

Edge AI-Based Real-Time Physiological Signal Processing for Wearable Health Monitoring

Neeraj Shetty

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

helloneeraj@clemson.edu

Abstract

Wearable health monitoring has moved from episodic measurement toward continuous, longitudinal observation of physiological state, generating new requirements for real-time signal processing, privacy preservation, and system-level governance. Edge artificial intelligence has emerged as a promising architectural strategy for processing electrocardiographic, photoplethysmographic, electrodermal, and inertial signals near the body rather than relying exclusively on centralized cloud infrastructure. This paper presents a system-oriented analysis of edge AI for wearable physiological monitoring, emphasizing architectural trade-offs, resource governance, robustness, fairness, deployment sustainability, and policy implications. Instead of treating edge AI as a narrow algorithmic advance, the paper examines it as a distributed sociotechnical infrastructure in which sensing, inference, communication, and clinical decision support must be jointly governed. The discussion addresses the layered structure of wearable edge systems, the challenge of motion artifact resilience in ambulatory settings, and the tension between local model autonomy and centralized learning. It further considers federated learning as a governance mechanism, the regulatory complexity of automated health inference, and the clinical integration difficulties that arise when models trained in controlled settings are exposed to heterogeneous real-world conditions. The paper argues that sustainable and equitable wearable edge-AI systems require adaptive model management, transparent subgroup evaluation, staged clinical validation, and clear accountability across device manufacturers, software providers, health systems, and regulators. Future research should focus on energy-aware inference, model drift monitoring, cross-device interoperability, and governance frameworks that account for the continuous and intimate nature of wearable physiological data.

Keywords

edge artificial intelligence, wearable health monitoring, physiological signal processing, real-time inference, federated learning, responsible deployment, system architecture.

1. Introduction

The rapid expansion of wearable sensing technologies has fundamentally altered the possibilities for continuous health monitoring outside clinical environments. Modern wristbands, chest patches, smart garments, and implantable monitors can now capture electrocardiographic signals, photoplethysmographic waveforms, skin temperature, electrodermal activity, and movement data over extended periods. The clinical value of these data depends less on raw acquisition than on the ability of the broader system to transform noisy, high-frequency physiological streams into actionable health indicators with minimal delay and robust privacy protections. The convergence of wearable hardware with artificial intelligence has therefore shifted attention from mere data collection to distributed inference architectures that can operate close to the human body. High-performance medicine

increasingly depends on such computational fabrics because centrally processed physiological data often arrive too late for meaningful intervention and raise substantial governance burdens [1]. Wearable sensors are now central to remote health monitoring because they support longitudinal observation of cardiovascular, respiratory, and neurological status in everyday settings [2]. At the same time, the heterogeneity of sensor modalities and the complexity of human movement require multi-sensor fusion strategies that can reconcile signals with differing sampling rates, noise profiles, and physiological significance [3].

Edge computing has emerged as a critical architectural response to the limits of cloud-centric processing for latency-sensitive and bandwidth-intensive wearable applications. Rather than transmitting every raw waveform to a distant data center, edge artificial intelligence places feature extraction, anomaly detection, and early classification functions on devices or nearby gateways. This approach reduces dependence on stable connectivity, strengthens privacy by limiting the movement of sensitive physiological data, and enables real-time interventions in acute or safety-critical contexts. The conceptual foundations of edge computing emphasize proximity to data sources, geographic distribution, and the need to balance local autonomy with cloud-scale orchestration [4]. Extending this vision, edge intelligence integrates machine learning inference directly into resource-constrained nodes, thereby shifting part of the learning pipeline from centralized clusters to heterogeneous field devices [5]. Deep learning models have already demonstrated cardiologist-level arrhythmia detection in ambulatory electrocardiograms, illustrating both the clinical promise and the computational challenge of deploying sophisticated models in continuous wearable settings [6]. Physiological signals acquired during physical activity, however, remain vulnerable to motion artifacts, and adaptive filtering methods have been proposed to improve signal recovery under such conditions [7]. This paper examines the system-level implications of these developments, focusing on architectural trade-offs, resource governance, robustness, fairness, privacy, sustainability, and policy rather than on isolated algorithmic improvements.

