npj | digital medicine

News & Views

Published in partnership with Seoul National University Bundang Hospital

https://doi.org/10.1038/s41746-025-01621-2

Advancing cardiovascular care through actionable AI innovation

Giuseppe Biondi-Zoccai, Arjun Mahajan, Dylan Powell, Mariangela Peruzzi, Roberto Carnevale & Giacomo Frati

Despite significant advances, the prevention and management of cardiovascular disease remain challenging, especially for ischemic heart disease (IHD). Current clinical decision-making relies heavily on physician expertise, guideline-directed therapies, and static risk scores, which often inadequately accommodate individual patient complexity. Machine learning (ML) and artificial intelligence (AI), particularly reinforcement learning (RL), may augment current physician-driven approaches and provide enhanced cardiovascular disease prevention and management. Indeed, offline RL refers to a class of ML algorithms that learn optimal decision-making policies from a fixed dataset of previously collected experiences—such as electronic health records or registries—without the need for active, real-time interaction with the clinical environment. This approach enables the safe development of treatment strategies in high-stakes domains where experimentation on live patients could be unethical or impractical. Notably, offline RL models hold the promise of optimizing decision-making in complex clinical settings, such as revascularization strategies for coronary artery disease. However, challenges remain in integrating AI into practice, ensuring interpretability, maintaining performance, and proving cost-effectiveness. Ultimately, validation, integration, and collaboration among clinicians, researchers, and policymakers are crucial for transforming AI-driven solutions into practical, patient-centered cardiovascular care improvements, pending prospective (and hopefully randomized) validation.

I am always ready to learn although I do not always like being taught
Winston Churchill

Despite significant advances in the prevention and treatment of cardiovascular disease (CVD), refined risk-stratification and personalized treatment strategies remain a challenge, particularly for ischemic heart disease (IHD)[1]. Current decision-making frameworks rely on a combination of physician expertise, risk scores, and guideline-directed therapies, yet these approaches often fail to accommodate the complexity of individual patient profiles, and remain fraught with subjectivity and imprecision[2]. Traditional predictive models, such as the Systematic Coronary Risk Evaluation (SCORE), based on a patient age, sex, smoking, blood pressure and cholesterol, can estimate risk and guide broad treatment recommendations, but remain static and lack the adaptability required for individualized care[3].

The increasing availability of extensive real-world data from electronic health records (EHRs) and administrative databases, imaging studies, and wearable devices, presents a unique opportunity to enhance decision-making by deploying and harnessing machine learning (ML) and artificial intelligence (AI)[4]. Many promising algorithms are available, but reinforcement learning (RL) offers a uniquely powerful, dynamic, and data-driven approach to treatment optimization by continuously refining strategies based on evolving clinical conditions (Table 1; Fig. 1)[5]. However, while ML and AI models have demonstrated impressive predictive capabilities, their successful implementation in routine cardiovascular care has remained limited with challenges in being able to translate into improved outcomes in real-world practice[6]. Additionally, the complexity of integrating AI into everyday clinical practice requires careful consideration of ethical, logistical, and regulatory aspects, all of which will influence its long-term viability and acceptance.

In their article in *npj Digital Medicine*, Ghasemi et al. begin to address several of these critical challenges with their offline RL model, demonstrating a practical approach to translating AI advances into meaningful clinical applications[7]. Here, we offer commentary on this seminal work—not with the intent to validate the clinical efficacy of offline reinforcement learning approaches, but rather to contextualize and interpret this important contribution within the broader landscape of medical AI.

