

AI-Assisted Early Detection of Cardiovascular Abnormalities Using Continuous Wearable Sensor Signals

Vijay Merthiy

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis,
OR, USA.

vijay1978@oregonstate.edu

Jiangtian Zhao

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

zhaojiang@missouri.edu

Abstract

The growing burden of cardiovascular disease has intensified interest in continuous monitoring paradigms that extend beyond episodic clinical encounters. Wearable devices now acquire high-frequency photoplethysmographic, electrocardiographic, accelerometric, and contextual signals over prolonged periods, creating opportunities for artificial intelligence systems to detect subtle anomalies before acute events. However, translating raw sensor streams into clinically trustworthy early warning requires careful consideration of sensing quality, system architecture, data governance, robustness, fairness, and deployment sustainability. This article presents a system-level examination of AI-assisted early detection of cardiovascular abnormalities using continuous wearable sensor signals. It discusses how motion artifact, signal degradation, and interindividual variability shape sensing infrastructure. The paper then analyzes architectural choices spanning on-device processing, edge aggregation, cloud inference, and adaptive model updating. It further addresses privacy-preserving learning, algorithmic bias, uncertainty quantification, regulatory validation, interoperability, and lifecycle sustainability. Rather than focusing on a single algorithm, the analysis emphasizes structural trade-offs that determine whether such systems can operate safely in real-world health ecosystems. The article argues that early detection platforms should be designed as socio-technical infrastructures with transparent governance, prospective clinical evidence, and mechanisms for continuous monitoring of model performance. These considerations are essential for moving beyond diagnostic accuracy toward equitable, robust, and sustainable cardiovascular care.

Keywords

continuous wearable sensing; cardiovascular abnormality detection; artificial intelligence; system architecture; health data governance; clinical deployment.

1. Introduction

Cardiovascular diseases remain a leading source of mortality and long-term disability worldwide, imposing substantial burdens on health systems and societies [1]. Traditional diagnostic pathways frequently depend on episodic measurements obtained in clinics, emergency departments, or short-term ambulatory monitoring programs. These encounters

provide only fragmented views of a patient's physiological state and may miss paroxysmal abnormalities, gradual deterioration, or early decompensation. Wearable sensors have emerged as a complementary infrastructure because they can collect physiological data continuously in uncontrolled everyday environments. Photoplethysmography and single-lead electrocardiography are now embedded in consumer wristbands, chest patches, earbuds, and medical-grade wearables, allowing longitudinal observation of cardiovascular function outside clinical settings [2]. Large-scale evaluation programs have demonstrated that smartwatch-based screening can identify undiagnosed atrial fibrillation in broad populations, although such programs also raise questions about false positives, follow-up burden, and clinical integration [3]. In parallel, advances in deep learning have shown that automated models can classify arrhythmias from ambulatory electrocardiographic recordings with performance approaching that of cardiologists under certain conditions [4]. These developments have generated enthusiasm for AI-assisted early detection, but they also reveal an important gap. High diagnostic accuracy in curated datasets does not automatically translate into safe, equitable, and sustainable performance in heterogeneous real-world environments. A broader system-level perspective is therefore required. This paper examines the sensing infrastructure, architectural design, governance mechanisms, robustness requirements, deployment considerations, and policy implications of AI-assisted cardiovascular early warning systems. The discussion focuses not on a single algorithm but on the structural trade-offs and socio-technical arrangements that shape whether continuous wearable sensing can become a trusted component of cardiovascular care.

