Royal College of General Practitioners (RCGP) organisational response to the Medicines and Healthcare products Regulatory Agency (MHRA) – National Commission into the Regulation of AI in Healthcare: call for evidence

January 2026

About RCGP

We are the professional membership body for GPs in the UK. Our purpose is to encourage, foster and maintain the highest possible standards in general medical practice. We support GPs through all stages of their career, from medical students considering general practice, through to training, qualified years and retirement.

Call for Evidence Questions

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

- No change: The framework should be maintained as it is
- Minor adjustments: The current framework works but requires small changes
- Significant reform: The current framework requires substantial changes
- X Complete overhaul: The overall framework should be replaced entirely

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

- X Strongly disagree
- Disagree
- Neither agree nor 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

- Strongly disagree
- Disagree
- X Neither agree nor 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

- Strongly disagree
- X Disagree
- Neither agree nor 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

- Strongly disagree
- Disagree
- Neither agree nor disagree
- X 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

- X Strongly disagree
- Disagree
- Neither agree nor disagree
- Agree
- Strongly agree

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

- Too restrictive [stifles innovation]
- Somewhat restrictive [creates some barriers]
- About right [balances safety and innovation]
- Somewhat loose [lacks necessary controls]
- X Too loose [risks patient safety]

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?

[MHRA: 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 UK's current regulatory framework is fragmented. Responsibility is shared across multiple bodies, including the MHRA and the CQC, which can lead to inconsistent oversight, duplication and unclear accountability. GPs often look to regional NHS bodies, such as integrated care boards (ICBs) and commissioners, for guidance on AI adoption. This guidance is not yet consistent in how AI tools should be assessed, approved and used across the UK.

The joint RCGP and Nuffield Trust report, 'How are GPs using AI? Insights from the front line' (Dec. 2025), highlights the consequences of this fragmentation. Among GPs surveyed, 'a lack of regulatory oversight on AI' was the second largest concern regarding AI adoption in general practice. This uncertainty is slowing adoption and leading to ungoverned use.

To improve the regulatory framework, we recommend the following:

- Clear, centralised UK-wide guidance on the specific use of AI in general practice. While guidance already exists, a consistent set of expectations for safety, governance, accountability and liability is needed to reduce reliance on variable local interpretation. This guidance should also address the environmental impact of AI, ensuring alignment with NHS net zero commitments. We would support a use-case-by-use-case approach, as demonstrated by the NHS England Ambient Voice Technology Self-Certified Supplier Registry. It is critical that the GP profession is appropriately engaged in the development of guidance, including via discussions with the Joint GP IT Committee of the BMA and the RCGP.

- Regulation should be risk proportionate. We agree with the recommendation from the MHRA that AI tools deployed in general practice should be considered by default as not less than Class II medical devices.

- Regulators should define what the intervals should be for when software updates, retraining, or model drift trigger revalidation or regulatory review.

- Improved clarity around how patient data used by AI tools is handled, including where data is stored, who controls it, retention policies, and whether it is used for secondary purposes. This is particularly important for tools that generate or rely on recordings or transcripts of clinical interactions, where it is often unclear whether such data forms part of the medical record and is subject to rights such as Subject Access Requests. In addition, to protect public trust, patients and clinicians must be given a meaningful ability to opt out of the use of AI tools, where AI involves recording/transcription, secondary use of their data, or direct patient-facing generation of content, with a clear non-inferior alternative pathway. For AI that is embedded in core operational or safety

functions, patients should instead be provided with transparency, accessible explanations, and a route to challenge/correct outputs.

To qualify our answer to Q.2.4, the clinical evidence that the MHRA requires to assure minimum viable product in a medical device is not a cause of concern for us. There are, however, deep concerns in relation to the research evidence base required to justify the proposed benefits of the current generation of tooling, which requires high level strategic assessment and is not in scope for regulation.

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

[MHRA: 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.]

