

Explainable Artificial Intelligence for Activity-Aware Cardiovascular State Assessment Using Wearable Sensors

Jordan L. Garrett

Department of Computer Science, University of Houston, Houston, TX, USA.

jordangarrett176@uh.edu

Albert Gregory

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.

albertwork@buffalo.edu

Blake Miles

Department of Computer Science, University of North Texas, Denton, TX, USA.

contactblake@unt.edu

Abstract

Wearable sensors now support continuous cardiovascular monitoring outside traditional clinical settings, yet physical activity creates substantial complexity for interpreting cardiac signals. Motion artifact, autonomic adjustments, and sensor displacement can mimic or mask pathological changes, making activity-aware cardiovascular state assessment a difficult systems problem. This paper examines how explainable artificial intelligence can support robust and accountable cardiovascular inference in wearable environments. It argues that explainability must move beyond local feature attribution and become a structural property of the sensing, inference, deployment, and governance layers. The paper reviews advances in clinical cardiovascular deep learning, wearable activity recognition, and explainable artificial intelligence, then develops a system-level account of architectural trade-offs. Key issues include edge-cloud partitioning, multimodal fusion, activity context conditioning, uncertainty communication, data representativeness, fairness across demographic and physiological groups, and regulatory accountability. The discussion emphasizes that activity-aware explanation should clarify not only which physiological features influenced a prediction but also how motion context, signal quality, and population-specific factors shaped the inference. The paper further analyzes deployment infrastructure, sustainability, alert fatigue, privacy, and policy integration. It concludes that clinical value depends on designing systems that remain interpretable, robust under real-world activity, equitable across populations, and governable throughout their operational lifecycle.

Keywords

explainable artificial intelligence, wearable sensors, cardiovascular state assessment, activity-aware inference, health data governance, deployment sustainability, fairness in clinical machine learning.

1. Introduction

Cardiovascular disease remains a leading cause of morbidity and mortality, and its management increasingly depends on the ability to detect subtle physiological changes before

acute events occur. Wearable sensors have shifted cardiovascular monitoring from episodic clinical measurements toward continuous ambulatory observation, thereby creating new possibilities for early detection, remote triage, and individualized intervention [1]. However, translating wearable cardiovascular monitoring into routine care requires attention to sensor reliability, data governance, user adherence, clinical workflow, and integration with existing health information systems [2]. Wearable devices capture not only cardiac signals such as photoplethysmography and electrocardiography but also accelerometry, gyroscopic motion, and environmental cues. Public benchmark datasets and evaluation protocols have supported the development of activity monitoring algorithms, enabling researchers to compare sensing configurations and recognition models across standardized conditions [3]. Deep convolutional and recurrent architectures further demonstrated that multimodal wearable data can be transformed into robust representations of physical motion, which is fundamental for separating physiological change from movement-induced artifact [4].

Activity-aware cardiovascular state assessment is a systems problem rather than a single-model classification task. Physical activity alters heart rate, autonomic tone, peripheral circulation, and respiration in ways that overlap with pathological changes. An explainable artificial intelligence framework is therefore necessary to clarify whether a model is responding to true cardiovascular risk or to contextual activity patterns. The broader field of explainable artificial intelligence has produced taxonomies and evaluation strategies that distinguish post hoc attribution from inherently transparent modeling [5]. There is also a vigorous debate about whether high-stakes medical decisions should rely on interpretable models rather than complex black-box predictors with post hoc explanations [6]. This paper examines the structural design of activity-aware cardiovascular assessment systems, with particular attention to architectural trade-offs, explainability requirements, robustness under motion, fairness across demographic groups, deployment infrastructure, sustainability, and policy implications. The discussion emphasizes system-level reasoning rather than algorithmic detail.

