BCS' response to the Call for Evidence: National Commission into the Regulation of AI in Healthcare

Medicines & Healthcare products Regulatory Agency

Feb 2026

Compiled by Claire Penketh, BCS Senior Policy and Public Affairs Manager

Table of Contents

Executive Summary

BCS, The Chartered Institute for IT argues that the UK requires **significant reform of its regulatory framework for AI in healthcare to ensure both patient safety and timely access to effective innovation**. The existing system, built around traditional medical devices, is not suited to modern, adaptive AI systems that evolve over time and influence clinical decisions, workflows, and population health management.

BCS identifies substantial shortcomings in current regulation, particularly in **safety standards, transparency, clinical evidence requirements, and post-market surveillance, all of which are judged insufficient**. The organisation calls for a shift to a lifecycle-based model of oversight, with progressive and proportionate assurance that includes robust pre-deployment evaluation, controlled real-world testing, and mandatory post-deployment monitoring for bias, drift, and usability. Post-market surveillance is highlighted as a critical weakness; BCS recommends requiring suppliers to demonstrate surveillance capability before approval and making monitoring a shared responsibility.

BCS emphasises the need for clearer rules on liability, arguing that current laws are not fit for AI-driven systems. Responsibility should align with the level of control and knowledge held by manufacturers, provider organisations, and clinicians at each stage of the AI lifecycle. This includes addressing automation bias, tracking repurposed use and updating governance.

The response also stresses that **safe deployment depends on workforce capability and capacaity, professional registration of Digital, Data and Technology staff, and consistent digital skills training across the NHS**. Finally, BCS urges MHRA to improve predictability for innovators through clear service-level timelines and aligned standards, helping ensure safe, trustworthy AI adoption at scale. The aim should be to align existing regulations to improve safety, not increase regulatory burden.

These are the views of BCS Faculty of Health and Care (FoHC) members and BCS Fellows working in digital, data, clinical and technical roles in health and social care. BCS thanks the following for their contributions to this consultation: Jeremy Wyatt, Emeritus, professor of Digital Healthcare, Univ. of Southampton, former clinical advisor on validation methods at NHSX AI Lab; Kieran Zucker, Health Data Scientist, Clinical Oncologist; Hema Purohit FBCS, Former CTO, Microsoft for Healthcare and Public Sector for Generative AI; Rowland Agidee, Chief Data and Analytics Officer, University Hospitals of Derby and Burton; Dr. Sebastian Alexander FBCS. MB. ChB. AFOM. FRCGP. FFCI, Chair of Clinical Safety Group & Clinical Informatic Manager; Paul Llewellyn FBCS, Chief Clinical Information Officer at Accenture and Avanade, LFEDIP, BSI Medical Devices Regulation Professional; Maggie Lay, Clinical Safety Expert, Clinical Informatician, Registered Nurse;

and Dr Avi Mehra FBCS, MBBS, Associate Partner & Clinical Safety Officer, IBM Healthcare & Life Sciences.

Q1. Which of the following best describes your view about the need to change the UK's framework for regulating AI in healthcare?

c) Significant reform: The current framework requires substantial changes. ✅

Q2.1 - Q2.5: To what extent do you agree or disagree that the current regulatory framework is sufficient in the following domains:

For each, select from: - Strongly disagree - Disagree - Neither agree nor disagree - Agree - Strongly agree

a) Safety & performance standards: *Strongly disagree*
b) Data privacy & data governance: *Disagree*
c) Transparency: *Strongly disagree*
d) Requirements for clinical evidence: *Disagree*
e) Post-market surveillance: *Strongly disagree*

Q3: How would you rate the current framework's impact on innovation?

b) Somewhat restrictive [creates some barriers] ✅

Q4: How might the UK's framework for regulation of AI in healthcare be improved to ensure the NHS has fast access to safe and effective AI health technology?

Currently, the framework, based on traditional medical device rules, doesn't address the adaptive, software-defined characteristics of AI systems in clinical workflows. The UK should shift from a static, product-based approach to a lifecycle, system-aware regulatory model for health and care.

AI systems and Software as a Medical Device (SaMD) often don't fit neatly within device categories. Workflow orchestration engines, population-level risk stratification models, or foundation model-based decision support may not qualify as devices but still influence clinical decisions.