The strengthening of post-market surveillance requirements for AI health technologies is welcome. However, to be effective in practice, frontline clinicians need clearer guidance on what constitutes an AI-related safety incident and when it should be reported. Without this, under-reporting is likely, particularly where harms are subtle, cumulative, or linked to model performance rather than a single clinical event.

Reporting mechanisms must be simple, well-signposted, and integrated into existing systems. At present, inefficient reporting mechanisms – often lacking clear consequences or feedback – can discourage reporting. While the MHRA's Yellow Card scheme is a key mechanism for identifying problems with AI tools classified as medical devices, we are aware that very few AI-related reports have been submitted to date, despite anecdotal reports of performance and safety concerns. We are concerned that GPs may not submit reports via the scheme, particularly where issues are perceived as minor or ambiguous. We therefore call for a transparent, central reporting mechanism for all digital clinical safety concerns, and for alignment between regulatory reporting and the contractual requirement for GPs, from 1 October 2025, to register with the Learn from Patient Safety Events (LFPSE) reporting system. Regulation and central reporting must be fully integrated.

Effective post-market surveillance for AI will also require a hybrid approach that combines clinician-reported incidents with automated system monitoring. Regulators should set expectations for manufacturers and deployers to monitor real-world performance, including detecting model drift, bias and unintended or off-label use, and ensure that relevant findings are shared in a timely and standardised way to support regulatory oversight and action. Integrating these technical monitoring

requirements with clinical reporting would allow for the earlier identification of risks and more proportionate regulatory intervention, supporting both patient safety and continued use of effective technologies.

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

- Sufficient: existing laws (e.g. medical negligence, Product Liability etc.) can adequately handle AI-related disputes
- Gaps exist: existing laws work for most cases but leave uncertainty in some scenarios
- X Insufficient: existing laws are unfit for AI
- 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?

[MHRA: 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.]

We suggest that manufacturers and developers would be best placed to hold primary responsibility for the safe design, development and maintenance of AI systems. We anticipate that this would include transparency around intended use, performance and limitations; accurate and realistic device classification; clear labelling; and the safe management of updates and retraining. It should be standard practice for manufacturers to involve clinicians in the design of AI tools.

Manufacturers should also actively monitor real-world performance post-deployment and communicate emerging risks. They must clearly describe known risks in all training and supporting materials, and design systems to mitigate automation bias, for example, through prompts that encourage clinicians to check outputs. Vendor-led demonstrations should not be treated as sufficient training or as a transfer of medico-legal responsibility. At least during early adoption, there may be value in standardised baseline AI safety training for users, refreshed annually, focused on limitations, accountability and reporting.

National guidance and registries will be critical in providing appropriate reassurance for commissioners and procurers of AI on the regulatory status and evidence base of tools. We therefore welcome the introduction of NHS England's Ambient Voice Technology Self-Certified Supplier Registry, which requires suppliers to comply with standards on clinical safety, technology and data protection. We hope this represents a step towards a comprehensive national approach to assurance and oversight of AI in general practice. In addition, procurers of AI tools, including ICBs and GP practices, should be encouraged, and supported, to challenge manufacturers' claims and

classifications, if they have concerns, and given resourcing to ensure that appropriate training is in place.

While national guidance and registries are necessary, the rollout of AI exposes the lack of applicability of the DCB0129 and DCB0160 standards, which were developed for secondary care organisations commissioning an electronic patient record. These standards are not fit for purpose for the deployment of AI systems and tooling at scale in general practice, where clinical safety can only be effectively managed through system-level, cross-organisational risk assessment and management. This requires appropriate resourcing at system level, rather than expecting every individual practice to undertake similar or identical risk assessment and management activities.

Healthcare professionals should remain responsible for using AI tools within their intended scope, applying professional judgement, and acting on or escalating concerns when outputs appear unsafe, inappropriate or misleading. Clear expectations are needed that AI supports, rather than replaces, clinical decision-making, and that accountability for patient care remains with the clinician.

