Position: Specialty Society–Led Meta-Governance Is Essential to Responsible Implementation of Generative AI in Cardiovascular Care

Yumin Gao[1], Jacqueline Sandling[1], Francis Arellano[1], Zahra Ahmad[1], Deeya Garg[1], Emily Chen[1] Aamir Javaid[2], Seth Martin[3], Francoise Marvel[3], Raj Ratwani[1, 4], Charles German[5]

[1]Georgetown University School of Medicine
[2]University of California San Francisco, Division of Cardiology
[3]Johns Hopkins Ciccarone Center for the Prevention of Cardiovascular Disease
[4]MedStar National Center for Human Factors Engineering in Healthcare
[5]MedStar Heart and Vascular Institute

yg351@georgetown.edu

Abstract

Generative artificial intelligence (AI), particularly large language models (LLMs), is rapidly emerging in cardiology, the leading global cause of death and a major determinant of population health. While the development of AI applications in cardiology has expanded steadily and now ranks second only to radiology in FDA-approved AI-enabled devices, the use of generative AI models in real-world clinical settings remains very limited. Current efforts focus primarily on exploratory studies related to documentation and patient education rather than direct clinical decision support and disease management. Four barriers define the current landscape: 1) insufficient frameworks for external and local validation, 2) sensitivity to contextual and user factors, 3) lack of structured post-deployment governance, and 4) misaligned reimbursement and regulatory incentives. Addressing these challenges will require coordination between regulators and specialty societies such as the American Heart Association (AHA) and American College of Cardiology (ACC). Society-led meta-governance anchored in validation standards, monitoring infrastructure, and reimbursement guidance is essential to ensure the safe, equitable, and effective implementation of generative AI in cardiovascular care.

1 Introduction

Generative artificial intelligence (AI) refers to machine learning systems trained on vast datasets that can produce novel outputs such as text, images, audio, or synthetic waveforms. Large language models (LLMs), the most prominent form of generative AI, are being explored in cardiology to streamline administrative tasks, generate patient education materials, and synthesize electrocardiographic and imaging data for training and analysis.

Cardiovascular applications demand exceptional precision because they rely on multimodal data such as ECGs, echocardiograms, and hemodynamic biometrics, and involve rapid decision cycles where model hallucination may directly affect outcomes like revascularization timing or arrhythmia detection. Cardiology is also among the most technologically active specialties, ranking second only to radiology in FDA-cleared AI-enabled devices [1]. Given that cardiovascular disease remains

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: The Second Workshop on GenAI for Health: Potential, Trust, and Policy Compliance.

the leading global cause of death [2], advances in AI for this field have broad population health implications. The scalability of generative AI can further accelerate patient education, triage, and decision support across millions of encounters, making cardiology a critical proving ground for safe and equitable AI governance.

Despite this promise, adoption in cardiovascular care remains limited. Current applications center on workflow efficiency and patient education rather than clinical validation or direct patient care [3]. As of August 2025, the Food and Drug Administration (FDA) has cleared six cardiovascular AI-enabled devices, all based on traditional machine learning rather than generative AI (Figure 1) [1]. Early studies have shown promise. For example, ChatGPT improved patient understanding of post-catheterization care compared with Google searches [4], and documentation tools reduced note time by 20% [5]. Yet these applications remain focused on low-complexity tasks, with limited integration into cardiovascular care pathways. As a result, they stay in an exploratory phase with limited impact on clinical outcomes, and their implementation in cardiovascular care continues to progress slowly.

LLMs also raise distinct safety and equity concerns, including bias amplification, confabulation, and sociodemographic stereotyping [6]. Unlike prior innovations in drugs, devices, or traditional machine learning, these risks demand new standards and governance. In this paper, we outline four key barriers, review current FDA guidance, and propose a framework for the responsible implementation of LLMs in cardiovascular care (Figure 2).

