Royal College of Physicians

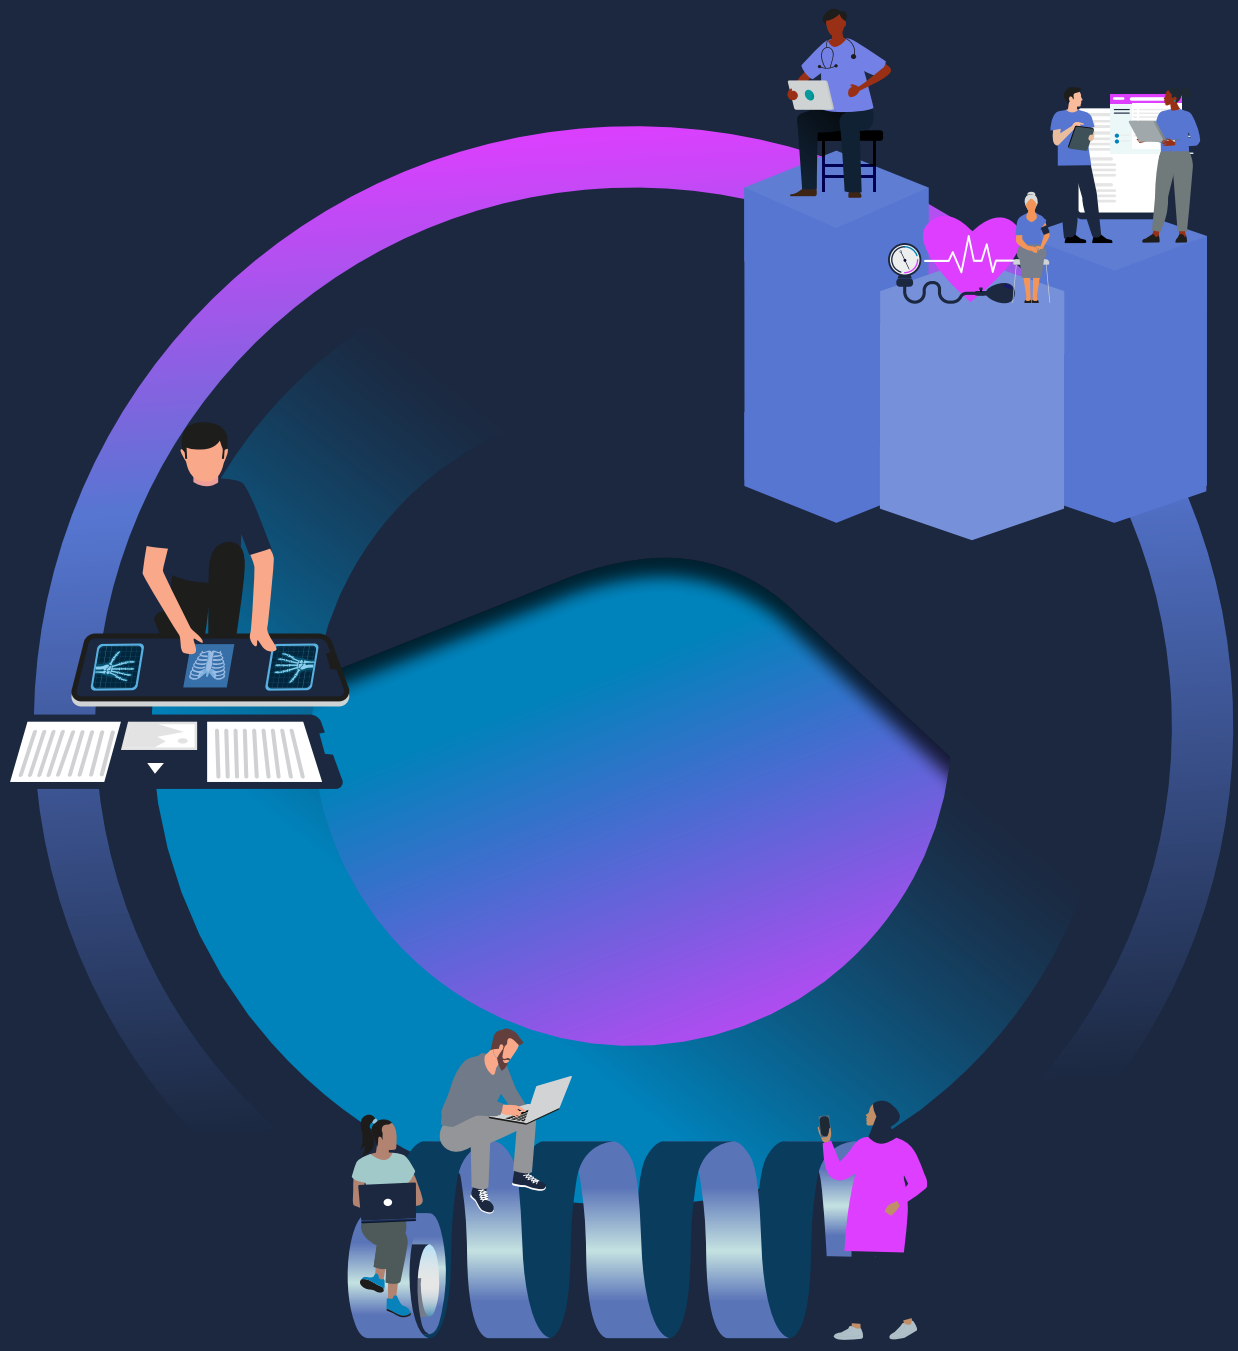

RCP view on digital and AI

January 2026

Contents

Introduction

The NHS is at a pivotal moment in its digital evolution. In an era of rapidly advancing technology, digital systems have become essential to healthcare delivery. There has been an exponential increase in the data available to clinicians about each individual patient. In recent years, artificial intelligence (AI) has transformed society and the way that we work, leading to its increasing use in healthcare, as well as increasing pressure to use it for everything else. As the NHS faces increasing demand, workforce pressures and the need to deliver more efficient care, digital systems and AI are potentially powerful tools to support 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.