2. Wearable Sensing Infrastructure and Signal Characteristics

The foundation of any AI-assisted early detection system is the quality and interpretability of the underlying sensor signal. Wearable cardiovascular monitoring typically relies on photoplethysmographic sensors that measure blood volume changes, often combined with electrodes for electrocardiographic acquisition and accelerometers for motion context. Photoplethysmography is particularly attractive for long-term use because it can be embedded in low-cost, low-power wrist-worn devices without requiring adhesive electrodes. However, the signal is highly sensitive to motion, ambient light, skin pigmentation, sensor placement, and peripheral perfusion. When users walk, exercise, or sleep, these factors introduce complex artifacts that can obscure clinically relevant information. Adaptive notch-filtration methods have been proposed to recover photoplethysmographic signals during physical activity by selectively attenuating motion-related frequency components while preserving physiological waveform features [5]. Such preprocessing approaches illustrate the need for robust signal conditioning before machine learning inference. Yet filtering alone cannot resolve all forms of degradation. Signal quality indices and data fusion strategies have therefore been developed to identify which segments of an ambulatory recording are suitable for automated interpretation and which should be rejected or downweighted [6]. These infrastructure-level decisions are critical because a model trained on clean hospital recordings may fail when exposed to noisy wearable data. A comprehensive review of wearable photoplethysmography emphasizes that measurement site, hardware design, sampling rate, and environmental conditions jointly determine the information available for cardiovascular inference [7]. Even when signal quality is acceptable, the clinical meaning of a detected anomaly depends on contextual factors such as physical activity, sleep state, and medication use. Long-term electrocardiographic patch studies have shown that continuous monitoring can uncover paroxysmal atrial fibrillation in patients who would otherwise remain undiagnosed [8]. Nevertheless, continuous data collection introduces new operational challenges, including

storage volume, battery life, data loss, and the need for time synchronization across modalities. The sensing layer must therefore be understood not as a fixed input source but as a dynamic subsystem that requires continuous quality management, adaptive filtering, and contextual awareness.

3. System Architecture for Continuous Cardiovascular Risk Inference

The architecture of an AI-assisted early detection system determines where computation occurs, how models are updated, and how alerts are generated. A fully cloud-centric architecture is conceptually simple because sensor data can be streamed to central servers for preprocessing, feature extraction, and deep model inference. Cloud computing offers powerful resources for training large models, storing long-term records, and aggregating population-level insights. However, continuous streaming of high-frequency physiological data imposes bandwidth, latency, privacy, and energy costs that may be unacceptable for everyday use. On-device inference reduces dependence on network connectivity and keeps raw signals closer to the user, but it is constrained by the memory, processor, and battery limitations of wearable hardware. Edge computing architectures introduce intermediate gateways, such as smartphones or home hubs, that can perform more sophisticated preprocessing and maintain local model copies. A promising structural trade-off is the use of hierarchical inference, in which lightweight models on the device detect obvious signal quality problems and only transmit uncertain or clinically suspicious segments for further analysis. This design reduces energy consumption and communication overhead while preserving diagnostic sensitivity. Cloud components can then support computationally intensive models, longitudinal pattern analysis, and integration with electronic health records. The choice of architecture also affects model updating. Continuous wearables generate nonstationary data because user behavior, sensor hardware, and physiological states change over time. A static model deployed without adaptation may drift and lose accuracy. Federated approaches enable model updates across many devices without transferring raw personal data to a central repository [9]. Such architectures align with the broader goal of building learning health systems that improve over time while respecting local data custody. However, federated learning introduces coordination complexity, communication costs, and the need to handle heterogeneous device populations. System designers must also decide how to represent temporal dependencies. Cardiovascular abnormalities may manifest as subtle changes in heart rate variability, morphology, or rhythm over minutes, hours, or days. The computational architecture must therefore support multi-resolution temporal modeling while managing limited onboard memory. These architectural choices are not purely technical; they have direct implications for usability, safety, and regulatory oversight.

4. Data Governance, Privacy, and Ethical Infrastructure

Continuous wearable sensing generates an unprecedented volume of intimate physiological data. These data can reveal not only cardiovascular status but also sleep patterns, stress, physical activity, and potentially other health conditions. The governance framework must address consent, data minimization, secondary use, retention, and patient access. A narrowly technical approach to privacy is insufficient because re-identification risks extend beyond raw waveforms to include derived features, activity signatures, and linkage with location or demographic data. Privacy-preserving architectures such as federated learning, secure aggregation, and differential privacy offer partial protections, but they also create trade-offs with model transparency and auditability [9]. Algorithmic bias represents another governance concern. Models trained on unrepresentative populations may systematically underdetect

