

Reinforcement Learning-Based Adaptive Noise Suppression for Real-Time Biomedical Signal Processing

Florian Wells

Department of Computer Science, Binghamton University, Binghamton, NY, USA.
wellsflorian@binghamton.edu

Jingtian Sheng

Department of Computer Science, George Mason University, Fairfax, VA, USA.
jing.shao@gmu.edu

Wenhaozhi Yu

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
wenhaozhi.yu517@colostate.edu

Janis Lawson

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
janis.lawson804@uab.edu

Abstract

Reinforcement learning offers a principled approach for adapting noise suppression policies in real-time biomedical signal processing. This paper presents a system-level analysis of adaptive noise suppression architectures that integrate reinforcement learning with conventional signal conditioning, physiological monitoring, and safety supervision. We examine the structural trade-offs between rapid adaptation and stability, between edge autonomy and centralized oversight, and between learning flexibility and clinical interpretability. The discussion emphasizes the design of acquisition pipelines, state representation, action constraints, reward shaping, and runtime safety monitors. Rather than proposing a single algorithm, the paper develops a governance-oriented framework for designing and deploying learning-based noise suppression systems in clinical and ambulatory settings. The analysis covers robustness to sensor degradation, motion artifact, electrode impedance changes, and population heterogeneity. We further address fairness, accountability, sustainability, and regulatory policy implications. The treatment is intentionally at the system level, connecting control-theoretic filtering, deep learning, clinical deployment, and infrastructure governance. It is argued that reinforcement learning can improve noise suppression only when embedded within a broader socio-technical architecture that includes transparent fallback mechanisms, continuous evaluation, and clear regulatory pathways. This perspective provides a foundation for subsequent empirical and clinical validation of real-time learning-based biomedical signal enhancement.

Keywords

reinforcement learning, adaptive noise suppression, biomedical signal processing, real-time systems, clinical AI governance.

1. Introduction

Biomedical signal processing underpins modern monitoring and diagnostic technologies. Electrocardiographic, electroencephalographic, and photoplethysmographic signals are inherently low amplitude and susceptible to multiple noise sources, including baseline wander, power line interference, muscle activity, sensor displacement, and motion artifact. Traditional noise suppression methods often assume fixed or slowly varying noise statistics. Such assumptions fail in ambulatory care, intensive care units, and wearable health monitoring, where noise properties change rapidly and signal morphology must be preserved for clinical interpretation. Reinforcement learning provides a conceptual and computational alternative by framing adaptive filtering as a sequential decision problem. The learning agent continuously selects filter configurations or suppression actions based on observed signal features and receives feedback from signal quality and clinically relevant objectives [1]. Deep reinforcement learning extends this approach through neural network function approximation, enabling complex mappings from high-dimensional biomedical inputs to control actions [2]. Similar strategies have been shown to handle complex sequential decisions in domains with large state spaces and delayed rewards [3]. Continuous control algorithms are especially relevant because noise suppression often requires smooth adjustment of filter coefficients rather than discrete mode switches [4]. These advances create opportunities for real-time systems that learn to suppress noise while respecting physiological constraints.

The significance of a system-level perspective lies in the gap between algorithmic success and clinical utility. Learning-based noise suppression cannot be treated as an isolated filter component. It is embedded in signal acquisition chains, real-time inference pipelines, alarm generation modules, and clinical decision support systems. Deep learning provides powerful representational capacities, but its deployment in biomedical infrastructure raises distinctive constraints related to latency, interpretability, failure behavior, and regulatory acceptance [5]. The present paper therefore develops a governance-oriented analysis of reinforcement learning-based adaptive noise suppression. It examines structural trade-offs across acquisition, feature representation, policy learning, runtime safety monitoring, data infrastructure, and deployment sustainability. The goal is not to advocate a specific model architecture, but to clarify how system design choices shape robustness, fairness, and accountability. This framing is necessary because adaptive systems may improve average signal quality while introducing rare but clinically significant failure modes that are difficult to detect through conventional benchmarking. A system-level discussion can inform engineers, clinical informaticians, regulators, and healthcare administrators about how learning components can be introduced safely into real-time biomedical processing chains.

2. Background and Related Work

Biomedical signal processing has long relied on adaptive filtering to address nonstationary interference. Deep learning has recently expanded the representational tools available for learning complex noise and signal transformations [6]. At the same time, conventional adaptive methods such as adaptive noise cancellation, least mean squares filters, recursive least squares filters, and notch filters remain widely used in clinical devices because of their interpretability and computational efficiency [7]. The foundational adaptive noise cancellation architecture separates a primary input from a reference noise input and adjusts filter coefficients to minimize residual interference [8]. Established adaptive filter theory provides convergence and stability conditions under certain statistical assumptions [9]. However, these methods often fail when noise statistics change abruptly, when the reference channel contains

signal leakage, or when motion artifact is poorly correlated with accelerometer-based references. Such conditions are common in wearable and bedside monitoring, especially during patient transport, exercise, or intensive care procedures.

