Multimodal Clinical Data Fusion for Perioperative Risk Stratification in Arthroscopic Knee Procedures Under Regional Anesthesia

Hugh Whitlock
Department of Biomedical Informatics, University of Missouri–Kansas City, Kansas City, Missouri, USA.
h.whitlock@umkc.edu

Peul Qllisan
Department of Anesthesiology, University of Vermont, Burlington, Vermont, USA.
p.qllisan@uvm.edu

Keith Redcliffe
School of Computing and Information Sciences, University of Maine, Orono, Maine, USA.
k.redcliffe@maine.edu

Abstract

The increasing adoption of regional anesthesia techniques in arthroscopic knee procedures has generated substantial opportunities for data-driven perioperative risk management. Contemporary perioperative environments produce large volumes of heterogeneous clinical information, including electronic health records, physiological monitoring streams, imaging data, laboratory results, anesthetic documentation, patient-reported outcomes, and operational workflow records. Despite this abundance of information, risk stratification processes remain fragmented and often rely on isolated indicators rather than integrated representations of patient status. Multimodal clinical data fusion offers a promising pathway for transforming perioperative decision-making by combining diverse information sources into comprehensive predictive frameworks capable of supporting individualized risk assessment and resource allocation.

This paper examines the system-level architecture, governance considerations, and deployment challenges associated with multimodal clinical data fusion for perioperative risk stratification in arthroscopic knee procedures performed under regional anesthesia. Particular attention is given to the integration of preoperative, intraoperative, and postoperative data streams within artificial intelligence-enabled clinical environments. The discussion explores data interoperability, infrastructure requirements, algorithmic robustness, fairness considerations, human-AI collaboration, and institutional governance mechanisms necessary for sustainable implementation. The paper further evaluates trade-offs between predictive accuracy, interpretability, scalability, and clinical usability while considering broader healthcare policy implications.

The analysis argues that successful deployment of multimodal risk stratification systems requires more than algorithmic innovation. Effective implementation depends upon coordinated sociotechnical infrastructures that align technological capabilities with clinical workflows, regulatory expectations, organizational governance, and patient-centered care objectives. Multimodal fusion frameworks have the potential to substantially improve perioperative safety and efficiency, but their long-term impact will depend on the development of trustworthy, equitable, and resilient healthcare AI ecosystems capable of supporting continuous learning and adaptation across diverse clinical settings.

Keywords

Clinical artificial intelligence; Multimodal data fusion; Perioperative risk stratification; Regional anesthesia; Arthroscopic knee surgery; Healthcare informatics; Clinical decision support; Digital health infrastructure.

1. Introduction

Arthroscopic knee procedures represent one of the most frequently performed orthopedic interventions worldwide and are increasingly conducted using regional anesthesia techniques because of their favorable safety profile, reduced opioid requirements, accelerated recovery trajectories, and enhanced patient satisfaction [1,2]. Advances in ultrasound-guided nerve block procedures have improved anesthetic precision and expanded the applicability of regional anesthesia across a wide range of knee interventions [3]. Nevertheless, perioperative risk remains highly heterogeneous among patients due to variations in physiological status, comorbidity burden, functional capacity, pain sensitivity, medication exposure, and healthcare utilization patterns [4].

Traditional perioperative risk assessment approaches frequently rely on limited subsets of clinical variables collected during preoperative evaluations. While these methods provide valuable baseline information, they often fail to capture the dynamic and multidimensional nature of patient risk. Modern perioperative environments generate extensive quantities of data throughout the surgical journey, including laboratory values, physiological monitoring records, imaging findings, medication administration logs, clinician documentation, and patient-reported experiences [5]. The challenge facing healthcare systems is not merely data acquisition but meaningful integration of heterogeneous information sources into coherent decision-support mechanisms.

