AI-Assisted Personalized Anesthetic Planning for Arthroscopic Knee Surgery Based on Peripheral Nerve Block Characteristics

Gachael Enderson
Department of Biomedical Informatics, University of Arkansas, Fayetteville, AR, USA.
mjanderson@uark.edu

Emily R. Foster
School of Computing and Information Sciences, University of North Florida, Jacksonville, FL, USA.
efoster@unf.edu

Letian Zou
Department of Industrial and Systems Engineering, University of Alabama in Huntsville, Huntsville, AL, USA.
dkwilson@uah.edu

Abstract

Arthroscopic knee surgery has become one of the most frequently performed orthopedic procedures worldwide due to its minimally invasive characteristics, reduced hospitalization requirements, and favorable postoperative recovery profiles. Despite these advantages, perioperative pain management remains a substantial clinical challenge because postoperative analgesic responses vary significantly across patients. Peripheral nerve block techniques, particularly femoral nerve block, sciatic nerve block, and combined approaches, have demonstrated substantial benefits in improving pain control and reducing opioid consumption. However, existing anesthetic planning strategies often rely heavily on generalized clinical guidelines and practitioner experience, limiting the ability to accommodate individual variability in anatomical structures, physiological conditions, procedural complexity, and recovery trajectories.

Recent advances in artificial intelligence have created opportunities for personalized anesthetic planning through the integration of heterogeneous clinical data sources. This study proposes a system-level framework for AI-assisted anesthetic planning in arthroscopic knee surgery based on peripheral nerve block characteristics. Rather than focusing exclusively on prediction accuracy, the paper examines the broader socio-technical ecosystem required to support personalized decision-making. The proposed framework integrates multimodal patient information, perioperative monitoring data, historical outcomes, and institutional knowledge repositories to generate individualized anesthetic recommendations. Particular attention is given to infrastructure architecture, governance mechanisms, model robustness, fairness considerations, deployment constraints, and long-term sustainability.

The analysis demonstrates that successful implementation requires coordination across clinical, computational, organizational, and regulatory domains. AI-assisted planning has the potential to improve analgesic effectiveness, reduce adverse events, optimize resource allocation, and support evidence-based anesthesia practice. Nevertheless, challenges related to data quality, algorithmic transparency, interoperability, clinician trust, and institutional governance remain significant. The paper concludes by outlining future research directions aimed at creating adaptive, trustworthy, and scalable anesthetic intelligence platforms capable of supporting personalized perioperative care.

Keywords

artificial intelligence, personalized anesthesia, arthroscopic knee surgery, peripheral nerve block, clinical decision support, healthcare informatics, perioperative management.

1. Introduction

Arthroscopic knee surgery has emerged as a cornerstone of contemporary orthopedic practice, serving a wide range of diagnostic and therapeutic purposes including ligament reconstruction, meniscal repair, cartilage restoration, and synovial interventions. The procedure's minimally invasive nature has contributed to increased utilization across healthcare systems, enabling shorter hospital stays and accelerated rehabilitation pathways. Nevertheless, postoperative pain remains one of the most significant determinants of patient satisfaction, functional recovery, healthcare utilization, and overall surgical outcomes [1].

Peripheral nerve block techniques have transformed perioperative analgesic management in knee arthroscopy. Clinical evidence has demonstrated that femoral nerve block, sciatic nerve block, and combined blockade strategies can substantially improve pain control while reducing systemic opioid requirements [2]. These approaches are particularly valuable in ambulatory surgery settings where rapid recovery and early discharge are prioritized. However, substantial inter-patient variability continues to challenge the effectiveness of standardized anesthetic protocols.

Traditional anesthetic planning often relies on population-level guidelines derived from aggregated clinical evidence. While such guidelines provide important foundations for safe practice, they may not adequately address individual differences in age, body composition, neurological characteristics, comorbidities, pain sensitivity, previous anesthetic experiences, psychological status, and procedural complexity. Consequently, patients receiving identical anesthetic interventions may experience markedly different analgesic outcomes.

Artificial intelligence technologies offer a potential pathway toward individualized anesthetic planning by leveraging large-scale clinical datasets and advanced analytical capabilities [3]. Machine learning systems can identify complex patterns that may not be readily apparent through conventional statistical methods or human observation alone. When integrated into perioperative workflows, such systems can support clinicians in selecting personalized nerve block strategies, optimizing anesthetic dosages, anticipating complications, and forecasting recovery trajectories.

