Royal College of Physicians

RCP view on digital and AI

Summary for policymakers

January 2026

The NHS is at a pivotal moment in its digital evolution. With increasing patient demand, workforce pressures and the need to deliver more efficient care, digital systems and AI are potentially powerful tools to support NHS clinicians and improve patient outcomes. But the promise of these technologies will only be fully realised if they are implemented thoughtfully, safely and inclusively, hand in hand with clinicians and with a relentless focus on patient safety.

Successful widespread AI adoption will require robust governance, clear accountability, clinician involvement and a commitment to equity and transparency. Crucially, AI must support – not replace – clinical judgement and solve real-world clinical challenges rather than being led by technological possibility. Clinical safety must be at the heart of AI development.

1 Digital

The future of healthcare depends on digital clinical systems that support clinicians to deliver safe care without adding extra burdens or risks. Digital tools have the potential to transform the NHS, highlighting patients at risk and reducing variation by enabling triage and supporting clinical decision making based on best practice guidelines. They are also the foundation for safely and successfully integrating AI into the NHS.

But a snapshot survey of RCP members conducted in June 2025 found that 68% of 548 respondents either somewhat (20%) or strongly (48%) disagreed that the NHS has the right digital infrastructure to support widespread introduction of AI that will make a difference.

Shifting from analogue to digital must include optimising existing care pathways and digital systems in the NHS. Fixing hardware and investing in infrastructure are critical but these alone will not deliver the shift. It should also mean building intuitive software, interoperable systems, complete datasets and a commitment to digital inclusion.

What do we mean by digital?

When we talk about digital healthcare, we mean a world where data and information relating to patients, staff and equipment can be stored and accessed on digital systems. Patient-level data are available to patients, clinicians and operational managers to organise and deliver care. Data that are entered in routine clinical practice allow analysis and interpretation to provide knowledge for improvement, audit and research. This would mean that we realise Tom Loosemore's definition of digital as 'applying the culture, processes, business models and technologies of the internet-era to respond to people's raised expectations'.

How analogue is the NHS?

Electronic patient record (EPR)

The core digital system at the heart of daily working for most clinicians is the EPR. For each patient, this should contain a complete longitudinal record of their full healthcare history, setting out a single version of the truth. But currently, most patients will have data in multiple EPRs, even within one organisation. Primary care EPR data are transmitted poorly into the secondary care record and vice versa. Significant manual effort is required to make sure that information is uploaded into the correct patient record.

The majority of acute NHS trusts have adopted digital systems, but around 6% still rely on paper-based records and lack an EPR system, despite commitment from government that all trusts should have a functioning EPR by 2025.

All clinicians should have access to digital information about their patients, but simply having an EPR in every trust will not deliver the digital ambitions of the 10 Year Health Plan. EPR systems vary widely across NHS trusts, leading to inefficiencies, safety risks, clinician frustration and unproductive wasted time. Inconsistent data formats and fragmented systems make it hard for clinicians to access and share vital patient information.

We must optimise existing digital systems to function as effectively as possible and to share data with other systems. RCP members have shared numerous stories of poorly functioning IT, including one hospital still using Windows 7 as its operating system, which stopped receiving technical updates and security support from Microsoft in 2020.

What are the impacts of poor usability?

Patient safety

The patient safety risks of poorly designed digital systems are relatively new and badly described. Developers and users often overlook the new or unique safety risks that digital systems introduce, especially if the risks are not present in paper-based approaches.

The current focus on digital maturity does not factor in the change in risks to patient safety between paper and digital, often due to fragmented implementation and limited co-production with clinicians of digital tools.

The National Early Warning Score (NEWS) chart is a UK national standard tool originally developed by the RCP that is fundamental to the detection of and response to clinical deterioration in adult patients. The digital NEWS chart eliminates calculation mistakes, but the lack of an agreed format in the NHS means that it loses the benefit of the standardised paper chart in clearly demonstrating deterioration.

The EPR is another example. Without standardisation in the EPR for structured data entry, each clinician can record patient information differently and for different purposes. This lack of standardised data entry means that data cannot be easily shared between systems, making it harder for clinicians to get a full picture of a patient's condition. Workarounds, such as copying from one note to the next, can lead to inaccuracies in the record and, in some cases, patient harm. Systems should enable clinicians to document easily, quickly and in a standardised way without the need for these potentially risky workarounds, underlining the importance of both user-centric design to improve the EPR and for clear national standards for EPRs deployed in the NHS.

Clinician experience

Poor usability adds to the operational strain on clinicians. There is an assumption that digitising workflows will deliver quicker, automated and/or more efficient working, but poor levels of usability in digital systems mean that this is often not the case. The adoption of an EPR by NHS trusts can often reduce productivity, as health professionals have to spend time navigating poorly designed systems to complete tasks.

