

AI-Driven Early Warning of Cardiovascular Abnormalities Using Continuous Wearable Biosensing

Cody Lindberg

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
codyl@colostate.edu

Viktor L. Mills

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

Abstract

Continuous wearable biosensing is shifting cardiovascular care from episodic hospital measurement toward longitudinal, ambient data collection, creating new possibilities for early detection of arrhythmias, ischemic patterns, and hemodynamic instability. However, the translation of this sensing capability into reliable early warning systems requires more than improved classification accuracy. It demands careful integration of sensing hardware, embedded computation, network infrastructure, learning pipelines, clinical workflows, and regulatory oversight. This paper presents a system-level analysis of artificial intelligence driven early warning architectures for cardiovascular abnormalities using wearable biosensors. It examines trade-offs between on-device and cloud inference, constraints on energy and data quality, the influence of motion artifacts and heterogeneous sensor modalities, and the organizational arrangements required for safe deployment. The discussion addresses robustness under dynamic ambulatory conditions, fairness across diverse populations, privacy-preserving learning approaches such as federated systems, and the governance mechanisms needed to maintain clinical trust. Rather than focusing on a single algorithm or sensor, the paper argues that sustainable early warning depends on a layered socio-technical infrastructure in which algorithmic performance is coupled with data stewardship, infrastructure resilience, and accountability. The analysis draws on current evidence from wearable photoplethysmography and electrocardiography studies, machine learning deployment research, and health data governance scholarship to articulate structural design principles for future cardiovascular biosensing systems.

Keywords

cardiovascular early warning, wearable biosensing, artificial intelligence, digital health governance, federated learning, robustness, health equity.

1. Introduction

Wearable sensors for remote health monitoring have expanded from simple activity trackers into multimodal physiological instruments capable of collecting photoplethysmographic, electrocardiographic, inertial, temperature, and bioimpedance signals over days or weeks [1]. Among these modalities, photoplethysmography has become particularly important because it enables opto-physiological monitoring through low-cost optical components, allowing continuous pulse waveform acquisition without clinical lead placement [2]. The clinical

appeal of such sensing lies in its potential to identify intermittent cardiovascular abnormalities that may not appear during short clinical examinations. Atrial fibrillation, paroxysmal arrhythmias, and early signs of autonomic or hemodynamic decompensation can be episodic and therefore invisible to conventional diagnostic snapshots. Wearable devices that maintain continuous contact with the body create a longitudinal observation window that is structurally different from hospital monitoring because it follows patients through sleep, exercise, work, and daily stressors.

Substantial evidence indicates that wrist-worn optical sensors can support population-scale arrhythmia screening. Large investigations using commercial smartwatches demonstrated that photoplethysmography combined with notification algorithms can identify atrial fibrillation in a substantial proportion of participants, although confirmatory electrocardiography remains necessary [3,4]. These studies show the public health potential of continuous screening, but they also reveal system-level challenges related to false positives, user adherence, data completeness, and the burden of downstream clinical confirmation. The early warning problem is therefore not solely a signal processing task. It includes how alerts are generated, transmitted, visualized, interpreted, and acted upon by patients and clinicians. A robust system must balance sensitivity against alarm fatigue, preserve privacy, function across heterogeneous devices, and remain trustworthy under real-world motion and environmental variability.

This paper approaches artificial intelligence driven early warning as an infrastructure problem rather than an isolated algorithmic challenge. It analyzes the structural trade-offs that arise across sensing, learning, deployment, governance, and long-term sustainability. The argument is that durable cardiovascular early warning requires a layered architecture in which embedded inference, cloud services, clinical data systems, and regulatory mechanisms operate together under conditions of uncertainty. The following sections develop this argument through detailed examination of sensing modalities, learning architectures, edge-cloud coordination, robustness, fairness, regulation, and policy.