2. Related Work and System-Level Challenges

Clinical deep learning has already shown that cardiovascular state can be inferred from electrocardiographic data with high accuracy. Deep neural networks have achieved cardiologist-level arrhythmia detection in ambulatory electrocardiograms, indicating that automated interpretation can operate outside strictly controlled clinical settings [7]. Related work demonstrated that an artificial intelligence-enabled electrocardiogram can detect left ventricular dysfunction even when obvious clinical signs are absent [8]. These successes motivate wearable extensions, but they also reveal a central gap. Hospital-grade electrocardiographic signals are relatively clean, whereas wearable signals are contaminated by motion, variable electrode contact, and changing tissue optical properties. Reviews of wearable sensor data mining in health monitoring highlight recurrent challenges such as noise, missing data, heterogeneous device formats, and the need for context-aware preprocessing [9]. Large clinical data resources such as the MIMIC-III critical care database have accelerated model development and reproducibility, yet they are not direct substitutes for ambulatory wearable data collected during everyday life [10].

Explainability in medicine requires more than model introspection. It must connect algorithmic outputs to concepts that are meaningful for clinicians and patients. The notion of causability emphasizes that an explanation must be useful for a human decision process, not merely descriptive of model internals [11]. Visual attribution techniques such as gradient-

weighted class activation mapping have been used to identify relevant input regions in cardiovascular imaging and signal analysis [12]. Axiomatic attribution methods provide a formal basis for assigning importance to input features, although their assumptions require careful validation in noisy physiological data [13]. Local model-agnostic explanation frameworks offer a complementary route for probing predictions, particularly when heterogeneous clinical and sensor features are combined [14]. A multidisciplinary perspective on explainability in healthcare stresses that organizational, legal, and human factors influence whether an explanation is sufficient for safe deployment [15]. In the specific context of wearable monitoring during physical activity, joint attention mechanism learning has been proposed to facilitate opto-physiological monitoring, allowing the model to attend jointly to motion context and physiological signal features [16]. This line of work is important because activity-aware inference cannot be achieved simply by removing motion artifact. The model must actively condition physiological interpretation on the type, intensity, and timing of physical activity.

3. Architectural Considerations for Wearable Cardiovascular Assessment

A system architecture for activity-aware cardiovascular state assessment should include at least five functional components: multimodal sensing, signal quality and activity context processing, physiological state estimation, explanation generation, and clinical decision communication. The sensing layer must reconcile different sampling rates, gain settings, and wearable placements. The activity context layer should estimate whether the user is resting, walking, climbing stairs, cycling, sleeping, or performing unstructured movement. This context is then used by the cardiovascular state estimator to normalize physiological features or to select an appropriate inference pathway. The explanation layer must generate outputs that are meaningful for different stakeholders, including clinicians, patients, and remote monitoring teams. A failure in any layer can produce misleading cardiovascular state estimates even when the underlying predictive model is statistically accurate.

The central structural trade-off is between edge computing and cloud-based inference. Edge deployment reduces latency and limits the transmission of sensitive physiological data, but it constrains model complexity and explanation fidelity because battery, memory, and thermal budgets are limited. Cloud deployment enables larger multimodal models, continuous retraining, and richer explanation artifacts, but it introduces connectivity dependencies, privacy risks, and higher infrastructure costs. Hybrid architectures are often most appropriate for wearable cardiovascular assessment. Signal quality assessment and activity recognition can operate on-device, while high-resolution physiological feature extraction and explainable state estimation can be performed in a secure cloud environment when connectivity is available. This split requires careful data versioning, device metadata management, and fallback behavior during offline periods.

Data infrastructure further shapes system behavior. Wearable datasets are often dominated by younger, healthier, higher-income, and technologically literate populations, which can bias physiological models [17]. If the training data do not adequately represent skin tones, body compositions, cardiovascular conditions, medications, or activity profiles, the resulting model may perform unevenly across groups. The wearable medicine literature has emphasized that technical accuracy is insufficient without attention to user engagement, reimbursement, interoperability, and evidence generation [18]. Regulatory frameworks are evolving to address artificial intelligence as a medical device, raising questions about how adaptive algorithms can be validated when model parameters change after deployment [19]. Ethical guidance

further argues that machine learning in medicine must preserve patient autonomy, fairness, and transparency throughout the data lifecycle [20]. These considerations should be embedded into architecture rather than appended after a model is trained.