The analogue to digital shift in the government's 10 Year Health Plan sets out an ambitious vision for a digitally enabled NHS in England. The government has also announced plans for NHS Online, an online hospital that will connect patients to clinicians anywhere in England, available from 2027. The 10 Year Health Plan promises to overhaul the NHS App so that it becomes a single front door to the NHS for patients, to make AI every clinician's 'trusted assistant' by automating tasks and supporting decision making, and to ensure that all NHS staff are trained to use AI through reforms to medical curricula and training. Central to this vision is the belief that AI will fundamentally transform healthcare, enabling more personalised care, better prevention and improved population health.

This report sets out the Royal College of Physicians' (RCP) view on the role of digital and AI in the NHS, drawing on the experiences of physicians across the UK. We frequently hear from our members about the frustrations of poorly functioning IT, and its impact on their wellbeing and ability to deliver the best possible clinical care to patients. As a fast-evolving technology where our understanding is developing, this report does not provide definitive views or solutions on AI, but examines its growing use in clinical practice, its potential and the risks that must be addressed. AI is already reshaping workflows to improve clinicians' working lives through tools such as ambient voice technology (AVT) and diagnostic support. Successful widespread 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 its development must be grounded in trying to solve real-world clinical challenges rather than in technological possibility. Clinical safety must be at the heart of AI development.

The 10 Year Health Plan made a welcome commitment to investing in digital infrastructure, but delivering a successful shift from analogue to digital will require more than that. Digital systems can enhance care when they are well designed, interoperable and responsive to clinical needs. But poor usability, fragmented infrastructure and lack of standardisation can undermine patient safety, clinician productivity, and patient experience and outcomes. Focusing on innovation without effective implementation risks undermining its benefits and introducing new risks to clinical and patient safety. Digital and AI transformation must be driven by clinical need, co-designed with patients and clinicians, and underpinned by strong clinical leadership, standards, infrastructure and training. We must optimise existing systems and understand how the use of different digital tools can help or hinder high-quality care.

Providing the NHS workforce with the right level of education and training will also be key, so that they can understand not only how best to co-design and use these technologies, but also their risks and limitations, and how to mitigate these.

There is significant potential for AI to address health inequalities by identifying patients who are at greater risk or who would benefit from targeted interventions; the risk is that, if we poorly implement AI and fail to educate clinicians on bias in existing datasets and AI tools, we will amplify and entrench inequality.

We must also be vigilant about the significant risk of optimism bias around AI in healthcare, especially in addressing workforce pressures. Many of the productivity gains which it is hoped that AI will deliver also rest on significant systemic change.

The risks of digital exclusion, as well as issues around health literacy, must be actively considered and addressed in policy interventions. The analogue to digital shift must take a 'digital plus' approach, rather than 'digital only', to ensure that people who can't use digital systems can still access care.

There are many important considerations in an increasingly digitised world and health system: for example the environmental cost of AI, the critical importance of robust cybersecurity to protect the NHS and patient safety from attacks, and the vulnerabilities of digital systems to extreme heat. It is far beyond the scope of this report to offer definitive solutions to all of these issues – but the report aims to outline the steps needed to ensure that digital and AI technologies work for the NHS workforce and the patients they serve, so that current and future innovation leads to safer, more effective and more equitable care.

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.

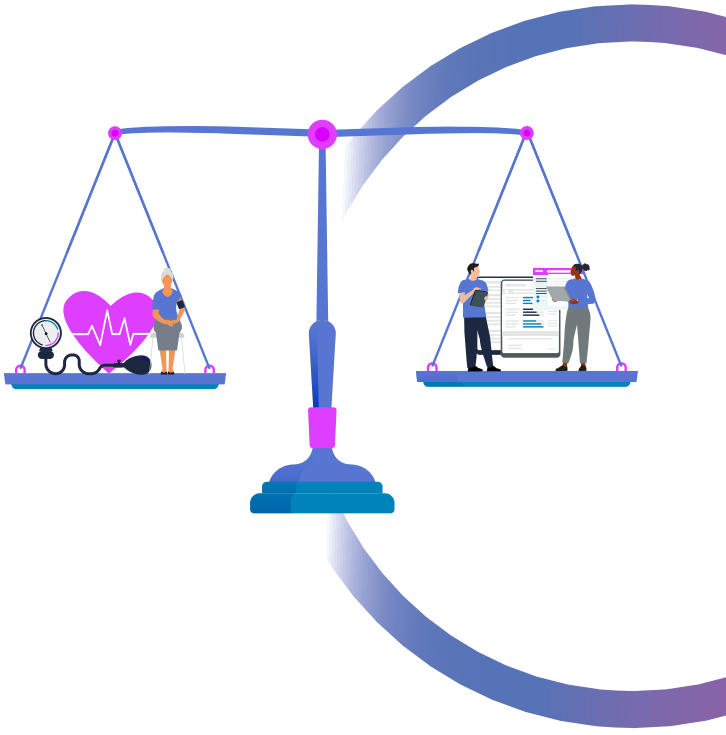

8 The government's promised 'roadmap for AI in the NHS' must set out a plan for effective and ethical implementation of AI in the NHS, including how AI can enable clinical research and tackle health inequalities by actively improving equity of access, experience and outcomes in the health service. DHSC 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.

68%

UK physicians either somewhat or strongly disagree that the NHS has the right digital infrastructure to support widespread introduction of AI that will make a difference.

We need to define the digital hardware needed in different clinical environments.

Digital systems don't talk to each other, aren't intuitive to use and hardware is often out of date or broken. We need to optimise existing digital systems.

We need to fix the NHS' digital infrastructure as the foundation for safely and successfully integrating AI into healthcare.

Clinicians must be part of the design and development of clinical digital systems and have training on how to use them.

Cybersecurity for NHS systems must be a priority for government.

There should be an electronic patient record model content specification standard for suppliers to meet.

Organisations should be 'digital plus' rather than 'digital only' to reduce digital exclusion and health inequalities.

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 clinical systems should support the delivery of seamless, productive care in the way that access to the picture archiving and communication system (PACS) has transformed access to imaging, improving reporting and access to results. However, too often, different digital systems are unable to share information. This means that information does not follow the patient, leading to siloed working. There are some pockets of the NHS where local organisations have supported innovation and implemented digital transformation well, showing the potential of what could be achieved – but this has not yet expanded across the country.

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'.

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. These are all essential components in enabling a functioning health service where clinicians can provide the best patient care. Functional and user-friendly digital systems that can bring evidence-based clinical knowledge to clinicians and patients are key to improving care standards, patient experience and clinician satisfaction.

