Federated Learning Framework for Cross-Institutional Prediction of Analgesic Effectiveness in Knee Arthroscopy

Derek Langford
Department of Biomedical Informatics, University of Arkansas, Fayetteville, AR, USA.
d.langford@uark.edu

Emily R. Foster
School of Computing, Southern Illinois University Carbondale, Carbondale, IL, USA.
erfoster@siu.edu

Julian Carrington
Department of Health Data Science, University of North Texas, Denton, TX, USA.
j.carrington@unt.edu

Abstract

The increasing adoption of regional anesthesia techniques in arthroscopic knee surgery has generated substantial volumes of heterogeneous perioperative data across healthcare institutions. Despite growing interest in predictive analytics for individualized pain management, the development of reliable analgesic effectiveness prediction systems remains constrained by fragmented data ownership, privacy regulations, institutional heterogeneity, and limited interoperability among healthcare providers. Federated learning has emerged as a promising paradigm capable of enabling collaborative model development while preserving patient confidentiality and institutional autonomy. This study proposes a comprehensive federated learning framework for cross-institutional prediction of analgesic effectiveness following knee arthroscopy. Rather than focusing solely on predictive accuracy, the research examines the broader socio-technical architecture required to support sustainable deployment across diverse healthcare environments.

The paper investigates the interaction among distributed machine learning infrastructures, perioperative clinical workflows, governance mechanisms, data standardization strategies, and ethical considerations. Particular attention is devoted to challenges associated with data heterogeneity, model fairness, communication efficiency, privacy preservation, and regulatory compliance. The proposed framework integrates local institutional learning, secure parameter aggregation, federated governance structures, and continuous model monitoring mechanisms. Through a system-level analysis, the study demonstrates how federated learning can facilitate collaborative knowledge generation without necessitating centralized patient databases. Furthermore, the paper explores the implications of cross-institutional predictive intelligence for personalized analgesia planning, resource allocation, clinical decision support, and healthcare quality improvement.

The findings suggest that successful implementation of federated learning in perioperative pain management requires a balanced integration of technical innovation, organizational coordination, and policy development. The framework provides a foundation for future large-

scale clinical AI ecosystems capable of improving postoperative outcomes while maintaining trust, transparency, and institutional independence.

Keywords

Federated Learning; Clinical Artificial Intelligence; Knee Arthroscopy; Analgesic Effectiveness Prediction; Regional Anesthesia; Distributed Healthcare Systems; Privacy-Preserving Machine Learning; Healthcare Informatics.

1. Introduction

Arthroscopic knee surgery has become one of the most frequently performed orthopedic procedures worldwide, reflecting advances in minimally invasive surgical techniques and increasing demand for musculoskeletal interventions. Although knee arthroscopy generally results in favorable clinical outcomes, postoperative pain management remains a significant challenge affecting patient satisfaction, recovery trajectories, rehabilitation adherence, healthcare utilization, and overall quality of care. Regional anesthesia approaches, particularly femoral nerve blocks and combined femoral-sciatic nerve block strategies, have demonstrated substantial benefits in reducing postoperative pain and opioid consumption while improving functional recovery [7]. Nevertheless, considerable variability persists in analgesic effectiveness across patient populations, healthcare institutions, and clinical practice environments.

Recent developments in artificial intelligence have introduced new opportunities for predicting postoperative pain outcomes and supporting individualized analgesic planning. Machine learning models can leverage large volumes of perioperative information, including demographic characteristics, surgical variables, anesthesia parameters, medication histories, physiological measurements, and recovery indicators. These predictive capabilities offer the potential to move beyond generalized treatment protocols toward precision analgesia strategies tailored to individual patients. However, achieving robust predictive performance requires access to diverse, representative, and large-scale datasets that exceed the capacity of many single institutions.

The conventional approach to clinical machine learning relies heavily on centralized data aggregation, whereby patient information from multiple healthcare organizations is transferred to a common repository for model development. While this strategy can facilitate large-scale analytics, it introduces substantial concerns related to privacy protection, regulatory compliance, cybersecurity risks, institutional competition, and governance complexity. Healthcare organizations are increasingly reluctant to share raw patient data due to legal obligations and reputational considerations. Consequently, data fragmentation has become a major barrier to the development of generalizable clinical AI systems.

