National Commission into the
Regulation of AI in Healthcare

National Commission into the Regulation of AI in Healthcare

Recommendations for a future regulatory framework

Published 10 September 2026

Contents

Produced by the Medicines and Healthcare products Regulatory Agency and National Commission into the Regulation of AI in Healthcare. www.gov.uk/mhra

The Medicines and Healthcare products Regulatory Agency (MHRA) established the National Commission into the Regulation of AI in Healthcare in September 2025, to review current regulations and provide recommendations for a new regulatory framework for AI in healthcare. As the sponsoring organisation, the MHRA commissioned and oversaw the delivery of the Commission's programme of work. This included the establishment of the Commission, governance, management of the research programme, stakeholder engagement, secretariat and operational support. The Commission operated independently in developing its advice and recommendations, and responsibility for the report's conclusions and recommendations rests solely with the Commission.

Foreword

Healthcare is about people. In healthcare we are interested in technology not for technology's sake, but because of its potential benefits to people. We want to see more accurate diagnostics, more effective treatments, and better access to sustainable, high-quality services. Any technology – including artificial intelligence (AI) – has to prove its worth. As with any technology we need to consider the opportunities and the risks. We require evidence of the benefits, the safety profile and downstream consequences. We need to ensure that our systems can be trusted to bring safe and effective technologies to patients and the health system, whilst keeping out those that are unsafe, ineffective or not yet ready.

The National Commission into the Regulation of AI in Healthcare is people-centred rather than tech-centred. It is about defining what we want our future AI-enabled healthcare system to look like, and how we use regulation to shape that future. In chairing this Commission, we have brought our experience as a hospital doctor and a general practitioner, to keep the discussions, however technical, grounded in the reality of healthcare, and focused on people. We have been supported in this endeavour by a huge number of patients, the public, healthcare professionals, innovators, regulators, policymakers, healthcare leaders and others, who have generously shared their perspectives on the opportunities and challenges presented by AI in healthcare. Contributing through individual conversations, formal public dialogues, large-scale surveys, 'Call for Evidence' submissions, expert working groups, and international engagements, all these insights have influenced the Commission and its recommendations.

What we heard was not a debate about *whether* AI should be used in healthcare, but rather *how* it should be used. People are simultaneously enthusiastic about its potential benefits, and worried about the risks. They want to see healthcare 'catch up' with other sectors where this would lead to faster, more personalised services; but they also want reassurance that any AI is introduced responsibly and equitably, that accountability remains clear, that concerns are heard and acted upon, and that the benefits of innovation are felt by patients and healthcare staff. Trust, in other words, cannot be assumed. It must be earned.

People want to know that any move towards greater use of AI in healthcare has been considered in terms of its impact not only on some people, or on the majority, but on everyone. We have viewed every discussion, and every recommendation through that lens. The changes described here are designed to improve safety for everyone, to accelerate beneficial innovation for everyone, and to do whatever we can to ensure that AI in healthcare reduces inequality gaps rather than worsens them. Every single one of our recommendations should be viewed through this lens.

In being people-centred, we have considered not only what AI means for all who *use* healthcare, but also what it means for all those who *provide* healthcare. Patients, carers and the wider public have consistently stressed the importance of human connection, and of human oversight. As healthcare professionals, our roles may change, but it is clear that we still have a role and that role is valued. The challenge from the public is to find a path for 'AI to do what AI is good at, and humans to do what they are good at'. Indeed, there is an

important opportunity here for AI to improve the working lives of healthcare staff and reduce the risk of burn-out. Healthcare professionals recognise many routine and frustrating tasks that simply get in the way of us delivering high-quality care, and we would welcome the opportunity for AI assistance on those tasks to allow us to spend more time on those parts of our role where we can make the most difference.

This report represents the culmination of one of the most extensive programmes of research and engagement undertaken on AI in healthcare. Our recommendations have been shaped by evidence gathered from across the four nations, alongside the expertise of Commission members, working groups and partner organisations. Most importantly, they have been informed by the people who will ultimately be affected by the decisions that follow.

We would like to express our sincere thanks to everyone who contributed to this work. To those who responded to the Call for Evidence, participated in public engagement activities, joined roundtables and working groups, challenged our assumptions and shared their expertise and experience: this report is stronger because of your contribution.

Publication of this report is not the end of a conversation. It is the beginning of the next phase. The recommendations set out here will require continued collaboration between government, regulators, healthcare providers, professionals, industry, patients and the public. Success will depend not on the actions of any single organisation, but on a shared commitment to delivering this together.

The regulatory framework set out in this report addresses both the opportunities and the risks of AI in healthcare. We aim to ensure that patients can benefit earlier from AI technologies that are safe, effective, and equitable; to provide greater confidence to health professionals in their delivery of healthcare in an AI-enabled world; and to give greater clarity to innovators as they develop these technologies. This is about creating a future healthcare system that is increasingly tech-enabled and always people-centred.

Professor Alastair Denniston
Chair, National Commission into the Regulation of AI in Healthcare
Professor of Regulatory Science and Innovation at the University of Birmingham
Honorary Consultant Ophthalmologist, University Hospitals Birmingham NHS Foundation Trust
Executive Director, Centre of Excellence for Regulatory Science in AI & Digital Health (CERSI-AI)

Professor Henrietta Hughes
Deputy Chair, National Commission into the Regulation of AI in Healthcare
Patient Safety Commissioner for England
General Practitioner and Visiting Professor, Institute of Medicine, University of Greater Manchester

Executive summary

Artificial Intelligence (AI) has the potential to transform healthcare. It can support earlier diagnosis, more personalised care, improved patient outcomes and better experiences for clinicians, while helping the NHS and wider health system meet rising demand. AI should augment, rather than replace healthcare professionals. Used responsibly, AI can automate routine and administrative tasks, streamline patient care and enable professionals to focus on communication, compassion and shared decision-making. However, software and AI products also present regulatory challenges that differ from those associated with traditional medical devices. Unlike many conventional devices, AI products may evolve over time, with their performance being context-dependent and sensitive to real-world deployment environments.

The National Commission into the Regulation of AI in Healthcare was established to advise government on a future regulatory framework for AI in healthcare. In doing so, it is considering not only the regulation of safe and effective software and AI-enabled medical devices but also wider issues such as accountability, transparency, clinical practice, organisational governance and system-wide level assurance. Its mission is to help ensure that patients can benefit from safe and effective AI technologies, support the government's ambition to make the NHS the most AI-enabled healthcare system in the world, and establish a globally competitive regulatory environment that attracts investment and supports innovation.

The Commission's recommendations are grounded in a research and engagement programme, including the Call for Evidence, public deliberation events, engagement with groups who are seldom-heard, professional and industry roundtables, specialist working groups and further engagement across government and the health system.

What the evidence showed

The evidence showed strong support for the use of AI in healthcare, but this support is conditional, rather than automatic.

People are not asking simply whether AI should be used in healthcare. They are asking how it can be used safely, fairly and transparently, with clear evidence of benefit, meaningful human oversight, accountability and continued public trust.

What needs to change

The Commission's central conclusion is that the regulation and assurance of AI in healthcare must become more proportionate, lifecycle-based and system-wide.

Current approaches were largely designed for products that are more static and easier to reliably assess at a single point in time. AI-enabled products may iterate rapidly, perform differently in different settings and depend on the data, workflows, people and organisations around them.

A framework that relies too heavily on one-off pre-market assessment will not be enough. Regulation must support safe, effective and trusted use throughout the lifecycle of a device, from development and deployment through to monitoring, updating and learning from real-world use. More continuous monitoring of these products, once deployed, will be needed to provide ongoing assurance across the lifecycle.

The United Kingdom (UK) needs a framework that is safe, fast and trusted: safe in protecting patients and managing risk; fast in enabling responsible innovation and timely access to beneficial technologies; and trusted in giving confidence to patients, professionals, providers, developers, manufacturers, regulators and government.

The Commission's overall vision

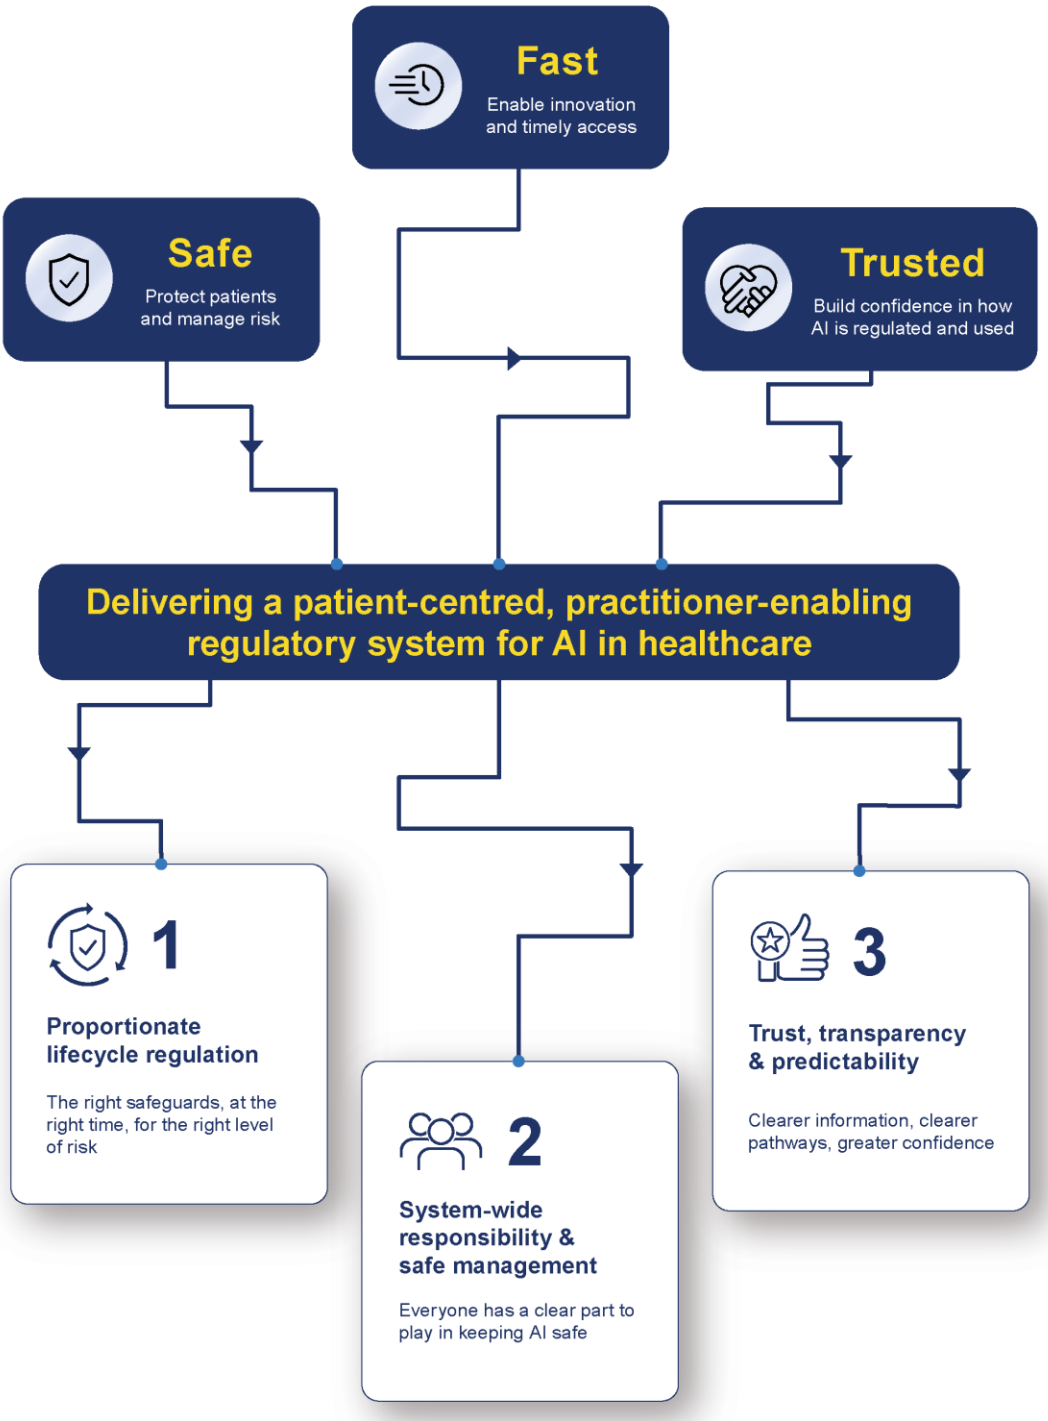

Delivering the vision

The Commission's recommendations are organised around three connected areas of change.

1. Proportionate lifecycle regulation

AI-enabled devices require regulatory approaches that reflect their distinctive characteristics. The future framework should support proportionate and tailored oversight across the product lifecycle, including clear routes to market, flexible mechanisms to support rapid device iteration and stronger use of real-world evidence to support ongoing assurance.

Key themes include:

- proportionate regulation focused on risk and benefit
- health equity as an ethical imperative and a critical aspect of safety and performance
- improved device qualification and classification pathways
- flexible mechanisms for safe device iteration and updating
- staged routes to market that support innovation while generating real-world evidence
- greater use of regulatory sandboxes
- enhanced enforcement mechanisms
- international leadership, harmonisation and recognition

2. System-wide responsibility and safe management

Safe and effective use of AI in healthcare depends on more than regulation alone. Regulation cannot be the only enabler of change; the safe and effective use of AI will require a collaborative, system-wide effort. Manufacturers, healthcare providers, healthcare professionals, regulators and policymakers all have a role to play. The future framework should support clear responsibilities, organisational readiness, workforce capability and effective governance across the healthcare system.

Key themes include:

- fostering a culture of safety for everyone in relation to AI
- AI readiness toolkit to support healthcare organisations

- clear responsibilities throughout the product lifecycle, including liability arrangements
- strong governance arrangements for implementation and monitoring
- workforce capability and AI literacy
- best practices for procurement, deployment and risk controls

3. Trust, transparency and predictability

Public confidence is essential to realising the benefits of AI in healthcare. The future framework should strengthen transparency, accountability and patient involvement, while providing clearer and more predictable pathways for innovators.

Key themes include:

- system approach to improved transparency of AI used in care pathways
- enhanced, inclusive patient and public engagement in policymaking processes
- stronger user-centred design and transparency for device users
- improved communication about safety concerns for patients and healthcare professionals
- educational resources and tools to demystify the regulatory journey
- predictable mechanisms for early engagement with regulators to support innovation

What this means in practice

Together, these three areas of change support the Commission's vision for a regulatory and assurance framework that is safe, proportionate, trusted, flexible and works for everyone. The recommendations are designed not only to address today's challenges, but also to create a framework that can evolve as technologies, evidence and patterns of use change over time. In doing so, they support the UK's ambition to become a global leader in AI-enabled healthcare innovation and regulation.

The Commission is therefore not proposing a single new process or a narrow set of technical fixes. It is recommending a modern framework that recognises the distinctive characteristics of AI-enabled products and supports safe innovation throughout the product lifecycle. The framework is intended to protect patients, support healthcare professionals, provide confidence for healthcare providers and create a predictable environment for innovators.

In practice, this means:

- **For patients and the public**, greater confidence that AI technologies used in healthcare are safe, effective, inclusive and properly governed, alongside greater transparency about when and how AI is used in their care and clearer routes to raise concerns when things go wrong.

- **For healthcare professionals**, clearer expectations regarding the safe use of AI, better access to training and support, stronger mechanisms for reporting concerns, and greater clarity regarding accountability and responsibilities.

- **For healthcare providers**, practical tools to assess organisational readiness, deploy technologies safely, monitor performance over time and demonstrate effective governance of AI-enabled products.

- **For developers and manufacturers**, a pragmatic, proportionate and innovation-friendly regulatory framework that provides clearer expectations, more predictable pathways and greater regulatory clarity throughout product development, deployment and improvement. Regulation should not operate as a barrier, but as a framework that helps innovators, regulators and healthcare systems work together to ensure new technologies are safe, effective and ready for real-world use.

- **For regulators and government**, a more joined-up system that can learn from real-world evidence, respond to emerging risks, adapt as technologies evolve, maintain public confidence and support the UK's long-term global leadership in AI-enabled healthcare innovation.

Next steps

AI presents one of the most significant opportunities to improve healthcare in a generation. Realising that opportunity will require a regulatory and assurance framework that evolves alongside the technology itself: protecting patients, enabling responsible innovation and maintaining public trust. The Commission's recommendations aim to provide a foundation for the Medicines and Healthcare products Regulatory Agency (MHRA), government and wider system partners to achieve that ambition.

A cross-government response will follow separately, setting out how government and system partners will consider and take forward the recommendations.

Introduction

Background

Through the 10 Year Health Plan for England and the Life Sciences Sector Plan, the UK Government has set out an ambition for the NHS to become one of the most AI-enabled healthcare systems in the world. AI has the potential to improve health outcomes, support healthcare professionals, increase productivity and contribute to growth in the UK's life sciences sector.

At the same time, AI is creating new questions about regulation, governance, accountability and public confidence[1]. These questions extend beyond any single regulatory framework or organisation and affect how technologies are developed, deployed, monitored and used across the health system.

AI used in healthcare may fall within different regulatory and governance regimes depending on its intended purpose, functionality and context of use. Some AI-enabled products will qualify as medical devices and be regulated under the UK Medical Devices Regulations 2002 (as amended) and associated MHRA requirements. Others may be governed primarily through data protection law, professional standards, organisational governance, consumer protection law or other sector-specific frameworks.

AI is used across a wide range of healthcare applications, from administrative tools to decision support for healthcare professionals and direct-to-consumer products, and not all applications are regulated in the same way. As adoption increases, questions have emerged around qualification and classification of software and AI-enabled medical devices, visibility of use across the health system, post-market monitoring in real-world settings, accountability and the suitability of existing regulatory approaches for technologies that may be frequently updated over time.

These developments have prompted consideration of whether regulatory and assurance approaches should evolve to better reflect how AI technologies are developed, deployed and used in practice.

[1] Bengio et al (2026) International AI Safety Report 2026.

Scope of the National Commission

The National Commission considered the wider regulatory ecosystem for AI in healthcare. This includes product regulation, professional regulation, organisational governance, patient safety and redress, public confidence, and the wider systems and institutions that support the safe adoption and oversight of AI in healthcare.

While many of the Commission's recommendations relate to software and AI-enabled medical devices, the Commission's remit was broader than UK medical device regulation alone. The report addresses AI products regulated as medical devices and considers wider uses of AI in healthcare where issues such as safety, accountability, governance, transparency and public confidence are relevant. Throughout this report, the term 'product' is used deliberately to encompass both AI-enabled medical devices and other AI-enabled products used in healthcare that are not subject to UK medical device regulation.

The Commission considered how different forms of regulation, assurance, governance and oversight should work together to support the safe, effective and responsible use of AI in healthcare. This includes the roles of government, regulators, healthcare providers, professionals, developers, manufacturers, patients and the public.

The Commission recognises that successful implementation of its recommendations will require regulatory bodies to be appropriately resourced and equipped with the necessary capability and capacity.

The National Commission and its approach

The National Commission into the Regulation of AI in Healthcare was established by the MHRA in September 2025 as an independent expert advisory body. It was asked to consider how the UK's regulatory and assurance framework should evolve to support the safe, effective and responsible use of AI in healthcare.

Chaired by Professor Alastair Denniston and Deputy Chair Professor Henrietta Hughes, the Commission brought together expertise from healthcare, technology, law, patient groups, the public, government, the NHS, and partners across England, Scotland, Wales and Northern Ireland. The Commission also sought international expertise, inviting stakeholders from partners in the United States (US), Europe and Singapore to support its work.

Its work was supported by a core group of experts, providing strategic oversight and challenge, and four specialist working groups:

Cross Whitehall Working Group Brought together government departments and public bodies to consider policy alignment and wider regulatory and system implications.	**Devolved Authorities Working Group** Brought together partners from across England, Scotland, Wales and Northern Ireland to reflect different health systems, delivery contexts and responsibilities.
Health Systems Working Group Considered how AI-enabled technologies operate within healthcare pathways, including governance, adoption, operational readiness and safe management.	**Technology Working Group** Provided technical insight on AI capabilities and their regulatory implications, helping ensure recommendations were grounded in how technologies are developed and deployed.

Together, these groups helped ensure that the recommendations were informed by practical delivery considerations as well as regulatory principle.

What we heard

The Commission's recommendations were informed by a UK-wide research and engagement programme coordinated on behalf of the Commission by the MHRA[2]. This included an open Call for Evidence, targeted engagement with patients, healthcare professionals and industry, public deliberation sessions delivered with The Health Foundation, engagement with seldom-heard groups supported by National Voices, and structured clinical engagement with healthcare leaders. Taken together, this represented one of the most extensive programmes of engagement undertaken on AI in healthcare in the UK.

The evidence highlighted both the opportunities presented by AI and the expectations that accompany its adoption. Stakeholders consistently emphasised the importance of maintaining patient safety, public confidence, transparency and accountability while enabling innovation and beneficial technologies to reach all patients, so that the benefits of AI are felt by everyone.

The recommendations that follow seek to respond to those expectations through a framework that is proportionate, lifecycle-based, flexible and focused on supporting equitable access to safe and effective technologies.

[2] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare: research, engagement and call for evidence findings.

How the Commission reached its recommendations

Evidence, expertise, engagement

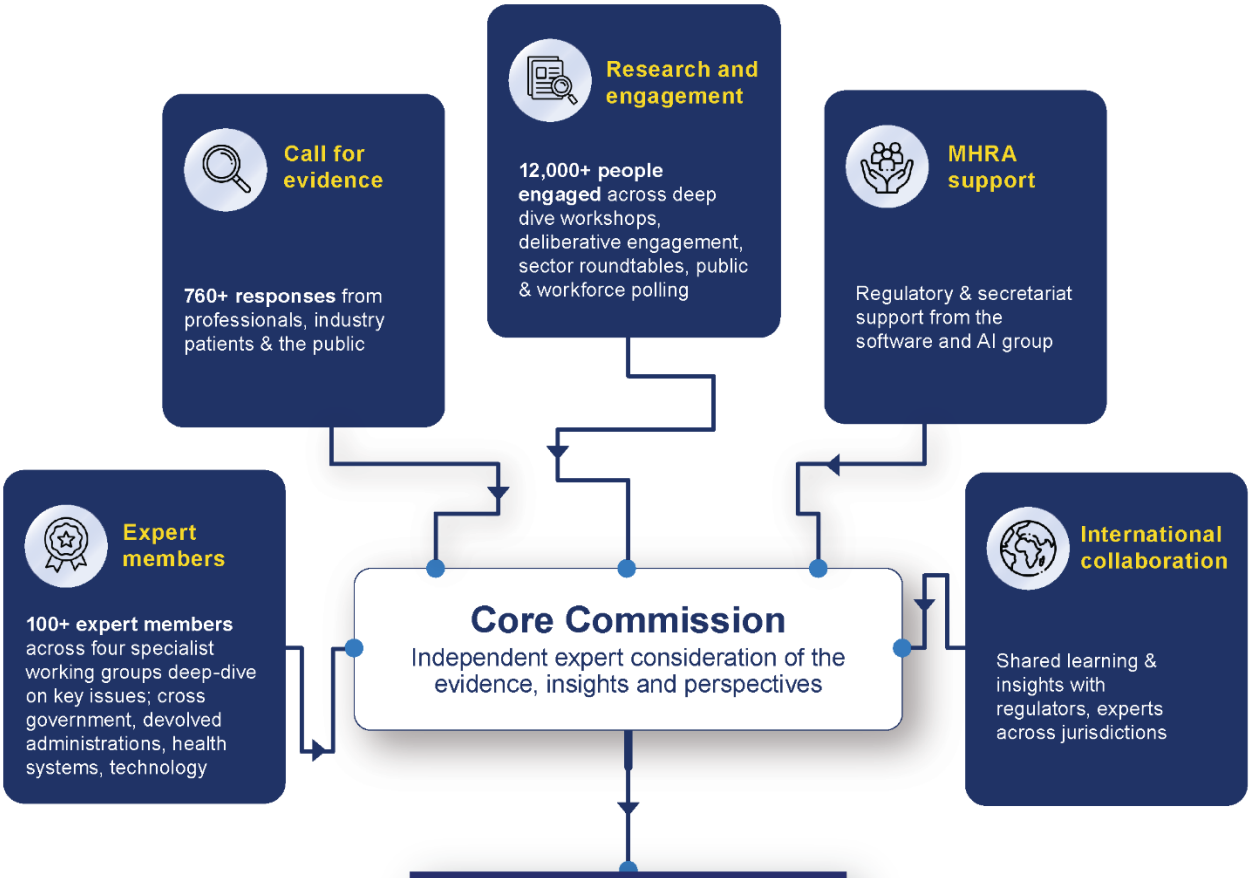

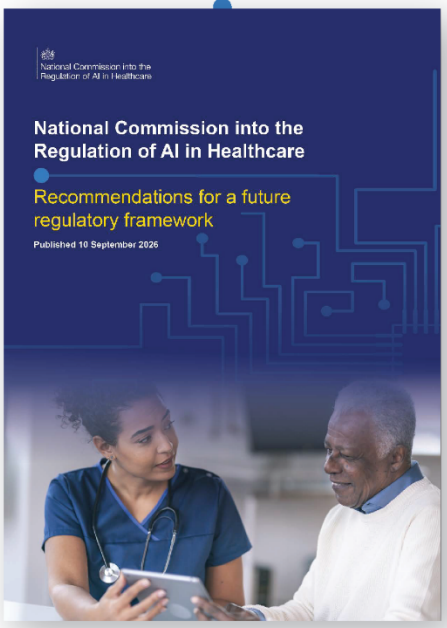

Purpose of this report

This report sets out the Commission's recommendations for the future regulation and assurance of AI in healthcare.