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. 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. Digital innovation is unlikely to deliver the benefits that it could if we don't get the basics right first.

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. Trusts and integrated care systems (ICSs) have been incentivised to achieve digital maturity for their EPR, using stages (1–7) of the Healthcare Information and Management Systems Society (HIMSS) Electronic Medical Record Adoption Model (EMRAM) as a benchmark. The HIMSS model focuses on the extent to which an organisation has transitioned from paper-based to digital records, with stages 6 and 7 also referring to its ability to manage data effectively.

‘Do you believe the NHS has the right digital infrastructure to support widespread introduction of AI that will make a difference?’

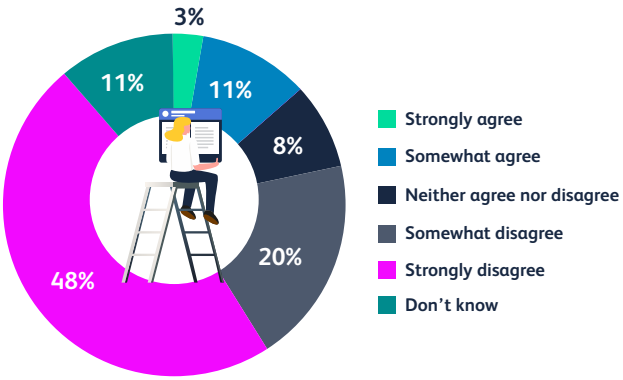

548 UK physician respondents, June 2025 snapshot survey

NHS England (NHSE), on the other hand, defines broader digital maturity as an organisation’s ability to respond to changes and trends in technology or its ‘state of readiness’ to be able to adapt to, and integrate with, these technologies.

There is no question that all clinicians should have access to digital information about their patients, but simply having an EPR in every trust will not deliver the digital ambition laid out in the 10 Year Health Plan. We must focus on optimising 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. 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.

Without optimising existing systems, we risk continued deployment of digital systems that are inefficient, compromise patient safety (especially in cross-sector working), burden clinicians and fail to enhance care.

It should not be a question of analogue or digital, but how ‘good’ the digital is. Technology or digital systems do not automatically equate to better productivity in the NHS – poorly functioning digital systems hold back clinicians and patient care. The objective should be for information to follow the patient in a consistent and intelligible format.

Prioritising digital usability

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 – 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 – shows how digital systems can resolve some issues associated with analogue systems, while introducing new challenges. A digital NEWS chart in the EPR eliminates calculation mistakes, but the lack of an agreed format for NHS digital NEWS charts means that it loses the benefit of the standardised paper chart in clearly demonstrating deterioration. There is currently no consensus on the best way for digital systems to alert clinicians to an abnormal NEWS score.

A lack of structured data, combined with poor design of observation charts and dashboards, makes it harder to identify patients at risk. Without standardisation in the EPR for structured data entry, each clinician can record patient information differently and for different purposes. In this way EPRs are used as ‘digital paper’, where clinicians record digital notes in the same way as they would on paper. 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.

Poorly designed systems and incomplete patient records also increase the risk of serious error, such as missed test results or misdiagnosis, and also increase the risk that electronic referrals and handovers get lost in the system. Workarounds to make timely care possible, such as copying from one note to the next, can lead to inaccuracies in the record and, in some cases, patient harm. While inaccuracies in documentation – such as those introduced by copy and paste – can cause harm and may be a breach of standards, we must recognise that poorly designed workflows will increase the chance that users develop workarounds that increase risks.

‘Digital systems vary widely between hospitals, and it really affects how we work. In one trust, I can use voice recognition to dictate an outpatient letter straight into the EPR. But in another, I type my note into the EPR, switch systems to dictate a letter, which a secretary edits then sends back for approval, which can take multiple rounds of edits. Even adding test results to a letter varies: some systems let me drop them in, while others require manual input by a secretary. These inconsistencies make processes much less efficient than they could be.’ – RCP fellow

Systems should enable clinicians to document easily and quickly 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.

User-centric design and clinician involvement are key to achieving standardised data entry without adding burden. These tools should make it easier for a clinician to enter frequently used terms (such as Neuro NAD), as well as supporting better entry of structured data that support interoperability.

Clinician experience

As well as posing a serious risk to patient safety, poor usability also adds to the operational strain on clinicians. We often hear clinicians’ frustrations about working with different digital systems and the differences between hospitals – what might be one process in one system in one hospital might be five or six processes across two systems in another.

There is an assumption that digitising workflows will deliver quicker, automated 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. Experienced physicians (RCP members and fellows) consistently tell us that while they are much more effective as clinicians due to their training and experience, they are significantly less productive as a result of poor digital systems.

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.

What determines usability?

In a resource-constrained environment, it is important that NHS trusts get maximum use from the systems that they pay for. Since 2021, NHSE 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. The 2021–22 survey also found that, relative to other global systems, the NHS is poor at training, infrastructure and clinical engagement – all key to usability.

There are a number of factors that can improve the usability of clinical digital systems.

User experience and design principles

Clinical input into the procurement, design and development of clinical digital systems is key to their functionality. Currently, trusts often fail to engage with user experience to improve systems. The NHS design principles set out how care processes and the digital systems that support them should be developed. They 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.

Personalisation

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, such as 'auto texts' to automatically input frequently used text and automatically drop results into the notes, or 'favourites' folders that allow rapid access to pre-completed orders for tests or blood tests. This can make it easier for clinicians to input and use structured data. As there will rarely be one approach that fits all, micro-personalisation means that the same workflow or template can be adapted for different clinicians, reducing design costs and system complexity.

Education and training

Usability can also be improved with training. The 2024 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.

Ideally, digital systems would be intuitive, supporting staff to deliver and document care in the best way possible, and not require extensive training to use. However, it is important that all clinicians feel competent and confident in using complex systems. A report by NHS Confederation found that workforce and training were key to ICSs being able to provide and deliver a quality frontline digitised service. This training in digital systems should take place for new starters and be ongoing for existing employees. Organisations should routinely analyse how staff are using the digital systems, to offer tailored additional training to those who are struggling with the system or not using it optimally.

Functional 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%). On the flipside, when we asked what would make the biggest improvement to physicians' wellbeing at work, well-functioning IT equipment was the most common response (43%), ahead of reduced clinical workload (33%).

Top three priorities for improving workplace wellbeing identified in the 2025 'Focus on physicians' survey. Based on 1,398 respondents.

