived features. The accuracy of wearable optical heart rate sensors is known to vary across skin tones, sensor placement, and activity conditions [21]. Respiratory signals derived from wearables may similarly be affected by motion and environmental noise [22]. These sources of variation must be explicitly addressed through domain adaptation, robust feature construction, and continuous monitoring of model performance.

Robustness in real-world deployment requires more than algorithmic adjustment. It demands feedback loops that can detect drift in sensor performance, changes in user behavior, and shifts in the underlying population. A model that performed well during a controlled validation study may degrade when deployed in a routine care setting with lower adherence, older devices, or different data collection patterns. Continuous postmarket surveillance is therefore essential. Regulatory frameworks for clinical decision support increasingly emphasize lifecycle management and real-world performance monitoring [23]. For wearable-based risk detection, this means that deployment should be treated as an ongoing system engineering process rather than a one-time release.

The temporal nature of cardiometabolic risk adds another layer of complexity. Early detection is not a single classification task but a longitudinal monitoring problem. Risk may evolve over months or years, and short-term fluctuations may be less important than sustained trends. A robust system must distinguish between transient physiological variation and meaningful deterioration. This requires models that can represent temporal context, handle irregular sampling, and incorporate prior clinical knowledge. It also requires careful definition of outcome labels, because the transition from risk to disease is often gradual and may not be clearly documented in administrative data. Multimodal systems must therefore be designed with longitudinal evaluation protocols that assess calibration, stability, and responsiveness to intervention.

Case illustrations from remote patient monitoring and critical care inform the design of wearable-based early warning systems. In intensive care, continuous physiological monitoring has enabled the development of early warning scores that synthesize multiple vital signs. However, these systems also illustrate the risk of alarm fatigue and the importance of contextualizing alerts within clinical workflows. Wearable-based cardiometabolic monitoring

operates in a much less controlled environment, with less frequent clinical verification. The lessons from critical care suggest that early detection systems should prioritize high-specificity alerts, transparent uncertainty communication, and integration with human judgment. They also suggest that model performance alone is insufficient to guarantee clinical benefit.

6. Fairness, Governance and Policy Implications

Fairness is a central concern for multimodal artificial intelligence in cardiometabolic prevention. Wearable data are not neutral representations of physiology; they are shaped by device access, sensor performance, user behavior, and social determinants of health. If certain populations are underrepresented in training data or experience systematically lower signal quality, models may produce inaccurate risk estimates for those groups. Algorithmic bias in healthcare has been shown to produce significant disparities even when models appear accurate on aggregate [11]. For wearable-based cardiometabolic risk detection, bias may arise from differences in skin tone, body composition, activity patterns, or access to higher-quality devices. A fair system must therefore be evaluated not only on overall accuracy but also on performance across demographic, clinical, and socioeconomic subgroups.

Governance mechanisms are needed to ensure that multimodal systems are developed and deployed responsibly. International guidance on the ethics and governance of artificial intelligence for health emphasizes human rights, data protection, transparency, and accountability [24]. These principles are particularly relevant for wearable data because such data are collected continuously in private settings and may reveal intimate details about behavior, sleep, and physiological status. Patients must have meaningful control over how their data are used, and institutions must establish clear lines of accountability for model errors or harmful recommendations. Governance frameworks should address the full lifecycle of the system, from data collection and model development to clinical deployment and decommissioning.

Regulatory policy also shapes the design and deployment of wearable-based risk detection systems. In many jurisdictions, software intended for diagnostic or therapeutic purposes may be regulated as a medical device. The boundaries between wellness applications and medical devices remain ambiguous for many wearable systems, particularly when a risk score is used to guide preventive care. Clarity in regulatory classification is important because it affects evidentiary requirements, quality management, and postmarket obligations. The governance of artificial intelligence in health is evolving rapidly, and policies must balance innovation with patient safety. A robust regulatory approach should require validation on clinically meaningful outcomes, transparent reporting of limitations, and continuous monitoring after deployment.