The significance of AI-assisted anesthetic planning extends beyond technical prediction tasks. It represents a broader transformation of healthcare decision-making infrastructures, where clinical expertise, digital intelligence, organizational governance, and regulatory oversight become increasingly interconnected. Understanding these interdependencies is essential for developing sustainable and trustworthy implementations capable of generating long-term clinical value.

This paper examines the development of AI-assisted personalized anesthetic planning systems for arthroscopic knee surgery based on peripheral nerve block characteristics. Emphasis is placed on system architecture, deployment considerations, governance mechanisms, fairness implications, and future healthcare integration strategies.

2. Clinical Foundations of Personalized Peripheral Nerve Block Planning

Peripheral nerve blocks occupy a unique position within modern anesthesia practice because they influence multiple dimensions of perioperative care simultaneously. Effective blockade contributes not only to pain reduction but also to enhanced mobility, decreased opioid consumption, improved patient satisfaction, and reduced healthcare costs [4]. Consequently, selecting an appropriate nerve block strategy represents a multidimensional optimization problem rather than a single therapeutic decision.

Femoral nerve block has traditionally served as a widely adopted approach for knee arthroscopy due to its effectiveness in reducing anterior knee pain. However, excessive motor blockade may delay ambulation and increase fall risk among certain patient populations. Conversely, sciatic nerve block can provide additional posterior knee analgesia but may introduce different recovery considerations and procedural complexities [5]. Combined approaches often offer superior pain relief but may increase procedural time and resource requirements.

The complexity of anesthetic planning becomes more pronounced when considering patient-specific variables. Elderly individuals may exhibit altered pharmacodynamic responses that affect anesthetic duration and recovery characteristics. Patients with obesity may present anatomical challenges influencing ultrasound-guided block performance. Individuals with diabetes may demonstrate variations in nerve physiology that affect block effectiveness and complication risk. Similarly, patients with chronic pain conditions may exhibit altered nociceptive processing mechanisms that influence postoperative analgesic requirements [6].

Clinical decision-making therefore involves balancing competing objectives. Maximizing analgesic effectiveness may conflict with goals related to early mobilization. Minimizing opioid use may require more extensive regional anesthesia interventions. Resource constraints may influence the feasibility of certain techniques within specific institutional environments. These competing priorities highlight the limitations of simplistic decision frameworks and create opportunities for AI-assisted optimization.

Importantly, personalization should not be interpreted merely as individualized prediction. Effective personalized anesthesia requires continuous adaptation throughout the perioperative continuum. Preoperative assessments, intraoperative monitoring, and postoperative recovery data collectively contribute to a dynamic understanding of patient status. AI systems capable of integrating information across these stages may provide substantially greater value than isolated predictive models operating at a single point in time.

Research involving knee arthroscopy has increasingly emphasized the importance of combined femoral and sciatic nerve blockade strategies in improving perioperative outcomes [7]. Such findings underscore the need for sophisticated decision-support systems capable of translating heterogeneous clinical evidence into patient-specific recommendations while accounting for local institutional practices and resource availability.

3. System Architecture for AI-Assisted Anesthetic Planning

The development of AI-assisted anesthetic planning requires a robust socio-technical architecture capable of integrating diverse data streams across multiple healthcare environments. Unlike isolated machine learning applications, personalized anesthetic intelligence systems must function within complex clinical ecosystems characterized by heterogeneous technologies, evolving workflows, and stringent safety requirements.

At the foundational level, data infrastructure serves as the primary enabling component. Relevant information originates from electronic health records, perioperative management systems, imaging repositories, physiological monitoring devices, laboratory databases, and patient-reported outcome platforms [8]. These sources often exhibit significant variability in structure, quality, completeness, and semantic consistency. Consequently, data harmonization becomes a critical prerequisite for effective AI deployment.

A multimodal integration layer facilitates the aggregation of structured and unstructured clinical information. Structured variables may include demographic characteristics, comorbidities, medication histories, procedural details, and historical anesthetic outcomes. Unstructured information may originate from clinical narratives, imaging interpretations, nursing documentation, and patient feedback records. Advanced natural language processing and representation learning techniques can assist in transforming these heterogeneous data sources into unified analytical frameworks [9].