The recommendations are intended to:

- support timely and inclusive access to safe and effective AI-enabled technologies
- strengthen patient safety and public confidence
- support healthcare providers and professionals in the safe adoption of AI
- provide greater clarity and predictability for developers and manufacturers
- help ensure that regulation and governance remain effective as technologies evolve

The report is directed at government, regulators, healthcare providers, professional bodies, developers, manufacturers and others with a role in shaping the future use of AI in healthcare.

Structure of this report

The report is organised around three connected themes.

Chapter 1: Proportionate lifecycle regulation

This chapter considers how regulatory approaches should evolve to reflect the characteristics of AI-enabled technologies. It explores how regulation can be proportionate to risk and benefit, informed by evidence generated across the product lifecycle, and designed to support timely and equitable access to safe and effective technologies.

Chapter 2: System-wide responsibility and safe management

This chapter examines how the safe adoption, management and oversight of AI-enabled technologies should be supported across the healthcare system. It considers how responsibilities should be shared across manufacturers, healthcare providers, professionals, regulators and other system partners, and how accountability, organisational readiness, workforce capability and governance can support their safe and effective use.

Chapter 3: Trust, transparency and predictability

This chapter explores how trust, transparency and predictability can be strengthened across the healthcare ecosystem. It considers how greater transparency, clearer accountability, meaningful stakeholder involvement and more navigable pathways can support the responsible development, adoption and use of safe and effective AI technologies.

Together, these chapters set out the Commission's recommendations for a regulatory and assurance framework that is safe, proportionate, trusted and flexible, supporting both innovation and public confidence in AI-enabled healthcare.

A key consideration for the Commission was to define its scope, in terms of both breadth and depth. In terms of breadth, the Commission centred on those AI-enabled technologies which are involved in direct patient care. Particular attention is given to technologies that qualify as medical devices, given their risks and regulatory requirements. In terms of depth, the Commission operated at four tiers, and the recommendations are aligned accordingly.

- **Tier 1** describes the overall intent in terms of what the new regulatory framework should be characterised to be: person-centred, safe, fast and trusted. This is the 'why'.
- **Tier 2** describes the three key regulatory principles that emerged during the Commission and starts to provide some degree of specificity to the key elements of the proposed new framework. Each chapter reflects one of these three principles. This is the 'what'.
- **Tier 3** describes the specific regulatory innovations and changes that are necessary in order to achieve the desired new framework. This has a much higher degree of specificity and is contained within the individual recommendations. This is the 'how'.
- **Tier 4** would be the detail of the policy change itself, and needs to sit with the relevant regulator, government department or other body responsible for those actions.

The Commission sought to balance its remit to set direction and shape the new regulatory framework, without stepping into the remit of the policymakers themselves; in responding to these recommendations, it will be important that policymakers have the flexibility to respond in line with a fast-moving technology field and a fast-moving policy landscape.

Testing the recommendations across the UK system

Healthcare is a devolved policy area in the UK. Responsibility for healthcare policy sits with the Department of Health and Social Care (DHSC) in England and the devolved health departments in Scotland, Wales and Northern Ireland, while healthcare is delivered through the respective national health systems. These health systems are further supported by an ecosystem of national and devolved bodies, including sector regulators, professional regulators, representative bodies, and national guidance organisations (see Glossary of terms and acronyms).

Unless otherwise stated, references to the relevant healthcare organisations should be understood to include all organisations that may have a role in implementing, supporting or overseeing recommendations, as appropriate within England and each of the devolved

authorities. The following terms are used to distinguish between different groups of organisations involved in implementing recommendations:

- ‘DHSC and the devolved health departments’ refers to the health policy functions of the UK Government and the devolved governments (see Glossary of terms and acronyms)

- ‘Relevant healthcare organisations from England and the devolved authorities’ is used where recommendations may affect organisations beyond the health departments themselves, including national health system bodies, sector regulators and professional regulators (see Glossary of terms and acronyms)

The Commission Chairs undertook further engagement with groups across government and the health system, taking a four-nation approach, to test the practical implications of the recommendations prior to publication. This engagement helped ensure that the recommendations took account of different health system arrangements, governance structures and implementation environments across England, Scotland, Wales and Northern Ireland.

The Commission believes it is important to maintain a four-nations approach to UK medical device regulation wherever feasible and appropriate. At the same time, implementation will need to reflect differences in governance, digital maturity, workforce capacity and service delivery across the four nations. This includes the distinct position of Northern Ireland under the Windsor Framework and the need to consider national requirements such as Welsh language standards.

Continued engagement with devolved authorities and wider health system partners will be important as the recommendations are taken forward.

Looking ahead

These recommendations support the creation of a regulatory framework that aims to address the needs of today, while introducing sufficient flexibility to engage with the needs of tomorrow.

The acceleration of AI innovation will bring further potential opportunities for patients and the health system, but also new regulatory challenges that will need to be addressed. Despite the flexibility we have provided here, we anticipate that updates will be needed over time to address new specific challenges and opportunities. We anticipate that the core principles underlying our recommendations will remain true, even if the mechanisms to achieve them may need to be updated to match advances in the technologies being considered.

The Commission has highlighted the extent to which the regulation and assurance of AI in healthcare is something that affects all of us, and that there is an enthusiasm across all groups to work together to shape what that AI-enabled future looks like.

There is a need to ensure that as the Commission finishes, the collaboration and dialogue continue. There needs to be a continuation of the very constructive discussions across the health system, government and wider society, as we seek together to iterate the UK's regulatory framework and wider system assurance, to responsibly engage with the emerging benefits of AI in healthcare.

Members of the National Commission into the Regulation of AI in Healthcare

Chair - Professor Alastair Denniston, Professor of Regulatory Science and Innovation at the University of Birmingham; Honorary Consultant Ophthalmologist at University Hospitals Birmingham NHS Foundation Trust, and Executive Director of the UK Centre of Excellence for Regulatory Science in AI & DigitalHealthTech (CERSI-AI).

Deputy Chair - Professor Henrietta Hughes OBE, Patient Safety Commissioner for England; practising GP and Visiting Professor at the Institute of Medicine, University of Greater Manchester.

Professor Neil Lawrence, DeepMind Professor of Machine Learning at the University of Cambridge; Chief Scientist at Trent AI.

Professor Catherine Sudlow OBE, Head of School of Population Health Sciences, University of Edinburgh; Director of UKRI Adolescent Health Study.

Dr Brian Anderson, CEO of the Coalition for Health AI (CHAI).

Dr Ricardo Baptista Leite, CEO of HealthAI.

Adjunct Professor Raymond Chua, CEO of Health Sciences Authority; Deputy Director-General of Health (Health Regulation) for Ministry of Health, Singapore.

Dame Jennifer Dixon, Chief Executive of the Health Foundation.

Dr Paul Goldsmith, Non-Executive Director of the MHRA; Chair of the MHRA's Regulatory and Safety Committee; visiting Professor at the Institute of Global Health Innovation, Imperial College London; Consultant Neurologist and Clinical Senate.

Professor Vish Ratnasuriya MBE, Practising GP; Chair of Our Health Partnership; co-founder of Primary Care Accelerator; Honorary Professor at the University of Birmingham.

Dr Gabriella Spinelli, Head of Brunel Design School, College of Engineering, Design and Physical Science, Brunel University of London.

Dr Barry Stein, Chief Clinical Innovation Officer and Chief Medical Informatics Officer for Hartford HealthCare.

Richard Stubbs, Chief Executive of Health Innovation Yorkshire & Humber.

Professor Richard Susskind CBE KC, President of the Society for Computers and Law.

Chapter 1. Proportionate lifecycle regulation

Key principle: Ensure regulatory requirements are proportionate to a device's risks and benefits, tailored to the evidentiary needs across the product lifecycle and support equitable access to safe and effective software and AI-enabled technologies.

Chapter summary	This chapter sets out recommendations to support a more proportionate, risk-based and lifecycle-focused approach to regulating software and AI-enabled medical devices, from understanding when a product is a regulated medical device through to pre-market evidence needs, efficient routes to market, post-market surveillance and taking action when something goes wrong.

As software and AI-enabled medical devices continue to develop, it is critical to have a proportionate and tailored regulatory approach to support their safe and effective use. This tailored approach needs to recognise that there are many types of software and AI-enabled devices, such that it should enable a right-sized approach to a given device's intended purpose, risks, and benefits. To face the needs of the rapidly evolving technological landscape, regulation and evidence generation for these devices must become more continuous across the product lifecycle, including for safety and performance.

The current regulatory framework for medical devices was not designed for software and AI-enabled medical devices, particularly devices whose design and performance can change over time. The current framework generally relies on performance and safety data at the pre-deployment stage and gathers insight on adverse outcomes through reactive post-market surveillance. The current framework is not fit-for-purpose for these technologies, as it lacks the necessary balanced lifecycle oversight mechanisms for software and AI-enabled medical devices. Specifically, the current approach is heavily weighted on pre-market assurance, which does not work well for these devices as they may show significant differences in performance between pre-market and real-world deployment. These devices also may experience performance drift over time, including as a result of context-dependent differences between sites and environments of use. Further, the current framework was not designed to practically accommodate the frequent updates or changes these devices may be subject to over their lifecycle. Patients, healthcare professionals[3], industry and

[3] For the purposes of this report, "healthcare professionals" includes clinicians and other health and care professionals, including doctors, nurses, pharmacists, allied healthcare professionals, healthcare scientists, dentists, and other professionals regulated by the relevant statutory regulators across the four nations.

policymakers who took part in the National Commission's research and engagement programme also highlighted that current regulatory frameworks are not well matched to AI[4]. Additionally, the emphasis on lifecycle oversight is consistent with public expectations, with a recent report from The Health Foundation highlighting strong support for continuous monitoring and real-world evaluation of AI-enabled technologies after deployment, rather than relying on initial regulatory approval alone[5].

Moving forward, a more tailored, flexible and proportionate approach to regulation is required. This can not only enhance safety but can foster an ecosystem that enables innovation through ensuring that the regulatory requirements are proportionate, and that they do not impose unnecessary regulatory burden that could hinder patient access to beneficial devices. This approach should enable MHRA to focus its time and resources on devices that benefit from closer scrutiny and oversight. In a lifecycle-based regulatory framework that has sufficient flexibility, levels of regulation for device types may need to be reviewed and updated over time, based on understanding of safety, performance and ongoing controls such as post-market surveillance taking place as a more integrated, real-time activity.

For medium and high-risk medical devices, the current Great Britain regulatory framework relies on conformity assessment by MHRA-designated Approved Bodies[6]. Following a successful assessment against the requirements set out in the UK Medical Device Regulations 2002[7] (SI 2002 No 618, as amended) (UK MDR 2002), an Approved Body issues a certification enabling a manufacturer to affix a UK Conformity Assessed (UKCA) marking and place the device on the market in Great Britain. The MHRA itself does not assess medical devices for conformity with the regulations. During the course of the National Commission process, the MHRA informed the Commission of the government's intention to address certain limitations with this model, including by exploring the ability to directly license medical devices. This reflects MHRA's wider ambition to develop a more flexible and responsive regulatory system that can better support emerging technologies, including AI-enabled medical devices, while maintaining high standards of patient safety and strengthening the long-term resilience of the UK's regulatory framework. This has been considered and taken into account by the National Commission as it developed this report.

The recommendations included in this section are designed to support a regulatory framework that is proportionate and can support manufacturers to manage performance and safety of devices across the entire lifecycle. In particular, it supports a clear and predictable framework for device qualification, device classification, generation of necessary evidence, management of changes, and more continuous understanding of device performance across the product lifecycle.

---

[4] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement programme for the National Commission into the Regulation of AI in healthcare, p.13.
[5] Thornton et al - The Health Foundation (2026) The public's views on the regulation of AI in health care: Findings from deliberative research and public polling and implications for policy, p.37.
[6] Medicines and Healthcare products Regulatory Agency (2023) Approved Bodies for Medical Devices.
[7] Legislation (2026) The Medical Device Regulations 2002.

From pre-market emphasis to proportionate lifecycle assurance

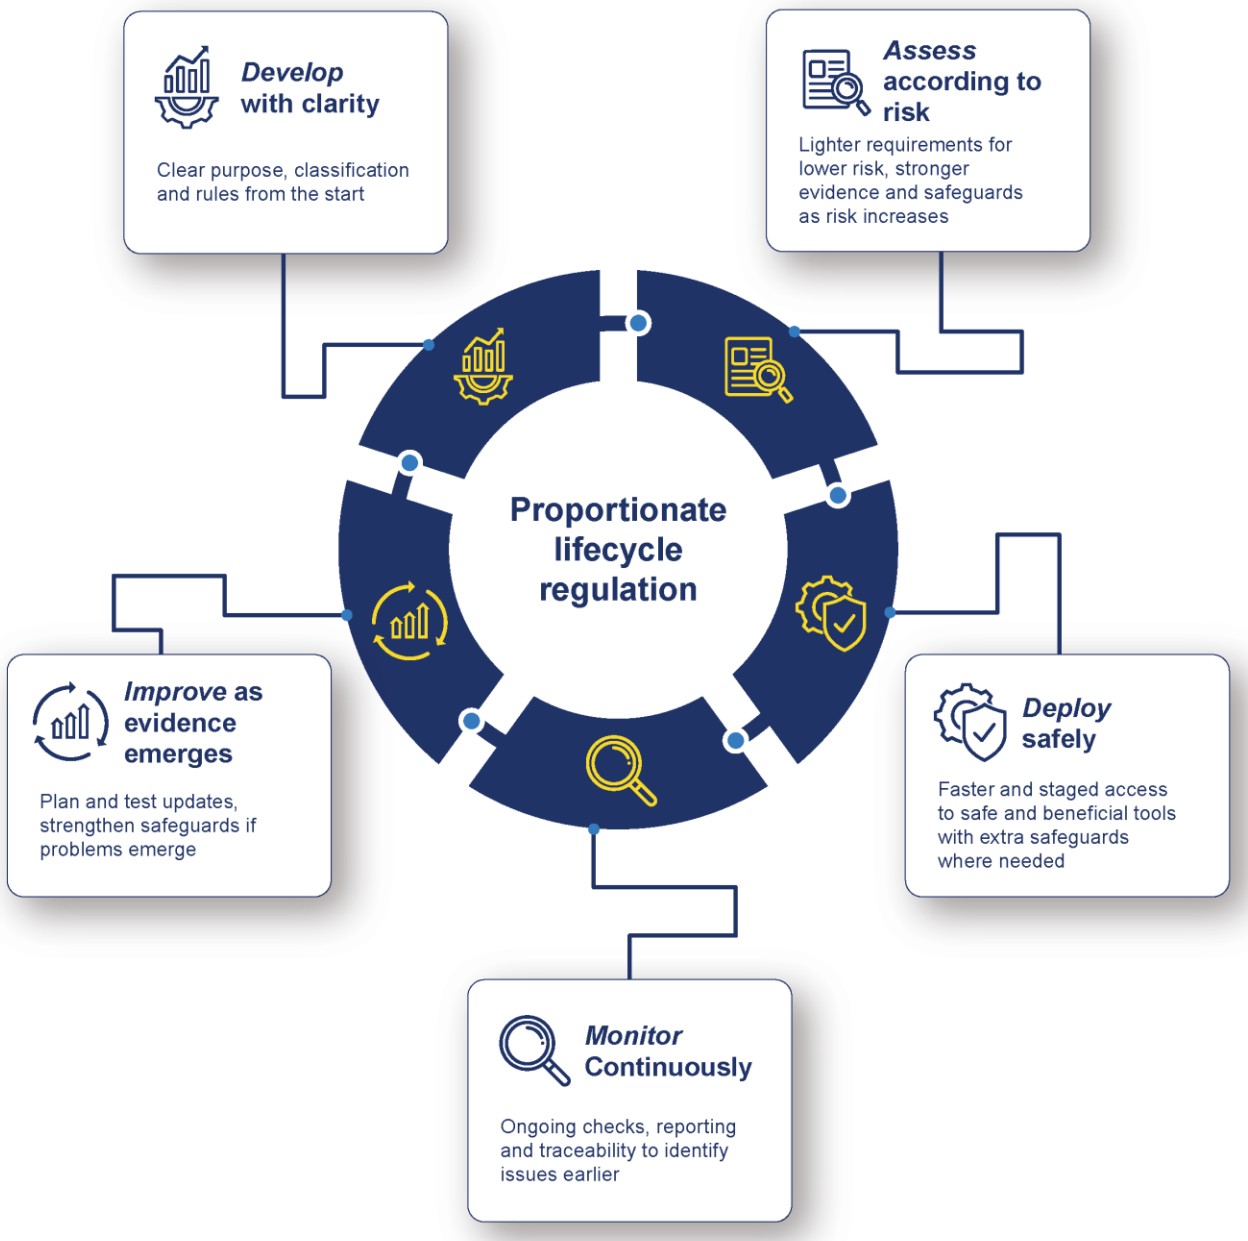

The right safeguards, at the right time, for the right level of risk

1.1 Clear and consistent device qualification and classification

There is a wide spectrum of software and AI-enabled products for use in healthcare, and not all of these products are regulated as medical devices. It is important that manufacturers and the health system have the clarity needed to understand when a product qualifies as a regulated medical device. This clarity is generally provided through the combination of legislation and guidance that outlines the interpretation of the legislation. The existing framework as set out in the UK MDR 2002 outlines that products with an intended purpose that meets the definition of a medical device are regulated medical devices. This foundational concept of device regulation should be maintained but refined moving forward.

For many software and AI-enabled products, decisions on qualification and classification in the current framework can be complex and subjective. It can be challenging to characterise a product's intended purpose while considering potential changes in capability, design and risk across the product lifecycle. For example, manufacturers may position AI-enabled medical devices for a wider range of applications or functions, which results in less specific intended-purpose statements and increased uncertainty around the necessary supporting evidence to validate such broad purposes.

The National Commission recognises that the MHRA has already taken steps to help explore and address this challenge. Through the AI Airlock regulatory sandbox programme8, the MHRA worked with innovators to understand how to define, maintain and enforce the intended purpose of an AI product that can evolve in functionality and clinical influence over time. This work identified the need for updated guidance to address the changing device landscape and considerations across the product lifecycle[9].

As manufacturers continue to develop such devices with different functions and broader uses that may change over the course of product lifecycle, it will therefore be even more important to consider and define intended purpose clearly. Intended purpose should explicitly include device design and functionality in addition to a manufacturer's claims and promotional materials. It is important to understand how and when the specifics of how a device has been designed to function is factored into understanding its intended use, especially if that design is inconsistent with specific claims or promotional materials. Clear guidance and resources are needed to support more consistent and ongoing evaluation and understanding of a device's intended purpose.

[8] Medicines and Healthcare products Regulatory Agency (2026) AI Airlock: the regulatory sandbox for AIaMD - GOV.UK.
[9] Medicines and Healthcare products Regulatory Agency (2026) AI Airlock Sandbox Phase 2 Programme Report.

To achieve a proportionate regulatory framework, several categories of products warrant careful consideration by MHRA as to how and when the UK medical device regulations should apply ('qualification'). Three categories of software and AI-enabled products which deserve specific consideration in this regard are those for: (1) administrative purposes, (2) general wellbeing purposes and (3) low risk clinical decision support purposes. Some of these products, including many of those belonging to the first two categories, are likely not medical devices within the existing definition. Some of these products, such as those that are intended to provide certain clinical decision support, may currently meet the definition of a device but are sufficiently low risk such that regulatory oversight may not add value. The MHRA should consider this issue, including the potential role of legislative change and new guidance to provide clarity on when such products should be subject to UK medical device regulation.

The current classification scheme also carries significant limitations for software and AI-enabled devices, as it was not designed for modern technologies of this kind. Most of these devices today are self-declared as Class I and this can result in limited understanding of the risks and benefits to patients. Further, this may lead to regulatory requirements that are not right-sized for a given device. Some of these products are very low risk and may not even qualify as medical devices, while others warrant more tailored oversight beyond self-registration prior to deployment, such as regulatory authorisation and ongoing surveillance.

Further, approaches to classification can create challenges for manufacturers and the health system in Northern Ireland, who comply with the EU's medical device regulations through the Windsor Framework. For example, if a software and AI-enabled product is not classed as a medical device in Great Britain but is in the EU, the manufacturer may not pursue Conformité Européenne (CE) marking, which could impact on access in Northern Ireland. Any changes to Great Britain's qualification and classification rules should include careful consideration of Northern Ireland impact. Where divergence occurs, the MHRA should consider other possible routes for innovative products, to ensure Northern Ireland can still access emerging technologies.

The following recommendations focus on updating existing UK medical device regulations and guidance documents to support clear, consistent, and proportionate qualification and classification processes for software and AI-enabled medical devices. These changes are recommended to provide the necessary clarity, oversight and predictability to help the ecosystem determine when software and AI products need to be classed as medical devices. Legislative change can take time, so it is important that MHRA consider what guidance and clarity can be provided in the near-term within the current framework in parallel to broader reform efforts.

Recommendation 1. Modernising the regulation

The MHRA should systematically review and update, as necessary, the existing UK Medical Devices Regulations to allow for a more tailored approach to software and AI-enabled medical devices. This should include:

a. updates to the definition of a medical device to provide more clarity on when certain software and AI-enabled products are regulated as medical devices. For example, such updates should explicitly clarify when software intended for administrative purposes in a healthcare setting, general wellbeing purposes, and certain decision support purposes are not medical devices.

b. a classification system and approach that considers the clinical risks, patient benefits, device lifecycle, and international harmonisation to provide a clear and resilient approach to classification of software and AI-enabled devices. This classification system should:

    i. address the known limitations of self-declared Class I devices in the current framework

    ii. appropriately apply requirements and/or oversight for devices where there is sufficient risk to patients

    iii. employ a proportionate approach to low-risk devices.

c. updates to the definition of intended purpose to ensure that device design and functionality are appropriately considered in addition to a manufacturer's claims and promotional materials.

Recommendation 2. Improving regulatory guidance

The MHRA should systematically review and update, as necessary, guidance and other resources to provide clarity and predictability for the ecosystem. The MHRA should develop and implement an ongoing process for reviewing and updating guidance and other resources to ensure supporting policies are updated and available in a timely fashion. This should include:

a. clarification of which types of software and AI-enabled products are, and are not, medical devices

b. clarification of the application of the risk-based classification system to relevant software and AI-enabled devices

c. clarification of how intended purpose is interpreted for software and AI-enabled devices, including practical examples relevant to both understanding device qualification and post-market compliance.

1.2 Tailored and risk-proportionate regulatory approaches

Traditionally, UK medical device regulation has lacked the necessary agility to meet the needs of a rapidly evolving field such as AI. It is well-understood that AI-enabled medical devices are challenging this established approach, and a framework is needed that is more clearly tailored to the benefits and risks of a given device across its lifecycle.