Recent developments in artificial intelligence and machine learning have accelerated interest in multimodal data fusion as a strategy for improving clinical prediction performance [6,7]. Rather than treating individual data streams independently, multimodal frameworks seek to construct unified representations that capture complex interactions among physiological, behavioral, procedural, and organizational factors. Such approaches are particularly relevant in regional anesthesia contexts where patient outcomes are influenced by both clinical characteristics and procedural execution variables.

The emergence of multimodal fusion technologies coincides with broader transformations in healthcare digitization. Hospitals increasingly operate as interconnected information ecosystems in which patient care depends upon coordination among diverse technical infrastructures, professional groups, and governance mechanisms [8]. Consequently, the development of perioperative risk stratification systems should be understood not solely as a technical endeavor but as a sociotechnical challenge involving organizational adaptation, regulatory compliance, and ethical oversight.

This paper investigates the role of multimodal clinical data fusion in perioperative risk stratification for arthroscopic knee procedures performed under regional anesthesia. The discussion emphasizes architectural considerations, implementation challenges, and long-term sustainability requirements that influence successful adoption within contemporary healthcare systems.

2. Clinical and Operational Context of Perioperative Risk Stratification

Perioperative risk stratification serves multiple functions within surgical care pathways. Beyond predicting adverse clinical events, risk assessment informs resource allocation, staffing decisions, postoperative monitoring intensity, discharge planning, and patient counseling. In arthroscopic knee surgery, clinicians must evaluate risks associated with anesthesia, pain management, functional recovery, and potential complications while balancing efficiency objectives within increasingly constrained healthcare environments [9].

Regional anesthesia introduces unique considerations into this risk landscape. Femoral nerve blocks, sciatic nerve blocks, adductor canal blocks, and combined approaches have demonstrated significant benefits in perioperative analgesia and rehabilitation outcomes [10]. Clinical evidence suggests that combined femoral and sciatic nerve block strategies can provide effective perioperative anesthesia and postoperative pain control in knee arthroscopy settings [11]. However, variability in patient responses remains substantial. Differences in anatomical characteristics, comorbid conditions, inflammatory status, psychological factors, and procedural complexity contribute to diverse recovery trajectories.

Healthcare organizations increasingly recognize that perioperative outcomes emerge from interactions among numerous interconnected variables rather than isolated risk factors. For example, postoperative pain experiences may be influenced simultaneously by preexisting chronic pain conditions, psychological stress levels, anesthetic technique selection, intraoperative medication patterns, and rehabilitation adherence. Similar complexity characterizes complications such as delayed recovery, unplanned admissions, and prolonged opioid utilization.

From an operational perspective, risk stratification has become central to value-based healthcare delivery. Hospitals seek to optimize patient outcomes while controlling costs and maximizing resource utilization efficiency. Accurate prediction of postoperative recovery pathways enables more effective scheduling, bed management, staffing allocation, and follow-up planning. Consequently, risk stratification systems increasingly function as organizational infrastructure components rather than purely clinical decision-support tools.

The complexity of modern perioperative environments has exposed limitations of conventional risk assessment methodologies. Static scoring systems often fail to account for temporal changes in patient status and may not adequately capture nonlinear relationships among clinical variables. As healthcare institutions adopt more comprehensive digital infrastructures, opportunities emerge for developing adaptive risk models capable of continuously incorporating new information throughout the perioperative journey.

3. Multimodal Clinical Data Fusion Architecture

The concept of multimodal clinical data fusion centers on integrating diverse information sources into unified analytical frameworks. Within arthroscopic knee surgery environments, relevant data modalities encompass structured electronic health records, laboratory measurements, imaging studies, physiological monitoring streams, medication records, clinician narratives, wearable sensor outputs, and patient-reported outcome measures [12].

A comprehensive fusion architecture typically operates across multiple temporal phases. During preoperative evaluation, demographic characteristics, medical histories, laboratory findings, imaging assessments, medication exposures, and functional status indicators contribute foundational risk information. Intraoperatively, real-time physiological monitoring data, anesthetic administration records, procedural duration metrics, and clinician observations generate dynamic indicators of patient condition. Postoperatively, recovery assessments, pain scores, mobility measurements, rehabilitation progress indicators, and follow-up outcomes provide additional layers of information relevant to ongoing risk assessment.