The structural diversity of biomedical signals further complicates noise suppression. Cardiac and neurological signals differ in bandwidth, amplitude, electrode interface, and artifact morphology [10]. Clinical electrocardiographic analysis requires preservation of subtle features such as ST-segment deviations and T-wave alternans, while electroencephalographic analysis must avoid distorting transient epileptiform patterns. Advanced signal processing methods have therefore moved beyond simple frequency filtering to incorporate nonlinear transforms, independent component analysis, and signal quality indices [11]. Reinforcement learning introduces a different perspective by learning policies that map observed signal context to filtering actions. Universal value function approximators enable a single policy to handle multiple related goals, such as reducing motion artifact while preserving R-peak amplitude or minimizing baseline wander without attenuating low-frequency ST-segment features [12]. This is conceptually attractive for biomedical noise suppression because different clinical scenarios may require different trade-offs rather than a uniform error minimization target.

The translation of reinforcement learning from controlled benchmarks to biomedical systems is complicated by sensitivity to hyperparameters, reward scale, and environment design. Reproducibility studies show that reinforcement learning performance is highly dependent on implementation details and evaluation protocols [13]. In biomedical settings, this sensitivity has direct safety implications. A noise suppression policy that performs well on one dataset may behave unpredictably on another patient population, sensor placement, or motion profile. Consequently, system-level validation must include robustness testing under distribution shift, prospective monitoring, and fallback comparisons against established adaptive filters. This need for rigorous evaluation connects algorithm design to infrastructural governance.

3. System Architecture

A clinically realistic reinforcement learning-based noise suppression system consists of several interacting components: signal acquisition, streaming feature extraction, policy inference, adjustable filtering, safety supervision, and data logging. The acquisition front end must provide synchronized signals from physiological sensors and auxiliary references such as accelerometers or electrode-tissue impedance. Feature extraction compresses raw waveforms into interpretable descriptors of noise level, signal quality, heart rate, motion intensity, and morphology. The policy then selects filter parameters or suppression actions. An adjustable filter applies these actions, and a safety monitor evaluates whether the filtered signal remains within physiologically plausible bounds. This modular separation is important because it allows conventional filters to remain active as a safe default while the learning component adjusts behavior only when sufficient confidence exists. Real-world reinforcement learning systems face challenges related to delayed feedback, partial observability, and nonstationary environments [14]. In biomedical applications these challenges are compounded by noise, inter-patient variability, and the need for deterministic worst-case safety guarantees.

The architecture must also address trade-offs between edge autonomy and centralized oversight. Wearable monitors often have limited power, memory, and communication bandwidth, making cloud-based policy inference undesirable for life-critical real-time processing. Edge deployment reduces latency and privacy exposure but restricts the

complexity of the learned policy. A hybrid architecture may train policies centrally on curated datasets and deploy lightweight inference engines on the device, with periodic encrypted updates and local adaptation. Runtime safety monitors should enforce maximum filter depth, saturation thresholds, and signal quality gating to prevent malformed outputs from reaching diagnostic or alarm modules. Artificial intelligence safety research emphasizes the importance of avoiding negative side effects and robustly handling distributional shift [15]. In noise suppression, negative side effects include introducing artificial morphology changes that could be misinterpreted as pathological events. Therefore architectural design should treat the reinforcement learning policy not as an autonomous controller but as an advisory layer within a constrained filtering stack.

4. Reinforcement Learning as a Sequential Control Problem

The noise suppression problem can be conceptualized as a continuous control process in which the system observes the current signal context, selects an adjustment to filtering parameters, and receives a reward based on signal quality improvement and preservation of clinically meaningful features. The state representation may include short-term spectral features, motion context, signal quality estimates, and previous filter settings. The action may be a vector of adjustment levels for cutoff frequencies, notch bandwidths, or adaptive step sizes. The reward combines multiple objectives, including reduction of artifact power, maintenance of waveform morphology, and avoidance of abrupt changes. This formulation connects clinical signal improvement to the broader medical artificial intelligence agenda, in which machine learning models support rather than replace clinical interpretation [16]. In practice, reward design is a central governance challenge because poorly chosen rewards may produce filters that aggressively remove artifact at the cost of flattening diagnostically relevant morphology. Recent work on adaptive notch filtration in photoplethysmographic signals during physical activity illustrates the value of targeted artifact suppression that preserves physiological waveform features [17]. This baseline demonstrates that frequency-adaptive filtering can improve signal recovery, but a learning-based layer can adapt more flexibly across movement states.