2. System Architecture and the Continuous Sensing Continuum

A cardiovascular early warning system based on wearable biosensing typically consists of multiple device-level sensors, local processing firmware, wireless communication links, a cloud or edge analytics layer, and a clinical interface. The sensing layer acquires photoplethysmographic and electrocardiographic waveforms together with accelerometer and gyroscope signals that provide contextual information about motion and posture. These signals differ substantially in their information density and noise sensitivity. Electrocardiographic recordings contain rich morphological information about atrial and ventricular electrical activity, but continuous electrode-based acquisition is less convenient for everyday use. Optical pulse signals are easier to collect continuously, yet they are affected by skin pigmentation, tissue perfusion, probe pressure, ambient light, and physical movement. Deep neural networks applied to ambulatory electrocardiograms have shown cardiologist-level arrhythmia detection [5], but the generalization of such models to low-quality wearable traces remains a major systems challenge.

The broader machine learning literature in health care has emphasized that model performance is highly dependent on the data pipeline, labeling quality, and the clinical context in which predictions are used [6]. For wearable cardiovascular monitoring, this insight translates into a multi-level architecture in which raw signals are cleaned, segmented, and transformed before inference. The placement of these processing functions determines latency,

energy consumption, privacy exposure, and resilience. A device that performs only raw data streaming simplifies hardware but imposes high wireless bandwidth and battery costs. A device that performs local feature extraction and anomaly detection can reduce communication load and enhance privacy, but it may lack the contextual knowledge available in the cloud. High-performance medicine increasingly depends on harmonizing these distributed computational roles rather than concentrating intelligence in a single location [7].

The concept of a continuous sensing continuum is useful for understanding this architecture. At one end are ultra-low-power wearable nodes that monitor basic rhythm regularity and motion context. At the other end are cloud-based clinical decision support systems that integrate longitudinal trend data, electronic health records, and population models. Between these endpoints lie edge gateways such as smartphones or home hubs that can aggregate multiple devices, perform medium-complexity inference, and manage connectivity. This continuum requires clear contracts for data representation, model versioning, alert thresholds, and fallback behavior when connectivity degrades. Without such contracts, the early warning pipeline becomes fragile because devices, gateways, and cloud services may evolve independently and produce inconsistent interpretations of the same physiological event.

A further architectural consideration is the temporal nature of cardiovascular risk. Many abnormalities develop over minutes, hours, or days rather than in a single heartbeat. Early warning therefore requires not only beat-level classification but also longitudinal pattern recognition across circadian rhythms, medication schedules, physical activity, and sleep. This shifts the system design toward stateful monitoring rather than stateless event detection. The learning infrastructure must maintain a patient-specific baseline and distinguish stable variation from clinically meaningful drift. Such dynamic baseline modeling introduces additional trade-offs because personalized models may overfit to short-term noise, while population-level models may miss rare individual patterns.

3. Learning Architectures and Clinical Inference Trade-Offs

Modern machine learning for cardiovascular biosensing relies on deep sequence models that learn temporal dependencies from raw or preprocessed waveforms. These architectures are attractive because they can automatically extract features that are difficult to specify manually, such as subtle morphological changes preceding atrial fibrillation or decompensation. However, their deployment in wearable systems raises important structural questions about model capacity, interpretability, and computational cost. A high-capacity model may yield stronger retrospective performance but may be too large for embedded hardware and too sensitive to sensor noise. A compact model may be deployable on a wristband but may miss complex arrhythmias. The system designer must therefore evaluate the point at which increased model complexity ceases to improve clinical utility and begins to undermine reliability.

Digital twin approaches have been proposed as a way to integrate computational models of cardiac function with patient-specific data, enabling simulation of interventions and disease progression [8]. In a digital twin paradigm, the wearable early warning model is not an isolated classifier but part of a continuously updated representation of the patient's cardiovascular state. This representation can support richer inference than simple event detection, including predictions about response to medication, fluid status, or activity tolerance. The implementation cost, however, is substantial because digital twins require consistent data standards, validated physiological simulators, and governance around model updates.

