

Motion Artifact Suppression in Photoplethysmography Using Adaptive Neural Networks

Jose Kelley

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

jose.work@clemson.edu

Abstract

Photoplethysmography has become a foundational sensing modality in wearable and clinical monitoring because it enables continuous estimation of cardiovascular and respiratory parameters through noninvasive optical interrogation of peripheral tissue. However, the clinical utility and real-world reliability of photoplethysmographic signals are severely limited by motion artifacts arising from physical activity, sensor displacement, and environmental vibration. Conventional filtering and signal decomposition methods offer partial mitigation, but they often fail under nonstationary motion and complex physiological variability. This paper presents a system-level analysis of adaptive neural network approaches for motion artifact suppression in photoplethysmography. Rather than concentrating solely on algorithmic novelty, the paper examines the architectural trade-offs between convolutional, recurrent, and encoder-decoder models, the infrastructure requirements for edge and cloud deployment, data governance and fairness considerations, and the long-term sustainability of learning-based physiological monitoring systems. The discussion integrates perspectives from wearable computing, biomedical signal processing, machine learning, and healthcare policy to demonstrate that motion artifact suppression should be treated not as an isolated signal processing problem but as a socio-technical challenge requiring robust data pipelines, representative training data, transparent validation, and regulatory alignment. The analysis further considers how adaptive neural models can be embedded into resource-constrained devices while maintaining clinical reliability and equitable performance across diverse patient populations.

Keywords

photoplethysmography; motion artifacts; adaptive neural networks; wearable systems; health monitoring; socio-technical infrastructure.

1. Introduction

Photoplethysmography is a low-cost optical technique that measures blood volume changes in the microvascular bed through light absorption and reflection. It has become widely used in pulse oximeters, wrist-worn devices, and remote patient monitoring systems because it supports continuous estimation of heart rate, oxygen saturation, respiratory rate, and vascular dynamics without requiring invasive instrumentation [1]. The expansion of wearable sensing platforms has increased the volume and diversity of photoplethysmographic data collected outside controlled clinical settings, but this expansion has also exposed fundamental limitations in signal quality and interpretability [2]. Among these limitations, motion artifact contamination is the most persistent barrier to reliable physiological inference.

Motion artifacts originate from mechanical disturbances that alter the optical path between the sensor and the tissue, including relative displacement, skin deformation, muscle contraction, and ambient light variation [3]. These disturbances produce large amplitude fluctuations that overlap with the frequency band of the physiological signal, making linear filtering and fixed-frequency decomposition insufficient. In response, researchers have proposed a range of adaptive signal processing methods, including time-frequency decomposition, spectral tracking, and blind source separation, but these approaches frequently rely on assumptions about stationarity or artifact periodicity that do not hold during free-living activity [4]. More recently, data-driven neural network models have shown considerable promise because they can learn complex temporal and spatial patterns directly from large datasets of corrupted and reference signals [5]. The present paper addresses this shift from model-based filtering to adaptive neural computation by examining the broader system architecture, governance, deployment, and policy dimensions that determine whether such models can translate into robust clinical and consumer health technologies.

The central argument is that motion artifact suppression in photoplethysmography should be understood as a systems problem rather than an isolated algorithmic task. An effective solution must balance model expressiveness with computational efficiency, integrate multimodal sensor context, ensure representative training data, preserve patient privacy, and demonstrate fairness across demographic and physiological differences. This paper develops that argument through a structured review of architectural options, infrastructure constraints, data governance challenges, deployment sustainability, and regulatory considerations. By connecting these domains, the discussion aims to inform both technical researchers and health system designers about the structural trade-offs that shape the practical impact of adaptive neural networks in wearable photoplethysmography.

2. The Motion Artifact Challenge in Photoplethysmography

The photoplethysmographic waveform contains a slowly varying direct current component related to tissue absorption and a pulsatile alternating current component related to cardiac synchronous blood volume changes. This waveform is highly sensitive to probe geometry, tissue perfusion, skin pigmentation, and venous pressure, all of which can vary with posture and movement. Under rest conditions, the morphology of the pulse wave is relatively stable, and conventional peak detection algorithms can estimate heart rate with acceptable accuracy. Under physical activity, however, the signal-to-noise ratio deteriorates rapidly because motion-induced changes in tissue compression and sensor contact modulate the optical signal at amplitudes that often exceed the physiological pulsatile component [6]. The resulting artifact is nonstationary and spectrally complex, frequently overlapping the heart rate frequency range and producing spurious peaks that can be mistaken for cardiac cycles.