The intelligence layer constitutes the core analytical component of the architecture. Rather than relying on a single predictive model, contemporary systems increasingly employ ensembles of specialized models addressing distinct clinical objectives. One model may estimate postoperative pain severity, another may predict block duration, while additional models evaluate complication risk, recovery timelines, or opioid consumption patterns. Integrating these outputs allows for comprehensive decision support that reflects the multidimensional nature of anesthetic planning.

Equally important is the decision orchestration layer responsible for translating analytical outputs into clinically meaningful recommendations. Clinicians require actionable insights rather than abstract probability estimates. Therefore, recommendation systems must present information in interpretable formats that align with existing workflows and cognitive processes. Explainability mechanisms become particularly important in high-stakes healthcare environments where transparency directly influences clinician trust and adoption [10].

Infrastructure scalability represents another critical architectural consideration. Large healthcare organizations may process thousands of surgical cases annually, requiring real-time analytical capabilities capable of supporting simultaneous clinical workflows. Cloud-native architectures, edge computing approaches, and hybrid deployment models offer potential solutions for balancing computational efficiency, latency requirements, and data governance obligations.

4. Governance, Trust, and Ethical Dimensions

The successful deployment of AI-assisted anesthetic planning depends as much on governance frameworks as on algorithmic performance. Healthcare organizations operate within highly regulated environments where patient safety, accountability, transparency, and ethical responsibility constitute fundamental institutional obligations.

Algorithmic transparency occupies a central position within trustworthy healthcare AI. Clinicians are unlikely to adopt recommendations generated through opaque computational processes, particularly when patient outcomes are directly affected. Explainable AI methodologies can provide insight into the factors influencing specific recommendations, thereby facilitating informed clinical judgment and enhancing professional confidence [11].

Fairness considerations are equally important. Clinical datasets frequently contain demographic imbalances that may introduce unintended biases into predictive models. Variations in healthcare access, socioeconomic status, geographic location, race, age, and gender can influence data distributions and subsequently affect algorithmic performance [12]. Without careful evaluation, AI systems may inadvertently reinforce existing healthcare disparities rather than reduce them.

Governance structures must therefore incorporate continuous monitoring mechanisms capable of identifying performance degradation, emerging biases, and unintended consequences. Regular auditing procedures should assess not only predictive accuracy but also equity outcomes, safety indicators, and operational impacts. Such monitoring transforms AI deployment from a one-time implementation activity into an ongoing organizational responsibility.

Regulatory considerations further complicate implementation efforts. Healthcare AI systems increasingly operate within evolving legal frameworks emphasizing transparency, accountability, data protection, and clinical validation. Institutions deploying personalized anesthetic planning systems must establish clear policies governing model development, validation procedures, clinician oversight, and incident response protocols [13].

Human-AI collaboration represents another essential governance dimension. Personalized anesthetic planning systems should augment rather than replace clinical expertise. Maintaining appropriate levels of clinician involvement helps preserve professional accountability while ensuring that contextual factors not captured within data models continue to inform decision-making. Effective governance frameworks recognize that AI recommendations constitute one component of a broader clinical reasoning process rather than definitive directives.

5. Robustness and Clinical Deployment Challenges

Although predictive performance frequently dominates discussions surrounding healthcare artificial intelligence, real-world deployment success depends heavily upon robustness under diverse operational conditions. Anesthetic planning systems must function reliably across different patient populations, institutional environments, clinical workflows, and technological infrastructures. The transition from research prototypes to operational clinical systems therefore introduces challenges that extend far beyond model development.

One of the most significant concerns involves dataset shift. Models trained using historical data may encounter future patient populations whose characteristics differ substantially from those represented during development. Changes in surgical techniques, anesthetic protocols, institutional policies, patient demographics, and healthcare delivery patterns can gradually reduce model effectiveness over time. In the context of peripheral nerve block planning, even relatively small modifications in ultrasound guidance techniques or postoperative rehabilitation protocols may alter relationships between predictors and outcomes [14]. Consequently, continuous monitoring and periodic recalibration become essential components of sustainable deployment strategies.

Data quality presents another persistent challenge. Electronic health records often contain missing values, inconsistent coding practices, documentation errors, and heterogeneous terminology. While controlled research datasets may undergo extensive preprocessing, operational healthcare environments rarely achieve comparable levels of standardization. AI systems must therefore be designed with resilience mechanisms capable of managing incomplete and imperfect information without generating unreliable recommendations. Robust data validation pipelines, uncertainty quantification methods, and confidence reporting frameworks contribute significantly to maintaining safe operational performance.

