

Energy-Efficient Algorithms for Continuous Physiological Monitoring in Resource-Constrained Wearable Devices

Sean L. Riley

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.

sean793@colostate.edu

Zachary Boleman

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

zacharywork@unt.edu

Abstract

Continuous physiological monitoring through wearable devices has become central to precision health, remote care, and early clinical intervention. However, the promise of uninterrupted sensing is constrained by the severe energy, computational, communication, and memory limitations of miniature devices. This paper examines energy-efficient algorithms for continuous physiological monitoring from a system-level perspective, integrating architectural trade-offs, signal processing, adaptive sampling, embedded inference, and communication scheduling. Rather than presenting a single optimization technique, the discussion situates algorithmic efficiency within the larger infrastructure of wearable sensing, edge computing, and health data governance. The paper analyzes how dynamic power management, compressed sensing, and adaptive duty cycling interact with clinical reliability and device longevity. It further explores structural trade-offs among local processing, communication offloading, and cloud aggregation, arguing that energy efficiency cannot be separated from robustness, fairness, and regulatory compliance. Using examples including motion-tolerant photoplethysmography, electrocardiogram monitoring, and neural network compression, the analysis highlights how algorithmic decisions influence device sustainability and equitable health outcomes. The paper concludes that future wearable systems should be designed around adaptive, context-aware resource allocation mechanisms that reconcile clinical validity with energy constraints while embedding fairness and privacy considerations across the full deployment lifecycle. The review offers a cross-domain perspective for engineers, clinical informatics researchers, and policy stakeholders seeking to understand the design space of resource-constrained continuous monitoring.

Keywords

energy-aware computing; wearable sensors; physiological monitoring; embedded machine learning; digital health governance.

1. Introduction

Continuous physiological monitoring through wearable devices has emerged as a foundational capability for remote health care, chronic disease management, and early intervention. Wearable systems now measure cardiac activity, oxygen saturation, movement, temperature, and other physiological variables outside clinical settings, generating longitudinal data that can support personalized treatment and population health initiatives.

Surveys of wearable sensor systems have documented the rapid expansion of sensing modalities and their clinical relevance [1]. At the same time, the practical deployment of these devices remains constrained by fundamental resource limitations, including battery capacity, processing power, memory, and wireless communication bandwidth [2]. These constraints are especially acute for devices intended to operate continuously over days or weeks on small batteries, often in ambulatory environments where connectivity and energy availability vary.

The central tension in designing energy-efficient algorithms for continuous monitoring is that clinical usefulness depends on sustained, high-quality observation, while device longevity depends on minimizing energy expenditure. Early work on dynamic power management established that system-level energy savings require coordinated decisions across sensing, computation, and communication rather than isolated optimizations [3]. In wearable physiological monitors, this coordination is more complex because physiological signals are nonstationary, context-dependent, and clinically sensitive. An algorithm that reduces energy consumption by discarding or compressing data may also reduce diagnostic fidelity or introduce population-specific biases. Energy efficiency must therefore be understood as a structural property of the entire monitoring infrastructure, not merely as a property of individual processing routines. This paper examines the design space of energy-efficient algorithms from a system-level perspective, emphasizing architectural trade-offs, robustness, fairness, governance, and sustainability. The discussion connects signal processing, embedded inference, communication scheduling, and policy considerations, and argues that sustainable continuous monitoring requires adaptive resource management that is aware of both physiological context and socioeconomic deployment conditions.

2. Related Work and System-Level Challenges

Research on energy-efficient wearable health monitoring has evolved from early sensor network designs to complex edge computing systems. A survey of context recognition systems for health care identified energy efficiency as a central challenge because continuous activity and physiological sensing often require processing multiple heterogeneous data streams [4]. That survey highlighted how decisions about feature extraction, classification complexity, and duty cycling influence overall energy consumption. Low-power sensor nodes designed for wireless research, such as the Telos platform, demonstrated that hardware sleep states and low-power communication protocols can substantially extend operational lifetime, but they also introduced latency and throughput constraints that are problematic for real-time physiological monitoring [5]. These platform-level lessons remain relevant because many wearable devices inherit the design assumptions of wireless sensor networks even though their clinical requirements are stricter.