Applying reinforcement learning to cardiovascular care

Ghasemi et al.'s model was based on a conservative Q-learning approach, a variant of offline RL, within a tabular setting based on observational data from a large registry of patients with or at risk of coronary artery disease (CAD), aiming to refine decision-making regarding coronary revascularization (i.e., percutaneous coronary intervention [PCI] or coronary artery bypass grafting [CABG])[7]. Their study highlights how RL can learn from past patient trajectories, identifying optimal treatment strategies that balance the risks and benefits of PCI versus CABG, as well as conservative

Table 1 | Reinforcement learning and similar artificial intelligence (AI) approaches for improving cardiovascular care

Approach	Algorithmic nature	Data usage	Key advantage	Key challenge	Example in cardiovascular care
Deep learning	Neural network-based, can be supervised or unsupervised	Requires large, high-quality datasets for training	Highly flexible; excels at complex tasks like imaging analysis	"Black box" effect, large computational demands, risk of overfitting	Automated interpretation of echocardiograms or coronary CT angiography
Hybrid or ensemble methods	Combination of multiple algorithms (e.g., supervised + RL)	Integrates diverse data sources; can switch algorithms dynamically	Balances strengths of different models, potentially increasing accuracy & robustness	Complexity, potential for overfitting, and higher computational cost	Multi-stage CAD risk assessment where an RL agent refines a supervised model treatment recommendations
Reinforcement Learning	Iterative decision-making process	Learns optimal actions from sequential patient data (e.g., CAD registry)	Actively improves treatment strategies by continuously updating policies based on outcomes	Requires careful offline validation and interpretability; may differ from standard physician patterns	Offline RL models for revascularization decisions, balancing PCI vs. CABG (Ghasemi et al.)
Supervised Learning	Predictive modeling	Relies on labeled datasets (e.g., mortality or MACE outcomes)	Excellent for classification (e.g., risk stratification), quick to train and deploy	Limited adaptation to changing data; focuses on static predictions rather than ongoing decision-making	Predicting likelihood of major adverse cardiac events post revascularization
Unsupervised learning	Clustering or dimensionality reduction	Uses unlabeled patient data (e.g., imaging, laboratory results)	Identifies hidden patterns or phenotypes without prior assumptions	Results can be harder to interpret; no direct linkage to specific outcomes	Segmenting heart failure subtypes or clustering CAD phenotypes

CABG Coronary Artery Bypass Grafting, *CAD* Coronary Artery Disease, *CT* computed tomography, *MACE* Major Adverse Cardiovascular Events, *PCI* Percutaneous Coronary Intervention

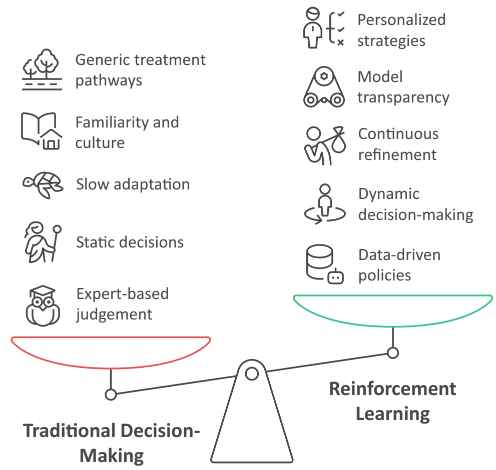

Fig. 1 | Comparison between traditional clinical decision-making and reinforcement learning-based clinical decision support systems. This diagram uses a balance scale to conceptually contrast traditional clinical decision-making (left) with reinforcement learning-enabled decision support (right). Traditional decision-making is rooted in physician expertise, clinical culture, and guidelines, typically characterized by static decisions, slower adaptation to new evidence, and broadly applicable treatment pathways. In contrast, reinforcement learning-based approaches are driven by data-derived policies and enable adaptive, sequential decision-making that continuously refines treatment strategies based on historical outcomes and patient-specific characteristics. Key attributes of each approach—such as familiarity and expert judgment on one side, versus transparency, personalization, and dynamic refinement on the other—highlight the theoretical and operational shift represented by reinforcement in cardiovascular care.

management. Unlike traditional supervised learning (SL) approaches, which rely on explicitly labeled datasets, RL actively develops its own decision policies by interacting with historical data, making it particularly suited for complex, sequential decision-making in medical applications stemming from real-world data.