A system-level formulation must also define the temporal horizon. Short-sighted suppression may overreact to transient artifact, while long-horizon optimization may delay stable recovery. The choice of horizon affects stability and monitoring complexity. In real-time monitoring, the system should operate with a receding window and bounded action magnitudes. Safety constraints can be encoded through reward penalties, action clipping, or external monitors. However, encoding all clinical constraints in reward functions is difficult. Therefore, the policy should be embedded within a hierarchical control framework in which a deterministic physiological plausibility monitor can override learned actions. This structure preserves learning benefits while maintaining a guaranteed fallback to conventional filtering. Beyond algorithmic design, the system must log the contextual state, policy action, filter response, and safety interventions. Such logs support retrospective review and continuous model governance, which are essential for regulated medical environments.

5. Data and Signal Infrastructure

Learning-based noise suppression requires representative data that capture physiological signal morphology, noise diversity, and clinical contexts. High-performance medical artificial intelligence depends on high-quality labeled datasets and clinically grounded evaluation [18]. For adaptive filtering, labels may include clean reference signals, noise annotations, or downstream diagnostic outcomes. Synthetic noise models can augment real data, but they

must reflect realistic electrode motion, muscle artifact, and environmental interference. Deep learning in healthcare further requires careful attention to data provenance, annotation reliability, and hidden stratification [19]. A robust infrastructure should include protected data lakes, versioned annotation pipelines, and reproducible preprocessing workflows. Without such infrastructure, learned policies may inherit annotation errors and fail silently when deployed on patients with underrepresented characteristics.

Data governance for adaptive noise suppression is not solely a technical issue. Responsible machine learning for health care requires clear accountability for how training data are collected, how performance is measured, and how models are updated after deployment [20]. Federated learning offers a promising infrastructure for collaborative model development across institutions without centralizing sensitive physiological data [21]. However, federated approaches introduce additional complexities related to heterogeneity of devices, communication constraints, and fair aggregation across sites. In noise suppression, local data distributions may differ due to electrode type, skin preparation, patient movement patterns, and clinical environment. Federated training can improve generalization if aggregation is carefully managed, but it can also mask local performance failures. Therefore, each deployment site should maintain local evaluation capabilities and have authority to adjust monitoring thresholds based on its patient population.

6. Robustness and Safety

Robustness in adaptive noise suppression must be evaluated across multiple axes: sensor degradation, electrode-skin impedance changes, motion intensity, environmental interference, and patient-specific waveform morphology. A policy that is optimal for one sensor placement may perform poorly when the device shifts, when the patient is diaphoretic, or when multiple noise sources interact. Real-world reinforcement learning deployments frequently encounter distributional drift that invalidates learned assumptions [14]. In biomedical contexts, such drift can produce worse signal quality than a fixed conventional filter. Robustness therefore requires runtime detection of out-of-distribution states and automatic fallback. The safety monitor should compare learned policy outputs with a conventional adaptive filter and reject actions that produce implausible waveform changes. This is consistent with broader artificial intelligence safety principles that emphasize safe exploration, oversight, and minimizing negative side effects [15]. The system should also be tested against adversarial or rare physiological conditions, including arrhythmias, low-amplitude signals, and rapid motion artifact bursts. These conditions may not be adequately represented in training data, so systematic stress testing is necessary before clinical use.

Another dimension of robustness is temporal stability. Learned policies may respond to small input changes with large action changes, producing visible artifacts in monitored waveforms. Such instability can confuse clinicians and trigger false alarms. To mitigate this, temporal smoothing, hysteresis, and action magnitude constraints can be included in the system design. Furthermore, the policy should provide confidence estimates or uncertainty indicators that allow the policy to provide confidence estimates or uncertainty indicators that allow downstream modules to assess reliability before using filtered output. Confidence estimation is particularly important in clinical settings where a silent failure may be more harmful than a visible but interpretable artifact. Uncertainty-aware filtering can route ambiguous states to conventional filters, request additional signal segments, or reduce the influence of learning-based actions. Such mechanisms align with safety-critical control practices in which learned components are supervised by formal monitors and conservative fallbacks.