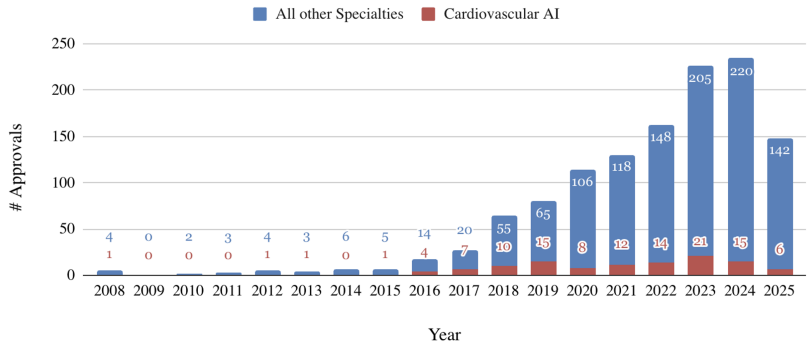

Figure 1: FDA approvals of AI-enabled medical devices by specialty, 2008–2025.

2 Barrier 1: Validation of LLMs

External validation, the traditional benchmark for cardiovascular risk scores such as the atherosclerotic cardiovascular disease Pooled Cohort equations, assumes stable inputs and outputs with reliable performance transfer across populations [7,8]. Generative models violate this assumption because their behavior depends on local documentation patterns, device vendors, and clinician prompting behaviors. In cardiology, this difference is clinically meaningful: an LLM summarizing echocardiography reports or ECG tracings may exhibit significant performance variation across health systems or patient demographics. In a multi-site study of echocardiography reports, a fine-tuned LLM showed a 5–10% performance drop when used at another site and a 25% drop when the note structure was altered [9]. Unlike static models, LLM behavior is shaped by prompt phrasing, user role, and context, so validation at one site does not ensure safety or utility elsewhere.

The FDA has issued draft guidance on lifecycle management and marketing submission of Artificial Intelligence–Enabled Device Software Functions (AI-DSFs), defined as software components that incorporate AI [2]. While the FDA does not directly validate such devices, the guidance requires both performance and human factors validation before marketing. Performance validation ensures a device meets intended use, while human factors validation assesses safe, effective user interaction. These broad recommendations focus on machine learning rather than the unique challenges of LLMs, and because compliance is not tracked, it remains unclear how consistently these recommendations are applied or how validation lapses influence model reliability in practice.

A stronger approach is recurring local validation, a practice for continuously monitoring performance in a specific context and retraining over time [10]. Unlike external validation, recurring local validation dynamically evaluates cost-effectiveness, workflow efficiency, and fairness. Model architectures, parameters, and weights can be adapted to changes in patient populations, clinical practice, and provider use [10]. One concern is the clinicians and engineering effort required to tailor models for each deployment [11]. Because the FDA lacks resources to oversee local validation across diverse health systems, many organizations have established internal AI oversight councils with custom strategies for safety, performance, and risk management.

This distinction aligns with the evolving dialogue between external validation (testing transportability across populations) and local validation (recurring, context-specific monitoring). For generative AI, both are necessary but serve different goals: external validation ensures generalizability across cardiovascular populations, while local validation safeguards day-to-day reliability within a specific cath lab, imaging core, or ambulatory practice.

Professional cardiology societies are uniquely positioned to set broader, evidence-based standards. For example, the American Heart Association (AHA) outlined best practices for real-world monitoring of AI/ML algorithms in its 2024 Circulation statement [12]. Similarly, the American College of Cardiology (ACC) recently launched its Artificial Intelligence Resource Center, which curates educational materials and resources to help members understand and apply AI in the digital transformation of cardiovascular care delivery [13]. Continuous validation is essential as patient populations, workflows, and models evolve, and professional societies can help guide its implementation by promoting recurring local validation frameworks and aligning evidence thresholds for use across diverse practice settings.

