Regulatory Science: A Condition of Innovation for AI as a Medical Device

Authors:	*Ellen Coughlan,[1,2] Charlotte Radovanovic[1,3] 1. Centre of Excellence for Regulatory Science and Innovation in AI and Digital Health (CERSI-AI), UK 2. University of Birmingham, UK 3. University Hospitals Birmingham NHS Foundation Trust, UK *Correspondence to e.coughlan@bham.ac.uk

Conflicts of interest:	The authors have declared no conflict of interest.
Gen AI use:	During the final edit of this work, the authors used Claude (Anthropic, San Francisco, California, USA) to check that author edits addressed the reviewer's comments. AI tools were not used to write this work, and after using this tool, the authors reviewed the content as needed and take full responsibility for the content of the final publication.
Peer review:	This article was accepted following double-blind peer review.
Received:	17.08.26
Accepted:	22.10.26
Keywords:	AI, governance, regulation, responsible AI.
Citation:	EMJ Innov. 2026; https://doi.org/10.33590/emjinnov/0269016E

Abstract

AI systems are increasingly being introduced into healthcare, many of them meeting the threshold for regulation as medical devices. However, the regulatory frameworks built to evaluate their safety and performance were designed for stable, tangible products rather than the highly context-dependent, frequently updating software systems whose risks emerge in real-world use. This creates a dilemma for health systems: the pressure to adopt AI at pace sits uneasily alongside the need to validate, monitor, and govern it vigilantly. This piece will describe how regulatory science offers a route through this tension to realise the benefits of innovation safely. Drawing on national infrastructure and case studies from the Centre of Excellence for Regulatory Science and Innovation in AI and Digital Health (CERSI-AI), this article sets out what regulatory science looks like in practice, and why building the governance and capability to deploy AI responsibly may prove harder than building the technology itself.

Key Points

1. The use of AI as a medical device is adaptive and vulnerable to drift, leaving health systems caught between adopting it quickly and governing it rigorously, a live tension in regulatory policy both in the UK and internationally.

2. Regulatory science can address this tension and help achieve a regulatory framework that positions regulators as lifecycle partners and accelerates the best innovations along the regulatory pathway.

3. Ensuring that regulatory science infrastructure in the UK remains responsive and translational depends upon sustained investment and strong links to governing institutions and regulators.

INTRODUCTION

Why Does AI Stretch Our Traditional Approach To Regulating Healthcare Tools?

The existing framework is built for stable, predictable products with well-established pre-market risk clarity and clear manufacturer responsibility, such as syringes or implants. AI systems differ from these other medical technologies because of their adaptability, context-dependence, and emerging risks related to a requirement for continuous monitoring and careful implementation, muddying the waters of accountability.

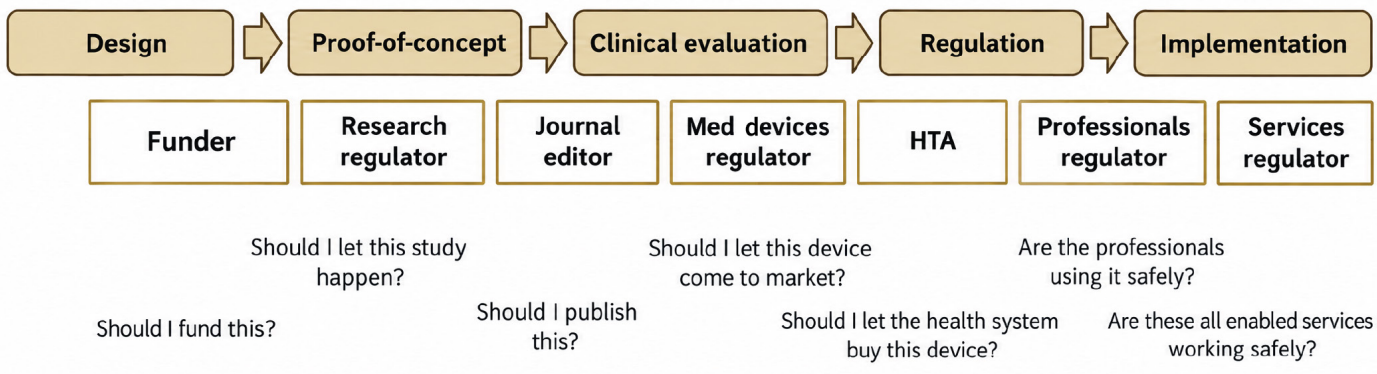
Not all AI systems possess these features, and not all AI systems require regulation; the intended purpose triggers the regulation rather than the technology itself. If the intended purpose qualifies as a medical purpose (as per the regulator's definition), then there are legal duties on the manufacturer, including the need to secure approval from the regulator to bring it to market. This will depend on providing evidence of safety and performance, and a commitment to continue to gather post-market performance data to demonstrate that it is continuing to operate safely and as intended in the real world (post-market surveillance). In the UK, the responsibility for maintaining and enforcing effective regulations for medical devices placed on the UK market sits with the Medicines and Healthcare products Regulatory Agency (MHRA). However, it is worth recognising that the role of the medical device regulator, which is essential to sell into a relevant market (MHRA in the UK; FDA in the USA, etc.), is not the only regulatory or governance gate in that innovation pathway. Each gate has a purpose: to protect the public. It needs to be sufficient to the task, preventing technologies that are unfit (or unready) from progressing. However, it also needs to be efficient, otherwise the cumulative delay along this pathway means that beneficial innovation will perish before it gets to patients. This is just one of many gates on the path from design to implementation for AI innovations, each with its own governing question (Figure 1).