The implementation of technology in the NHS has also led to 'task shifting', where routine tasks that were previously carried out by administrative roles have been transferred to clinicians to undertake alongside their clinical roles.

We welcome the commitments in the 10 Year Health Plan to introduce single sign-on (SSO) for NHS software to remove duplication. Small changes like this could make a big difference to the working lives of doctors.

Usability

Since 2021, NHS England has partnered with KLAS Research and Ethical Healthcare Consulting on two usability surveys on the EPR. The first survey, conducted in 2021–22, found that implementation of the EPR was more important than its functionality. Around two-thirds of user experience is dependent on how a system is implemented (and only one-third is associated with the particular EPR). This includes clinician involvement in design, implementation and the ability to iteratively improve workflows, as well as the extent to which organisations invest in building users' knowledge, skills and confidence with the technology. The second 2024–25 survey found that an organisation's ability to provide a stable, available and fast system was the foundation to higher EPR user satisfaction.

Clinical input into the procurement, design and development of clinical digital systems is key to their functionality. The NHS design principles emphasise the need for engagement with patients and staff: 'put people at the heart of everything you do'. Clinicians need to be able to report workflow and patient safety issues, and organisations need to be able to respond with improvements to the digital process that improve care.

Users of digital systems must be able to add their own 'micro-personalisation' – designing processes that allow rapid completion of tasks which they undertake frequently.

Usability can also be improved with training. The 2024–25 Ethical Health Consulting and KLAS Research EPR usability survey found that 60% of clinicians wanted more education on the EPR, with 44% reporting that they had received no ongoing EPR education. They found that the ideal training package would be 3–5 hours of initial training, followed by 1–2 hours of annual training post-implementation.

It is important that all clinicians feel competent and confident in using complex systems. Organisations should routinely analyse how staff are using their digital systems, to offer tailored additional training to those who are struggling with the system or not using it optimally.

Hardware

Hardware in the NHS is often out of date, broken or poorly designed for the task. This is a big source of frustration and burnout for physicians, increasing the time taken to perform basic tasks and making patient care harder. In the 2025 'Focus on physicians' survey of UK consultant physicians, when asked about issues negatively affecting wellbeing at work, poorly functioning IT equipment was the second most common response (44%).

We need to get the basics right. This means functioning computers, laptops and other hardware, working Wi-Fi across NHS estates and an ability for all digital devices to communicate with the EPR, for example machines taking observations, electrocardiograms (ECGs) and intravenous (IV) pumps. It also includes background infrastructure (such as servers) to ensure quick system response times, that relevant data can be stored and shared, and that the system is able to complete clinical tasks such as outbounding letters.

Thought must be given to the hardware needed to deliver the vision in the 10 Year Health Plan. Complex noisy environments, such as the emergency department and the ward, are likely to require mobile computers, handheld devices and specialised microphones to allow the use of ambient voice technology (AVT). Systems are 'on' 24/7, so robust kit and battery life are fundamental to good care.

There is also a need to consider and define the optimal hardware required for different clinical environments, from ward rounds and digitised hospital at home services to delivering care in patients' homes.

Systems that work well together

Implementing unified digital platforms, such as the NHS app and the proposed single patient record – where patient records, test results and appointments are centralised – would simplify access to relevant information for patients and healthcare providers.

The single patient record announced in the 10 Year Health Plan promises to bring together data from multiple sources including the EPR, personal health data supported by the Federated Data Platform (FDP), to act as a 'patient passport' that will make care more seamless. It promises an interoperable dataset that brings together all patient data in one place, which can be accessed anywhere in the health system. Over time, it is planned that the data included in these records will expand to include a personalised account of health risk, by drawing on lifestyle demographic and genomic data. If realised, this would create an interoperable database of citizen healthcare data that can be integrated into, rather than replacing, the EPR.

However, the RCP believes that there is an optimism bias in the rhetoric in the 10 Year Health Plan on this. The ambition to deliver fully personalised care – drawing on information from genomics, existing healthcare data, and lifestyle data (including data from wearables) – is praiseworthy, but there are risks around public confidence in data sharing and being able to use the insights that the data may deliver. Success should be built on learning from existing shared care records.

Improving the interoperability of data between different EPRs to allow clinicians and patients to see a seamless record of care is vital to enable healthcare to shift closer to the patient and to shift from treatment to prevention.

