Machine Learning-Enabled Respiratory Rate Estimation from Wearable Optical Signals

Enzo Alvarez
Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.
enzo.alvarez@ku.edu

Krish Sinha
School of Computing, Clemson University, Clemson, SC, USA.
ksinha@clemson.edu

Jack C. Virtanen
Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
jack.work@uab.edu

Chengyu Ning
Department of Computer Science, University of Central Florida, Orlando, FL, USA.
ningc@ucf.edu

Abstract

Respiratory rate is a vital sign whose clinical significance is often underestimated, yet deterioration in respiratory status frequently precedes other indicators of physiological instability. Wearable optical sensors, particularly those based on photoplethysmography, offer a noninvasive and continuous pathway for respiratory monitoring outside clinical settings. However, extracting respiratory rate from wearable optical signals remains challenging because respiratory modulations are subtle and easily corrupted by motion, sensor displacement, low perfusion, and population heterogeneity. Machine learning has emerged as a powerful tool to address these challenges by learning robust representations from complex signal patterns. This paper presents a system-level examination of machine learning-enabled respiratory rate estimation from wearable optical signals. It considers the structural trade-offs associated with sensor placement, data infrastructure, edge and cloud computing, model selection, deployment, uncertainty management, and continuous monitoring. The discussion extends beyond algorithmic accuracy to include robustness, fairness, privacy, governance, and policy implications. Attention is given to public physiological datasets, federated learning, synthetic data, and responsible machine learning frameworks. The paper argues that long-term clinical and public health value depends on the integration of technical performance with operational accountability and regulatory readiness. Future directions include personalized adaptive models, open benchmarking under diverse conditions, and sustainable deployment architectures that balance clinical utility with resource consumption and data protection.

Keywords

respiratory rate estimation, wearable photoplethysmography, machine learning, edge computing, health equity, digital health governance, physiological monitoring.

1. Introduction

Wearable optical sensing has become a central modality in the expansion of continuous physiological monitoring beyond hospital walls. Photoplethysmography, which measures blood volume changes through light absorption or reflection in perfused tissue, provides a low-cost, noninvasive, and energy-efficient foundation for wrist-worn, finger-worn, earlobe, and chest-based devices [1]. Although photoplethysmography was initially associated with heart rate and oxygen saturation measurements, its signal contains additional physiological information, including respiratory-induced modulations that arise from mechanical, hemodynamic, and autonomic effects [2]. Respiratory rate is among the most clinically informative vital signs because it can indicate early deterioration in conditions such as sepsis, pneumonia, heart failure, and chronic obstructive pulmonary disease. Despite this importance, respiratory rate is frequently measured intermittently or inaccurately in busy clinical environments.

The continuous and unobtrusive estimation of respiratory rate from wearable optical signals has therefore attracted considerable research interest. Traditional extraction techniques based on time-domain filtering, frequency-domain analysis, and time-frequency decomposition can isolate respiratory components under stable conditions, but they often fail when signals are contaminated by motion artifacts or when respiratory patterns change rapidly [3]. Machine learning has been proposed as a means of learning complex nonlinear relationships from labeled physiological data, allowing the model to disentangle respiratory information from noise and to adapt to varied measurement contexts [4]. In parallel, advances in wearable technology and data-driven health care have expanded the possibilities for real-world deployment [5]. The convergence of high-performance artificial intelligence and clinical medicine has created expectations that algorithmic systems can support early detection and personalized care [6].

However, translating laboratory-level accuracy into dependable wearable respiratory monitoring requires more than a well-performing model. The problem is embedded within an end-to-end system that includes optical sensor design, data acquisition, signal quality management, preprocessing, inference, decision support, and user interaction. Each layer imposes constraints that shape model performance and clinical usefulness. This paper therefore adopts a system-level perspective. It examines the structural trade-offs that arise across sensor physiology, data infrastructure, machine learning architecture, deployment, robustness, fairness, and governance. The goal is not to present a single algorithmic contribution but to articulate the broader design considerations that determine whether machine learning-enabled respiratory rate estimation can become a trustworthy component of digital health infrastructure.