Federated learning offers an alternative to centralized model training by allowing models to learn from distributed devices without transferring raw physiological data to a central repository [9]. This approach has structural advantages in preserving privacy and reducing data transfer, but it introduces challenges in coordinating model updates across devices with different hardware capabilities, data distributions, and user behaviors. Federated systems also require robust aggregation methods to prevent poor local updates from degrading the global model and must address the possibility of malicious or inadvertently biased contributions. Privacy concerns remain prominent in medical big data, particularly because wearable signals can reveal sensitive behavioral and health information beyond the intended cardiovascular target [10]. Consequently, privacy preservation cannot be treated as a purely technical encryption problem; it also involves data minimization, consent, and institutional controls.

Fairness is another central concern in learning architectures. Algorithmic bias in health care has been shown to produce systematically worse outcomes for certain populations when models rely on proxy variables that reflect structural inequities [11]. Ensuring fairness requires active evaluation across demographic groups, device types, skin tones, and clinical subgroups [12]. The learning pipeline must include pre-deployment fairness audits and post-deployment monitoring because distributions may shift once a system is used in new environments. Dataset shift is particularly problematic for wearable biosensing because sensor models, firmware updates, seasonal activity patterns, and population differences can change the relationship between inputs and labels [13]. A model that performs well in a controlled study may degrade when deployed at scale without mechanisms for recalibration.

Recent work on joint attention mechanism learning for opto-physiological monitoring during physical activity illustrates how advanced learning strategies can improve robustness under motion and signal corruption [14]. Attention mechanisms allow models to focus on the most informative segments of a noisy waveform and to suppress motion-related distortions. Such techniques are valuable because physical activity is a defining feature of ambulatory monitoring rather than an exceptional condition. Cardiovascular early warning must function while users walk, climb stairs, exercise, and sleep. The integration of attention mechanisms with multimodal sensing can therefore reduce the trade-off between continuous wearability and diagnostic signal quality.

4. Edge-Cloud Coordination and Infrastructure Deployment

The deployment of cardiovascular early warning models requires careful division of labor between wearable devices, edge gateways, and cloud services. Edge computing brings machine learning closer to data sources, reducing latency and dependence on network connectivity [15]. In cardiovascular monitoring, low-latency local inference can be important for detecting dangerous bradycardia, tachycardia, or extreme rhythm irregularities that require immediate user feedback. Energy-efficient context recognition remains a persistent constraint for wearable health systems because continuous sensing and inference drain batteries and reduce user adherence [16]. System designers must balance sampling frequency, signal resolution, model complexity, and transmission intervals against the practical limits of battery chemistry and device thermal constraints.

Deployment also involves managing software updates and model drift over time. Machine learning systems develop hidden technical debt when data dependencies, model parameters, and serving infrastructure become entangled without clear abstraction boundaries [17]. In a cardiovascular early warning system, a model update that improves detection sensitivity for one arrhythmia may inadvertently alter alert thresholds for another condition or interact

unpredictably with a new device firmware. This interdependence requires disciplined model versioning, staged rollout, shadow mode evaluation, and automated monitoring of false positive rates. Without these practices, continuous wearable systems may become less reliable as they scale, even when individual models improve.

Clinical deployment must also account for regulatory expectations. Artificial intelligence based medical devices are increasingly subject to frameworks that require validation, risk classification, and post-market surveillance [18]. A wearable early warning system that provides diagnostic suggestions may be regulated differently from a wellness application that only reports heart rate trends. The boundary between screening and diagnosis is not always stable, and system design choices such as the language used in notifications can influence regulatory obligations. Early warning systems should therefore be developed with a regulatory pathway in mind, including clinical validation, software as a medical device considerations, and adverse event reporting.

5. Data Governance, Fairness, and Accountability

Governance for continuous cardiovascular biosensing extends beyond data security to include questions of ownership, consent, algorithmic accountability, and clinical responsibility. Wearable devices generate high-frequency personal health data that may reveal information about sleep, stress, medication adherence, and physical capacity. The use of these data for early warning creates benefits but also risks if inferences are shared with employers, insurers, or commercial platforms. Privacy-preserving architectures such as federated learning reduce some risks, but they do not eliminate the need for transparent data governance and meaningful consent mechanisms [10].