Federated learning offers an alternative paradigm that fundamentally reconfigures the relationship between data ownership and collaborative model development [1]. Instead of moving data to algorithms, federated learning distributes algorithms to data sources, allowing participating institutions to train local models while sharing only model parameters or updates. This architecture enables collective learning across geographically dispersed organizations without exposing raw patient information. The resulting framework has attracted significant attention across healthcare domains, including medical imaging, disease prediction, population health analytics, and personalized medicine [2,3].

Within the context of knee arthroscopy, federated learning presents unique opportunities for constructing predictive models capable of estimating analgesic effectiveness across diverse patient populations and clinical settings. Such systems may enhance evidence generation, improve treatment personalization, reduce outcome disparities, and strengthen institutional collaboration. However, successful deployment requires more than technical feasibility. It demands a comprehensive understanding of organizational structures, infrastructure requirements, governance mechanisms, fairness considerations, and long-term sustainability.

This paper develops a federated learning framework specifically designed for cross-institutional prediction of analgesic effectiveness in knee arthroscopy. The study emphasizes system-level analysis rather than algorithmic optimization alone. By examining interactions among technology, healthcare operations, policy environments, and clinical practice, the paper seeks to establish a holistic foundation for future implementation efforts.

2. Background and Conceptual Foundations

The evolution of perioperative pain management has increasingly reflected broader transformations occurring throughout healthcare systems. Traditional approaches to analgesic planning relied heavily on clinician experience, generalized treatment protocols, and population-level evidence. While these methods remain valuable, growing recognition of interpatient variability has motivated the search for more individualized decision-support mechanisms. Factors such as age, sex, comorbidities, genetic predispositions, psychological characteristics, procedural complexity, and anesthetic techniques contribute to substantial differences in postoperative pain experiences.

Machine learning technologies have emerged as promising tools for capturing these complex multidimensional relationships. Predictive models can identify subtle patterns within large clinical datasets that may not be readily apparent through conventional statistical methods. In orthopedic surgery, researchers have increasingly investigated machine learning applications for outcome prediction, complication detection, rehabilitation planning, and resource management [4,5]. Nevertheless, many existing models suffer from limited generalizability due to narrow institutional datasets and insufficient demographic diversity.

Federated learning addresses these limitations by enabling collaborative model development across organizational boundaries. The concept was originally introduced to support privacy-preserving machine learning within distributed environments where direct data sharing was impractical [1]. In healthcare contexts, federated learning aligns closely with emerging regulatory expectations emphasizing patient privacy, institutional accountability, and responsible AI development.

From a systems perspective, federated learning represents a form of distributed intelligence infrastructure. Rather than concentrating analytical capabilities within a centralized repository, federated architectures create networks of interconnected learning nodes. Each participating institution contributes local knowledge derived from its own patient population while retaining operational control over sensitive information. The collective model gradually accumulates generalized insights through iterative parameter aggregation processes.

The relevance of federated learning extends beyond technical performance. Healthcare systems increasingly operate within complex ecosystems characterized by fragmented governance structures, heterogeneous technologies, and competing organizational incentives. Federated learning provides a mechanism for balancing collaboration and autonomy. Institutions can participate in collective knowledge generation without relinquishing ownership of strategic data assets. This balance may prove particularly important in perioperative medicine, where clinical practices vary substantially across hospitals and regional healthcare networks.

The clinical context of knee arthroscopy further highlights the importance of distributed learning approaches. Analgesic effectiveness depends upon numerous interacting variables spanning preoperative, intraoperative, and postoperative phases of care. Differences in patient selection criteria, anesthetic techniques, rehabilitation protocols, and institutional workflows generate substantial heterogeneity. A federated framework can leverage this diversity as a source of learning rather than treating it solely as a methodological obstacle.

Moreover, contemporary discussions regarding trustworthy AI emphasize transparency, fairness, accountability, and explainability as essential dimensions of clinical deployment [6]. Federated learning introduces additional governance challenges related to model ownership, update validation, auditing procedures, and collaborative decision-making. Consequently, technical architecture must be complemented by robust organizational structures capable of sustaining long-term cooperation among participating institutions.