2. Wearable Optical Signal Foundations and Respiratory Modulation

Photoplethysmography operates by illuminating the skin and measuring the transmitted or reflected light that varies with pulsatile blood flow. The resulting waveform includes a fast cardiac component and slower respiratory-related fluctuations [1]. Respiratory modulation of the photoplethysmographic signal occurs through several interacting mechanisms. First, respiration alters intrathoracic pressure and venous return, which affects stroke volume and pulse amplitude. Second, chest and abdominal movement can introduce mechanical perturbations in the optical path, especially for reflective sensors worn on the wrist or torso. Third, respiratory sinus arrhythmia and autonomic regulation can modulate heart rate and vascular tone in synchrony with breathing [2]. These mechanisms are not equally present across all individuals, body positions, and activity levels. Their relative contributions can shift

with sensor location, skin temperature, hydration, and cardiovascular state, making respiratory rate estimation a moving target.

Because the respiratory signal is embedded within multiple photoplethysmographic features, robust estimation often benefits from combining amplitude modulation, frequency modulation, baseline wander, and pulse wave morphology features. Time-domain approaches may track breath-to-breath intervals in the amplitude envelope, while frequency-domain approaches search for a distinct respiratory peak in the signal spectrum. However, when motion artifacts overlap with the respiratory frequency band, spectral methods may produce misleading estimates. This has motivated the development of fusion approaches that combine multiple respiratory-induced modulations and assign confidence based on signal quality [7]. Multiparameter estimation strategies have shown that redundancy across photoplethysmographic features can reduce the influence of any single unreliable modulation [8]. Further robustness can be achieved by fusing photoplethysmography with electrocardiography, which provides complementary information about cardiorespiratory coupling [9].

These signal processing foundations reveal a central design principle for machine learning-enabled systems. The respiratory information available in wearable optical signals is inherently distributed across multiple weak and interacting components. No single feature is sufficient under all conditions. Therefore, model architectures and preprocessing pipelines should be designed to exploit physiological redundancy rather than relying on a narrow representation. At the same time, redundancy alone cannot overcome severe sensor corruption. Signal quality assessment must operate alongside the respiratory model so that low-quality periods can be identified, flagged, or excluded from clinical interpretation. This coupling between quality estimation and respiratory inference is a defining feature of dependable wearable monitoring systems.

3. System Architecture and Data Infrastructure

A machine learning-enabled respiratory estimation system can be organized as a layered architecture. The lowest layer includes optical emitters, photodetectors, analog front ends, and embedded processing units that manage sampling, filtering, and power consumption. The middleware layer performs segmentation, artifact detection, signal quality assessment, feature extraction, and model inference. The upper layer includes local or cloud-based storage, longitudinal analytics, clinical dashboards, and model update services. This layered structure creates important trade-offs between local autonomy and centralized intelligence. On-device inference reduces latency and privacy exposure while enabling operation during connectivity interruptions. However, embedded processors constrain model size, memory use, and energy consumption. Cloud-based inference allows larger and more flexible models, but it creates dependence on communication networks and raises concerns about raw physiological data transfer.

Data infrastructure is equally important. Public datasets such as the Multiparameter Intelligent Monitoring in Intensive Care II database have provided a foundation for reproducible algorithm development and benchmarking in critically ill populations [10]. The PhysioBank and PhysioNet ecosystem has further supported access to standardized physiological recordings and computational tools [11]. These resources have accelerated methodological progress, but they do not fully represent the distribution of wearable optical signals encountered in free-living conditions. Inpatient recordings often have different artifact profiles, sensor placements, and patient states compared with ambulatory wrist-worn sensing.

A recent study on rapid vital sign extraction for real-time opto-physiological monitoring at varying physical activity intensity levels illustrates the importance of evaluating respiratory estimation under dynamic activity conditions rather than only at rest [12]. This kind of evaluation is essential because physical activity changes motion artifact structure and cardiorespiratory coupling in ways that are difficult to simulate in controlled laboratory protocols.