The AI Dilemma and the Regulatory Science Solution

Two things are being asked of the system: adopt AI faster, and ensure that it is safe and effective. In practice, this means to deploy quickly, validate thoroughly, iterate rapidly, and oversee change safely. Historically, the system hasn't been designed to do both at once.

Regulatory science, the development and application of scientific methods, tools, standards, and approaches that support the evaluation and regulation of healthcare technologies, can help create

Figure 1: Progress along the innovation pipeline depends on providing evidence needed to pass each gate.



HTA: Health Technology Assessment.

CC BY-NC 4.0 Licence • Copyright © 2026 EMJ

a system that optimises the innovation pathway and the regulatory gates behind it. The approach allows us to accelerate safe health technologies to patients and improve healthcare while still applying a keen vigilance to safety, quality, and performance standards. It repositions regulators as lifecycle partners rather than sentries at the gates, creating the conditions for speed and safety to be pursued in parallel. The use of regulatory sandboxes to support effective regulation internationally[1] and the early adoption of the Centres of Excellence for Regulatory Science and Innovation (CERSI) model in the USA[2] have realised public health benefits, informed policies, and improved standards,[3] demonstrating that regulatory science can successfully accelerate safe innovation with the right infrastructure.

The Infrastructure Underpinning Regulatory Science

In the UK, there are five main contributors to regulatory science in relation to AI and digital health technologies.

Academic-led initiatives include formal CERSI and academic groups that undertake relevant work in this area. CERSIs, including the Centre of Excellence for Regulatory Science and Innovation in AI and Digital Health (CERSI-AI),[4] supply the connective research infrastructure and focus on solving regulatory and policy challenges at pace. Enabled by diversity of expertise, these sit across the entire pathway and invite colleagues from all sectors to collaborate using a range of scientific methodologies to understand, evaluate and optimise the system.

Regulator-led initiatives, of which the MHRA Airlock is the most prominent example, can generate the evidence base[5] for AI systems by allowing supervised, real-world testing to inform decisions about which gates are appropriate for different AI tools. Other regulators, including the National Institute for Health and Care Excellence (NICE), the professional regulators and service quality inspectors, may also undertake regulatory science approaches, such as systematically reviewing and evaluating approaches to regulation that may inform their future policy.

There are also national reviews. The National Commission on the Regulation of AI in Healthcare, while not a traditional regulatory science initiative, was directly informed by regulatory science approaches including quantitative and qualitative methodologies. The Commission reviewed the entire journey, calling for a modern regulatory framework with proportionate lifecycle regulation and system-wide responsibility underpinned by trust, transparency and predictability.[6] Research and engagement[7] commissioned by the AI Commission surfaced equity and proportionality as values driving its evidence base, signalling that the review asked both what works in the regulatory pathway and what is fair. This work has directly led to policy change in the UK's approach to the regulation of AI in healthcare.

The AI and Digital Regulations Service[8] is a collaboration between NICE, the Health Research Authority (HRA), the Care Quality Commission (CQC), the MHRA, and the NHS, built around clarity. Supported by CERSI-AI, this is a forum in which regulators can work together to ensure that innovators and adopters understand and navigate the pathway efficiently. It provides a single place to understand which regulatory gates apply to them, and in which order.