3 Barrier 2: Sensitivity to Contextual Factors

Another barrier to safe and consistent LLM use in cardiology is the lack of standardized definitions for task type, data format, and user expertise. Without this clarity, performance metrics may be misleading: GPT-4 scored nearly 75% on board-style cardiovascular medicine questions [14] but dropped to 42% for bradycardia management in ACLS simulations [15]. These tasks require very different reasoning and inputs. Domain-tuned models such as Med-PaLM 2, which incorporate medical training data and retrieval augmentation, improved long-form clinical output [16], underscoring the importance of task framing and data coverage. User factors further complicate reliability: because prolonged routine AI reliance may erode baseline clinical skills [17] and non-specialists are more prone to automation bias [18], there is concern that trainees or generalists using LLMs to interpret ECGs, echocardiograms, or risk scores may lack the depth needed to catch subtle errors.

For marketing submissions, the FDA requests examples of output format, a video demonstration, and an overview of device operation [19]. However, current guidance does not require LLM-based devices to define or validate task type, data format, or intended user expertise, which are three factors central to performance variability in cardiology. Section VI ("User Interface and Labeling") in their Marketing Submission Recommendations outlines how interfaces should present instructions and alerts but does not address how models should behave when non-identical inputs are given, or how to disclose optimal use conditions. Since reliability can shift when tasks change (multiple choice vs an ACLS scenario), inputs differ (free-text echo report vs structured dataset), or users vary (trainee vs specialist), FDA guidance could be strengthened by requiring manufacturers to specify, test, and report these parameters in submissions.

Professional cardiology societies could complement FDA action by defining and standardizing these variables for clinical AI tools. They already track research on task performance, data quality, and human factors, and convert evidence into guidelines. For instance, the American College of Cardiology's AI Resource Center already provides educational materials on clinical trial implementation and disease screening [13]. By extending this framework generative AI, societies could create standards for task definitions (e.g., diagnostic classification vs decision support), data input formats (e.g., structured fields vs narrative notes), and user profiles (e.g., cardiology fellow vs general internist). Embedding such standards into training, accreditation, and quality measures would reduce ambiguity, improve reproducibility, and support safe deployment across diverse practice settings.

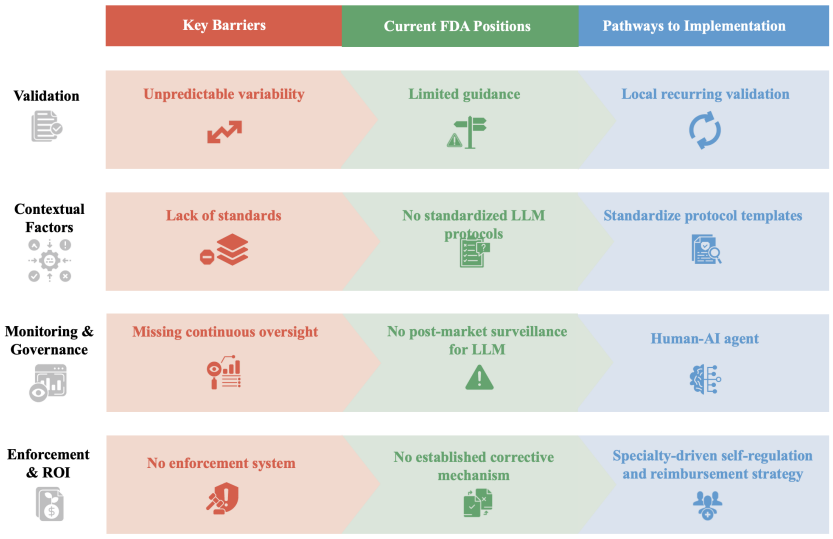

Figure 2: Key Barriers, current FDA positions, and pathways to implementation for large language models in cardiology.