Wearable optical signals are also sensitive to skin tone, sensor pressure, ambient light, temperature, and anatomical placement. These factors have been shown to influence the accuracy of optical heart rate sensors [13]. Respiratory rate estimation is likely to be even more vulnerable because respiratory modulations are smaller in amplitude and can be easily obscured by the same noise sources. When data infrastructure does not systematically record these covariates, machine learning models may develop spurious associations that perform well in development datasets but fail across diverse populations and device fleets. Architectural responses to these challenges include federated learning, in which model updates are computed across decentralized data sources without centralizing raw physiological signals [14]. Federated approaches could support personalization and privacy-preserving learning, but they introduce challenges in model aggregation, communication efficiency, data heterogeneity, and label alignment. System designers must therefore balance the benefits of centralized data resources against the privacy and operational advantages of distributed learning.

4. Machine Learning Modeling and Signal Processing Trade-offs

Machine learning models for respiratory rate estimation span a broad spectrum from shallow regression on engineered features to deep neural networks that learn representations directly from raw or lightly processed waveform segments. The model choice encodes trade-offs in interpretability, sample efficiency, computational cost, and robustness. Models built on handcrafted respiratory features can be more transparent and easier to validate, but they may discard subtle information that is not captured by predefined amplitude or frequency modulations. End-to-end deep models can learn richer and more flexible representations, but they often require large and diverse datasets to avoid overfitting. In wearable deployment, these models must also satisfy strict constraints on memory, processor load, and battery consumption. Quantization, pruning, knowledge distillation, and lightweight architecture design are therefore central to the translation of research models into operational systems.

The nonstationary nature of photoplethysmographic signals poses a further modeling challenge. Respiratory rate can change over seconds, and motion artifacts can appear intermittently. Temporal models such as recurrent neural networks, gated recurrent units, and convolutional architectures with dilation can capture dependencies across multiple breathing cycles. Attention mechanisms can help the model focus on cleaner signal segments while suppressing corrupted regions. These architectural choices can improve robustness, but they also increase the risk of learning shortcuts from device-specific noise or motion patterns that are not physiologically meaningful. Multi-task learning offers a useful mitigation strategy. By jointly predicting respiratory rate, heart rate, and signal quality, the model can develop shared representations that are less likely to rely on spurious correlations. This alignment with clinical practice, where respiratory rate is interpreted alongside other vital signs, makes the model more useful in real-world decision support.

Uncertainty estimation is an essential but often neglected model output. In continuous monitoring, a model that silently produces a plausible respiratory estimate during severe

motion artifact may be more harmful than one that reports low confidence or requests a repeat measurement. Probabilistic models, ensembles, and Bayesian approximations can provide measures of predictive uncertainty, but they also increase computational complexity. The system must decide how to act on uncertainty, whether by displaying a quality flag, suppressing an alert, or prompting the user to adjust the sensor. These decisions cannot be made by the model alone. They require clinical and operational protocols that translate model outputs into safe actions. In this sense, machine learning-enabled respiratory rate estimation is best understood as a socio-technical system rather than an isolated inference task.

5. Deployment, Robustness, and Operational Governance

Deployment of wearable respiratory rate estimation systems involves more than installing a model on a device. Operational robustness depends on how the system handles missing data, sensor displacement, battery depletion, connectivity loss, and software updates. Continuous monitoring can generate large volumes of data, but much of it may be redundant or of insufficient quality for clinical use. Edge-based preprocessing can reduce data transmission and energy use while preserving the most informative segments. However, aggressive preprocessing may remove information that a more capable cloud model could have used. Hybrid deployment architectures therefore require careful negotiation between local filtering and centralized analysis. The system should also include mechanisms for model versioning, audit logs, and retrospective performance monitoring.

Alarm design is a critical governance issue. Respiratory rate alarms in wearable systems can alert users or clinicians to possible deterioration, but excessive false alarms lead to alarm fatigue and reduced trust. The deployment architecture should support tiered alerts based on confidence, trend, and clinical context. For example, a single noisy estimate may warrant a request for re-measurement, while a sustained trend of elevated respiratory rate with high signal quality may trigger a clinical notification. This tiered approach depends on robust signal quality indicators and a clear understanding of model uncertainty. It also requires collaboration between engineers, clinicians, and regulatory specialists to define acceptable false positive and false negative rates for different use cases.