3. System Architecture of the Proposed Federated Learning Framework

The proposed federated learning framework is conceptualized as a multi-layer socio-technical architecture integrating clinical data environments, distributed machine learning processes, governance structures, and operational monitoring mechanisms. Rather than viewing federated learning as a standalone algorithmic solution, the framework positions it as a component within a broader healthcare intelligence ecosystem.

At the institutional level, participating hospitals maintain local data repositories containing perioperative information relevant to knee arthroscopy procedures. These repositories include demographic characteristics, comorbidity profiles, anesthesia records, surgical details, medication administration logs, pain assessment scores, recovery indicators, and follow-up outcomes. Local data remain within institutional boundaries and are governed according to existing privacy policies and regulatory requirements.

Embedded within each institution is a local learning environment responsible for model training and validation. The environment performs data preprocessing, feature harmonization, quality assessment, and model optimization using locally available information. Importantly, no raw patient records are transmitted outside institutional infrastructure. Instead, only abstract model updates are prepared for external communication.

A central coordination layer serves as the aggregation mechanism connecting participating institutions. This layer receives encrypted model parameters from local nodes, performs secure aggregation procedures, and generates updated global models. The resulting models are redistributed to participating institutions for subsequent training cycles. Through repeated iterations, collective intelligence emerges without direct data sharing.

One of the most significant architectural challenges involves addressing data heterogeneity. Healthcare organizations frequently employ different electronic health record systems, coding standards, documentation practices, and measurement protocols. Without careful harmonization, such variability can undermine model performance and introduce systematic bias. The proposed framework incorporates semantic interoperability mechanisms designed to map institution-specific variables into common analytical representations while preserving local operational flexibility.

Another essential architectural component involves communication efficiency. Large-scale federated learning networks may include dozens or hundreds of participating institutions, generating substantial communication overhead during model synchronization. Excessive network demands can reduce scalability and increase operational costs. Therefore, the framework incorporates adaptive communication strategies that prioritize meaningful model updates while minimizing unnecessary data transmission.

Security considerations occupy a central role within the architecture. Healthcare organizations remain attractive targets for cyberattacks, making secure communication channels essential. The framework integrates encryption protocols, authentication mechanisms, secure aggregation techniques, and continuous threat monitoring capabilities. These protections aim to reduce vulnerabilities associated with distributed model training environments while maintaining operational efficiency.

Beyond technical infrastructure, governance architecture forms a critical component of the overall framework. Participating institutions require transparent mechanisms for decision-making, model validation, performance evaluation, and conflict resolution. Federated governance councils can provide oversight regarding participation criteria, update schedules, ethical standards, and auditing procedures. Such structures help establish trust among stakeholders and support sustainable collaboration over extended periods.

The integration of these architectural layers creates a resilient foundation for cross-institutional prediction of analgesic effectiveness. By combining distributed intelligence, privacy preservation, interoperability support, and collaborative governance, the framework seeks to overcome longstanding barriers limiting the development of scalable clinical AI systems.

4. Data Governance, Privacy Preservation, and Institutional Collaboration

The long-term viability of federated learning in healthcare depends not only on computational performance but also on the governance structures that regulate data stewardship, organizational participation, and accountability. In cross-institutional analgesic effectiveness prediction, governance becomes particularly important because participating hospitals may differ substantially in their operational priorities, technological maturity, legal obligations, and strategic interests. Consequently, successful federated learning initiatives require governance frameworks capable of balancing institutional autonomy with collective objectives.

Traditional centralized research infrastructures often rely on extensive data-sharing agreements that specify ownership rights, usage limitations, liability provisions, and compliance requirements. Federated learning alters these relationships by allowing institutions to retain control of patient records while contributing to collaborative model development. Although this approach reduces many privacy concerns, it does not eliminate governance challenges. Participating organizations must still establish consensus regarding model ownership, intellectual property rights, update procedures, performance standards, and acceptable use policies.

The proposed framework incorporates a multi-layer governance structure designed to support trust and transparency among participants. At the operational level, each institution maintains responsibility for local data quality assurance, model training compliance, and ethical oversight. At the network level, a federated governance board coordinates model validation, auditing activities, dispute resolution mechanisms, and policy harmonization efforts. Such arrangements distribute responsibility across stakeholders while avoiding excessive centralization.