4 Barrier 3: Monitoring & Governance

LLMs are not static devices. Their behavior shifts with model updates, evolving patient populations, and different clinical workflows. This creates a governance barrier: once deployed, there is no established mechanism to ensure outputs remain accurate, safe, and equitable over time. In cardiology, where guidelines evolve rapidly and clinical stakes are high, the absence of continuous monitoring risks silent degradation and bias. The FDA minimizes this risk for traditional drugs and devices using post-market surveillance through registries or Risk Evaluation and Mitigation Strategies (REMS). For example, Mavacamten, the first FDA-approved myosin inhibitor for obstructive hypertrophic cardiomyopathy, carries a REMS program requiring prescriber certification and monitoring via echocardiography. The FDA mandated echocardiograms every 12 weeks until April 2025, when the FDA approved twice-yearly monitoring in stable patients with preserved ejection fraction [20].

However, no comparable FDA governance infrastructure exists for LLMs, despite their potential to influence equally high-stakes treatment decisions. Current FDA activity in governing clinical LLM use is minimal aside from published calls for stakeholder input, leaving operational oversight largely unaddressed. Moreover, while human oversight remains critical for ensuring accuracy, ethical compliance, and patient safety, exclusive reliance on human review at the FDA is operationally unsustainable given the scale and complexity of modern AI applications [21].

One possible solution is to introduce a human–AI monitoring agent that is distinct from the AI technology used for direct patient care [22, 23]. In this model, clinicians contribute contextual judgment and clinical expertise [5], while AI systems enable real-time surveillance, pattern detection, and large-scale data analysis [24]. This combined oversight approach can increase monitoring efficiency but also raises new governance requirements to ensure that the supervisory process itself remains valid, unbiased, and aligned with clinical objectives. Without structured governance at this higher level, human–AI monitoring systems risk introducing new vulnerabilities, including automation complacency, error propagation, and selective deskilling [22].

These "meta-governance" requirements call for an authoritative body with clinical expertise, established ties to technical and industry stakeholders, and the ability to balance operational details with system-level objectives. Professional cardiology societies are well positioned to convene this role. They routinely set evidence-based standards, develop training resources, and maintain networks spanning regulators, researchers, and AI experts [25, 26]. Building on their existing work in AI guidelines, societies can also provide governance frameworks in key domains such as education and competency (credentialing, CME), evidence thresholds for use, oversight and monitoring (registries, safety task forces), and regulatory alignment. Importantly, this requires engagement of a multidis-

ciplinary coalition that includes clinicians, allied health professionals, data scientists, engineers, ethicists, regulators, and patient representatives to ensure that oversight remains clinically meaningful, ethically sound, and sustainable.

5 Barrier 4: Enforcement & ROI

Currently, no enforcement system ensures prevention of unsafe or ineffective LLM use in U.S. healthcare. The FDA's authority, rooted in the FD&C Act and the Safe Medical Devices Act of 1990, emphasizes postmarket surveillance and recalls [27]. Although the FDA can mandate recalls, most are company-led, such as Boston Scientific's 2020 pacemaker recall for accidental safety mode activation [28-31]. For AI-enabled devices, no rules define which violations would trigger penalties, or what penalties would entail.

The European Union's AI Act offers a comparison [32]. Like the FDA, it requires premarket and lifecycle assessments, but it goes further by imposing substantial fines for noncompliance and banning "unacceptable risk" applications such as biometric identification, social scoring, and behavioral manipulation [33]. While no fines have been issued yet for AI-enabled medical devices, the Act suggests that clear rules backed by enforcement can drive compliance.

The Safe Medical Devices Act set precedent for FDA enforcement of recalls and civil penalties, suggesting the agency similarly could define expectations for AI-enabled LLMs. However, no action has been taken. A more practical and timelier U.S. path may be market-driven self-regulation through professional societies. For example, the Coalition for Health AI is working with The Joint Commission on AI best practices [34]. While lacking certification carries no formal penalties, it risks insurer scrutiny and patient distrust, creating incentives to comply. Expanding specialty-driven certification programs could help ensure safety and effectiveness.