Setting standards

We recommend clinical digital standards across the following areas:

EPR suppliers: a standard for model content for an NHS EPR that adheres to NHS clinical and operational guidelines. This should include:

- minimum expectations for the configuration of the EPR to reduce variability, including templates for letters and for the visualisation of results – for example, there should be a standardised direction for timelines of observations and results
- standardised visual formatting of observations charts (such as previous paper NEWS chart).

Digital systems beyond the EPR: Standards for the procurement of digital systems, linking to standards for suppliers to guide how and which systems are purchased for use in the NHS.

Data interoperability: There should be interoperable data standards for clinical and operational data to support the single patient record.

Health apps used by patients: A standard for the minimum set of evidence for safety and efficacy that patient-facing apps must meet for clinicians to recommend them to patients.

Standards for digital systems in the NHS will decrease cognitive load for clinicians and improve patient safety. The ability to do this at scale will determine how effectively clinicians will be able to use AI and digital decision-making support in practice.

The currently poor interoperability, particularly between primary and secondary care systems, means that few digital or AI systems are able to process or display a complete longitudinal record of a patient's health. Having this standardised set of requirements in place means that all NHS trusts generate the same data, allowing them to draw better conclusions about services to allocate resource more effectively.

System resilience

As NHS processes become increasingly digitised and the availability of data continues to grow, cybersafety and cyber-resilience are of paramount importance, and cyber-resilience standards should be embedded across NHS digital systems.

The NHS is national infrastructure, and a reliance on digital or AI systems without cybersecurity and robust protections poses a significant risk. NHS trusts, clinicians and suppliers need to be aware of these risks and their important role in work to mitigate them, and cybersecurity for NHS systems must be a priority for government and the NHS in the analogue to digital shift. Without secure systems that are resilient to outside threats, and robust contingency plans, the increasing digitisation of the NHS will pose an ongoing risk to patient safety.

Patient involvement and digital literacy

Digital exclusion often correlates with social exclusion and those who are in more vulnerable groups. The Good Things Foundation has found that 7.9 million people lack basic digital skills and, of those, 69 % have a disability or impairment, 47 % have no basic qualifications and 77 % are over the age of 65. 5 % of the UK population lack access to the internet, and studies show that a significant proportion of the population (31 % of UK adults) don't access health services online.

Organisations should follow a 'digital plus' rather than a 'digital only' approach or 'digital by default' assumption, recognising that, even with support, digital solutions will not work for the entire patient population. Other routes of access must remain available to patients – and a 'digital plus' approach allows a greater focus on supporting people who can't use digital systems through other approaches that work best for them.

The use of the NHS app as a single digital front door to the NHS should make it easier for patients to access their health records and manage their conditions. We welcome approaches that support patients to understand and manage their own health, whatever their level of health literacy. Recognising the importance of, and barriers to, health literacy that exist for patients – whether they are accessing information online or not – is critical for the shifts to digital and community. Access to digital tools will be insufficient if health literacy is an issue. Action needs to be taken to improve both health and digital literacy across the population, and clinicians will need to be confident in discussing health misinformation with their patients.

2 AI

AI, and in particular machine learning, has been part of healthcare for decades. Recent breakthroughs, especially in large language models (LLMs) and generative AI (genAI), have significantly expanded the possibilities for AI in the NHS, with the potential to support clinical decision making, enhance administrative workflows and, crucially, improve patient safety, diagnosis, disease management and patient experience.

A challenge is avoiding optimism bias, so we are realistic about AI's potential, and confront technical, clinical, ethical and regulatory barriers to effective implementation of high-quality AI in a way that makes a tangible difference to clinicians and clinical care.

An overarching strategy for how AI can support the NHS and improve health is needed. A clear rationale for the adoption of AI in the NHS is required to form the basis of all development, alongside a set of standards for AI developers as a prerequisite to be used in the NHS that includes how suppliers can ensure the safe use of data, as well as how the AI tools work and the data that they are trained on. It should also set out patient safety thresholds, guidance on the use of AI in clinical practice that is not organisationally led, and how the infrastructure is going to be created to make the vision for AI possible.

To ensure that AI delivers meaningful impact in the NHS, we must avoid simply chasing emerging innovative technologies and instead focus on optimising systems and thoughtfully integrating new solutions into current clinical and operational pathways.

Nearly one-third of the 571 UK physician respondents reported using AI tools in their clinical practice either every day (16%) or weekly (15%). One-third (33%) said that they use it rarely and 30% never use it.