Those commissioned to evaluate AI health technologies should share responsibility for ensuring evaluations meet recognised ethical and scientific standards. Balanced, transparent evaluation is essential to support safe adoption, informed decision-making and public confidence in the use of AI in general practice.

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?

[MHRA: 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).]

The framing of AI outputs as ‘correct’ or ‘incorrect’ is problematic and does not reflect the realities of clinical decision-making. Medical practice is holistic and involves probability and judgment that extends beyond what most current AI systems can capture.

A core part of clinical expertise, particularly in general practice, is the ability to interpret guidelines and tools in the context of the individual patient, including patient preferences and social factors. Shared decision-making with patients complicates the notion of a single 'correct' answer. If a clinician has exercised reasonable professional judgment and can justify why the AI recommendation was not followed, there must not be automatic expectation of any liability simply for overriding the tool. AI systems must be viewed as decision-supporters, not decision-makers.

For a clinician to reasonably rely on an AI tool, there must be transparency about its accuracy, limitations and likelihood of error. Currently, many AI outputs are presented with an unwarranted sense of certainty, without confidence intervals, error rates or clear communication of risk, and the eloquence effect elevates the risk of acceptance of erroneous information. If an AI system is deployed without adequate validation, monitoring and disclosure of its performance characteristics, then developers, vendors and healthcare organisations must share responsibility for any resulting harm.

Accountability mechanisms are essential. AI systems integrated into healthcare should only be adopted when there is clear evidence of their accuracy, an understanding of how often they fail and defined consequences for when they do. As a practical safety measure, AI tools should include explicit human-in-the-loop safeguards: clear labelling of AI-generated recommendations, audit trails showing clinician review and preserved clinician autonomy to challenge or reject AI outputs. Doctors must not be pressured to follow AI recommendations.

Assigning liability based on whether an AI answer was 'correct' or 'incorrect' oversimplifies clinical reality and risks undermining both patient safety and professional judgment.

In the system-level approach to clinical safety management, liability is passed to the organisation(s) that have the ability to best mitigate the risks.

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.

The joint RCGP and Nuffield Trust report, 'How are GPs using AI? Insights from the front line' (Dec. 2025) shows the proportion of GPs in the UK currently using AI in their clinical practice and explores how they are using it. We sought to understand the range of tasks for which GPs use AI, the perceived benefits and concerns, and the barriers and enablers to greater adoption of AI in general practice.

- RCGP and Nuffield Trust. How are GPs using AI? Insights from the front line. 2025 Dec. (https://www.nuffieldtrust.org.uk/sites/default/files/2025-12/How%20are%20GPs%20using%20AI_%20final.pdf)

We consider it essential for there to be an evidence base for the use of AI in healthcare, and appropriate clinical trials to evaluate AI tools. The scientific method must be applied before widescale deployment: predefined hypotheses, external validation and prospective trials are needed where patient outcomes are at stake.

- NIHR. Driving safe evaluation of AI systems in healthcare (CONSORT-AI). 2025 Mar 3. (https://www.nihr.ac.uk/story/driving-safe-evaluation-ai-systems-healthcare):

  "For any new treatment, test or other health intervention, we want to know that it is safe, clinically effective, and cost effective. This should be true for all interventions, including those involving artificial intelligence (AI) systems.

  There is a need to move from the initial 'proof-of-concept' studies, which show the potential of the technology, to clinical trials that evaluate these AI systems within their intended clinical pathways, and which measure the real benefits (or harms) to patients.

  This evaluation in the real world is critical if we want to rigorously and robustly evaluate the safety, clinical effectiveness and cost effectiveness of AI health interventions."

Environmental impacts of AI: Data-centre cooling requires substantial water use; procurement should require disclosure of energy and water consumption.

- Diagnostic and Interventional Imaging. Climate change and artificial intelligence in healthcare: Review and recommendations towards a sustainable future. 2024 Nov. (https://www.sciencedirect.com/science/article/pii/S2211568424001384)

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.

N/A