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. Hospital working is vastly different from office working. 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 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. Hospital at home and the rollout of remote monitoring will require the design and implementation of wearable remote monitors, as well as devices that allow clinicians to access and update the EPR from the patient's home.

Systems that work well together

Good interoperability is crucial to ensuring that different systems, platforms and technologies can communicate and share data seamlessly across the NHS. A lack of interoperability means that clinicians struggle to get a full picture of a patient's medical condition. This fragmentation can lead to gaps in care or missed critical information, as well as an inability to undertake population health management. Patients may have to repeat their medical history to each part of the health service that they encounter, which can lead to wasted time, and poorer experience and treatment quality. Doctors want systems where patient records are visible to all areas of the healthcare system rather than the current siloed approach, and patients think that we already have that. Improving the data available in the EPR and sharing data between multiple EPRs will enable the use of longitudinal patient data for individual patients and integrated population data to detect risk. It will be easier to track patients' health and journey through the NHS, improving continuity of care and reducing waste. 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. Interoperability is critical to delivering the vision of the 10 Year Health Plan. The single patient record announced in the 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 so they not only include health and care records, but 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

Setting digital standards for clinical processes in the NHS would improve safety, usability and reliability of systems, and reduce the cognitive load on clinicians. In this context, standards refer to established rules and guidelines which aim to provide consistent expectations for a digital technology to ensure that patients remain safe and that standards of care are upheld and improved.

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.

Having standards in place for digital systems in the NHS will decrease cognitive load for clinicians and improve patient safety. For example, the EPR should make it easy to collect data generated from routine care to use in audit, research and improvement. 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. Standardising datasets may improve our ability to do this. The Emergency Care Data Set (ECDS), a dataset with an information standard that is collected about people attending emergency departments and the treatment that they receive, may be a good example. There should be particular focus on collecting datasets for patient care that crosses the interface where a range of teams and professionals are involved, such as outpatient care, to improve patient pathways.

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.

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. Just as usability and safety standards guide system design and procurement, cybersecurity standards must ensure that systems are resilient to threats. The 2024 ransomware attack on Synnovis laboratories, which disrupted NHS services across south-east London and was linked to patient harm, highlights the serious consequences of cyberattacks. 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

While the move to digital services in the NHS can streamline processes for staff and make interactions with services more efficient for patients, we must ensure that all patients are brought on the journey. 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.

Digital exclusion often correlates with social exclusion and those who are in more vulnerable groups, such as people with disabilities, those living in more deprived areas and those with limited digital literacy. 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. These groups are often more likely to experience worse health outcomes – understanding this relationship and patients' general health literacy are key to ensuring equitable healthcare advancements, and that health inequalities are not exacerbated by the analogue to digital shift.

We should not make assumptions about which groups are less likely be digitally literate, for example based on age or access to digital devices. Digital tools should be used to lessen inequalities, not worsen them. We must recognise that the shift from analogue to digital risks excluding people who either have no access to technologies or data, or people who cannot successfully navigate digital systems because of the appropriateness of those systems for their needs, their confidence, capability or motivations.

To ensure that digital services meet patient needs, it is crucial to engage patients throughout their design, development and implementation. The Darzi report into the state of the NHS found that the patient voice is often not sufficiently heard in the design of services: 'the NHS could do better at involving real experts (those living with an ongoing health condition) in how care was provided'. The RCP was pleased to see that patient engagement and co-creation were identified as ways to maximise the inclusive potential of digital technology in the 10 Year Health Plan. The NHS needs to engage meaningfully with patients when designing services, including groups that are most at risk of health inequalities, who can be harder to reach. The King's Fund paper on creating inclusive digital services in collaboration with people and communities recognises that while time and money are currently major constraints in the health system, good engagement is still possible, and should not be overlooked.

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. The 10 Year Health Plan promises that inclusion will be designed into the NHS app by default. This includes tailoring health information to meet patient need and proactively identifying people who have lower digital literacy to offer support. While we need to see more detail on implementation, we welcome approaches that support patients to understand and manage their own health, whatever their level of 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.

Summary

Digital transformation in the NHS must prioritise introducing and optimising systems that are safe, usable and interoperable, enabling clinicians to deliver high-quality care without added burden. Usability depends not only on system design but on having the right infrastructure, thoughtful implementation, and training that is agile and iterative. Standards must be strengthened across procurement, design and data to reduce variation and support system-wide consistency. Moreover, digital inclusion needs to be actively pursued through co-designed services and tailored support, ensuring that technological advancement reduces, rather than reinforces, health inequalities. If these principles are realised, digital tools can support clinical decision making, reduce administrative burden on clinicians, and enhance patient experience.

70%

of physicians said that they were either very or somewhat supportive of AI tools being implemented widely in the NHS.

Good implementation and understanding dataset biases are essential to ensure that AI tackles, rather than entrenches, health inequalities.

AI tools must solve real-world clinical problems with clinicians and patients involved in development.

AI must support – not replace – the clinical judgement of doctors.

Successful widespread AI adoption will require robust governance, clear accountability, clinician involvement and a commitment to equity and transparency.

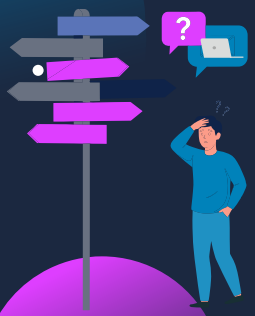

Avoiding optimism bias is essential.

73%

of physicians said that their biggest concern about AI in clinical practice was the risk of error.

Training and education are essential to build an AI literate workforce and digital medical leaders.

There must be a bank of NHS-approved AI tools and apps, robust, joined-up regulation and a clear government plan for implementing AI in the NHS.

2 AI

AI, and in particular machine learning, have 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. In certain specialties, such as radiology and pathology, AI has already demonstrated positive impact – machine learning has improved diagnostic accuracy, as well as efficiency. Beyond this, AI presents an opportunity to address many of issues in the EPR, which have been explored in the first part of this report.

The exponential growth in patient- and service system-level data available to clinicians and NHS organisations from electronic systems, wearable devices and digital observations has created the basis for AI to transform data into actionable knowledge. AI is already rapidly improving the analysis of data, including biobanks, experimental data and routinely collected clinical interactions, which can convert patient data into insights for operational decision making, service improvements, research and innovation.