70% of physicians said that they were either very (29%) or somewhat (41%) supportive of AI tools being implemented widely in the NHS. Much like the advent of digital, the presence of AI in the NHS is varied, concentrated in pockets where particular trusts or engaged clinicians have spearheaded initiatives. While the absence of a centralised overall vision for AI in the NHS may stimulate local innovation, it risks variation and conflicting approaches, in turn mimicking the problems that we now see – with different digital systems used across the NHS that are not interoperable and cannot enable systematic processes for patient care. Thought should be given to an NHS 'approve and scale' model that would encourage local innovation while ensuring the right safeguards to prevent pockets of variation and conflicting approaches between trusts. Such a model would allow the NHS to provide robust evaluation locally and provide approved tools or systems with a standardised route for national scale-up. Again, greater standardisation of the EPR models used in each trust will be critical to enable the wider deployment of useful digital tools.

Supporting, not replacing, our workforce

AI has often been posed as a 'silver bullet' to improve productivity. When asked to select up to three biggest benefits of using AI in clinical practice, reduced admin burden (69%), time savings (62%) and improved diagnostic accuracy (34%) were most commonly cited by UK physicians.

The existence of, or access to, AI tools alone will not improve productivity. AI tools that are poorly designed, poorly implemented and have poor usability risk making clinicians less productive.

There certainly is an opportunity for AI to reduce the time taken to complete some tasks by making EPRs more easily searchable, summarising notes for discharge summaries, creating letters by listening to consultations, or automating processes such as appointment booking/scheduling and doctors' rotas. AVT could give doctors more time for meaningful engagement with patients, allowing clinicians to capture a conversation rather than spending the majority of an appointment making notes. Technology facilitating interactions that feel more human could significantly improve patient experience – but as we digitise healthcare tasks, we need to recognise the full spectrum of 'purpose' that exists in paper or analogue forms. For example, recording information is not the sole function of note taking; it is part of clinicians' thinking and considering a patient's symptoms. It is also unlikely that we will achieve the full potential of AVT until it is fully integrated with the EPR and able to act as a full agentic AI.

Productivity gains from AI should enable doctors to deliver vital supervision, education and service transformation, or to undertake portfolio projects across areas such as clinical research. Being able to deliver more patient-facing and professional development activities will contribute to improved patient care and job satisfaction.

The 10 Year Health Plan says that AI and technology will mean that 'world-class care can be delivered without inexorable growth in staffing numbers' as 'evidence shows as much as 60% of what an individual NHS staff member does can be freed up by technology'. AI should not be seen as a complete solution to solving staffing pressures. Technologies, including AI, being leveraged to free up doctors to use their unique skillsets to deliver care that only they can provide would be welcome. But capacity issues are unlikely to be resolved by technology alone, and we need to be realistic about what the technologies are capable of.

Getting the right governance and regulation

Robust governance and regulation are essential to ensure the safe, ethical and effective use of AI in the NHS.

The 10 Year Health Plan announced a new regulatory framework for medical devices including AI to be published in 2026, alongside an NHS AI strategic roadmap 'that will enable clear ethical and governance frameworks for AI'. The National Commission into the Regulation of AI in Healthcare, a non-statutory advisory body established by the Medicines and Healthcare products Regulatory Agency (MHRA), was launched in September 2025.

Respondents to the RCP's June 2025 snapshot survey expressed a strong sentiment that more robust regulation is needed on AI. When asked about main barriers to the deployment of clinical AI systems in the NHS, 36% of respondents noted a lack of regulation, double the proportion of respondents (18%) who said too much regulation.

One area where there is clear need for stronger regulation is AVT. NHS England's national chief clinical information officer wrote to NHS organisations in June 2025 to clarify guidance on the use of AVT tools, instructing trusts and individuals to stop implementation of any non-NHS-compliant solutions. Some of the most widely available and widely used AVT systems do not currently integrate with NHS clinical systems. Where AI tools are not fully integrated, we will not realise the full benefits of what AI has to offer for patient care, clinician experience and productivity.

AVTs are a prime example of how complex the regulation of AI devices in the NHS can be. All ambient scribes must be registered as an MHRA Class I device and those that provide clinical decision-making support are likely to be classified as MHRA Class II devices, with increased regulatory requirements. Regulation of a device itself sits with the MHRA; organisational use is regulated by the Care Quality Commission (CQC); individual use and its safety sit with the General Medical Council (GMC); and data security sits with the Information Commissioner's Office (ICO). Likewise, tools such as ChatGPT are not regulated for use in healthcare, but using ChatGPT for health purposes makes it a medical device.

Liability needs to be distributed proportionally across the various actors involved, including vendors, purchasers and users, and governance and regulation should set this out. But it is also critical that the NHS can provide central oversight and guidance on the regulation of AI systems. Clinicians will need to be protected by their organisation and the wider NHS from taking on increased liability from AI tools.

Engagement

Meaningful iterative development with clinicians and patients is essential throughout design, testing and implementation. AI tools must be developed with a purpose, and clinicians are uniquely positioned to highlight on-the-ground issues and to work directly with AI developers to design tools that are grounded in clinical practice.