Deployment must also confront fairness and representation. Performance disparities can emerge across demographic groups because of differences in skin pigmentation, perfusion, body mass index, device fit, and underlying comorbidities. Algorithms trained on unrepresentative populations may reproduce structural inequities when deployed at scale [15]. The wearable respiratory monitoring system should therefore include mechanisms for subgroup evaluation and ongoing performance auditing. If certain populations experience systematically lower signal quality or higher estimation error, the system should flag these limitations rather than presenting the same confidence to all users. Governance structures should ensure that model updates do not inadvertently worsen disparities, and that evaluation datasets include meaningful demographic and clinical diversity.

6. Fairness, Privacy, and Policy Implications

Responsible machine learning for health care requires an explicit account of intended use, validation evidence, and harm mitigation [16]. For wearable respiratory rate estimation, the intended use may range from consumer wellness tracking to clinical early warning. These contexts carry different regulatory burdens and risk profiles. A model that is acceptable for general wellness may be inadequate forclinical decision support without additional validation and regulatory approval. In the United States, the Food and Drug Administration has provided

guidance for software as a medical device, and in the European Union, the Medical Device Regulation imposes obligations for clinical evaluation and post-market surveillance. These frameworks require manufacturers to demonstrate that wearable respiratory monitoring systems perform safely across intended populations and use conditions, including ambulatory settings with motion artifact [17]. In many jurisdictions, algorithms that provide diagnostic or therapeutic advice are subject to stricter oversight than general wellness tools. This distinction has direct consequences for respiratory rate estimation because the same model could be deployed as a consumer feature or as a clinical early warning system with different evidentiary requirements.

Privacy governance is another central concern. Wearable optical signals are identifiable physiological data that may reveal cardiovascular, respiratory, and even behavioral information. Under frameworks such as the European General Data Protection Regulation and sector-specific health privacy laws, continuous photoplethysmographic recordings must be treated as sensitive health data [18]. Raw waveform data can be re-identified or linked to other health information, and centralized data aggregation increases the risk of secondary use beyond the user's consent. Federated learning and on-device processing offer architectural responses that reduce raw data movement, but they also complicate transparency and auditability. Systems designers should document where data are stored, how long they are retained, and how model updates are triggered. Privacy impact assessments should be conducted before deployment in settings where vulnerable populations may be monitored at scale.

The operational implementation of wearable respiratory monitoring hinges on interoperability and clinical workflow integration. Effective systems do not merely deliver a respiratory rate value; they organize data around meaningful episodes, integrate with electronic health records, and support timely intervention without overwhelming clinicians [19]. This requires standardized data formats, semantic interoperability, and role-based access controls. Alarm thresholds should be configurable according to clinical context, and the user interface must communicate signal quality and confidence clearly. When a respiratory estimate is based on a short clean segment rather than a full minute, the system should indicate that the value may be less stable. In high-risk surveillance scenarios, automated algorithms must be embedded within broader safety protocols that include human review and escalation pathways.

The environmental and economic sustainability of machine learning-enabled wearable monitoring deserves greater attention. Although wearable devices themselves are low-power, the backend infrastructure for continuous data storage, model retraining, and cloud inference can be energy-intensive. Training deep models on large physiological datasets has an associated carbon footprint that scales with model size and retraining frequency [20]. Edge computing can reduce network traffic and centralized energy use, but it may require more frequent on-device model updates and dedicated hardware. Sustainable deployment should therefore include energy audits, model compression, and selective retraining based on performance drift rather than arbitrary periods. Economic sustainability also depends on equitable access. If advanced respiratory monitoring is available only through premium consumer devices, the technology may widen health disparities rather than reduce them.