The social and ethical dimensions of early cardiometabolic detection extend beyond individual-level fairness. Population-level screening with wearable devices could create new forms of risk stratification that affect insurance, employment, or access to services. If a model labels someone as high risk based on noisy or biased data, that label may carry consequences even if the underlying prediction is uncertain. Ethical machine learning in healthcare requires attention to these downstream effects and to the ways in which risk information is communicated [25]. A system-level perspective therefore includes not only the predictive model but also the institutional policies, consent processes, and communication strategies that surround it.

7. Sustainability and Long-Term Maintenance

The long-term sustainability of multimodal artificial intelligence for cardiometabolic risk depends on organizational, technical, and financial factors. Many artificial intelligence projects in healthcare fail to progress beyond pilot studies because they are not integrated into routine operations or because they require continuous data curation that institutions cannot sustain. Wearable-based systems are particularly vulnerable to sustainability challenges because they depend on consumer devices with short lifecycles, changing software platforms, and variable user adherence. A sustainable system must be designed for maintainability, with mechanisms for updating models as devices and populations change. It must also include realistic cost models that account for data storage, computational infrastructure, clinician time, and technical support.

Continuous learning introduces further sustainability considerations. Unlike static clinical decision rules, machine learning models may need periodic retraining to maintain performance. However, repeated retraining can introduce instability if it is not governed by clear protocols. Federated learning can support ongoing improvement across distributed sites without centralizing data, but it requires durable technical governance and sustained collaboration among participating institutions [12], [13]. The sustainability of federated networks depends on trust, shared standards, and equitable distribution of benefits. Without these organizational supports, technical interoperability alone will not ensure long-term viability.

The environmental and computational footprint of large multimodal models also merits attention. Training deep learning models on high-frequency wearable data can be resource-intensive, particularly when models are retrained frequently or deployed across many devices. Although edge computing can reduce data transfer and central processing demands, it introduces trade-offs in model complexity and energy consumption. Sustainable design should consider lightweight architectures, efficient sampling strategies, and selective data transmission. The goal is not only to maximize predictive performance but also to ensure that the system can operate within the resource constraints of real health systems.

Long-term success also requires meaningful clinical engagement. Early detection systems will only improve outcomes if clinicians and patients can act on the information they provide. This requires education, workflow redesign, and clear guidance on how to interpret risk estimates. It also requires ongoing evaluation of whether alerts lead to appropriate preventive actions without causing unnecessary anxiety or overmedicalization. The sustainability of cardiometabolic risk detection therefore depends on its ability to become embedded in a broader culture of preventive care, supported by evidence, infrastructure, and governance. Technical innovation alone cannot deliver this transformation.

8. Conclusion

Multimodal artificial intelligence has substantial potential to improve the early detection of cardiometabolic risk using wearable data. Continuous physiological and behavioral sensing can reveal patterns that are invisible in isolated clinical measurements, and advanced learning methods can integrate these patterns into coherent risk estimates. However, the translation of this potential into dependable health systems requires attention to far more than predictive accuracy. The design of multimodal architectures involves complex trade-offs among data synchronization, model interpretability, computational cost, and robustness. Data infrastructure must support integration across fragmented device ecosystems and clinical

records while preserving privacy and enabling meaningful clinical interpretation. Deployment in real-world settings demands continuous surveillance, subgroup-level evaluation, and responsiveness to device and population drift. Fairness and governance frameworks are essential to ensure that early detection does not amplify existing disparities or create new forms of risk. Ultimately, wearable-based cardiometabolic risk detection should be understood as a sociotechnical system in which sensing devices, learning models, clinical workflows, regulatory policies, and patient expectations interact. Progress will depend on building institutions and infrastructures that can support such systems over time, not merely on producing more complex models. The future of early cardiometabolic risk detection lies in the careful integration of multimodal artificial intelligence into equitable, explainable, and sustainable preventive care pathways.