2. Architectural Foundations of Wearable Edge-AI Systems

The system architecture for wearable edge artificial intelligence is best understood as a layered continuum spanning sensing nodes, personal edge devices, local gateways, and cloud services. The distinguishing property of this continuum is that computational authority can be dynamically allocated according to clinical criticality, energy availability, privacy sensitivity, and network conditions. The rise of edge computing has introduced a class of infrastructure that sits between sensor-rich environments and centralized clouds, enabling interactive and data-intensive services without requiring all data to traverse wide-area networks [8]. Mobile edge computing further refines this model by embedding computational resources within the radio access network or in close proximity to mobile users, thereby reducing end-to-end latency and backhaul congestion for physiological monitoring applications [9]. The broader Internet of Things ecosystem provides the protocols, discovery mechanisms, and middleware foundations that allow diverse wearable devices, smartphones, and environmental sensors to coordinate their sensing and actuation activities [10]. Within such an ecosystem, the wearable monitoring system operates as a specialized cyber-physical subsystem in which body area networks must be managed through lightweight communication stacks and service-oriented abstractions [11].

A central architectural challenge is deciding which processing functions should be performed on the wearable device, which should be delegated to a nearby edge gateway, and which should be uploaded to a cloud environment for population-scale learning. On-sensor

processing offers the strongest privacy and latency guarantees but is constrained by battery capacity, thermal dissipation limits, memory bandwidth, and the computational intensity of modern deep learning models. Edge gateways such as smartphones, home hubs, and bedside routers provide intermediate computational capacity that can support more complex inference while still keeping raw signals within a local trust boundary. Cloud resources remain necessary for retrospective model retraining, longitudinal cohort analysis, and cross-institutional aggregation. This tiered arrangement is not simply a technical convenience; it reflects a governance decision about where physiological data can flow, how long they can reside, and which parties can access them. The management of streaming data across these tiers introduces significant big data challenges because physiological signals arrive continuously, vary in resolution, and must be combined with contextual metadata to be clinically meaningful [12]. Early wearable sensor-based health monitoring systems already highlighted the need for systematic approaches to sensor placement, artifact handling, feature extraction, and decision support to translate raw signals into reliable health prognostics [13].

The architectural choices made at this level have direct implications for system robustness and adaptability. A tightly coupled architecture that assumes continuous connectivity to a cloud inference engine may fail during wireless outages or in rural settings, whereas a fully autonomous wearable node may miss the benefits of population-level model updates. Effective deployments therefore adopt a hybrid control model in which local models maintain a baseline of monitoring capability, while asynchronous synchronization mechanisms permit periodic refinement using aggregated observations. This design requires careful attention to versioning, model rollback, sensor calibration drift, and the coexistence of multiple inference pipelines that may operate at different temporal granularities. The remainder of this paper examines how these architectural tensions shape physiological signal processing, resource governance, privacy, fairness, and long-term sustainability.

3. Real-Time Physiological Signal Processing and Artifact Resilience

Real-time processing of physiological signals in wearable systems must contend with a range of nonstationary noise sources that are rarely present in controlled clinical monitoring. Motion artifacts, sensor displacement, skin impedance changes, environmental electromagnetic interference, and intermittent skin contact can all corrupt the underlying signal. Electrocardiographic and photoplethysmographic waveforms are particularly sensitive to physical activity because the same bodily movements that make ambulatory monitoring valuable also generate substantial baseline wander and high-frequency artifact. If these corruptions are not addressed before inference, downstream models may produce false alarms or miss clinically relevant episodes. Adaptive filtering methods have been developed to suppress periodic and quasiperiodic interference while preserving the morphological features needed for subsequent analysis [7]. Such preprocessing is not a separate academic concern but an essential front end for edge artificial intelligence because the performance of any downstream classifier is bounded by the quality of the signal representation it receives.