Privacy preservation remains a central motivation for federated learning adoption [1,2]. Healthcare data are among the most sensitive forms of personal information, and regulations such as the Health Insurance Portability and Accountability Act (HIPAA) in the United States impose stringent requirements regarding patient confidentiality. Federated learning aligns with these regulatory objectives because patient-level information remains within local institutional environments. Nevertheless, researchers have demonstrated that model updates themselves may occasionally reveal sensitive information under certain attack scenarios [8].

To address these risks, the proposed architecture incorporates multiple layers of privacy protection. Secure aggregation mechanisms prevent central coordinators from observing individual institutional updates. Differential privacy techniques introduce carefully controlled statistical noise that reduces the possibility of patient re-identification. Furthermore, continuous security auditing and anomaly detection systems monitor network activities for potential adversarial behavior. These measures collectively strengthen trustworthiness while preserving analytical utility.

Institutional collaboration introduces additional considerations related to incentive alignment. Healthcare organizations may differ in their willingness to contribute computational resources, personnel time, and infrastructure investments. Some institutions may possess larger datasets, while others contribute specialized expertise or unique patient populations. Governance frameworks must therefore establish equitable participation models that recognize diverse forms of contribution while maintaining fairness across the network.

The success of federated learning initiatives ultimately depends on sustained cooperation rather than short-term experimentation. Consequently, governance structures should be viewed as strategic infrastructure rather than administrative overhead. By fostering transparency, accountability, and mutual trust, effective governance enables distributed learning ecosystems capable of generating clinically meaningful insights while respecting institutional independence and patient privacy.

5. Fairness, Bias Mitigation, and Ethical Considerations

The increasing use of artificial intelligence in healthcare has intensified concerns regarding fairness, equity, and ethical responsibility. Predictive models developed from historical healthcare data may inadvertently reproduce existing disparities related to socioeconomic status, race, ethnicity, geographic location, insurance coverage, and healthcare accessibility [9]. In the context of postoperative pain management, such disparities can have direct consequences for treatment quality, patient experience, and recovery outcomes.

Federated learning introduces both opportunities and challenges for fairness enhancement. On one hand, cross-institutional collaboration expands the diversity of training data by incorporating patients from multiple healthcare environments. Greater population diversity may improve model generalizability and reduce biases associated with single-institution datasets. On the other hand, federated networks may also amplify disparities if participating institutions contribute unevenly distributed populations or exhibit significant differences in clinical practices.

Analgesic effectiveness prediction presents a particularly complex fairness challenge because pain itself is influenced by biological, psychological, cultural, and social factors. Numerous studies have documented variability in pain assessment, treatment allocation, and outcome reporting across demographic groups [10]. Consequently, predictive models trained solely on historical treatment patterns may inadvertently reinforce existing inequities rather than improve clinical decision-making.

The proposed framework incorporates fairness monitoring mechanisms throughout the model lifecycle. Rather than evaluating performance exclusively through aggregate accuracy metrics, the system examines predictive behavior across multiple demographic and institutional subgroups. Continuous monitoring enables early detection of performance disparities and facilitates corrective interventions before deployment-related harms emerge.

Transparency represents another critical ethical consideration. Clinicians are unlikely to trust predictive systems that operate as opaque black boxes, particularly when recommendations influence pain management decisions. Explainable AI methodologies can enhance interpretability by identifying factors that contribute to individual predictions [11]. Within the federated context, explainability also supports accountability by enabling stakeholders to understand how distributed learning processes influence model behavior.

Ethical governance extends beyond technical implementation. Questions regarding responsibility, liability, and clinical oversight remain central to healthcare AI deployment. If predictive recommendations contribute to adverse outcomes, determining accountability may become challenging within distributed learning environments. The proposed governance framework therefore emphasizes human-centered decision support rather than full automation. Predictive outputs are intended to augment clinical judgment rather than replace professional expertise.

Patient trust constitutes another essential ethical dimension. Public acceptance of AI-enabled healthcare systems depends upon confidence that technologies are deployed responsibly and transparently. Federated learning may strengthen trust by reducing the need for centralized data repositories, thereby addressing common concerns regarding surveillance, commercialization, and unauthorized information sharing. However, trust cannot be achieved through technical mechanisms alone. Ongoing stakeholder engagement, public communication, and ethical oversight remain necessary components of responsible implementation.