Regulation will also need to become more flexible and responsive to modern, emerging technologies, with new approaches to better manage device changes while assuring safety and effectiveness over time. Regulators will additionally need to consider how this flexibility works for more increasingly capable AI products, even at some future point addressing the challenges of Artificial General Intelligence (AGI).

When new types or risks of software and AI-enabled medical devices emerge, regulators need flexibility to introduce appropriate requirements or controls to manage specific risks. Those requirements should apply to any subsequent devices of that type to ensure fair and consistent regulation. It is also important to have a regulatory framework that accounts for a regulator's learning over time. It may be necessary and appropriate to change (increase, amend or reduce) the requirements for a given device type, so they remain proportionate as the technology, and its use, mature. For example, various regulatory jurisdictions address this for certain low risk software and AI-enabled devices by applying enforcement discretion or not focusing their regulatory oversight on such products.

Although the MHRA can learn from the deployment of these devices and refine its approaches, updating the underlying regulatory requirements across the product lifecycle can be slow. This can make it difficult to adjust controls as evidence develops and as the device becomes better understood.

To support innovation while maintaining appropriate safeguards, the MHRA should consider ways to make its regulatory framework more flexible and adaptable, enabling it to respond effectively to the ongoing development of software and AI-enabled medical devices. This could include clearer mechanisms to amend regulatory requirements and, where justified by evidence, to strengthen or reduce requirements as signals and safety data emerge. These powers should be applicable across all risk classes.

Recommendation 3. Proportionate regulatory oversight

The MHRA should be empowered to improve the flexibility of its application of the regulatory framework, allowing for tailored approaches over time to allocate its resources and oversight to best assure the public health. For example:

a. adding safety-related requirements for devices where new/available evidence shows a need for this beyond which is covered by existing classification requirements to support patient safety.

b. reducing or limiting requirements and/or oversight for devices where the best available evidence shows that the risk is sufficiently low (for example, applying concepts such as enforcement discretion, where appropriate). Where such actions are proposed, it is essential that this be signalled clearly, applied consistently, and aligned to risk proportionality and patient benefit.

As software and AI-enabled medical devices become more complex and capable, they often include multiple functions within a single product. More of these multifunction products are becoming available, and they can have both regulated medical device functions as well as functions that are not regulated as medical devices. For example, an Ambient Voice Technology (AVT) product may include functions such as summarisation and transcription that are not classed as medical devices, alongside functions that provide advanced decision support that may be classed as medical devices[10].

However, it would be disproportionate and inefficient for regulators like the MHRA to oversee non-medical device functions within the product instead of focusing on the specific functions or purposes that are covered by UK medical device regulations. For example, if a manufacturer of a general-purpose large language model introduces a feature intended for a medical purpose, only that feature should be the focus of MHRA's oversight as a medical device. The entire model and all of its other functionalities that do not qualify as a medical device should not be the focus of oversight.

The MHRA should therefore explore options for medical device regulation in the UK that focuses on the safety and performance of specific functions that qualify as medical devices. This would enable a more proportionate, targeted approach and reduce unnecessary compliance with whole-product or system-based regulation. The MHRA should also work

[10] Medicines and Healthcare products Regulatory Agency (2026). Ambient voice technology-enabled products - GOV.UK.

with other sector regulators to ensure any function-based approach is coordinated and does not overlap with existing regimes.

The introduction of a more function-based approach should also contribute to efforts to level the regulatory playing field between large, established manufacturers and small and medium-sized enterprises. Function-based regulation should be flexible to meet the needs of manufacturers of all sizes.

Recommendation 4. Function-based regulation

The MHRA should consider a function-based approach to the regulation of software and AI-enabled medical devices, recognising increasing product complexity and supporting proportionate regulation. It is critical that regulatory oversight be tailored to and focused on the medical device functionality, rather than any non-medical device functionalities present in products where there may be multiple functionalities.

The nature of many AI-enabled medical devices – including their context sensitivity, frequent (and potentially continuous) updating, and performance drift – means that it is becoming increasingly difficult to demonstrate and assure real world performance and safety at the pre-market stage. Increasingly, post-market evidence will become more necessary and indicative of actual safety and performance.

To address this challenge of pre-market performance not being a great predictor of real-world performance, the MHRA should consider how to recalibrate the evidentiary requirements to be more weighted to the post-market phase. This could enable a stronger and more tailored lifecycle-based oversight approach that is more responsive and better suited to these devices. To support this, the MHRA should ensure approaches to rebalancing requirements are not one-size-fits-all and are based on the specific clinical risk profile, technology and intended use of each device. Further, the MHRA should also work with other organisations who collect evidence on medical devices, such as the National Institute for Health and Care Excellence (NICE), to ensure evidence requirements are aligned and do not result in duplication, wherever possible, for manufacturers and healthcare providers.

To demonstrate this proportionate approach to regulation, the MHRA should provide clarity by setting out what types of evidence (for example, clinical investigations, real-world data, pre-clinical testing, bench testing, synthetic data, etc.) and processes (for example, risk management, change management, post-market surveillance, etc.) will be sufficient to

support a device across its lifecycle. This could include the evidence required to demonstrate the safety and performance of an AI device at the pre-market stage, and what data and processes are required to demonstrate that the device remains safe and effective when deployed in different settings. By having additional, ongoing post-market oversight and understanding of device safety and performance, the pre-market evidence needs may be reduced, as deemed appropriate.

Recommendation 5. Lifecycle-based evidence and oversight

The MHRA and relevant partners should consider how to appropriately balance oversight and evidence expectations across the product lifecycle. In particular, this should include clear identification of sufficient pre-market data and post-market surveillance for software and AI-enabled devices. This should be in line with the expectations and best interests of the public, clinical risks, technological characteristics, and the intended purpose of a given device. For example, the MHRA should provide clarity regarding what pre-market evidence is needed to support the safety and effectiveness of an adaptive AI-enabled device while considering the real-world data monitoring and reporting that will be needed to reduce overall uncertainty.

As more adaptive and context-dependent AI-enabled devices are developed, it will become more difficult and impractical to identify and assure changes during the pre-market stage within the current framework. Ensuring AI models continue to work safely and effectively requires ongoing upkeep and timely updating of the models. The current regulatory framework does not enable this best practice easily. Regulation should support and not hinder responsible developers' efforts to maintain a model to support patient safety and benefit.

This current limitation of the framework will hinder the timely deployment of these devices and innovation that can benefit patients, with two-thirds of industry respondents and half of healthcare providers who responded to the National Commission's Call for Evidence viewing the current framework as restricting innovation[11]. To therefore support the ongoing deployment and iteration of these devices across their product lifecycle, it is important that the MHRA provides clear regulatory mechanisms that allow change to be prospectively planned, managed and authorised.

[11] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of Findings from the Call for Evidence, p.31.

Such regulatory mechanisms to help manage change in software and AI-enabled medical devices include Predetermined Change Control Plans (PCCPs)[12]. PCCPs have been developed with international partners as a mechanism to enable manufacturers and regulators to agree on prespecified changes to a medical device following initial authorisation and deployment. Such changes can then be implemented in accordance with the approved PCCP and the manufacturer's quality management system. This may be overseen through appropriate mechanisms, including compliance.

When PCCPs were initially designed, manufacturers were expected to fully describe and define specific, planned changes, test methods, acceptance criteria and associated impact assessments. While this approach has been successful, it does not clearly accommodate more adaptive technologies or other changes that have not been specifically defined.

The MHRA should therefore continue to consider opportunities and mechanisms to better support manufacturers to manage changes to their software and AI-enabled devices. There is an opportunity to build upon previous work from the MHRA and the International Medical Device Regulators Forum (IMDRF)[13], whilst recognising that further policy development is needed to take this forward and operationalise it. This policy development could include expanding the design and implementation of PCCPs to better accommodate device changes while assuring safety and effectiveness. A refreshed and practical approach could include inviting manufacturers to define the boundaries in which an AI-enabled medical device can operate instead of requiring specific changes to be defined prospectively. These boundaries would need to be accompanied by clearly established testing and monitoring methodologies to assure continued safety and performance. PCCP policy and guidance should be developed to address some of the emerging challenges with these devices. This should include clear paths to setting out how a device can be fine-tuned to meet specific site or sub-population specifications and enabling devices to be deployed with a broader intended purpose.

The PCCP policy and guidance should also consider learnings relating to evidence-based use expansion from the MHRA's AI Airlock regulatory sandbox programme. With some models that were tested through the AI Airlock demonstrating consistent diagnostic performance across several use cases, it may be reasonable for a device's permitted use case to be expanded through a PCCP, based on additional targeted validation taking place.

---

[12] Medicines and Healthcare products Regulatory Agency (2023) Predetermined change control plans for machine learning-enabled medical devices: guiding principles.
[13] International Medical Device Regulators' Forum (2026) Essential Principles and Content of Predetermined Change Control Plans.

The AI Airlock programme also demonstrated that the significance of a change cannot be assessed solely on the technical modification itself. The change often depends on factors such as intended purpose, clinical function, user interaction, level of autonomy, deployment environment, impact on the benefit-risk profile, and degree of human oversight (including consideration of automation bias). These findings suggest that future PCCP frameworks should adopt a contextual, risk-based approach to change assessment, supported by clear guidance to promote consistent interpretation and application.

Recommendation 6. Predetermined Change Control Plans (PCCPs)

The MHRA should consider robust approaches to change management, including providing clarity on the scope and implementation of PCCPs for software and AI-enabled devices. This should include ensuring PCCPs provide a clear path and support for:

a. adaptive and continuously learning AI-enabled devices by moving toward an approach that allows for clear boundaries and guardrails to define the allowable scope of change rather than requiring specific, planned changes

b. site or sub-population specific tuning

c. devices intended to operate with broad intended purposes or goal-versus-task defined intended purposes (for example, broad medical domain applicability, regulated AI agents)

d. evidence-driven use expansion (for example, detecting/characterising additional radiology image findings to support clinician diagnosis or validating additional genes as they are characterised for an AI-driven diagnostic) that does not alter the device's intended purpose.

Another important consideration for the new regulatory framework is how it should address the growing number of devices that are being developed by leveraging general-purpose AI models, which include platforms, foundation models, and other types of underlying technologies. These models are trained on broad datasets and are developed to be capable of a wide range of tasks and applications by upstream developers.

While these general-purpose AI models are usually not intended for medical purposes (and are not regulated as medical devices), there is the potential for them to be adapted for such a purpose, either by the original developer or a third party who then seeks authorisation for that product as a medical device. This can create challenges since a downstream device manufacturer will often be limited in the technical information regarding the underlying model, which would likely be needed to support regulatory submissions.

The MHRA should therefore consider introducing an optional 'Master File' approach for general-purpose AI models, which builds upon similar work undertaken by the U.S. Food and Drug Administration (FDA)[14], Health Canada[15] and Japan's Pharmaceutical and Medical Devices Agency[16]. As a voluntary regulatory tool, a 'Master File' would enable upstream developers to confidentially provide information and technical specifications that support medical device manufacturers to seek authorisation for AI-enabled medical devices that are underpinned by a general-purpose AI model.

Although the 'Master File' is an established regulatory concept, its application to AI models is new, and will require development by the MHRA in dialogue with key stakeholders including those who build and oversee general-purpose AI models, the device manufacturers who seek to build off those models, and independent technical, clinical and patient experts. The aim is that this process optimises the sharing of relevant data between those who need access to that data to support safe innovation. Notably: (1) it maximises the sharing of information from AI developer to regulator where this can provide the regulator with the necessary confidence to grant market authorisation (and increase public confidence that the regulator had all necessary information to make that decision); and (2) it minimises the sharing of commercially sensitive data where there would be no additional benefit to the regulatory decision-making and enhancement of patient safety.

With the specifications and capabilities of general-purpose AI models developing at an accelerated rate, it is also important that upstream developers understand what information they need to include in 'Master Files'. This will help ensure medical device manufacturers can provide sufficient assurance that their AI-enabled medical device is safe, effective, and continues to perform as intended. The MHRA should consider engaging with technical, clinical and informatics experts, to help determine what types of information needs to be included in 'Master Files'. The MHRA could also provide guidance, which outlines the general steps that manufacturers should take in developing 'Master Files'.

Recommendation 7. Master files for general-purpose AI models

The MHRA should consider establishing an opt-in 'Master File' approach for general-purpose platforms, foundation models or underlying technologies used by medical device developers. These 'Master Files' can be used to support the assessment of medical devices leveraging such platforms or models. This should include:

a. engagement with technical experts to explore the kinds of proportionate information that would provide long term value within a 'Master File', such as

[14] United States Food and Drug Administration (2023) Device Master Files.
[15] Government of Canada (2026) Guidance on procedures and administrative requirements for master files: Overview.
[16] Pharmaceutical and Medical Devices Agency (2026) Master File System.

benchmarks, model cards, internal guardrails and firewalled mechanisms of providing additional information as available and necessary. The information should aim to support assurance that a given device leveraging the underlying technology will continue to function as intended

b. development of MHRA voluntary guidance setting out the process for a 'Master File' and the types of information a 'Master File' should contain for such underlying technologies to adequately support device developers and regulatory decision making.

An additional consideration is that increasing numbers of AI products built on general-purpose foundation models may rely on a small number of external providers. This challenges resilience and creates a sovereignty risk when those models are owned outside the UK. Just as there is a need to enhance the level of transparency regarding the foundation model themselves (for example, through the 'Master File' to support device applications), there is also a need to enhance transparency regarding the products built from those models, in particular the provenance and nature of the model, dependencies, related risks and their risk controls. This information should be provided at the time of regulatory submission (for AI-enabled medical devices) and at procurement (for all AI-enabled products).

There is also a need to monitor, at a system level, the cumulative risk exposure arising from a large number of medical devices being critically dependent on a relatively small number of foundation models, through the aggregation of this information across the health system.

Recommendation 8. Reporting of product dependencies on underlying general-purpose models

The MHRA, DHSC and the devolved health departments should consider how to obtain the necessary information when a medical product has an underlying dependency on a general-purpose model. Specifically, the manufacturer should be expected to transparently report this dependency, any related risks and mitigations (including continuity plans). This should be an expectation as part of regulatory submissions (for AI-enabled medical devices) and through procurement and contract terms (for AI-enabled products).

As the use of AI-enabled medical devices continues within healthcare, users[17] of devices and members of the public will have questions on how these devices operate and remain safe following deployment. An example will be general-purpose AI models that are fine-tuned and serve as the basis of a medical device, such as a chatbot intended to provide medical advice to users to support diagnosis and treatment. If clear and accessible information is not available, this could reduce user confidence and hinder the safe use of such a device. This view came through clearly, in the National Commission's research and engagement programme, with clear communication highlighted as key for enabling trustworthy deployment[18].

The MHRA should therefore consider mechanisms to enable manufacturers to build transparency and confidence in how AI-enabled medical devices operate and remain safe. This could include building on known mechanisms, such as model cards and good user interface design, so manufacturers can improve the transparency of their device while providing users with up-to-date information about its performance and limitations. Dynamic labelling could also be explored as an approach to improve transparency across the lifecycle of a device. To support this, the MHRA should provide guidance on the types of information that manufacturers should provide to users to improve transparency and support the safe use of such devices.

Mechanisms or tools such as model cards should provide adequate transparency by setting out: where an AI-enabled medical device is being used, what the risks and benefits are, a summary of how the product was developed and tested, ongoing performance monitoring, and the level of autonomy of the device.

Recommendation 9. Transparent information for users and the public

The MHRA should consider issuing guidance to support developers in providing clear, transparent information to users and the public, such as through the use of model cards, careful user-centred interface design and dynamic labelling to support the safe use of software and AI-enabled devices.

[17] For the purposes of this report, "users" refers to individuals who directly interact with, operate, or rely upon an AI-enabled product. Depending on the context and intended use of the technology, this may include healthcare professionals, practitioners working in health and care roles that are not subject to statutory regulation, administrative or operational staff, patients, service users, or consumers relying on devices at home.
[18] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement programme of the National Commission into the Regulation of AI in Healthcare, p.14.

As healthcare providers continue to become more digitally mature, they will increasingly utilise a range of interconnected data pipelines, deployment platforms and digital health systems. This interconnectivity and integration of these devices into healthcare pathways can increase the risks of cyber-attacks and subsequent patient harm. Indeed, senior healthcare leaders in England have warned that cyber-attacks now pose a bigger threat to the NHS than another pandemic[19]. Examples of cyber-attacks include data poisoning, in which training data is deliberately corrupted to compromise an AI-enabled medical device's outputs, and the infiltration of malware, which can damage, disrupt and lock wider software infrastructure and networks. Such cyber risks can lead to patient harm due to errors in patient diagnosis, treatment and clinical decision making.

Cybersecurity is an inherent part of medical device safety and performance for increasingly connected devices. To support the resilience and continuity of healthcare systems, it is important that software and AI-enabled medical devices are developed with safeguards to withstand cyber threats and to continue to fulfil their intended purpose. The MHRA should therefore look to provide guidance for manufacturers on specific cybersecurity expectations for software and AI-enabled medical devices. This guidance should support manufacturers to develop devices with safeguards embedded within its operating model from the design phase and with controls that enable the device to function as intended.

Recommendation 10. Cybersecurity guidance and resources

The MHRA should consider providing clear guidance and educational resources on cybersecurity expectations across the product lifecycle for software and AI-enabled devices, which are increasingly connected across networks and with other devices. Cybersecurity is a critical factor of device safety as cybersecurity threats can directly compromise the safety or effectiveness of devices, including the potential to cause rapid harm to many people.

Increasingly, there is a shift in where and how healthcare is accessed and delivered. With many software and AI-enabled medical devices now being provided as products which people can access directly, without involvement of a healthcare practitioner, these products (known as direct-to-consumer or DTC) can empower individuals to better manage their healthcare needs. Emerging direct-to-consumer products have the potential to support prevention and earlier disease detection and to provide healthcare access to people who may otherwise not receive it. As the use of such devices increases, it will be important that

[19] Clover (2026) Mackey: Cyber-attack risk 'dramatically accelerating'.

their use is supported by clear regulatory pathways and expectations, which are proportionate to the benefits and risks of these products.

To support the safety of consumers, direct-to-consumer products that qualify as medical devices are regulated as medical devices. This should continue to be the case under any new framework. Although the underlying principles of the regulatory framework remain the same, the direct nature of the direct-to-consumer relationship between manufacturer and consumer, and the lack of a human professional as an intermediary may have implications for manufacturer and regulatory practice.

For example, there are additional expectations regarding the need for the manufacturer to design the device, its user interface and all supporting materials in a way that is suitable for a general user, and without being able to rely on a health professional or health provider for any risk controls. Conversely, the direct relationship to the consumer (and their device) may allow greater control of information exchange, supporting near-real-time return of performance and safety data, rapid update cycles, user feedback and a high level of agency for the individual.

Through guidance, the MHRA should set out clear, proportionate regulatory expectations and considerations for direct-to-consumer medical devices. This could include providing clarity on when a direct-to-consumer product is classed as a medical device, evidence expectations across the device lifecycle, and how manufacturers can design such devices with the end user in mind.

Recommendation 11. Direct-to-consumer devices

The MHRA should consider providing clarity through guidance on regulatory expectations for direct-to-consumer applications and wearables to fully harness their potential for the UK such as by supporting a shift towards preventative care. It is important that careful consideration be given to how to apply a proportionate approach to empower people with access to information about their health while understanding the unique benefits and risks of direct-to-consumer use.

The performance and safety of many software and AI-enabled devices are dependent on how the user interacts with them. This is not unique to AI, and indeed the role of human factors in medical device design and testing, and the need to consider this within its regulation is well-recognised. However, the nature of AI – and the breadth of functions that newer devices may be approved to undertake – does underscore both the importance and the difficulty in doing this well.

With this in mind, manufacturers need to ensure they have appropriately considered user-centred design and human factors across the full breadth of users and settings in which their device is intended for use, and that their design and risk controls incorporate these human factor considerations. Similarly, the regulatory process needs to consider the extent to which evidence is provided to address this, with a view to enhancing safe use.

The MHRA should therefore provide guidance to manufacturers on how human factors and usability need to be considered as part of the design, development and maintenance processes for software and AI-enabled medical devices. This should include proportionate expectations around sufficient testing, best practices, and user interface and labelling considerations to support safe use.

Recommendation 12. Human factors and usability expectations

The MHRA should consider providing clarity through guidance on human factors, usability considerations and expectations to support safe use as part of the design, development and testing of software and AI-enabled medical devices.

Devices need to be safe and effective for their intended population. This means the future regulatory framework for software and AI-enabled medical devices needs to clearly address health equity as a critical aspect of safety and performance. It is important these devices are developed, evaluated, deployed, and monitored with appropriate consideration of the diversity of the population, which may include characteristics such as age, gender, sex, race, ethnicity, and geographical location[20]. These devices may interact with, and inform the care outcomes of, population groups with certain characteristics or vulnerabilities in different ways.

Concerns around potential exclusion and inequity arising from the use of AI in healthcare were a recurrent theme highlighted through the National Commission's research and engagement programme. Participants raised concerns about various aspects of bias, including selective underperformance of the algorithm and digital exclusion. Participants emphasised that equity must be considered not just at the outset, but throughout the use of the device (i.e. across the lifecycle). This issue was also prioritised for in-depth targeted engagement with the public, including with underserved groups to explore their views on the regulation of AI in healthcare, including specific trade-offs, and how to handle different scenarios in which unequal performance might be observed between groups. Participants

[20] International Medical Device Regulators Forum (2025) Good machine learning practice for medical device development: Guiding principles.

stressed that equity and inclusion are fundamental for safety: if AI-enabled medical devices do not work well for all populations, trust in AI-enabled care will be weakened and existing inequalities may be exacerbated[21]. This was further reflected in recent deliberative public research, which found that participants had a clear red line that AI should never lead to a worse outcome for any population group, and that additional safeguards should be implemented where disparities in performance are identified[22].

To help address these concerns, the MHRA should work with partners to provide guidance for manufacturers on how to demonstrate that their software and AI-enabled medical devices are safe and deliver equitable performance and care outcomes for diverse population groups. This guidance should support manufacturers with specific considerations and regulatory expectations across the product lifecycle, including ongoing monitoring of subgroups, to advance health equity.

Recommendation 13. Recognising the diversity of the intended population

The MHRA should consider working with relevant partners to provide guidance for manufacturers on health equity expectations for software and AI-enabled medical devices across the product lifecycle, including consideration of the diversity of the intended population, with particular regard to underserved population groups.

[21] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement programme for National Commission into the Regulation of AI in Healthcare, p.39.
[22] Thornton et al - The Health Foundation (2026) The public's views on the regulation of AI in health care: Findings from deliberative research and public polling and implications for policy, p.30.

1.3 Clear and timely paths to market, with sufficient risk control

A future regulatory framework for AI should be not only effective but also efficient. It must protect citizens and the healthcare system from products that are unsafe and ineffective, but it should also accelerate access to products that are safe and effective. The National Commission's Call for Evidence highlighted concerns that the current regulatory framework is not fit for AI, with two-thirds of industry and half of health providers stating that it is 'somewhat restrictive' or 'too restrictive', with a negative impact on innovation[23]. Getting this right is important for patients, for those who work in healthcare and for the health tech sector.

Moving forward, authorisation routes and deployment pathways will need to better accommodate the increasingly adaptive and iterative nature of the software and AI-enabled medical devices they are supporting. These pathways need to be designed to enable manufacturers to generate evidence of performance, safety and outcomes in real world settings across the product lifecycle. It is also important that these pathways are supported with robust risk controls that are tailored to the unique characteristics of software and AI-enabled devices.

AI tools are becoming increasingly sophisticated, iterative and context-dependent, as exemplified by recent advances in adaptive technologies, generative AI and AI agents. This highlights both how the current regulatory system is increasingly unable to deal with these technologies, and that any future regulatory framework will need in-built flexibility to address these challenges. To better support manufacturers, the process should be reimagined in a way that carefully considers device benefits, risks, sufficient evidence needs for a given technology and use prior to deployment, and necessary controls/mitigations to support ongoing safety and effectiveness in the intended use environment.

To better support manufacturers of software and AI-enabled medical devices with an emerging use case, the MHRA should enable appropriate suppliers to deploy their device through a staged authorisation pathway. Such a pathway would allow an AI-enabled medical device to be deployed within a tightly controlled scope, based on MHRA review of initial evidence to support device safety and agreed upon risk controls and reporting. The authorisation can be expanded once the device meets certain prespecified evidence thresholds for performance and safety. Broad support for phased or conditional authorisation

[23] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement programme for National Commission into the Regulation of AI in Healthcare, p.18.

was observed during the National Commission's research and engagement programme, including from industry participants[24].