The real-time signal processing pipeline in a wearable edge-AI system typically involves multiple stages that must be scheduled within strict latency and energy budgets. After acquisition, the signal must be filtered and segmented into windows that are long enough to capture diagnostic features but short enough to support timely intervention. Quality assessment algorithms discard or downweight corrupted segments, while feature extraction layers transform filtered waveforms into compact representations suitable for local inference. In some configurations, lightweight neural networks operate directly on raw or minimally

processed signals to reduce hand-engineered preprocessing, but this approach can increase computational cost and may be less interpretable to clinical engineers. Hybrid pipelines that combine classical signal conditioning with learned feature extractors often offer a better balance for wearable hardware because they allow domain-specific constraints such as sampling rate, sensor noise characteristics, and physiological plausibility to be embedded in the front end. The choice among these alternatives is deeply architectural, as it determines the division of labor between signal processing accelerators, general-purpose processors, and neural processing units within the wearable device.

Different physiological modalities present distinct trade-offs. Electrocardiographic monitoring requires precise preservation of timing and morphology, which means that aggressive filtering can distort rhythm classification. Photoplethysmography is more tolerant of low-frequency trend removal but remains susceptible to venous pulsation and motion artifact. Electrodermal activity evolves over much longer timescales and is often used for stress or arousal estimation rather than acute event detection. Multi-sensor fusion can use accelerometer data to identify movement periods and adapt filtering parameters accordingly, but this requires synchronized timing and careful management of sensor drift. In all cases, the real-time processing goal is not to maximize a single signal quality metric in isolation but to ensure that the information reaching the inference stage is sufficiently stable for the clinical or wellness context in which the system operates. This context dependence is precisely why edge-AI systems must be designed as adaptive sociotechnical infrastructures rather than as static algorithmic pipelines.

4. Resource Governance and Inference Trade-Offs

The deployment of real-time physiological signal processing on wearable devices introduces a persistent tension between model sophistication and resource availability. Wearable devices are constrained by small batteries, limited thermal dissipation, constrained memory, intermittent connectivity, and highly variable user contexts. Deep learning models that perform well on retrospective datasets may be far too expensive for continuous on-device execution if they require numerous multiply-accumulate operations per second or large activation buffers. System designers must therefore adopt a portfolio of efficiency strategies including model quantization, pruning, knowledge distillation, and the use of recurrent or attention-based architectures specifically tailored for streaming physiological data. Each strategy entails a trade-off among accuracy, latency, energy consumption, memory footprint, and interpretability. A quantized model may preserve diagnostic sensitivity while reducing energy demand, but it may also exhibit subtle calibration shifts when deployed across heterogeneous hardware generations. Resource governance is thus not a one-time optimization but an ongoing process of monitoring performance drift, energy budgets, and device health in the field.

Federated learning has become an influential framework for updating models across populations of wearable devices without transferring raw physiological data to a central repository. This approach allows local devices to compute model updates on private data and transmit only those updates to an aggregation server, thereby reducing the privacy risks associated with centralized data accumulation [14]. However, federated learning in wearable health monitoring is not free from governance challenges. Participant devices may have heterogeneous sensing hardware, inconsistent sampling rates, different user demographics, and skewed distributions of health conditions. The aggregation server must manage non-independent and non-identically distributed data, detect malicious or faulty updates, and avoid model divergence when local data distributions shift. Furthermore, the communication

efficiency of federated learning remains a concern because model updates can still be large relative to the bandwidth available to wearable gateways. These considerations reveal that privacy-preserving learning is not a single algorithmic choice but a socio-technical infrastructure that must be governed across technical, organizational, and regulatory layers.