Clinician expertise and input mean that AI tools are more likely to integrate properly into existing workflows and decision-making processes, rather than needing to retrofit and add additional steps for clinicians. In the RCP's June 2025 snapshot survey, when asked about barriers to the deployment of clinical AI systems in the NHS, 70% of UK physician respondents identified the inability to integrate AI tools with other system such as the EPR, and 65% said poor interoperability of systems. It is critical that time is prioritised for doctors to be engaged in digital tool development and be part of iterative testing.

A patient perspective in the development of AI tools is essential, building trust and understanding about how and why the tools can support better care.

An approach of learning from failure and thorough implementation processes, where there is a focus on developing, testing, adapting and trying again to get things right, is critical, rather than rushing to scale. AI tools need to go through real-world testing and evaluation beyond validation to be effective and safe.

Liability and explainability

73% of responding physicians reported that their biggest concern about using AI in their clinical practice was the risk of error. Respondents were next most concerned about liability risks (54%), the risk of de-skilling clinicians (52%), risk of model drift (meaning that the AI algorithm changes over time, 48%), risk of bias (48%) and explainability risks (meaning that it's not possible to know how the AI produces its output, 47%).

It should not be assumed that human oversight will always catch errors made by an AI system. Human–AI interactions are complex and something that we are still learning about, and experience (clinically and of the AI) and workload can increase margins of error. The use of AI to support clinical decision making raises questions for liability when the AI is incorrect. Many LLMs rely on training data that are intentionally undisclosed, ambiguous or commercially protected, which can make it more difficult to align the use of LLMs with the principles of evidence-based medicine. Clinicians should not be held liable for decisions made by algorithms that they cannot fully understand or interrogate.

Doctors are trained to hold risk and liability for their decisions. In an AI context, this means that clinicians need to have meaningful understanding and control of the full decision-making process, including where AI has been used, so that they have appropriate information from the AI system and understanding of where that information and data originated.

The years of training that doctors undertake equip them to make the best judgement decisions based on clinical evidence; AI makes recommendations based on the evidence, guidelines or datasets available to it. An AI system's recommendations may be wrong or even pose a risk to patient safety – but in practice, they may also sometimes be suboptimal, meaning that the AI is not recommending the best possible patient care. There is not always an obvious answer to complex clinical questions. Clinicians should not just defer to an AI's output.

Given that diagnostic expertise is built during training, the increasing widespread use of diagnostic interpretation raises questions about how we best support resident doctors to both harness AI and develop expertise to make their own assessments independent of it. Prioritising high-quality medical training for resident doctors is essential.

Clear lines of responsibility and robust governance frameworks are needed so that clinicians do not become 'liability sinks', absorbing all responsibility for patient harm even when an AI system is the major contributing cause. Explicitly addressing the challenge of liability must be part of future regulations.

Use of personal AI tools not provided by the NHS

In response to the RCP's June 2025 snapshot survey, almost seven in 10 (69%) of the 305 physician respondents said that they were using personal access to genAI tools like ChatGPT and Microsoft Copilot for clinical questions. 15% said that they were using a medical-specific AI tool for diagnosis and 21% were using a personal ambient AI tool.

These findings suggest that the NHS is not moving quickly enough to provide clinicians with AI tools that are useful, efficient and safe. This is a risk and must be recognised as such. Widely available AI tools such as ChatGPT are not designed or regulated for use in healthcare, and using them in this way comes with risk. Doctors and NHS organisations need clear guidance on what AI tools are safe to use in healthcare, along with more agile NHS procurement processes to bring approved and effective technologies into their organisations quickly. NHS organisations must also educate clinicians about how AI tools work and the data that they are trained on, so they understand the potential risks of using non-NHS-approved AI tools to clinical and patient safety.

AI will be most beneficial when it is designed to support, not to replace, clinical judgement. It should assist clinicians in providing better patient care, but the complexities and nuances of providing care mean that a human-centred, empathetic approach will always be needed. By putting clinical safety at the centre of AI development, the NHS can harness the potential of these technologies while maintaining trust, accountability and quality of care.

Data

Standardised approaches to data access are currently lacking across the NHS. Each system collects slightly different datapoints to generate datasets, making it difficult to develop and deploy AI systems across NHS organisations. We need to collect the right data, make sure that data are available across routine clinical practice and ensure that those data are interoperable.

Many datasets lack diversity or exclude certain subsets of the population, particularly certain ethnic or socioeconomic groups. Populations with limited access to digital technologies are also under-represented in the data used to develop AI tools. Incomplete datasets can lead to biased algorithms. Having unrepresentative models and algorithms can therefore limit the effectiveness of health interventions for groups that are

under-represented in the data, such as minoritised ethnic groups or women, further widening health inequalities.