A staged authorisation pathway should be supported with appropriate deployment routes, such as partnerships with healthcare providers who can demonstrate sufficient AI readiness. This means having the necessary knowledge, capabilities and capacity to deploy the relevant device safely, and to implement any required risk controls[25]. During this period, additional risk controls and evidence gathering are likely to be necessary conditions specified as part of the staged authorisation. Additional risk controls will support safety, while additional evidence gathering will support full regulatory authorisation. Both will require effective cooperation between the manufacturer and healthcare provider.

To support successful implementation, the MHRA should ensure that staged authorisation pathways demonstrate several characteristics:

- They should be flexible and tailored to the benefits and risks of the medical device seeking staged authorisation. This process should include consideration of technological maturity, the potential clinical benefit, and the extent to which the available data and proposed risk controls would ensure patient safety.

- The MHRA should ensure staged pathways are temporary only, and that there is a clear and efficient path to full authorisation that appropriately leverages robust, real-world evidence.

- The deployment of a device through a staged authorisation pathway should also be clearly defined and communicated in an accessible way to patients and health system partners, with regular updates provided to patients and professionals in settings where these devices are being deployed. This will help build public and professionals' awareness of, and confidence in, the use and outcomes provided by these devices.

- Staged deployment pathways should address the limitations of current approaches, which are heavily weighted to the single, point-in-time pre-market authorisation. Staged deployment can help manufacturers better demonstrate real-world performance over time, ensure uncertainties in performance can be proactively managed, and enable oversight and assurance activities to be balanced across the lifecycle of the medical device through greater emphasis on post-market monitoring and surveillance.

[24] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement programme for the National Commission into the Regulation of AI in Healthcare, p.36.
[25] Medicines and Healthcare products Regulatory Agency (2026) Summary minutes: Sixth Meeting of the National Commission into the Regulation of AI in Healthcare (20th March 2026).

The National Commission additionally welcomes the efforts of the new National Healthtech Access Programme (NHAP) that was launched earlier in 2026[26]. This can accelerate access to high-impact technologies, including a limited number of promising AI-enabled medical devices, with the first wave including devices for prostate and breast cancer detection on histopathology slides.

When considering opportunities for staged authorisation approaches, the MHRA should work with NICE and other NHAP partners to ensure coherence and clarity to external stakeholders.

Recommendation 14. Staged authorisations

The MHRA should enable staged authorisations, supported by appropriate deployment pathways that support evidence generation and infrastructure to enable sustainable, long-term patient access to beneficial devices. These authorisation pathways should be:

a. sufficiently flexible to allow tailoring of the authorisations and pathways to support those devices, where a staged authorisation is appropriate. This should include consideration of technical maturity, clinical needs and necessary risk controls to be put in place to ensure patient safety

b. transparently defined and reported to ensure patient awareness and support public confidence

c. designed to address the recognised limitations of single-timepoint assessment of software and AI-enabled devices, including demonstrating real world performance, managing uncertainty, and providing lifecycle assurance of device safety and benefit

d. clearly defined processes and pathways from staged to full market authorisation to enable predictable, uninterrupted delivery of beneficial devices to patients.

Regulatory sandboxes are an established mechanism for manufacturers and regulators to work together to test emerging technologies in controlled environments and to refine regulatory approaches. Examples include the MHRA's AI Airlock[27]. Through its pilot with four companies, the AI Airlock identified 40 recommendations to enhance patient safety and support innovation in medical devices[28]. The Financial Conduct Authority (FCA) also

[26] National Institute for Health and Care Excellence (2026) Faster, fairer access to Healthtech under new national programme.
[27] Medicines and Healthcare products Regulatory Agency (2024) AI Airlock: the regulatory sandbox for AIaMD.
[28] Medicines and Healthcare products Regulatory Agency (2025) AI Airlock Sandbox Programme Report, p.72.

operates a regulatory sandbox, which has worked with over 700 businesses and accelerated time to market by 40% on average[29].

As increasingly capable, autonomous and adaptive AI-enabled medical devices emerge, it will be important to develop and evaluate new approaches to their regulation and safe deployment. Proposals include the development of approaches that better reflect the current performance of the device such as real time monitoring and more continuous or staged regulatory authorisation.

Sandboxes can play a key role here. Although sandboxes have traditionally been developed to simulate real-world environments, they are also now being established in real-world clinical settings, where more representative, robust evidence of performance, safety, and impact can be collected. This includes the development of the London Region I sandbox[30]. As a collaboration between the MHRA, NHS England and London Health Innovation Networks, this sandbox will support up to 10 AI-enabled medical devices to develop additional evidence of benefit in controlled real-world settings.

While the launch of the London and Manchester regional regulatory sandboxes is a very welcome step in supporting the advancement of promising AI-enabled medical devices, there is more the MHRA can do. Moving forward, the MHRA should continue to be proactive and forward-looking in how it leverages regulatory sandboxes to support safe and effective software and AI-enabled devices as the landscape continues to rapidly evolve. For example, MHRA should ensure that regulatory sandboxes have a clear focus on developing corresponding innovative regulatory mechanisms and accelerated routes to market with robust evidence generation to demonstrate safety and performance. This work should also be supported by ongoing evaluation, which helps build evidence of how these approaches facilitate innovation and support better patient outcomes. This will help ensure that promising devices can be brought to market more quickly and are overseen by responsive and proportionate regulation once they have been deployed.

The MHRA should also consider opportunities to work with regional integrated care systems[31] or devolved authorities to develop sandboxes which are configured to local populations, geographies or provider maturity. This could include working with the Northern Ireland Executive to develop a sandbox which utilises the country's size, scalability and integrated systems, such as its single healthcare record. This type of sandbox would give

[29] Cambridge University Press (2021) A Sandbox Approach to Regulating High-Risk Artificial Intelligence Applications.
[30] Medicines and Healthcare products Regulatory Agency (2026) Pioneering AI health innovations regulatory sandbox launched.
[31] What is integrated care? (accessed 2026) NHS England » What is integrated care?

manufacturers access to the data needed to test and refine their device, which can then be deployed at greater scale.

Recommendation 15. Regulatory sandboxes

The MHRA, in collaboration with system partners, should continue using regulatory sandboxes and extend their remit to support the development and testing of regulatory approaches, best practice and safe deployment pathways for novel technologies that provide a clear path to market authorisation.

The MHRA is one of several medical device regulators developing approaches to ensure that regimes governing AI-enabled medical devices are proportionate, enable innovation and focus on patient safety. The MHRA has been working closely with like-minded agencies around the world, including those in the US, Canada, the EU, Singapore, Australia, Switzerland and India, among others.

To support the development of the UK market, and to facilitate patient access to novel AI-enabled medical devices developed and authorised in other jurisdictions, the MHRA should build upon the findings from its recent routes to market and in vitro diagnostic devices consultation. With 59% of consultation respondents agreeing that software-based medical devices should be added to international reliance routes[32], the MHRA should develop and introduce recognition pathways for software and AI-enabled medical devices with selected international regulators where it is in the best interest of patients. These pathways should include review points and break clauses, so they remain responsive to emerging capabilities and risks, while supporting investment and market growth and ensuring patient safety.

Further, the MHRA should continue to build on its efforts and success as a strong collaborator and advocate for international harmonisation. As the new regulatory framework is developed, it is important that the MHRA works to share lessons learned with other trusted regulators while seeking to gain consensus and harmonisation on best practice and regulatory approaches for these devices. This harmonisation directly benefits public health by providing clear paths to delivering beneficial devices to patients across the globe.

[32] Medicines and Healthcare products Regulatory Agency (2025) Response to the international reliance, UKCA marking and in vitro diagnostic devices consultation proposal, p.13.

Recommendation 16. Recognition and reliance pathways and international harmonisation

The MHRA should continue its efforts to advance international harmonisation of the regulation of software and AI-enabled devices by:

- sharing lessons learned and best practices for proportionate, lifecycle-based regulation to align where possible and support responsible innovation for patients
- introducing recognition or reliance pathways for relevant devices with trusted international regulators, where it is in the best interests of patients.

1.4 Proportionate post-market assurance

Traditionally, post-market assurance for medical devices focused primarily on investigating incidents and adverse event reports once a device was already in use. In June 2025, the introduction of the GB Post Market Surveillance Regulations[33] significantly strengthened this approach by introducing more systematic and proactive post-market surveillance requirements. Manufacturers are now required to collect and analyse real-world safety and performance data, identify emerging risks and trends, take corrective action where necessary, and report relevant information to the MHRA. These reforms represent an important step forward in strengthening post-market assurance and improving the evidence available to support regulatory decision-making.

Many of the principles underpinning effective post-market assurance remain the same for software and AI-enabled medical devices. However, as increasingly iterative and context-dependent devices are deployed some aspects of assurance will need to evolve further. This increased emphasis on effective post-market assurance is central to the new framework, supporting a more balanced, lifecycle-based approach that enables earlier access to beneficial devices.

The National Commission's Call for Evidence found that 65% of respondents disagreed or strongly disagreed that current approaches to post-market surveillance for medical devices (including AI-enabled medical devices) are sufficient[34]. Evidence gathered through the Commission's research and engagement programme also highlighted support for more dynamic approaches to assurance, including greater use of real-world evidence, more continuous monitoring and improved information sharing across the healthcare ecosystem[35].

While building on the foundations established by the 2025 reforms, the recommendations that follow focus on how post-market assurance can be further strengthened while remaining proportionate and responsive for AI-enabled medical devices. In particular, they seek to support a more continuous understanding of real-world performance, and better sharing and use of information across the ecosystem throughout the product lifecycle.

The increasing breadth of AI-enabled medical devices ranging in function, complexity, autonomy and adaptive nature, will generate different types of evidence of performance, present different risk profiles and will behave differently depending on the environment in which they are used. A one-size-fits-all approach to post-market surveillance is therefore unlikely to be sufficient.

---

[33] Medicines and Healthcare products Regulatory Agency (2025) Medical Devices: Post-market Surveillance Requirements.
[34] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of Findings from the Call for Evidence, p.30.
[35] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of Findings from the Call for Evidence, p.34.

Stakeholders noted this in the National Commission's Call for Evidence, calling for more responsive and effective mechanisms to track performance, monitor safety and support compliance across the device lifecycle[36].

The MHRA should therefore work to move towards a more tailored and context-based approach to post-market surveillance, with requirements based on the risks, benefits and uncertainties associated with the device. This will ensure regulatory requirements at the post-market surveillance stage are proportionate and sufficient and do not introduce unnecessary resource constraints or duplicative burden for manufacturers.

The MHRA should consider a toolkit of regulatory approaches, such as:

- **Post-market surveillance plans assessed at the point of regulatory authorisation**: The MHRA should consider the evaluation of a robust post-market surveillance plan within a manufacturer's pre-market application as an important way to reduce overall uncertainty for a given device at the time of authorisation. Specifically, by providing planned risk controls and ongoing monitoring commitments, more complex and adaptive technologies may be able to be approved for earlier market access for example as part of staged authorisation pathways. Such monitoring plans can help supplement available evidence at the time of pre-market assessment. This is particularly necessary for technologies that cannot practically be tested comprehensively pre-market (such as those with broad intended purposes or adaptive underlying technology). Expectations for post-market surveillance plans should be proportionate, taking account of the device, its intended purpose, risks and benefits.
- **Real-world post-market studies (as needed):** In addition to post-market surveillance, prospective post-market studies may be considered to address areas of residual uncertainty, for example, to gather additional evidence of performance and safety in specific patient groups, in specific settings or under certain conditions where there is perceived to be an increased risk of device failure ('failure modes'). The results of any such required studies should be made available to device users through updates to labelling and other mechanisms to support safe use.
- **Regular reporting on the performance of AI-enabled medical devices:** Although manufacturers already provide information to the MHRA when adverse incidents or outcomes occur, they are not routinely required to provide regular reports on overall device performance following deployment. More regular, proactive and ongoing performance reporting could provide additional assurance to regulators, healthcare providers, professionals and the public that AI-enabled medical devices remain safe and perform as intended in real-world settings. The MHRA could therefore consider incorporating proportionate performance reporting requirements as a complementary

[36] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of findings from the Call for Evidence, p.33.

regulatory mechanism to, or embedded within, PCCPs for appropriate AI-enabled medical devices. This would allow manufacturers to rapidly iterate and update their medical devices without additional regulatory submissions, while providing an appropriate level of visibility on the continued performance in the real world. This approach would help mitigate concerns around passive reporting that only happens after problems have occurred.

- **Processes for escalating performance degradation**: As AI-enabled medical devices evolve throughout their lifecycle, performance may improve or deteriorate over time. The MHRA should consider processes to define how and when manufacturers should report relevant or material deterioration or unexpected changes in performance, helping to identify emerging risks and build a better understanding of device performance across different use cases and stages of deployment. This could include ensuring there are prespecified performance thresholds that may indicate a meaningful change in performance.

As the use of AI-enabled medical devices increases in healthcare, the MHRA may also experience a corresponding increase in the volume of data, evidence and signals it receives through post-market surveillance. This provides an opportunity for the MHRA to upgrade its approaches and infrastructure to support a more responsive and real-time approach to post-market surveillance. This could include using AI-enabled post-market surveillance tools to automate or enhance data collection and signal detection. This would enable the MHRA to identify issues and trends with devices more proactively and take timely action to maintain patient safety.

Recommendation 17. Tailored post-market surveillance

The MHRA should continue to tailor post-market surveillance requirements to be proportionate with a device's risks, benefits and related uncertainty to ensure effective, resource-efficient oversight. This should include a menu or set of mechanisms that manufacturers may be required to implement, including:

a. post-market surveillance plans and other real-world data collection/monitoring plans assessed at the time of device authorisation to reduce uncertainty

b. real-world post-market studies where needed

c. more regular, ongoing performance reporting made available publicly and/or to regulators, as appropriate, to provide continued assurance of device safety and effectiveness

d. established processes for escalation where performance degradation is detected, even if a reportable incident has not occurred.

Reporting of device experience, including failures and adverse incidents, is a critical part of a well-functioning framework. Different stakeholders within the healthcare ecosystem have varying insights on, and levels of reporting of, the performance, safety and outputs of AI-enabled medical devices. This information is not easily or routinely shared with others across the healthcare ecosystem. Reporting often occurs in silos, with respondents to the National Commission's Call for Evidence commenting that fragmented data infrastructure makes it harder to use AI-enabled medical devices and that strengthened frameworks for data sharing are needed[37]. All of which makes it more difficult to take a more ongoing, lifecycle view of how device performance and safety changes over time.

To help realise a more joined-up and collaborative approach, the MHRA should work with partners across the health ecosystem, including manufacturers, healthcare providers, regulators and patient groups to review and refine existing mechanisms for reporting and sharing information. With respondents to the National Commission's Call for Evidence calling for standardised reporting mechanisms for AI-enabled medical devices[38], this work should focus on ensuring that the information shared across the ecosystem is clear, accessible and supports robust oversight of device safety and performance across the lifecycle.

Further, as healthcare increasingly extends beyond hospitals and other formal clinical settings, effective oversight and assurance cannot rely solely on information generated within formal clinical settings. Clinically relevant data and AI-driven decisions will also arise from home and community-based technologies, such as wearables, remote monitoring tools and consumer-facing apps, and may later need to be integrated into clinical workflows. Safety signals may therefore emerge from a wider range of sources than conventional clinical reporting alone, including patient complaints, digital reporting systems, consumer feedback platforms and online communities, as well as local assurance activity within healthcare providers. If these signals remain fragmented, important patterns of harm or unintended consequences may be missed.

This requires a post-market monitoring approach that draws together multiple data streams across settings. This can be achieved by combining structured clinical data and manufacturer reporting with patient-reported experiences and broader public signals, to support earlier detection of harm, learning from real-world use, and continued improvement over time. Healthcare professionals and patients using these devices should also be supported with clear messaging to critically assess outputs from decision-support tools and to understand what action to take if they are uncertain, to reduce overreliance on these tools.

[37] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of findings from the Call for Evidence, p.16.
[38] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of findings from the call for evidence, p.17.

There are several steps which the MHRA should consider for reviewing and refining these mechanisms. As discussed for manufacturers, this could include embedding monitoring and reporting requirements into PCCPs for a given device. This could also include requiring more regular reporting on ongoing device performance, including any signals of changes to performance. Manufacturers should also consider designing mechanisms to allow product users to seamlessly share feedback during use, undertaken in line with user-centred design principles. Software and AI-enabled medical devices should also include appropriate logging to detect performance drift over time. Manufacturers should monitor and report to the MHRA any changes that could cause harm, even where no harm has occurred. Embedded reporting routes, such as a Yellow Card button designed into a user interface, could make this easier.

For regulators, this could include enhanced partnerships and data sharing agreements to enable information to be shared in a more end-to-end way. For providers and professionals, the focus could be on how they are supported to work closely with manufacturers, so they can easily share information from deployment sites with manufacturers in a more real-time manner. For the public, there could be a focus on specific improvements to the Yellow Card scheme[39]. This could include building on findings from the National Commission's Call for Evidence through increased awareness of the system[40], improving specific reporting mechanisms for software and AI-enabled medical devices and strengthening feedback. Whenever possible, this information on device performance should be transparent and made available to the public to increase trust and confidence across the ecosystem.

To strengthen public confidence in AI-enabled medical devices, it is also important that patients and the public can better understand what adverse incidents look like, in practice, and what they mean for the delivery of healthcare. The MHRA should therefore consider how to make adverse incident reporting data more accessible to the public with an appropriate level of detail, improving transparency around safety and performance issues involving software and AI-enabled medical devices. This should include access to information on reports submitted by manufacturers, as required, in addition to reports submitted voluntarily by others across the system, such as professionals and patients. For example, the MHRA could build on approaches such as its interactive Drug Analysis Profiles (iDAPs)[41] and introduce a similar tool for devices. The MHRA should also consider the best practices of other international partners, such as the FDA's MAUDE Database[42], which

[39] Medicines and Healthcare products Regulatory Agency (2026) Yellow Card – Making medicines and medical devices safer.
[40] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of findings from the Call for Evidence, p.34.
[41] Check drug analysis profiles (iDAPs) (2026) Check drug analysis profiles (iDAPs) - GOV.UK.
[42] United States Food and Drug Administration (2026) Manufacturer and User Facility Device Experience (MAUDE) Database.

enables members of the public to search by fields such as manufacturer, device name and specified timeframe, among others.

Recommendation 18. Reporting and information sharing across the lifecycle

The MHRA, manufacturers and relevant healthcare organisations from the devolved authorities should work with partners across the healthcare ecosystem to review and refine mechanisms for reporting and sharing information (including best practice) across the ecosystem to ensure that they are clear, accessible, and where appropriate, standardised to enable more robust insight into ongoing performance and safety. This should support timely sharing of information from deployment sites and greater awareness of available reporting routes.

For example, the MHRA should consider how the Yellow Card scheme can be improved for AI and software-enabled devices, including for different users (such as patients) who may need to interact with the system. Likewise, manufacturers could support this process by focusing on user-centred design that enables real-time feedback, including the reporting of adverse incidents and malfunctions.

Recommendation 19. Sharing information on adverse incidents with the public

The MHRA should consider developing a database or tool that enables members of the public to easily search for adverse incidents related to a specific software or AI-enabled medical device. This should build on initiatives such as MHRA's interactive Drug Analysis Profiles (iDAPs) and include functionality that enables members of the public to search and filter by manufacturer, device name and timeframe, among other variables.

1.5 Infrastructure, monitoring and enforcement

As the adoption of software and AI-enabled medical devices continues, these devices, based on their application, will become more embedded into healthcare pathways and workflows. Over time, it is anticipated that these devices will operate with higher levels of independence and may have a greater impact on patient outcomes.

When there are higher levels of integration, it can be challenging to clearly understand the role a device has played in a care outcome, which can make reporting more challenging and potentially burdensome. As a result, this also leads to challenges in holding manufacturers to account for outcomes resulting from use of their device. It is also recognised that these issues apply in other sectors, such as finance and law. Lessons from these sectors have been considered throughout the Commission's work and should continue to be considered as case law, regulation, policy and best practice continues to evolve. To therefore maximise the positive impact of AI-enabled medical devices and to maintain strong performance and safety, there also needs to be a focus on the infrastructure, processes and activities which support the use of a device.

As more advanced software and AI-enabled medical devices emerge, it is important that the framework enables clear traceability and monitoring. With devices able to interact with, and inform, decisions at different points of the healthcare pathway, it may not always be clear when a device has been involved in a patient's care pathway and the subsequent impact it has had. This is particularly important for when adverse care outcomes occur, with respondents to the National Commission's Call for Evidence calling for stronger approaches to tracking performance across the lifecycle of an AI-enabled medical device[43].

The MHRA should therefore explore how to make the identification of AI-enabled medical devices and the associated patient outcomes easier to track once they have been deployed. This could include building upon proposals set out in the Pre-Market Statutory Instrument[44] to require all software and AI-enabled medical devices that operate in the UK to be issued with a unique device identifier (UDI), which follows the device through its lifecycle. This would support providers to better understand how a device informed a patient's care outcomes. To effectively enable the use of UDIs, the MHRA should also consider interoperability and the data maturity of different providers, as these factors can affect the extent to which identifiers can be used at scale.

The MHRA could also consider how to adequately incorporate version control into such mechanisms. This would allow an identifier to be potentially updated where, for example, a device's capability, intended purpose or risk classification changes. UDI or other traceability

[43] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of findings from the Call for Evidence, p.32.
[44] Medicines and Healthcare products Regulatory Agency (2026) Standard – Implementation of the future regulations.

information about a device would be most useful to the system, including Yellow Card reporting, if captured in a structured field in the patient health record.

Recommendation 20. Traceability and monitoring

The MHRA should introduce mechanisms to identify AI-enabled medical devices once deployed, including appropriate version control, to provide clear traceability throughout a device's lifecycle. For example, a unique device identifier (UDI) should be explored as a potential mechanism for traceability within the patient record, enabling healthcare professionals and providers to identify and audit the outcomes of patients whose care involved a particular device. The MHRA should also work closely with DHSC and the devolved health departments to support implementation, ensuring that device traceability information, including version, is routinely recorded as part of the patient record.

Digital and AI approaches may also enable automation of some aspects of the reporting process. This opportunity was also highlighted during an industry roundtable, which the National Commission held with techUK, with attendees calling for more system level telemetry or automation of reporting[45].

The MHRA should therefore explore opportunities to support manufacturers to automate lower-burden reporting of events that they are required to submit to the MHRA and other health system stakeholders as part of regulatory requirements. This could include exploring opportunities to integrate automated processes into existing reporting systems, such as the Yellow Card scheme, and clinical systems, including electronic patient records[46], where appropriate. To support this, the MHRA should set out the reporting requirements that can be automated and provide assurance that these are considered sufficiently low risk and low burden. These reporting requirements should also be kept under review as regulations evolve and are updated.

Any reporting automation will need to be accompanied by advanced tools for MHRA assessors to effectively manage a likely increase in reports. Such tools would support the transformation of regulatory decision-making capabilities by improving efficiency and enabling assessors to focus their time and complex judgement on higher-value aspects of the process. Such automation and tooling would benefit the wider ecosystem by enabling

[45] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of findings from the Call for Evidence, p.34.
[46] Purpose of the GP electronic health record (2025) NHS England » Purpose of the GP electronic health record.

more efficient and timely signal detection and action, helping to maintain public safety and trust in software and AI-enabled medical devices.

Recommendation 21. Automating low-burden reporting

The MHRA should identify opportunities for manufacturers to automate low-burden reporting requirements, including by integrating AI into existing reporting mechanisms.

When considering the product lifecycle, it is essential that the MHRA be empowered with the necessary resources, tools and enforcement mechanisms to ensure patient safety and device compliance issues are addressed. The current regulatory system is not designed with clear, strong incentives for proper compliance. A future regulatory framework should encourage a culture in which manufacturers take primary responsibility for compliance, quality management and post-market surveillance, supported by clear regulatory expectations and engagement from the MHRA. Building on existing regulatory systems and best practice, the aim should be to make compliance easier for responsible organisations while ensuring that non-compliance carries meaningful consequences.