Future research should address several systemic gaps. Synthetic physiological data generation can augment limited clinical datasets and improve representation of rare respiratory patterns, but synthetic signals must be physiologically faithful and validated against real recordings [21]. Explainability remains an unmet need for clinical adoption. Clinicians are more likely to

trust a respiratory rate estimate when they can understand which signal features contributed to the inference and whether the estimate was influenced by artifact [22]. Explainability should be integrated with uncertainty reporting so that clinicians can distinguish between a confident estimate based on clean respiration and a fragile estimate based on marginal signal quality. Longitudinal wearable studies should evaluate respiratory rate estimation as part of real-time disease monitoring across large and diverse cohorts [23].

Finally, the broader governance of artificial intelligence in health requires a learning system that monitors performance after deployment. Remote patient monitoring has demonstrated the value and the limits of wearable data in public health response, particularly during respiratory disease outbreaks [24]. A governance model should include model registries, performance dashboards, bias audits, and channels for users to report erroneous estimates. This aligns with national and international calls for trustworthy artificial intelligence in health care, which emphasize the need for safety, effectiveness, transparency, and accountability throughout the algorithm life cycle [25]. Respiratory rate estimation is not a static problem; it evolves with new sensor generations, population characteristics, and computational methods. Governance structures must therefore be adaptive and evidence-based.

7. Conclusion

Machine learning-enabled respiratory rate estimation from wearable optical signals represents a meaningful convergence of sensor physiology, data infrastructure, and algorithmic inference. The respiratory modulation embedded in photoplethysmography is subtle and distributed across multiple physiological mechanisms, and its reliable extraction demands more than isolated model optimization. This paper has emphasized a system-level perspective, examining the architecture of wearable sensing platforms, the importance of signal quality assessment, the trade-offs between on-device and cloud inference, and the role of multicomponent data fusion. Robustness cannot be achieved by algorithm design alone; it requires signal processing pipelines, quality metrics, and deployment protocols that respond coherently to motion, low perfusion, device variation, and population heterogeneity.

The clinical usefulness of these systems depends on their integration into broader health care workflows. A respiratory rate estimate is only valuable when it is accompanied by context, confidence, and appropriate governance. Fairness concerns, including performance disparities across skin tones and demographic groups, must be treated as design constraints rather than post-deployment observations. Privacy architectures such as federated learning and on-device inference are important, but they must be balanced against the need for auditability and clinical oversight. Regulatory frameworks provide evolving guidance, yet many operational decisions remain in the hands of system designers and deploying institutions. Sustainable and equitable deployment requires attention to energy consumption, data stewardship, and user-centered alarm design.

Future progress will likely depend on high-quality longitudinal datasets that reflect real-world variability, personalized adaptive models that respect individual physiology, and open benchmarks that evaluate performance under activity and artifact conditions. The development of synthetic data, explainable inference, and multi-task learning architectures can strengthen the foundation for trustworthy respiratory monitoring. However, these technical advances must be paired with robust governance, continuous performance auditing, and transparent uncertainty communication. If these systemic requirements are met, machine learning-enabled respiratory rate estimation can move beyond proof-of-concept demonstrations toward a dependable element of preventive, acute, and chronic care.

References

1. Allen, J. (2007). Photoplethysmography and its application in clinical physiological measurement. Physiological Measurement, 28(3), R1–R39.

2. Meredith, D. J., Clifton, D., Charlton, P., Brooks, J., Pugh, C. W., & Tarassenko, L. (2012). Photoplethysmographic derivation of respiratory rate: A review of relevant physiology. Journal of Medical Engineering & Technology, 36(1), 1–7.

3. Charlton, P. H., Bonnici, T., Tarassenko, L., Clifton, D. A., Beale, R., & Watkinson, P. J. (2016). An assessment of algorithms to estimate respiratory rate from the electrocardiogram and photoplethysmogram. Physiological Measurement, 37(4), 610–626.

4. Pimentel, M. A. F., Johnson, A. E. W., Charlton, P. H., Birrenkott, D., Watkinson, P. J., Tarassenko, L., & Clifton, D. A. (2017). Toward a robust estimation of respiratory rate from pulse oximeters. IEEE Transactions on Biomedical Engineering, 64(8), 1914–1923.