Importantly, Ghasemi et al.'s findings demonstrate that, in retrospective simulations, offline RL-derived treatment strategies achieved better expected cardiovascular outcomes compared to average physician-assigned decisions within the dataset. While this suggests that RL-based recommendations may surpass conventional strategies in modeled scenarios, it does not directly and conclusively imply their superiority over clinician judgment in real-world practice. Notably, the ability of RL to simulate numerous treatment pathways and identify optimal strategies for individual patients highlights its potentially expanding role as a decision-support tool. However, translating these retrospective gains into tangible clinical improvements will require rigorous prospective validation, arguably in a dedicated randomized trial. Indeed, RL models must be tested in real-world clinical practice to determine whether their recommendations can lead to better patient outcomes, fewer complications, and longer overall and event-free survival, while maintaining or improving efficiency and workflow integration[8]. Furthermore, AI applications such as RL must prove adaptable across diverse healthcare systems and patient populations to ensure its generalizability and broad clinical impact. This requirement is crucial given that clinical decision-making often extends beyond a binary choice of whether to revascularize a patient, encompassing decisions about laboratory testing, imaging studies, and pharmacologic treatment regimens[9].

Fig. 2 | Implementation of artificial intelligence (AI) in cardiovascular care. This diagram outlines the essential domains for the safe and effective integration of AI into cardiovascular care, following a total product lifecycle perspective. Technical foundations include interoperability planning and early attention to privacy, security, and regulatory alignment. Validation efforts must assess clinical effectiveness and economic value. Successful clinical integration depends on compatibility with existing workflows, including compatibility with available electronic health records (EHR) and clinician trust. Finally, ongoing oversight ensures regulatory compliance, transparency, and bias mitigation. Addressing these domains holistically is critical to realizing AI potential in real-world cardiovascular practice.

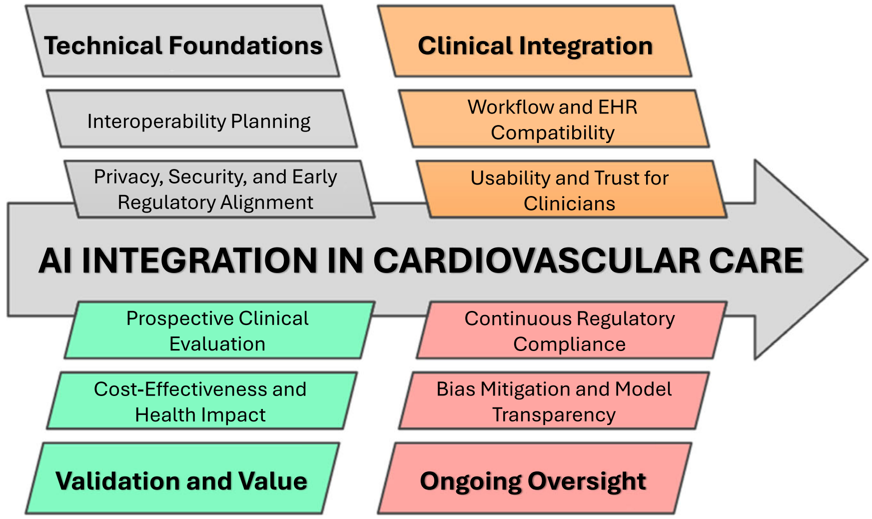

Making artificial intelligence work for clinicians

The success of AI in cardiology will depend not only on predictive accuracy, but also on interpretability and explainability and ease of integration into busy clinical workflows[10]. Physicians are unlikely to trust or adopt AI-generated recommendations without understanding their rationale and eventually trusting them as they would an experienced and competent colleague. The "black-box" nature of many machine learning models—particularly proprietary systems, which may be restricted by intellectual property protections, commercial use or limited documentation—remains a major barrier to clinical implementation, as clinicians require transparency in decision-making to align AI recommendations with their own expertise (Fig. 2).