As healthcare systems increasingly adopt distributed AI infrastructures, fairness and ethics must be integrated into design processes rather than treated as secondary considerations. Sustainable innovation requires balancing predictive performance with societal values, ensuring that technological progress contributes to equitable and trustworthy healthcare delivery.

6. Clinical Deployment and Operational Integration

Bridging the gap between research innovation and clinical practice represents one of the most persistent challenges in healthcare artificial intelligence. Many predictive models demonstrate promising performance under experimental conditions but fail to achieve meaningful impact in real-world environments. Successful deployment requires careful integration with clinical workflows, organizational processes, and healthcare infrastructure.

Within knee arthroscopy pathways, analgesic effectiveness prediction has potential applications across multiple stages of perioperative care. During preoperative assessment, predictive models may help clinicians identify patients at elevated risk of inadequate pain control. Such information can inform anesthesia planning, patient counseling, resource allocation, and postoperative monitoring strategies. During recovery, predictive insights may support personalized pain management interventions aimed at improving patient outcomes while reducing unnecessary medication exposure.

The proposed federated learning framework is designed to function as a clinical decision-support ecosystem rather than an isolated analytical tool. Integration with electronic health record platforms enables predictive outputs to be delivered within existing clinician workflows. This approach minimizes disruption while increasing the likelihood of practical adoption.

Operational integration requires careful attention to usability and workflow compatibility. Healthcare professionals operate within highly demanding environments characterized by time constraints, information overload, and competing priorities. Predictive systems that generate excessive alerts or require substantial additional effort may encounter resistance regardless of technical performance. Consequently, user-centered design principles play a critical role in deployment success.

Institutional readiness also influences implementation outcomes. Hospitals vary considerably in terms of digital infrastructure, workforce capabilities, leadership support, and organizational culture. Federated learning initiatives must therefore accommodate heterogeneous operational contexts. Flexible deployment architectures allow institutions to participate according to local capabilities while maintaining interoperability across the broader network.

Continuous monitoring represents another essential component of operational deployment. Clinical environments are dynamic, with patient populations, treatment protocols, and healthcare policies evolving over time. Models that perform well during initial development may experience performance degradation due to shifting clinical conditions [12]. The proposed framework incorporates lifecycle management mechanisms capable of detecting drift, initiating retraining processes, and maintaining model relevance.

Economic considerations further influence deployment feasibility. Healthcare organizations increasingly face pressure to demonstrate value from technology investments. Federated learning infrastructures require expenditures related to computing resources, cybersecurity, personnel training, and governance activities. However, these costs may be offset through improved pain management outcomes, reduced complications, enhanced patient satisfaction, and more efficient resource utilization. Comprehensive evaluation frameworks should therefore assess both clinical and organizational impacts.

Ultimately, the success of federated learning in perioperative medicine depends upon its ability to integrate seamlessly into healthcare operations. Technical excellence alone is insufficient. Sustainable deployment requires alignment among technology, clinical practice, organizational strategy, and human factors.

7. Future Directions, Policy Implications, and Sustainability

The emergence of federated learning reflects a broader transition toward distributed healthcare intelligence ecosystems. As healthcare organizations generate increasing volumes of digital information, collaborative analytical infrastructures will likely become essential for addressing complex clinical challenges. Prediction of analgesic effectiveness in knee arthroscopy represents only one example of a much larger transformation involving data-driven healthcare delivery.

Future federated learning systems may evolve beyond traditional hospital networks to include ambulatory surgery centers, rehabilitation facilities, primary care providers, wearable device platforms, and patient-generated health data sources. Such expansion could enable more comprehensive understanding of perioperative recovery trajectories while supporting continuous learning across the entire continuum of care.

Advances in interoperability standards are expected to play a significant role in facilitating this evolution. Greater adoption of standardized healthcare data models may reduce integration barriers and improve collaborative learning efficiency. Simultaneously, developments in privacy-enhancing technologies will likely strengthen protections against emerging cybersecurity threats while maintaining analytical effectiveness.

Policy frameworks will substantially influence the trajectory of federated healthcare AI. Regulatory agencies increasingly recognize both the opportunities and risks associated with machine learning technologies [13]. Future regulations may address issues such as model auditing, algorithmic transparency, cross-institutional accountability, and continuous performance monitoring. Policymakers must balance innovation promotion with patient protection, ensuring that regulatory environments encourage responsible technological development.