References

1. World Health Organization. (2021). Cardiovascular diseases (CVDs) fact sheet. World Health Organization. [https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-\(cvds\)](https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-(cvds))
2. Roth, G. A., Mensah, G. A., Johnson, C. O., Addolorato, G., Ammirati, E., Baddour, L. M., Barengo, N. C., Beaton, A. Z., Benjamin, E. J., Benziger, C. P., Bonny, A., Brauer, M., Brodmann, M., Cahill, T. J., Carapetis, J., Catapano, A. L., Chugh, S. S., Cooper, L. T., Coresh, J., ... Fuster, V. (2020). Global burden of cardiovascular diseases and risk factors, 1990–2019: Update from the GBD 2019 study. *Journal of the American College of Cardiology*, 76(25), 2982–3021. <https://doi.org/10.1016/j.jacc.2020.11.010>
3. 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>
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. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *New England Journal of Medicine*, 380(14), 1347–1358. <https://doi.org/10.1056/NEJMr1814259>
6. 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.
7. 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>
8. Witt, D. R., Kellogg, R. A., Snyder, M. P., & Dunn, J. (2019). Windows into human health through wearables data analytics. *Current Opinion in Biomedical Engineering*, 9, 28–46. <https://doi.org/10.1016/j.cobme.2019.01.001>
9. Smuck, M., Odonkor, C. A., Wilt, J. K., Schmidt, N., & Swiernik, M. A. (2021). The emerging clinical role of wearables: Factors for successful implementation in healthcare. *npj Digital Medicine*, 4, 45. <https://doi.org/10.1038/s41746-021-00418-3>
10. Arnett, D. K., Blumenthal, R. S., Albert, M. A., Buroker, A. B., Goldberger, Z. D., Hahn, E. J., Himmelfarb, C. D., Khera, A., Lloyd-Jones, D., McEvoy, J. W., Michos, E. D., Miedema, M. D., Muñoz, D., Smith, S. C., Virani, S. S., Williams, K. A., Yeboah, J., & Ziaeian, B. (2019). 2019 ACC/AHA guideline on the primary prevention of

cardiovascular disease. *Circulation*, 140(11), e596–e646.
<https://doi.org/10.1161/CIR.0000000000000678>

11. 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>
12. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 54, 1273–1282.
13. 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>
14. Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, 2(6), 305–311. <https://doi.org/10.1038/s42256-020-0186-1>
15. Price, W. N., & Cohen, I. G. (2019). Privacy in the age of medical big data. *Nature Medicine*, 25(1), 37–43. <https://doi.org/10.1038/s41591-018-0272-7>
16. Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. *PLoS Medicine*, 15(11), e1002689. <https://doi.org/10.1371/journal.pmed.1002689>
17. Amann, J., Blasimme, A., Vayena, E., Frey, D., & Madai, V. I. (2020). Explainability for artificial intelligence in healthcare: A multidisciplinary perspective. *BMC Medical Informatics and Decision Making*, 20, 310. <https://doi.org/10.1186/s12911-020-01332-6>
18. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
19. Ghassemi, M., Oakden-Rayner, L., & Beam, A. L. (2020). The false hope of current approaches to explainable artificial intelligence in health care. *The Lancet Digital Health*, 2(12), e745–e750. [https://doi.org/10.1016/S2589-7500\(20\)30240-7](https://doi.org/10.1016/S2589-7500(20)30240-7)
20. Matheny, M. E., Israni, S. T., Ahmed, M., & Whicher, D. (Eds.). (2019). *Artificial intelligence in health care: The hope, the hype, the promise, the peril*. National Academy of Medicine.
21. 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. <https://doi.org/10.1038/s41746-020-0226-6>
22. Aliverti, A. (2017). Wearable technology: Role in respiratory health and disease. *Breathe*, 13(2), e27–e36. <https://doi.org/10.1183/20734735.008417>
23. U.S. Food and Drug Administration. (2022). *Clinical decision support software: Guidance for industry and Food and Drug Administration staff*. U.S. Department of Health and Human Services.

24. World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. World Health Organization.
25. Chen, I. Y., Pierson, E., Rose, S., Joshi, S., Ferryman, K., & Ghassemi, M. (2021). Ethical machine learning in healthcare. *Annual Review of Biomedical Data Science*, 4, 123–144. <https://doi.org/10.1146/annurev-biodatasci-092820-114757>