5. Dunn, J., Runge, R., & Snyder, M. (2018). Wearables and the medical revolution. Personalized Medicine, 15(5), 429–448.

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

7. Charlton, P. H., Birrenkott, D., Bonnici, T., Pimentel, M. A. F., Johnson, A. E. W., Alastruey, J., Pueyo, E., & Clifton, D. A. (2018). Breathing rate estimation from the electrocardiogram and photoplethysmogram: A review. IEEE Reviews in Biomedical Engineering, 11, 2–20.

8. Karlen, W., Raman, S., Ansermino, J. M., & Dumont, G. A. (2013). Multiparameter respiratory rate estimation from the photoplethysmogram. IEEE Transactions on Biomedical Engineering, 60(7), 1946–1953.

9. Orphanidou, C., Bonnici, T., Tarassenko, L., & Clifton, D. A. (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.

10. Saeed, M., Villarroel, M., Reisner, A. T., Clifford, G., Lehman, L.-W., Moody, G., Heldt, T., Kyaw, T. H., Moody, B., & Mark, R. G. (2011). Multiparameter Intelligent Monitoring in Intensive Care II: A public-access intensive care unit database. Critical Care Medicine, 39(5), 952–960.

11. Goldberger, A. L., Amaral, L. A. N., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C.-K., & Stanley, H. E. (2000). PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. Circulation, 101(23), e215–e220.

12. 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.

13. Sjoding, M. W., Dickson, R. P., Iwashyna, T. J., Gay, S. E., & Valley, T. S. (2020). Racial bias in pulse oximetry measurement. New England Journal of Medicine, 383(25), 2477–2478.

14. Rieke, N., Hancox, J., Li, W., Milletarì, 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. Nature Digital Medicine, 3, 119.

15. 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.

16. Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V. X., Doshi-Velez, F., Jung, K., Heller, K., Kale, D., Saeed, M., Ossorio, P. N., Thadaney-Israni, S., & Goldenberg, A. (2019). Do no harm: A roadmap for responsible machine learning in health care. Nature Medicine, 25(9), 1337–1340.

17. Gerke, S., Minssen, T., & Cohen, G. (2020). Ethical and legal challenges of artificial intelligence-driven healthcare. Artificial Intelligence in Medicine, 103, 101703.

18. Price, W. N., & Cohen, I. G. (2019). Privacy in the age of medical big data. Nature Medicine, 25(1), 37–43.

19. Smuck, M., Odonkor, C. A., Wilt, J. K., Schmidt, J., & Swerdlow, M. A. (2021). The emerging clinical role of wearables: Factors for successful implementation in healthcare. npj Digital Medicine, 4, 45.

20. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics.

21. Chen, R. J., Lu, M. Y., Chen, T. Y., Williamson, D. F. K., & Mahmood, F. (2021). Synthetic data in machine learning for medicine and healthcare. Nature Biomedical Engineering, 5(6), 493–497.

22. Tjoa, E., & Guan, C. (2020). A survey on explainable artificial intelligence (XAI): Toward medical XAI. IEEE Transactions on Neural Networks and Learning Systems, 32(11), 4793–4813.

23. Guk, K., Han, G., Lim, J., Jeong, K., Kang, T., Lim, E.-K., & Jung, J. (2019). Evolution of wearable devices with real-time disease monitoring for personalized healthcare. Nanomaterials, 9(6), 813.

24. Seshadri, D. R., Davies, E. V., Harlow, E. R., Hsu, J. J., Knighton, S. C., Walker, T. A., Voos, J. E., & Drummond, C. K. (2020). Wearable sensors for COVID-19: A call to action to harness our digital infrastructure for remote patient monitoring and virtual assessments. Frontiers in Digital Health, 2, 8.

25. Matheny, M. E., Whicher, D., & Thadaney Israni, S. (2020). Artificial intelligence in health care: The hope, the hype, the promise, the peril. National Academy of Medicine.