Compressed sensing has been proposed as a means to reduce the amount of data that must be processed and transmitted for biosignals such as electrocardiogram and electromyogram recordings [6]. The general principle is that many physiological signals have sparse or compressible representations, allowing substantial dimensionality reduction before transmission. However, compressed sensing introduces reconstruction overhead and can be sensitive to motion artifacts and signal nonstationarity. More broadly, healthcare Internet-of-Things architectures have drawn attention to the need for clinical validation and reliable data flow across sensing, aggregation, and decision layers [7]. Body area network surveys also emphasize that wireless channel conditions, interference, and human mobility create uneven energy demands that static algorithms cannot fully address [8]. These bodies of work collectively indicate that energy efficiency in wearable monitoring is not simply a matter of

selecting low-power components or compressing data; it requires continuous negotiation among clinical requirements, environmental variability, and available computational resources.

3. Architectural Foundations and Resource Constraints

The architecture of a continuous physiological monitoring system determines where and how algorithmic decisions consume energy. Hierarchical architectures that combine wearable devices, local gateways, fog nodes, and cloud services can reduce some energy burdens by moving computationally intensive operations away from the body-worn device. Fog computing approaches have been explored for electrocardiogram feature extraction, showing that processing at nearby infrastructure can reduce latency and bandwidth while preserving clinically relevant information [9]. Nevertheless, reliance on off-device processing creates dependencies on network availability and introduces privacy risks, particularly when raw physiological data are transmitted beyond the user's immediate control. Energy-aware microsensor systems have long advocated for distributed coordination among sensing, processing, and communication tasks to reduce overall power consumption [10]. In these systems, the energy cost of communication often outweighs the cost of local computation, which suggests that selective local processing can be beneficial if the algorithms are sufficiently efficient.

At the signal conditioning stage, motion artifact management illustrates how algorithmic choices affect both energy and clinical quality. Wearable photoplethysmography and electrocardiography are frequently contaminated by movement, and robust artifact suppression can require filtering methods that adapt to changing noise conditions. Adaptive notch-filtration has been shown to recover photoplethysmographic signals during physical activity while maintaining low computational complexity [11]. Such adaptive filtering approaches are particularly important for resource-constrained devices because they avoid the high energy consumption associated with more complex source separation or deep reconstruction models, while still improving signal quality under dynamic real-world conditions. The placement of these signal conditioning routines in the processing pipeline, the frequency with which they are invoked, and the precision of their internal parameters all influence battery life and diagnostic reliability. A system-level architecture must therefore allocate energy budgets across acquisition, filtering, feature extraction, inference, and communication with explicit awareness of the clinical consequences of each allocation decision.

4. Energy-Aware Signal Acquisition and Motion Artifact Management

Signal acquisition is often assumed to be a fixed cost, but it can be a major source of energy consumption in wearable devices, particularly when high sampling rates are used continuously. The assumption that all signal detail must be captured at all times is increasingly questioned in energy-constrained sensing. Compressed sensing strategies for electrocardiogram and electromyogram biosensors demonstrate that lower effective acquisition rates can still preserve sufficient information for many clinical tasks [6]. However, the applicability of compressed sensing depends on the signal class, the presence of transient events such as arrhythmias, and the tolerance of downstream algorithms for reconstruction error. In continuous monitoring, a missed transient may have greater clinical consequences than increased energy consumption, so acquisition strategies must be dynamically adjusted according to patient risk and activity context. Motion artifact management adds another layer of complexity because movement affects sensor contact, optical coupling, and electrode

impedance, leading to bursts of noise that can obscure the underlying physiology. Adaptive filtering techniques, including adaptive notch-filtration, can suppress periodic motion components and recover photoplethysmographic signals during physical activity without the cost of high-dimensional reconstruction [11].