4. Explainability and Activity-Aware Interpretation

Explainability in activity-aware cardiovascular assessment must address three distinct questions: why a cardiovascular state estimate was produced, how much physical activity influenced the estimate, and what would change if the activity context were different. This is more demanding than static explanation because the model operates on time-varying inputs with complex dependencies. A useful explanation should therefore distinguish motion-induced changes in heart rate from pathological changes, identify periods of poor signal quality, and describe whether the system is confident or uncertain about its output. Confidence information must be activity-aware. The same heart rate estimate may be highly reliable during stationary rest but less reliable during high-intensity motion.

Several explainability mechanisms can be combined in a layered fashion. Attribution methods that propagate relevance to raw signal time points can reveal whether the model is focusing on the QRS complex, motion artifact, or irrelevant sensor noise. Local surrogate models can summarize the contribution of activity features, demographic variables, and physiological features for a specific prediction. Concept-based explanations can map latent representations to clinically meaningful categories such as exercise intolerance, abnormal recovery, atrial ectopy, or poor peripheral perfusion. In all cases, explanation must align with the cognitive workflow of the clinical user [21]. Trustworthy artificial intelligence policy proposals emphasize that high-risk systems should provide appropriate transparency, human oversight, and technical robustness [22]. Machine learning systems in clinical medicine require not only predictive accuracy but also careful attention to workflow integration, alert fatigue, and clinician trust [23].

The trade-off between interpretable models and post hoc explanation remains especially relevant. Interpretable models such as regularized linear models, decision trees, or risk scores may be sufficient when the relationship between activity-adjusted physiological features and cardiovascular state is well understood. However, wearable signals are high-dimensional and nonstationary, so more expressive models are often needed. The challenge is not to avoid complex models altogether but to construct an explanation layer that can interrogate them under realistic deployment constraints. Rather than treating explainable artificial intelligence as a set of post hoc tools, the system should be designed so that activity context, signal quality, and cardiovascular state are represented as explicit variables in the reasoning pipeline. For example, a cardiovascular state estimator may receive an activity category from a wrist-worn device and a signal quality index from the photoplethysmography channel. The explanation module can then report how each element contributed to the final estimate, thereby allowing a clinician to distinguish a true physiological alert from an artifact-driven output. Activity-aware explanation should also support counterfactual reasoning. A clinician may need to know whether an abnormal cardiovascular estimate would have arisen if the patient had been resting rather than walking. This requires the system to maintain alternative activity-conditioned inference pathways or to estimate the sensitivity of the predicted state to activity variables. Although counterfactual explanation is computationally demanding, it is valuable in wearable monitoring because physical activity is a controllable and commonly reported context.

Uncertainty communication is another essential component. Wearable systems often exhibit higher uncertainty during motion because sensor contact degrades and tissue perfusion varies with movement. A system that reports only a single cardiovascular state without uncertainty can create false confidence. The explanation layer should therefore combine model uncertainty, data quality, and activity intensity into a single contextual confidence statement. For instance, a prediction of elevated heart rate during a period of vigorous walking should be accompanied by a note that the estimate is based on relatively clean photoplethysmography data but that activity intensity may still amplify heart rate independently of any pathological process. This type of explanation is not simply a technical refinement; it is central to clinician acceptance and patient safety. Clinical users are more likely to trust a system that acknowledges its own limitations and provides reasons for why a particular patient context may reduce reliability [21]. Moreover, explanations must be dynamically generated. A static summary of the entire monitoring session is less useful than temporally localized explanations that highlight only the segments where cardiovascular state changes or where activity episodes begin and end. These localized explanations allow clinicians to inspect transitions, such as recovery after exercise, which have diagnostic value in cardiology.