There are examples of AI in clinical practice: in neurology for example, AI is automating and supporting image interpretation tasks, as well as being used as a tool in neuroscience research. Intelligent liver function tests (iLFTs) at Ninewells Hospital in Dundee use algorithmic processes within blood sciences systems to facilitate the correct testing of patients with possible liver disease, as well as advising primary care on the right actions for the results. But there is no widescale successful rollout of AI in medicine in the NHS apart from image analysis.

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. Issues such as model accuracy, reliability, bias, health inequality and automation bias in clinical decision-making support tools must be carefully considered and managed. Clinicians and organisations will need the skills and knowledge, and governance frameworks must be in place, to help AI tools become more safe, effective and equitable.

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. The Health Foundation has suggested a twin-track approach to future AI development: setting out high-level priorities for AI use and supporting the testing and spread of these tools, as well as supporting the most promising innovation that is taking place locally.

AI has significant potential, but it is not a panacea. To ensure that it 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. The primary measure of success for the implementation of any clinical AI tool must be that it improves outcomes for patients.

How widely is AI being used by physicians?

In our June snapshot survey, 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.

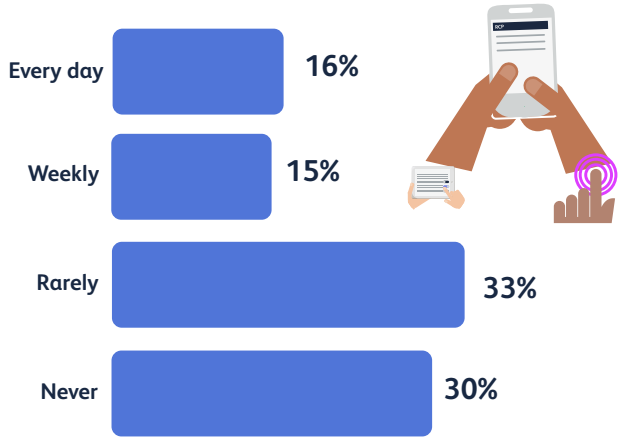

When asked what clinical tasks were supported by AI in their organisation, radiology and pathology interpretation was most commonly cited (42% of 317 respondents). This was followed by using ambient AI for letters or notes in outpatient settings (29%) and AI to support clinical decision making (19%). The use of explicit LLMs within the EPR, and using AI to predict clinical changes like deterioration, did not attend (DNA) rates and discharge, were each reported by fewer than 23 respondents.

70% of physicians said that they were either very (29%) or somewhat (41%) supportive of AI tools being implemented widely in the NHS. 5% were somewhat unsupportive and 7% not at all supportive (10% were neither supportive nor unsupportive, and 9% said that they didn't know). With the government's ambitions for AI, such as all hospitals being 'fully AI enabled' within the lifetime of the 10 Year Health Plan, the use of AI tools in the NHS will continue to grow.

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. The analogue to digital shift will speed up the presence and role of AI in the health system, but there has not been a concerted rollout of AI tools across the NHS so far. 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, causing clinician and patient frustration and patient safety. Having a diversity of digital systems and competition between providers, combined with procurement at an individual trust level, has been a failure and is something that we cannot afford to repeat with AI. We must take advantage of the NHS being a national health system.

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. The government's commitment to produce an NHS AI strategic roadmap is a vital opportunity to set out a coherent vision and approach. There is learning from the introduction of digital tools in the NHS that should be applied to AI. Better coordination and evaluation of approaches is needed at regional and national levels, with the most successful then shared for wider implementation in systems nationally.

Supporting, not replacing, our workforce

AI has often been posed as a 'silver bullet' to improve productivity. The NHS workforce is under pressure, working hard to reduce waiting times and meet demand. Increased productivity is an even bigger priority in the context of the government's promise to restore the 18-week treatment target. It is easy to see why the potential of AI is positioned as a quick or easy fix.

Clinicians themselves feel the most optimistic about productivity savings when it comes AI. 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. Only one in five clinicians selected better outcomes for patients (20%), and 8% said that there were no benefits. The existence of, or access to, AI tools alone will not improve productivity. In fact, much like digital, 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 act as an enabler, for example by reducing 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. There are lessons to be learned from other sectors, like hospitality, where AI can automatically offer cancelled slots to those on waiting lists. The 10 Year Health Plan pledges that 'AI-backed ambient voice technology will automate clinicians' note-taking', eradicating the 'need for tasks like clinical note taking, letter drafting and manual data entry'. Locally, there are some encouraging case studies of AVT aiding clinicians in their consultations and boosting productivity. In these cases, AVT has shown promise for note taking, but there is still a way to go in automatically integrating or coding these data into records. NIHR RSET (Rapid Service Evaluation Team) is currently carrying out an evaluation of the use of AVT in the NHS to determine the extent of its benefits, which is due to be completed by the end of 2026.

Great Ormond Street Hospital carried out a pilot project using AVT in their outpatient appointments, which was then expanded to a number of NHS environments across London. The London pilot evaluated over 18,000 patient encounters and found that the use of AVT in patient encounters resulted in an average time saving of 8%. In emergency departments, these immediate time savings indicated the potential for clinicians to see 10% more patients, while in outpatient and GP settings, the assistance of AVT meant that clinicians could spend around 15% more of the appointment solely focused on the patient, as less time was spent on writing notes. However, the pilot also highlighted how the implementation of AI tools needs to take a phased approach, carried out with clinicians, rather than it being done to them. This includes giving time for new technologies to be embedded into work processes before reviewing capacity, and creating personalised templates for different specialties and services. For the latter, it was found that a generic template across services did not work as it did not capture all information needed, whereas tailored templates can help to capture the essential information that clinicians need – this ability to personalise digital and AI technologies is what will support successful implementation.

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. We need to remember this as we design workflows in tools like AVT, and consider how we can capture every useful part of a process.

It is 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. This would allow AI to support the ordering of relevant tests, documenting in the correct place in the record, structured coding that can be shared between systems, and drafting letters to patients and clinical colleagues. This would improve patient care and reduce burden on clinicians. In high-demand areas like emergency departments, this could lead to significant productivity gains. In outpatient settings, AVT could enable clinicians to complete all tasks related to an appointment within the scheduled time – something that most struggle with. Only 34% of respondents in a February 2025 RCP snapshot survey said that their job plans included time for outpatient work beyond the appointment itself. If the time saved can be used for more patient contact and less administration, AVT has the potential to transform clinical practice, allowing doctors to focus on care and decision making, rather than administrative tasks.

