

Introduction to the National Commission into the Regulation of AI in Healthcare

The UK government wants to create a healthcare system that is fit for the future, using artificial intelligence (AI) to help improve care. To support this, new proposals are being considered for regulating medical devices that use AI.

AI has the potential to make healthcare safer, faster, and more personalised. It could help NHS staff work more efficiently and support the growth of health technology companies in the UK. However, AI also brings new challenges for how we regulate and oversee these technologies.

The MHRA, as the regulator for medicines, medical devices, and blood products, is continuing to consider what the regulatory rules should be for AI as a medical device. We have established the National Commission into the Regulation of AI in Healthcare to support this goal.

The [National Commission](#) brings together experts from technology, healthcare, law, patient groups, the public, government, and the NHS. The National Commission will tackle the most important questions about how AI should be regulated and to give recommendations to the MHRA, which will be published in 2026.

Topics considered in the Call for Evidence

The Call for Evidence invites evidence and views on:

- whether the UK's framework for regulating AI in healthcare is sufficient
- how the UK's regulatory framework may need to be improved to ensure fast access to safe and effective AI medical devices
- approaches to checking safety once AI medical devices are in use
- how responsibility and liability are managed between different parties involved in the deployment of AI medical devices.

Question 1: Which of the following best describes your view about the need to change the UK's framework for regulating AI in healthcare? (optional)

No change

Minor adjustments

Significant reform

Complete overhaul

Unsure

Question 2.1: To what extent do you agree or disagree that the current regulatory framework is sufficient in the following domains: Safety and Performance Standards (optional)

Strongly disagree

Disagree

Neither agree or disagree

Agree

Strongly agree

Question 2.2: To what extent do you agree or disagree that the current regulatory framework is sufficient in the following domains: Data Privacy and Data Governance (optional)

Strongly disagree

Disagree

Neither agree or disagree

Agree

Strongly agree

Question 2.3: To what extent do you agree or disagree that the current regulatory framework is sufficient in the following domains: Transparency (optional)

Strongly disagree

Disagree

Neither agree or disagree

Agree

Strongly agree

Question 2.4: To what extent do you agree or disagree that the current regulatory framework is sufficient in the following domains: Requirements for clinical evidence (optional)

Strongly disagree

Disagree

Neither agree or disagree

Agree

Strongly agree

Question 2.5: To what extent do you agree or disagree that the current regulatory framework is sufficient in the following domains: Post Market Surveillance (optional)

Strongly disagree

Disagree

Neither agree or disagree

Agree

Strongly agree

Question 3: How would you rate the current framework's impact on innovation? (optional)

Too restrictive

Somewhat restrictive

About right

Somewhat loose

Too loose

Question 4: How might the UK's framework for regulation of AI in healthcare be improved to ensure the NHS has fast access to safe and effective AI health technology? (optional)

You may wish to consider some or all of the following in your response: (1) Gaps and other limitations of the existing regulatory framework, (2) Innovative and effective approaches to AI regulation used in other sectors, and other jurisdictions, (3) Ensuring public and patient safety whilst minimising the cost of complying with regulations (in terms of time and resource), (4) The boundaries of regulation, including the ways AI can qualify as a 'medical device', and how such devices are classified according to risk.

The Royal Pharmaceutical Society strongly supports a regulatory framework for AI in healthcare that is safe, efficient and trusted. It is important that the foundation of risk-based classification principles of regulation of medical devices can be built upon to accommodate the needs of new digital technologies. A focussed consultative process to analyse the impact and risk assessment criteria which takes into account new and emerging functionality is urgently required and must be repeated at intervals commensurate with the pace of technology advancement.

Pre-market assessment and scrutiny of information governance, security and privacy must be incorporated into the medical device regulation analysis and not be considered separately – this is a vital component of safe deployment of AI in healthcare as a medical device.

Studies indicate that, with many AI systems - in particular, predictive healthcare modelling software and clinical diagnostics systems – prior validation of the system is vital. There is also a need for ongoing monitoring for safety (due to potential unintended outcomes) and to gather evidence of clinical effectiveness. Consequently, the system for regulation of AI in healthcare will need to more closely approximate the system for medicines licensing, with rigorous initial assessment, licensing and post-marketing surveillance. As well as clinical outcomes and safety, the system would need to assess equity in system use (ensuring the system treats both users and data subjects with equity). Unlike the medicine licensing system, an AI device licensing system would need to focus on “real-world” design and implementation, which are crucial factors in AI device use.

This will be a significant and onerous development for software suppliers, but the adoption process can be made easier by the development of clear approval processes, and appropriate guidance on the system for software suppliers. The collaborative multi-agency approach to guidance for developers and adopters is a positive offering and must continue to evolve as digital transformation progresses.

In terms of the boundaries of regulation, as generative AI devices will adapt to data exposure over time, the regulation system will need to decide when a re-validation trigger is met by the device – is that triggered as a pre-specified time interval, model update threshold or other pre-determined factor? The point being that post-marketing surveillance relying on reports of harm is not sufficient for this technology.