There is no such thing as a completely unbiased dataset, as existing biases are already present in the data that we have. Trying to mitigate this as much as possible, including educating clinicians to recognise this and understand the limitations of datasets more widely, is key to delivering equitable healthcare.

We also need to be clear about what data are not safe to share with AI during training. Data previously considered to be anonymous are becoming increasingly identifiable. Clinicians need to confidently understand which data that are considered anonymous today may 'function as a fingerprint' in future, so they can make judgements about what they should and shouldn't share with AI.

Building the digital medical leaders and workforce of the future

In our June 2025 snapshot survey of RCP members, 31% of respondents said that they were somewhat confident in using AI tools in clinical practice, 24% were not at all confident, 17% were neither confident nor unconfident, and 15% were somewhat unconfident. Only 8% of respondents said that they were very confident in using AI tools in their clinical roles.

That survey also found that a considerable majority said that they need training in clinical AI tools, but a majority do not have access to it.

There is clearly a need for education and training to build knowledge, skills and confidence in using clinical AI tools. The foundations for building an AI-confident medical workforce and leaders of the future are in understanding digital healthcare. Doctors' education and training must reflect the digitised NHS that they are already learning and working in. This is vital if clinicians are to be able to drive the development of clinical systems and be leaders in a digital- and AI-enabled NHS. The 10 Year Health Plan commitment to reform curricula to include comprehensive training in AI and digital tools is therefore welcome.

NHS clinicians must be offered training so that they have the confidence and capability to use digital and AI systems. This training must go beyond how to use AI tools, and ensure that clinicians have a comprehensive understanding of an AI tool's intended purpose, its limitations and its potential biases. Given the rapid development of these tools, training and education on AI need to be iterative and embedded at all career stages.

Population health and health inequalities

When designed properly, AI can actively help to reduce health inequalities. AI tools can make engagement with the healthcare system easier, for example, by tailoring communications with patients based on factors such as their literacy levels.

AI is a disruptive technology that is likely to democratise healthcare by widening the group of people who are able to access what was previously specialised knowledge, including patients' knowledge about their own health. A poll commissioned by Healthwatch England published in November 2025 found that around one in five men under the age of 35 is likely to use ChatGPT or another AI tool to find out about health conditions or check symptoms. For women of the same age, it was around 10–15%. This use of AI by patients will require changes in the way that doctors work and communicate with them. The expertise and experience of senior doctors will continue to be vital.

AI can also help to identify underserved communities or predict individuals at higher risk of poor health outcomes. AI tools have already been used to flag patients likely to miss appointments based on past behaviour or socioeconomic barriers, enabling targeted interventions such as personalised reminders or transport support. This, in turn, reduces healthcare inequalities by improving access to services.

Recommendations

1. **Government and the NHS must invest in well-functioning digital infrastructure and up-to-date IT systems, so that clinicians have access to digital tools that work. This should include investment in the improvement and optimisation of digital systems, data and the electronic patient record (EPR).**

Investment is needed in digital infrastructure, including hardware, software and connectivity, and in staff capability at both technical and leadership levels. This will ensure that organisations have the people and tools that they need to optimise their digital systems. Outdated systems need to be upgraded, alongside ensuring usability and speed, to avoid adding to clinician workload. The analogue to digital shift depends on realising the full potential of all digital systems. Without prioritising the optimisation of existing digital systems, the NHS will continue to fail to meet basic digital requirements and will be unable to deliver the ambitions of the 10 Year Health Plan. By getting the basics right, the NHS can create a digital foundation that supports safe care, improves clinician productivity and enables future innovation, including the use of AI.

2. **The NHS should set an EPR model content specification standard that EPR providers must meet to ensure that their products meet NHS requirements.**

There is variability in each EPR across NHS trusts, meaning that even if trusts have an EPR from the same provider, their functionalities and appearance can be drastically different. It also means that trusts currently have to pay each time to ensure that the design of the EPR meets NHS requirements. Implementing a model content specification in secondary care trusts, in the way that robust standards have improved convergence in primary care EPRs, would address this and resolve issues with functionality and appearance, which negatively impact usability for clinicians and slow down the rate at which care can be provided. The NHS model should include minimum expectations for EPR configuration, and how suppliers collect and use data and standards for elements like timelines, results visualisation and letter templates. The EPR should be designed to underpin AI decision support, automation and data sharing, and AI tools must integrate well with EPRs.