Productivity gains from AI should also 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 ultimately contribute to improved patient care and job satisfaction.

We must be vigilant about the significant risk of optimism bias when it comes to the role of AI in addressing workforce pressures. 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.

The use of AI in healthcare should be driven by its potential to improve patient care and free up clinicians to undertake other vital tasks. Realising these gains rests on several things: ensuring that AI tools are designed to solve real efficiency and clinical challenges in healthcare, having the right data to train AI tools, and ensuring that high-quality AI tools are integrated into wider NHS systems and that clinicians are confident to use them. The tools must be integrated alongside the right people, processes and systems.

Making AI tools useful for clinicians and safe for patients

AI solutions in healthcare have often been driven by technological possibility rather than clinical need, leading to tools that have to be retrofitted into existing workflows. As with all digital transformation, AI will be most effective when it responds to clinical need or tries to solve real problems in clinical processes, designed with and by clinicians and patients.

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 government and the NHS must grip this issue. With the rapid evolution of AI technologies, the performance, usability and impact of these tools need to be monitored in clinical settings to ensure that AI systems are delivering the intended outcomes. The UK is currently behind in this regulatory process – the EU AI Act came into force in 2024, introducing strict compliance measures alongside performance monitoring, and a number of regulations have been introduced in the USA. Meanwhile, monitoring practices have been underdeveloped in the NHS. But the government's 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) to review current regulations and provide recommendations for a new regulatory framework for AI in healthcare, was launched in September 2025.

The work of this commission, alongside the NHS AI roadmap, will be key to ensuring that we have the right governance and regulation to protect and reassure doctors and patients about the use of these tools. 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 number of respondents who thought that too much regulation was a barrier (18%). This shows that there is a demand for clearer, more consistent governance frameworks that can support safe innovation in the NHS. Robust outcome measures are needed to continuously check that the AI is working to its intended purpose and to ensure that people interact with the tools in a safe way.

One area where there is clear need for stronger regulation is AVT. NHSE'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, given risks to clinical safety and data protection. This letter also identified the risk of fragmentation to the broader NHS digital strategy. 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. This will depend on better regulation, good contracts with vendors and good design standards.

One of the key challenges to the safe oversight of AI is ensuring that NHS organisations are equipped with the right processes to feed in high-quality data and monitor AI performance over time. To do this, a dual approach is needed, where responsibility is placed on AI developers and suppliers and NHS organisations. The suppliers should provide built-in monitoring tools, while NHS providers maintain oversight and ensure safe interactions with AI systems.

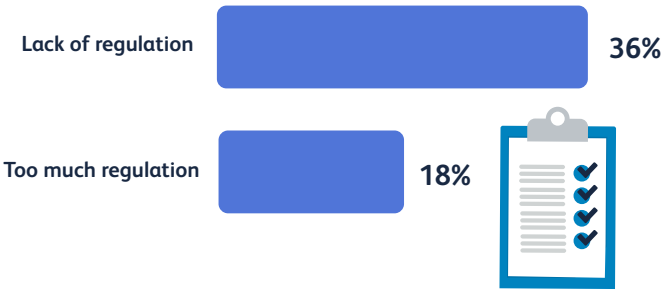

541 UK physician respondents, June 2025 RCP snapshot survey, on the main barriers to deployment of clinical AI systems in the NHS

Engagement

Meaningful iterative development with clinicians and patients is essential throughout design, testing and implementation. Clinical need and safety must be placed at the centre of any technological development. 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 systems such as the EPR, and 65% said poor interoperability of systems. Retrofitting tools after development, rather than designing them to fit from the outset, has been a recurring issue in digital innovation, often increasing cognitive load on clinicians and inadvertently introducing patient safety risks. In multiple systems, time has not been taken to include clinicians and patients in the design, testing and implementation of tools, meaning that the rollout has failed or the tools do not solve the issues they were designed to solve. We must learn from this and ensure that the development of new AI tools is driven by clinical need and fits the workflow. Clinician involvement, alongside robust regulation, is key to ensuring that medical professionals trust (and therefore adopt) the tools.

It is critical that, in the context of NHS staffing constraints, time is prioritised for doctors to be engaged in digital tool development and be part of iterative testing. This level of involvement is central to effective implementation and realising the benefits of digital and AI tools to patient care. If resource constraints and culture within the NHS do not prioritise digital engagement as part of clinicians' work, the success of the shift will be put at risk.

Patients should ultimately have ownership over their health and healthcare. 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. This is critical, as patients are more likely to share data for AI purposes if they have trust and understanding. The NIHR Patient and Public Involvement and Engagement approach provides a good example of how listening to patients can ensure that technologies are aligned with patient need.

The incremental development of AI tools, led by clinicians who understand the complexities of everyday healthcare delivery, is also vital for patient safety. Using AI tools to support clinical decision making must come with a commitment to evaluate and iteratively improve performance. 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

The June 2025 snapshot survey of RCP members highlighted concerns about liability and responsibility when AI is used in clinical practice, and the impact on clinical skills of an over-reliance on AI.

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%).

541 UK physician respondents to the RCP's June 2025 snapshot survey on the main risks they are worried about in using AI in clinical practice.

Clinical safety is clearly a priority for physicians. The potential for AI to cause error or be inaccurate, or worries about this being a likelihood, could be a potential blocker to the widespread use of AI in clinical practice. Model drift or explainability are concerns that apply to systems which are in use now. Action needs to be taken to ensure that post-deployment surveillance and processes are in place to mitigate these issues in current and future systems.

It should also 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. This is especially concerning as some AI systems, particularly those based on LLMs, operate as 'black boxes' with limited explainability. Beyond this, 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, with senior consultants and experienced medics doing this every day in clinical practice. 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.

Doctors need to maintain control and choose whether to agree or disagree with an AI's output. AI is there to support, not to replace, the expert clinical judgement of doctors. Good patient care depends on not just the diagnostic algorithm but a wide range of contextual factors, not least patient preferences and choices as part of shared decision making, or taking account of the impact of frailty or the approach of the end of life. For this reason, it is critical that clinicians can maintain their diagnostic expertise, built over years of training, rather than becoming dependent on AI for decision making. 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.