Regulatory expectations add another dimension to resource and model governance. The General Data Protection Regulation and similar frameworks impose obligations regarding transparency, data minimization, and automated decision-making that interact with edge-AI architectures in complex ways. Legal scholarship has critically examined whether individuals possess a right to explanation of automated decisions and has shown that such a right is less straightforward than often assumed [15]. For wearable health systems, this means that designers cannot simply rely on model accuracy as a proxy for regulatory compliance. They must consider how to communicate the behavior of local inference engines to users, clinicians, and regulators without revealing proprietary model details or overwhelming non-specialist stakeholders. Responsible machine learning for health care requires proactive attention to safety, efficacy, equity, and accountability, particularly when algorithms influence clinical decisions outside controlled hospital environments [16]. Edge-AI deployments must therefore embed governance mechanisms directly into the system lifecycle, from data acquisition and model training to update distribution and incident reporting.

5. Privacy, Fairness, and Responsible Clinical Integration

Edge-AI wearable monitoring changes the privacy landscape by keeping more data local, but it does not eliminate privacy risk. Local processing can reveal sensitive health states through model outputs, alarm behaviors, timing patterns, or even energy consumption signatures. A device that continuously infers arrhythmias, sleep stages, or stress levels is itself a disclosure risk if accessed by unauthorized parties. The privacy benefits of edge processing must therefore be evaluated alongside device security, secure element storage, encrypted communication, and robust identity management. Machine learning in medicine has the potential to improve diagnosis and monitoring, but this potential depends on access to high-quality data, rigorous validation, and transparent evaluation across diverse subpopulations. When edge-AI models are trained predominantly on data from populations with limited demographic diversity, they may encode spurious associations between sensor artifacts, skin tone, body mass, movement patterns, and underlying health states. Algorithmic bias in health care is not hypothetical; population-level analyses have documented systematic disparities in the allocation of care driven by biased predictive models [17]. In wearable monitoring, such biases can manifest as differential sensor signal quality, variable false alarm rates, or unequal access to model updates across device generations. Fairness auditing must therefore become part of the deployment lifecycle, with disaggregated evaluation across demographic groups, activity levels, skin phototypes, and comorbid conditions. Without such auditing, local inference engines may appear accurate in aggregate while failing on the very subpopulations most in need of continuous monitoring.

Explainability is closely related to fairness but should not be conflated with a single technical method. In wearable health monitoring, clinicians and patients may need to understand why a particular alarm was generated, what physiological features triggered a classification, and how confident the system is under current signal conditions. A multidisciplinary perspective on explainability suggests that the goals of local agency, clinical trust, and regulatory accountability require more than heatmaps or salience maps [18]. They require structured explanations that can be interpreted within an existing clinical workflow and updated as the

model changes. Edge-AI systems, with their constrained user interfaces and real-time operation, present particularly difficult explanation challenges because long textual rationales may be inappropriate during an acute event. The explanation must be actionable and context-sensitive rather than technically exhaustive, which in turn requires close collaboration between interface designers, clinical informaticians, and machine learning engineers.

From a regulatory standpoint, edge-based physiological signal processing that influences diagnosis or treatment will generally be subject to medical device requirements in many jurisdictions. The practical implementation of artificial intelligence in medicine demands staged validation, post-market surveillance, and clear definitions of intended use [19]. For wearable edge-AI, intended use is complicated by the fact that the same device may provide wellness estimates in one market, remote patient monitoring in another, and decision support in yet another. Regulatory agencies are increasingly exploring frameworks for adaptive machine learning software, but the distributed nature of edge inference creates additional challenges for version control, adverse event attribution, and recall management. A system that updates its model across thousands of devices must be able to trace which model version generated which inference, a requirement that pushes wearable deployments toward stronger data governance and configuration management than is typically found in consumer-grade hardware. The responsible clinical integration of edge-AI monitoring therefore depends on an institutional commitment to audit trails, clinician training, and clear escalation pathways when automated inferences conflict with patient symptoms or local clinical judgment.