Clinical workflow integration represents an equally important determinant of deployment success. Many healthcare AI initiatives fail not because predictive models are inaccurate, but because recommendations arrive at inappropriate moments, disrupt existing practices, or impose excessive cognitive burdens on clinicians. Personalized anesthetic planning systems must align with established perioperative processes, supporting rather than complicating clinical decision-making. Recommendation interfaces should provide concise, interpretable, and contextually relevant information that complements existing professional expertise [15].

Institutional variation further complicates deployment efforts. Hospitals differ substantially in staffing models, technological capabilities, patient populations, and resource availability. An anesthetic planning framework that performs effectively within a large academic medical center may require substantial adaptation before deployment in smaller community hospitals or ambulatory surgical facilities. Scalable architectures capable of accommodating local customization while preserving core analytical functionality are therefore particularly valuable.

Cybersecurity considerations have also become increasingly important as healthcare systems undergo digital transformation. AI-assisted anesthetic planning platforms frequently depend upon interconnected databases, cloud computing resources, and real-time communication networks. These dependencies introduce potential vulnerabilities that may compromise patient privacy, disrupt clinical operations, or undermine institutional trust. Comprehensive security architectures incorporating encryption, access controls, anomaly detection, and incident response protocols are essential components of responsible deployment strategies.

From an organizational perspective, clinician acceptance remains a critical factor influencing long-term success. Healthcare professionals often exhibit understandable caution toward technologies that affect patient care decisions. Building trust requires rigorous validation, transparent communication, participatory design approaches, and demonstrable clinical value. Institutions that actively involve anesthesiologists, surgeons, nurses, and informatics specialists throughout system development and implementation processes are generally better positioned to achieve meaningful adoption outcomes.

6. Infrastructure Sustainability and Health System Integration

The long-term value of AI-assisted anesthetic planning depends upon its ability to function as part of a sustainable healthcare infrastructure rather than as an isolated technological intervention. Sustainability encompasses technical, economic, organizational, and societal dimensions that collectively determine whether innovations can generate enduring benefits across healthcare systems.

From a technical perspective, sustainable infrastructures require interoperability. Modern healthcare organizations utilize numerous digital platforms developed by different vendors and implemented over extended periods. Personalized anesthetic planning systems must interact effectively with electronic health records, operating room management platforms, pharmacy systems, imaging repositories, and postoperative monitoring applications. Interoperability standards facilitate information exchange while reducing implementation complexity and supporting future scalability [16].

Economic sustainability is equally important. While AI technologies often promise efficiency gains and improved outcomes, implementation costs can be substantial. Infrastructure investments may include computational resources, software development, integration activities, cybersecurity enhancements, workforce training, and ongoing maintenance. Healthcare organizations must therefore evaluate not only initial deployment expenses but also long-term operational requirements. Demonstrating measurable improvements in outcomes, resource utilization, patient satisfaction, and complication reduction can help justify continued investment.

Workforce sustainability represents another critical consideration. The emergence of AI-assisted clinical decision support creates new professional responsibilities related to model oversight, data governance, performance evaluation, and system maintenance. Developing these capabilities requires educational initiatives that extend beyond traditional clinical training. Future anesthesiology professionals may increasingly require competencies in data literacy, algorithmic interpretation, and digital healthcare management [17].

Environmental sustainability has also emerged as an important topic within healthcare technology discussions. Large-scale machine learning systems can consume significant computational resources, particularly during model training and continuous updating processes. As healthcare organizations pursue broader sustainability objectives, attention to computational efficiency and responsible infrastructure design may become increasingly relevant.

At the health system level, personalized anesthetic planning offers opportunities to improve resource allocation and operational efficiency. More accurate prediction of postoperative pain trajectories may support optimized staffing decisions, recovery room management, and discharge planning. Improved analgesic effectiveness may reduce readmissions, decrease opioid-related complications, and facilitate earlier functional recovery. These system-wide benefits highlight the importance of evaluating AI implementations from a population health perspective rather than focusing exclusively on individual patient outcomes.

Healthcare systems also benefit from knowledge accumulation effects. As anesthetic planning platforms process increasing volumes of clinical data, they can contribute to institutional learning by identifying emerging patterns, evaluating protocol effectiveness, and supporting continuous quality improvement initiatives. This capacity transforms AI systems from passive decision-support tools into active components of organizational intelligence infrastructures.