Ghasemi et al. addressed this concern by implementing conservative Q-learning (CQL), which ensures that AI-driven recommendations align with established clinical patterns while still optimizing patient care[7]. By constraining the recommendations generated by the RL model in order to keep them within the bounds of observed physician decisions, CQL reduces the likelihood of overly aggressive, unrealistic, or unfamiliar treatment strategies. This hybrid approach helps balance innovation with physician trust, and, while Ghasemi et al. do not formally demonstrate nor claim the following, we may reasonably infer that their methods can be leveraged to ensure that AI enhances rather than replaces clinical expertise in specific clinical scenarios.

Looking toward the future, subsequent iterations of RL models should incorporate intuitive visualization tools and transparent decision pathways, enabling physicians to quickly and critically assess AI-driven suggestions rather than accepting them uncritically. Such features, including the ability to adjust key clinical variables, may help mitigate automation bias by reinforcing the physician role in judgment and oversight while leveraging the model strengths[11]. Additionally, embedding explainable AI techniques—such as feature attribution or saliency maps—may support clinicians in understanding how a specific model arrived at specific recommendations. Complementary tools like counterfactual explanations, while not strictly part of explainability, may further reduce automation bias by enabling clinicians to explore "what if" scenarios and critically assess alternative decisions[12]. Increased interdisciplinary collaboration among data scientists, physicians, and healthcare administrators will also be essential to refining these systems and addressing potential biases that could impact clinical recommendations.

Overcoming barriers to clinical implementation

In order to ensure that RL-based models or similarly complex AI tools are actually implemented and eventually impact routine cardiovascular care, they must be embedded within EHR systems to provide real-time, context-aware recommendations. This would allow clinicians to access AI-driven insights at the point of care, integrating predictive analytics into existing workflows[13]. However, this integration faces substantial logistical challenges, including healthcare system interoperability and compatibility between AI tools and diverse data structures to ensure seamless adoption across departments, institutions and regions. Data privacy and security concerns are also paramount, as AI models must comply with stringent regulations to protect patient confidentiality, particularly when training on detailed and sensitive medical records, which amount to a clinical fingerprint[14]. Similarly, clinician usability and workflow integration must be ensured, with AI recommendations presented in an intuitive, actionable, and time-efficient format that simplifies rather than complicates already demanding clinical environments, all while maintaining reasonable purchase and maintenance costs[15,16]. Additionally, RL models (like all sophisticated AI tools) will require regular retraining with updated patient data to remain relevant in evolving clinical settings. A model trained on past data may become outdated as treatment guidelines, procedural techniques, and population characteristics change over time, as exemplified by the significant shifts in care patterns witnessed during recent pandemic conditions[17]. Addressing these challenges will likely impact the eventual uptake of RL as a transformative tool in cardiovascular care.

Validation, cost-effectiveness, and the future of artificial intelligence in cardiology

Beyond technical implementation, AI-driven tools must demonstrate real-world efficacy and economic viability[6]. While Ghasemi et al. provide compelling retrospective evidence, prospective trials are sorely needed to evaluate whether RL-based decision-making reduces mortality and major adverse cardiovascular events, thereby mitigating the risk of drawing erroneous conclusions from potentially confounded observational associations. Randomized controlled trials (RCTs), such as those with a cluster and factorial design, can be structured to compare AI-augmented clinical decisions to standard physician judgment and will be crucial in validating the utility of RL and similar AI approaches in clinical practice[18].

Beyond clinical efficacy, AI adoption will also depend on demonstrating across-the-board cost-effectiveness[19]. Healthcare systems currently face increasing financial pressures, and new technologies must prove that they enhance care without introducing unsustainable costs. Novel RL-based interventions must show that they improve patient outcomes in a way that justifies implementation costs, reduce unnecessary procedures, hospitalizations, and resource utilization, and optimize healthcare efficiency by streamlining decision-making and reducing physician workload[20,21]. Indeed, AI should not only enhance clinical decision-making but also contribute to the overall sustainability and accessibility of cardiovascular care[22]. These considerations will be crucial in determining the extent to which AI-based decision-making will be embraced by healthcare institutions and regulatory bodies.