6. Adaptive Model Management and Drift in Field Deployments

Continuous health monitoring systems operate in environments that are far more heterogeneous than the controlled datasets used for initial model development. Sensor degradation, seasonal changes, medication use, skin hydration, device placement habits, and the progression of chronic disease can all alter the statistical relationship between raw signals and physiological states. A model that performs well at deployment may gradually lose calibration or develop unexpected failure modes as the population and device fleet evolve. Adaptive model management therefore requires mechanisms for detecting distribution shift, deciding whether a local model should be updated, and verifying that updates do not degrade performance on previously stable subgroups. The computational constraints of edge devices make model updates especially challenging. Even when a more accurate model is available from the cloud, shipping and executing it on millions of heterogeneous wearables may be expensive and may disrupt service continuity. Techniques such as pruning, quantization, and weight sharing can reduce the memory and energy requirements of deep networks, allowing sophisticated models to operate within the computational budgets of low-power devices [20]. However, compressed models may not transfer equally well across hardware generations, and the interaction between compression artifacts and physiological signal noise is not yet well understood. A responsible update strategy should therefore incorporate staged rollouts in which new models are first deployed to a small, diverse set of devices under enhanced monitoring before broader distribution.

The governance of model adaptation cannot be separated from broader ethical frameworks for health-related artificial intelligence. International guidance emphasizes that autonomy, safety, transparency, and accountability must shape the entire algorithm lifecycle, from design through deployment and decommissioning [21]. In the edge context, these principles imply that model updates cannot be treated as routine software patches but must be subjected to health-specific risk assessment, particularly when updates alter the sensitivity or specificity of

diagnostic alarms. Drift management thus operates under a broader mandate that includes human oversight, algorithmic transparency, and the protection of users who may not have the technical literacy to evaluate changes in device behavior. Continual learning research has shown that sequential adaptation can lead to catastrophic forgetting when earlier knowledge is overwritten by newer data [22]. In wearable health monitoring, such forgetting is clinically untenable because a model update that improves detection for a newly prevalent condition must not sacrifice sensitivity for a rare but life-threatening arrhythmia. Model drift monitoring should therefore combine statistical process control over input features, periodic evaluation on held-out reference data, and clinician review of alarm patterns.

Drift management also exposes a governance tension between local autonomy and centralized oversight. If every wearable device were permitted to adapt its model continuously based on local data, the fleet would quickly diverge, making it difficult to audit clinical performance or reproduce adverse events. Conversely, if adaptation is too centralized, devices may remain stale between infrequent update cycles and fail to respond to emerging physiological changes. The most viable systems are likely to adopt a hierarchical approach in which local devices detect drift and send summary statistics or small calibration updates to an edge orchestrator, while substantive model retraining occurs in the cloud under controlled validation protocols. This hierarchy requires clear metadata about sensor type, firmware version, demographic context, and episode labels, but such metadata can itself be privacy sensitive and must be handled under the governance principles described earlier. Ultimately, adaptive model management should be treated as a clinical safety function rather than a purely engineering task, with defined roles for regulatory review, clinical oversight, and user notification when significant model changes occur.

7. Deployment Sustainability and Interoperability

Sustainability in wearable edge-AI encompasses not only energy consumption but also hardware lifecycle, software maintenance, data stewardship, and organizational capacity. Energy-efficient inference is a precondition for long-term wearability, but the broader environmental footprint of edge-AI includes the production and disposal of sensors, the energy used during cloud retraining, and the carbon costs of model experimentation. System designers should adopt a life-cycle perspective that considers how device obsolescence, firmware update policies, and security patch commitments affect users who may depend on the technology for continuous monitoring. A wearable monitor that requires frequent cloud retraining may be clinically effective but environmentally and economically unsustainable if the underlying cloud infrastructure is not powered by renewable energy or if model updates are too large for users in low-bandwidth settings. Interoperability is another critical dimension of sustainable deployment. Wearable health devices often operate within fragmented ecosystems where proprietary data formats, communication protocols, and closed application programming interfaces prevent seamless integration with electronic health records and clinical decision support systems. The adoption of modern health interoperability standards can reduce this fragmentation and allow physiological signals and inferred states to be exchanged across systems without loss of context [23]. However, interoperability must not be treated as a purely technical compatibility problem. It also requires governance agreements about data provenance, consent, liability, and the clinical meaning of inferred variables. Without shared semantic models, a heart rate variability estimate produced by one edge-AI system may not be directly comparable to that produced by another, even if both devices export the same field name. Sustainable deployment therefore demands both technical

standardization and institutional coordination among manufacturers, health systems, and regulators.