Sustainability considerations extend beyond technical performance and regulatory compliance. Long-term success requires stable funding mechanisms, workforce development strategies, governance maturity, and organizational commitment. Federated learning networks should be viewed as enduring infrastructure investments rather than temporary research projects. Their value derives not only from individual predictive models but also from the collaborative ecosystems they create.

International collaboration represents another promising direction. Pain management practices vary across healthcare systems, cultural contexts, and geographic regions. Global federated networks could facilitate comparative analyses that generate broader clinical insights while respecting jurisdiction-specific privacy requirements. Such initiatives may contribute to more inclusive and representative forms of medical knowledge production.

Importantly, future development efforts should maintain a human-centered orientation. Artificial intelligence technologies are most effective when they augment rather than replace clinical expertise. Federated learning systems should support clinicians, empower patients, and strengthen healthcare institutions while preserving the interpersonal dimensions of medical care. By aligning technological innovation with ethical principles and societal values, federated learning can contribute to more resilient and equitable healthcare systems.

8. Conclusion

Federated learning offers a transformative approach to cross-institutional prediction of analgesic effectiveness in knee arthroscopy by enabling collaborative model development without requiring centralized patient data repositories. This paper has proposed a comprehensive system-level framework integrating distributed machine learning architecture, privacy-preserving technologies, governance mechanisms, fairness monitoring processes, clinical deployment strategies, and long-term sustainability considerations.

The analysis demonstrates that successful implementation depends on far more than predictive accuracy. Robust federated ecosystems require coordinated attention to interoperability, organizational collaboration, regulatory compliance, cybersecurity, ethical accountability, and operational integration. By addressing these interconnected dimensions, healthcare institutions can develop predictive intelligence capabilities that enhance personalized pain management while maintaining patient privacy and institutional autonomy.

As healthcare systems continue their digital transformation, federated learning provides a promising foundation for scalable, trustworthy, and collaborative clinical AI infrastructures. The framework presented here contributes to ongoing efforts to create data-driven healthcare environments capable of improving patient outcomes, supporting clinician decision-making, and fostering responsible innovation across institutional boundaries.

References

1. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Aguera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 1273–1282.

2. 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(1), 119.

3. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3), 50–60.

4. Senders, J. T., Staples, P. C., Karhade, A. V., Zaki, M. M., Gormley, W. B., Broekman, M. L. D., Smith, T. R., & Arnaout, O. (2018). Machine learning and neurosurgical outcome prediction. World Neurosurgery, 109, 476–486.

5. Cabitza, F., Rasoini, R., & Gensini, G. F. (2017). Unintended consequences of machine learning in medicine. JAMA, 318(6), 517–518.

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

7. 金子, 吴川, 王秀丽, & 刘朋. (2014). 股神经联合坐骨神经阻滞用于膝关节镜诊治术. 实用医学杂志, 30(4), 666-667.

8. Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. IEEE Symposium on Security and Privacy, 3–18.

9. 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.

10. Meghani, S. H., Byun, E., & Gallagher, R. M. (2012). Time to take stock: A meta-analysis and systematic review of analgesic treatment disparities. Pain Medicine, 13(2), 150–174.

11. Samek, W., Wiegand, T., & Müller, K. R. (2017). Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. arXiv preprint arXiv:1708.08296.

12. Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. BMC Medicine, 17(1), 195.

13. Benjamens, S., Dhunnoo, P., & Meskó, B. (2020). The state of artificial intelligence-based FDA-approved medical devices and algorithms. NPJ Digital Medicine, 3(1), 118.

14. 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.

15. Kaissis, G., Makowski, M., Rückert, D., & Braren, R. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. Nature Machine Intelligence, 2(6), 305–311.

16. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, J., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P., & Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210.

17. Wiens, J., & Shenoy, E. S. (2018). Machine learning for healthcare: On the verge of a major shift in healthcare epidemiology. Clinical Infectious Diseases, 66(1), 149–153.

18. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. New England Journal of Medicine, 380(14), 1347–1358.

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

20. Beam, A. L., & Kohane, I. S. (2018). Big data and machine learning in health care. JAMA, 319(13), 1317–1318.