These issues are already present in some digital systems – for example in electronic prescribing, where there may be pre-filled 'order sentences' recommending certain doses of a drug. This means that they can become a default, even if that is not the intention.

In this context, prioritising high-quality medical training for resident doctors is essential. For example, 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. Doctors will need to able to make clinical decisions based on their skills and judgement about whether to follow the advice of AI clinical decision-making support tools. In the face of increasing service and staffing pressures, we must not be tempted to use AI at the expense of supporting residents to develop clinical judgement and expertise that will allow them to use AI decision-making support tools safely in clinical practice.

Clear lines of responsibility and robust governance frameworks are essential 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. The regulatory framework for AI in the NHS should provide sufficient safeguards to ensure that all use of in AI healthcare is safe and ethical.

Use of personal AI tools not provided by the NHS

The availability of free AI tools in apps or browsers means that clinicians are not bound by what is offered by their organisation. 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 generative AI 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. There must be clear guidance about acceptable and appropriate use of AI in the NHS. The NHS is working to provide this: as highlighted earlier in this report, NHSE warned trusts in June 2025 about halting the use of AVT that did not meet minimum standards, including those for personal use. There must be further guidance from the NHS about a wider range of AI tools.

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. It is vital that NHS organisations can meet clinicians' AI needs so that they can harness and benefit from this technology safely. 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 AI tools to clinical and patient safety. For example, the Welsh government has published interim guidance on the safe and responsible adoption of AVT in clinical settings to support organisations and individuals in how they can be used, while detailed guidance is in development. Clinicians must understand the differences between using NHS and non-NHS AI. It may be beneficial and safer to focus on implementing AI tools that support the NHS workforce with non-clinical tasks such as admin, which may carry less risk in terms of patient and data safety, but still provide benefits in terms of productivity or improved work experience. More innovative clinical AI tools should continue to be developed, with those that have been through the most rigorous testing with the most positive outcomes then introduced across the NHS.

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 as an enabler, not an obstacle

The availability of complete clinical datasets is essential to making AI usable and effective – AI is only as good as the data that go into it. Over half (51%) of respondents to the RCP's June 2025 snapshot survey said that data were the main barrier to the deployment of clinical AI systems in the NHS.

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. This is essential to avoid exacerbating health inequalities and bias. The inability to generate accurate comprehensive structured data from EPRs will make it much harder to realise the full potential benefits of AI tools. In addition, as patients directly access their health records more regularly, systems need to be designed to export data to the patient-facing portal in a way that is accurate and easy for patients to understand.

One of the biggest data concerns is the quality and representativeness of datasets used to develop AI systems. 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.

Alongside having complete datasets, biased algorithms can be avoided by developing AI with a holistic, human-centred design approach and ensuring that AI developers themselves reflect a range of backgrounds and experiences. Research from the University of Oxford and Imperial College London has underscored the importance of accurate, representative data, particularly in regards to ethnicity, in building equitable AI systems.

The NHS Research Secure Data Environment Network has been designed to provide secure and fast access to health data for research and development purposes, including for the development of AI algorithms. The network is relatively new, operating in all areas of England since March 2025, but it could help to ensure that the development of AI tools is based on accurate, curated data that can create better algorithms. The network also aims to improve interoperability by aggregating data at national and regional scales, which will reduce data silos and landscape complexity. The success of OpenSAFELY, a programme that provides a safe platform with protected GP records for the entire England population to NHS researchers, is an example of good practice for how patient data can be utilised and safely shared for research and innovation.

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. In future, as the range and volume of data held on individuals expand, AI may be able to identify patients from data items like an ECG, which is currently considered to be non-identifiable data. 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

A skilled digital- and AI-literate NHS workforce is essential. 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 (79% of respondents) said that they need training in clinical AI tools. When asked if they had access to training, 66% said no, 28% said that they did not know, and only 6% said yes.

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. All doctors must understand the new ways of working that come from a digitally enabled NHS. For example, almost all prescribing in the NHS is now digital, the vast majority of documentation in the NHS is digitised, and the remote monitoring of patients through digital tools is becoming increasingly common. Doctors' education and training must reflect the digitised NHS that they are already learning and working in – understanding digital prescribing risks and alerts, data collection in electronic systems for audit, QI and research, medtech, and holding risk in remote healthcare. 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. There is a need for a digital curriculum at all levels and routes into careers as clinical informaticians.

The 10 Year Health Plan commitment to reform curricula to include comprehensive training in AI and digital tools is therefore a welcome move. Almost three-quarters (74%) of respondents to the June 2025 snapshot survey of RCP members said that a lack of clinical expertise in AI was the main barrier to deployment of clinical AI systems in the NHS.

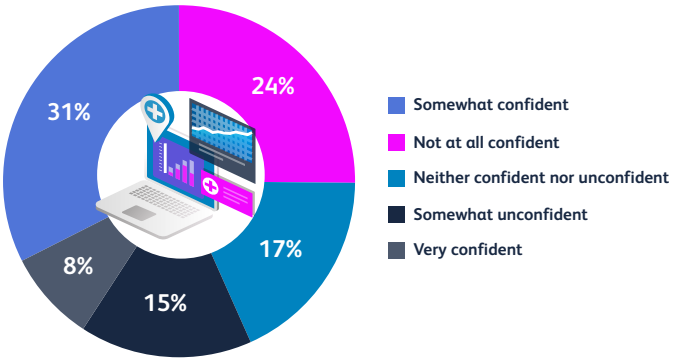

578 UK physician respondents on their confidence in using AI tools in clinical practice.

Embedding AI into the medical curriculum will equip the next generation of doctors with better understanding and confidence to use these tools in their clinical work (including their limitations)and become AI confident clinical leaders. This is critical to the next generation of doctors being AI-confident digital clinical leaders who can shape AI policy and practice in the NHS, and support and lead innovation. As AI becomes more prevalent, training is also crucial for those already working in the NHS. We welcome the AI upskilling programmes for the NHS workforce that were announced in the 10 Year Health Plan. With proper training, healthcare professionals can interpret AI outputs more effectively, identify when AI might be incorrect or biased, and combine AI insights with clinical judgement for better outcomes. It is important that clinicians understand the risks posed by sharing data with AI tools to ensure that patient confidentiality and safety are maintained. Doctors should have the skills to recognise and mitigate risks and have a deeper understanding of how AI works in clinical care, which should lead to safer and more informed use of AI in the NHS. We welcome that the government wants the NHS to have the most AI-enabled workforce in the world. 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. Ongoing training, tailored to individual need, will improve clinical efficiency and patient safety.