Robustness evaluation must also consider the interaction between learned suppression and downstream diagnostic algorithms. A waveform may appear visually smooth after aggressive filtering while losing subtle features that are essential for detecting ischemia, seizure onset, or respiratory variability. The system should therefore include morphology preservation metrics that go beyond generic signal-to-noise ratio. These metrics may quantify R-peak amplitude stability, ST-segment deviation, T-wave shape, or specific electroencephalographic band power. By incorporating these clinically oriented metrics into evaluation and reward design, the system can be aligned with diagnostic priorities rather than purely statistical noise reduction. The tension between artifact removal and morphology preservation is a recurring theme in biomedical signal processing; it cannot be resolved by algorithmic sophistication alone but requires explicit clinical specification of acceptable trade-offs. A governance board or clinical engineering committee should approve the quantitative bounds used for such trade-offs and review them periodically as clinical evidence evolves.

7. Fairness and Population Heterogeneity

Adaptive noise suppression systems may perform unevenly across patient subgroups because signal morphology, skin properties, motion patterns, and noise sources vary with age, body habitus, skin tone, underlying disease, and sensor placement. Machine learning algorithms in healthcare have demonstrated that models trained on aggregate data can embed and amplify disparities when performance is not evaluated within subpopulations [22]. In the context of physiological monitoring, underperformance may lead to missed alarms, misleading waveforms, or unnecessary clinical interventions. Fairness cannot be reduced to a single metric because clinical relevance depends on the downstream task, the patient population, and the deployment environment. A system may exhibit acceptable average error reduction while failing for patients with darker skin tones, low peripheral perfusion, or tremor. Therefore, fairness-aware evaluation should be integrated throughout the lifecycle, from training data composition to post-market surveillance.

Population heterogeneity also affects the definition of a clean reference signal. What is considered normal waveform morphology differs across age groups, athletic conditioning, medication use, and comorbidities. A learning-based filter trained predominantly on young healthy volunteers may penalize physiological variants as noise. This can lead to inappropriate filtering of low-amplitude QRS complexes or the loss of delta activity in older adults. To address this, training data should be stratified by demographic and clinical covariates, and performance reports should include subgroup analyses. Institutional review boards and medical device regulators increasingly expect evidence that machine learning components do not introduce or worsen disparities [23]. In noise suppression, this expectation translates into mandatory reporting of subgroup-specific signal quality metrics and the establishment of equity thresholds before clinical deployment. The system architecture must support collecting the necessary metadata without compromising patient privacy or introducing excessive documentation burden for clinicians.

8. Deployment and Sustainability

Sustainable deployment of learning-based noise suppression requires attention to computational cost, energy consumption, model updates, and lifecycle maintenance. Edge devices in wearable monitors often operate under strict battery and thermal constraints, and continuous policy inference may conflict with long-term wearability. Model compression, quantization, and hardware-aware design can reduce resource demand, but these techniques may also affect robustness and interpretability. The environmental footprint of training deep

reinforcement learning policies is nontrivial, especially when training data are large and hyperparameter search is extensive. Reporting energy and computational costs alongside clinical performance supports transparent decision-making about whether a learning-based approach provides sufficient benefit over a conventional adaptive filter [24]. A sustainable system should also include a clear decommissioning pathway when performance degrades or when new evidence shows that the learned policy is inferior to simpler alternatives.

Lifecycle management encompasses post-deployment monitoring, periodic recalibration, and secure software updates. Distribution drift may occur because of changes in sensor manufacturing, electrode materials, hospital protocols, or patient demographics. Continuous monitoring of signal quality outcomes can detect drift before it produces clinical harm. However, drift detection requires infrastructure for collecting and evaluating real-world data, which raises privacy and governance questions. Federated mechanisms can support distributed model improvement without transferring raw physiological waveforms, but they still require standardized data schemas, robust aggregation protocols, and local validation capabilities. The deployment architecture should allow practitioners to disable learning-based adaptation and rely on conventional filtering for sites or periods in which monitoring capacity is limited. This fallback is not an admission of failure but a recognition that adaptive systems must remain accountable to local operational realities.

9. Governance and Regulatory Policy

The introduction of reinforcement learning into biomedical signal processing raises governance challenges that extend beyond technical validation. Medical device regulators generally require predetermined performance specifications, traceable software changes, and evidence of safety under intended use. Continuously adapting policies complicate these requirements because the system behavior may change over time even without explicit software updates. One governance approach is to restrict online learning to a narrow set of bounded adjustments and to freeze policy parameters during clinical use, with periodic offline retraining and regulatory review. Another approach is to allow local adaptation but maintain a deterministic safety envelope that constrains the action space and can be inspected independently. The choice between these approaches depends on the risk classification of the device, the clinical use case, and the maturity of the learning infrastructure. High-risk monitoring applications should favor conservatism until empirical evidence demonstrates that learning-based adaptation improves outcomes without introducing unacceptable failure modes [25].