From a structural perspective, acquisition and artifact management should not be optimized independently. A device that reduces its sampling rate too aggressively may save energy but produce signals that are harder to interpret accurately, requiring more downstream computation or transmission to resolve ambiguity. Conversely, an adaptive filtering scheme that runs continuously at maximum precision may consume more energy than it saves by reducing retransmission or rejection of corrupted data. The appropriate balance depends on the signal type, the monitoring objective, and the operational context. For example, in a low-motion overnight monitoring scenario, a fixed filtering schedule may be sufficient and highly energy-efficient because the noise environment is relatively stable. In contrast, during exercise or rehabilitation, the system may need to activate adaptive filtering more frequently, accept short-term energy costs, and temporarily increase the sampling rate to preserve diagnostic information. These trade-offs are not merely technical; they structure the clinical and economic viability of prolonged monitoring programs.

5. Adaptive Sampling and Duty Cycling Strategies

Duty cycling is one of the most widely used mechanisms for reducing energy consumption in wearable devices. The principle is to alternate between active modes in which sensing and processing occur and sleep modes in which most subsystems are powered down. Dynamic power management frameworks have shown that effective duty cycling requires predictive knowledge of upcoming workloads and latency constraints [3]. In physiological monitoring, sleep states can reduce energy consumption dramatically, but they also create gaps in observation that may be unacceptable for detecting sporadic events such as seizures, falls, or dangerous arrhythmias. Consequently, duty cycling algorithms for medical wearables must be more conservative than those designed for environmental sensor networks. Low-power wireless platforms have demonstrated that hardware-level sleep scheduling can extend lifetime, but these gains may be offset by the need for frequent synchronization, clock drift correction, and rapid wake-up paths [5]. Energy-aware microsensor network designs further emphasize that communication scheduling should be tightly coupled with sensing duty cycles to avoid redundant wake-ups [10].

Adaptive sampling offers a more nuanced alternative to fixed duty cycles. Rather than simply turning the sensor on and off, adaptive sampling changes the rate or resolution of acquisition based on the estimated information content of the signal. In periods of physiological stability, the system can reduce sampling frequency and spend more time in low-power states. When elevated variability or clinically significant patterns are detected, the system can increase the sampling rate and activate additional processing modules. This approach has strong conceptual appeal for continuous monitoring because it aligns energy expenditure with clinical need. However, adaptive sampling introduces control complexity and can produce biased datasets if the adaptation rule is poorly calibrated. For example, if the system only increases sampling when motion is high, it may overrepresent noisy periods and underrepresent calm periods, complicating longitudinal analysis and model training. Therefore, adaptive sampling policies must be developed in conjunction with data governance practices that track how sampling decisions shape the resulting record.

6. On-Device Inference and Communication Trade-offs

The growing use of machine learning in physiological monitoring has shifted attention from signal processing to the energy cost of embedded inference. Deep neural networks can extract subtle patterns from physiological time series, but their computational demands often exceed the capabilities of low-power wearable processors. Efficient processing techniques such as hardware-aware algorithm design, reduced precision arithmetic, and specialized dataflows have substantially lowered the cost of deep learning inference [12]. These methods enable some clinical classification tasks to be performed locally, reducing the need to transmit raw data continuously. Network pruning further reduces memory and computation by removing redundant connections while preserving accuracy, making it possible to deploy compact models on resource-constrained devices [13]. Tiny deep learning frameworks have extended this direction by co-designing network structures and inference engines for microcontroller-class hardware [14]. These advances are important because they shift the energy balance from communication toward local computation, which can improve latency, privacy, and resilience to network outages.