5. Deployment Infrastructure and Operational Sustainability

Operational deployment of activity-aware cardiovascular assessment requires a robust infrastructure that extends beyond the model itself. Device provisioning, secure data transport, quality assurance, firmware updates, and interoperability with electronic health records are all part of the system. Wearable cardiovascular monitoring has generated considerable enthusiasm, but clinical adoption depends on the ability to integrate these tools into existing care pathways without overwhelming clinicians with alerts [24]. A sustainable system should differentiate between urgent cardiovascular alerts, non-urgent physiological abnormalities, and artifact-driven anomalies. Alert fatigue is a well-documented hazard in remote monitoring, and an explainable system can mitigate this by providing context-rich reasons for alerts, allowing care teams to triage more effectively. For example, an alert that includes activity context, signal quality, and a concise explanation may be actionable, whereas an alert without such context may be ignored or create unnecessary anxiety [18].

Edge-cloud partitioning strongly affects sustainability. Battery life, data transmission volume, and storage costs are not neutral parameters; they determine whether patients will continue wearing the device and whether health systems can afford to operate the service. On-device activity recognition can reduce bandwidth and energy consumption by triggering high-resolution physiological analysis only when clinically relevant activity or cardiac events are detected. Cloud resources can then be reserved for model retraining, longitudinal summarization, and clinician-facing explanation displays. However, edge deployment also introduces challenges in maintaining software consistency across heterogeneous device platforms. A version mismatch between the activity recognition model on the device and the physiological state estimator in the cloud can produce contradictory context labels and undermine explanation quality. Governance processes must therefore include synchronized model versioning, data schema validation, and audit logs that record which versions produced a given cardiovascular state estimate.

Data governance further constrains deployment. Wearable cardiovascular data are highly sensitive because they reveal not only health status but also patterns of movement, sleep, and daily activity. Reidentification risks remain significant even when obvious identifiers are removed [20]. Systems should adopt privacy-preserving practices such as federated learning

where feasible, although federated architectures can complicate the generation of global explanations. A practical compromise is to keep raw sensor data on the device or in a trusted local hub while transmitting only aggregated physiological features and activity contexts to the cloud. Explanation generation can then operate on high-level features without exposing raw photoplethysmography or electrocardiography waveforms. This approach reduces privacy risks but may limit the depth of post hoc attribution if raw signal relevance is needed. The system architect must therefore balance privacy against explanatory fidelity, a trade-off that should be made explicit in system documentation and regulatory submissions.

6. Fairness and Population-Specific Considerations

Algorithmic fairness in wearable cardiovascular assessment cannot be reduced to statistical parity across demographic groups. The relationship between activity, physiological signals, and cardiovascular risk is heterogeneous across age, sex, body composition, medication use, and comorbid conditions. A model trained predominantly on healthy young adults may learn activity normalization rules that are inappropriate for older adults with reduced chronotropic responses or for patients on beta-blockers. The resulting system may systematically overestimate or underestimate cardiovascular risk in these groups. This is not merely a data collection problem; it is a structural issue in how activity context is modeled. If activity categories are derived primarily from sensor patterns learned from young populations, the system may misclassify the movement of older adults or people with gait abnormalities, thereby distorting the activity-aware cardiovascular inference [17].

Wearable data additionally contain historical and social biases. Access to wearable devices is uneven, and the populations most likely to contribute training data are often not representative of those with the highest cardiovascular morbidity [18]. Skin tone can affect photoplethysmography signal quality, and movement patterns can vary by occupation, environment, and mobility limitations. If these factors are not explicitly measured and included in the explanation layer, the system may produce outputs that appear objective but hide unjustified disparities. Fairness in this context should therefore be understood as a property of the entire sociotechnical system rather than a single fairness metric applied to model outputs. It requires ongoing monitoring of subgroup performance, contextual accuracy, and the distributional consequences of clinical decisions informed by the system [20]. Explanations should be sensitive to population variation by reporting whether the current patient is well represented in the training data and whether the prediction is affected by known signal quality biases.