Regulatory policy must also clarify roles and responsibilities when adaptive systems behave unexpectedly. The hospital, the device manufacturer, the algorithm developer, and the integrating platform may all share responsibility for a single patient harm event. Transparent logging of policy actions, input features, fallback activations, and human overrides can support incident investigation and accountability. Such logs should be time-stamped, encrypted, and access-controlled, with retention periods defined by institutional policy and applicable law. Auditing adaptive systems is not merely a technical exercise; it requires governance structures that connect engineering teams, clinical safety officers, data protection specialists, and patient representatives. Without clear lines of authority and well-defined escalation procedures, adaptive noise suppression may create diffuse accountability that undermines trust in both the technology and the healthcare institution.

10. Conclusion

Reinforcement learning offers a promising framework for adaptive noise suppression in real-time biomedical signal processing, but its clinical value depends on system-level integration rather than isolated algorithmic performance. This paper has argued that learning-based filters should be embedded within modular architectures that combine conventional adaptive filtering, streaming feature extraction, safety monitoring, and fallback mechanisms. The structural trade-offs between edge autonomy and centralized oversight, between adaptation speed and stability, and between artifact reduction and morphology preservation must be addressed explicitly through design and governance. Data infrastructure, fairness evaluation, sustainability planning, and regulatory accountability are not peripheral concerns but core components of a responsible learning-based monitoring system. The discussion throughout has emphasized that reinforcement learning should function as an advisory layer within a constrained clinical filtering stack, rather than as an autonomous replacement for established physiological signal processing methods. Future work should move beyond small-scale demonstrations toward prospective clinical evaluations that measure subgroup-specific outcomes, failure modes, and lifecycle performance under realistic deployment conditions. Only through such rigorous and transparent evaluation can adaptive noise suppression realize its potential to improve the quality, equity, and reliability of continuous physiological monitoring.

References

1. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.
2. Arulkumaran, K., Deisenroth, M. P., Brundage, M., & Bharath, A. A. (2017). A brief survey of deep reinforcement learning. *IEEE Signal Processing Magazine*, 34(6), 26–38.
3. Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484–489.
4. Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., & Wierstra, D. (2016). Continuous control with deep reinforcement learning. *International Conference on Learning Representations*.
5. Topol, E. J. (2019). High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56.
6. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
7. Haykin, S. (2002). *Adaptive filter theory* (4th ed.). Prentice Hall.
8. Widrow, B., Glover, J. R., McCool, J. M., Kaunitz, J., Williams, C. S., Hearn, R. H., Zeidler, J. R., Dong, E., & Goodlin, R. C. (1975). Adaptive noise cancelling: Principles and applications. *Proceedings of the IEEE*, 63(12), 1692–1716.
9. Sayed, A. H. (2008). *Adaptive filters*. John Wiley & Sons.
10. Sörnmo, L., & Laguna, P. (2005). *Bioelectrical signal processing in cardiac and neurological applications*. Elsevier Academic Press.

11. Orphanidou, C., Bonnici, T., Charlton, P., Clifton, D., Vallance, D., & Tarassenko, L. (2015). Signal-quality indices for the electrocardiogram and photoplethysmogram: derivation and applications to wireless monitoring. *IEEE Journal of Biomedical and Health Informatics*, 19(3), 832–838.
12. Schaul, T., Horgan, D., Gregor, K., & Silver, D. (2015). Universal value function approximators. *Proceedings of the 32nd International Conference on Machine Learning*, 1312–1320.
13. Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., & Meger, D. (2018). Deep reinforcement learning that matters. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), 3207–3214.
14. Dulac-Arnold, G., Levine, N., Mankowitz, D. J., Li, J., Paduraru, C., Gowal, S., & Hester, T. (2021). Challenges of real-world reinforcement learning: Definitions, benchmarks and analysis. *Machine Learning*, 110(9), 2419–2468.
15. Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.
16. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *New England Journal of Medicine*, 380(14), 1347–1358.
17. 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.
18. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., & Dean, J. (2019). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24–29.
19. Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine*, 17, 195.
20. 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.
21. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. *npj Digital Medicine*, 3, 119.
22. 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.
23. Char, D. S., Shah, N. H., & Magnus, D. (2018). Implementing machine learning in health care—addressing ethical challenges. *New England Journal of Medicine*, 378(11), 981–983.
24. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650.

25. Gerke, S., Minssen, T., & Cohen, G. (2020). Ethical and legal challenges of artificial intelligence-driven healthcare. *Artificial Intelligence in Healthcare*, 295–336.