Equity in deployment is deeply shaped by these sustainability and interoperability considerations. Users in well-resourced health systems may benefit from continuous model updates, seamless integration with clinical teams, and early access to new sensing modalities, while users in lower-resource settings may be left with older devices, limited connectivity, and less clinical follow-up. Edge-AI architectures that reduce dependence on continuous connectivity can mitigate some of these disparities, but only if model updates and device support are also designed for low-bandwidth and intermittent environments. Fairness in edge-AI therefore extends beyond algorithmic bias to include the distribution of computational resources, maintenance burdens, and clinical integration capacities. A system that is technically robust in a laboratory but too expensive to maintain in community health settings cannot be considered equitable or sustainable. Long-term success will require procurement policies, open standards, and public health investments that enable edge-AI monitoring to serve diverse populations without reinforcing existing digital divides.

8. Governance and Policy Implications

The governance of wearable edge-AI for physiological monitoring must address the continuous, intimate, and distributed nature of the data and inferences involved. Traditional medical device regulation tends to assume a relatively stable product with a defined intended use, but edge-AI systems combine hardware, software, model updates, and cloud services in ways that blur the boundaries between products and ongoing services. Policy frameworks must therefore recognize that the safety of a wearable monitoring system depends not only on a single clearance decision but on the entire lifecycle of model management, signal processing updates, cybersecurity patches, and data governance practices. Regulatory authorities will need to develop mechanisms for certifying update pipelines, auditing model versions, and requiring proactive post-market surveillance that can detect performance degradation across subpopulations and clinical contexts. These mechanisms should not be so burdensome that they stifle innovation, but they must be rigorous enough to protect patients who rely on continuous monitoring for arrhythmia detection, seizure warning, or glycemic management.

Privacy governance in edge-AI wearable systems cannot be reduced to a single compliance checklist. Although local processing reduces the movement of raw physiological data, it also creates local repositories of sensitive inferences that may be vulnerable to device theft, unauthorized access, or coercive sharing. Data protection law generally requires data minimization and purpose limitation, but these principles are complicated when a single device serves multiple purposes across wellness, research, and clinical care. Clear data governance architectures are needed to separate consent domains, manage third-party model updates, and ensure that users retain meaningful control over how their physiological signals are used. At the same time, overly fragmented consent architectures can undermine clinical utility by making it difficult to combine data across devices or share necessary information with emergency responders. Policy frameworks should therefore support context-dependent data sharing that is technically enforced at the edge rather than relying solely on lengthy legal notices.

Accountability is a further policy challenge in distributed edge-AI ecosystems. When a wearable device produces a false negative for a serious arrhythmia, responsibility may be distributed across the sensor manufacturer, the algorithm developer, the edge gateway provider, the hospital that prescribed the device, and the cloud operator that distributed the

model update. This diffusion of responsibility can make it difficult for patients to seek redress and for regulators to identify the appropriate party for corrective action. Policies should encourage clear contractual and technical mechanisms for attributing decisions to specific model versions and data governance actions. A common traceability infrastructure, combined with transparent incident reporting and independent audits, would improve accountability without necessarily centralizing all processing in the cloud. The future of wearable edge-AI will depend on the ability of technologists, clinicians, regulators, and communities to co-produce governance arrangements that preserve the latency, privacy, and resilience benefits of edge computing while ensuring that health monitoring remains safe, equitable, and aligned with public interest.