Traditional approaches to motion artifact suppression include adaptive filtering using accelerometer reference signals, independent component analysis, singular spectrum analysis, and time-frequency masking [7]. Some systems use a wrist-worn accelerometer to estimate the motion component and subtract it from the photoplethysmographic signal, but the relationship between acceleration and optical artifact is nonlinear and time-varying. Spectral filtering methods attempt to reconstruct the heart rate trajectory by tracking dominant frequency components over time [8]. Although these methods can work well in controlled protocols, they often require manual parameter tuning and degrade when motion intensity changes rapidly or when the heart rate is close to the cadence of repeated movement [9]. The limitations of classical methods have motivated the exploration of learning-based approaches

that can adapt internal representations to the specific characteristics of individual users and activity contexts.

The data acquisition context also influences the nature of motion artifacts. In clinical settings, patients may be relatively stationary, but sensor repositioning, bed movement, and shivering can still introduce transient disturbances. In ambulatory settings, walking, running, typing, and household tasks generate diverse motion profiles that are difficult to capture in a single training distribution. Moreover, respiratory effects, vasomotion, and low-frequency oscillations can be confused with motion artifacts when only the optical signal is available. This confusion is particularly relevant when attempting to extract respiratory rate from photoplethysmography, because respiratory modulation is subtle and can be obliterated by movement [10]. A robust suppression system must therefore understand not only the artifact itself but also the physiological signal components that co-occur with the artifact and the contextual information available from companion sensors.

3. Adaptive Neural Network Architectures for Artifact Suppression

Neural network architectures for photoplethysmographic processing have evolved from simple multilayer perceptrons to specialized temporal and spatial models. Convolutional neural networks are attractive because they can learn invariant features from raw or spectrogram-based representations of the signal. Recurrent neural networks, particularly long short-term memory networks, are suited to capturing long-range temporal dependencies and sequence-to-sequence mapping from corrupted input to clean output [11]. Encoder-decoder structures with skip connections, such as the U-Net architecture originally developed for biomedical image segmentation, have been adapted for signal denoising because they preserve fine-grained waveform structure while removing large-scale disturbances [12]. These architectures can be trained to map noisy photoplethysmographic segments to clean reference segments acquired simultaneously from electrocardiography or a motion-stable sensor.

The key advantage of neural network models is their capacity to learn nonlinear mappings that are difficult to specify analytically. In photoplethysmography, the relationship between accelerometer signals, optical signals, and artifact components is not fixed; it depends on sensor placement, skin properties, motion direction, and wearable mechanics. A neural model can learn these relationships from data and adapt its internal parameters to different operating conditions [13]. Some architectures incorporate attention mechanisms that weight the relevance of different time points or sensor channels, allowing the model to focus on reliable segments and ignore heavily corrupted intervals. This adaptive behavior is essential because motion artifacts are intermittent and vary in severity across a recording.

However, architectural expressiveness introduces important system-level trade-offs. Deeper and wider networks may achieve lower reconstruction error on benchmark datasets, but they also demand more memory, energy, and computational throughput. These demands conflict with the resource constraints of wearable devices, where battery life, heat dissipation, and real-time operation are critical. Models that require high-resolution input or multi-second context windows may introduce latency that is unacceptable for continuous monitoring or closed-loop interventions. Researchers have therefore explored lightweight convolutional designs, temporal convolutional networks, and recurrent models with reduced hidden states to balance accuracy and efficiency [5][6][7]. The architectural choice cannot be separated from the target deployment environment, the required update frequency, and the availability of cloud connectivity.

Multimodal fusion is another structural consideration. Many wearable systems include accelerometers, gyroscopes, and sometimes electrocardiographic or pressure sensors. Incorporating these signals into the neural network can improve artifact suppression by providing an explicit representation of motion. Fusion can occur at the input level, through concatenated feature vectors, or at intermediate layers, where modality-specific encoders produce a shared latent representation [14]. The design of fusion pathways affects both the robustness of the model and the complexity of the data acquisition infrastructure. A tightly coupled multimodal model may fail gracefully when one sensor is missing or degraded, so architectures must include mechanisms for handling missing modalities. This requirement is particularly relevant in real-world deployment, where sensor availability and sampling rates may vary across device generations and manufacturers.