Another barrier to cardiology generative AI adoption is payment [35]. Multiple reimbursement models exist, broadly utilization based (fee for service) or performance based (value based), but the models are difficult to compare, lack standardized terminology, and present different tradeoffs [36]. This complexity risks diverting resources away from implementation and toward administrative waste. Interestingly, despite recent additions to the Current Procedural Terminology (CPT) code set introduced by the American Medical Association (AMA) to describe AI-enabled services, major challenges persist in implementation and integration [26,37]. The payment infrastructure remains immature, as most AI codes are Category III and therefore not reimbursed, while regulatory alignment among the American Medical Association, FDA, and the Centers for Medicare & Medicaid Services (CMS) remains incomplete. Moreover, the definition of physician "work" becomes ambiguous as models progress toward greater autonomy, complicating both liability and valuation.

The absence of clear regulatory guidance highlights the role of professional societies in advancing this discussion. In cardiology, societies could build on their existing work in AI guidelines by providing frameworks for vendor evaluation, return on investment (ROI) guidance [38], standardized contract language, and computable templates to harmonize billing [39]. By convening stakeholders across academia, industry, and payer organizations, societies can clarify value capture, reduce administrative burden, and support the efficient and equitable integration of generative AI into cardiovascular care.

6 Conclusion

Generative AI and LLMs can transform cardiovascular care but challenge current frameworks for validation, governance, and reimbursement. FDA oversight remains necessary but is not sufficient. Recurrent local validation, context-specific standards, structured monitoring, and sustainable reimbursement models are required to ensure safety and equity. Professional cardiology societies can bridge these gaps by setting technical and clinical standards, convening multidisciplinary task forces, and developing certification and value frameworks. Aligning regulation with society-led meta-governance will create the accountability needed to responsibly implement LLMs, advancing cardiovascular care while safeguarding patient outcomes.

Acknowledgments and Disclosure of Funding

This work was supported by collaborations through H2AI at Georgetown University, a platform advancing innovation, research, and education at the intersection of AI and medicine.

References

[1] U.S. Food and Drug Administration. "Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations." https://www.fda.gov/regulatory-information/search-fda-guidance-documents/artificial-intelligence-enabled-device-software-functions-lifecycle-management-and-marketing. Accessed Aug. 31, 2025.

[2] S. S. Martin et al., "2025 Heart Disease and Stroke Statistics: A Report of US and Global Data From the American Heart Association," Circulation, vol. 151, no. 8, pp. e41-e660, Feb 25 2025, 10.1161/CIR.0000000000001303.

[3] T. M. Maddox, P. Embi, J. Gerhart, J. Goldsack, R. B. Parikh, and T. C. Sarich, "Generative AI in Medicine - Evaluating Progress and Challenges," *N Engl J Med*, vol. 392, no. 24, pp. 2479-2483, Jun 26 2025, https://doi.org/10.1056/NEJMsb2503956.

[4] C. Krittanawong, H. U. H. Virk, S. L. Kaplin, Z. Wang, S. Sharma, and H. Jneid, "Assessing the Potential of ChatGPT for Patient Education in Cardiac Catheterization Care," *JACC Cardiovasc Interv*, vol. 16, no. 12, pp. 1551-1552, Jun 26 2023, https://doi.org/10.1016/j.jcin.2023.04.042.

[5] M. J. Duggan et al., "Clinician Experiences With Ambient Scribe Technology to Assist With Documentation Burden and Efficiency," *JAMA Netw Open*, vol. 8, no. 2, p. e2460637, Feb 3 2025, https://doi.org/10.1001/jamanetworkopen.2024.60637.

[6] G. Quer and E. J. Topol, "The potential for large language models to transform cardiovascular medicine," *Lancet Digit Health*, vol. 6, no. 10, pp. e767-e771, Oct 2024, https://doi.org/10.1016/S2589-7500(24)00151-1.

[7] P. Muntner et al., "Validation of the atherosclerotic cardiovascular disease Pooled Cohort risk equations," *JAMA*, vol. 311, no. 14, pp. 1406-15, Apr 9 2014, https://doi.org/10.1001/jama.2014.2630.