<u>1.Focus regulation on the magnitude and nature of clinical risk</u>

Regulation should focus on clinical risk—likelihood and severity of harm not just device type. Frameworks for health apps exist and can help. Clearer boundaries for classification, intended use, and user skills are needed. Where software spans clinical and non-clinical domains, explicit triggers for reclassification - Class 1, 2, etc - should be defined and

depend on use in each context. This will reduce inconsistency and misleading performance claims. Aligning SaMD and AI regulation would also help.

2.Adopt progressive, proportionate assurance

Traditional pre-market clinical trials are inadequate for adaptive AI. The UK should adopt progressive, proportionate assurance that combines:

- Robust pre-deployment validation tailored to model behaviour and risk.
- Controlled real-world evaluation before scaling;
- Mandatory post-deployment monitoring for drift, bias, and usability issues.

This approach echoes international practices such as the EU AI Act's transparency and the FDA's process-based oversight of trusted manufacturers. Mandated standards - evaluation frameworks, model cards, and structured reporting - would improve comparability. Transparency must be clearly defined – is it supplier, algorithmic, or clinical-level? Clinicians should know when AI is active, its version, data sources, update date, and intended use. QR codes or "about" sections can help.

Ending restrictive NDAs would support learning from failures and improve safety. Healthcare organisations should be part of a regulated ecosystem, backed by a code of conduct and assurance models to prevent shifting risk to clinicians.

3. Make post-market surveillance a core regulatory pillar

Make post-market surveillance a core regulatory pillar - see our answer to Q5 - the current post-market regime, including the Yellow Card scheme, is inadequate. Each AI system should have standardised reporting and clear intervention thresholds.

4.Reduce regulatory burden while improving safety

At present, loopholes allow under classification (e.g., self-certification as Class I when Class IIa/IIb may be warranted), while limited capacity among approved bodies and associated costs create barriers, especially for SMEs, where pilot projects using small data sets face excessive burdens. To improve safety and reduce burden, MHRA clinical risk documentation should be aligned with NHS DCB 0129/0160, with a clear requirement to share the clinical risk management file with deploying organisations.

5. Recognition that all regulated health and care organisations belong to the ecosystem

Safe AI use depends on organisational readiness and workflow integration; regulation should establish accountability and liability, reflecting how AI is developed, procured, implemented and used in clinical settings.

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

From the outset, pre-market assessment should require suppliers to have a robust mandatory post-market surveillance (PMS) plan beyond incident reporting. Making demonstrable PMS capability a condition of approval, rather than an expectation after deployment, is one of the most tangible changes regulators could make that would directly improve safety without slowing access.

PMS for AI must be shared across manufacturers and healthcare providers. It should have continuous monitoring, with automated metrics, clear up-to-date governance, and escalation thresholds. As AI models can evolve and be repurposed, evidence requirements should move towards staged assurance, including real-world evaluation before scale-up.

1. Explicit update and learning declarations

AI products must disclose if they learn pre-deployment, post-deployment, or not at all. Products updating post-deployment must publish update governance, define re-certification triggers, and ensure clear versioning and labelling. A clear definition of material changes reduces uncertainty for suppliers and health and care deployers.

2.Outcome-focused surveillance for foundation and agentic systems

Foundation models and agentic systems introduce system-level risks that cannot be fully anticipated at approval, especially if the model underpins several tasks. PMS should monitor outcomes, not intentions, and require task-specific evidence and monitoring for each use case, with clear definitions of what constitutes a major change needing further evaluation.

3.Shared accountability for use beyond the approved intended use.

Regulations must cover how AI outputs are described, including instructions and assumptions about data, workflows, and user competence. If an AI system is used beyond its approved intended purpose, responsibility should rest with the party that has both the authority to make that decision and the expertise to understand and manage the associated risks, rather than defaulting to a single stakeholder.

Regulators must mandate:

- Documentation of intended-use boundaries.
- Escalation pathways for divergence.

Shared accountability between manufacturer and provider for known or facilitated repurposing.

4. In-house tools require proportionate oversight.

AI developed in regulated health and care organisations should not be exempt. The same risks apply regardless of provenance. Proportionate assurance. covering validation, monitoring, update control, and incident reporting, should apply to in-house systems and commercial products alike.

5.Build a regulatory regime that reduces burden while increasing safety.

This can achieved through standardised assurance artefacts (e.g., model cards, risk cases, audit logs), sandboxes with clear transition paths to production, and mutual recognition of high-quality international evidence. The aim is not to add further layers, but to align and simplify expectations for easier compliance and greater focus on real-world risk.