4. System-Level Design Considerations and Infrastructure

Deploying adaptive neural networks for photoplethysmography requires more than a trained model. The surrounding infrastructure must support data acquisition, signal synchronization, model inference, firmware updates, and quality assessment. Wearable systems have long emphasized low-power sensing and wireless communication, but the addition of deep learning changes the balance between local computation and cloud offloading [15]. On-device inference reduces privacy risks and network dependence but imposes strict limits on model size and arithmetic precision. Cloud inference enables larger models and centralized monitoring but introduces latency, bandwidth constraints, and concerns about continuous transmission of physiological data. Hybrid architectures may perform lightweight artifact detection at the edge and reserve more complex reconstruction for the cloud when connectivity permits.

The quality of the output from an adaptive model depends heavily on the quality of the input data pipeline. Photoplethysmographic signals must be sampled consistently, timestamped, and synchronized with motion references. Missing samples, clock drift between sensors, and variable dynamic range can undermine even the most sophisticated neural network. Real-time quality assessment layers are therefore essential to flag corrupted segments, suppress unreliable predictions, and trigger sensor repositioning or recalibration [20]. These layers often rely on simple morphological features or lightweight classifiers that can operate at the sampling rate. In a well-designed system, the artifact suppression engine is embedded within a larger framework that includes signal quality indexing, contextual activity recognition, and confidence estimation.

The physical infrastructure of the wearable device also influences artifact characteristics. Sensor placement on the wrist, fingertip, ear, or chest changes the mechanical coupling between the body and the optical sensor. Wrist-worn devices are more susceptible to arm swing and rotational motion, while chest-worn devices experience respiratory movement and torso rotation. The optical components, including light-emitting diode wavelength, photodetector geometry, and skin contact pressure, all affect the signal-to-noise ratio and the nature of motion artifacts [21]. Adaptive neural networks trained on one form factor may not transfer directly to another because the input distribution shifts with sensor design. System designers must therefore consider domain adaptation, transfer learning, and device-specific fine-tuning as part of the deployment strategy.

Scalability and maintenance also shape the long-term viability of adaptive photoplethysmographic monitoring. Cloud-based systems can aggregate data across thousands of devices to improve model training, but they require consistent data governance

and secure storage. Federated learning offers an alternative that allows models to be updated using decentralized data without transmitting raw physiological waveforms to a central server. This approach improves privacy and reduces communication overhead but introduces additional complexity in aggregation, versioning, and evaluation. The choice between centralized and federated training is not merely technical; it is also a governance decision that affects accountability, auditability, and the ability to detect performance drift across subpopulations.

5. Data Governance, Robustness, and Fairness

The data used to train motion artifact suppression models determine how well those models perform under real-world conditions. Many publicly available photoplethysmography datasets contain recordings from healthy young adults under controlled movement protocols. This narrow distribution limits generalization to older adults, patients with cardiovascular disease, individuals with darker skin tones, and users with peripheral vascular conditions. Because optical absorption depends on melanin concentration and tissue composition, models trained on homogeneous populations may exhibit systematically higher error rates for underrepresented groups [2]. This performance disparity is not merely a technical inconvenience; it can lead to inaccurate vital sign estimation and inappropriate clinical decisions if the technology is deployed without fairness auditing.

Robustness must be evaluated across a range of motion intensities, sensor attachment conditions, and environmental factors. A model that performs well on treadmill walking may fail during cycling, typing, or sleep. Furthermore, the presence of arrhythmias or low perfusion states can alter the waveform morphology in ways that motion artifact suppression networks may misinterpret as noise. The interaction between physiological abnormality and motion artifact is particularly important for clinical applications, where the very populations most in need of monitoring may also present atypical signal patterns [13]. Robustness therefore requires training data that include clinical populations and pathophysiological variability, not only healthy volunteers.