This is particularly notable for software and AI-enabled devices that often challenge traditional approaches. With the potential for sudden changes in performance at scale, there is a need for mechanisms by which safety signals are rapidly detected and then rapidly acted on. The MHRA should therefore explore the extent to which their existing reporting systems and enforcement mechanisms equip them to do this, and any gaps should be addressed.

Where quality signals, changes in performance or new safety concerns that could impact patients are identified, there should be clearly established processes and expectations for sharing these with healthcare providers, professionals and the public. The sharing of these signals, near misses and emerging performance concerns should be encouraged, recognising that these can provide important opportunities to identify and address risks before patient harm occurs.

This should be supported by strengthened national and international information sharing protocols, to enable the sharing of information between stakeholders who use different systems. This would help build upon findings from the National Commission's research and engagement programme, where members of the public called for greater transparency and

accountability over when AI is involved in care, including clearer communication when things go wrong[47].

The Call for Evidence also heard clearly from stakeholders that manufacturers need to be held responsible for the safety of their device at the point of sale, but also after the device has been deployed into healthcare settings[48]. Where manufacturers identify performance and safety issues and do not take the necessary steps to address these issues or do not follow corrective actions issued by the MHRA, there should be clear consequences (such as financial penalties or fines). The MHRA needs to be empowered to take appropriate compliance and enforcement action to demonstrate to manufacturers that devices that fail to comply with the regulations and pose harm to patients will not be tolerated. The scale of consequences (such as financial penalties) should be determined by the seriousness of the compliance issue and patient impact in question.

This approach would allow the MHRA to focus its resources on identifying, investigating and addressing poor practice, while creating strong incentives for responsible manufacturers to maintain robust quality management systems, promptly share relevant signals and take corrective action when issues arise.

Recommendation 22. Enhanced enforcement mechanisms

The MHRA should develop and introduce enhanced enforcement mechanisms that enable responsive and timely engagement with manufacturers of software and AI-enabled devices to address emerging issues and support patient access to safe, effective and high-quality devices. For example, this could include mechanisms such as:

a. more timely and transparent sharing of early quality and safety signals for a given device or device type with providers, healthcare professionals and the public. It is important that these stakeholders have early awareness of potential device safety concerns that may alter the benefit-risk profile

b. the issuance of financial penalties or fines to penalise manufacturers that fail to comply with legal requirements and place patients at risk.

[47] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of findings from the Call for Evidence, p.21.
[48] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of findings from the Call for Evidence, p.36.

As technology evolves, it is important that approaches to compliance and engagement across the lifecycle evolve as well. Newer AI-enabled medical devices challenge the historical approaches to device compliance for reasons of complexity, context dependency and frequent updating throughout the lifecycle. Some challenges can be addressed through more clarity and predictability, to remove ambiguity from the system about when UK medical device regulations apply to a given product.

This will mean that these software and AI-enabled medical devices will need a more modern and ongoing approach to compliance. As noted in the National Commission's Call for Evidence, stakeholders called for new regulatory approaches that are continuous and ongoing and help to manage compliance across the lifecycle of an AI-enabled medical device[49].

The MHRA could therefore increase its efforts to work with manufacturers to define and undertake compliance tests of more responsive approaches. These efforts could focus on working with a range of manufacturers to design and refine modern approaches to compliance, that are better suited to AI-enabled medical devices whose functionality and performance can change across the lifecycle. This work should focus on exploring compliance approaches that assure devices are safe and effective while reducing unnecessary friction in the compliance process as much as possible. In line with the regulatory framework as a whole, compliance approaches should be proportionate and recognise a device's intended purpose, benefits and risks. Learnings from these efforts should be transparently shared and used to inform future approaches to lifecycle-based compliance and regulation for AI-enabled medical devices.

Recommendation 23. Compliance approaches

The MHRA should work with a range of stakeholders, such as manufacturers and Approved Bodies, to identify, test and refine how more responsive approaches to engagement will support a more holistic, lifecycle-based approach, driving and incentivising regulatory compliance, while assuring patients have access to safe, effective and high-quality devices.

[49] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of findings from the Call for Evidence, p.12.

Chapter 2. System-wide responsibility and safe management

Key principle: *Consider regulation and safe management of software and AI-enabled health technologies as a system-wide function, with each part of the system delivering its responsibilities so these technologies can be used effectively and responsibly*

Chapter summary	This chapter sets out how responsibility should be allocated across manufacturers, healthcare providers, professionals and system partners, including patient redress, operational conditions for safe deployment, contracts, AI readiness, training and reporting culture.

AI-enabled health technologies require a system-wide approach to assuring safety and accountability. Safe and effective use depends not only on the technology itself but on the actions of stakeholders across the healthcare system. The proposals in this section seek to provide greater clarity and certainty for manufacturers, healthcare providers and professionals, while avoiding accountability gaps that could undermine risk mitigation and patient redress if incidents occur.

Unlike many traditional health technologies, the performance of AI technologies can be influenced by various factors in their use and implementation in real world settings and may evolve over time as data, clinical practice and deployment environments change. As a result, safe adoption depends on the wider systems, processes and people that support its procurement, deployment, oversight and ongoing use.

As described in Chapter 1, safety cannot be adequately assured at a single regulatory point or through product regulation alone. Because of their context dependency and changing performance over time, safety needs to be an ongoing process informed by feedback on how these technologies perform in the real world. Responsibility for safe use therefore cannot sit with one actor. Instead, manufacturers, healthcare providers, professionals, regulators and other system partners should each have distinct but complementary responsibilities for ensuring AI-enabled health technologies are used safely, effectively and responsibly.

This shared-responsibility model reflects the complexity of the healthcare ecosystem, including the diversity of clinical care needs, AI technologies, data sources and deployment settings, ranging from hospitals and community care settings to patients' homes.

The recommendations in Chapter 2 are intended to strengthen linkages across the ecosystem in a purpose-led way. They seek to improve information sharing and cooperation between stakeholders to support safe clinical use, the timely identification and escalation of risks, effective accountability through clarity of roles and responsibilities, and continuous improvement over time.

Shared responsibility and system-wide safe management

Responsibility should sit with those best placed to understand, control and manage each risk.

2.1 Clarity of roles and responsibilities

It is important that the new framework clarifies roles and responsibilities to support a more effective and collaborative approach to responsible development, deployment and oversight of AI technologies in healthcare. This was reflected in the National Commission's Call for Evidence, where respondents highlighted uncertainty about how roles, responsibilities and liability are defined and applied in practice. Patients and the public called for a comprehensive approach to accountability that addresses current gaps; healthcare professionals emphasised the need to maintain clinical accountability while clarifying shared responsibilities across the system; and healthcare providers highlighted the importance of robust governance structures and clear organisational accountability[50]. Further, recent deliberative public research found that participants viewed human oversight as an important mechanism for maintaining accountability and ensuring the accuracy of AI-supported decisions[51].

Increased clarity is necessary to address some of the most pressing tensions arising from the use of AI in healthcare. For example, under the current negligence framework, claims are disproportionately likely to be brought against healthcare professionals and healthcare providers, as these actors have the clearest duty of care. This contributes to perceived 'liability sinks', where responsibility for errors may be transferred onto healthcare professionals and providers without sufficiently recognising the influences of AI systems and wider system design. To address this, healthcare professionals have called for consistent and clear guidance on where their responsibilities start, end and depend on others. This should be supplemented with explicit guidance to help healthcare professionals make more informed decisions on when reliance on AI advice is appropriate or not.

Similarly, the National Commission's research and engagement activities found that greater clarity in approaches to liability is likely to support safer and more impactful adoption of AI-enabled products in healthcare[52]. Across these responses, there was consensus on the need for a more structured approach to understanding the distribution of liability that reflects the different roles played by developers, manufacturers, healthcare providers, professionals and regulators.

Liability for AI is an area of active exploration, both in healthcare and beyond. The UK Jurisdiction Taskforce has recently issued their 'Legal Statement on Liability for AI Harms'[53],

[50] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement Programme for the National Commission into the Regulation of AI in Healthcare p.19.
[51] Thornton et al – The Health Foundation (2026) The public's views on the regulation of AI in health care: Findings from deliberative research and public polling and implications for policy, p.26.
[52] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement Programme for the National Commission into the Regulation of AI in Healthcare p.15.
[53] The UK Jurisdiction Taskforce (2026) Legal Statement on Liability for AI Harms under the private law of England and Wales.

which explores how liability would fall under the common law of England and Wales. This features a number of healthcare examples. They highlight that the existing law is ‘a well-developed flexible common law system, [that] has frequently accommodated novel and disruptive technical developments and demonstrated the ability to provide certainty and predictability in the context of technological innovation.’ The National Commission echoes their observation that ‘the perception of uncertainty.... can discourage the responsible use of technology – even where that perception is inaccurate’ and welcomes their efforts to dispel uncertainty where the law is clear, but also to highlight where genuine uncertainty relating to AI remains.

During the National Commission's research and engagement work, a key area of interest was the uncertainty arising from what the 'required standard of care' may look like in the context of AI technologies. The National Commission welcomes the work of the Professional Standards Authority and the healthcare professional regulators and Accredited Registers it oversees, as they work together to explore this issue, including guidance, position statements and other resources. The General Medical Council (GMC) has already published a learning resource covering how to apply its standards in the context of AI[54], and a number of organisations have work in development. The National Commission endorses their efforts to work together to not only bring clarity within the individual organisations, but also consistency across the organisations.

In principle, responsibility for managing a particular risk should sit with those best placed to manage it, including those with the ability to prevent harm, access relevant information, or respond when issues arise. Conversely, individuals including healthcare professionals and users of AI should not be asked to carry accountability for risks they cannot plausibly mitigate. Those held accountable should also have the resources, authority and incentives needed to manage those risks in practice. Without this clarity, there is a risk that responsibilities fall between organisations or are pushed onto stakeholders who are not well placed to manage them. This could undermine safe deployment, weaken accountability and make it harder for patients to understand who is responsible if harm occurs. The health system and developers should therefore work together to ensure that responsibility is allocated in a way that is effective and clear at each stage of the product lifecycle.

[54] For example: General Medical Council Artificial intelligence and innovative technologies.

Recommendation 24. Clear allocation of responsibility

DHSC and relevant healthcare organisations from England and the devolved authorities should work together to ensure that distribution of responsibility at each stage of the AI product lifecycle is clear and appropriate to their respective roles.

As we, as a society, begin to explore what high-quality AI-enabled healthcare looks like, and establish expectations and norms around its use, it is important that patients and the public feel involved and informed, including understanding what to expect and how to seek answers, raise concerns, or obtain redress where they have suffered harm.

The Commission has undertaken a comprehensive programme of work, including legal analysis, stakeholder engagement and a series of roundtables across industry, professional bodies and regulators, to understand different perspectives on the suitability of the current legal framework for liability in AI-enabled healthcare. In particular, professional stakeholders raised concerns that AI could create practical barriers to patients accessing compensation where harm occurs[55].

Liability for AI in healthcare is complex and evolving. Wider legal reform may be needed, is likely to take time, and will require cooperation beyond the National Commission. Any future legal reform should create the conditions that encourage the responsible innovation of AI-enabled healthcare technologies, while providing appropriate and fair routes to redress so affected parties can challenge harmful AI-enabled outcomes. In the meantime, the health system should take practical steps to ensure that the adoption of AI technologies does not weaken patients' access to redress. This includes supporting patients to understand when and how AI has influenced their care, ensuring relevant evidence can be obtained where harm occurs, and supporting access to redress where an actor has failed to meet their duty of care to patients.

[55] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement Programme for the National Commission into the Regulation of AI in Healthcare, p.39.

Recommendation 25. Patient communication and redress

DHSC should convene relevant healthcare organisations from England and the devolved authorities to work together to ensure that, as AI technology is adopted, patients have access to redress if they receive a standard of care that falls below what is reasonably expected.

Relevant organisations that invite feedback on care, including the CQC and equivalent organisations in the devolved authorities56, should work with patients (or other service users) to define what high-quality care looks like in an AI-enabled health system; and clearly communicate how to report concerns and access redress if they believe the standard of care has fallen below what is reasonably expected.

[56] Please see Glossary of terms and acronyms for Sector Regulators, Professional Regulators and National Guidance Authorisations.

2.2 System-wide commitment to delivering on risk controls with shared understanding of roles

Under the UK Medical Devices Regulations 2002, manufacturers are required to provide the information needed for a device to be used safely and in line with its intended purpose. This includes labelling, instructions for use, warnings, precautions, and information about the intended users, clinical purpose and limitations of the device. These provisions establish an important baseline by requiring manufacturers to communicate the information users need to understand how the product should be used and what risks or limitations remain.

For software and AI-enabled medical devices, safe use often depends on the implementation and effective management of particular conditions or controls in the user environment, such as appropriate cybersecurity controls, user training, monitoring processes, escalation routes or wider institutional readiness to deploy and oversee such devices. While this type of information is often included in a manufacturer's risk management file, the current framework does not consistently require manufacturers to clearly explain these dependencies in a structured and operationally usable way, or to identify what controls must be delivered by healthcare providers, healthcare professionals or consumers in practice.

This lack of clear responsibility creates a risk, such that if something goes wrong, the cause of an incident may be unclear, including whether it relates to the underlying AI-enabled device, an update, the workflow, or how the system was used. This ambiguity can limit patients' ability to receive timely answers and seek appropriate redress where harm has occurred.

To address this, manufacturers – who are best placed to understand a device's performance, limitations and uncertainties – should be responsible for communicating this information in a form the health system can act on, including by making it explicit where safe use depends on actions or conditions outside the product itself. Handover points are also needed so that duties of care and responsibilities are clear throughout the care pathway.

Pre-market submissions should include clear descriptions of the responsibilities of the manufacturer, healthcare provider and professionals, and the patient or consumer (for direct-to-consumer products), in managing risks set out in the device's risk management documentation. These measures should be assessed as part of pre-market authorisation before products are placed on the market in the future framework.

Illustrative examples of such risk controls may include:

- **Performance monitoring and detection of drift**: indicators of underperformance, minimum monitoring frequency, and required response pathways.
- **Training and onboarding**: minimum training outcomes for professional users, and onboarding for direct-to-consumer products to support safe use.

Lengthy instructions for use and labelling are often impractical and ineffective in real-world clinical workflows. Where risk controls place expectations on users, they should, wherever possible, be embedded into the device workflow and informed by human factors considerations, so the risk controls can be acted on reliably in practice. To support best practice, MHRA should publish redacted examples aimed to help translate this mechanism into practical regulatory expectations, support high-quality submissions and encourage effective risk controls that are usable in real-world care settings.

Recommendation 26. Operational conditions for safe deployment

Building on existing risk management expectations, the MHRA should require manufacturers to specify, in pre-market submissions and risk management files, the operational conditions needed for safe deployment and use. This includes the controls that must exist in the user environment – healthcare provider, health professional, or consumer context, as applicable – such as those related to cybersecurity, user training, and institutional AI readiness.

Recommendation 27. Sharing best practice examples of risk controls

To support manufacturers, the MHRA should publish guidance, including redacted illustrative examples of risk controls that demonstrate good practice and support more consistent adoption of this mechanism.

Healthcare providers play a critical role in the governance, safe deployment and use of software and AI-enabled technologies. This includes ensuring a clear understanding of, and commitment to, the ongoing arrangements needed to operate different tools safely, as informed by manufacturers' risk controls. The National Commission's research and engagement programme also highlighted the importance of this issue, with legal

professionals identifying a need for clearer contractual expectations to better support healthcare professionals and ensure that responsibilities are clearly understood in practice[57].

Where a manufacturer sells to a healthcare provider, the contract should set out the risk controls and regulatory commitments that each party is responsible for delivering as part of deployment and use, in a 'contract to deploy and use'. This means that all risk controls should be allocated before deployment and documented within the contract. The manufacturer and the healthcare provider can decide how to allocate these responsibilities (and this may vary for the same product between deployments), but no risk controls should be unaccounted for, and each party should know how they are going to deliver on them. This also means that if a manufacturer is limited in their post-market surveillance activities because the healthcare provider is not sharing data, this is identified early and a solution is found. Consideration should also be given to developing contracting principles for use between software and AI product manufacturers and NHS trusts, building on other existing approaches[58,59].

Clearer contractual allocation of responsibilities related to risk controls may help ensure frontline staff are not left carrying the full weight of post-deployment responsibilities in practice without access to the information and support needed to deliver them effectively. This approach could support consistency across the system, promote best practice and facilitate patient access to safe and effective technologies.

Recommendation 28. Contractual allocation of responsibilities

The contract between the manufacturer and healthcare provider should include an explicit allocation of responsibilities for delivering required risk controls, including those specified in the risk management file, and for meeting relevant regulatory commitments.

[57] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement Programme for the National Commission into the Regulation of AI in Healthcare p.40.
[58] National Health Service England (2026) NHS Standard Contract.
[59] National Health Service England Procurement Framework Strategy Recommendations.

2.3 A healthcare system that is equipped and empowered to safely deploy AI technologies

The public expects the health system to only deploy a technology when they can do so safely. The National Commission's research and engagement programme found that support for the use of AI in healthcare is consistently conditional, with use being linked to trust in how systems are used and governed in practice[60]. The safe and effective deployment and adoption of AI technologies in healthcare is highly dependent on the capabilities, experience, and established governance processes of healthcare providers. The level of this 'AI readiness' may be quite variable between different healthcare providers and needs to be taken seriously as a key enabler for supporting the safe deployment of these technologies. Where additional capability is needed, manufacturers may provide implementation support, or healthcare providers may draw on third-party support, to help put the required risk controls in place to support safe use of AI technologies. However, external support should not obscure the provider's responsibility to understand whether it can safely procure, deploy and oversee a particular technology in its own care context. These technologies should be deployed only where it is safe to do so, based on what is set out by the manufacturer and agreed by the regulator, where applicable.

To support this, 'AI readiness' could be characterised as a toolbox of governance measures, capabilities, and best practices for health providers to draw from in order to ensure that they can (1) make appropriate decisions about when to deploy an AI technology, and then (2) safely implement a specific AI technology. The 'AI Readiness Project' describes this as 'decision-making readiness' and 'implementation readiness' respectively[61]. It is important to note that an organisation's 'AI readiness' is usually built on existing processes and capabilities (for example, good governance processes, digital maturity, etc.). The rationale for the 'toolbox' concept is to recognise that many of these important tools may already be in place, and that whilst some are mandated and should be present in all organisations (for example, the use of DCB129 and DCB160 in England), other tools may reflect local differences. The main aim is to ensure that, taken together, risks are identified, appropriate controls are in place, and that any AI deployments are undertaken safely. Such an 'AI readiness toolbox' would also help establish a shared language between developers/manufacturers and healthcare providers and support appropriate procurement and deployment decisions.

[60] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement programme for the National Commission into the Regulation of AI in Healthcare, p.12.
[61] CERSI-AI (2026) The AI readiness checklist.

Recommendation 29. AI readiness toolbox

DHSC and the devolved health departments should co-develop an adaptable ‘AI readiness’ toolbox that can be tailored to the implementation and ongoing use of specific medical products. This would involve providers sharing their evolving approaches to meeting the operational requirements and risk controls necessary for the safe deployment and use. In return, they would benefit from access to a repository of peers’ emerging practices. Such a resource should also help providers recognise their obligations and capabilities with regard to the safe deployment of AI technologies, including:

a. an ability to self-assess their 'readiness' to deploy a particular AI product

b. an ability to deliver the risk controls necessary for a particular AI product

c. an ability to articulate the governance structures in place to safely deploy AI products.

Healthcare providers should treat ‘AI readiness’ as an essential component of good governance and clinical safety. This means being able to assess whether they are ready to adopt a specific AI-enabled technology, put the necessary operational arrangements in place, and monitor its safe use after deployment. Providers would need to understand whether an AI-enabled product is suitable for their local context and also whether they have the governance, resources, monitoring processes and escalation routes needed to manage it effectively over time.

However, providers should not be expected to develop these arrangements in isolation or on an inconsistent local basis. Local governance should be tailored to the specific care pathway and organisational context, but there is also a need for stronger national direction from DHSC and the devolved health departments on the baseline capabilities providers should aim to have in place for ‘AI readiness’. This should include nationally defined expectations of good practice as well as mechanisms for sharing learning across trusts.

A national learning function or observatory would help ensure that experience from local implementation is captured and used to improve ‘AI readiness’ across the system. As providers begin to deploy AI-enabled technologies in different clinical and organisational contexts, they are likely to encounter common challenges. The introduction of a mechanism to collect and share these lessons would help to reduce duplication and avoid repeatable mistakes. A national function could therefore identify emerging patterns across providers and support them to build on successful implementation models.

These national resources should support and not inhibit the safe adoption of AI technologies across the ecosystem. It is important that the implementation of such an approach does not result in a greater disparity between trusts; instead, it should accommodate differences in readiness and incentivise continued improvement through practical implementation experience.

Beyond 'AI readiness', it is also important to support the ecosystem with national best practice on lifecycle management of AI technologies, including procurement, implementation, monitoring, reporting and end-of-life management. These best practices should recognise and accommodate the broad spectrum of complexity and risk presented by AI technologies deployed in healthcare settings, while also harmonising with (and not duplicating) regulatory requirements, where applicable. By creating national resources informed by the real-world challenges and successes experienced across the system, providers will be better equipped to safely adopt innovative products at pace.

Recommendation 30. National leadership for AI readiness and adoption

DHSC and the devolved health departments should provide the national direction to minimise local disparities in AI readiness, adoption processes and governance approaches. To this end, they should:

a. define national best practice and expectations for organisational capabilities in AI readiness, including any baseline governance expectations, to ensure providers are equipped to enhance their AI readiness consistently across the system

b. establish a national learning function (or observatory) to collect learning from across the system, ensuring that implementation successes and challenges are shared across trusts to minimise siloed working. This could include centralising and maintaining the AI readiness toolbox for providers to access and contribute to.

Recommendation 31. National best practice on lifecycle management

DHSC and the devolved health departments should provide national best practice on the lifecycle management of emerging AI-enabled technologies (for example, agentic systems), including procurement, implementation, monitoring, reporting and end-of-life management, to support providers across the ecosystem.

Healthcare professionals play an important role in the safe use of AI-enabled health technologies. They apply clinical judgement to contextualise AI system outputs, recognise when performance may be drifting or when advice is inappropriate for a given patient, and act as a safeguard against harm where needed.

However, professionals cannot be expected to manage their responsibilities without appropriate support. Many clinicians perceive themselves to be overly burdened with the responsibility of monitoring for drift and intercepting inaccuracies, while their ability to do so is constrained by limited access to critical information and context. To perform, and be held accountable for, fulfilling their duty of care, healthcare professionals must be supported with clear guidance, training and escalation pathways that set out expectations on when and how concerns should be raised.

Clinicians and other healthcare professionals also have a responsibility to become competent in technologies that form part of their role, consistent with expectations when new technologies are introduced into health services. This includes being supported to undertake the training required to use new technologies, such as AI-enabled technologies, safely and to understand their appropriate use, limitations, and risks in clinical practice.

Training should be proportionate to professional role and level of interaction with AI-enabled products. Baseline literacy could be considered as part of mandatory training to support a common understanding of AI use in healthcare. Additional technology-specific training should then be provided where professionals are expected to use, interpret or oversee AI-enabled products. Training would be intended to support healthcare professionals to use AI-enabled products safely and confidently but does not displace professional judgement or accountability for clinical decisions. It is important that training opportunities and obligations are viewed as a necessary system investment to support the adoption of innovative technologies that have the potential to transform patient care.

Initiatives are emerging at national and local levels that seek to address the knowledge and confidence gap among healthcare professionals with regard to AI. The National Commission

welcomes the work of the NHS Digital Academy's Clinical AI Fellowships, the AI Ambassador Network of more than 14,000 members from NHS England, and ongoing work from the medical Royal Colleges, including the coordinating role of the Academy of Medical Royal Colleges. Continued investment is, however, required to meet growing demand, provide equitable access and scale these initiatives so that all healthcare professionals are equipped and empowered to use and oversee these technologies competently and confidently.

Recommendation 32. Coordinated approach to training professionals and wider staff

DHSC and the devolved health departments, working with professional regulators, those responsible for setting and delivering education and training for the healthcare professions (including the medical Royal Colleges), and healthcare education providers, should establish a coordinated approach to building professional capability in the safe and effective use of AI in healthcare. This should span initial education, postgraduate training (where relevant) and continuing professional development, and should complement setting-specific training on technologies. Expectations should be expressed as defined capabilities so that they remain applicable across the four nations and adaptable to the differing needs of professional groups and providers.