However, local inference introduces new structural trade-offs. A model that runs efficiently on a wearable device may have lower representational capacity than a cloud-based model, potentially missing rare clinical events. Communication scheduling then becomes a mechanism for managing uncertainty: the device may perform initial screening locally and transmit only selected segments or features when confidence is low or when an alert condition is detected. Federated learning offers another architectural option by allowing multiple devices to contribute model updates without sharing raw physiological data [15]. Although federated learning reduces communication volume relative to centralized raw data upload, it still requires synchronization and can impose computational overhead on each device. Moreover, the energy cost of local training may be prohibitive for very small wearables unless the training workload is carefully scheduled during charging periods or offloaded to a companion device. The selection among local inference, edge offloading, and federated collaboration is therefore not a single technical decision but a system-level negotiation involving energy budgets, clinical accuracy, data protection, and operational sustainability.

7. Robustness, Fairness, and Clinical Validity Across Populations

Energy-efficient algorithms can inadvertently introduce or amplify health inequities if they are optimized primarily for average performance. Fairness in clinical machine learning requires explicit attention to how algorithmic decisions perform across different patient groups, because models trained on biased data may fail for underrepresented populations [16]. In the context of wearable monitoring, energy-saving mechanisms can interact with fairness in subtle ways. Adaptive sampling or noise suppression algorithms may be tuned using data from healthy young adults with regular motion patterns, leading to degraded performance for older adults, persons with gait disorders, or individuals with darker skin tones where optical sensing signals may differ. If an energy-efficient algorithm discards data more aggressively when signal quality is low, it may systematically under-record episodes for groups whose signals are less easily measured by the sensor technology. This can reduce clinical validity and deepen disparities in diagnosis and treatment.

Robustness is also central because continuous monitoring devices operate in uncontrolled environments. Energy-efficient algorithms that are tested under laboratory conditions may fail under varying humidity, temperature, clothing, and activity patterns. The ethical and governance literature on algorithms emphasizes that system designers must consider unintended consequences and the distribution of risks across stakeholders [17]. For

physiological monitoring, risk distribution is particularly important because the burden of false positives and false negatives is not shared equally. A false negative for a high-risk cardiac patient may delay urgent care, while frequent false positives in low-risk users may generate unnecessary alarm fatigue and health system burden. Energy constraints can worsen these outcomes if the system is forced to choose between operating modes that reduce diagnostic sensitivity or increase latency. Thus, robustness and fairness should be treated as first-order design constraints, not as post-hoc validation criteria, in the development of energy-efficient monitoring algorithms.

8. Governance, Data Protection, and Lifecycle Sustainability

Continuous physiological monitoring generates sensitive personal data that are subject to stringent legal and ethical obligations. Data protection frameworks such as the European Union General Data Protection Regulation establish requirements for transparency, purpose limitation, data minimization, and user rights [18]. Energy-efficient algorithms intersect with these obligations in important ways. On-device processing that reduces data transmission can support data minimization by limiting the amount of raw physiological information that leaves the user's control. However, if adaptive sampling and compression are not documented clearly, users may not understand what data were collected and what clinical events might have been missed. Governance mechanisms must therefore ensure that energy-driven algorithm decisions are auditable and explainable. This requires metadata about sampling schedules, filtering parameters, and inference confidence, which itself consumes storage and energy, creating a need for careful metadata design.

Lifecycle sustainability extends beyond individual privacy to the environmental footprint of wearable devices. Green computing principles emphasize that energy-efficient operation is only one component of overall sustainability; manufacturing, disposal, software maintenance, and network infrastructure also contribute to environmental impact [19]. Prolonging battery life can reduce the frequency of charging and the rate of battery degradation, thereby extending device lifespan and reducing electronic waste. However, if energy-efficient algorithms reduce clinical utility and require more frequent device replacement or clinical follow-up, the net sustainability effect may be negative. Global health policy frameworks for mobile health suggest that digital monitoring technologies should be evaluated not only for technical performance but also for their feasibility and equity in diverse health systems [20]. Energy efficiency is therefore embedded in a broader policy question about how wearable monitoring can be scaled responsibly without increasing environmental harm or excluding resource-limited settings.