Fairness and equity are structural concerns that must be addressed throughout the system life cycle. Differences in device access, smartphone ownership, health literacy, and connectivity can create disparities in who benefits from early warning. Digital health equity requires attention to affordability, language, usability, and the availability of clinical follow-up for underserved populations [19]. If a wearable system detects a potential arrhythmia in a user who cannot access timely confirmatory care, the early warning may produce anxiety without clinical benefit. System designers should therefore consider the full care pathway, not only the detection algorithm.

Accountability is especially important when machine learning systems operate continuously and autonomously. Clear allocation of responsibility among device manufacturers, algorithm developers, health care providers, and patients is required to manage harm when predictions are incorrect. Ethical frameworks for machine learning in medicine emphasize the importance of transparency, fairness, and human oversight [20]. In practice, these principles require that advisory outputs be interpretable enough for clinicians to contest or validate, that models be evaluated for subgroup performance, and that patients receive understandable information about the meaning and limits of alerts. Without such mechanisms, early warning systems may erode trust even when their aggregate statistical performance is strong.

6. Robustness, Sustainability, and Long-Term Operational Resilience

Robustness in continuous wearable biosensing is a multidimensional property that includes signal quality, algorithmic stability, infrastructure resilience, and clinical safety under heterogeneous conditions. Motion artifacts, sensor displacement, skin tone variation, and ambient temperature can alter signal distributions and reduce inference reliability. Attention-based multimodal learning is one approach to improving robustness by teaching models to

ignore irrelevant signal components and emphasize informative waveform regions [14]. However, robustness also depends on the data acquisition strategy. Devices that adapt their sampling rate, adjust optical gain, or request an electrocardiographic confirmation when optical signal quality is poor can maintain performance in real-world conditions more effectively than devices that rely on fixed sensor settings.

Sustainability involves the long-term maintenance of hardware, software, data infrastructure, and human expertise. Continuous monitoring systems require ongoing model monitoring, sensor calibration, cybersecurity updates, and clinical workflow adaptation. The failure to budget for these activities can lead to performance degradation over time, even if the initial deployment is successful. Hidden technical debt accumulates when predictions depend on unstable data sources or when model updates cannot be traced to specific evaluation datasets [17]. Sustainable deployment therefore requires institutional commitment to longitudinal quality assurance rather than one-time validation.

Infrastructure resilience is also critical because early warning systems may be needed during high-stress events such as pandemics, natural disasters, or health system surges. Cloud dependency can become a vulnerability when connectivity is interrupted or when server capacity is overwhelmed. Edge-first architectures that preserve essential alerting functions locally can provide graceful degradation under adverse conditions. At the same time, local-only inference may reduce access to longitudinal trend analysis and population-level context. A resilient system should therefore be designed with redundant communication paths, local data buffering, and clear priorities for which alerts must be generated immediately and which can be delayed.

7. Policy Implications and Future Directions

Policymakers face the challenge of encouraging innovation in wearable cardiovascular monitoring while ensuring safety, privacy, and equitable access. Regulatory frameworks must become more adaptive to continuous learning systems that evolve after initial approval. This may require new forms of post-market surveillance, real-world performance reporting, and clear rules for model updates. At the same time, reimbursement policies must address the cost of wearable monitoring, confirmatory testing, and clinician review. Without sustainable payment models, early warning systems may remain confined to affluent populations or research studies rather than becoming part of routine cardiovascular care.

Future research should examine the interaction between personalized baselines and population-level early warning models, the role of multimodal fusion in reducing false alarms, and the use of federated evaluation to measure fairness without centralizing sensitive data. There is also a need for studies that follow patients over long periods to assess whether continuous warning systems improve meaningful outcomes such as stroke prevention, hospitalization reduction, or quality of life. The field would benefit from open benchmarks that include diverse sensor types, motion conditions, skin tones, and clinical endpoints, because current evaluations often rely on limited datasets that do not reflect deployment variability.