Future research should prioritize pragmatic clinical trials, human-in-the-loop testing, and cost-utility assessments to assess how AI models, including offline RL, perform under actual decision-making constraints, acknowledging implementation issues due to patients' preferences and healthcare delivery subtleties (e.g., payment models).

Limitations

While Ghasemi et al.'s offline RL model provides promising retrospective insights, several drawbacks should temper overinterpretation. First, their analysis was a simulation-based exercise, and did not involve real-time physician interaction, prospective testing, or randomized allocation. As such, the observed superiority of RL-derived strategies remains largely hypothetical until externally validated. Second, the use of observational registry data introduces the risk of unmeasured confounding and selection bias, limiting internal as well as external validity. Third, RL models may underperform when applied to populations or settings not reflected in the training dataset, raising concerns about generalizability and, again, external validity. Finally, technical, regulatory, cost and ethical barriers—including data silos, privacy requirements, and model explainability—must be addressed before these tools can be safely and equitably integrated into routine cardiovascular care.

Conclusion

The integration of RL and similar AI tools in cardiovascular medicine represents a major opportunity to refine patient care, offering data-driven insights that could personalize treatment strategies more effectively than ever before. However, to transition from promising retrospective analyses to real-world clinical impact, AI models must undergo extensive validation, address workflow integration challenges, and prove their cost-effectiveness. Ultimately, AI will succeed in advancing care in cardiology and medicine at large only if it enhances clinician decision-making, improves the outcomes of patients while protecting all their rights, and remains accessible and ethical in its deployment.

Data availability

No datasets were generated or analysed during the current study.

Giuseppe Biondi-Zoccai[1,2] ✉, Arjun Mahajan[3], Dylan Powell[4], Mariangela Peruzzi[2,5], Roberto Carnevale[1,6] & Giacomo Frati[1,6]

[1]Department of Medical-Surgical Sciences and Biotechnologies, Sapienza University of Rome, Latina, Italy. [2]Maria Cecilia Hospital, GVM Care & Research, Cotignola, Italy. [3]Harvard Medical School, Boston, MA, USA. [4]Faculty of Health Sciences & Sport, University of Stirling, Stirling, UK. [5]Department of Clinical Internal, Anesthesiological and Cardiovascular Sciences, Sapienza University of Rome, Rome, Italy. [6]IRCCS NEUROMED, Pozzilli, Italy. ✉e-mail: giuseppe.biondizoccai@uniroma1.it