Privacy and consent are central governance concerns. Photoplethysmographic signals can contain information about cardiovascular health, stress, and potentially individual identity. Continuous collection of such data raises risks related to re-identification, secondary use, and surveillance. Even when raw waveforms are not transmitted, features extracted by neural networks may leak sensitive information. Ethical frameworks for algorithmic systems emphasize the need for transparency, purpose limitation, and meaningful consent [24]. In practice, this means that data collection for model training must be governed by clear protocols, and users must be informed about how their physiological data will be used, stored, and shared. The governance infrastructure should also address data deletion rights and the conditions under which models may be updated or replaced.

Algorithmic fairness in health monitoring is closely tied to representation and evaluation. If training datasets do not include sufficient diversity, model performance will reflect the biases embedded in those datasets. Fairness cannot be achieved solely by collecting more data; it requires deliberate sampling strategies, subgroup analysis, and transparent reporting of performance across demographic and clinical categories. The regulatory environment increasingly expects manufacturers to demonstrate that machine learning medical devices are safe and effective for their intended populations [25]. For photoplethysmography-based systems, this means evaluating motion artifact suppression not only in aggregate but also within subgroups defined by skin pigmentation, age, sex, body mass index, and disease status.

These evaluations should be built into the development lifecycle rather than performed retrospectively after deployment.

6. Deployment and Sustainability in Healthcare Environments

The transition from a validated algorithm to a sustainable healthcare deployment involves integration with clinical workflows, electronic health records, and existing monitoring infrastructure. A motion artifact suppression model may perform well in isolation but fail when embedded in a clinical environment where sensor placement, nurse workflows, and network conditions differ from the training context. Hospitals are dynamic environments with electromagnetic interference, frequent patient movement, and diverse clinical conditions. The model must therefore be paired with a robust device management layer, calibration procedures, and clinical decision support that interprets the output in context [1]. Without such integration, the technology may generate additional workload rather than reduce clinical burden.

Sustainability also involves the environmental and economic costs of data processing. Large-scale neural network training consumes energy and computational resources, and continuous inference on millions of wearable devices can have significant aggregate energy demand. Model compression, quantization, and pruning can reduce the carbon footprint of inference, while hardware acceleration and low-power microcontrollers can extend battery life. These efficiency measures are not secondary concerns; they influence whether adaptive neural networks can be deployed in low-resource settings where reliable electricity and connectivity are limited [21]. A sustainable design considers the entire lifecycle of the device, from manufacturing and training to operation and eventual disposal.

Interoperability is another critical deployment factor. Photoplethysmographic devices from different manufacturers use different wavelengths, sampling rates, and signal processing pipelines. A motion artifact suppression model trained on one device may produce unreliable results when applied to another. Standardized data formats, open benchmarks, and shared evaluation protocols can improve interoperability, but they also require coordination among manufacturers, researchers, and regulators [10]. The development of robust reference datasets with synchronized electrocardiography and motion signals is essential for benchmarking and for ensuring that reported performance metrics are meaningful across studies.

Real-time rapid vital sign extraction has become a focus of wearable health research because clinical and consumer applications increasingly require immediate feedback under varying activity levels. Systems that can estimate heart rate, respiratory rate, and oxygen saturation in real time while suppressing motion artifacts are more useful for postoperative monitoring, rehabilitation, and remote triage [18]. However, rapid extraction must not compromise accuracy. Deploying such systems requires careful attention to algorithmic latency, buffering strategies, and alert thresholding. A false alarm caused by residual motion artifact can lead to alarm fatigue, while a missed event can delay necessary intervention. The deployment strategy must therefore include mechanisms for uncertainty communication, context-aware alarm suppression, and clinician override.

7. Policy Implications and Clinical Translation

The regulatory landscape for adaptive neural networks in medical devices is evolving. In many jurisdictions, machine learning models that perform physiological monitoring may be classified as medical devices, particularly when they are used for diagnosis, treatment, or clinical decision support. Manufacturers must provide evidence of safety and effectiveness,

which includes demonstrating that motion artifact suppression does not obscure clinically significant physiological events. The challenge is that adaptive models can change over time if they are continuously updated, requiring new regulatory approaches that accommodate learning systems without sacrificing patient safety [25]. Predetermined change control plans, post-market surveillance, and algorithmic impact assessments are increasingly discussed as mechanisms for governing adaptive software.