7. Comparative Analysis with Other AI-Enabled Perioperative Systems

The evolution of AI-assisted anesthetic planning can be better understood by examining its relationship to broader developments in perioperative intelligence systems. Across healthcare domains, artificial intelligence has increasingly been applied to risk stratification, outcome prediction, resource management, imaging interpretation, and clinical workflow optimization [18]. Personalized peripheral nerve block planning represents one component within this larger transformation of perioperative care.

Comparisons with AI-driven surgical risk prediction systems reveal both similarities and important distinctions. Like anesthetic planning platforms, risk prediction tools integrate heterogeneous patient information to support clinical decision-making. However, anesthetic planning typically involves more dynamic interactions between physiological processes, procedural characteristics, and patient-specific recovery objectives. Consequently, decision-support frameworks must account for a broader range of contextual factors and competing priorities.

Analogous developments can be observed in critical care medicine, where predictive analytics increasingly support early detection of physiological deterioration. Both domains require continuous monitoring, uncertainty management, and integration into high-stakes clinical environments. Lessons learned from intensive care deployments emphasize the importance of transparency, workflow alignment, and human oversight, all of which are directly applicable to personalized anesthesia systems [19].

AI-enabled precision medicine initiatives provide another useful comparison. Precision oncology, for example, leverages genomic information to personalize therapeutic interventions. While anesthetic planning generally relies more heavily on physiological, procedural, and historical clinical data than genomic information, both fields share a common objective: tailoring medical decisions to individual patient characteristics. The success of precision medicine initiatives demonstrates the broader feasibility of personalized healthcare approaches while also highlighting challenges related to data integration, evidence generation, and equitable access.

Importantly, anesthetic planning differs from many other healthcare AI applications because recommendations influence immediate perioperative events with direct implications for patient comfort, mobility, and safety. The temporal proximity between decision and outcome creates opportunities for rapid feedback and continuous learning. Such characteristics may enable faster model refinement and more responsive adaptation than is possible in domains where outcomes emerge over much longer time horizons.

Cross-domain comparisons also reveal the importance of governance consistency. Whether applied to radiology, oncology, intensive care, or anesthesiology, successful AI implementations share common requirements involving transparency, accountability, fairness, validation, and clinician engagement. These recurring themes suggest that future progress may depend as much upon institutional learning and governance innovation as upon advances in machine learning methodology itself.

8. Future Directions

The future development of AI-assisted personalized anesthetic planning is likely to be shaped by advances across multiple technological and organizational domains. One promising direction involves the integration of real-time physiological sensing technologies capable of providing continuous information regarding patient responses before, during, and after nerve block administration. Such capabilities may enable adaptive anesthetic strategies that evolve dynamically in response to changing clinical conditions rather than relying solely on preoperative predictions.

Digital twin technologies represent another emerging opportunity. By creating computational representations of individual patients that continuously incorporate new clinical information, healthcare systems may eventually simulate alternative anesthetic strategies before implementation. These virtual environments could support more sophisticated personalization while reducing uncertainty associated with complex clinical decisions [20].

Federated learning approaches may further enhance model development by enabling collaboration across institutions without requiring centralized data sharing. Given increasing concerns regarding privacy, security, and regulatory compliance, federated architectures offer a potentially valuable pathway toward large-scale evidence generation while preserving local data governance. Such approaches could be particularly beneficial for anesthetic planning because outcome variability often necessitates large and diverse datasets.

Multimodal foundation models may also transform future capabilities. Current systems frequently rely on specialized models addressing specific prediction tasks. Emerging architectures capable of integrating text, images, physiological signals, clinical histories, and procedural information within unified representations may support more comprehensive decision-making frameworks. These developments could improve both predictive performance and contextual understanding.

Equity considerations will likely become increasingly important as deployment expands. Future research should prioritize evaluation across diverse demographic populations, healthcare settings, and socioeconomic environments. Ensuring that personalized anesthetic planning contributes to reducing rather than exacerbating healthcare disparities represents a major ethical and policy objective.

Finally, the evolution of regulatory frameworks will significantly influence future implementation pathways. As healthcare authorities establish more comprehensive standards governing clinical AI systems, developers and healthcare organizations will need to balance innovation with accountability. Transparent validation methodologies, post-deployment surveillance mechanisms, and standardized reporting practices may become essential requirements for widespread adoption.