Question 5: How should the regulatory framework manage post-market surveillance for AI health technologies? (optional)

You may wish to consider some or all of the following in your response: (1) The challenges posed by novel and emerging types of AI, including foundation models and highly capable agentic AI, (2) AI systems which are capable of continuous learning and/or updating, (3) AI systems that are used for other purposes beyond the original intended use, (4) AI systems which are developed by a single institution for in-house use only, (5) Information sharing between healthcare provider organisations and manufacturers for the purposes of post-market surveillance.

There needs to be a formal post-marketing reporting system for all incidents relating to use of AI devices in healthcare contexts (not just those that may be designated as serious ones), which will enable unintended consequences arising from a) autonomous behaviour of the device, b) continuous learning by the device, or c) use of the device outside manufacturer's original remit. Under-reporting of incidents is a common issue with pharmacovigilance with medicines, but this in theory could be helped by systems within the software to enable the reporting of issues. Software suppliers should also monitor issues reported in academic and professional literature.

The regulator and developers need to know when a re-validation trigger is met by the device – is that triggered as a pre-specified time interval, model update threshold or other pre-determined factor? The point being that post-marketing surveillance relying on reports of harm is not sufficient for this technology.

Question 6: Which statement best reflects your view on the current legal framework for establishing liability in healthcare AI tools? (optional)

Sufficient

Gaps exist

Insufficient

I am unsure

Question 7: How could manufacturers of AI health technologies, healthcare provider organisations, healthcare professionals, and other parties best share responsibility for ensuring AI is used safely and responsibly? (optional)

You may wish to consider some or all of the following in your response: (1) The specific duties for each party, and (2) any duties which are shared.

Responsibility for the safe and responsible use of AI in healthcare should be shared across three groups — manufacturers, healthcare providers, and clinicians — with clearly defined duties for each. Manufacturers should ensure their AI systems meet strict safety, performance, transparency, and monitoring standards. Healthcare provider organisations should integrate AI safely into clinical workflows, maintain appropriate governance and oversight, and ensure staff are trained to use AI tools correctly. Clinicians should apply professional judgment when using AI-supported decisions, understand system limitations, and escalate safety concerns.

Some responsibilities — such as monitoring performance, reporting safety incidents, protecting patient data, and ensuring AI is used within its intended purpose — should be shared across all parties. Establishing these duties formally, potentially through statute, could mirror the structured regulatory model used in the Human Fertilisation and Embryology Act 1990, providing clarity, accountability, and consistent safeguards across the system.

Question 8: In the event of an adverse patient outcome where an adverse patient outcome involved an AI tool, where do you think liability should lie? (optional)

You may wish to first consider the following scenarios: (1) When the AI tool gives the correct answer, but is incorrectly overridden by the healthcare professional, (2) When the AI tool gives the incorrect answer and the healthcare professional follows it (i.e. they incorrectly choose to trust the AI).

From current evidence, AI devices are likely to be more accurate in some respects but misleading in others – validation and reporting will be vital to determine the difference. On one hand, many well-trained AI diagnostic tools are more accurate than a human operator. On the other, however, generative AI that relies on data collation from a range of sources are prone to the “hallucination” effect, where an output looks right, but is not.

Studies indicate that, with conventional decision support systems in prescribing (Straichman YZ, Kurnik D et al. Prescriber response to computerized drug alerts for electronic prescriptions among hospitalized patients. *International Journal of Medical Informatics*. 2017; 107: 70-5), clinicians regularly override alerts correctly, suggesting

these alerts are of limited value. There is therefore a “culture” among clinicians of overriding alerts, which may not be helpful if an AI-generated alert is accurate but where its provenance cannot be easily demonstrated.

Moreover, with generative AI tools, significant factual accuracy issues have been demonstrated. For example, testing of ambient scribe tools indicates that they may not be accurate with medicine names (Wang H, Yang R et al. An evaluation framework for ambient digital scribing tools in clinical applications. *NPJ Digital Medicine*. 2025; 8: 358), which is especially worrying in clinical context. This is where the risk classification of AI technology in healthcare is important – ambient scribe tools may be classified as low risk, but it could be argued that these require post-marketing surveillance to ensure ongoing safety and performance.

Question 9: Do you have any other evidence to contribute? You can submit written evidence in the comment box. Note: please confirm that you have the necessary permissions prior to sharing any documents in this way. (optional)

Developments on AI regulation should be based on and extend the scope of guidance such as:

NHS England. Guidance on the use of AI-enabled ambient scribing products in health and care settings. At <https://www.england.nhs.uk/long-read/guidance-on-the-use-of-ai-enabled-ambient-scribing-products-in-health-and-care-settings/>.

And

NHS England. Planning and implementing real-world artificial intelligence (AI) evaluations: lessons from the AI in Health and Care Award. At <https://www.england.nhs.uk/long-read/planning-and-implementing-real-world-ai-evaluations-lessons-from-the-ai-in-health-and-care-award/>.

And our own Royal Pharmaceutical Society policy on Artificial Intelligence in Pharmacy Practice. [Artificial Intelligence \(AI\) | RPS](#)

All these documents are in the public domain.

Question 10: You can upload documents to be considered as part of this call for evidence. Note: please confirm that you have the necessary permissions prior to sharing any documents in this way. (optional)