Population health and health inequalities

When designed properly, AI can actively help to reduce health inequalities. Equitable access to innovation must be prioritised across the UK. Reliance on individual champions or specific trusts risks concentrating development in a few centres, leading to population bias and limiting broader impact.

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. Poor health literacy means that patients find it harder to understand their own health, and it affects their experience of the health system. In the UK, 7.1 million adults read and write at or below the level of a 9-year-old child, and studies have found that 43% of adults do not understand written health information. AI can summarise and simplify complex health information for patients, and explain both conditions and issues, as well as how to manage them.

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. AI tools have also been used to help manage waiting lists, by using predictive analytics to prioritise patients on elective waiting lists who are most at risk of deterioration due to factors such as inequity and the wider determinants of health. In 2023, three trusts in Cheshire and Merseyside used an AI-enabled patient tracking list to help prioritise elective care. The tracking list used data from more than 200 million records from 46 countries over the past 30 years, including social determinants of health, to review individual patient health profiles and they type of treatment that they were waiting for, to estimate risk of complications during or after treatment. An evaluation found that the tool accurately predicted the risk of mortality and complication, with bed days freed up and a reduction in the number of long-waiters and those with the highest urgency.

Summary

The NHS has a significant opportunity to harness AI to improve care delivery and reduce inequalities. Realising this potential requires improving digital literacy across the workforce by investing in targeted training, and prioritising tools that reduce administrative burden on clinicians and enhance the patient journey. Decision makers and the health sector need to proactively recognise and address the potential risks of AI and not be blinded by optimism bias, instead creating and operating within robust governance frameworks that foster safe use and innovation to improve patient care. Co-developing explainable, clinically relevant algorithms with patients and professionals will be key to safe and effective implementation.

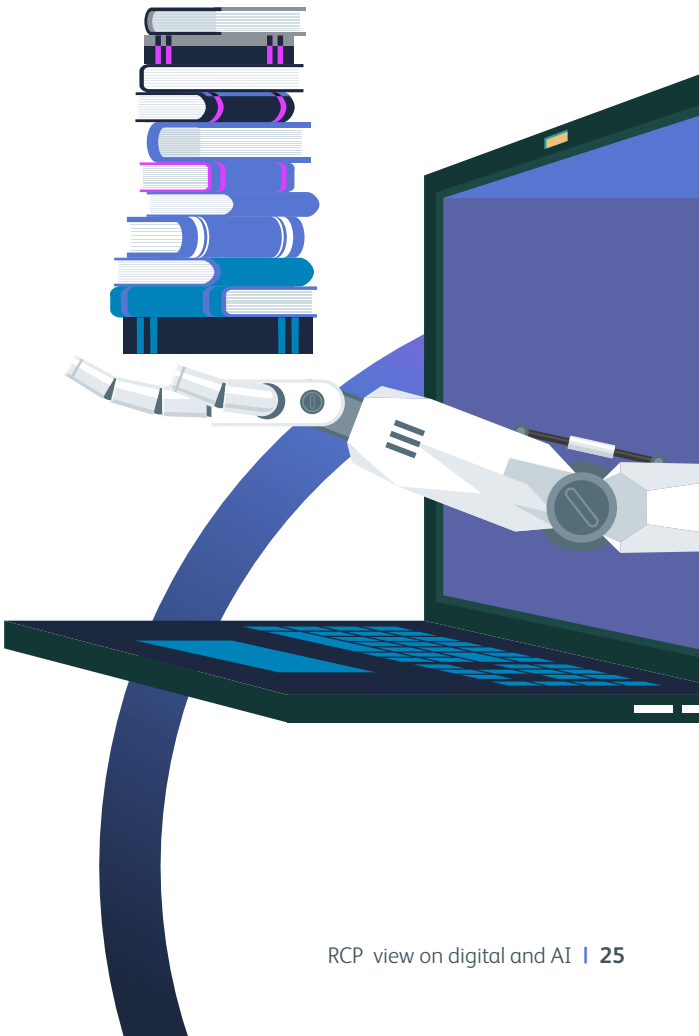

Glossary

AI	Artificial intelligence
AVT	Ambient voice technology
CPD	Continuing professional development
CQC	Care Quality Commission
DHSC	Department of Health and Social Care
Digital by default	Systems that are digital as standard. Digital by default should mean digital services that are so straightforward and convenient that all those who can use them will choose to do so, while those who can't are not excluded
Digital maturity	An organisation's ability to respond to changes and trends in technology
Digital only	Services or tools that only exist in a digital form – such as datasets used by the federated data platform or patients' access to GP notes in the NHS App (without difficulty and bureaucracy)
Digital plus	Systems that are developed as an enhancement to existing pathways to protect staff time for patients whose needs or preferences make digital tools unsuitable
DNA	Did not attend
ECDS	Emergency Care Data Set
EMRAM	Electronic Medical Record Adoption Model
EPR	Electronic patient record: the core digital system at the heart of daily working for most clinicians. Most patients will have data in multiple EPRs, even within one organisation
Explainability	It's not possible to know how the AI produces its output
FDP	Federated Data Platform
GenAI	Generative AI
GMC	General Medical Council
HIMSS	Healthcare Information and Management Systems Society
ICO	Information Commissioner's Office
ICS	Integrated care system
iLFTs	Intelligent liver function tests
Interoperability	Different digital systems, platforms and technologies being able to communicate and share data seamlessly
LLM	Large language model
MESH	Message Exchange for Social Care and Health
MHRA	Medicines and Healthcare products Regulatory Agency
Model drift	The AI algorithm changes over time
NEWS	National Early Warning Score
NHSE	NHS England
OpenSAFELY	A programme that provides a safe platform with protected GP records for the entire England population to NHS researchers
Optimism bias	The tendency to overestimate chances of positive experiences and underestimate chances of negative experiences
Single patient record	A new record proposed in the 10 Year Health Plan to bring together data from multiple sources, including from the EPR and the Federated Data Platform (FDP) to act as a 'patient passport' for seamless care
PACS	Picture archiving and communication system
RCP	Royal College of Physicians
RSET	Rapid Service Evaluation Team
SSO	Single sign-on
Usability	How easy or hard it is to use a digital product to achieve the intended goal

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