Beyond the individual regulators, the UK government, through the NHS and the Department of Health and Social Care (DHSC), has an overarching view of the safety of its citizens, including how the regulatory components should fit together. They also draw on regulatory science publications to ensure that they can deliver evidence-based policy.

These initiatives complement each other well but have relied variously on policy interest and short-term investment from government and funders to drive them forward. Other regulatory collaborations, including the Digital Regulation Cooperation Forum in the UK and the European Platform for Regulatory Science Research in the EU, could be seen as instructive models for long-term

investment in regulatory science and regulatory capacity.

How Does Regulatory Science Steady the Journey Along the Regulatory Pathway?

There are plenty of excellent developers with transformational ideas in the UK and elsewhere. To ensure that their ideas can grow from infancy to a safe product, a close examination of regulatory challenges and development of cross-sector solutions is required.

One key challenge for AI health technology developers is access to appropriate data. Optimising secure data environments for AI health technology development is an important step to enable safe and trusted access to NHS data for the best innovators. University Hospitals Birmingham and University of Birmingham worked together with colleagues across academia, regulators, the NHS, patients, and the public to develop a set of recommendations[9] to balance innovation, privacy, and governance. Patients were clear that transparency, proportionality, and a clear articulation of benefits were vital components of this work. Technical colleagues produced an analysis of model egress and disclosure risk, describing where further research is needed. By addressing governance concerns, technical requirements, and policy challenges in concert, we can use regulatory science to facilitate the development of innovations that are effective, trusted, and trained on large, high-quality datasets.

The difference between the way a syringe is procured and used in a community hospital in Bristol compared to a teaching hospital in London is unlikely to be vast; syringes are predictable, their configuration does not change overnight, and they do not tend to work differently on different patient groups or over time. AI systems used in healthcare, on the other hand, may be remotely updated in the absence of clear processes for transparency and traceability;[10] their performance may drift, resulting in waning performance for particular patients. These systems require thoughtful implementation and careful governance in order to manage emerging risks, which may differ between different healthcare settings or geographies. System readiness is a vital regulatory consideration for healthcare providers at the implementation gate of the innovation pathway. Both overcautious and overconfident organisations may cause harm; readiness is not a yes or no, but rather the capacity to accurately assess risk for that product in that setting and control it.

The AI Readiness Framework[11] was developed through funding from the Health Foundation and the NHS, and offers practical concepts and principles to apply in healthcare settings to evaluate readiness for AI tools. It aims to reduce variation across NHS providers, prevent avoidable risks, and enable AI innovations to scale in a way that delivers value for patients and the wider system. Delivered by CERS-AI and co-designed with individuals from a range of NHS and social care settings, it accelerates the benefits of innovations by addressing the 'last mile' problem, where technologies often fail at implementation.[12]

Despite best efforts to design, develop, and implement AI systems that can be used without issue every time, AI as a medical device may still contribute to patient safety events like any other clinical tool. However, unlike medicines, the system lacks both a clear response model to failure and robust regulatory guidance regarding liability. As AI as a medical device becomes more commonplace, responsibility should be shared appropriately across stakeholders, performance information should be transparent, and patients should be actively involved rather than passively informed. The Professional Standards Authority for Health and Social Care (PSA) has responded to this challenge by committing to co-producing a set of high-level principles to guide the regulation of AI use by healthcare professionals.[13] This followed a series of workshops with CERSI-AI to understand the gaps exposed by existing standards and guidance regarding the allocation of responsibility, and the UK Government has committed to building on this work.[14] Furthermore, researchers at the Universities of Glasgow and Birmingham are currently

CC BY-NC 4.0 Licence • Copyright © 2026 EMJ

co-designing a rapid assessment framework with regulators and health system leaders for healthcare organisations to respond to AI failure. Together, these initiatives promote the importance of investing in the regulatory science approach that enables regulators to perform a lifecycle partner role, responding to emerging challenges and allowing the development of new standards and tools to drive innovation.

CONCLUSION

Building AI requires expertise, hard work, and resources. It is difficult. Building the organisations, governance, and capability needed to deploy it responsibly, working with patients and the public as partners, may prove harder still. The National Commission into the Regulation of AI in Healthcare published recommendations that describe a better regulatory framework that reminds us that the work is not done. The pace of technological change means that regulatory science cannot only look backwards, it must also look forwards, anticipating what disruptive technology is coming next, building the frameworks before the tools arrive, and ensuring that the system is never so far behind that innovation moves faster than the ability to keep people safe.

