Curia

ukai

AI Is Ready. Is the System?

From Pilot to Infrastructure: Delivering AI at Scale Across the NHS and UK Healthcare

2025–2026

Updated:

Q4 2025: Integrated Care Systems Dual-Index Framework for Innovation Readiness and System Performance Pressure Analysis

This report brings together six months of policy work for UKAI and Curia. This report is editorially independent and does not reflect the view of any single member of either organisation. The parliamentary meetings were kindly hosted by Member of Parliament for Morecambe and Lunesdale, Lizzi Collinge MP, and Shadow Secretary of State for Wales, Shadow Minister (Women), and Member of Parliament for East Grinstead and Uckfield, Mims Davies MP.

Contents

Member of Parliament for Morecambe and Lunesdale, Lizzie Collinge MP hosted the Parliamentary meeting

Forewords

From Innovation to Impact: Delivering AI-enabled Healthcare Across the NHS

It was a pleasure to host a parliamentary discussion on artificial intelligence (AI), healthcare, and life sciences. As a Member of Parliament in the North West with a strong personal and professional connection to the NHS, I was struck not only by the expertise in the room, but by the shared sense of urgency.

The UK has long been a global leader in life sciences. The sector contributes more than £100 billion to our economy and employs hundreds of thousands of people across the country, including in regions such as the North West.[1] Our universities conduct world-class research. Our clinicians are respected internationally. Our health data, collected across the life course, is an extraordinary national asset.

Yet leadership in research does not automatically translate into leadership in delivery.

The central challenge explored in this roundtable is not whether AI can improve healthcare – it is already doing so. AI is supporting earlier cancer detection, improving diagnostic accuracy, reducing administrative burden on clinicians, and enabling new approaches to prevention through genomics and predictive modelling. The question is whether we can embed these advances consistently, safely, and equitably across the NHS.

From my perspective as both a parliamentarian and someone with close family ties to the health service, this is not an abstract policy debate. When innovation works, it means faster diagnoses, better outcomes, and more time for clinicians to focus on patients. When it stalls, it means continued pressure on staff, longer waits, and avoidable inequalities.

Several themes stood out clearly.

First, infrastructure matters. We cannot expect modern digital tools to function effectively in environments constrained by legacy systems and fragmented data. If we are serious about delivering the Secretary of State for Health and Social Care's vision – from analogue to digital, and from reactive to preventative care – then investment in interoperable, secure, and resilient infrastructure must be treated as essential, not optional.

Second, regulation must evolve alongside technology. Safety must always remain paramount. However, we also heard that processes designed for previous generations of medical devices do not always fit adaptive AI systems. Policymakers face a delicate balance: creating flexibility without weakening safeguards. Parliament has a responsibility to ensure that legislation enables innovation responsibly, rather than inadvertently freezing progress.

Third, prevention requires structural reform. Predictive medicine offers the potential to identify risk years before disease develops. Yet funding models remain weighted toward treating illness once it occurs. If we genuinely want to shift toward prevention, our budgetary frameworks and accountability mechanisms must reflect that ambition.

Trust also featured prominently in our discussions. Communities who have historically experienced poorer outcomes or unequal treatment must have confidence that new technologies will not deepen disparities. Transparency around data, inclusion in research, and clear communication are not peripheral considerations. They are central to public legitimacy.

Finally, we must recognise the economic dimension. The UK has the potential to build sovereign capability in AI-enabled healthcare. That requires aligning NHS transformation with industrial growth, ensuring that startups and scaleups can navigate data access, procurement, and regulation without insurmountable barriers. Done correctly, this is not a trade-off between public service and economic growth; it is an opportunity to advance both.

As policymakers, our role is to listen carefully, challenge assumptions and ensure that strategy translates into delivery. The NHS faces immense pressure. Innovation is not a luxury – it is part of ensuring long-term sustainability. However, progress must be thoughtful, inclusive, and accountable.

These discussions demonstrate that there is no shortage of expertise or commitment. What is required now is coordination, and determination to deliver. If we align infrastructure, regulation, and funding with our ambition, the UK can lead not only in research, but in responsible deployment – improving outcomes for patients while strengthening our life sciences ecosystem for the future.

I am grateful to all who contributed to these conversations. The work now is to turn insight into action.

Lizzie Collinge MP
Member of Parliament for
Morecambe and Lunesdale

Chair of Curia's Health, Care and Life Sciences Research Group, Rt Hon Andrew Stephenson CBE chaired both meetings in Parliament

From Pilot to Practice: Closing the Gap Between AI Innovation and NHS Deployment

Across government, the NHS, academia, and industry, there is broad agreement that AI and advanced data science will transform healthcare and life sciences. The science is advancing rapidly. British universities are among the strongest in the world. Our life sciences sector contributes more than £100 billion to the economy and supports hundreds of thousands of skilled jobs.[2] We have longitudinal health records, high GP registration, and a unique ability to link data across the life course. In many respects, we are better positioned than almost any nation to lead.

And yet, as these discussions make clear, the central challenge is no longer invention. It is implementation.