One of the primary architectural challenges involves harmonizing data characterized by differing structures, scales, frequencies, and quality levels. Physiological monitoring systems may generate high-frequency time-series data, whereas laboratory tests appear intermittently and clinical narratives exist as unstructured text. Effective fusion frameworks require sophisticated interoperability mechanisms capable of transforming heterogeneous inputs into coherent representations suitable for predictive modeling [13].

Contemporary fusion architectures increasingly employ layered infrastructures that separate data ingestion, normalization, representation learning, prediction generation, and clinical interface functions. Such modular designs improve scalability and facilitate adaptation to evolving clinical requirements. Furthermore, layered architectures support governance objectives by enabling transparent monitoring of data quality, algorithm performance, and decision-support outputs.

Another critical consideration involves temporal alignment across modalities. Patient risk evolves continuously throughout perioperative care, necessitating analytical approaches capable of incorporating both historical and real-time information. Dynamic fusion models can capture evolving relationships among variables and generate updated risk estimates as new data become available. This capability is particularly valuable in regional anesthesia settings where physiological responses and analgesic effectiveness may change rapidly during recovery.

The emergence of foundation models and advanced representation learning techniques has expanded possibilities for multimodal integration [14]. These approaches enable extraction of meaningful patterns from complex datasets while reducing dependence on manually engineered features. However, increased model complexity also introduces challenges related to interpretability, computational requirements, and governance oversight.

4. Artificial Intelligence and Predictive Modeling Frameworks

Artificial intelligence provides the analytical foundation underlying contemporary multimodal fusion systems. Traditional statistical methods remain valuable for certain clinical prediction tasks, but the growing volume and complexity of healthcare data have encouraged adoption of machine learning approaches capable of identifying subtle interactions across numerous variables [15].

Within perioperative environments, predictive modeling objectives extend beyond single-outcome forecasting. Healthcare organizations increasingly seek comprehensive risk profiles encompassing postoperative pain trajectories, recovery duration, complication likelihood, readmission risk, opioid consumption patterns, and patient satisfaction outcomes. Multimodal fusion enables development of integrated prediction frameworks that account for interactions among these outcomes while generating more holistic assessments of patient status.

Deep learning architectures have demonstrated particular promise for multimodal healthcare applications because they can process diverse data types simultaneously. Structured clinical variables, physiological waveforms, medical images, and narrative documentation can be incorporated into unified analytical frameworks that capture cross-modal relationships inaccessible through conventional methods [16]. Such capabilities may improve prediction performance in complex perioperative scenarios characterized by multifactorial risk mechanisms.

However, predictive accuracy alone is insufficient for clinical adoption. Healthcare professionals must understand and trust model outputs before integrating them into decision-making processes. Consequently, explainability and interpretability have emerged as critical design requirements. Clinicians often require insights regarding the factors contributing to specific risk predictions, particularly when recommendations influence treatment planning or resource allocation decisions [17].

The tension between predictive performance and interpretability represents a persistent challenge in healthcare AI. Highly complex models may achieve superior accuracy but provide limited transparency regarding underlying reasoning processes. Simpler models often offer greater interpretability yet may fail to capture intricate relationships present within multimodal datasets. Successful implementation therefore requires careful balancing of technical capabilities and clinical usability considerations.

Robustness represents another essential requirement. Clinical environments are characterized by incomplete records, variable documentation practices, missing data, and changing patient populations. Predictive systems must maintain acceptable performance despite these realities. Moreover, healthcare institutions frequently experience workflow modifications, technology upgrades, and demographic shifts that can alter underlying data distributions. Continuous monitoring and model adaptation mechanisms are therefore necessary to ensure sustained reliability over time.