[8] C. D. Rodriguez-Ariza et al., "External validation of the ACC/AHA ASCVD risk score in a Colombian population cohort," *Sci Rep*, vol. 13, no. 1, p. 6139, Apr 15 2023, https://doi.org/10.1038/s41598-023-32668-4.

[9] C. J. Chao et al., "Evaluating large language models in echocardiography reporting: opportunities and challenges," *Eur Heart J Digit Health*, vol. 6, no. 3, pp. 326-339, May 2025, https://doi.org/10.1093/ehjdh/ztae086.

[10] A. Youssef, M. Pencina, A. Thakur, T. Zhu, D. Clifton, and N. H. Shah, "External validation of AI models in health should be replaced with recurring local validation," *Nat Med*, vol. 29, no. 11, pp. 2686-2687, Nov 2023, https://doi.org/10.1038/s41591-023-02540-z.

[11] A. Rodman, L. Zwaan, A. Olson, A. K. Manrai, "When It Comes to Benchmarks, Humans Are the Only Way," *NEJM AI*, vol. 2, no. 4, 2025, https://doi.org/10.1056/AIe2500143.

[12] A. A. Armoundas et al., "Use of Artificial Intelligence in Improving Outcomes in Heart Disease: A Scientific Statement From the American Heart Association," *Circulation*, vol. 149, no. 14, pp. e1028-e1050, Apr 2 2024, https://doi.org/10.1161/CIR.0000000000001201.

[13] American College of Cardiology. "Artificial Intelligence (AI) Resource Center." https://www.acc.org/About-ACC/Innovation/Artificial-Intelligence#sort=%40commonsortdate%20descending. Accessed Aug. 31, 2025.

[14] S. Milutinovic et al., "Evaluating Performance of ChatGPT on MKSAP Cardiology Board Review Questions," *Int J Cardiol*, vol. 417, p. 132576, Dec 15 2024, https://doi.org/10.1016/j.ijcard.2024.132576.

[15] C. Pham, R. Govender, S. Tehami, S. Chavez, O. E. Adepoju, and W. Liaw, "ChatGPT's Performance in Cardiac Arrest and Bradycardia Simulations Using the American Heart Association's Advanced Cardiovascular Life Support Guidelines: Exploratory Study," *J Med Internet Res*, vol. 26, p. e55037, Apr 22 2024, https://doi.org/10.2196/55037.

[16] K. Singhal et al., "Toward expert-level medical question answering with large language models," *Nat Med*, vol. 31, no. 3, pp. 943-950, Mar 2025, https://doi.org/10.1038/s41591-024-03423-7.

[17] K. Budzyn et al., "Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study," *Lancet Gastroenterol Hepatol*, Aug 12 2025, https://doi.org/10.1016/S2468-1253(25)00133-5.

[18] C. Reverberi et al., "Experimental evidence of effective human-AI collaboration in medical decision-making," *Sci Rep*, vol. 12, no. 1, p. 14952, Sep 2 2022, https://doi.org/10.1038/s41598-022-18751-2.

[19] S. Gaube et al., "Do as AI say: susceptibility in deployment of clinical decision-aids," *NPJ Digit Med*, vol. 4, no. 1, p. 31, Feb 19 2021, https://doi.org/10.1038/s41746-021-00385-9.

[20] Bristol Myers Squibb. "U.S. Food and Drug Administration Updates CAMZYOS® (mavacamten) Label to Reduce Echocardiography Monitoring Requirements and Contraindications." https://news.bms.com/news/details/2025/U-S--Food-and-Drug-Administration-Updates-CAMZYOS-mavacamten-Label-to-Reduce-Echocardiography-Monitoring-R default.aspx. Accessed Aug. 31, 2025.

[21] H. J. Warraich, T. Tazbaz, and R. M. Califf, "FDA Perspective on the Regulation of Artificial Intelligence in Health Care and Biomedicine," *JAMA*, vol. 333, no. 3, pp. 241-247, Jan 21 2025, https://doi.org/10.1001/jama.2024.21451.