Clinical translation also depends on the ability to communicate model outputs in ways that are meaningful to healthcare providers. A pulse rate estimate accompanied by a confidence score derived from artifact suppression performance can help clinicians decide whether to trust the value or obtain an alternative measurement. However, many wearable systems currently present only a single smoothed number, obscuring the uncertainty introduced by motion. Policy and design standards should encourage the reporting of signal quality and artifact burden alongside physiological estimates. This would support more informed clinical interpretation and reduce the risk of overreliance on imperfect sensor data [24].

Equity is a policy issue as much as a technical issue. If adaptive neural network systems are trained and validated primarily on data from affluent, healthy, and light-skinned populations, they may widen health disparities by providing less reliable monitoring for those who already face barriers to care. Reimbursement policies, procurement decisions, and regulatory guidance can incentivize manufacturers to address these disparities by requiring subgroup performance evidence. Government health systems and private insurers can also play a role by preferring devices that meet fairness benchmarks and by funding independent validation studies in diverse populations [20]. These mechanisms are necessary because market forces alone may not reward investment in equity.

The future of motion artifact suppression will likely involve tighter coupling between physiological sensing, context awareness, and personalization. Adaptive models may learn individual-specific characteristics of skin, tissue, and movement patterns over time, improving accuracy while also raising new questions about data retention and user autonomy. Personalized models can blur the boundary between a standardized medical device and a continuously evolving system tailored to a single user. Regulatory frameworks will need to address version control, explainability, and the right to inspect or challenge algorithmic decisions. In this context, photoplethysmography serves as a useful case study for the broader challenges of deploying learning-based systems in health care, because it combines a simple sensor modality with complex physiological variability and significant public health potential.

8. Conclusion

Motion artifact suppression in photoplethysmography has moved from deterministic filtering toward adaptive neural network architectures that can learn complex artifact patterns from multimodal physiological data. This evolution offers substantial potential for improving the reliability of wearable health monitoring in real-world conditions, but it also introduces structural trade-offs that extend beyond algorithm design. The choice of network architecture, the placement of computation between edge and cloud, the quality and representativeness of training data, the governance of physiological information, and the integration of uncertainty into clinical decision support are all central to the success of these systems. A system-level perspective shows that the most accurate model is not necessarily the most deployable or equitable model. Instead, sustainable and robust photoplethysmographic monitoring requires coordinated attention to sensor infrastructure, data governance, model fairness, energy efficiency, and regulatory compliance. As wearable sensing becomes more deeply embedded

in clinical and public health systems, adaptive neural networks will play an important role, but only if they are designed, evaluated, and governed within a broader socio-technical framework that prioritizes patient safety, equity, and long-term system resilience.

References

1. Allen, J. (2007). Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3), R1–R39. <https://doi.org/10.1088/0967-3334/28/3/R01>
2. Elgendi, M. (2012). On the analysis of fingertip photoplethysmogram signals. *Current Cardiology Reviews*, 8(1), 14–25. <https://doi.org/10.2174/157340312801215782>
3. Asada, H. H., Shaltis, P., Reisner, A., Rhee, S., & Hutchinson, R. C. (2003). Mobile monitoring with wearable photoplethysmographic biosensors. *IEEE Engineering in Medicine and Biology Magazine*, 22(3), 28–40. <https://doi.org/10.1109/MEMB.2003.1213624>
4. Zhang, Z., Pi, Z., & Liu, B. (2015). TROIKA: A general framework for heart rate monitoring using wrist-type photoplethysmographic signals during intensive physical exercise. *IEEE Transactions on Biomedical Engineering*, 62(2), 522–531. <https://doi.org/10.1109/TBME.2014.2359372>
5. Temko, A. (2017). Accurate heart rate monitoring during physical exercises using PPG. *IEEE Transactions on Biomedical Engineering*, 64(9), 2015–2024. <https://doi.org/10.1109/TBME.2017.2676243>
6. Reiss, A., Indlekofer, I., Schmidt, P., & Van Laerhoven, K. (2019). Deep PPG: Large-scale heart rate estimation with convolutional neural networks. *Sensors*, 19(14), 3079. <https://doi.org/10.3390/s19143079>
7. Biswas, D., Everson, L., Liu, M., Panwar, M., Van Helleputte, N., & Van Hoof, C. (2019). CorNET: Deep learning framework for PPG-based heart rate estimation and biometric identification. *IEEE Journal of Biomedical and Health Informatics*, 23(3), 1014–1022. <https://doi.org/10.1109/JBHI.2018.2872010>
8. Salehizadeh, S. M. A., Dao, D., Bolkhovskiy, J., Cho, C., Mendelson, Y., & Chon, K. H. (2016). A novel time-varying spectral filtering algorithm for reconstruction of motion artifact corrupted heart rate signals. *IEEE Transactions on Biomedical Engineering*, 63(5), 955–966. <https://doi.org/10.1109/TBME.2015.2480132>
9. Jarchi, D., & Casson, A. J. (2017). Towards photoplethysmography-based estimation of instantaneous heart rate during physical activity. *IEEE Transactions on Biomedical Engineering*, 64(10), 2324–2334. <https://doi.org/10.1109/TBME.2017.2665063>
10. Charlton, P. H., Birrenkott, D. A., Pimentel, M. A. F., Johnson, A. E. W., Alastruey, J., Ward, S., ... & Kyriacou, P. A. (2018). Breathing rate estimation from the electrocardiogram and photoplethysmogram: A review. *IEEE Reviews in Biomedical Engineering*, 11, 2–20. <https://doi.org/10.1109/RBME.2017.2763681>
11. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
12. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted*