9. Conclusion

Edge artificial intelligence has the potential to transform wearable health monitoring from a passive data collection activity into a responsive, privacy-preserving, and clinically useful infrastructure. The architectural placement of inference at or near the body reduces latency, lowers dependence on continuous connectivity, and limits unnecessary exposure of raw physiological signals. However, these technical advantages do not automatically translate into safe, fair, or sustainable deployment. Real-time physiological signal processing must contend with motion artifact, sensor heterogeneity, resource constraints, and the challenge of maintaining model validity across diverse populations and changing conditions. Adaptive notch filtering and related signal conditioning techniques are important components of the preprocessing pipeline, but they must be integrated into a broader governance architecture that manages model updates, monitors drift, and ensures accountable clinical integration.

The system-level analysis presented in this paper shows that wearable edge-AI should be treated as a distributed sociotechnical infrastructure rather than a collection of optimized algorithms. Architectural trade-offs between local inference and cloud retraining are simultaneously technical, clinical, economic, and regulatory choices. Federated learning can reduce privacy risks, but it introduces new challenges in heterogeneity, communication efficiency, and auditability. Fairness audits, explainable interfaces, staged validation, and lifecycle sustainability are essential to responsible deployment. Policy frameworks must evolve to accommodate adaptive model updates, distributed accountability, and interoperability across fragmented device ecosystems.

Future research should address energy-aware model compression under physiological workload constraints, robust drift detection for ambulatory monitoring, cross-device and cross-population generalization, and governance mechanisms that preserve user autonomy without undermining clinical safety. Particularly important is the development of open evaluation benchmarks that reflect the heterogeneity of real-world wearable use, including diverse skin tones, body types, movement patterns, and connectivity conditions. Without such benchmarks, the field risks optimizing for narrow datasets while failing to deliver equitable benefit. The path forward lies in interdisciplinary collaboration that treats edge AI not as an isolated computational advance but as a complex health infrastructure whose ultimate value must be judged by its impact on patient outcomes, system resilience, and social equity.

References

1. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56.

2. Majumder, S., Mondal, T., & Deen, M. J. (2017). Wearable sensors for remote health monitoring. *Sensors*, 17(1), 130.
3. Gravina, R., Alinia, P., Ghasemzadeh, H., & Fortino, G. (2017). Multi-sensor fusion in body sensor networks: State-of-the-art and research challenges. *Information Fusion*, 35, 68–80.
4. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646.
5. Zhou, Z., Chen, X., Li, E., Zeng, L., Luo, K., & Zhang, J. (2019). Edge intelligence: Paving the last mile of artificial intelligence with edge computing. *Proceedings of the IEEE*, 107(8), 1738–1762.
6. 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.
7. 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.
8. Satyanarayanan, M. (2017). The emergence of edge computing. *Computer*, 50(1), 30–39.
9. Mach, P., & Becvar, Z. (2017). Mobile edge computing: A survey on architecture and computation offloading. *IEEE Communications Surveys & Tutorials*, 19(3), 1628–1656.
10. Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Generation Computer Systems*, 29(7), 1645–1660.
11. Chen, M., Gonzalez, S., Vasilakos, A., Cao, H., & Leung, V. C. (2011). Body area networks: A survey. *Mobile Networks and Applications*, 16(2), 171–193.
12. Chen, M., Mao, S., & Liu, Y. (2014). Big data: A survey. *Mobile Networks and Applications*, 19(2), 171–209.
13. 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.
14. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 1273–1282.
15. Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76–99.
16. 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.

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. 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(1), 310.
19. He, J., Baxter, S. L., Xu, J., Xu, J., Zhou, X., & Zhang, K. (2019). The practical implementation of artificial intelligence technologies in medicine. *Nature Medicine*, 25(1), 30–36.
20. Sze, V., Chen, Y. H., Yang, T. J., & Emer, J. S. (2017). Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12), 2295–2329.
21. World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. World Health Organization.
22. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., & Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526.
23. Mandel, J. C., Kreda, D. A., Mandl, K. D., Kohane, I. S., & Ramoni, R. B. (2016). SMART on FHIR: A standards-based, interoperable apps platform for electronic health records. *Journal of the American Medical Informatics Association*, 23(5), 899–908.