3. **The NHS must establish robust clinical national standards for the procurement of digital systems and data interoperability.**

Without clear procurement standards, NHS trusts may adopt systems that are incompatible, hard to use, or fail to meet clinical needs. There must be a requirement to demonstrate the clinical safety of digital systems and devices. Where clinical risk is caused by systems or devices, there should be a system of national reporting that can inform procurement standards. Standards should also include a requirement for interoperability that allows data to transfer into and out of the EPR and between records, allowing structured data to land in the right place in the record (such as Message Exchange for Social Care and Health (MESH)).

4. **The Department of Health and Social Care (DHSC) must develop standards for how data in the NHS are recorded and create complete, standardised, accurate databases to ensure that data are usable, consistent, secure and representative.**

Understanding real-life patient data and service use should be key to service transformation. Datasets in the NHS are often siloed, fragmented, inconsistent or incomplete. Having the right data that can be integrated into algorithms and digital systems will largely determine the usefulness and accuracy of digital and AI tools. DHSC needs to provide national oversight to standardise how data are collected, formatted and shared in the NHS; to ensure that datasets are accurate, secure and representative of their populations; and to reduce duplication or data silos. Optimising the usability of digital systems in the NHS will also be key to ensuring that structured data can be shared easily between systems, to allow complete clinical records in each EPR and appropriate prioritisation of patients.

5. **The DHSC and the NHS should establish central banks of NHS-approved algorithms, AI tools and patient-facing apps that meet national standards.**

A central repository of NHS-approved digital tools would ensure that only safe technologies are used in clinical practice, support equitable access across organisations, reduce duplication, and give clinicians confidence in using AI tools that have demonstrated positive outcomes in NHS settings. The bank must be regularly updated and accessible to all NHS organisations. A similar system for patient-facing apps would ensure that clinicians feel comfortable recommending them to patients. Clinicians have the same duty in recommending health apps as they do when prescribing medication, but are much less well prepared and informed. Patient-facing apps should have to demonstrate clinical effectiveness and ease of use for most of the population.

6. **NHS organisations should follow the NHS design principles for all digital transformation, including AI tools – prioritising user experience, and engaging with clinicians and patients from the outset to ensure that digital and AI solutions address real-world challenges, improve clinical workflows and experiences, and support safe, patient-centred care.**

Digital systems that are designed with clinicians and patients in mind are safer, more efficient and easier to use. Clinicians bring essential insights to patient care, system pressures and practical challenges that developers may otherwise be unaware of or overlook. Involving them from the outset will ensure that tools are designed to solve real-world clinical problems and meet real-world clinical need, rather than being led by technical possibility. It will also ensure that new tools integrate smoothly into existing workflows, and identify important contextual system pressures and practical challenges that developers may otherwise overlook. Clinicians must be given time to engage in digital transformation. Meaningful engagement with clinicians and patients will help to foster trust, improve adoption and, ultimately, lead to better outcomes for staff and patients.

7. **To deliver digital clinical leaders of the future, the government must meaningfully engage and work in partnership with medical royal colleges on its reforms to medical curricula, to include competencies and teaching on digital and AI, alongside embedding digital and AI competencies for NHS clinicians at all career stages in continuing professional development (CPD).**

Education must train clinicians to work in and lead the digitised, AI-enabled NHS that the government hopes to create. The new AI competencies developed as part of the promised updates to curricula must support the development of the digital clinical leaders of the future. The curriculum refresh should aim to translate existing competencies that either already are, or increasingly will be, delivered digitally into the digital skills that doctors will need to deliver modern medicine. It needs to cover how AI algorithms function, including limitations, explainability and potential biases, so that clinicians can understand how AI clinical decision support systems work, and how to use and respond to them appropriately. It should also include regulation and what this means for clinical accountability and patient safety, alongside teaching on data safety. Training and education should include foundational literacy, clinical application and ethical awareness, to ensure safe and effective use of AI tools in clinical practice.