abnormalities in groups that differ in age, sex, skin tone, body mass, or comorbid conditions. Evidence from other health algorithms demonstrates that bias can arise even when sensitive variables are excluded, because proxy variables and unequal data quality can reproduce structural disadvantage [10]. Responsible machine learning frameworks therefore emphasize the need for prospective monitoring, clear model ownership, and defined accountability mechanisms [11]. In the context of continuous cardiovascular monitoring, bias may appear not only in diagnostic accuracy but also in alert frequency, sensor performance, and patient adherence. For example, a wrist-worn photoplethysmographic sensor may produce lower signal quality for individuals with darker skin pigmentation, leading to differential rejection of data and unequal benefit. Governance processes must formalize how these disparities are measured and mitigated. Key challenges for clinical impact include establishing robust validation across sites, defining meaningful endpoints, and ensuring that AI outputs are interpretable to clinicians [12]. The ethical infrastructure must also address the psychological burden of false positives and the potential for overdiagnosis. Early detection systems can create anxiety, increase healthcare utilization, and alter self-perception even when abnormalities are not clinically significant. Data governance should therefore extend beyond privacy and security to include the responsible communication of risk and the patient's right to meaningful explanation.

5. Robustness, Fairness, and Clinical Generalization

Robustness is a central requirement for AI-assisted early detection because failure may occur silently across months of continuous monitoring. A model may achieve excellent performance on retrospective datasets yet degrade when confronted with novel sensor hardware, different skin types, artifact patterns, or physiological conditions. Uncertainty quantification is essential to distinguish confident predictions from ambiguous ones and to route uncertain cases for human review [13]. Without calibrated uncertainty estimates, clinicians cannot determine which alerts warrant immediate action. Explainability further supports clinical trust by contextualizing model outputs in terms that align with clinical reasoning. However, explanations must be designed for the actual workflows of healthcare providers rather than presented as generic technical visualizations [14]. Robustness also depends on the quality and representativeness of training data. Multi-site data collection, standardized annotation, and external validation are needed to ensure that performance estimates are not inflated by hidden selection biases. Digital health validation frameworks emphasize the need to define the intended use, specify the target population, and test the system under realistic conditions before broad deployment [15]. Interoperability is another structural requirement. Wearable data must be integrated with electronic health records, clinical decision support tools, and follow-up pathways. Without common data models and semantic standards, AI outputs become fragmented alerts that are difficult to act upon [16]. Fairness cannot be addressed solely through data balancing or post hoc calibration. It requires analyzing the entire pipeline, including sensor access, signal quality, enrollment criteria, clinician response, and patient outcomes. A system that detects abnormalities equally across groups in a controlled laboratory study may still produce unequal real-world benefit if follow-up infrastructure is unevenly distributed. Clinical generalization therefore demands recursive evaluation that connects model performance with care delivery and social context. The goal is not merely to maximize average sensitivity but to ensure that early detection systems do not widen existing health disparities.

6. Deployment, Integration, and Sustainability

Deployment of AI-assisted cardiovascular monitoring must account for organizational workflows, professional roles, and the long-term maintenance of technical infrastructure. Many digital health innovations fail to scale despite promising pilot results because they ignore the social and organizational dimensions of implementation. The nonadoption and abandonment of health technologies often stem from misaligned incentives, unclear clinical ownership, excessive workload, and lack of interoperability [17]. Wearable early warning systems introduce a continuous stream of information into clinical environments that are already saturated with alerts and administrative demands. A successful deployment therefore requires careful design of triage protocols, escalation pathways, and feedback loops. Ethical mapping of AI in healthcare further suggests that issues such as responsibility, fairness, privacy, and the clinician-patient relationship must be evaluated alongside technical performance [18]. Sustainability is an often neglected dimension. The computational demands of training and maintaining deep learning models, the energy consumption of continuous sensing, and the lifecycle of wearable hardware all have environmental and economic costs. Sustainable AI frameworks call for explicit attention to how algorithmic systems contribute to resource use and how their benefits can be maintained without continuous growth in computation [19]. Synthetic data generation offers one strategy to reduce the need for large proprietary datasets and to simulate rare cardiovascular events, but synthetic data must be carefully validated to avoid introducing new biases or overestimating generalization [20]. Deployment also requires ongoing monitoring because model performance may drift as populations, device firmware, and clinical guidelines evolve. Continuous post-market surveillance should track alert rates, positive predictive value, user adherence, and downstream outcomes. Without such infrastructure, early detection systems may perform well in a regulatory approval study but gradually become unreliable in routine practice. The sustainability of these systems thus depends on governance arrangements that ensure accountability, funding, and technical maintenance over time.