9. Conclusion

Energy-efficient algorithms for continuous physiological monitoring occupy a complex intersection of signal processing, embedded computing, clinical reliability, and social governance. This paper has argued that energy efficiency cannot be treated as an isolated optimization goal. It is a system-level property shaped by the interaction of acquisition schedules, motion artifact management, adaptive sampling, local inference, communication strategies, and data protection obligations. Architectural decisions such as where computation occurs, how often data are transmitted, and how devices enter low-power states have direct consequences for battery lifetime, clinical sensitivity, privacy, and fairness. Adaptive filtering techniques that recover physiological signals under motion make localized processing more viable, but their benefits depend on appropriate integration with broader resource management policies. The evolution of tiny machine learning and federated learning suggests

that more intelligence can reside on the device, yet these methods introduce new governance and sustainability challenges. Future research should focus on context-aware resource management frameworks that dynamically balance energy expenditure with clinical need and population-specific equity. These frameworks must be accompanied by transparent data governance, lifecycle assessment, and policy mechanisms that ensure continuous monitoring systems remain reliable, fair, and sustainable outside the laboratory.

References

1. 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.
2. Majumder, S., Mondal, T., & Deen, M. J. (2017). Wearable sensors for remote health monitoring. *Sensors*, 17(1), 130.
3. Benini, L., Bogliolo, A., & De Micheli, G. (2000). A survey of design techniques for system-level dynamic power management. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 8(3), 299-316.
4. 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(3), 485-505.
5. Polastre, J., Szewczyk, R., & Culler, D. (2005). Telos: Enabling ultra-low power wireless research. In *IPSN 2005: Fourth International Symposium on Information Processing in Sensor Networks* (pp. 364-369). IEEE.
6. Dixon, A. M. R., Allstot, E. G., Gangopadhyay, D., & Allstot, D. J. (2012). Compressed sensing system considerations for ECG and EMG wireless biosensors. *IEEE Transactions on Biomedical Circuits and Systems*, 6(2), 156-166.
7. Habibzadeh, H., Dinesh, K., Shishvan, O. R., Boggio-Dandry, A., Sharma, G., & Soyata, T. (2020). A survey of healthcare Internet-of-Things (HIoT): A clinical perspective. *IEEE Internet of Things Journal*, 7(1), 53-71.
8. Chen, M., Gonzalez, S., Vasilakos, A., Cao, H., & Leung, V. C. M. (2011). Body area networks: A survey. *Mobile Networks and Applications*, 16(2), 171-193.
9. Gia, T. N., Jiang, M., Rahmani, A. M., Westerlund, T., Liljeberg, P., & Tenhunen, H. (2015). Fog computing in healthcare Internet of Things: A case study on ECG feature extraction. In *2015 IEEE International Conference on Computer and Information Technology; Ubiquitous Computing and Communications; Dependable, Autonomic and Secure Computing; Pervasive Intelligence and Computing (CIT/IUCC/DASC/PICOM)* (pp. 356-363). IEEE.
10. Sinha, A., & Chandrakasan, A. (2001). Dynamic power management in wireless sensor networks. *IEEE Design & Test of Computers*, 18(2), 62-74.
11. 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.
12. 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.

13. Han, S., Mao, H., & Dally, W. J. (2015). Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. arXiv preprint arXiv:1510.00149.
14. Lin, J., Chen, W.-M., Lin, Y., Cohn, J., Gan, C., & Han, S. (2020). MCUNet: Tiny deep learning on IoT devices. *Advances in Neural Information Processing Systems*, 33, 11711-11722.
15. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, PMLR 54, 1273-1282.
16. Chen, I. Y., Pierson, E., Rose, S., Joshi, S., Ferryman, K., & Ghassemi, M. (2021). Ethical machine learning in healthcare. *Annual Review of Biomedical Data Science*, 4, 123-144.
17. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 2053951716679679.
18. European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union*, L119, 1-88.
19. Murugesan, S. (2008). Harnessing green IT: Principles and practices. *IT Professional*, 10(1), 24-33.
20. World Health Organization. (2011). *mHealth: New horizons for health through mobile technologies: second global survey on eHealth*. World Health Organization.