6. Secure system-level information sharing

Regulators should mandate secure, standardised mechanisms for cross-site learning, including sharing performance metrics, adverse events, and corrective actions. This should be aggregated across the NHS while protecting confidentiality. Restrictive NDAs that block the sharing of safety-critical information should be discouraged. Effective PMS for AI must be continuous, data-driven, collaborative, and supported by explicit, up-to-date governance and protocols.

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

Insufficient: existing laws are unfit for AI

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

Responsibility must be distributed and explicit, aligned with control and knowledge at each stage of the AI lifecycle. Vague or contractual handovers, especially attempts to shift responsibility solely onto end-clinician users, are unacceptable and unsafe.

1.Manufacturers (including in-house NHS builders) should:

- design and validate AI systems with clear intended uses, known failure modes, and contraindications.
- publish transparent performance claims, including limitations, data prerequisites, and sub-population results.
- provide update governance (versioning, change logs, rollback paths) and support post-market surveillance (monitoring tools, reporting channels, data schemas).

- ensure clinical usability and human-factors engineering to mitigate automation bias and over-reliance, including clear "AI active" signalling and accessible documentation.

2.Healthcare provider organisations should:

- conduct local validation before deployment, confirming performance on local data and workflows.
- ensure competent integration into clinical pathways with role-appropriate training, access controls, and clinical oversight.
- maintain operational monitoring (drift, bias, safety signals), appoint designated AI leads, and implement escalation procedures for adverse events or unexpected behaviour.
- manage repurposing risks: when tools are used beyond intended boundaries, conduct impact assessments, and engage with manufacturers and regulators.

3.Healthcare professionals should:

- use AI within approved indications and trained contexts, exercising clinical judgement and resisting automation bias
- report concerns, near misses, and adverse events through streamlined mechanisms.

4.Shared responsibilities include:

- participating in continuous learning about updates, limitations, and local performance.
- post-market surveillance (data sharing, incident reporting, corrective and preventive actions).
- risk management across the socio-technical system, recognising that safety emerges from the interaction of model, interface, workflow, and human decision-making.

5. Transparency to patients and staff about AI's role in care and operationalising this model, MHRA should:

- require role-specific duties in regulatory submissions and PMS plans.
- promote standardised assurance artefacts (model cards, risk cases, audit logs) to support clear allocation of tasks, and c) encourage service-level expectations for incident response and update communication.

6. Ensuring that staff who use or engage with AI in clinical settings are accountable, competent, and ethical is crucial.

Professional registration of all Digital, Data and Technology staff in the NHS and beyond is necessary. Being on a public registry, such as the BCS, The Chartered Institute for IT's Chartered IT Professional (CITP) validates a person's skills and helps build trust. Professional bodies such as BCS can also provide continuous professional development, career development, and knowledge sharing, for instance, through its members led Health and Care Faculty.

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

The guiding principle is that the party with the most control and best knowledge of a risk carries the primary duty to mitigate and disclose it, while all parties contribute to monitoring and safe operation. Defining how uncertainty is represented in clinical decision support tools is a safety-critical design choice that directly impacts automation bias. Since AI safety depends on interactions among the model, interface, workflow, and human judgment, case-by-case liability apportionment is inevitable. However, predictable default allocations can guide practice and minimise disputes.

Scenario A: AI correct, clinician overrides appropriately, but harm occurs for other reasons. If the AI output was correct and reasonable to override within professional standards (e.g., due to context the AI could not assess), liability typically does not rest with the manufacturer. Responsibility would focus on the clinical decision-making context and system factors.

Scenario B: AI is incorrect, clinician follows it (automation bias). Automation bias is a documented human factors risk. Liability should reflect shared duties:

- manufacturers' are responsible for foreseeable and preventable errors, such as inadequate data, misleading interfaces, lack of uncertainty disclosure, or poor labelling of intended uses.
- health organisations share responsibility if deployment lacks local validation, training, guardrails, or monitoring for drift.
- clinicians are liable when acting outside training or approved use, or ignoring clear warnings
- liability shifts toward the manufacturer if design choices increase automation bias (e.g., authoritative phrasing, hidden changes, or ambiguous indications).

<u>1.Predictable vs. unforeseeable errors</u>

Regulation should distinguish predictable failure modes (e.g., data gaps, sub-population risks, calibration limits) from unforeseeable anomalies. If failure was predictable but not disclosed or mitigated, liability falls on the manufacturer (and possibly the provider). If

failure was unforeseeable despite diligence, liability may be shared, focusing on detection, disclosure, and corrective action.