7. Policy Implications and Future Directions

Policy frameworks for AI-assisted cardiovascular early detection are still evolving. Regulators face the challenge of evaluating software that changes over time, operates in uncontrolled environments, and interacts with consumer devices that were not originally designed for medical use. Wearable-based screening programs have produced evidence that continuous monitoring can detect undiagnosed atrial fibrillation, but they have also raised questions about the clinical significance of detected episodes and the appropriate threshold for intervention [21]. Clear guidance is needed on risk classification, software updates, post-market monitoring, and the division of responsibility between device manufacturers, algorithm developers, and healthcare providers. Reimbursement policies also shape adoption. If continuous monitoring and AI interpretation are not reimbursed as part of chronic disease management, health systems may be reluctant to invest in the necessary infrastructure. Future systems should move toward prospective clinical trials that evaluate not only diagnostic accuracy but also patient outcomes, resource utilization, and equity. Deep learning in healthcare continues to expand its capabilities, yet the translation of these capabilities into sustained clinical benefit requires integrated data platforms, regulatory clarity, and a trained workforce [22]. The next generation of continuous cardiovascular monitoring may combine multiple sensor modalities, including photoplethysmography, electrocardiography, bioimpedance, and cuffless blood pressure estimation. This multimodal integration could improve robustness, but it also increases system complexity and data governance burdens. Policy should encourage modular architectures that allow independent validation of sensor

components, machine learning modules, and clinical decision support functions. It should also promote interoperability standards that prevent vendor lock-in and enable data portability. As early detection systems become more autonomous, questions of liability, consent, and human oversight will require ongoing interdisciplinary deliberation. The ultimate policy goal is to create an environment in which innovation can proceed without undermining patient safety, privacy, or equity.

8. Conclusion

AI-assisted early detection of cardiovascular abnormalities using continuous wearable sensor signals represents a significant but demanding frontier in digital health. The clinical promise of detecting paroxysmal arrhythmias, silent ischemic changes, or early decompensation must be balanced against the realities of noisy sensor data, architectural complexity, and sociotechnical risk. This article has argued that system-level thinking is essential. Sensing infrastructure must be designed for continuous quality management rather than treated as a fixed input. System architectures must balance on-device efficiency, edge processing, cloud inference, and adaptive learning in ways that respect privacy and operational constraints. Governance frameworks must address not only data protection but also fairness, bias, accountability, and the communication of risk. Robustness requires uncertainty quantification, external validation, and interoperability across clinical systems. Deployment and sustainability raise further questions about clinical workflow, resource use, and long-term maintenance. Policy must evolve to provide regulatory clarity, reimbursement pathways, and post-market oversight for adaptive software systems. The central insight is that the most important challenges are not located solely within machine learning models. They are distributed across sensors, networks, institutions, clinicians, patients, and regulatory regimes. Early detection systems will only fulfill their potential if they are designed as accountable, equitable, and sustainable infrastructures rather than as isolated algorithmic products. This perspective calls for sustained interdisciplinary research, prospective evaluation, and governance innovation that can keep pace with the rapid evolution of wearable sensing and artificial intelligence.