5. Governance, Fairness, and Ethical Infrastructure

The deployment of multimodal clinical data fusion systems within perioperative care environments raises important governance and ethical considerations that extend beyond technical performance. Healthcare organizations increasingly recognize that artificial intelligence systems function as institutional actors whose recommendations may influence clinical decisions, operational planning, and patient experiences. Consequently, governance frameworks must address issues of accountability, transparency, fairness, and regulatory compliance throughout the lifecycle of model development and deployment.

One of the most significant governance challenges involves data stewardship. Multimodal fusion systems depend upon aggregation of information originating from multiple clinical departments, information systems, and external sources. Electronic health records, imaging repositories, anesthesia information management systems, rehabilitation databases, and patient-generated health data frequently operate under different governance structures and quality standards. Effective implementation therefore requires coordinated institutional policies that establish clear responsibilities for data access, validation, security, and maintenance [18].

Data quality governance is particularly important because predictive performance is directly influenced by the reliability of underlying information assets. Missing values, documentation inconsistencies, coding variations, and incomplete records may introduce systematic biases that affect model outputs. In perioperative environments, these issues are amplified by the involvement of numerous clinical teams operating across different phases of care. Governance mechanisms must therefore include continuous auditing procedures capable of identifying data quality degradation before adverse impacts on prediction accuracy occur.

Algorithmic fairness has emerged as another critical concern in healthcare artificial intelligence. Numerous studies have demonstrated that predictive models may inadvertently reproduce historical inequities embedded within healthcare datasets [19]. In perioperative risk stratification, biased predictions could influence access to surgical services, postoperative monitoring resources, rehabilitation interventions, or pain management strategies. Such outcomes would undermine both clinical effectiveness and ethical legitimacy.

The complexity of multimodal fusion systems creates additional fairness challenges because bias may emerge through interactions among multiple data sources rather than individual variables alone. Socioeconomic status, geographic location, healthcare utilization history, and insurance-related factors may become indirectly encoded within seemingly neutral clinical indicators. Consequently, fairness evaluation requires comprehensive assessment across demographic groups, clinical populations, and institutional contexts.

Transparency constitutes another essential governance requirement. Clinicians, patients, and healthcare administrators increasingly expect meaningful explanations regarding how predictive recommendations are generated. While complete interpretability may not always be achievable in highly complex models, organizations should establish mechanisms that allow stakeholders to understand model objectives, limitations, data sources, and performance characteristics. Transparency supports trust formation and facilitates informed oversight of automated decision-support systems.

Regulatory compliance further shapes governance requirements. Healthcare institutions operate within evolving legal environments that increasingly address artificial intelligence deployment, patient privacy protection, and automated decision-making. Organizations implementing multimodal risk stratification systems must therefore develop governance structures capable of adapting to changing regulatory expectations while maintaining operational flexibility [20].

Importantly, governance should not be viewed solely as a compliance function. Effective governance infrastructures contribute directly to system sustainability by promoting stakeholder trust, facilitating organizational learning, and reducing implementation risks. Institutions that integrate governance considerations into system design processes are often better positioned to achieve long-term success than those that treat governance as an afterthought.

6. Deployment and Clinical Integration Challenges

Despite rapid advances in machine learning research, translation of multimodal fusion technologies into routine clinical practice remains challenging. Many predictive systems demonstrate strong performance under experimental conditions yet fail to achieve meaningful adoption within operational healthcare environments. Understanding this implementation gap requires examination of the broader sociotechnical context in which clinical technologies are deployed.

Workflow integration represents one of the most important determinants of adoption success. Perioperative clinicians operate within time-sensitive environments characterized by significant cognitive demands and complex coordination requirements. Decision-support systems that increase workload, disrupt established routines, or generate excessive alerts may encounter resistance regardless of their predictive capabilities [21]. Consequently, successful multimodal fusion platforms must be designed around clinical workflows rather than expecting workflows to adapt to technological constraints.