Received: 13 March 2025; Accepted: 8 April 2025;
Published online: 05 May 2025

References

1. Krittanawong, C. et al. Strategies for chronic coronary disease: Abrief guide for clinicians. *npj Cardiovasc. Health* **1**, 6 (2024).
2. Lourida, K. G. & Louridas, G. E. Constraints in clinical cardiology and personalized medicine: interrelated concepts in clinical cardiology. *Cardiogenetics* **11**, 50–67 (2021).
3. Conroy, R. M. et al. Estimation of ten-year risk of fatal cardiovascular disease in Europe: the SCORE project. *Eur. Heart J.* **24**, 987–1003 (2003).
4. Ye, J., Woods, D., Jordan, N. & Starren, J. The role of artificial intelligence for the application of integrating electronic health records and patient-generated data in clinical decision support. *AMIA Jt Summits Transl. Sci. Proc.* **31**, 459–467 (2024).
5. Khezeli, K. et al. Reinforcement learning for clinical applications. *Clin. J. Am. Soc. Nephrol.* **18**, 521–523 (2023).
6. Armoundas, A. A. et al. Use of artificial intelligence in improving outcomes in heart disease: a scientific statement from the American Heart Association. *Circulation* **149**, e1028–e1050 (2024).
7. Ghasemi, P. et al. Personalized decision making for coronary artery disease treatment using offline reinforcement learning. *npj Digit. Med.* **8**, 99 (2025).
8. Jayaraman, P., Desman, J., Sabounchi, M., Nadkami, G. N. & Sakhuja, A. A primer on reinforcement learning in medicine for clinicians. *npj Digit. Med.* **7**, 337 (2024).
9. Alowais, S. A. et al. Revolutionizing healthcare: the role of artificial intelligence in clinical practice. *BMC Med. Educ.* **23**, 689 (2023).
10. Karalis, V. D. The integration of artificial intelligence into clinical practice. *Appl. Biosci.* **3**, 14–44 (2024).
11. Žlahtič, B. et al. Trusting AI made decisions in healthcare by making them explainable. *Sci. Prog.* **107**, 368504241266573 (2024).
12. Amann, J. et al. Z-Inspection initiative. To explain or not to explain?-Artificial intelligence explainability in clinical decision support systems. *PLOS Digit. Health* **1**, e0000016 (2022).
13. Kawamoto, K., Finkelstein, J. & Del Fiol, G. Implementing machine learning in the electronic health record: checklist of essential considerations. *Mayo Clin. Proc.* **98**, 366–369 (2023).
14. Murdoch, B. Privacy and artificial intelligence: challenges for protecting health information in a new era. *BMC Med. Ethics* **22**, 122 (2021).
15. Elhaddad, M. & Hamam, S. AI-driven clinical decision support systems: an ongoing pursuit of potential. *Cureus* **16**, e57728 (2024).
16. Labkoff, S. et al. Toward a responsible future: recommendations for AI-enabled clinical decision support. *J. Am. Med. Inform. Assoc.* **31**, 2730–2739 (2024).
17. Krishnan, G. et al. Artificial intelligence in clinical medicine: catalyzing a sustainable global healthcare paradigm. *Front. Artif. Intell.* **6**, 1227091 (2023).
18. Plana, D. et al. Randomized clinical trials of machine learning interventions in health care: a systematic review. *JAMA Netw. Open* **5**, e2233946 (2022).
19. Kastrup, N., Holst-Kristensen, A. W. & Valentin, J. B. Landscape and challenges in economic evaluations of artificial intelligence in healthcare: a systematic review of methodology. *BMC Digit. Health* **2**, 39 (2024).
20. Khanna, N. N. et al. Economics of artificial intelligence in healthcare: diagnosis vs. treatment. *Healthcare* **10**, 2493 (2022).
21. Mahajan, A. & Powell, D. Generalist medical AI reimbursement challenges and opportunities. *NPJ Digit. Med.* **8**, 125 (2025).
22. Ramezani, M. et al. The application of artificial intelligence in health policy: a scoping review. *BMC Health Serv. Res.* **23**, 1416 (2023).

Acknowledgements

This manuscript was drafted and illustrated with the assistance of artificial intelligence tools, including ChatGPT 4 (OpenAI, San Francisco, CA, USA) and Napkin AI (Napkin AI, Palo Alto, CA, USA), in keeping with established best practices (Biondi-Zoccai G, editor. ChatGPT for Medical Research. Torino: Edizioni Minerva Medica; 2024). The final content, including all conclusions and opinions, has been thoroughly revised, edited, and approved by the authors. The authors take full responsibility for the integrity and accuracy of the work and retain full credit for all intellectual contributions. Compliance with ethical standards and guidelines for the use of artificial intelligence in research has been ensured.

Author contributions

G.B.Z., M.P., R.C., and G.F. developed the manuscript concept. G.B.Z., M.P., R.C., G.F., and A.M. wrote the main manuscript. D.P. provided oversight in drafting and editing of the manuscript. All authors read and approved the final manuscript.

Competing interests

Giuseppe Biondi-Zoccai has consulted, lectured and/or served as advisory board member for Abiomed, Advanced Nanotherapies, Aleph, Amarin, AstraZeneca, Balmed, Cardionovum, Cepton, Crannmedical, Endocore Lab, Eukon, Guidotti, Innovheart, Meditrial, Menarini, Microport, Opsens Medical, Synthesa, Terumo, and Translumina, outside the present work. The other authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Giuseppe Biondi-Zoccai.

Reprints and permissions information is available at http://www.nature.com/reprints