8. Conclusion

AI-driven early warning of cardiovascular abnormalities using continuous wearable biosensing represents a transformative opportunity to detect disease earlier and to shift care toward prevention. However, the successful realization of this vision depends on systemic design that integrates sensing hardware, embedded inference, cloud infrastructure, learning

architectures, clinical workflows, and governance mechanisms. Signal quality under physical activity, energy constraints, model drift, fairness across populations, and privacy protection are not secondary concerns but central determinants of whether such systems can be trusted and sustained. The analysis in this paper has emphasized that reliable early warning is not achieved solely by improving a single classifier. It requires a coherent socio-technical infrastructure in which models are continuously evaluated, devices operate robustly under real-world conditions, patients understand alerts, and clinicians can act on them. Future progress will depend on interdisciplinary collaboration among engineers, clinicians, regulators, and patients to align algorithmic innovation with health system capacity and equity.

References

1. Majumder, S., Mondal, T., & Deen, M. J. (2017). Wearable sensors for remote health monitoring. *Sensors*, 17(1), 130.
2. Allen, J. (2007). Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3), R1–R39.
3. Tison, G. H., Sanchez, J. M., Ballinger, B., Singh, A., Olgin, J. E., Pletcher, M. J., Vittinghoff, E., Lee, E. S., Fan, S. M., Gladstone, R. A., Mikell, C., Sohoni, N., Hsieh, J., & Marcus, G. M. (2018). Passive detection of atrial fibrillation using a commercially available smartwatch. *JAMA Cardiology*, 3(5), 409–416.
4. 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.
5. Hannun, A. Y., Rajpurkar, P., Haghpanahi, M., Tison, G. H., Bourn, C., Turakhia, M. P., & Ng, A. Y. (2019). Cardiologist-level arrhythmia detection with ambulatory ECG. *Nature Medicine*, 25(1), 65–69.
6. 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.
7. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56.
8. Corral-Acero, J., Margara, F., Marciniak, M., Rodero, C., Loncaric, F., Feng, Y., Gilbert, A., Fernandes, J. F., Bukhari, H. A., Wajdan, A., Martinez, M. V., Santos, M. S., Shamohammdi, M., Luo, H., Westphal, P., Leeson, P., DiAchille, P., Gurev, V., Mayr, M., ... Lamata, P. (2020). The digital twin to enable the vision of precision cardiology. *European Heart Journal*, 41(48), 4556–4564.
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. Price, W. N., & Cohen, I. G. (2019). Privacy in the age of medical big data. *Nature Medicine*, 25(1), 37–43.

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.
12. Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2018). Ensuring fairness in machine learning to advance health equity. *Annals of Internal Medicine*, 169(12), 866–872.
13. Subbaswamy, A., & Saria, S. (2020). From development to deployment: Dataset shift, causality, and shift-stable models in health AI. *Biostatistics*, 21(2), 345–352.
14. Zheng, X., Hu, S., Dwyer, V., Barrett, L., & Derakhshani, M. (2026). Joint attention mechanism learning to facilitate opto-physiological monitoring during physical activity. *Biomedical Signal Processing and Control*, 113, 108949.
15. Murshed, M. G. S., Murphy, C., Hou, D., Khan, N., Ananthanarayanan, G., & Hussain, F. (2021). Machine learning at the network edge: A survey. *ACM Computing Surveys*, 54(8), 1–37.
16. Rault, T., Bouabdallah, A., Challal, Y., & Marin, F. (2017). A survey of energy-efficient context recognition systems using wearable sensors for healthcare. *Personal and Ubiquitous Computing*, 21(4), 607–619.
17. Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J. F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. In *Advances in neural information processing systems* 28 (pp. 2503–2511).
18. Benjamins, S., Dhunnoo, P., & Meskó, B. (2020). The state of artificial intelligence-based FDA-approved medical devices and algorithms: An online database. *NPJ Digital Medicine*, 3, 118.
19. Lyles, C. R., Wachter, R. M., & Sarkar, U. (2021). Focusing on digital health equity. *JAMA*, 326(18), 1795–1796.
20. Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. *PLoS Medicine*, 15(11), e1002689.