Human-AI collaboration provides a useful framework for understanding deployment requirements. Rather than replacing clinical expertise, multimodal risk stratification systems should augment professional decision-making by providing additional information and analytical support. This perspective recognizes that clinicians possess contextual knowledge, ethical judgment, and interpersonal skills that cannot be fully captured within computational models. Effective systems therefore facilitate collaboration between human expertise and machine intelligence rather than privileging either component exclusively.

Institutional readiness also influences implementation outcomes. Healthcare organizations differ substantially in terms of digital maturity, technical infrastructure, workforce capabilities, and leadership support. Hospitals with advanced interoperability frameworks and established data governance processes may be well positioned to adopt multimodal fusion technologies, whereas resource-constrained institutions may face significant implementation barriers. These disparities raise important questions regarding equitable access to emerging healthcare innovations.

The integration of perioperative risk stratification systems with existing health information technologies presents additional challenges. Many healthcare organizations continue to operate fragmented information environments composed of legacy systems developed by multiple vendors. Achieving seamless interoperability often requires substantial investments in data standardization, interface development, and infrastructure modernization [22].

Economic considerations further complicate deployment decisions. Although predictive analytics may generate long-term efficiency gains through improved outcomes and resource utilization, implementation costs can be substantial. Expenses associated with software development, infrastructure acquisition, workforce training, cybersecurity enhancement, and ongoing maintenance must be carefully balanced against anticipated benefits. Healthcare administrators therefore require robust evidence demonstrating organizational value before committing resources to large-scale deployment initiatives.

Clinical validation remains another crucial requirement. Models developed within specific institutions may not generalize effectively across different patient populations, healthcare settings, or operational environments. External validation studies are therefore necessary to establish reliability and identify contextual factors that influence performance. Continuous evaluation mechanisms should also be implemented to detect performance degradation following deployment.

Importantly, deployment should be understood as an ongoing process rather than a discrete event. Healthcare systems are dynamic environments characterized by changing patient populations, evolving clinical practices, and technological innovation. Sustainable implementation therefore requires continuous monitoring, adaptation, and stakeholder engagement throughout the operational lifespan of the system.

7. Sustainability, Resilience, and Policy Implications

The long-term value of multimodal clinical data fusion depends upon sustainability and resilience characteristics that extend beyond immediate predictive performance. Healthcare organizations increasingly recognize that artificial intelligence systems must remain reliable, adaptable, and economically viable over extended periods if they are to contribute meaningfully to patient care.

Sustainability begins with infrastructure design. Multimodal fusion systems frequently require substantial computational resources, data storage capacity, and network connectivity. As healthcare datasets continue to expand, organizations must develop scalable architectures capable of accommodating future growth without excessive cost escalation. Cloud-based infrastructures, federated learning approaches, and distributed analytics frameworks offer potential pathways toward sustainable deployment models [23].

Resilience represents an equally important consideration. Healthcare delivery environments are subject to numerous disruptions, including cybersecurity incidents, equipment failures, staffing shortages, and public health emergencies. Risk stratification systems must therefore be designed with redundancy, fault tolerance, and recovery capabilities that ensure continuity of operation during adverse conditions. The COVID-19 pandemic demonstrated the importance of resilient digital infrastructures capable of supporting healthcare adaptation under rapidly changing circumstances [24].

Policy implications emerge at multiple levels. At the institutional level, healthcare organizations must establish policies governing data sharing, model oversight, performance evaluation, and accountability. At the regional and national levels, policymakers face decisions regarding regulatory frameworks, reimbursement mechanisms, interoperability standards, and ethical requirements. These policy choices will significantly influence the pace and direction of healthcare AI adoption.

The development of standardized data models represents a particularly important policy priority. Fragmentation across healthcare information systems continues to impede large-scale data integration efforts. Common standards can facilitate interoperability, reduce implementation costs, and promote broader access to advanced analytical capabilities. Standardization may also enhance reproducibility and support collaborative research initiatives spanning multiple institutions.