Intervention – MICCAI 2015 (pp. 234–241). Springer. https://doi.org/10.1007/978-3-319-24574-4_28

13. 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. <https://doi.org/10.1038/s41591-018-0268-3>
14. Vullings, R., de Vries, B., & Bergmans, J. W. M. (2011). An adaptive Kalman filter for ECG signal enhancement. *IEEE Transactions on Biomedical Engineering*, 58(4), 1094–1103. <https://doi.org/10.1109/TBME.2010.2099226>
15. Majumder, S., Mondal, T., & Deen, M. J. (2017). Wearable sensors for remote health monitoring. *Sensors*, 17(1), 130. <https://doi.org/10.3390/s17010130>
16. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
17. Chen, W., Jaques, N., Taylor, S., Sano, A., Fedor, S., & Picard, R. W. (2015). Wavelet-based motion artifact removal for electrodermal activity. In *2015 IEEE 12th International Conference on Wearable and Implantable Body Sensor Networks (BSN)* (pp. 1–6). IEEE. <https://doi.org/10.1109/BSN.2015.7299391>
18. Zheng, X., Dwyer, V. M., Barrett, L. A., Derakhshani, M., & Hu, S. (2023). Rapid vital sign extraction for real-time opto-physiological monitoring at varying physical activity intensity levels. *IEEE Journal of Biomedical and Health Informatics*, 27(7), 3107–3118.
19. Liang, Y., Elgendi, M., Chen, Z., & Ward, R. (2018). An optimal filter for short photoplethysmogram signals. *Scientific Data*, 5, 180076. <https://doi.org/10.1038/sdata.2018.76>
20. Naeini, E. K., Azimi, I., Rahmani, A. M., Liljeberg, P., & Dutt, N. (2020). A real-time PPG quality assessment approach for IoT-based smart health. *IEEE Transactions on Biomedical Circuits and Systems*, 14(5), 1184–1196. <https://doi.org/10.1109/TBCAS.2020.3013845>
21. Heikenfeld, J., Jajack, A., Rogers, J., Gutruf, P., Tian, L., Pan, T., Li, R., ... & Wang, J. (2018). Wearable sensors: Modalities, challenges, and prospects. *Lab on a Chip*, 18(2), 217–248. <https://doi.org/10.1039/C7LC00914C>
22. Vollmer, M., Bläsing, D., & Kaderali, L. (2020). Alignment of multi-sensor data for robust estimation of respiratory rate from wrist-worn devices. *IEEE Journal of Biomedical and Health Informatics*, 24(11), 3161–3171. <https://doi.org/10.1109/JBHI.2020.2984741>
23. Poh, M. Z., McDuff, D. J., & Picard, R. W. (2010). Non-contact, automated cardiac pulse measurements using video imaging and blind source separation. *Optics Express*, 18(10), 10762–10774. <https://doi.org/10.1364/OE.18.010762>
24. 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. <https://doi.org/10.1177/2053951716679679>
25. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56. <https://doi.org/10.1038/s41591-018-0300-7>