9. Conclusion

AI-assisted personalized anesthetic planning based on peripheral nerve block characteristics represents a significant advancement in the evolution of perioperative care for arthroscopic knee surgery. By integrating heterogeneous clinical data, historical outcomes, physiological information, and institutional knowledge resources, intelligent decision-support systems can help address longstanding challenges associated with variability in analgesic effectiveness and recovery outcomes.

The value of such systems extends beyond predictive accuracy. Successful implementation requires comprehensive socio-technical infrastructures encompassing data integration, interoperability, governance, cybersecurity, clinician engagement, fairness assessment, and organizational learning. Personalized anesthetic planning therefore should be understood not merely as a machine learning application but as a component of a broader transformation toward intelligent healthcare ecosystems.

Although substantial challenges remain, including issues related to transparency, robustness, regulatory compliance, and equitable deployment, the potential benefits are considerable. Improved pain management, reduced opioid dependence, enhanced operational efficiency, and more individualized patient care collectively support continued investment in this area. Future developments involving federated learning, digital twins, real-time sensing technologies, and multimodal foundation models may further expand capabilities and accelerate clinical impact.

Ultimately, the successful integration of artificial intelligence into anesthetic planning will depend upon balancing technological innovation with ethical responsibility, clinical expertise, and institutional governance. Through such balanced approaches, personalized anesthesia has the potential to become a foundational element of next-generation perioperative medicine.

References

1. [1] Kehlet, H., & Dahl, J. B. (2003). Anaesthesia, surgery, and challenges in postoperative recovery. The Lancet, 362(9399), 1921–1928.
2. [2] Hadzic, A. (2017). Textbook of regional anesthesia and acute pain management (2nd ed.). McGraw-Hill Education.
3. [3] Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. Nature Medicine, 25(1), 44–56.
4. [4] Mariano, E. R., & Ilfeld, B. M. (2018). Contemporary peripheral nerve blocks for orthopedic surgery. Anesthesiology Clinics, 36(3), 377–390.
5. [5] 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), 1049–1061.
6. [6] Gan, T. J. (2017). Poorly controlled postoperative pain: Prevalence, consequences, and prevention. Journal of Pain Research, 10, 2287–2298.
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. [8] Jensen, P. B., Jensen, L. J., & Brunak, S. (2012). Mining electronic health records: Towards better research applications and clinical care. Nature Reviews Genetics, 13(6), 395–405.
9. [9] Esteva, A., Robicquet, A., Ramsundar, B., et al. (2019). A guide to deep learning in healthcare. Nature Medicine, 25(1), 24–29.
10. [10] Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv Preprint arXiv:1702.08608.
11. [11] Amann, J., Blasimme, A., Vayena, E., Frey, D., & Madai, V. I. (2020). Explainability for artificial intelligence in healthcare: A multidisciplinary perspective. BMC Medical Informatics and Decision Making, 20(1), 310.
12. [12] 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.
13. [13] European Commission. (2024). Artificial Intelligence Act: Regulation laying down harmonised rules on artificial intelligence. Publications Office of the European Union.
14. [14] Finlayson, S. G., Subbaswamy, A., Singh, K., et al. (2021). The clinician and dataset shift in artificial intelligence. New England Journal of Medicine, 385(3), 283–286.
15. [15] Sutton, R. T., Pincock, D., Baumgart, D. C., et al. (2020). An overview of clinical decision support systems: Benefits, risks, and strategies for success. NPJ Digital Medicine, 3(17), 1–10.
16. [16] Mandel, J. C., Kreda, D. A., Mandl, K. D., Kohane, I. S., & Ramoni, R. B. (2016). SMART on FHIR: A standards-based, interoperable apps platform for electronic health records. Journal of the American Medical Informatics Association, 23(5), 899–908.

17. [17] Char, D. S., Shah, N. H., & Magnus, D. (2018). Implementing machine learning in health care: Addressing ethical challenges. New England Journal of Medicine, 378(11), 981–983.
18. [18] Hashimoto, D. A., Rosman, G., Rus, D., & Meireles, O. R. (2018). Artificial intelligence in surgery: Promises and perils. Annals of Surgery, 268(1), 70–76.
19. [19] Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C., & Faisal, A. A. (2018). The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. Nature Medicine, 24(11), 1716–1720.
20. [20] Corral-Acero, J., Margara, F., Marciniak, M., et al. (2020). The 'Digital Twin' to enable the vision of precision cardiology. European Heart Journal, 41(48), 4556–4564.