Safety-critical design choices should be reflected in regulation, not left to local implementation. For example, a system generates probabilistic outputs but presents them as binary decisions or definitive recommendations, thereby increasing automation bias and obscuring uncertainty. Best practice is to make model confidence explicit by using risk scores or probability ranges, with guidance and training.

2.Repurposed or out-of-scope use

When AI is used beyond its intended purpose, liability shifts to the deploying organisation and user, unless the manufacturer encouraged or failed to prevent foreseeable misuse. Clear intended-use boundaries and triggers reduce ambiguity.

3.Version changes and updates

If harm results from an update applied without sufficient testing, notification, or rollback, liability falls on the party controlling the update, usually the manufacturer. Transparent versioning and pre-rollout testing are essential.

4.Operationalising fairness

To minimise disputes, MHRA should publish scenario-based guidance mapping typical cases to default liability, require documented human factors mitigations for automation bias, and ensure new AI norms dovetail with existing guidance.

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

To strengthen AI regulation and accelerate safe NHS adoption, MHRA should draw on the following evidence streams:

1. Safety-critical industries

High-hazard sectors like aviation, nuclear, and rail use mature approaches to system safety, incident reporting, and assurance. Adapting techniques such as safety cases, change control, and just culture learning to healthcare AI can help the NHS avoid pitfalls such as automation bias and mode confusion. Comparative reviews of monitoring and transparent failure reporting would reduce the recurrence of mistakes.

2. Candid evaluation of NHS AI deployments

The UK should curate a repository of independent evaluations, including unsuccessful projects, to understand where benefits were not realised. Both quantitative and qualitative methods should be emphasised, including their impact on clinical decision-

making. Removing restrictive NDAs between testing centres will enable system-level learning from failures.

3. International regulatory innovation

Regulatory experiments such as adaptive approvals, sandboxes, and proportionate oversight of trusted manufacturers can reduce time to value without compromising safety. Alignment with emergent reporting frameworks, for example, model cards or "nutrition labels" for clinical AI, and adoption of standardised assurance artefacts (risk cases, audit logs, change histories) would make evidence more comparable and reduce duplicative effort across applications and sites. Where appropriate, mutual recognition of high-quality international evidence should be considered to limit delays and start-up cash burn.

4. Standards

Relevant standards, such as British Standards for AI governance and ISO/IEC 42001, provide a foundation for compliance. Mapping regulatory expectations to these standards clarifies requirements.

5. Workforce capacity

Professional registration and defined competencies for clinical AI teams will strengthen public trust and ensure ethical, accountable practice. There is a workforce capability and capacity gap, meaning that, despite regulations, people may not be able to comply. Skills gaps prevent many Trusts from progressing beyond small-scale pilots, underscoring the need for a long-term, 10-year digital transformation plan that embeds sustained investment in people and capability. Current national training is fragmented, with no unified curriculum and limited access to funded development. Digital and AI responsibilities are often added informally to existing roles, leading to unclear accountability. Empowering CIOs, CCIOs, and wider digital teams with the authority and resources to lead change is essential to building a confident, future-ready workforce. Establishing structured digital career pathways, supporting staff to access professional registration, and registering digital professionalism will be central to NHS transformation, and BCS has expertise in these areas.

6. Timeliness and service levels

Finally, the MHRA and notified bodies should adopt service-level targets for key touchpoints, such as scientific advice, classification queries, significant change reviews, and approvals. Predictable timelines reduce uncertainty and cash burn for innovators, enabling SMEs to participate without compromising rigour. Together, these evidence sources and process improvements would support a lifecycle-based, system-aware approach to AI regulation.

Who we are

BCS is the UK's Chartered Institute for Information Technology. The purpose of BCS as defined by its Royal Charter is to promote and advance the education and practice of computing for the benefit of the public.

We bring together industry, academics, practitioners, and government to share knowledge, promote new thinking, inform the design of new curricula, shape public policy and inform the public.

As the professional membership and accreditation body for Information Technology we serve over 70,000 members including practitioners, businesses, academics, and students, in the UK and internationally.

We also accredit the computing degree courses in over ninety universities around the UK. As a leading information technology qualification body, we offer a range of widely recognised professional and end-user qualifications.

BCS
The Chartered Institute for IT
3 Newbridge House,
Newbridge Square,
Swindon SN1 1BY
BCS is a registered charity: No 292786