The integration of large-scale machine learning into health care has demonstrated that well-performing algorithms can still perpetuate or exacerbate existing disparities when deployed without population-level oversight [25]. Activity-aware cardiovascular assessment shares this risk because wearable data are collected under highly uncontrolled conditions. An effective fairness program should include representative recruitment strategies, device validation across skin tones and body types, activity protocol diversity, and external auditing by groups not involved in model development. The explanation interface should allow clinicians to see not only the model features that influenced a prediction but also the demographic and activity profile of the training data most relevant to the current patient. Such a design supports critical judgment rather than blind reliance on algorithmic output. Without this, activity-aware systems may become another source of hidden structural bias in cardiovascular care.

7. Policy and Regulatory Implications

Regulatory evaluation of activity-aware cardiovascular assessment must account for the continuous and context-dependent nature of wearable inference. Traditional medical device regulation often assumes a relatively fixed relationship between inputs and outputs, but wearable systems may update activity recognition models or physiological state estimators over time. Adaptive algorithms raise difficult questions about when a model modification requires revalidation. A change in the activity recognition component may indirectly alter the clinical performance of the cardiovascular state estimator even if the cardiovascular model itself is unchanged. Regulatory frameworks should therefore require evidence of system-level performance under realistic motion conditions, not merely component-level accuracy [19]. This implies testing across multiple activity intensities, device placements, and patient subgroups.

Governance in this domain also intersects with data protection and patient rights. The continuous collection of physiological and motion data creates a rich record that can be used for purposes beyond immediate clinical assessment, including insurance risk stratification, employer monitoring, and secondary research. Patients may not fully understand how their data will be used or how inferences about their cardiovascular state will be shared. Trustworthy artificial intelligence frameworks emphasize human agency, oversight, and accountability [22]. These principles are especially relevant when wearable devices are used outside clinical settings where patients may have less direct supervision. Systems should provide accessible explanations to patients as well as clinicians, enabling individuals to understand why a device issued a cardiovascular alert and what actions they should take. Patient-facing explanations must be carefully designed to avoid unnecessary alarm while still conveying appropriate urgency.

Policy mechanisms should also encourage longitudinal safety monitoring. Postmarket surveillance for wearable cardiovascular systems should include not only adverse clinical events but also systematic audits of explanation quality, alert accuracy, and subgroup performance. Health systems should be required to maintain audit logs that link each alert to the model version, activity context, signal quality, and explanation provided at the time. Such logs enable retrospective review when a patient experiences a missed event or a false alarm. They also provide a basis for learning system improvements. Finally, accountability must be distributed across manufacturers, health systems, and clinical users in ways that do not obscure responsibility. An explainable system can support accountability by making the basis of a decision visible, but it cannot replace professional clinical judgment. The goal is to create a structured environment in which activity-aware cardiovascular assessment is transparent, contestable, and continuously improved.

8. Conclusion

Activity-aware cardiovascular state assessment using wearable sensors offers substantial promise for moving cardiovascular care from episodic clinical encounters to continuous ambulatory observation. However, physical activity introduces profound complexity that cannot be addressed by predictive accuracy alone. Explainable artificial intelligence should be treated as a system-level requirement that spans sensing, activity context modeling, physiological inference, uncertainty communication, and clinical governance. The paper has argued that activity-aware explanations must clarify how motion context and signal quality influence cardiovascular state estimates, thereby enabling clinicians and patients to distinguish real physiological changes from artifact-driven variation.