Workforce development constitutes another critical policy domain. Successful implementation of multimodal fusion technologies requires professionals possessing expertise in clinical medicine, data science, informatics, ethics, and organizational management. Educational institutions and healthcare organizations must therefore invest in interdisciplinary training programs capable of preparing future practitioners for increasingly data-driven clinical environments.

From a societal perspective, multimodal perioperative risk stratification reflects broader transformations occurring throughout healthcare systems. The transition toward learning health systems, precision medicine, and digitally enabled care models is reshaping relationships among patients, clinicians, organizations, and technologies. Effective policy frameworks should support innovation while ensuring that technological advancement remains aligned with public health objectives and patient welfare.

8. Future Research Directions

Future research should move beyond narrow evaluations of predictive accuracy toward comprehensive assessments of real-world clinical impact. While machine learning performance metrics remain important, healthcare organizations increasingly require evidence demonstrating improvements in patient outcomes, workflow efficiency, resource utilization, and health equity.

One promising direction involves the integration of emerging data sources into multimodal fusion frameworks. Wearable devices, mobile health applications, rehabilitation sensors, and patient-generated health records offer opportunities to extend risk assessment beyond traditional clinical settings. Such capabilities may enable continuous monitoring of recovery trajectories and facilitate more personalized postoperative care strategies [25].

Digital twin technologies also represent an emerging area of interest. By constructing individualized computational representations of patients and care pathways, digital twins may enable simulation-based risk assessment and personalized intervention planning. Integration of digital twin concepts with multimodal fusion architectures could significantly enhance predictive and prescriptive capabilities within perioperative medicine.

Federated learning approaches warrant further investigation because they address important privacy and governance challenges associated with cross-institutional collaboration. Rather than centralizing sensitive patient data, federated frameworks enable distributed model development while preserving local data control. Such approaches may facilitate creation of more robust and generalizable risk stratification systems capable of learning from diverse populations and clinical environments.

Future research should also examine the organizational consequences of AI-enabled risk stratification. Questions regarding workforce adaptation, professional roles, patient trust, and institutional governance remain insufficiently understood. Addressing these issues will require interdisciplinary collaboration among clinicians, engineers, social scientists, ethicists, and policymakers.

Ultimately, the future of multimodal clinical data fusion depends upon the ability of healthcare systems to integrate technological innovation with human-centered care principles. Research agendas that balance technical advancement with ethical, organizational, and societal considerations are likely to generate the most sustainable and impactful outcomes.

9. Conclusion

Multimodal clinical data fusion represents a transformative approach to perioperative risk stratification in arthroscopic knee procedures performed under regional anesthesia. By integrating heterogeneous data sources across preoperative, intraoperative, and postoperative phases of care, multimodal frameworks offer opportunities to develop more comprehensive and dynamic representations of patient risk than conventional assessment methods. The growing availability of electronic health records, physiological monitoring systems, imaging data, and patient-reported outcomes provides a rich foundation for advanced predictive analytics capable of supporting individualized clinical decision-making.

However, successful implementation requires more than sophisticated machine learning algorithms. Effective deployment depends upon robust data governance, interoperability infrastructures, workflow integration strategies, fairness protections, and organizational readiness. Healthcare institutions must address challenges related to transparency, accountability, sustainability, and resilience while ensuring that technological innovation remains aligned with patient-centered care objectives.

The future development of multimodal perioperative risk stratification systems will likely be shaped by advances in artificial intelligence, federated learning, digital health technologies, and healthcare policy frameworks. As these technologies mature, their impact will increasingly depend upon the quality of sociotechnical ecosystems supporting deployment and continuous adaptation. Organizations that successfully integrate technical innovation with governance excellence and clinical collaboration will be best positioned to realize the benefits of data-driven perioperative care.

In this context, multimodal data fusion should be viewed not merely as a predictive technology but as an enabling infrastructure for the next generation of intelligent, equitable, and learning-oriented healthcare systems.

References

1. Ackerman, W. E., Ahmad, M., & Sellers, S. (2005). The efficacy of lumbar plexus block for outpatient knee arthroscopy. Anesthesia & Analgesia, 101(2), 603–605.