Too often, innovation in the NHS remains confined to pilots. Promising technologies are tested, evaluated and praised – and then stall before reaching routine practice. The reasons are not mysterious. They include fragmented procurement, unclear lines of accountability, data governance complexity, workforce pressures, and, at times, a cultural instinct to equate the status quo with safety.

We must challenge that assumption directly. Doing nothing is not risk-free. In a health system facing rising demand, workforce shortages, and widening inequalities, standing still carries consequences of its own. Innovation, when properly evaluated and safely deployed, is not a threat to patient safety – it is an essential route to improving it.

Throughout these roundtables, three themes emerged repeatedly.

First, system readiness matters as much as technological capability. AI tools do not deploy themselves. They depend on infrastructure, workforce confidence, digital maturity, and governance that is proportionate rather than paralysing. Building the bridge from innovation to adoption requires effort on both sides.

Second, regulatory reform must strike the right balance. We must be safe, but we must also be fast and trusted. The regulatory framework for medical devices was not designed with adaptive AI in mind. We therefore need approaches that allow continuous monitoring, post-market learning, and clear accountability, without creating unnecessary delay. Businesses need clarity, boards need assurance, and patients need confidence.

Third, we must align economic growth with NHS transformation. Startups and scaleups cannot wait two years for data access while their runway expires. Equally, trusts cannot be expected to underwrite infrastructure without national support. If we want Britain to retain sovereign capability in AI-enabled healthcare, we must ensure that procurement pathways, data environments, and funding models enable responsible domestic innovation to scale.

Prevention and predictive medicine are particularly instructive. The science in genomics and risk stratification is advancing rapidly. Yet our funding structures remain weighted toward treating illness rather than anticipating it. If we are serious about shifting from reactive to preventative healthcare, budgetary and accountability frameworks must evolve accordingly.

As Chair of Curia's Health, Care and Life Sciences Research Group, my focus is on turning policy into practice. It is not enough to announce strategies. We must translate ambition into delivery at trust, system, and national level.

The discussions summarised in this report are candid – but practical. They highlight barriers, but they also demonstrate appetite for change. Clinicians want tools that give them more time with patients. Innovators want clarity and partnership. Policymakers want solutions that improve outcomes and support growth.

Our task now is to convert shared insight into coordinated action. If we do so, the UK can become not just a leader in AI research, but a leader in responsible, system-wide deployment – improving patient care, strengthening our healthcare and life sciences sectors, and ensuring the NHS remains sustainable for generations to come.

Rt Hon Andrew Stephenson CBE
Chair, Curia Health, Care and Life Sciences Research Group

UKAI Life Sciences Working Group member, Adama Ibrahim chaired one of the Parliamentary panel sessions

From Strategy to Delivery: Building the Foundations for AI-enabled Healthcare

When we convened these discussions in Parliament, the objective was to move the conversation from pilots to scalable infrastructure, from policy statements to practical delivery.

The UK has not lacked ambition since the publication of the Government's AI Opportunities Action Plan.3 We have committed to AI at scale and articulated a vision for research excellence and preventative healthcare. We have world-class scientists, innovative companies, and a health system rich in data. But ambition alone does not deliver outcomes. The pathway from strategy to routine clinical practice is complex – and, at times, uncomfortable.

AI does not simply automate existing systems. It exposes their imperfections.

Throughout these sessions, we heard repeatedly that the barriers to scale are not rooted in a lack of technology. They lie in fragmented governance, legacy infrastructure, unclear ownership, and cultural hesitation. In secondary care, clinicians face multiple logins and siloed data. In primary care, digital maturity varies by locality. Startups encounter data access processes that outlast their funding cycles. Boards hesitate because liability is not always clearly defined.

None of these issues are insurmountable, but they require coordinated leadership.

One of the strongest messages from industry during the course of this programme was that deployment, not invention, is the bottleneck. Even if no new AI tools were built from tomorrow, it would take years to embed what already exists. That should focus our attention. This means that by the time they are deployed, they may largely become out of date. Measurement, evidence generation, and clear pilot exit criteria are essential. Trust is built through lived experience in real clinical settings, not through PowerPoint presentations.

At the same time, the global environment is moving rapidly. International markets are experimenting with direct-to-patient AI-enabled tools. Investment is flowing toward jurisdictions that provide clarity, interoperability, and predictable regulation. If the UK allows caution to become delay, we risk losing not only economic advantage but also the opportunity to shape standards aligned with our values of fairness, transparency, and safety.

Prevention illustrates the scale of the prize. Genomic risk scoring and predictive analytics can identify vulnerability years before disease manifests. Yet prevention is still funded from the same constrained budgets as acute care. If we are serious about shifting from reactive to preventative medicine, our funding and evaluation models must reflect that intent.

Trust must also remain central. Concerns about bias and representation are legitimate. AI systems learn from data, and if that data fails to reflect the diversity of our population, outcomes will differ. Transparency in training data, subgroup performance monitoring, and meaningful public engagement are not optional extras – they are foundational to confidence.