The structural design choices discussed in this paper include edge-cloud partitioning, multimodal fusion, self-supervised representation learning, explainability interfaces, fairness monitoring, and regulatory alignment. These choices trade off latency, privacy, model capacity, explanation fidelity, and operational sustainability. A robust system must be explicit about these trade-offs rather than optimizing a single technical metric. Future research should develop standardized protocols for activity-aware explanation, including benchmark datasets with diverse populations and realistic motion artifacts, as well as evaluation frameworks that measure both clinical utility and societal impact. By embedding explainability throughout the system architecture, activity-aware wearable cardiovascular monitoring can become not only more accurate but also more trustworthy, equitable, and clinically sustainable.

References

1. Kim, J., Campbell, A. S., de Avila, B. E.-F., & Wang, J. (2019). Wearable biosensors for healthcare monitoring. *Nature Biotechnology*, 37(4), 389–406.
2. Steinhubl, S. R., Muse, E. D., & Topol, E. J. (2013). The emerging field of mobile health. *Science Translational Medicine*, 5(196), 196rv2.
3. Anguita, D., Ghio, A., Oneto, L., Parra, X., & Reyes-Ortiz, J. L. (2013). A public domain dataset for human activity recognition using smartphones. In *Proceedings of the 21st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN 2013)* (pp. 437–442). Bruges, Belgium.
4. Ordonez, F. J., & Roggen, D. (2016). Deep convolutional and LSTM recurrent neural networks for multimodal wearable activity recognition. *Sensors*, 16(1), 115.
5. Barredo Arrieta, A., Diaz-Rodriguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115.
6. 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.
7. 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.
8. Attia, Z. I., Kapa, S., Lopez-Jimenez, F., McKie, P. M., Ladewig, D. J., Satam, G., Pellikka, P. A., Enriquez-Sarano, M., Noseworthy, P. A., Munger, T. M., Asirvatham, S. J., Scott, C. G., Carter, R. E., & Friedman, P. A. (2019). Screening for cardiac contractile dysfunction using an artificial intelligence-enabled electrocardiogram. *Nature Medicine*, 25(1), 70–74.
9. Pantelopoulos, A., & Bourbakis, N. G. (2010). A survey on wearable sensor-based systems for health monitoring and prognosis. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 40(1), 1–12.
10. Johnson, A. E. W., Pollard, T. J., Shen, L., Lehman, L. H., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L. A., & Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3, 160035.

11. Holzinger, A., Langs, G., Denk, H., Zatloukal, K., & Muller, H. (2019). Causability and explainability of artificial intelligence in medicine. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9(4), e1312.
12. Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision* (pp. 618–626).
13. Sundararajan, M., Taly, A., & Yan, Q. (2017). Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning* (pp. 3319–3328).
14. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144).
15. Tjoa, E., & Guan, C. (2020). A survey on explainable artificial intelligence (XAI): Toward medical XAI. *IEEE Transactions on Neural Networks and Learning Systems*, 32(11), 4793–4813.
16. 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.
17. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453.
18. Majumder, S., Mondal, T., & Deen, M. J. (2017). Wearable sensors for remote health monitoring. *Sensors*, 17(1), 130.
19. Gerke, S., Babic, B., Evgeniou, T., & Cohen, I. G. (2020). The need for a system view to regulate artificial intelligence/machine learning-based software as a medical device. *npj Digital Medicine*, 3, 53.
20. Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. *PLoS Medicine*, 15(11), e1002689.
21. Tonekaboni, S., Joshi, S., McCradden, M. D., & Goldenberg, A. (2019). What clinicians want: Contextualizing explainable machine learning for clinical end use. In *Proceedings of Machine Learning Research*, 106, 359–375.
22. High-Level Expert Group on Artificial Intelligence. (2019). Ethics guidelines for trustworthy AI. European Commission.
23. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *New England Journal of Medicine*, 380(14), 1347–1358.
24. Isakadze, N., & Martin, S. S. (2020). Wearable technology in cardiology: current evidence and future applications. *Trends in Cardiovascular Medicine*, 30(8), 480–488.
25. Beam, A. L., & Kohane, I. S. (2018). Big data and machine learning in health care. *JAMA*, 319(13), 1317–1318.