Recommendation 33. Provision of training by healthcare providers

Healthcare providers must ensure that professionals and their wider staff are well-equipped to use these technologies by undertaking appropriate training, including technology-specific training aligned with the AI technologies (including those with automated decision-making functionality) used in their practice.

Reporting of device experience, including adverse outcomes and device malfunctions, is a critical aspect of an effective regulatory framework. It is well-understood that the current framework is reactive and often characterised by significant underreporting of events. The reporting of adverse incidents and other safety issues is a core part of delivering high-quality care. It is an essential feedback loop which depends not only on the manufacturer, but also healthcare professionals, providers and sometimes the patient or carer. For some healthcare professionals, this is already enshrined in their set of professional standards. For example, in Good Medical Practice, the GMC specifically note that medical practitioners should

'contribute to adverse event recognition' and 'report adverse incidents involving medical devices (including software, diagnostic tests, and digital tools) that put the safety of a patient or another person at risk or have the potential to do so'.

A frequent observation during roundtables with professional organisations and public dialogue was that healthcare professionals and the public were usually not aware that the Yellow Card scheme applies beyond medicines to medical devices. In addition to the earlier recommendation to improve reporting mechanisms, significant culture change is therefore needed to improve the consistency and quality of reporting across the ecosystem from those who interact with AI-enabled technologies in practice.

Important information about usability, underperformance and unintended consequences, even where formal post-market surveillance systems exist, may not be captured unless healthcare professionals, providers, and patients are encouraged and enabled to report their experience. Fundamental change is needed to impact a culture shift that incentivises reporting.

Reporting should therefore be understood not only as a compliance activity, but also as a core mechanism for learning, accountability and continuous improvement across the lifecycle of AI technologies. Deliberate efforts should be made to promote a culture that values and prioritises reporting and the sharing of feedback to support safe and effective deployment of AI technologies.

Recommendation 34. Culture of reporting product experience

DHSC and the devolved health departments should promote a strong culture of reporting product experience, including adverse outcomes and product malfunctions, throughout the product lifecycle to support continuous improvement and ongoing awareness of product performance in real-world settings. This should include robust education, training and continuing professional development on reporting mechanisms and expectations for healthcare professionals as part of their duty of care.

Chapter 3. Trust, transparency and predictability

Key principle: *Design a framework that enhances trust, transparency and predictability for all stakeholders to enable responsible development, adoption and use of safe and effective AI technologies.*

Chapter summary	This chapter focuses on clearer routes for stakeholder input, more usable information for patients, professionals and providers, and a more navigable lifecycle pathway for developers and manufacturers.

Trust, transparency and predictability are fundamental for successful, safe and effective adoption of AI in healthcare. For AI-enabled technologies to be developed, adopted and used responsibly, stakeholders across the product lifecycle need confidence that the system is clear, consistent and responsive in how these technologies are designed, evaluated, deployed, monitored, maintained and regulated.

Trust may look different for different stakeholders. Patients and the public need confidence that AI-enabled products used in their care are safe, reliable, transparent, and explainable at an appropriate level and that they protect privacy. Healthcare professionals and providers need confidence that tools are accurate, supported by robust evidence, suitable for their workflows and designed to support, rather than replace, professional judgement. Developers and manufacturers need predictable requirements and clear routes through regulatory and assurance processes so they can innovate responsibly and understand what is expected across the product lifecycle.

Building this trust requires strong collaboration. Patients, the public, professionals, providers, developers, manufacturers and regulators should have meaningful and ongoing opportunities to shape how AI governance develops and improves over time. The National Commission's research and engagement programme found that trust is a key enabler of AI adoption in healthcare across system actors, patients and the public. However, the evidence also suggests that trust is not yet consistently present in practice. For example, some engagement highlighted significant gaps in the monitoring of AI in real-world use, which stakeholders viewed as a critical barrier to safe and effective deployment of AI technologies[62].

[62] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement programme for the National Commission into the Regulation of AI in Healthcare. p.32.

Additionally, current routes for engagement and information-sharing could be further strengthened. It is not always clear how stakeholder inputs are considered or how they inform regulatory decisions, which is a key factor in fostering public trust and confidence in these technologies.

In practice, different stakeholders engage with the system at different points and through different mechanisms. Patients and the public typically engage through post-market surveillance mechanisms, such as safety reporting (for example, the Yellow Card scheme), periodic MHRA engagement activities and ad hoc consultations. Healthcare professionals typically contribute through vigilance reporting and professional bodies, while healthcare providers engage primarily through local deployment assurance and procurement frameworks (for example, DCB0160[63] and DTAC[64]), with variable levels of national coordination and support. Developers and manufacturers often engage earlier in the lifecycle, during development and pre-market phases, through MHRA guidance and advice services (including the AI & Digital Regulations Service (AIDRS)[65] and AI Airlock[66]) and through engagement with Approved Bodies, while also navigating requirements across data protection, cybersecurity, NHS assurance, and evaluation.

Taken together, there is an opportunity to strengthen trust, transparency and predictability by creating clearer routes for stakeholder input, improving access to information, and making regulatory and assurance pathways easier for developers of all sizes and levels of experience to navigate.

The National Commission's recommendations aim to support a framework for software and AI-enabled health technologies that is trusted, transparent and predictable across the product lifecycle. Drawing on engagement with a wide range of stakeholders, including regulators, the recommendations focus on making it easier for patients, professionals, providers, developers and manufacturers to understand what is expected of them, how decisions are made, and how feedback and evidence shape the system over time. In practice, this means:

---

[63] National Health Service England (2023) DCB0160: Clinical Risk Management: its Application in the Deployment and Use of Health IT Systems.
[64] National Health Service England (2026) Digital Technology Assessment Criteria (DTAC) Guidance for Buyers and Suppliers.
[65] National Institute for Health and Care Excellence; Medicines and Healthcare products Regulatory Agency; Care Quality Commission; NHSE; Health Research Authority (ND) Understanding Regulations of AI and Digital Technology in Health and Social care.
[66] Medicines and Healthcare products Regulatory Agency (2024) AI Airlock: The Regulatory Sandbox for AIaMD - GOV.UK.

- Inclusive, routine opportunities for patients, the public and professionals to shape policy and regulatory approaches, supported by proportionate and accessible information about the use and safety of AI in care.

- Transparent and usable information for practitioners and providers on intended use, limitations, performance (including across different subgroups), updates and reporting routes, alongside standardised and comparable assurance information to support adoption decisions.

- A clear, navigable and proportionate lifecycle pathway for developers and manufacturers of AI-enabled health technologies, aligned across regulators and healthcare providers' assurance routes, and supported by appropriate training and competency expectations.

The recommendations below set out actions to support the responsible development, adoption and use of safe and effective AI-enabled technologies in healthcare, while strengthening trust, transparency and predictability across the system.

Trust, transparency and predictability

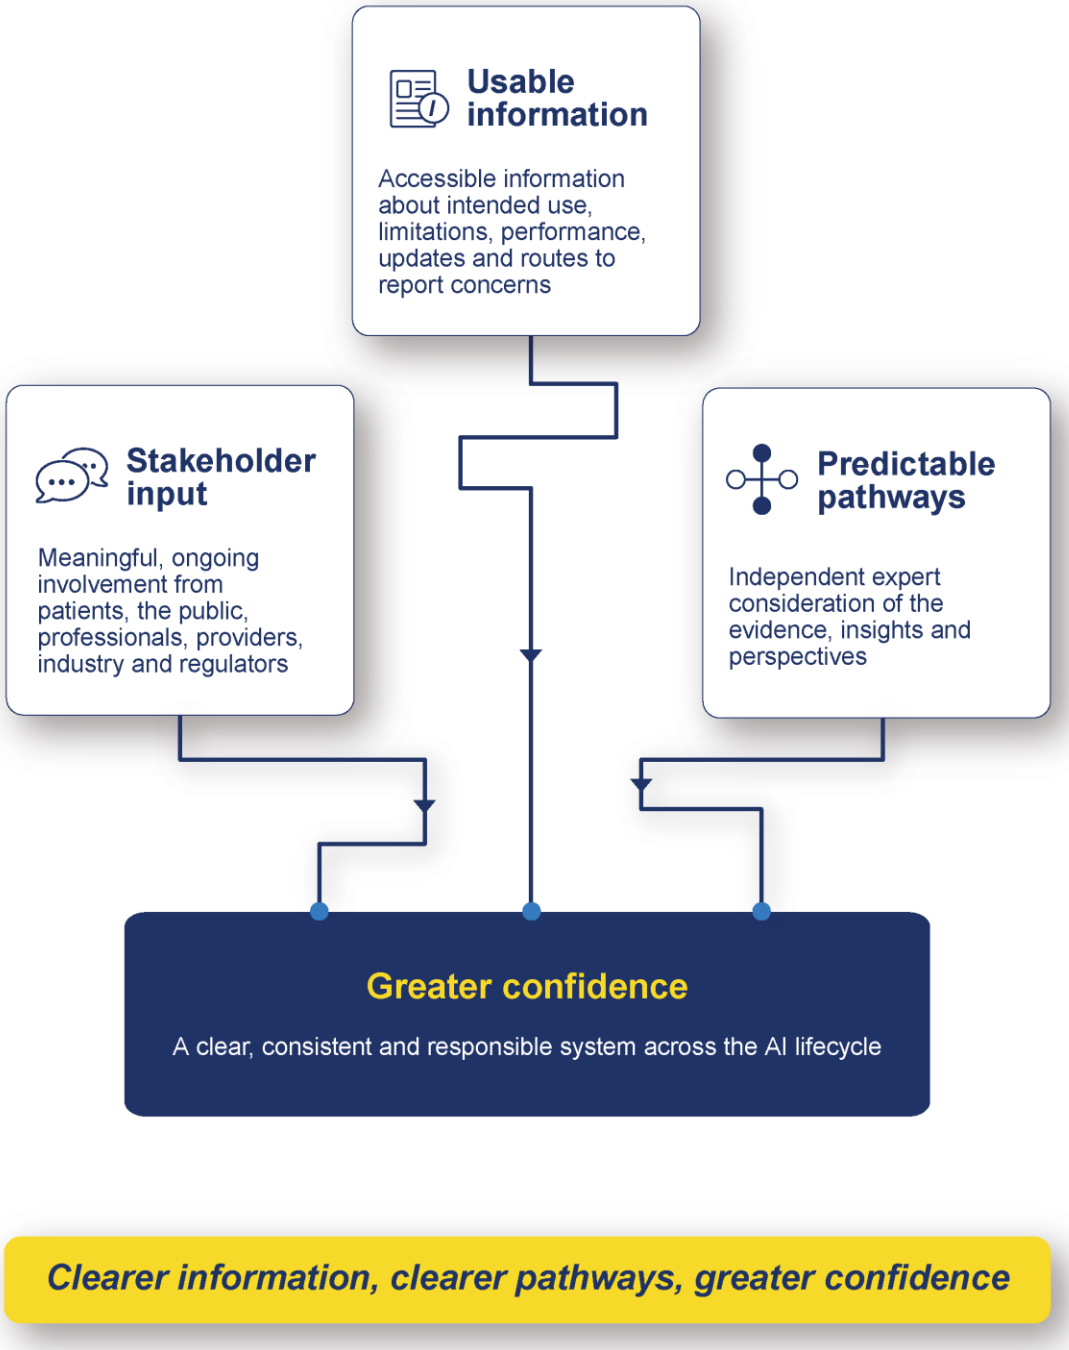

3.1 Transparency and engagement as a foundation for public trust

The National Commission's Call for Evidence found that transparency in how an AI product works and the interpretability of its outputs will be key for building patient trust in AI-enabled products in healthcare[67]. Recent deliberative research reinforces this finding, with a majority of participants believing that patients should be informed when AI has been used in their care[68]. This is particularly important as, under the current framework, there is no universal requirement or expectation for healthcare professionals to routinely inform patients every time AI is used in their care. However, existing professional standards recognise the importance of informed discussion with patients where innovative technologies are used, including the discussion of alternative options and any associated uncertainties or limitations. Taken together, the evidence suggests that, while routine disclosure of every AI use may not always be practical, clear and proportionate communication is essential to maintaining patient and public trust.

In practice, this cannot depend solely on individual healthcare professionals making disclosure decisions during every patient interaction. A system-level approach would provide a more consistent baseline of public awareness, while reserving more detailed disclosure and discussion for high-risk or more sensitive, novel uses of AI. This would align with established expectations for machine learning-enabled medical devices, where patients should receive clear and relevant information about a device's purpose, outputs and risks to support understanding and informed decision-making.

At the same time, clinical practice does not seek explicit consent for every supporting technology used in care (for example, electronic patient record systems), and requiring consent at every instance of AI use would be impractical and could potentially be detrimental to the system being able to deliver high-quality, equitable care. A proportionate approach should therefore combine general notice and public awareness with more active discussion where appropriate. This recognises that AI-enabled products sit across a spectrum of risk and benefit, and that transparency approaches should reflect the context and intended use of the technology rather than applying a one-size-fits-all approach. Clear, system-wide best practice for communication about when and how AI supports care would help patients

[67] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of Findings for the Call for Evidence p.15.
[68] Thornton et al - The Health Foundation (2026) The public's views on the regulation of AI in health care: Findings from deliberative research and public polling and implications for policy p.32.

understand what to expect, support informed conversations, and build patient and public confidence in the health system.

Recommendation 35. Patient transparency and trust

Relevant healthcare organisations from England and the devolved authorities should adopt a proportionate, system-level approach to transparency about the use of AI-enabled products in patient care, including addressing patients' reasonable expectation of being informed of the use of such technologies in their care and having the ability to opt out where possible or appropriate.

Trust in AI-enabled healthcare will depend on whether patients and the public feel that these technologies are developed, regulated and used in ways that reflect both their technical capabilities and the experiences, expectations and values of those affected by them[69]. Patients and the public need meaningful opportunities to shape how AI is governed, clear expectations about how AI may affect their care and confidence that concerns about safety and accessibility will be heard and addressed. This is further reflected in recent research examining public views on the regulation of AI in healthcare, which highlights ongoing, meaningful public engagement as a key feature of a trustworthy and responsive regulatory system[70]. Such opportunities for engagement can also provide valuable insights into how a proportionate approach should be applied as novel technologies and capabilities continue to emerge, directly shaping care delivery.

Building trust also requires collaboration with a broad range of communities, particularly those who may be most affected by poor design, digital exclusion, bias or unequal access. Trust and transparency are therefore closely linked to equity: if AI-enabled products underperform for some groups, or if people lack access to or understanding of these technologies, trust in AI-enabled care is likely to decline. This is particularly important for communities that may already have lower trust in health services or who express more cautious attitudes towards the use of AI in healthcare.

Patient and public engagement should therefore be more proactive, routine and inclusive of a diverse range of participants, particularly groups that are typically underrepresented or underserved. The system should use a mix of engagement mechanisms, tailored to different audiences, purposes and levels of impact. Using such mechanisms to engage directly with patients on emerging questions and considerations for AI-enabled products would foster a transparent, patient-centred and learning regulatory system. The MHRA should build on the success and learning from previous engagement initiatives, including its electronic patient

[69] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement programme for the National Commission into the Regulation of AI in Healthcare p.4.
[70] Thornton et al - The Health Foundation (2026) The public's views on the regulation of AI in health care: Findings from deliberative research and public polling and implications for policy, p.45.

leaflet programme, which benefited from significant patient involvement during its development.

Recommendation 36. Patient and public collaboration

The MHRA and relevant healthcare organisations from England and the devolved authorities should enhance mechanisms for ongoing inclusion of patient and public perspectives to support a learning regulatory system. These mechanisms include a Patient Engagement Advisory Committee, workshops and deliberative events.

Public trust in AI-enabled healthcare depends not only on the safety of these technologies, but also on people having access to clear, timely and meaningful information about safety when concerns arise.

Current UK medical devices regulations and MHRA processes already provide public-facing routes for safety information. There is, however, an opportunity to consider ways to further strengthen the overall patient-facing information system, especially for technologies that are rapidly evolving.

The MHRA publishes medical device safety information, including alerts and guidance where action is needed[71], and patients and the public can report suspected safety concerns through the Yellow Card scheme. Manufacturers may also issue Field Safety Notices[72] where corrective action is required, although these are often aimed at healthcare organisations and professional users rather than patients.

As a result, public-facing safety communication can be fragmented, technical and difficult for patients to navigate. This creates a particular challenge for rapidly evolving software and AI-enabled medical devices, where safety signals and concerns may emerge throughout the product lifecycle due to changes in performance, the introduction of software updates, changes in deployment context or new evidence from real-world use, among others.

As monitoring increasingly takes place across the product lifecycle, patients and the public will need accessible information about emerging safety signals, performance changes, corrective actions and any implications for their care. This information should be designed with public needs in mind, rather than solely from a technical or regulatory perspective. The MHRA should explore new mechanisms for communicating safety information to better meet the needs of the public. Further, patient and public groups should be involved in shaping how this safety information is communicated, so that it is clear, relevant and responsive to the needs of different communities. This would improve transparency and predictability,

[71] The Medicines and Healthcare products Regulatory Agency (2025) Safety Communications Concerning Medicines, Medical Devices and Other Healthcare Products.
[72] Medicines and Healthcare products Regulatory Agency (2025) Field Safety Notices: Guidance for Manufacturers.

strengthen trust and help ensure that people understand both how safety concerns are identified and how the system responds to them.

Recommendation 37. Public communication on safety concerns

The MHRA should work with patient and public groups and other relevant partners to develop robust mechanisms for providing public-facing safety information that is clear, transparent and tailored to the needs of the public.

Healthcare professionals currently use AI in a range of ways within clinical settings, including for administrative support and patient care activities such as screening, diagnosis and treatment decisions[73].

Understanding professional sentiment regarding AI-enabled technologies within healthcare is critical, because the healthcare workforce plays a central role in adopting, using and overseeing AI-enabled products in practice. The system is unlikely to be viewed as trustworthy if those directly using the tools lack confidence in either the tools themselves or the regulatory framework that governs them. Therefore, understanding healthcare professionals' trust, experience and confidence over time would provide valuable insight into barriers to adoption of AI in healthcare, implementation risks and the conditions needed for safe and effective use.

The National Commission's Call for Evidence found that 77% of healthcare professionals who responded called for a 'complete overhaul' or 'significant reform' of the regulatory framework for AI in healthcare[74]. Separate survey evidence found that 48% of staff considered AI safe for patient care, compared with 86% for administrative purposes[75]. These findings suggest that while many healthcare professionals recognise the potential value of AI-enabled healthcare, confidence varies according to the use case and the level of clinical risk. They also indicate that professional support for AI cannot be assumed. Trust varies depending on how AI is used within healthcare, with higher-risk applications in patient care needing stronger governance and clearer regulatory frameworks to support professional confidence.

Existing surveys, including those conducted by the Health Foundation, provide valuable insight into public and workforce attitudes towards AI. Building on this success, an AI sentiment census led by a government health department could further complement this by

[73] Medicines and Healthcare products Regulatory Agency (2026) Summary – Information on How AI is Used in Health and the Current Regulatory Frameworks.
[74] Medicines and Healthcare products Regulatory Agency (2026) National Commission into the Regulation of AI in Healthcare – Summary of Findings from the Call for Evidence, p.25.
[75] Binesmael et al - The Health Foundation (2026) Attitudes to Technology and AI in Health care.

monitoring trust, understanding and real-world experience of AI use in healthcare over time. This would provide an ongoing feedback mechanism to ensure policy decisions are informed by the experiences and perspectives of both patients and healthcare professionals.

A periodic ‘AI sentiment census’ could provide a structured way to monitor changes in patients’ trust, understanding, experience and acceptance of AI in healthcare. Regular, repeatable surveys would enable partners across the system to monitor trends, highlight areas where trust remains low and enable more targeted implementation decisions and tailored approaches to transparency and communication. Making findings visible, and demonstrating how they inform policy and regulatory decisions, would further strengthen transparency and trust. The ‘AI sentiment census’ could build on existing survey infrastructure, including the annual Health Foundation survey of public and patient attitudes to technology and AI in healthcare.

Recommendation 38. AI sentiment census

DHSC and the devolved health departments should consider implementing a periodic ‘AI sentiment census’ that includes both patients and professionals, so ongoing policymaking is informed by both patient and workforce experience and perspectives.

Transparency about the use of AI-enabled health technologies plays a critical role in supporting their safe and effective adoption within a clinical setting. Device users need to understand what an AI tool is intended to do, how to interpret its outputs, and when it may not be appropriate to rely on them. Design choices made early in development strongly shape whether a tool supports this understanding in practice or instead introduces new risks, particularly in complex clinical workflows.

Existing UK medical device regulations require manufacturers to provide information on a device's intended purpose, safe use, performance, warnings and precautions. However, current regulatory expectations do not explicitly set out user-centred design and transparency considerations for AI-enabled devices. As a result, users may lack the information needed to understand, interpret and use AI products confidently, or have access to clear mechanisms for sharing feedback and concerns.

To address this gap, the MHRA should work with relevant partners to develop guidance on user-centred design considerations that embed transparency considerations from the outset. This would help manufacturers develop devices that are designed with people in mind. It would also support users to better understand how AI-enabled products work, how outputs

should be interpreted, what limitations apply, and how these technologies can be used safely and effectively.

Such guidance should also promote device design with simple and intuitive reporting mechanisms for users to provide feedback and raise concerns directly with manufacturers. While existing processes such as the Yellow Card scheme support the reporting of safety issues to the MHRA, additional feedback and transparency mechanisms, such as more seamless user reporting and dynamic labelling, as suggested in Chapter 1, could capture broader user experience, support continuous improvement and help to build user trust in AI-enabled products. As a matter of good practice, manufacturers should incorporate safety considerations into product design from the outset.

Recommendation 39. User-centred design and transparency

The MHRA should work with relevant partners to develop guidance for developers on user-centred design considerations that support transparency for device users, including patients, professionals and the public, and enable simple and intuitive reporting of user feedback on the device experience.

The MHRA's existing medical device regulatory framework focuses on safety and performance, rather than on comparative value between products. As a result, organisations responsible for adopting AI technologies, such as NHS bodies and healthcare providers, may lack consistent and independent information to help them choose between different available AI-enabled products designed for the same intended use.

Where multiple AI products are available for the same use case, healthcare providers would benefit from clear and comparable evidence to determine which tool best meets their clinical, operational and system needs. While regulatory authorisation of software and AI-enabled devices provides assurance that a device meets relevant regulatory requirements, it does not necessarily enable comparisons between competing products or indicate which solution is most appropriate for a particular setting. This challenge can be greater for AI products that are not subject to device regulation. Supporting informed adoption decisions may therefore warrant consideration of additional mechanisms to synthesise or communicate comparative evidence. These approaches should not be introduced as a new requirement that could inadvertently create a system bottleneck but rather should be considered, where helpful, according to the scale, risk and strategic importance of a use case. Potential mechanisms include:

- centralised intelligence gathering and evidence synthesis, drawing together existing data and evidence, without the need to undertake new testing

- a more market-led consumer intelligence approach, akin to an independent product comparison and ratings service, delivered either by multiple third-party providers or centrally commissioned for NHS use

- a centralised testing function, which could deliver or commission vendor-independent head-to-head testing on behalf of the NHS when deemed necessary.

These mechanisms could be supported by a clearer and more established NICE pathway for evaluating and adopting AI-enabled technologies, building on NICE's existing work in digital health evidence standards[76] and health technology assessment[77], and reflecting its more developed approach to medicines. Wherever possible, the system should utilise existing organisations, expertise and programmes, including NICE, the National Institute for Health and Care Research (NIHR), and the DHSC/NHS England Digital Policy Unit.

Providing adopting organisations with clear, comparable evidence would strengthen trust, transparency and predictability. It would give healthcare providers greater confidence that adoption decisions are based on robust and relevant information, make the basis for those decisions more transparent to professionals and patients, and give developers a clearer understanding of the types of evidence that matter for real-world adoption. Over time, a more consistent approach to comparative evidence could support procurement decisions, reduce duplication across providers and help ensure AI-enabled products are adopted where they are most likely to deliver safe, effective and appropriate care.