[22] J. Adler-Milstein, D. A. Redelmeier, and R. M. Wachter, "The Limits of Clinician Vigilance as an AI Safety Bulwark," *JAMA*, vol. 331, no. 14, pp. 1173-1174, Apr 9 2024, https://doi.org/10.1001/jama.2024.3620.

[23] S. Johri et al., "An evaluation framework for clinical use of large language models in patient interaction tasks," *Nat Med*, vol. 31, no. 1, pp. 77-86, Jan 2025, https://doi.org/10.1038/s41591-024-03328-5.

[24] J. Zou and E. J. Topol, "The rise of agentic AI teammates in medicine," *Lancet*, vol. 405, no. 10477, p. 457, Feb 8 2025, https://doi.org/10.1016/S0140-6736(25)00202-8.

[25] American Heart Association. "Our Policy Positions." https://www.heart.org/en/get-involved/advocate/policy-research/our-policy-positions. Accessed Sep. 1, 2025.

[26] D. Dey et al., "Proceedings of the NHLBI Workshop on Artificial Intelligence in Cardiovascular Imaging: Translation to Patient Care," *JACC Cardiovasc Imaging*, vol. 16, no. 9, pp. 1209-1223, Sep 2023, https://doi.org/10.1016/j.jcmg.2023.05.012.

[27] U.S. Government Publishing Office. "Public Law 101-629—Safe Medical Devices Act of 1990." https://www.govinfo.gov/content/pkg/STATUTE-104/pdf/STATUTE-104-Pg4511.pdf. Accessed Sep. 1, 2025.

[28] U.S. Food and Drug Administration. "Class 1 Device Recall ACCOLADE DR." https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfRes/res.cfm?id=211663. Accessed Sep. 1, 2025.

[29] U.S. Food and Drug Administration. "Class 1 Device Recall PROPONENT DR." https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfRes/res.cfm?id=211664. Accessed Sep. 1, 2025.

[30] U.S. Food and Drug Administration. "Class 1 Device Recall ESSENTIO DR." https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfRes/res.cfm?id=211665. Accessed Sep. 1, 2025.

[31] U.S. Food and Drug Administration. "Class 1 Device Recall ALTRUA 2 DR." https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfRes/res.cfm?id=211667. Accessed Sep. 1, 2025.

[32] EU Artificial Intelligence Act. "Implementation Timeline." https://artificialintelligenceact.eu/implementation-timeline/. Accessed Sep. 1, 2025.

[33] EU Artificial Intelligence Act. "Article 5: Prohibited AI Practices." https://artificialintelligenceact.eu/article/5/. Accessed Sep. 1, 2025.

[34] American Hospital Association. "Joint Commission announces partnership to develop best practices for AI in health care." https://www.aha.org/news/headline/2025-06-12-joint-commission-announces-partnership-develop-best-practices-ai-health-care. Accessed Sep. 1, 2025.

[35] A. Shah, "Navigating the value proposition of artificial intelligence in cardiovascular disease prevention," *Digit Health*, vol. 11, p. 20552076251357409, Jan-Dec 2025, https://doi.org/10.1177/20552076251357409.

[36] Bessemer Venture Partners. "Roadmap: Healthcare AI." https://www.bvp.com/atlas/roadmap-healthcare-ai. Accessed Sep. 1, 2025.

[37] A. Robeznieks. "CPT codes offer the language to report AI-enabled health services." https://www.ama-assn.org/practice-management/cpt/cpt-codes-offer-language-report-ai-enabled-health-services. Accessed Sep. 1, 2025.

[38] C. E. Binkley et al., "An early pipeline framework for assessing vendor AI solutions to support return on investment," *NPJ Digit Med*, vol. 8, no. 1, p. 368, Jun 17 2025, https://doi.org/10.1038/s41746-025-01767-z.

[39] A. Z. Jabri, J. Sandling, K. Schulman, D. Scheinker, "Variation and Standardization in Prior Authorization Requirements," *medRxiv* [Preprint], 2025, https://doi.org/10.1101/2025.02.07.25321895.