References

1. 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.
2. Allen, J. (2007). Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3), R1–R39.
3. Perez, M. V., Mahaffey, K. W., Hedlin, H., Rumsfeld, J. S., Garcia, A., Ferris, T., Balasubramanian, V., Russo, A. M., Rajmane, A., Cheung, L., Hung, G., Lee, J., Kowey, P., Talati, N., Nag, D., Gummidipundi, S. E., Beatty, A., Hills, M. T., Desai, S., ... Turakhia, M. P. (2019). Large-scale assessment of a smartwatch to identify atrial fibrillation. *New England Journal of Medicine*, 381(20), 1909–1917.
4. 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.

5. Zheng, X., Dwyer, V. M., Barrett, L. A., Derakhshani, M., & Hu, S. (2022). Adaptive notch-filtration to effectively recover photoplethysmographic signals during physical activity. *Biomedical Signal Processing and Control*, 72, 103303.
6. Clifford, G. D., Behar, J., Li, Q., & Oeff, M. (2012). Signal quality indices and data fusion for determining acceptability of electrocardiograms collected in noisy ambulatory environments. *Physiological Measurement*, 33(9), 1419–1437.
7. 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., & Kyriacou, P. A. (2022). Wearable photoplethysmography for cardiovascular monitoring. *Proceedings of the IEEE*, 110(3), 355–381.
8. Steinhubl, S. R., Waalen, J., Edwards, A. M., Ariniello, L. M., Mehta, R. R., Ebner, G. S., Carter, C., Baca-Motes, K., Felicione, E., Sarich, T., & Topol, E. J. (2018). Effect of a home-based wearable continuous ECG monitoring patch on detection of undiagnosed atrial fibrillation: The mSToPS randomized clinical trial. *JAMA*, 320(2), 146–155.
9. 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.
10. 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.
11. 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.
12. Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine*, 17, 195.
13. Leibig, C., Allken, V., Ayhan, M. S., Berens, P., & Wahl, S. (2017). Leveraging uncertainty information from deep neural networks for disease detection. *Scientific Reports*, 7, 17816.
14. Tonekaboni, S., Joshi, S., McCradden, M. D., & Goldenberg, A. (2019). What clinicians want: Contextualizing explainable machine learning for clinical end use. *JMIR Medical Informatics*, 7(4), e14359.
15. Mathews, S. C., McShea, M. J., Hanley, C. L., Ravitz, A., Labrique, A. B., Cohen, W. J., & Green, J. (2019). Digital health: A path to validation. *npj Digital Medicine*, 2, 38.
16. Lehne, M., Sass, J., Essenwanger, A., Schepers, J., & Thun, S. (2019). Why digital medicine depends on interoperability. *npj Digital Medicine*, 2, 79.
17. Greenhalgh, T., Wherton, J., Papoutsi, C., Lynch, J., Hughes, G., A'Court, C., Hinder, S., Fahy, N., Procter, R., & Shaw, S. (2017). Beyond adoption: A new framework for

theorizing and evaluating nonadoption, abandonment, and challenges to scale-up of health and care technologies. *Journal of Medical Internet Research*, 19(11), e367.

18. Morley, J., Machado, C. C. V., Burr, C., Cows, J., Joshi, I., Taddeo, M., & Floridi, L. (2020). The ethics of AI in health care: A mapping review. *Social Science & Medicine*, 260, 113172.
19. van Wynsberghe, A. (2021). Sustainable AI: AI for sustainability and the sustainability of AI. *AI and Ethics*, 1(3), 213–218.
20. Chen, R. J., Lu, M. Y., Chen, T. Y., Williamson, D. F. K., & Mahmood, F. (2021). Synthetic data in machine learning for medicine and healthcare. *Nature Biomedical Engineering*, 5(6), 493–497.
21. Dörr, M., Nohturfft, V., Brasier, N., Bosshard, E., Djurdjevic, A., Gross, S., Raichle, C. J., Rhinisperger, M., Stöckli, R., & Steck, H. (2019). The WATCH AF trial: SmartWATCHes for detection of atrial fibrillation. *JACC: Clinical Electrophysiology*, 5(2), 199–208.
22. Miotto, R., Wang, F., Wang, S., Jiang, X., & Dudley, J. T. (2018). Deep learning for healthcare: Review, opportunities and challenges. *Briefings in Bioinformatics*, 19(6), 1236–1246.