Recommendation 40. Comparative evidence to support adoption decisions

DHSC and the devolved health departments should work with NICE and equivalent organisations in devolved authorities, to consider ways to provide adopting organisations with additional information regarding clinical effectiveness and value to support decisions on whether to adopt an AI health tool and where appropriate, which product best fits their needs.

[76] National Institute for Health and Care Excellence (2022) Evidence Standards Framework for Digital Health Technologies.
[77] National Institute for Health and Care Excellence (2026) NICE Technology Appraisal and Highly Specialised Technologies Guidance: The Manual.

3.2 Regulatory clarity and predictability to support the innovation journey

Bringing software and AI-enabled devices to market requires developers and manufacturers to navigate a complex regulatory landscape spanning multiple organisations, requirements and stages of the product lifecycle[78]. Depending on the nature of the technology, this may involve engagement with the MHRA, NICE, the Health Research Authority (HRA), the Care Quality Commission (CQC), the Information Commissioner's Office (ICO) and other bodies, alongside requirements relating to UK medical device regulation, evidence generation, health technology assessment, NHS assurance, and data governance. While services such as the AIDRS (AI and Digital Regulations Service) provide valuable support, innovators - particularly first-time and small to medium-sized enterprise (SME) developers - continue to report uncertainty about how requirements fit together, what applies to their product, and when and how to engage with regulators and other system partners. This was reflected in the National Commission's research and engagement programme, where stakeholders highlighted the need for clearer guidance on classification, safety standards and regulatory processes, alongside clarity on how evolving AI systems fit within existing regulatory frameworks[79].

To support and balance responsible innovation and regulatory processes, expectations should be clear and predictable across the product lifecycle. When developers and manufacturers face uncertainty or inconsistency, regulation can deter responsible innovation rather than enable it. The MHRA can reduce this uncertainty by offering a more predictable and collaborative regulatory experience for manufacturers of all sizes. Developers and manufacturers should therefore feel confident that they are operating in a clear, predictable regulatory environment to support the development of safe, effective and high-quality devices.

Building on AIDRS, developers would benefit from a user-friendly, interactive, public-facing resource that maps the regulatory and assurance journey across the lifecycle for AI-enabled products. This resource could signpost the key guidance, standards and templates required at each stage, and help users identify which requirements are relevant to their product. The resource should also address common FAQs and clearly signpost available engagement opportunities for innovators across the system. By improving clarity on what is required,

[78] National Institute for Health and Care Excellence (2026) Artificial Intelligence (AI) and Digital Regulations Service.
[79] Medicines and Healthcare products Regulatory Agency (2026) Report on the research and engagement programme for the National Commission into the Regulation of AI in Healthcare, p.18.

when, and by whom, it would increase transparency and predictability for innovators, support higher-quality submissions and reduce avoidable delays and rework[80].

Recommendation 41. Lifecycle journey map for developers and manufacturers

The MHRA, CQC, NICE, HRA, ICO, and equivalent organisations in devolved authorities, and relevant organisations from the devolved authorities should develop clear, accessible, public-facing educational resources and tools within their scope of activities, with the overall intent of making it easier for developers and manufacturers to navigate the regulatory and assurance journey for AI health technologies. These resources should demystify the journey and support new developers by reducing ambiguity across the lifecycle.

Early uncertainty about regulatory status, classification, evidence requirements and applicable routes to market can create significant downstream challenges for developers, including generating evidence that fails to meet regulatory expectations, undertaking inappropriate study designs, delays in assessment and costly rework later in the product lifecycle. Developers of software and AI-enabled devices would therefore benefit from earlier clarity and engagement opportunities to ensure that product development and evidence generation plans align with regulatory expectations from the outset.

Greater predictability within the regulatory framework would increase confidence in the pathway to market, support investment in innovation, and strengthen the UK's attractiveness to developers and manufacturers as a destination for developing and deploying such devices. This should be achieved in a way that supports innovators of all sizes, avoiding regulatory approaches that could unintentionally create barriers for smaller developers or reinforce the position of larger incumbents.

Providing clearer, earlier and more predictable opportunities for regulatory engagement before formal submission can help developers to be well positioned for success. In particular, structured pre-submission engagement, which results in clear written feedback from the MHRA and relevant partners, would give manufacturers greater confidence in their proposed regulatory approach and allow them to invest in appropriate evidence generation and study design[81]. Where complexity or novelty warrants it, such engagement could also

[80] United States Food and Drug Administration (2025) Regulatory Accelerator.
[81] United States Food and Drug Administration (2025) Requests for Feedback and Meetings for Medical Device Submissions: The Q-Submission Program; Guidance for Industry and Food and Drug Administration Staff.

include a meeting to discuss key issues directly. This kind of approach could build on existing regulatory advice meetings.[82]

Such engagement mechanisms should include clear, timely opportunities for developers to interact with the MHRA on a specific product and its intended purpose. As described in Chapter 1, 'intended purpose' can be a particularly difficult regulatory concept to apply to software and AI-enabled products. For example, a responsible developer of a large language model (LLM)-enabled product, that is not intended to be used for a medical purpose, should be able to approach the MHRA to request feedback on whether their proposed risk mitigations and controls are sufficient in the view of the regulator for that product to be outside the scope of UK medical device regulation. Similarly, it should be possible for a developer to have early engagement with regard to the evidence requirements for a device, to support market authorisation; it is noted that NICE already provides such a service within their remit[83]. It would be helpful if the MHRA and NICE consider opportunities to build on existing joint scientific advice arrangements to provide tailored support for manufacturers of software and AI-enabled devices.

The MHRA should also explore mechanisms for providing formal qualification and classification confirmation based on an assessment of information about a given product submitted by a manufacturer. Having documentation of a product's qualification and classification from the regulator can be a critical part of the development journey for a manufacturer as it provides clarity and confidence for the company, investors and other stakeholders impacted.

Recommendation 42. Early regulatory advice and pathway confirmation

The MHRA should introduce clearer, earlier and more predictable engagement mechanisms for developers and manufacturers, in collaboration with relevant partners. These mechanisms should include a structured route that provides written feedback (and, where necessary, meetings) on topics such as device qualification, classification, the appropriate path to market, evidence requirements and post-market expectations.

[82] Medicines and Healthcare products Regulatory Agency (2026) Medical devices: ask for a regulatory advice meeting from the MHRA - GOV.UK.
[83] National Institute for Health and Care Excellence NICE Advice Service.

Recommendation 43. Formal classification confirmation service

The MHRA should also consider offering a formal ‘classification confirmation’ service, enabling manufacturers to submit a short package and receive a written determination on a product’s regulatory status, reducing avoidable delay and rework.

AI-enabled products will increasingly rely on real-world evaluation and monitoring after deployment to demonstrate that they remain safe and effective in local settings. For software and AI-enabled devices, UK MDR 2002 imposes post-market surveillance obligations on manufacturers, including monitoring device performance, investigating incidents and taking corrective action. However, there is less clarity for developers on when local evaluation, assurance and monitoring activities should be treated as research requiring HRA approval, and when they form part of routine post-market evaluation, clinical governance, service evaluation, audit or quality improvement.

This distinction is particularly important for AI-enabled products, where safe deployment may be even more dependent on local testing, ongoing monitoring and controls before, during and after implementation. As a result, developers may face uncertainty about which evaluation and monitoring activities require formal research approvals, and which require governance and oversight through other mechanisms. This uncertainty can delay deployment, discourage proportionate safety checks and lead to inconsistent approaches across organisations.

Clearer HRA guidance and decision-support tools, aligned with MHRA and wider system guidance, would help organisations distinguish research from routine post-market evaluation, local assurance and clinical governance. This would support proportionate oversight, enable appropriate safety and performance checks in real-world settings, and reduce unnecessary regulatory burden.

Recommendation 44. Post-market evaluation clarity

As AI-enabled technologies increasingly rely on post-market monitoring and evaluation, the HRA, and equivalent organisations in devolved authorities, should build on its existing guidance and decision-support tools to help organisations determine when development, deployment, evaluation and monitoring activities constitute research and require research governance, and where they may sit outside of research, requiring governance and oversight through other mechanisms. This work should be aligned with MHRA and wider system guidance as the regulatory framework develops.

Acknowledgements

The National AI Commission and MHRA would like to thank everyone who contributed to the Commission's work and the development of this report.

We are particularly grateful to the patients, members of the public, carers and representatives of communities across the four nations who generously shared their views, experiences and expectations for how AI should be used in healthcare through the Commission's research and engagement programme. Their contributions have been essential in ensuring that the recommendations remain grounded in patient safety, public confidence and the responsible use of AI in healthcare. We also extend our thanks to the many individuals and organisations who contributed their time and insight, including healthcare professionals, industry representatives, regulators and health and care system partners.

We also wish to express our sincere thanks to the Commission members and Working Group members. Through their expertise, challenge, insight and commitment, they have played a central role in shaping the recommendations contained in this report. Their contributions helped ensure that the recommendations are evidence-based, informed by practical delivery considerations and reflective of perspectives from across the healthcare, regulatory, technology and policy landscape.

Glossary of terms and acronyms

This list contains a glossary of commonly used words, phrases, and abbreviations from this report. Please note that this glossary does not have a legal function, rather is intended as a tool to assist readers in understanding of this report, including by providing context for technical language and terminology. While some of the definitions below are set out in medical device regulations, such as the UK MDR and EU MDR, others are set out in guidance or reflect the National AI Commission's interpretation and use of these terms.

Term	Definition
Accountability (in AI systems)	The assignment of responsibility for the design, deployment, performance and outcomes of an AI system, including who is responsible when things go wrong across its lifecycle.
Accredited Registers	Registers of health and care practitioners who are not regulated by law but whose register-holding organisations are accredited by the Professional Standards Authority.
Adaptive technologies	Technologies that can continuously learn after deployment, modifying their behaviour or outputs over time in response to new data.
Adverse incident	An event involving a product or device that has caused, or could cause, harm to a patient, user or another person.
Agentic AI systems	AI systems capable of acting autonomously to pursue goals, often coordinating multiple tasks or tools with limited direct human oversight.
AI-enabled medical device	AI software that performs a medical intended purpose and is regulated as a medical device under the UK medical device regulations.
AI Airlock	An MHRA regulatory sandbox that allows developers, regulators and system partners to test innovative AI-enabled medical devices in a controlled environment to generate evidence, identify regulatory challenges and support safe deployment.
AI readiness	The governance, capabilities, resources, processes and organisational arrangements needed to safely adopt, deploy,

Term	Definition
	monitor and manage AI technologies within a healthcare setting.
AI sentiment census	A repeated survey used to understand how patients, the public and healthcare professionals feel about AI in healthcare, including their trust, experience and concerns.
Ambient Voice Technology (AVT)	AI-enabled technology that captures and processes conversations, often to generate summaries, clinical documentation or other outputs to support healthcare delivery.
Approved Body	An organisation designated by the MHRA to assess whether a medical device and its manufacturer meet the requirements of the UK Medical Device Regulation before a device can be registered to be placed on the market.
Artificial General Intelligence (AGI)	A theoretical AI system capable of performing the full range of intellectual tasks that humans can perform, including learning, reasoning, problem-solving, and adapting to new situations across multiple domains.
Artificial intelligence (AI)	Technologies that simulate human intelligence by performing complex tasks through the use of algorithms or models which may learn from data, recognise patterns, make predictions, and generate outputs.
Automation bias	The tendency to over-rely on or favour AI system outputs, even when they may be incorrect.
Clinical decision support	Software or AI tools that help healthcare professionals make decisions about patient care, without replacing professional judgement.
Conformité Européenne (CE) marking	A product marking indicating that a manufacturer considers a medical device to meet applicable EU requirements, allowing it to be placed on the relevant European market.
Context-dependent performance	Variation in how an AI system performs depending on factors such as local data, deployment environment, user capabilities or organisational settings.

Term	Definition
Continuous oversight	Ongoing monitoring and review of an AI system's safety and performance after deployment, rather than relying solely on pre-market checks.
Conformity assessment	The process used to demonstrate that a medical device meets the applicable regulatory requirements for safety, performance and quality before it can be placed on the market.
Cybersecurity	The state in which a medical device and its associated systems are protected so that they continue to function safely and as intended, preserving the integrity and availability of data and device functionality throughout the product lifecycle.
Device classification	The process of assigning a medical device to a risk class, which determines the regulatory requirements and level of conformity assessment that apply.
Device qualification	The process of determining whether a product falls within the legal definition of a medical device and is therefore subject to medical device regulation.
Devolved health departments	Government departments responsible for health policy within the devolved authorities of the UK. This includes the Scottish Government Health and Social Care Directorates, the Welsh Government Health Care and Prevention Group, and the Department of Health in Northern Ireland.
DHSC and the devolved health departments	The health policy functions of the UK Government and the devolved governments.
Direct-to-consumer (DTC) products	Products or services that are marketed and supplied directly to members of the public without necessarily involving a healthcare professional in their selection, use or interpretation.

Term	Definition
Dynamic labelling	A mechanism for providing users with regular, up-to-date information about a device, such as performance, limitations, safety information or updates.
Enforcement discretion	When a regulator does not intend to enforce some or all applicable requirements for specified products, generally where risks are considered sufficiently low.
EU Medical Device Regulations (EU MDR)	Regulation (EU) 2017/745
Failure mode	A specific way in which a medical device, software or AI product may fail to perform as intended, resulting in incorrect, degraded, delayed or unavailable functionality. Identification of failure modes is used in risk management to assess potential hazards, their causes, and their impact on patients, users or other stakeholders.
Field Safety Notice	A communication from a manufacturer about action needed to reduce a risk relating to a medical device already in use.
Foundation models	General-purpose AI models trained on extensive datasets that can be adapted for a wide range of tasks and downstream applications
Hallucinations (in AI)	Incorrect or misleading outputs generated by AI systems, often presented as if accurate or factual
Healthcare provider	An organisation responsible for delivering health or care services, including NHS organisations, hospitals, primary care providers, independent healthcare organisations and other care settings
Human oversight	The involvement of a human in reviewing, validating or controlling AI outputs
Intended purpose	The use to which the device is intended for according to the data supplied by the manufacturer on the labelling, the instructions for use and/or the promotional materials

Term	Definition
Large Language Models (LLMs)	AI systems trained on large volumes of text data that can interpret and generate natural language (for example, chatbots)
Liability	Legal responsibility for harm or adverse outcomes, including determining who is accountable under common law
Lifecycle regulation / oversight	An approach to regulation that considers an AI product across its full lifecycle, including design, training, development, deployment, use, ongoing monitoring and decommissioning
Manufacturer / Developer	(a) the person with responsibility for the design, manufacture, packaging and labelling of a device before it is placed on the market under their own name, regardless of whether these operations are carried out by that person themselves or on their behalf by a third party; or (b) any other person who assembles, packages, processes, fully refurbishes or labels one or more ready-made products or assigns to them their intended purpose as a device with a view to their being placed on the market under their own name, apart from a person who assembles or adapts devices already on the market to their intended purpose for an individual patient.
Master File	A confidential repository of technical information relating to a product or technology that can be referenced when assessing devices that rely on that product or technology.
Medical device	Any instrument, apparatus, appliance, software, material or other article, whether used alone or in combination, together with any accessories, including the software intended by its manufacturer to be used specifically for diagnosis or therapeutic purposes or both and necessary for its proper application, which – (a) is intended by the manufacturer to be used for human beings for the purpose of-

Term	Definition
	(i) diagnosis, prevention, monitoring, treatment, or alleviation of disease, (ii) diagnosis, monitoring, treatment, alleviation of or compensation for an injury or handicap, (iii) investigation, replacement, or modification of the anatomy or of a physiological process, or (iv) control of conception; and (b) does not achieve its principal intended action in or on the human body by pharmacological, immunological or metabolic means, even if it is assisted in its function by such means, and includes devices intended to administer a medicinal product or which incorporate as an integral part a substance which, if used separately, would be a medicinal product and which is liable to act upon the body with action ancillary to that of the device.
Model Cards	Documents that summarise an AI model's intended purpose, performance, limitations, and other key information to support transparency, accountability and appropriate use.
Model drift	Changes in an AI system’s performance over time due to shifts in data, environment or use, which may reduce accuracy or reliability.
National guidance organisations	Organisations that develop evidence-based clinical guidance to support healthcare practice. Examples include the National Institute for Health and Care Excellence (NICE), the Scottish Intercollegiate Guidance Network (SIGN), and Health Technology Wales (HTW).
National health systems	The publicly funded health services responsible for delivering healthcare in each nation, including NHS England, NHS Scotland, NHS Wales and Health and Social Care Northern Ireland (HSCNI). Arrangements in England are currently subject to reform though the Health Bill[84].

[84] Department for Health and Social Care (2026) Health Bill.

Term	Definition
Performance drift	A change in a device's performance over time, including deterioration or unexpected variation arising from changes in data, users, settings or operating conditions.
Post-market surveillance (PMS)	Activities carried out by manufacturers to proactively collect, and review experience gained from devices placed on the market or put into service for the purposes of identifying any need to apply corrective or preventive actions. These include requirements, such as post-market surveillance plans, introduced by the Medical Devices (Post-market Surveillance Requirements) (Amendment) (Great Britain) Regulations 2024[85].
Post-market evaluation	Activities undertaken after deployment to assess the performance, safety or effectiveness of a technology in real-world settings.
Predetermined Change Control Plan (PCCP)	A regulatory mechanism that allows for certain changes to be implemented for a medical device after authorisation, without requiring a new regulatory submission for each change.
Pre-market assessment / authorisation	Regulatory evaluation a medical device must undergo before it can be placed on the market to ensures the device meets the requirements for safety and performance
Professional and representative bodies	Organisations that represent, support and provide leadership for healthcare professionals and organisations, including the Medical Royal Colleges and NHS Alliance.
Professional regulators	Organisations responsible for setting standards for professional practice, for the education and training of the professions they regulate, for maintaining registers of qualified professionals, and for maintaining confidence in regulated professions. Examples include the General Medical Council (GMC), the Nursing and Midwifery Council (NMC), and the Health and Care Professions Council (HCPC). These regulators are overseen by the Professional

[85] Medical devices: post-market surveillance requirements (2025) Medical devices: post-market surveillance requirements - GOV.UK.

Term	Definition
	Standards Authority (PSA). The PSA also accredits and sets standards for organisations holding registers of health and care practitioners not regulated by law, known as 'Accredited Registers'.
Real-time monitoring	Continuous observation of an AI product's performance as it operates, enabling early detection of issues or risks.
Real-world evidence	Data and insights generated from the use of technologies in real-world settings, as opposed to controlled testing or trials, used to assess performance and safety over time.
Recognition and reliance pathways	Regulatory routes that use or take account of assessments or approvals from trusted regulators in other jurisdictions, while retaining domestic responsibility for the final regulatory decision.
Redress	The mechanisms available to individuals to seek remedy, compensation, explanation or corrective action when harm, loss or adverse outcomes have occurred.
Relevant healthcare organisations from England and the devolved authorities	Term used where recommendations may affect organisations beyond the health departments themselves, including national health system bodies, sector regulators and professional regulators
Risk-based regulation	A regulatory approach in which the level of scrutiny applied depends on the potential risk of the system to patients or users.
Risk management file	The collection of records and documentation maintained by a manufacturer that identifies, assesses, controls and monitors risks associated with a medical device throughout its lifecycle.
Sandbox	A collaborative, time-bound environment where innovators can test novel technologies, generate evidence and refine practices within existing regulatory frameworks under regulator supervision, before wider deployment.

Term	Definition
Sector regulators	Organisations that regulate, inspect, or oversee healthcare services. Examples include the Care Quality Commission (CQC), Healthcare Improvement Scotland (HIS), the Healthcare Inspectorate Wales (HIW), the Regulation and Quality Improvement Authority (RQIA), and the Health Research Authority (HRA).
Service evaluation	An assessment of how a service is working in practice, usually to improve local delivery rather than to produce generalisable research findings.
Staged Authorisation	A regulatory approach that allows a technology to be deployed initially within a defined scope, with expansion contingent on additional evidence generation and monitoring.
System-level approach	An approach that considers how AI interacts with people, organisations and processes, rather than treating it as a standalone product.
Traceability	The ability to identify and follow a device, product or version through its lifecycle, including where and when it has been used.
Transparency (in AI systems)	Openness about how and when AI is used, including its role, limitations, and how decisions are made or supported.
UKCA marking	A product marking indicating that a medical device has met the applicable requirements for placement on the Great Britain market following the relevant conformity assessment.
UK Medical Device Regulations (UK MDR)	The UK Medical Devices Regulations 2002 (SI 2002 No 618, as amended).
Unique Device Identifier (UDI)	A unique code assigned to a medical device to support identification, traceability and monitoring throughout its lifecycle.
Windsor Framework	The arrangements governing the application of certain EU rules in Northern Ireland, under which EU medical device

Term	Definition
	regulations continue to apply to devices placed on the Northern Ireland market.
Yellow Card scheme	The UK system for collecting and monitoring information on safety concerns such as suspected side effects or adverse incidents involving medicines and medical devices.

Annex A. List of Recommendations

Key Principle: Ensure regulatory requirements are proportionate to a device's risks and benefits, tailored to the evidentiary needs across the product lifecycle and support equitable access to safe and effective software and AI-enabled technologies.

Clear and Consistent Device Qualification and Classification

Recommendation 1. The MHRA should systematically review and update, as necessary, the existing UK Medical Devices Regulations to allow for a more tailored approach to software and AI-enabled medical devices. This should include:

a. updates to the definition of a medical device to provide more clarity on when certain software and AI-enabled products are regulated as medical devices. For example, such updates should explicitly clarify when software intended for administrative purposes in a healthcare setting, general wellbeing purposes, and certain decision support purposes are not medical devices.

b. a classification system and approach that considers the clinical risks, patient benefits, device lifecycle, and international harmonisation to provide a clear and resilient approach to classification of software and AI-enabled devices. This classification system should:

   i. address the known limitations of self-declared Class I devices in the current framework

   ii. appropriately apply requirements and/or oversight for devices where there is sufficient risk to patients

   iii. employ a proportionate approach to low-risk devices.

c. updates to the definition of intended purpose to ensure that device design and functionality are appropriately considered in addition to a manufacturer's claims and promotional materials.

Recommendation 2. The MHRA should systematically review and update, as necessary, guidance and other resources to provide clarity and predictability for the ecosystem. The MHRA should develop and implement an ongoing process for reviewing and updating guidance and other resources to ensure supporting policies are updated and available in a timely fashion. This should include:

a. clarification of which types of software and AI-enabled products are, and are not, medical devices

b. clarification of the application of the risk-based classification system to relevant software and AI-enabled devices

c. clarification of how intended purpose is interpreted for software and AI-enabled devices, including practical examples relevant to both understanding device qualification and post-market compliance.

Tailored and Risk Proportionate Regulatory Approaches

Recommendation 3. The MHRA should be empowered to improve the flexibility of its application of the regulatory framework, allowing for tailored approaches over time to allocate its resources and oversight to best assure the public health. For example:

a. adding safety-related requirements for devices where new/available evidence shows a need for this beyond which is covered by existing classification requirements to support patient safety.

b. reducing or limiting requirements and/or oversight for devices where the best available evidence shows that the risk is sufficiently low (for example, applying concepts such as enforcement discretion, where appropriate). Where such actions are proposed, it is essential that this be signalled clearly, applied consistently, and aligned to risk proportionality and patient benefit.

Recommendation 4. The MHRA should consider a function-based approach to the regulation of software and AI-enabled medical devices, recognising increasing product complexity and supporting proportionate regulation. It is critical that regulatory oversight be tailored to and focused on the medical device functionality, rather than any non-medical device functionalities present in products where there may be multiple functionalities.

Recommendation 5. The MHRA and relevant partners should consider how to appropriately balance oversight and evidence expectations across the product lifecycle. In particular, this should include clear identification of sufficient pre-market data and post-market surveillance for software and AI-enabled devices. This should be in line with the expectations and best interests of the public, clinical risks, technological characteristics, and the intended purpose of a given device. For example, the MHRA should provide clarity regarding what pre-market evidence is needed to support the safety and effectiveness of an adaptive AI-enabled device while considering the real-world data monitoring and reporting that will be needed to reduce overall uncertainty.