The regulatory challenges posed by large language models have not yet been resolved, and agentic AI is already on the horizon. Ambient voice technologies are reaching the market faster than national guidance can be developed for the organisations expected to deploy them responsibly. This is the nature of a field moving at this speed, but it is also why regulatory science matters. It is one of the most powerful tools we have to ensure that current and future AI tools reach patients safely, equitably, and with the trust of the public intact.

The call to action is clear: regulatory science must be treated not as a response to innovation, but as a condition of it. This must be underpinned by clear links to governing institutions and effective translation of findings to guidance and regulatory practice,[3] lest valuable work remains practically unrealised. With investment in the right infrastructure from the government and sustained investment in cross-sector collaboration by funders and regulators, regulatory science can build the foundations of how AI in healthcare is developed, deployed, and governed.

References

1. Gilbert S. Digital medicine's international race for regulatory sandboxes and voluntary alternative pathways picks up tempo. NPJ Digit Med. 2026;9(1):190.
2. Weichold F. The US Food and Drug Administration centres for regulatory science and innovation: current activities and future promise to accelerate innovations. Clin Pharmacol Ther. 2019;105(4):792-4.
3. Schönfelder A et al. European Union innovation needs regulatory science excellence centres. Nat. Biomed. Eng. 2026;10(9):1802-6.
4. The Centre of Excellence for Regulatory Science and Innovation in AI and Digital Health (CERSI-AI). Available at: https://www.cersi-ai.org/. Last accessed: 4 August 2026.
5. Medicines and Healthcare products Regulatory Agency (MHRA). Guidance: AI Airlock Sandbox Phase 2 Programme Report. 2026. Available at: https://www.gov.uk/government/publications/ai-airlock-sandbox-phase-2-programme-report. Last accessed: 4 August 2026.
6. Athni TS. The United Kingdom's new blueprint for regulation AI in health care. J Med Internet Res. 2026;28:e112992.
7. Medicines and Healthcare products Regulatory Agency (MHRA). Report on the research and engagement programme for the National Commission into the Regulation of AI in Healthcare. 2026. Available at: National_Commission_into_the_Regulation_of_AI_in_Healthcare_-_Research___Engagement_Report.pdf. Last accessed: 4 August 2026.
8. National Healthcare System (NHS). AI and digital regulations service for health and social care. Understanding regulations of AI and digital technology in health and social care. Available at: https://www.digitalregulations.innovation.nhs.uk/. Last accessed: 4 August 2026.
9. The Centres of Excellence for Regulatory Science and Innovation in AI and Digital Health (CERSI-AI). Optimising Secure Data Environments (SDEs) for AI Health Technology Innovation. 2026. Available at: https://www.cersi-ai.org/wp-content/uploads/2026/01/v4-Optimising-Secure-Data-Environments-SDEs-2025-V2.pdf. Last accessed: 4 August 2026.

10. Rotalinti Y, Ordish J et al. Identifying and understanding significant change due to drift when assessing AI models in healthcare: a narrative review. BMJ Digit Health Ai. 2026;2(1):e000085.
11. The Centres of Excellence for Regulatory Science and Innovation in AI and Digital Health (CERSI-AI). The AI Readiness Checklist. 2026. Available at: https://www.cersi-ai.org/the-ai-readiness-checklist/. Last accessed: 4 August 2026.
12. Cabitza F et al. Bridging the "last mile" gap between AI implementation and operation: "data awareness" that matters. Ann Transl Med. 2020;8(7):501.
13. Professional Standards Authority. Joint statement of intent on supporting the safe, effective and ethical use of AI by health and social care professionals. 2026. Available at: https://www.professionalstandards.org.uk/publications/joint-statement-intent-supporting-safe-effective-and-ethical-use-ai-health-and-social. Last accessed: 16 October 2026.
14. Medicines and Healthcare products Regulatory Agency (MHRA). Government response to the National Commission's recommendations on the regulation of AI in healthcare. 2026. Available at: https://www.gov.uk/government/publications/government-response-to-the-national-commissions-recommendations-on-the-regulation-of-ai-in-healthcare/government-response-to-the-national-commissions-recommendations-on-the-regulation-of-ai-in-healthcare. Last accessed: 16 October 2026.

CC BY-NC 4.0 Licence • Copyright © 2026 EMJ