2. Abdallah, F. W., & Brull, R. (2011). Is sciatic nerve block advantageous when combined with femoral nerve block for postoperative analgesia following knee surgery? Canadian Journal of Anesthesia, 58(11), 1040–1046.

3. Barrington, M. J., & Kluger, R. (2013). Ultrasound guidance reduces the risk of local anesthetic systemic toxicity following peripheral nerve blockade. Regional Anesthesia and Pain Medicine, 38(4), 289–299.

4. Glance, L. G., Lustik, S. J., Hannan, E. L., Osler, T. M., Mukamel, D. B., Qian, F., & Dick, A. W. (2012). The surgical mortality probability model. Anesthesiology, 117(4), 696–705.

5. Kheterpal, S., Martin, L., Shanks, A. M., & Tremper, K. K. (2009). Prediction and outcomes of impossible mask ventilation. Anesthesiology, 110(4), 891–897.

6. 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.

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

8. Friedman, C. P., Rubin, J. C., & Sullivan, K. J. (2017). Toward an information infrastructure for global health improvement. Yearbook of Medical Informatics, 26(1), 16–23.

9. Memtsoudis, S. G., Poeran, J., Cozowicz, C., Zubizarreta, N., Ozbek, U., Mazumdar, M., & Mariano, E. R. (2018). The impact of peripheral nerve blocks on perioperative outcome in hip and knee arthroplasty. Anesthesiology, 128(4), 719–727.

10. Mariano, E. R., & Ilfeld, B. M. (2018). Contemporary peripheral nerve blocks for orthopedic surgery. Anesthesiology Clinics, 36(3), 437–452.

11. 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.

12. Jensen, P. B., Jensen, L. J., & Brunak, S. (2012). Mining electronic health records. Nature Reviews Genetics, 13(6), 395–405.

13. Miotto, R., Wang, F., Wang, S., Jiang, X., & Dudley, J. T. (2018). Deep learning for healthcare. Briefings in Bioinformatics, 19(6), 1236–1246.

14. Moor, M., Banerjee, O., Abad, Z. S. H., Krumholz, H. M., Leskovec, J., Topol, E. J., & Rajpurkar, P. (2023). Foundation models for generalist medical artificial intelligence. Nature, 616(7956), 259–265.

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

16. Huang, S. C., Pareek, A., Seyyedi, S., Banerjee, I., & Lungren, M. P. (2020). Fusion of medical imaging and electronic health records using deep learning. NPJ Digital Medicine, 3(136), 1–9.

17. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv Preprint arXiv:1702.08608.

18. Raghupathi, W., & Raghupathi, V. (2014). Big data analytics in healthcare. Health Information Science and Systems, 2(3), 1–10.

19. 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.

20. World Health Organization. (2021). Ethics and governance of artificial intelligence for health. World Health Organization.

21. Sutton, R. T., Pincock, D., Baumgart, D. C., Sadowski, D. C., Fedorak, R. N., & Kroeker, K. I. (2020). An overview of clinical decision support systems. NPJ Digital Medicine, 3(17), 1–10.

22. Adler-Milstein, J., Holmgren, A. J., Kralovec, P., Worzala, C., Searcy, T., & Patel, V. (2017). Electronic health record adoption in US hospitals. Health Affairs, 36(8), 1416–1423.

23. 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), 1–7.

24. Budd, J., Miller, B. S., Manning, E. M., Lampos, V., Zhuang, M., Edelstein, M., Rees, G., Emery, V. C., Stevens, M. M., Keegan, N., Short, M. J., Pillay, D., Manley, E., Cox, I. J., Heymann, D., Johnson, A. M., & McKendry, R. A. (2020). Digital technologies in the public-health response to COVID-19. Nature Medicine, 26(8), 1183–1192.

25. Steinhubl, S. R., Muse, E. D., & Topol, E. J. (2015). The emerging field of mobile health. Science Translational Medicine, 7(283), 283rv3.