Recommendation 6. The MHRA should consider robust approaches to change management, including providing clarity on the scope and implementation of PCCPs for software and AI-enabled devices. This should include ensuring PCCPs provide a clear path and support for:

a. adaptive and continuously learning AI-enabled devices by moving toward an approach that allows for clear boundaries and guardrails to define the allowable scope of change rather than requiring specific, planned changes

b. site or sub-population specific tuning

c. devices intended to operate with broad intended purposes or goal-versus-task defined intended purposes (for example, broad medical domain applicability, regulated AI agents)

d. evidence-driven use expansion (for example, detecting/characterising additional radiology image findings to support clinician diagnosis or validating additional genes as they are characterised for an AI-driven diagnostic) that does not alter the device's intended purpose.

Recommendation 7. The MHRA should consider establishing an opt-in 'Master File' approach for general-purpose platforms, foundation models or underlying technologies used by medical device developers. These 'Master Files' can be used to support the assessment of medical devices leveraging such platforms or models. This should include:

a. engagement with technical experts to explore the kinds of proportionate information that would provide long term value within a 'Master File', such as benchmarks, model cards, internal guardrails and firewalled mechanisms of providing additional information as available and necessary. The information should aim to support assurance that a given device leveraging the underlying technology will continue to function as intended

b. development of MHRA voluntary guidance setting out the process for a 'Master File' and the types of information a 'Master File' should contain for such underlying technologies to adequately support device developers and regulatory decision making.

Recommendation 8. The MHRA, DHSC and the devolved health departments should consider how to obtain the necessary information when a medical product has an underlying dependency on a general-purpose model. Specifically, the manufacturer should be expected to transparently report this dependency, any related risks and mitigations (including continuity plans). This should be an expectation as part of regulatory submissions (for AI-enabled medical devices) and through procurement and contract terms (for AI-enabled products).

Recommendation 9. The MHRA should consider issuing guidance to support developers in providing clear, transparent information to users and the public, such as through the use of model cards, careful user-centred interface design and dynamic labelling to support the safe use of software and AI-enabled devices.

Recommendation 10. The MHRA should consider providing clear guidance and educational resources on cybersecurity expectations across the product lifecycle for software and AI-enabled devices, which are increasingly connected across networks and with other devices. Cybersecurity is a critical factor of device safety as cybersecurity threats can directly compromise the safety or effectiveness of devices, including the potential to cause rapid harm to many people.

Recommendation 11. The MHRA should consider providing clarity through guidance on regulatory expectations for direct-to-consumer applications and wearables to fully harness their potential for the UK such as by supporting a shift towards preventative care. It is important that careful consideration be given to how to apply a proportionate approach to empower people with access to information about their health while understanding the unique benefits and risks of direct-to-consumer use.

Recommendation 12. The MHRA should consider providing clarity through guidance on human factors, usability considerations and expectations to support safe use as part of the design, development and testing of software and AI-enabled medical devices.

Recommendation 13. The MHRA should consider working with relevant partners to provide guidance for manufacturers on health equity expectations for software and AI-enabled medical devices across the product lifecycle, including consideration of the diversity of the intended population, with particular regard to underserved population groups.

Clear and Timely Paths to Market (With sufficient risk controls)

Recommendation 14. The MHRA should enable staged authorisations, supported by appropriate deployment pathways that support evidence generation and infrastructure to enable sustainable, long-term patient access to beneficial devices. These authorisation pathways should be:

a. sufficiently flexible to allow tailoring of the authorisations and pathways to support those devices, where a staged authorisation is appropriate. This should include consideration of technological maturity, clinical needs and necessary risk controls to be put in place to ensure patient safety;

b. transparently defined and reported to ensure patient awareness and support public confidence;

c. designed to address the recognised limitations of single-timepoint assessment of software and AI-enabled devices, including demonstrating real world performance, managing uncertainty, and providing lifecycle assurance of device safety and benefit; and

d. clearly defined processes and pathways from staged to full market authorisation to enable predictable, uninterrupted delivery of beneficial devices to patients.

Recommendation 15. The MHRA, in collaboration with system partners, should continue using regulatory sandboxes and extend their remit to support the development and testing of regulatory approaches, best practice and safe deployment pathways for novel technologies that provide a clear path to market authorisation.

Recommendation 16. The MHRA should continue its efforts to advance international harmonisation of the regulation of software and AI-enabled devices by:

- sharing lessons learned and best practices for proportionate, lifecycle-based regulation to align where possible and support responsible innovation for patients; and
- introducing recognition or reliance pathways for relevant devices with trusted international regulators, where it is in the best interests of patients.

Proportionate post-market assurance

Recommendation 17. The MHRA should continue to tailor post-market surveillance requirements to be proportionate with a device's risks, benefits and related uncertainty to ensure effective, resource-efficient oversight. This should include a menu or set of mechanisms that manufacturers may be required to implement, including:

a. post-market surveillance plans and other real-world data collection / monitoring plans assessed at the time of device authorisation to reduce uncertainty

b. real-world post-market studies where needed

c. more regular, ongoing performance reporting made available publicly and/or to regulators, as appropriate, to provide continued assurance of device safety and effectiveness

d. established processes for escalation where performance degradation is detected, even if a reportable incident has not occurred.

Recommendation 18. The MHRA, manufacturers and relevant healthcare organisations from the devolved authorities should work with partners across the healthcare ecosystem to review and refine mechanisms for reporting and sharing information (including best practice) across the ecosystem to ensure that they are clear, accessible, and, where appropriate,

standardised to enable more robust insight into ongoing performance and safety. This should support timely sharing of information from deployment sites and greater awareness of available reporting routes.

For example, the MHRA should consider how the Yellow Card scheme can be improved for AI and software-enabled devices, including for different users (such as patients) who may need to interact with the system. Likewise, manufacturers could support this process by focusing on user-centred design that enables real-time feedback, including the reporting of adverse incidents and malfunctions.

Recommendation 19. The MHRA should consider developing a database or tool that enables members of the public to easily search for adverse incidents related to a specific software or AI-enabled medical device. This should build on initiatives such as MHRA's interactive Drug Analysis Profiles (iDAPs) and include functionality that enables members of the public to search and filter by manufacturer, device name and timeframe, among other variables.

Infrastructure, monitoring and enforcement

Recommendation 20. The MHRA should introduce mechanisms to identify AI-enabled medical devices once deployed, including appropriate version control, to provide clear traceability throughout a device's lifecycle. For example, a unique device identifier (UDI) should be explored as a potential mechanism for traceability within the patient record, enabling healthcare professionals and providers to identify and audit the outcomes of patients whose care involved a particular device. The MHRA should also work closely with DHSC and the devolved health departments to support implementation, ensuring that device traceability information, including version, is routinely recorded as part of the patient record.

Recommendation 21. The MHRA should identify opportunities for manufacturers to automate low-burden reporting requirements, including by integrating AI into existing reporting mechanisms.

Recommendation 22. The MHRA should develop and introduce enhanced enforcement mechanisms that enable responsive and timely engagement with manufacturers of software and AI-enabled devices to address emerging issues and support patient access to safe, effective and high-quality devices. For example, this could include mechanisms such as:

a. more timely and transparent sharing of early quality and safety signals for a given device or device type with providers, healthcare professionals and the public. It is important that these stakeholders have early awareness of potential device safety concerns that may alter the benefit-risk profile

b. the issuance of financial penalties or fines to penalise manufacturers that fail to comply with legal requirements and place patients at risk.

Recommendation 23. The MHRA should work with a range of stakeholders, such as manufacturers and Approved Bodies, to identify, test and refine how more responsive approaches to engagement will support a more holistic, lifecycle-based approach, driving and incentivising regulatory compliance, while assuring patients have access to safe, effective and high-quality devices.

Key Principle: Consider regulation and safe management of software and AI-enabled health technologies as a system-wide function, with each part of the system delivering its responsibilities so these technologies can be used effectively and responsibly

Clarity of roles and responsibilities

Recommendation 24. DHSC and relevant healthcare organisations from England and the devolved authorities should work together to ensure that distribution of responsibility at each stage of the AI product lifecycle is clear and appropriate to their respective roles.

Recommendation 25. DHSC should convene relevant healthcare organisations from England and the devolved authorities to work together to ensure that, as AI technology is adopted, patients have access to redress if they receive a standard of care that falls below what is reasonably expected.

Relevant organisations that invite feedback on care, including the CQC and equivalent organisations in the devolved authorities[86], should work with patients (or other service users) to define what high-quality care looks like in an AI-enabled health system; and clearly communicate how to report concerns and access redress if they believe the standard of care has fallen below what is reasonably expected.

System-wide commitment to delivering on risk controls with shared understanding of roles

Recommendation 26. Building on existing risk management expectations, the MHRA should require manufacturers to specify, in pre-market submissions and risk management files, the operational conditions needed for safe deployment and use. This includes the controls that must exist in the user environment – healthcare provider, health professional, or

[86] Please see Glossary of terms and acronyms for Sector Regulators, Professional Regulators and National Guidance Authorisations.

consumer context, as applicable – such as those related to cybersecurity, user training, and institutional AI readiness.

Recommendation 27. To support manufacturers, the MHRA should publish guidance, including redacted illustrative examples of risk controls that demonstrate good practice and support more consistent adoption of this mechanism.

Recommendation 28. The contract between the manufacturer and healthcare provider should include an explicit allocation of responsibilities for delivering required risk controls, including those specified in the risk management file, and for meeting relevant regulatory commitments.

A healthcare system that is equipped and empowered to safely deploy AI technologies

Recommendation 29. DHSC and the devolved health departments should co-develop an adaptable 'AI readiness' toolbox that can be tailored to the implementation and ongoing use of specific medical products. This would involve providers sharing their evolving approaches to meeting the operational requirements and risk controls necessary for the safe deployment and use. In return, they would benefit from access to a repository of peers' emerging practices. Such a resource should also help providers recognise their obligations and capabilities with regard to the safe deployment of AI technologies, including:

a. an ability to self-assess their 'readiness' to deploy a particular AI product

b. an ability to deliver the risk controls necessary for a particular AI product

c. an ability to articulate the governance structures in place to safely deploy AI products.

Recommendation 30. DHSC and the devolved health departments should provide the national direction to minimise local disparities in AI readiness, adoption processes and governance approaches. To this end, they should:

a. define national best practice and expectations for organisational capabilities in AI readiness, including any baseline governance expectations, to ensure providers are equipped to enhance their AI readiness consistently across the system

b. establish a national learning function (or observatory) to collect learning from across the system, ensuring that implementation successes and challenges are shared across trusts to minimise siloed working. This could include centralising and maintaining the AI readiness toolbox for providers to access and contribute to.

Recommendation 31. DHSC and the devolved health departments should provide national best practice on the lifecycle management of emerging AI-enabled technologies (for example, agentic systems), including procurement, implementation, monitoring, reporting and end-of-life management, to support providers across the ecosystem.

Recommendation 32. DHSC and the devolved health departments, working with professional regulators, those responsible for setting and delivering education and training for the healthcare professions (including the medical Royal Colleges), and healthcare education providers, should establish a coordinated approach to building professional capability in the safe and effective use of AI in healthcare. This should span initial education, postgraduate training (where relevant) and continuing professional development, and should complement setting-specific training on technologies. Expectations should be expressed as defined capabilities so that they remain applicable across the four nations and adaptable to the differing needs of professional groups and providers.

Recommendation 33. Healthcare providers must ensure that professionals and their wider staff are well-equipped to use these technologies by undertaking appropriate training, including technology-specific training aligned with the AI technologies (including those with automated decision-making functionality) used in their practice.

Recommendation 34. DHSC and the devolved health departments should promote a strong culture of reporting product experience, including adverse outcomes and product malfunctions, throughout the product lifecycle to support continuous improvement and ongoing awareness of product performance in real-world settings. This should include robust education, training and continuing professional development on reporting mechanisms and expectations for healthcare professionals as part of their duty of care.

Key Principle: Design a framework that enhances trust, transparency and predictability for all stakeholders to enable responsible development, adoption and use of safe and effective AI technologies.

Transparency and engagement as a foundation for public trust

Recommendation 35. Relevant healthcare organisations from England and the devolved authorities should adopt a proportionate, system-level approach to transparency about the use of AI-enabled products in patient care, including addressing patients' reasonable expectation of being informed of the use of such technologies in their care and having the ability to opt out where possible or appropriate.

Recommendation 36. The MHRA and relevant healthcare organisations from England and the devolved authorities should enhance mechanisms for ongoing inclusion of patient and

public perspectives to support a learning regulatory system. These mechanisms include a Patient Engagement Advisory Committee, workshops and deliberative events.

Recommendation 37. The MHRA should work with patient and public groups and other relevant partners to develop robust mechanisms for providing public-facing safety information that is clear, transparent and tailored to the needs of the public.

Recommendation 38. DHSC and the devolved health departments should consider implementing a periodic 'AI sentiment census' that includes both patients and professionals, so ongoing policymaking is informed by both patient and workforce experience and perspectives.

Recommendation 39. The MHRA should work with relevant partners to develop guidance for developers on user-centred design considerations that support transparency for device users, including patients, professionals and the public, and enable simple and intuitive reporting of user feedback on the device experience.

Recommendation 40. DHSC and the devolved health departments should work with NICE and equivalent organisations in devolved authorities, to consider ways to provide adopting organisations with additional information regarding clinical effectiveness and value to support decisions on whether to adopt an AI health tool and where appropriate, which product best fits their needs.

Regulatory clarity and predictability to support the innovation journey

Recommendation 41. The MHRA, CQC, NICE, HRA, ICO, and equivalent organisations in devolved authorities, and relevant organisations from the devolved authorities should develop clear, accessible, public-facing educational resources and tools within their scope of activities, with the overall intent of making it easier for developers and manufacturers to navigate the regulatory and assurance journey for AI health technologies. These resources should demystify the journey and support new developers by reducing ambiguity across the lifecycle.

Recommendation 42. The MHRA should introduce clearer, earlier and more predictable engagement mechanisms for developers and manufacturers, in collaboration with relevant partners. These mechanisms should include a structured route that provides written feedback (and, where necessary, meetings) on topics such as device qualification, classification, appropriate path to market, evidence requirements and post-market expectations.

Recommendation 43. The MHRA should also consider offering a formal 'classification confirmation' service, enabling manufacturers to submit a short package and receive a written determination on a product's regulatory status, reducing avoidable delay and rework.

Recommendation 44. As AI-enabled technologies increasingly rely on post-market monitoring and evaluation, the HRA, and equivalent organisations in devolved authorities, should build on its existing guidance and decision-support tools to help organisations determine when development, deployment, evaluation and monitoring activities constitute research and require research governance, and where they may sit outside of research, requiring governance and oversight through other mechanisms. This work should be aligned with MHRA and wider system guidance as the regulatory framework develops.

Annex B. Declarations of Interest

The MHRA ensured that the AI Commission's approach to managing conflicts of interest was aligned with the standards and processes applied to other independent expert groups. The Commission identified, recorded and managed actual, potential and perceived conflicts of interest throughout its work.

All Core Commission members, Working Group members and Commission Office staff were required to declare relevant financial and non-financial interests. Declarations were centrally recorded, monitored for completeness and retained to support transparency and appropriate oversight. Regular reminders were provided by the respective Chairs during Core Commission and Working Group meetings, prompting members to complete a Register of Interests and update their declarations where circumstances had changed.

Recognising that many declared interests arose directly from members' expertise, professional experience and leadership roles within relevant sectors, declarations did not automatically preclude participation. The Commission's role was advisory, providing recommendations to inform MHRA decision-making on future regulatory frameworks.

All declarations were reviewed and managed through agreed oversight arrangements. Following review of declared interests, no member was required to be recused from any Commission or Working Group meeting or activity.

A range of proportionate mitigation measures was applied where relevant. To minimise the risk of any actual or perceived commercial advantage, access to developing recommendations and draft report material was carefully controlled. Core Commission members had the opportunity to review and provide feedback on draft versions of the report and its recommendations.

In addition, draft or near-final versions of the report were also shared with a limited number of Working Group members and related stakeholders with a role in implementing the recommendations, including government officials, representatives of the devolved authorities, and key health system partners. Their engagement was intended to ensure factual accuracy, operational feasibility and implementation considerations were appropriately reflected before publication and to inform the draft Government response that will follow receipt of the recommendations.

Industry representatives on the Technology Working Group participated in an advisory capacity and were not provided with draft or near-final versions of the report, nor consulted on the final recommendations prior to publication.

As an advisory body, the Commission did not involve decisions relating to the allocation of public funding, procurement, grants or other financial resources. Recommendations informed by the Commission's work do not in themselves determine funding allocations within the MHRA, nor the allocation of grants or contracts to industry, academic institutions or individuals. This separation between advice and decision-making formed an important part of the Commission's conflict of interest management arrangements.

In the interests of transparency and public accountability, Declarations of Interest for Core Commission members are published alongside this report. Working Group members were appointed to provide specialist technical, clinical, operational and policy expertise to support the Commission's deliberations. They did not hold a governance or decision-making role within the Commission and were not responsible for agreeing the final recommendations.

Taking a proportionate approach to transparency and recognising the advisory role of Working Group members, individual Working Group declarations have not been published. However, all Working Group members were required to complete declarations of interest, which were reviewed and managed in accordance with the Commission's conflict of interest arrangements.

Collectively, the disclosure requirements, ongoing oversight and proportionate mitigation measures provided a robust framework for identifying and managing actual, potential and perceived conflicts of interest. This helped safeguard the integrity, independence and credibility of the Commission's work.

Alongside the arrangements described above for Commission participants, all MHRA staff members are required to follow MHRA conflicts of interest policies and procedures. Staff are required to declare relevant interests on appointment, annually and whenever circumstances change. In addition, MHRA staff supporting the Commission were also asked to register their interests, in line with the approach taken for Commission and Working Group members.

Commissioners' Declarations of Interest – September 2026

Professor Alastair Denniston, Professor of Regulatory Science and Innovation at the University of Birmingham; Honorary Consultant Ophthalmologist at University Hospitals Birmingham NHS Foundation Trust, and Executive Director of the UK Centre of Excellence for Regulatory Science in AI & Digital Health (CERSI-AI).

Organisation	Nature of interest / role	Paid?	Current?
CERSI-AI	Executive Director	No	Yes

Professor Henrietta Hughes OBE, Patient Safety Commissioner for England; practising GP and Visiting Professor at the Institute of Medicine, University of Greater Manchester.

Organisation	Nature of interest / role	Paid?	Current?
IBM	Family member employed / Shareholder	Yes	Yes

Professor Neil Lawrence, DeepMind Professor of Machine Learning at the University of Cambridge; Chief Scientist at Trent AI.

Organisation	Nature of interest / role	Paid?	Current?
AstraZeneca	Family member employed	Yes	Yes
Cambridge Innovation Capital	Provides AI advice and informal investment advice	No	Yes
IMU Biosciences	Provides advice, including being listed on pitch decks	No	Yes
Trent	Founder / Chief Scientific Officer	Yes	Yes

Professor Catherine Sudlow OBE, Head of School of Population Health Sciences, University of Edinburgh; Director of UKRI Adolescent Health Study.

No relevant financial or non-financial interests, roles or investments were declared.

Dr Brian Anderson, CEO of the Coalition for Health AI (CHAI).

Organisation	Nature of interest / role	Paid?	Current?
Coalition for Health AI	Chief Executive Officer	Yes	Yes

Dr Ricardo Baptista Leite, CEO of HealthAI.

No relevant financial or non-financial interests, roles or investments were declared.

Adjunct Professor Raymond Chua, CEO of Health Sciences Authority; Deputy Director-General of Health (Health Regulation) for Ministry of Health, Singapore.

No relevant financial or non-financial interests, roles or investments were declared.

Dame Jennifer Dixon, Chief Executive of the Health Foundation.

No relevant financial or non-financial interests, roles or investments were declared.

Dr Paul Goldsmith, Non-Executive Director of the MHRA; Chair of the MHRA's Regulatory and Safety Committee; visiting Professor at the Institute of Global Health Innovation, Imperial College London; Consultant Neurologist and Clinical Senate.

Organisation	Nature of interest / role	Paid?	Current?
Cambridge University ARIA NeuroWorks Scientific Advisory Board (SAB)	Scientific Advisory Board member	No	Yes

Closed Loop Medicine Ltd	Director and shareholder	Yes	No
Foundation for Evolution and Mental Health	Trustee	No	Yes
Ieso Health	Share options	No	Yes
Institute of Global Health Innovation (IGHI), Imperial College, London	Visiting Professor	No	Yes
Lanthor Ltd	Book publishing and medico-legal reports	Yes	Yes
Morchella Medical (Drug dosing optimisation, including the use of AI)	Consultant	Yes	Yes
NHS	Clinical Senate Member	No	Yes
NHS	Consultant Neurologist	Yes	Yes
Radix Big Tent Foundation	Trustee	No	Yes

Professor Vish Ratnasuriya MBE, Practising GP; Chair of Our Health Partnership; co-founder of Primary Care Accelerator; Honorary Professor at the University of Birmingham.

Organisation	Nature of interest / role	Paid?	Current?
Our Health Partnership	Chair	Yes	Yes
Primary Care Accelerator	Co-founder	Yes	Yes

Dr Gabriella Spinelli, Head of Brunel Design School, College of Engineering, Design and Physical Science, Brunel University of London.

Organisation	Nature of interest / role	Paid?	Current?
RADIANT-CERSI Centre for Excellence in Regulatory Science & Innovation in Digital Health & Healthcare AI	Director	No	Yes
Home Office Scientific Advisory Committee	Scientific Advisor	Yes	Yes

Dr Barry Stein, Chief Clinical Innovation Officer and Chief Medical Informatics Officer for Hartford HealthCare.

Organisation	Nature of interest / role	Paid?	Current?
Aidoc	Executive Advisory Board	No	Yes
Athos Therapeutics	Advisory Board	Yes	Yes
Beckers Healthcare	CIO Advisory Board	No	Yes
Coalition for Healthcare AI (CHAI)	Advisory Board member	No	Yes
Flare Capital Partners	Industry Advisory Board & AI Advisory Council	No	Yes
Global Center for Glycogen Storage Disease	Founding Member and Board of Trustees	No	Yes
Greater New York Hospital Association	AI Advisory Board	No	Yes
Healthier Capital	Limited Partner Advisory Committee	No	Yes
Holistic Hospital Optimization	Board of Directors	No	Yes
K Health	Advisory Board	No	Yes
Massachusetts Institute of Technology (MIT)-Hartford HealthCare Clinical and Operational Analytics Research Partnership	Hartford HealthCare Principal	No	Yes
MIT-LINQ Accelerator Programs (Catalyst, IDEA2)	Independent Advisory Faculty	No	Yes
Meridian Health Ventures	Limited Partner Advisory Committee	No	Yes

Richard Stubbs, Chief Executive of Health Innovation Yorkshire & Humber.			
Organisation	Nature of interest / role	Paid?	Current?
Health Foundation NHS Productivity Commission Expert Advisory Panel	Member	No	Yes
Westfield Health	Non-executive Director	Yes	Yes
Yorkshire and Humber Partners Academic Health Science Network Ltd.	Director	Yes	Yes

Professor Richard Susskind CBE KC, President of the Society for Computers and Law.
No relevant financial or non-financial interests, roles or investments were declared.

Other Commission meeting attendees:

Lawrence Tallon, CEO of Medicines and Healthcare products Regulatory Agency.			
Organisation	Nature of interest / role	Paid?	Current?
PharosAI (publicly funded, not for profit capacity building initiative)	External Advisory Board member	No	Yes

Dr Joseph Alderman, NIHR academic clinical lecturer in anaesthetics and intensive care; Postdoctoral researcher in AI and digital health at both the University of Birmingham and CERSI-AI.			
Organisation	Nature of interest / role	Paid?	Current?
The Alan Turing Institute	Co-organiser of clinical AI interest group	No	Yes
Faculty of Intensive Care Medicine (FICM)	Member of the professional affairs and safety committee	No	Yes

The Medicines and Healthcare products Regulatory Agency (MHRA) established the National Commission into the Regulation of AI in Healthcare in September 2025, to review current regulations and provide recommendations for a new regulatory framework for AI in healthcare. As the sponsoring organisation, the MHRA commissioned and oversaw the delivery of the Commission's programme of work. This included the establishment of the Commission, governance, management of the research programme, stakeholder engagement, secretariat and operational support. The Commission operated independently in developing its advice and recommendations, and responsibility for the report's conclusions and recommendations rests solely with the Commission.