Within the UKAI Life Sciences Working Group, our role is to bridge sectors. We bring together small- and medium-sized businesses (SMEs), established industry, clinicians, and policymakers to ensure that solutions are co-designed and grounded in real need. Collaboration must replace transaction. Safety case methodology must evolve with adaptive AI. Infrastructure must be treated as a national asset, not a discretionary upgrade.

Above all, culture matters. Fear based environments do not innovate. Curiosity, openness, and shared learning accelerate adoption. We must normalise iterative improvement rather than demanding perfection before first deployment.

This report does not represent the end of a conversation. It marks the beginning of a more coordinated phase of work over the course of the year ahead. The challenge is not whether AI will shape healthcare. It is already doing so. The question is whether we shape that transformation deliberately – in line with British principles of safety, accountability, and equity – or allow it to happen unevenly.

The opportunity is significant. So too is the responsibility.

If we align infrastructure, regulation, funding, and culture, the UK can lead not only in discovery, but in delivery – improving outcomes for patients while strengthening our life sciences ecosystem for the future.

Adama Ibrahim
UKAI Life Sciences Working Group
(Panel Chair)

Executive Summary

This report brings together the insights from two parliamentary roundtables examining the deployment of AI across healthcare and life sciences. The two roundtables were followed by a meeting of experts, which produced a response to The National Commission into the Regulation of AI in Healthcare, run by The Medicines and Healthcare products Regulatory Agency (MHRA) (the submission can be found in Appendix 1).[4]

Jointly convened by independent policy institute, Curia, and trade association for the UK AI economy, UKAI, these discussions assembled government, parliamentarians, NHS leaders, regulators, academics, innovators, and investors to confront a pressing national challenge: how to convert ambition into delivery.

The core question framing both sessions was how do we move from pilots and policy statements to national infrastructure and routine implementation of AI in healthcare?

The UK does not lack scientific leadership. It is among the world's strongest in AI-related clinical trials, with late-stage research close to market deployment. Its life sciences sector contributes billions to the economy and supports hundreds of thousands of skilled jobs. The NHS holds longitudinal, cradle-to-grave health data that most countries cannot replicate. British companies are building clinically grounded, safety-led AI tools that are already demonstrating measurable impact.

Yet scale remains slow and varies within the sector.

The central conclusion of these discussions is that the constraint is not invention, but integration. The UK possesses the science and industrial capability to lead in responsible AI-enabled healthcare. What is slowing progress is systemic friction across governance, infrastructure, funding, and culture.

Six interlocking themes explain this friction.

1. **System readiness** remains inconsistent. Digital maturity varies across NHS trusts and systems. Fragmented procurement, duplicated information governance processes, and legacy infrastructure slow deployment. Pilots often lack clear exit criteria, leading to stagnation rather than scale. Workforce confidence and operational alignment matter as much as the technical capability of AI systems.

Safety Case Methodology

While national regulatory reform provides the overarching framework for AI in healthcare, safe deployment ultimately depends on how technologies are implemented within local health systems. Insights from the Cornwall health system, as set out by Chief Information Officer at Royal Cornwall Hospitals NHS Trust, Kelvyn Hipperson, offered a practical perspective on how safety can be embedded into AI adoption without creating unnecessary bureaucracy or slowing progress.

He highlighted that safety should be re-centred around established digital clinical safety methodologies rather than layered through successive checklists and overlapping governance approvals. By strengthening existing safety frameworks rather than creating parallel processes[15], systems can maintain robust assurance while avoiding procedural complexity that delays adoption.

Re-centring Digital Clinical Safety

Digital clinical safety standards already exist within the NHS and have been developed to assess and mitigate risks associated with health information technology systems. These frameworks provide structured approaches to identifying hazards, assessing risk, and ensuring appropriate mitigation measures are in place before technologies are deployed in clinical environments.

However, as AI tools enter the healthcare landscape, organisations can sometimes default to creating additional governance layers on top of these established processes. While the intention is to ensure safety, the result can be fragmented assurance structures that duplicate effort without strengthening risk management.

The experience in Cornwall suggests a different approach. Rather than establishing parallel approval tracks specifically for AI, systems can strengthen and modernise the clinical safety case methodology already embedded within NHS digital governance. A structured safety case clearly articulates potential hazards, mitigation strategies, lines of accountability, and ongoing monitoring processes. This provides a disciplined and transparent way to evaluate risk without requiring governance structures to be rebuilt for each new technology.

By placing clinical risk management at the centre of deployment, the safety case methodology ensures that patient safety remains integral to decision-making rather than an additional requirement applied late in the process.

Avoiding Checklist Accumulation

A widespread challenge across many NHS organisations is the gradual accumulation of regulatory and governance requirements. Information governance assessments, data protection impact assessments, cybersecurity reviews, medical device compliance checks, equality impact assessments, and local policy approvals can quickly combine into lengthy and duplicative processes.

Each of these components serves a legitimate purpose. However, when they are applied independently without coordination, they can create a fragmented approval landscape that generates delay and uncertainty for both healthcare organisations and technology suppliers.

Published March 2026.
8 Duncannon Street, London, WC2N 4JF.
Email: team@curiauk.com | contact@ukai.co

www.ukai.co

www.curiauk.com