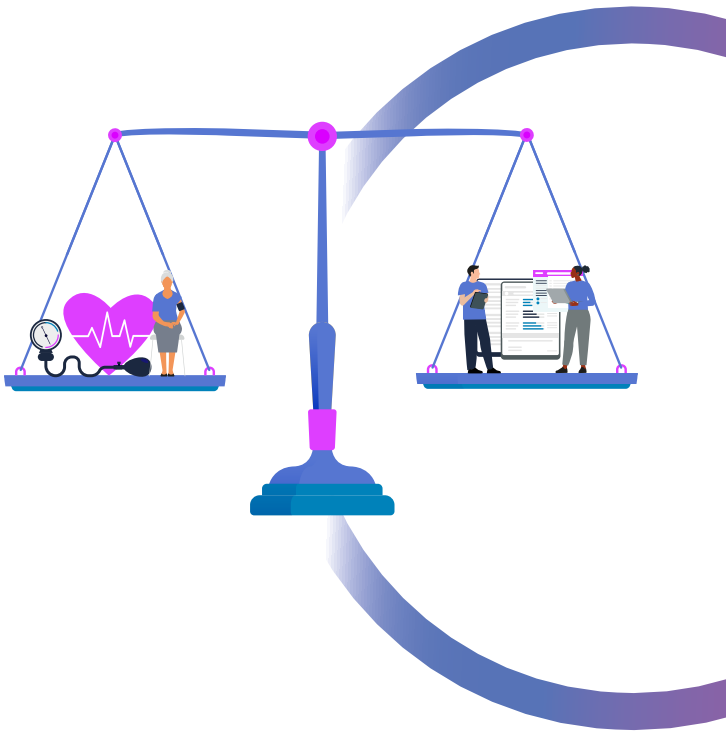

8. **The government's promised 'roadmap for AI in the NHS' must set out a plan for effective and ethical implementation of AI in the NHS, including how AI can enable clinical research and tackle health inequalities by actively improving equity of access, experience and outcomes in the health service. DHSC must consult with patients and doctors to develop its roadmap, particularly those from or working in deprived or underserved communities.**

The government's promised roadmap for AI must set out a coherent implementation plan for how AI will be used in the NHS, and why. It should seek to avoid conflicting approaches and different infrastructure across systems, and instead facilitate and prize interoperability, shared learnings and innovation. It should incentivise the co-design, development and deployment of AI tools with clinicians and patients, including conversations about informed consent for use of AI in patient care. It must identify how AI will be used to reduce healthcare inequalities: AI can help to identify or predict groups at higher risk of poor health outcomes, who therefore may benefit from targeted interventions such as personalised reminders or transport support for appointments. The role of AI in improving clinical research must also be addressed, from identifying eligible participants for trials and analysing large datasets to speeding up processes to enable clinical trials to get off the ground more quickly and efficiently. Consultation with clinicians and patients will be key, especially those from or working in deprived or underserved communities.

9. **Government and the NHS must deliver robust and joined-up regulatory frameworks that put the necessary safeguards in place to ensure the safe and ethical use of digital clinical systems and AI.**

Robust regulation is essential to ensuring that digital and AI tools are clinically safe. The National Commission into the Regulation of AI in Healthcare recommendations for a new regulatory framework should ensure sufficient safeguards to ensure that all AI use in healthcare is safe, going beyond technical standards to address clinical accountability. Regulation should require transparency from AI developers about how algorithms work and the data that they're trained on; clear guidance for clinicians on when and how to utilise AI outputs, retaining clinical judgement; and ongoing monitoring and re-evaluation of AI systems to ensure that they remain fit for purpose as technologies evolve. As far as is possible, regulation should take a principles-based approach, with the aim of covering future developments in technology so that it is not constantly playing catch-up.

10. **NHS organisations and application providers must develop strong governance and safety mechanisms, including collecting and responding to safety incidents, to mitigate risk and ensure the privacy of patient data in AI systems.**

Digital systems can introduce new and poorly understood risks to patient safety, such as miscommunication, data fragmentation and over-reliance on automated outputs, which are often overlooked in system design and implementation. Governance mechanisms need to be put in place to create standardised processes to mitigate patient risk, including systematically collecting and analysing safety incidents, sharing findings across trusts and suppliers to inform safer design and procurement practices, creating feedback loops for clinicians to report potential risks and improve usability, and developing testing mechanisms to mitigate risks before deployment. Investigations into patient safety events must look critically at where digital systems and tasks create risk, and the best ways of mitigating this. This learning should be widely shared and used to develop procurement standards and robust regulation. The combination of iterative learning, robust evaluation and appropriate regulation will mitigate the risks to patient and clinician safety. NHS organisations need to have the capacity and skills to carry out this evaluation at scale. These mechanisms should also ensure the safe use of patient data in AI systems, with clear measures and internal communications on data sharing, and transparency and communication with patients about the use of AI in healthcare, including clarification that patients retain ownership of their data.

Royal College of Physicians

Published January 2026

Development of the RCP view on digital and AI was led by the RCP's digital health clinical lead Dr Anne Kinderlerer. The report draws on the findings of a June 2025 RCP all-member snapshot survey, insights from a group of physician digital experts and a focus group on AI, and input from the clinical vice president and academic vice president. The focus group on AI brought together a range of AI experts, including physicians, the RCP's clinical vice president, special adviser on population health, the Royal College of Radiologists and patient representatives. The report was approved by RCP Council prior to publication.

Contact: policy@rcp.ac.uk